message: /home/dlegare/Mail/inbox/12
-- using template mhl.format --
Date: Mon, 10 Apr 95 10:56:16 EDT
To: Denis Collange <collange@sophia.cnet.fr>
cc: proceedings@CNRI.Reston.VA.US

From: proceedings@CNRI.Reston.VA.US
Subject: Re: old references

In-reply-to: Your message of "Fri, 07 Apr 95 13:27:16 -0000."
<9504070827.AA02631@avignon>
Sender: dlegare@CNRI.Reston.VA.US

> I'm searching two references :
>
> Jacobson, "Modified TCP Congestion Control Algorithm", April 30 1990,
> end2end-interest mailing list
>
> Jacobson, "Berkeley TCP Evolution from 4.3-Tahoe to 4.3-Reno", Proceedings
> of the Eighteenth Internet Engineering Task Force, p365 ,Sept 1990,
> University of British Columbia, Vancouver, B.C.

Denis,

I have a hard copy version of the 18th IETF Proceedings which you can purchase for $35.00. All you need to do is fill in the request form included at the end of this message and mail it to us with your check.

You might also try contacting Van directly <van@ee.lbl.gov>. He may know if these materials are available on-line.

Regards,

Debra

========

Attendees of each respective IETF Meeting receive courtesy copies of the Proceedings. However, the non-attendee document charge is $35.00 per document.

TO ORDER IETF PROCEEDINGS PLEASE COMPLETE THE ITEMS BELOW. POSTAL MAIL THIS FORM AND A CHECK FOR $35.00 (MADE PAYABLE TO CORPORATION FOR NATIONAL RESEARCH INITIATIVES) TO:

Corporation for National Research Initiatives
Attn: Accounting Department - IETF Proceedings
1895 Preston White Drive, Suite 100
Reston, VA 22091

Please print or type:

MONTH/YEAR OF IETF PROCEEDINGS REQUESTED:

NUMBER OF COPIES:

PERSON(S)/ADDRESS DOCUMENT(S) AND RECEIPT SHOULD BE MAILED TO:

PROCEEDINGS OF THE EIGHTEENTH INTERNET ENGINEERING TASK FORCE

UNIVERSITY OF BRITISH COLUMBIA
July 30 - August 3, 1990

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

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

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

ACKNOWLEDGEMENTS

The Eighteenth IETF Plenary meeting was held at the beautiful University of British Columbia. This first international meeting was terrific! Many thanks go to John Demco and his staff at the University, especially Marilyn Martin for their help in making this, the largest ever IETF meeting, a success and a pleasure.

I would like to acknowledge the support of CA*net and CDNnet in supporting this meeting. Terminal access, and more importantly, network bandwidth, for all attendees was provided by CA*net. It was due to the support of the CDNnet Executive Committee and member organizations that John Demco and Marilyn Martin were able to devote the time required to put on this meeting.

Thanks also to the Technical Speakers: Brian Handspicker, Thomas VonDeak, Mike Roberts, Mike Hrybyk, Paul Tsuchiya, Mark Crispin, Van Jacobson, Erik Huizer, and Rudiger Volk. Tony Hain, Kathleen Huber, Zbigniew Opalka, Jeffrey Burgan and Dale Johnson provided informative network status briefings. Special thanks to Dennis Ferguson for his introduction to CA*net, an important network in the international Internet.

Special thanks also to Megan Davies who, as a new member of the IETF team, has diligently worked hard to improve these Proceedings. Through tireless proofreading she has insured the information in these Proceedings is consistent and accurate, and most importantly, timely.
Greg Vaudreuil/CNRI

Contents

Contents

Chairman's Message

The Eighteenth IETF Meeting

The Eigtheenth IETF was held at the University of British Columbia in Vancouver on July 31-August 3rd. Let me add my thanks to our hosts John Demco and Marilyn Martin (UBC).

This meeting had numerous firsts. It was the first international IETF meeting. It was also the largest to date, with approximately 300 attendees. Approximately 38 of the current 45 Working Groups met in 49 separate sessions.

We were also very pleased to have the Privacy and Security Research Group (PSRG) meet with the IETF in Vancouver. In addition to conducting its own business, the PSRG met in joint session with several security related Working Groups (e.g., Site Security Policy Handbook and SNMP Authentication). This interaction was very productive. In the future, such interactions with the IRTF would prove to be quite beneficial to IETF efforts.

We were especially pleased to have a delegation from the European networking association RARE at the IETF. Erik Huizer (Surfnet, Netherlands), Rudiger Volk (RIPE, Dortmund Univ), Fernando Liello (Italy), and Olivier Martin (CERN, Switzerland). Erik and Rudiger gave a presentation on networking activities in Europe. The Federal Engineering Planning Group (FEPG) of the FNC met in joint session with the RARE delegation regarding joint US-European activities.

The ANSI X3S3.3 group also took the opportunity to meet in Vancouver. There is a growing amount of joint interest between IETF and relevant ANSI groups (particularly, X3S3.3, which focuses on the transport and network layer of the OSI model). We have tentatively agreed with X3S3.3 that we will attempt to schedule meetings during the same location and date whenever it is convenient for both groups to do so. The next such occasion will be in Boulder, Colorado in December.

The IETF meeting was very full and productive. The final meeting agenda is given in the main body of these Proceedings. However, even a brief list of the highlights would have to include:

- An excellent report on CA*net, the Canadian national network, by Dennis Ferguson (U. Toronto).
- Review of a draft proposal for IP over SMDS.
- A decision to draft a Link Layer Requirements document separate from the Router Requirements document. There was also a proposal to consolidate all IP specific issues into a separate document for the Router Requirements (and perhaps future editions of Host Requirements) to reference. These new efforts would continue to be an initiative of the Router Requirements Working Group.
- Near closure by the Router Discovery Working Group.
- Reorganization of the PDN Routing Working Group to include other public networks besides X.25 (e.g., SMDS).
- Discussions within the Network Joint Monitoring Working Group for common monitoring and report formats. We also discussed how the IETF Operations area could most effectively be organized and utilized (see Operations Area report).
- Announcement of the IAB recommendation, and the FNC agreement, to eliminate the notion of "connected status" for NIC assigned network numbers. With the FNC accepting this recommendation, MERIT will anounce how this will affect their policy for registering networks in their policy routing database.

Reminder – Next IETF Meeting

The Nineteenth IETF meeting will be December 3-7, 1990 in Boulder, Colorado. The meeting will be jointly sponsored by Westnet and NCAR. Carol Ward (University of Colorado) and Don Morris (NCAR) are the local hosts.

A very interesting technical agenda on high performance transport protocols is beginning to take shape.

Trial Modifications to the IETF Meeting Format

The attendance at IETF meetings has increased from under 100 to nearly 300 in the last 2 years. During this period the number of Working Groups has increased from 12 to over 45. Although I generally take this growth as a positive sign of success, it should also alert us to look for ways to make sure this new activity is integrated smoothly into the existing IETF structure.

In this section, I report some trial modifications for the Boulder IETF meeting format, based on suggestions by attendees. In the next section, I note ways for Working Group Chairs and attendees to keep meetings highly productive.

In the FSU IETF plenary in February 1990, we decided to reduce the number of IETF meetings from four to three per year. At the next two IETF plenaries (Pittsburgh/May 1990, Vancouver/August 1990), we decided to expand the current 3.5 day meeting format to 4.5 days on a trial basis.

Together with the new 4.5 day format, we have decided to include several other suggestions in the trial. We have divided each full day into 3 periods, resulting in 13 periods total. Eight of these will be Working Group sessions, three periods will be devoted to technical presentations, one for an open plenary/IESG session, and one for a reporting session on Friday morning. Total time for Working Groups is increased by a third, while time for other regular IETF features (e.g., technical presentations, reports, IESG) remain about the same.

The trial format for the Boulder meeting will look like:

	Mon	Tues	Weds	Thu	Fri
9:00-12:00	WG	WG	WG	WG	Reports
1:30-3:30	WG	WG	WG	Tech	
4:00-6:00	WG	Tech	Tech	IESG	

Early Registration will be on Sunday evening.

Again in response to suggestions, we will offer more technical presentations by incorporating some within Working Group sessions. For example, we will move the network status reports into the NJM or TEWG Working Group sessions.

We will be looking forward to comments on these new features.

Actions to Encourage Working Group Productivity

As a reminder to Working Group Chairs and attendees, there are some specific actions that can be done to help make Working Group meetings more productive.

Working Group Chair Actions:

1. Working Group Chairs are asked to provide "charters" and meeting reports, both of which are openly available online and in IETF Proceedings. The purpose of charters and reports is to help prospective attendees understand the

objectives and status of the groups, so that they can come to meetings prepared.

2. Chairs can further assist prospective attendees in preparing for each Working Group meeting by providing an agenda and document reading list. Not only will this help attendees prepare for meetings, but having an explicit agenda helps the Working Group focus the meeting and keep it on track.

<u>Working Group Attendee Actions:</u>

IETF Working Group meetings are technical *working* sessions. Active, informed, constructive participation is welcomed and encouraged. Observers are also welcome. Working Group meetings are generally fully open (although some sessions may be open only to document reviewers).

To get the most out of Working Group attendance (for yourself and for the group), attendees should come to meetings with a good understanding of the Working Group background and progress-to-date.

Attendees can become familiar with the current status and progress of Working Groups in several ways.

1. Objectives and notes from previous meetings are available online. For retrieval instructions (send to `ietf-manager@nri.reston.va.us`).
2. Objectives and notes from previous meetings are also reproduced in the hardcopy Proceedings (to order Proceedings, send to `proceedings@nri.reston.va.us`).
3. Agendas and reading lists for Working Group meetings will also be posted to the respective Working Group mailing lists.

IESG Standards Management

The IAB and IETF were founded, and continue to function, as technical development groups for Internet networking technology. Out of necessity over the past several years, the IAB and IETF have evolved a standards-making component to more rigourously define the protocols and procedures used in the Internet. Although this standards process is now reasonably well defined (see RFC1140), there are certain aspects of the process (and the procedures to implement the process) that are still "ad hoc". This is particularly true in the way that the IESG treats new work and the way that IESG makes recommendations to the IAB regarding standards actions.

In order to develop clearer IETF/IESG standards procedures, I have asked Dave Crocker (DEC) to establish a new IETF Standards Management position on the IESG. The specific charter of this new position will be to:

1. Write down new and existing IESG standards practices in a "IETF Standards

Practices Handbook".

2. Propose new or amended practices, where needed, to fill out a fully-developed IETF/IESG standards practice (up through the the recommendation to the IAB).
3. Act as coordinator to help move specific protocols through the IETF/IESG standards process.
4. Act as the liaison between the IESG and IAB on standards activity. This might include developing new general procedures for IAB/IESG interaction, helping to conduct "Technical Reviews" when needed, or generally tracking IESG recommendations through the IAB.

I envision that much responsibility will still belong to each Area Director for specific standards actions (e.g., primary responsiblity for advancing work to the IESG from IETF Working Groups, providing "Technical Summaries", etc.). Greg Vaudreuil (CNRI), as IESG Secretary, will continue to act as the agent for most of the specific actions (e.g., formulating the actual IESG recommendations and forwarding to the IAB, etc.).

However, in addition to recommending and codifying the standard practices, the new IESG Standards Manager will act as a backstop to make sure the process is followed in an expeditious manner, and nothing gets lost in the cracks.

In order to provide adequate focus on this new activity, Dave Crocker will give over his role as Network Management Area Director to the newly organized IESG Network Management Directorate. This is the subject of the next topic.

New IESG Network Management Directorate

One of Dave Crocker's goals as NM AD has been to form a NM "Review Board". The goal of this board would be to provide broad community perspective and input to IETF network management development decisions. For example, such a board would perform the key role of guiding and reviewing Internet MIB development activity. With Dave's moving from the NM area to a new Standards management role, it became clear that the time for forming this group had arrived.

We are now forming an "IETF Network Management Directorate". The NM Directorate will be a composed of approximately 9 persons. Its Chair will also serve as the IESG NM Area Director.

I am very pleased to announce that Chuck Davin (MIT) will be able to serve as the new NM AD and chair of the NM Directorate. He and I, after consultation with the IESG and other participants in the NM area, hope to be able to announce the complete membership of the new NM directorate at the next IETF meeting.

IETF Standards Procedures

The IESG is called upon to make recommendations to the IAB on Internet standards activity. The most common example is when an IETF Working Group wishes to submit a protocol document to the IAB for standardization. In such a case, the Working Group Chair forwards the protocol document to the IESG via the relevant Area Director. The IESG then forwards a recommendation to the IAB (usually after open discussion at an IETF meeting). All IESG recommendations to the IAB are cc'ed to the IETF mailing list.

In the future, the IESG will furnish a "Technical Summary" as part of all standards recommendation packages. A "Technical Summary" will include a brief overview of the document, and explain the motivation for the particular technical approach taken.

"Technical Summaries" are different from "Technical Reviews". The more concise "Technical Summary" is meant to give a brief overview of the main technical points, and will become a routine part of all future IESG recommendations to the IAB. "Technical Summaries" will be provided by the document author or Area Director.

"Technical Reviews" would generally be a more thorough, but less frequent, review conducted by a separate group drawn together by the IESG and/or by the relevant Area Director. So far, there have only been "Reviews" for CMOT, BGP, (and less formally) PPP and MTU Discovery.

To summarize (and, hopefully, to help clarify our evolving process):

- The IESG makes recommendations to the IAB on Internet standards actions. These IESG recommendations are usually formulated after discussion at an open plenary session of the IETF. The final recommendation is always cc'ed to the IETF mailing list.
- In the future, the IESG will include a "Technical Summary" as part of the recommendation package. The "Summary" will be provided by the document author or the appropriate AD.
- A wider, more comprehensive "Technical Review" may also be requested by either the IESG or IAB, but this is expected to be a less frequent occurance.

Phill Gross
IETF Chair

Final Agenda of the Eighteenth IETF

(July 31-August 3, 1990)

TUESDAY, July 31

9:00-9:15 am	"Introduction to the Privacy and Security Research Group" (Steve Kent/BBN) "Introduction to the Privacy Enhanced Mail Demonstration" (James Galvin/TIS)
9:15-12:00 noon	Morning Working Group Sessions • Privacy and Security Research Group - Open Meeting (DEC Distributed Systems Security Architecture (DSSA)) • SMNP, Transmission MIB, and Bridge MIB (Marshall Rose/PSI, John Cook/Chipcom and Fred Baker/Vitalink) • Router Discovery (Steve Deering/Xerox PARC) • Interconnectivity (Guy Almes/Rice) • User Services (Joyce K. Reynolds/ISI) • IP over SMDS (Mike Fidler/OSU and George Clapp/Ameritech) • Connection IP (Claudio Topolcic/BBN) • OSI General (Ross Callon/DEC and Rob Hagens/UWisc) • Network Printing Protocol (Leo McLaughlin/Wollongong)
1:00-4:00 pm	Afternoon Working Group Sessions • Joint SNMP Authentication (Jeff Schiller/MIT) and PSRG (Steve Kent/BBN) • IP over FDDI (Dave Katz/Merit) • Telnet (Dave Borman/Cray Research) • Multicast OSPF (Steve Deering/Xerox PARC) • Router Requirements (Philip Almquist/Stanford and Jim Forster/cisco) • IP over SMDS (Mike Fidler/OSU and George Clapp/Ameritech) • Connection IP (Claudio Topolcic/BBN) • Remote Lan Monitoring (Mike Erlinger/Micro Technology) • Network Joint Management (Phill Gross/CNRI)

4:15-5:45 pm Network Status Briefings

- "ESnet" (Tony Hain/LLNL)
- "Nasa Sciences Internet"
- "Mailbridge Report" (Zbigniew Opalka/BBN)
- "CA*net" (Dennis Ferguson/UToronto)
- "NSFnet" (Dale Johnson/Merit)

WEDNESDAY, August 1

9:15-12:00 noon — Morning Working Group Sessions

- Privacy and Security Research Group - Open Meeting (Privacy Enhanced Mail (PEM))
- Special Host Requirements (Bob Stewart/Xyplex)
- Management Services Interface (Oscar Newkerk/DEC)
- Topology Engineering (Scott Brim/Cornell)
- LAN Manager MIB (Dave Perkins/3Com)
- Call Accounting (Cyndi Mills/BBN)
- Site Security Policy Handbook (Joyce K. Reynolds/ISI and Paul Holbrook/CERT)
- Connection IP (Claudio Topolcic/BBN)
- Network Fax (Mark Needleman/UC Berkeley)
- IS-IS Routing (Ross Callon/DEC)

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

- Privacy and Security Research Group - Members Only
- Dynamic Host Configuration (Ralph Droms/Bucknell)
- FDDI MIB (Jeff Case/UTenn)
- Network Information Services Infrastructure (Dana Sitzler/Merit)
- Security Policy (Richard Pethia/CERT)
- Router Requirements (Philip Almquist/Stanford and Jim Forster/cisco)
- IP over SMDS (Mike Fidler/OSU and George Clapp/Ameritech)
- Connection IP (Claudio Topolcic/BBN)
- OSI NSAP Assignment (Richard Colella/NIST)
- DDN Interconnectivity (Zbigniew Opalka and Kathy Huber/BBN)

4:15-5:30 pm — IETF Protocol and Technical Presentations

- CMIP over TCP (Brian Handspicker/DEC)
- ACTS Satellite (Thomas vonDeak/NASA)

THURSDAY, August 2

9:15-12:00 noon — Morning Working Group Sessions

- Joint Security Policy, Site Security Handbook and PSRG
- IP over Appletalk (John Veizades/Apple)
- Point-to-Point Protocol Extentions (Stev Knowles/FTP)
- Call Accounting (Cyndi Mills/BBN)
- User Connectivity (Dan Long/BBN)
- Benchmarking Methodology (Scott Bradner/Harvard)
- DecNet IV MIB (Jon Saperia/DEC)
- OSI Internet Management (Lee LaBarre/Mitre and Brian Handspicker/DEC)
- Character MIB (Bob Stewart/Xyplex)
- Connection IP (Claudio Topolcic/BBN)
- PDN Routing (Carl-Herbert Rokitansky/Fern University of Hagen)
- OSI X.400 (Rob Hagens/UWisc)

1:00-4:15 pm — IETF Technical Presentations

- Engineering the CREN (Mike Roberts and Mike Hrybyk/Educom)
- Scaling and Policy Using Multiple Hierarchical Addresses (Paul Tsuchiya/Bellcore)
- IMAP Services (Mark Crispin/UWashington)
- Berkeley TCP Evolution from 4.3-Tahoe to 4.3-Reno (Van Jacobson/LLNL)
- Perspectives on Research Networks in Europe (Erik Huizer and Rudiger Volk/Rare,Ripe)

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

FRIDAY, August 3

9:00-11:30 am	Working Group Area and Selected Working Group Presentations • Host and User Services Area (Craig Partridge/BBN and 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 Interoperability Area (Ross Callon/DEC and Rob Hagens/UWisc) • Operations Area (Interim - Phill Gross/CNRI) • Network Management Area (Dave Crocker/DEC)
11:30-12:00 noon	Concluding Remarks (Phill Gross/CNRI)
12:15 pm	Adjourn

Chapter 1

IETF Overview

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

The IETF mission includes:

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

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

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

IETF and IESG Chair:	Phill Gross/CNRI
Applications:	Russ Hobby/UC-Davis
Host and User Services:	Craig Partridge/BBN
Internet Services:	Noel Chiappa/Consultant
Routing:	Robert Hinden/BBN
Network Management:	James Davin/ MIT
OSI Integration:	Rob Hagens/U-Wisc and Ross Callon/DEC
Operations:	Phill Gross/CNRI (interim)
Security:	Steve Crocker/TIS
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 quarterly plenary sessions composed of Working Group sessions, technical presentations and network status briefings. The meetings are currently three 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 mail is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 On Line IETF Information

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

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

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

1.1.1 The IETF Directory

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

Files prefixed with a 0 contain information about upcoming meetings. Files prefixed with a 1 contain general information about the IETF, the Working Groups, and the Internet Drafts.

FILE NAME	
0mtg-agenda	the current agenda for the upcoming quarterly IETF plenary, which contains what Working Groups will be meeting and at what times, and the technical presentations and network status reports to be given.
0mtg-logistics	the announcement for the upcoming quarterly IETF plenary, which contains specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and travel directions.
0mtg-rsvp	a standardized RSVP form to be used to notify the support staff of your plans to attend the upcoming IETF meeting.
0mtg-schedule	current and future meeting dates and sites for IETF plenaries.
1id-abstracts	the Internet Drafts currently on-line in the Internet-Drafts directory.
1id-guidelines	instructions for authors of Internet Drafts.
1ietf-overview	a short description of the IETF, the IESG and how to participate.
1wg-summary	a listing of all current Working Groups, the Working Group 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.

1.1.2 The Internet-Drafts Directory

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

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

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

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

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

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

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

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

1.1.3 Directory Locations

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

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

cd internet-drafts
cd ietf

Individual files can then be retrieved using the GET command:

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

NSF Network Service Center Address: nnsc.nsf.net

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

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

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

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

Pacific Rim Address: munnari.oz.au

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

Europe Address: nic.nordu.net (192.36.148.17)

1.2 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, but we strongly encourage a submission of a matching ascii version (even if figures must be deleted) for readers without postscript printers and for online searches.

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

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

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

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

1.3 IETF Working Group Summary (by Area)

Applications

Russ Hobby
rdhobby@ucdavis.edu

Domain Name System (dns)
Chair(s): Philip Almquist almquist@jessica.stanford.edu
WG mail: namedroppers@nic.ddn.mil
To Join: namedropped-request@nic.ddn.mil
Status: continuing

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

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

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

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

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

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

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

Internet Draft: "Telnet Linemode Option", 08/08/1990, Dave Borman <draft-ietf-telnet-linemodeoption-02.txt>

Host and User Services

Craig Partridge
craig@nnsc.nsf.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	
Status:	continuing	

Dynamic Host Configuration (dhc)

Chair(s):	Ralph Droms	droms@sol.bucknell.edu
WG mail:	host-conf@sol.bucknell.edu	
To Join:	host-conf-request@sol.bucknell.edu	
Status:	continuing	

Internet User Population (iup)

Chair(s):	Craig Partridge	craig@nnsc.nsf.net
WG mail:	ietf@venera.isi.edu	
To Join:	ietf-request@venera.isi.edu	
Status:	continuing	

Network Information Services Infrastructure (nisi)

Chair(s):	Dana Sitzler	dds@merit.edu
WG mail:	nisi@merit.edu	
To Join:	nisi-request@merit.edu	
Status:	continuing	

Special Host Requirements (shr)

Chair(s):	Bob Stewart	rlstewart@eng.xyplex.com
WG mail:	ietf-hosts@nnsc.nsf.net	
To Join:	ietf-hosts-request@nnsc.nsf.net	
Status:	new	

User Connectivity (ucp)

Chair(s): Dan Long `long@bbn.com`
WG mail: `ucp@nic.near.net`
To Join: `ucp-request@nic.near.net`
Status: continuing

User Services (uswg)

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

Internet Services

Noel Chiappa
jnc@ptt.lcs.mit.edu

Connection IP (cip)

Chair(s): Claudio Topolcic topolcic@bbn.com
WG mail: cip@bbn.com
To Join: cip-request@bbn.com
Status: continuing

Internet Draft: "Internet Stream Protocol", 09/04/1990, C Topolcic <draft-ietf-cip-st2-00.txt>

IP MTU Discovery (mtudisc)

Chair(s): Jeff Mogul mogul@decwrl.dec.com
WG mail: mtudwg@decwrl.dec.com
To Join: mtudwg-request@decwrl.dec.com
Status: continuing

IP over Appletalk (appleip)

Chair(s): John Veizades veizades@apple.com
WG mail: apple-ip@apple.com
To Join: apple-ip-request@apple.com
Status: continuing

IP over FDDI (fddi)

Chair(s): Dave Katz dkatz@merit.edu
WG mail: FDDI@merit.edu
To Join: FDDI-request@merit.edu
Status: continuing

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

IP over Switched Megabit Data Service (smds)

Chair(s):	George Clapp	meritec!clapp@bellcore.bellcore.com
	Michael Fidler	ts0026@ohstvma.ircc.ohio-state.edu
WG mail:	smds@nri.reston.va.us	
To Join:	smds-request@nri.reston.va.us	
Status:	continuing	

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

Point-to-Point Protocol Extentions (pppext)

Chair(s):	Stev Knowles	stev@ftp.com
WG mail:	ietf-ppp@ucdavis.edu	
To Join:	ietf-ppp-request@ucdavis.edu	
Status:	continuing	

Router Discovery (rdisc)

Chair(s):	Steve Deering	deering@pescadero.stanford.edu
WG mail:	gw-discovery@gregorio.stanford.edu	
To Join:	gw-discovery-request@gregorio.stanford.edu	
Status:	continuing	

Router Requirements (rreq)

Chair(s):	James Forster	forster@cisco.com
	Philip Almquist	almquist@jessica.stanford.edu
WG mail:	ietf-rreq@Jessica.Stanford.edu	
To Join:	ietf-rreq-request@Jessica.Stanford.edu	
Status:	continuing	

Internet Draft: "Requirements for Internet IP Routers", 09/17/1990, Philip Almquist <draft-ietf-rreq-iprouters-00.txt>

Network Management

Dave Crocker
dcrocker@nsl.dec.com

Alert Management (alertman)

Chair(s): Louis Steinberg louiss@ibm.com
WG mail: alert-man@merit.edu
To Join: alert-man-request@merit.edu
Status: continuing

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

Bridge MIB (bridge)

Chair(s): Fred Baker baker@vitalink.com
WG mail: bridge-mib@nsl.dec.com
To Join: bridge-mib-request@nsl.dec.com
Status: new

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
Status: new

DECnet Phase IV MIB (decnetiv)

Chair(s): Jonathan Saperia saperia%tcpjon@decwrl.dec.com
WG mail: phiv-mib@jove.pa.dec.com
To Join: phiv-mib-request@jove.pa.dec.com
Status: continuing

FDDI MIB (fddimib)

Chair(s): Jeffrey Case case@utkux1.utk.edu
WG mail:
To Join:
Status: new

Internet Accounting (acct)

Chair(s): Cyndi Mills `cmills@bbn.com`
WG mail: `accounting-wg@bbn.com`
To Join: `accounting-wg-request@bbn.com`
Status: continuing

LAN Manager (lanman)

Chair(s): David Perkins `dave_perkins@3com.com`
WG mail: `lanmanwg@cnd.hp.com`
To Join: `lanmanwg-request@cnd.hp.com`
Status: continuing

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

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

Management Services Interface (msi)

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

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

OSI Internet Management (oim)

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`
Status: continuing

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

Internet Draft: "OSI Internet Management: Management Information Base", 08/17/1990, Lee LaBarre <draft-ietf-oim-mib2-02.txt>

Remote LAN Monitoring (rlanmib)

Chair(s): Mike Erlinger `mike@mti.com`
WG mail: `rlanmib@decwrl.dec.com`
To Join: `rlanmib-request@decwrl.dec.com`
Status: new

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`
Status: continuing

Internet Draft: "Definitions of Managed Objects for the T1 Carrier Interface Type", 04/23/1990, C Kolb, Fred Baker <draft-ietf-snmp-t1mib-01.txt>

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

Internet Draft: "Towards Concise MIB Definitions", 09/05/1990, Marshall Rose, Keith McCloghrie <draft-ietf-snmp-mibdefinitions-01.txt>

Internet Draft: "A Convention for Defining Traps for use with the SNMP", 09/05/1990, Marshall Rose <draft-ietf-snmp-traps-01.txt>

Internet Draft: "Extensions to the Generic-Interface MIB", 09/12/1990, Keith McCloghrie <draft-ietf-snmp-interfacemibext-00.txt>

Internet Draft: "IEEE 802.4 Token Bus MIB", 09/26/1990, Keith McCloghrie, Richard Fox <draft-ietf-snmp-tokenbusmib-00.txt>

Internet Draft: "Definitions of Managed Objects for the Ethernet-like Interface Types", 09/26/1990, John Cook <draft-ietf-snmp-ethernetmib-00.txt>

Internet Draft: "IEEE 802.5 Token Ring MIB", 09/26/1990, Keith McCloghrie, Richard Fox, Eric Decker <draft-ietf-snmp-tokenringmib-00.txt>

Transmission Mib (transmib)

Chair(s): John Cook cook@chipcom.com
WG mail: unknown
To Join: unknown
Status: continuing

OSI Integration

Ross Callon
callon@bigfut.enet.dec.com
Rob Hagens
hagens@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
Status: continuing

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

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
Status: continuing

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
Status: continuing

OSI X.500 (osix500)
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
Status: new

Operations

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

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

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

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

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

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

Routing

Bob Hinden
hinden@bbn.com

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
Status: continuing

Internet Draft: "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", 08/27/1990, Ross Callon <draft-ietf-isis-spec-01.ps>

Interconnectivity (iwg)

Chair(s): Guy Almes almes@rice.edu
WG mail: iwg@rice.edu
To Join: iwg-request@rice.edu
Status: continuing

Internet Draft: "Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2)", 07/17/1990, Steven Willis, John Burruss <draft-ietf-iwg-bgp-mib-01.txt>

Multicast Extentions to OSPF (mospf)

Chair(s): Steve Deering deering@pescadero.stanford.edu
WG mail: mospf@devvax.tn.cornell.edu
To Join: mospf-request@devvax.tn.cornell.edu
Status: continuing

Open Systems Routing (orwg)

Chair(s): Martha Steenstrup msteenst@bbn.com
WG mail: open-rout-interest@bbn.com
To Join: open-rout-request@bbn.com
Status: continuing

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

Routing and Address Resolution over X.25 and SMDS (pdnarp)

Chair(s): George Clapp `meritec!clapp@bellcore.bellcore.com`

WG mail:

To Join:

Status: new

Security

Steve Crocker
crocker@tis.com

IP Authentication (ipauth)

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

Internet Security Policy (spwg)

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

SNMP Authentication (snmpauth)

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

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

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

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

Site Security Policy Handbook (ssphwg)

Chair(s):	J. Paul Holbrook	`ph@sei.cmu.edu`
	Joyce K. Reynolds	`jkrey@venera.isi.edu`
WG mail:	`ssphwg@cert.sei.cmu.edu`	
To Join:	`ssphwg-request@cert.sei.cmu.edu`	
Status:	continuing	

1.4 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.

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

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

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

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

"Internet Cluster Addressing Scheme", Carl-Herbert Rokitansky, 11/01/1989

<draft-ietf-pdn-clusterscheme-00.txt>

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

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

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

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

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

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

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

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

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

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

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

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

This document specifies an addressing convention to be used in conjunction with other protocols.

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

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

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

We present an architecture for policy routing among administrative domains within the Internet. The objective of inter-domain policy routing is to synthesize and maintain routes between source and destination administrative domains, providing user traffic with the

requested service within the constraints stipulated by the administrative domains transited. The architecture is designed to accommodate an Internet with tens of thousands of administrative domains.

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

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

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

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

"Definitions of Managed Objects for the T1 Carrier Interface Type", C Kolb, Fred Baker, 09/26/1990
<draft-ietf-snmp-t1mib-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 T1-carrier objects.

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

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

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

"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.

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

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

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

This memo is a product of the IETF Lan Manager MIB Working Group. It defines management objects to support the translation of LAN Manager alerts to SNMP traps. It is a companion document to Management Information Base for LAN Manager Management, which defines a base set of management objects for LAN Manager.

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

This memo provides a Management Information Base (MIB) for management of LAN Manager nodes with TCP/IP-based network management protocols. Together with documents describing the structure of management information (RFC 1155) and the Simple Network Management Protocol (RFC 1157) this document provides a specification for managing LAN Manager nodes in a TCP/IP environment.

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

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

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

This memo defines an experimental portion of the Management

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

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

Simple Network Management Protocol (SNMP) specification allows for the authentication of management operations by a variety of authentication algorithms. This memo defines two strategies for administering SNMP communities based upon either the SNMP authentication algorithm or the SNMP authentication and privacy algorithm. Insofar as the administration of SNMP communities based upon the trivial authentication algorithm may be realized by straightforward application of familiar network management techniques, administration of such communities is not directly addressed in this memo.

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

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

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

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

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

This memo discusses and defines a number of terms that are used in describing performance benchmarking tests and the results of such tests.

The terms defined in this memo will be used in additional memos to define specific benchmarking tests and the suggested format to be used in reporting the results of each of the tests.

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

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

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

"Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2)", Steven Willis, John Burruss, 09/21/1990 <draft-ietf-iwg-bgp-mib-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 the Border Gateway Protocol.

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

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

"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.

"The OSPF Specification, Version 2", John Moy, 07/24/1990
<draft-ietf-ospf-ospf2-00.txt>

This document is a specification of the Open Shortest Path First (OSPF) internet routing protocol. OSPF is classified as an Internal Gateway Protocol (IGP). This means that it distributes routing information between routers belonging to a single Autonomous System. The OSPF protocol is based on SPF or link-state technology. This is a departure from the Bellman-Ford base used by traditional internet routing protocols.

"X.25 Call Setup and Charging Determination Protocol (XCDP)", Carl-H Rokitansky, 07/27/1990
<draft-ietf-pdnrout-x25call-00.txt>

Therefore, the X.25 Call Setup and Charging Determination Protocol (XCDP)", described in this document, has been developed, to support global Internet connectivity via the system of X.25 Public Data Networks PDN (even via VAN-gateways preventing local charges), by providing a pseudo-reverse charging option, which is indicated in the Call User Data (CUD) field of the call request. In addition, information about the source and destination address of the Internet datagram to be transmitted, can also be indicated in the user data field of the call request.

"X.121 Address Resolution for IP Datagram Transmission Over X.25 Networks", Carl-Herbert Rokitansky, 07/27/1990
<draft-ietf-pdn-xarp-01.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>

"Telnet Linemode Option", Dave Borman, 08/08/1990
<draft-ietf-telnet-linemodeoption-02.txt>

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

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

"OSI Internet Management: Management Information Base", Lee LaBarre, 08/17/1990 <draft-ietf-oim-mib2-02.txt>

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

"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.

"Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", Ross Callon, 08/27/1990
<draft-ietf-isis-spec-01.ps>

This Internet Draft specifies an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments and dual environments. This specification was developed by the IS-IS Working Group of the Internet Engineering Task Force. Comments should be sent to isis@merit.edu.

"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.

"The Finger User Information Protocol", David Zimmerman, 09/04/1990
<draft-zimmerman-finger-03.txt>

The predecessor to this memo was RFC742, a description of the original Finger protocol. Currently, the development and use of the protocol has deviated from the imprecise RFC742 specifications, and the examples in RFC742 are woefully out of date.

"Internet Stream Protocol", C Topolcic, 09/04/1990
<draft-ietf-cip-st2-00.txt>

This memo defines the Internet Stream Protocol, Version 2 (ST-II), an IP-layer protocol which provides end-to-end guaranteed service across an internet. This specification obsoletes IEN 119 "ST - A Proposed Internet Stream Protocol" written by Jim Forgie in 1979, the previous specification of ST. ST-II represents some relatively minor changes to Version 1 of the protocol and is intended to fill in some of the areas left unaddressed, to make it easier to implement, and to support a wider range of applications. However, ST-II is not compatible with the previous version of ST.

"Towards Concise MIB Definitions", Marshall Rose, Keith McCloghrie, 09/26/1990
<draft-ietf-snmp-mibdefinitions-01.txt>

This memo describes a straight-forward approach toward producing concise, yet descriptive, MIB modules. Use of this approach is in every way fully consistent with the Internet-standard network management framework.

"A Convention for Defining Traps for use with the SNMP", Marshall Rose, 09/26/1990
<draft-ietf-snmp-traps-01.txt>

This memo describes a straight-forward approach toward defining traps used with the SNMP. It is specifically intended for use by the authors of experimental MIBs, and emphasizes a concise descriptive approach. Use of this approach is fully consistent with the Internet-standard network management framework.

"Experimental Definitions of Managed Objects for the Point-to-Point Protocol", Frank Kastenholz, 09/11/1990
<draft-ietf-ppp-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.

"Extensions to the Generic-Interface MIB", Keith McCloghrie, 09/12/1990
<draft-ietf-snmp-interfacemibext-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 managed object types as experimental extensions to the generic interfaces structure of MIB-II.

This memo does not specify a standard for the Internet community. However, after experimentation, if sufficient consensus is reached in the Internet community, then a subsequent revision of this document may be incorporated into the Internet-standard MIB.

"Transmission of IP Datagrams and ARP Packets over ARCNET Networks", Don Provan, 09/17/1990
<draft-provan-iparcnet-00.txt>

This draft document specifies a standard method of encapsulating Internet Protocol (IP) and Address Resolution Protocol (ARP) datagrams using the ARCNET Packet Header Definition Standard. This draft should obsolete RFC-1051. RFC-1051 used a different ARCNET packet header which is incompatible with most modern ARCNET software.

"Requirements for Internet IP Routers", Philip Almquist, 09/17/1990 <draft-ietf-rreq-iprouters-00.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.

"IEEE 802.4 Token Bus MIB", Keith McCloghrie, Richard Fox, 09/26/1990 <draft-ietf-snmp-tokenbusmib-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 managed objects used for managing subnetworks which use the IEEE 802.4 Token Bus technology.

"Definitions of Managed Objects for the Ethernet-like Interface Types", John Cook, 09/26/1990 <draft-ietf-snmp-ethernetmib-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 ethernet-like objects.

"IEEE 802.5 Token Ring MIB", Keith McCloghrie, Richard Fox, Eric Decker, 09/26/1990 <draft-ietf-snmp-tokenringmib-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 managed objects used for managing subnetworks which use the IEEE 802.5 Token Ring technology.

Chapter 2

Steering Group Report

2.1 Minutes of the August 2nd Meeting

The Internet Engineering Steering Group met during the open plenary on Thursday August 2nd. The topics of discussion included information on IETF growth and productivity, the IP address space problem, and network management.

2.1.1 Growth in IETF Participation and Activities

Phill Gross presented statistics on IETF attendance and Working Group progress over the last eighteen months.

Slide #1 (and the table below) shows attendance at the last seven IETF meetings. This appears to represent a steady growth from just over 100 to about 300 attendees in the last year and a half. This represents attendance by 500 different persons from 166 different organizations.

Attendance for the last six IETF meetings:

12th IETF	Jan 1989	University of Texas	121
13th IETF	Apr 1989	Kennedy Space Center	112
14th IETF	Jul 1989	Stanford University	215
15th IETF	Oct 1989	University of Hawaii	138
16th IETF	Feb 1990	Florida State Univ.	191
17th IETF	May 1990	PSC/SEI/CMU	243
18th IETF	Aug 1990	UBC, Vancouver	300

Repeat attendance by individuals is high enough to show a dedicated core of key IETF contributors (Slide # 2). There are 23 persons who have attended all of the previous 6 meetings, 28 who have attended 5 meetings, and 46 who have attended at least 4 of the last 6 IETF meetings. That is almost 100 persons who have attended at least 2/3 of the recent meetings. This is impressive when considered that attendance was only just over 100 for 3 of those meetings.

When the attendees are grouped by categories (Slide # 3) it shows that about 1/3 were from vendors, about 1/3 from government (DoD and civilian agencies), and over 1/4 from universities and regional network operators.

IETF Technical Progress and Contributions

A more important measure of IETF activity is the number of active Working Groups and the number of technical documents (e.g., RFCs) produced over the same period. Slide # 4 (also shown in the table below) shows both the total number of Working Groups, and the number which actually met at each of the last 6 IETF meetings. Notice that the number of Working Groups has shown a sharp increase since the creation of the IESG last fall. The IESG first met at the University of Hawaii in October 1989. Notice that after that meeting, the number of active Working Groups doubled.

Date	Location	Total WGs	# met
Jan 1989	University of Texas	12	12
Apr 1989	Kennedy Space Center	19	17
Jul 1989	Stanford University	20	18
Oct 1989	University of Hawaii	19	18
Feb 1990	Florida State Univ.	38	32
May 1990	PSC/SEI/CMU	40	33
Aug 1990	UBC, Vancouver	45	38

During this approximate period, there were over 80 RFCs published relating to Internet technical activities. Of those RFCs, around 30 pertained to Internet standards (Slide # 5). The IETF accounted for almost 30 percent of the total RFCs published and for almost 55 percent of all RFCs pertaining to standards. The IAB itself, together with the IRTF, accounted for almost another 30 percent, meaning that the IAB as an organization (i.e., including IETF and IRTF) accounted for almost 60 percent of all RFCs published in this period.

A very powerful conclusion can be drawn from these figures – the IETF has developed into a productive body for Internet technical development, and it is continuing to grow in positive ways.

IP number allocation and connected status

In the past, the NIC was delegated the responsibility for giving "connected status", upon request, when it assigned network numbers. The purpose of "connected status" was to show sponsorship from a federal agency to pass traffic across the federal backbones, In RFC 1174, in an effort to distribute responsibilities in the growing international Internet, and to recognize the growing role of non US-Federal networks in providing network service, the IAB recommended, and the FNC approved, that the notion of "connected status" be dropped. This has impact on the way that NSFnet registers and records its routing database.

Two presentations of the growth of network numbers were made in the IESG session. Zaw-Sing Su (SRI) showed the history of network number assignments by the NIC. His slides accompany this report. His numbers show a distinction between "connected" and "unconnected" networks. In the future, this distinction will not be made.

Sue Hares and Dale Johnson (MERIT) showed the growth of "configured" networks numbers in the NSFnet routing database. This indicates which networks have permission to send traffic across the NSFnet. Sue and Dale were instrumental in helping to define and explain these various "network number concepts", and how MERIT used these concepts in establishing its routing database.

Using this information, and information from BBN, Frank Solensky (Racal-Interlan), presented a statistical analysis on the rate of utilization of IP address space. He showed that the growth is exponential. See the accompanying slides for his projections when the IP address space become depleted (assuming continued exponential growth).

Van Jacobson pointed out that exponential growth cannot continue forever, and that we should look at the well-known "S" curve of finite population growth for a better model. In this model, the exponential growth slows, flattens out, and approaches the population limit asymptotically. That was a comforting observation. However, Frank's analysis seem to show that we have not yet left the exponential growth portion of the "S" curve.

2.1.2 Network Management

Dave Crocker led a discussion on Network Management issues. He began with a more general discussion on the problems of determining whether a specification is appropriate for standardization, indicating that it needs both technical review and sufficient constituency. According to Crocker, Standards have costs, and there must be a reason for the cost to be paid, a constituency. He was looking for "Market Research Focus Groups". Among possible sources for information, he identified the IETF Plenary as the most promising forum.

Topics for specific focus and feedback, during this Open IESG were:

- Lan Manager MIB
- MIB II
- Proxy Agent
- CMOT
- Alert Management

Lan Manager Mib

This is an example of a proprietary protocol which has had a public MIB specification done. Is it appropriate to make the MIB a public standard, even though the protocol managed by the MIB remains proprietary.

There was a feeling that this should be a private mib, not an internet standard. There is a region of the mib tree reserved specifically for this thing. Discussion continued on IETF standardization of vendor protocols.

MIB II

The minumum time in grade timer has expired, making it possible to consider promotion of MIB II to Draft status. Feedback on the document was solicited. Send comments on the draft to the `snmp-wg@nisc.nyser.net` mailing list.

Proxy Agents

There are two approaches to referencing proxy management agents:

- View MIB mechanism serves multiple purposes
  - MIB Access Control
  - MIB Variables "aliasing"
  - Proxy
  - Trap Destinations
- Community-string "source routing"
  - Addressing structure added

Model differences between the two approaches:

- Domain Name Vs. Source Routed
- Simplicity
  - Adequate?
  - Deceptive?

There was no closure.

CMOT

A new version of CMOT will be submitted for standards status. It contains a number of technical changes, and has received some implementation and testing experience. There is a question as to the status that should be assigned to it.

- Constituency
- Technical adequacy - Stability (It is a stable specification)
- IETF Preference?
  - Draft (Current state)
  - Proposed (sufficiently different as to start over)
  - Experimental (Not a dumping ground, place for further development)

It was felt that it would be a bad precedent to say that any significant revision of the specification would move the protocol back to the proposed standard level. Backward compatability was felt to be a criterion for demoting a protocol to proposed standard.

CMOT will run into the 2 year limitation for time in grade. A new version at the same level would reset the timer.

An informal poll was taken, asking those present what status they would prefer. About half wanted retention at Draft, one-fourth wanted Proposed and one-fourth wanted Experimental.

Alert Management

Skipped due to lack of time.

IETF Attendance (1/89-7/90)
Total Attendance per meeting

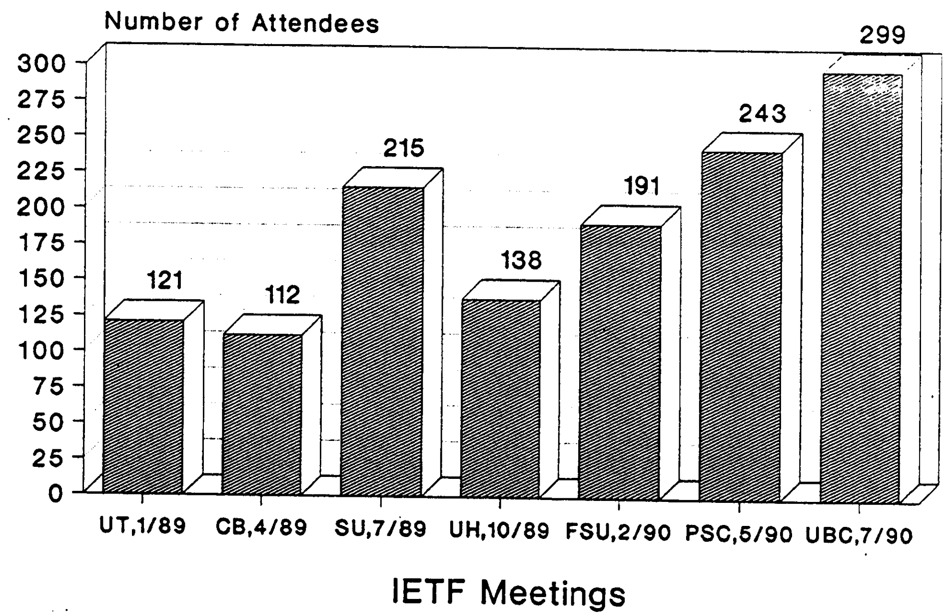

Phill Gross/CNRI

1.

IETF Attendance (1/89-5/90)

Repeat Attendees

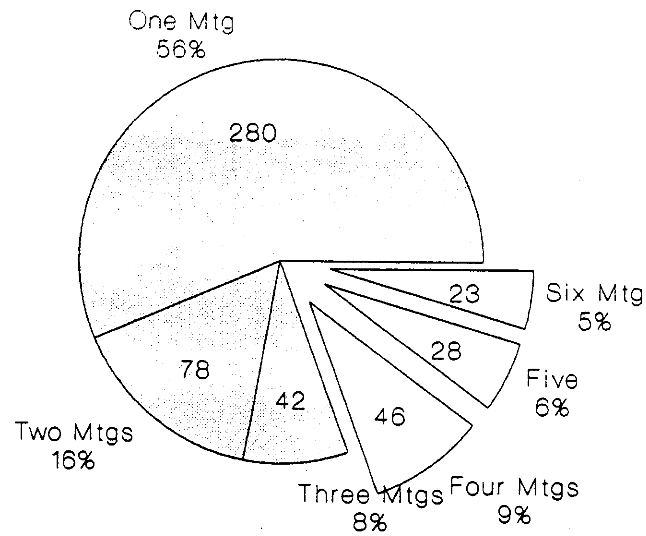

Phill Gross/ CNRI

2.

IETF Attendance (1/89-5/90)

Grouped by Category

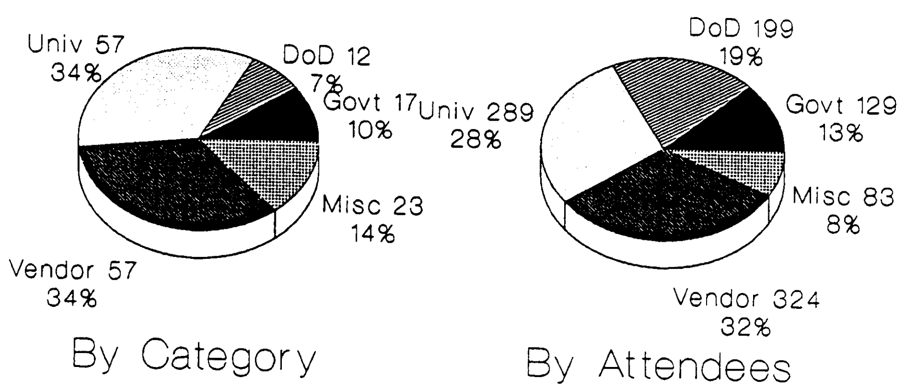

Phill Gross/CNRI

3.

IETF Working Groups
Total and Met (for 1/89-7/90)

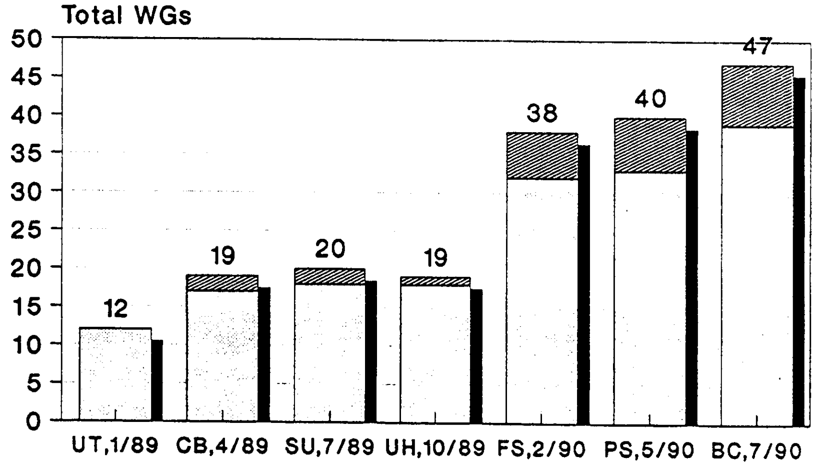

	UT,1/89	CB,4/89	SU,7/89	UH,10/89	FS,2/90	PS,5/90	BC,7/90
Series 2	0	2	2	1	6	7	8
Series 1	12	17	18	18	32	33	39

IETF Meetings

Phill Gross/CNRI

4.

Internet Evolution Activities
Sept 1988 - June 1990

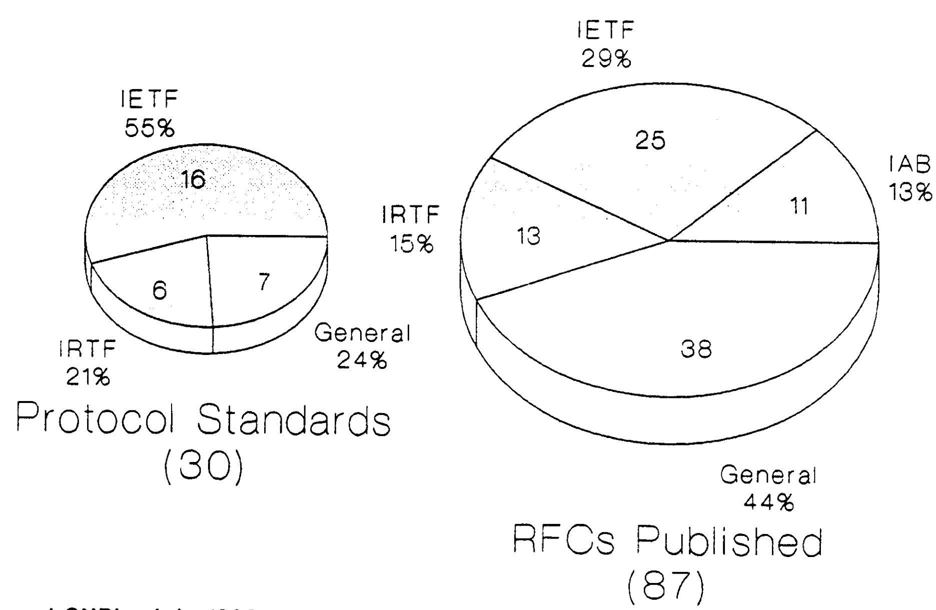

Phill Gross / CNRI July 1990

5.

Internet Growth (Continued):

Continued Internet Growth

Frank Solensky
Racal Interlan
solensky@interlan.com

- A preliminary analysis of data presented earlier in the conference projects the "size" of the Internet on several metrics assuming continued exponential growth.
  - NIC Assigned Network Numbers
  - NIC "Connected" Status Nets
  - BBN's snapshots
  - NSFnet Policy Routing DataBases
- As was mentioned during the discussion period, a logistic curve would likely be a more realistic model: this will be the subject of further analysis. Note, however, that the limit that this approaches *may* turn out to be beyond the capacity of the class A-B-C numbering scheme

NIC

"Connected" IP Network Numbers

- Assigned Numbers RFC defines connected networks as connected to research and operational internet.
- Does not reflect whether the net is, in fact, entered in any routing table

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan. 1983

	Class A	Class B	Class C	Class A-B
β	12.069	24.412	887.879	3032.211
α	.012163	.040721	.011690	.013467
growth rate per yr.	15.613%	61.440%	14.497%	17.413%
y	125	16,382	2,097,150	49,147
$\hat{x}$	192.193	159.839	664.438	206.846
	Jan 6, 1999	Apr 26, 1996	May 14, 2038	Mar 27, 2000
r	.9293	.9870	.7942	.9548

Assignment of IP Network Numbers

- Reflects organizations' desire for IP Address assignment; that is, to be listed in RFC-1166.
- Does not reflect "connectivity"

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan '83

	Class A	Class B	Class C	Class A-B
β	11.883	21.446	1531.793	2899.462
α	.013175	.049411	.027187	.015587
growth rate per yr.	17.007%	78.38%	37.973%	20.394%
y	125	16,382	2,097,150	49,147
$\hat{x}$	178.605	134.35	265.64	181.58
	(Nov 19, 1997)	(Mar. 11, 1994)	(Feb. 18, 2005)	(Feb 17, 1998)
r	.9491	.9842	.9800	.9749

Assigned Class "B" Network Numbers

$y = \beta e^{\alpha t}$ $\alpha = .049411$, $\beta = 21.446$

- Growing at 78.3%/year
- Doubles every 14.37 months

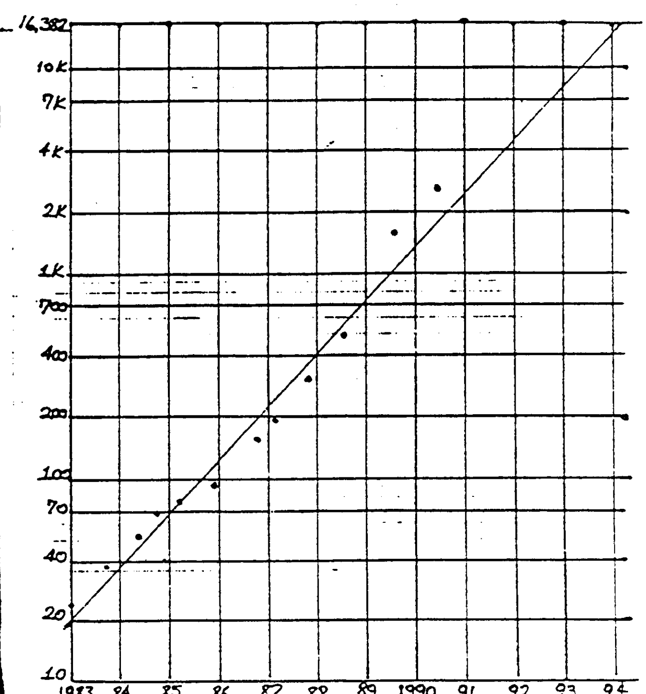

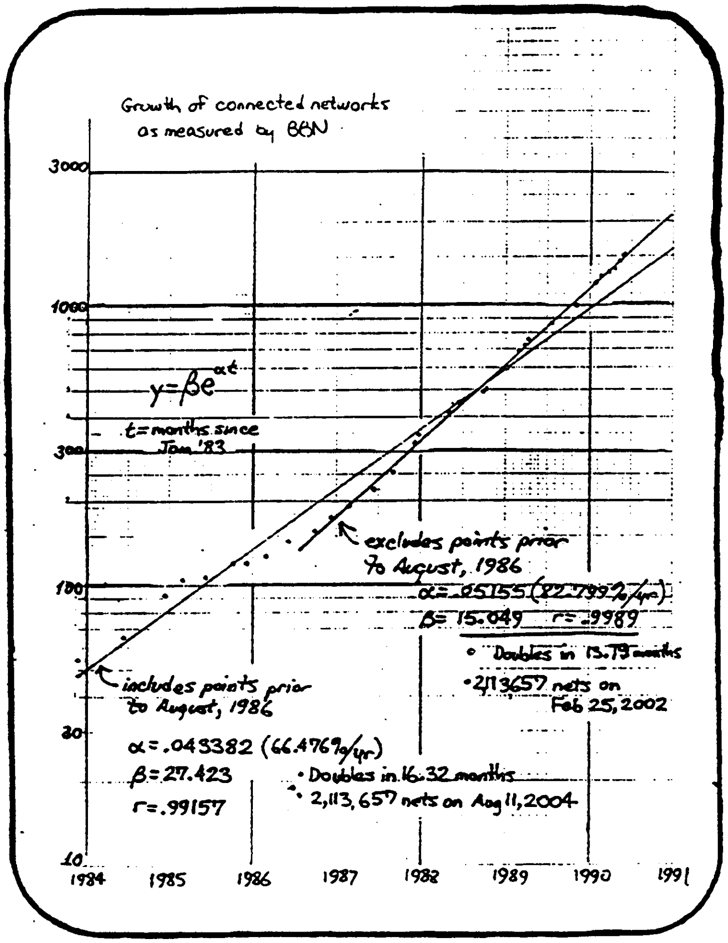
Growth of connected networks
as measured by BBN
$y=\beta e^{\alpha t}$
t= months since Jan '83
excludes points prior to August, 1986
α=.05155 (82.799%/yr)
β=15.049 r=.9989
• Doubles in 13.79 months
• 2,113,657 nets on Feb 25, 2002
includes points prior to August, 1986
α=.043382 (66.476%/yr)
β=27.423
r=.99157
• Doubles in 16.32 months
• 2,113,657 nets on Aug 11, 2004
3000
1000
300
100
30
10
1984
1985
1986
1987
1988
1989
1990
1991

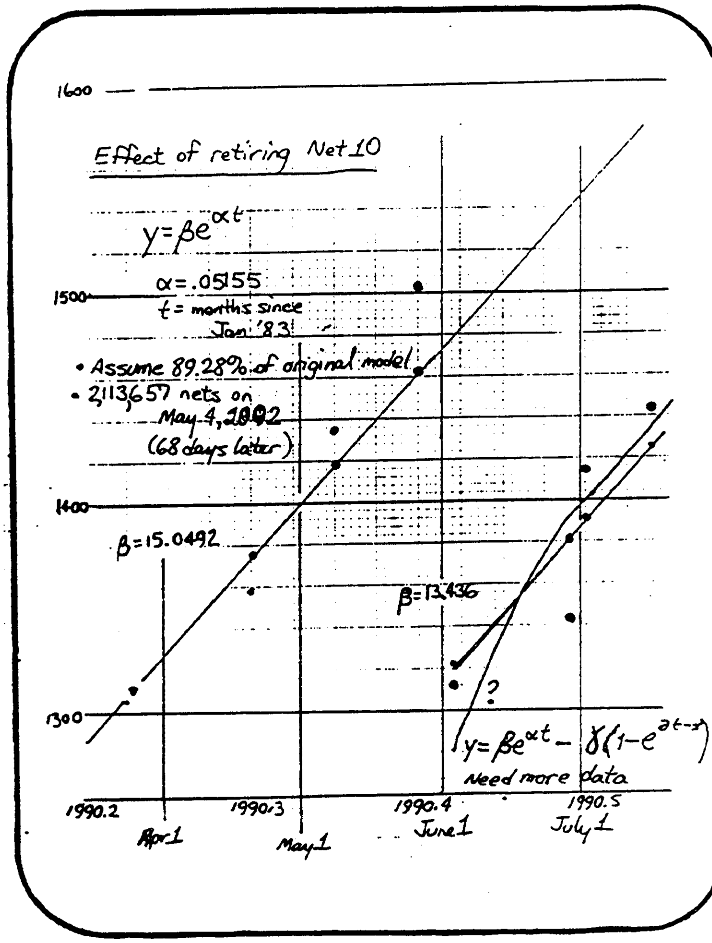
Effect of retiring Net 10
$y=\beta e^{\alpha t}$
α=.05155
t= months since Jan '83
• Assume 89.28% of original model
• 2,113,657 nets on May 4, 2002 (68 days later)
β=15.0492
β=13.436
?
$y=\beta e^{\alpha t}-\gamma(1-e^{\partial t-\tau})$
Need more data
1600
1500
1400
1300
1990.2
1990.3
1990.4
1990.5
Apr 1
May 1
June 1
July 1

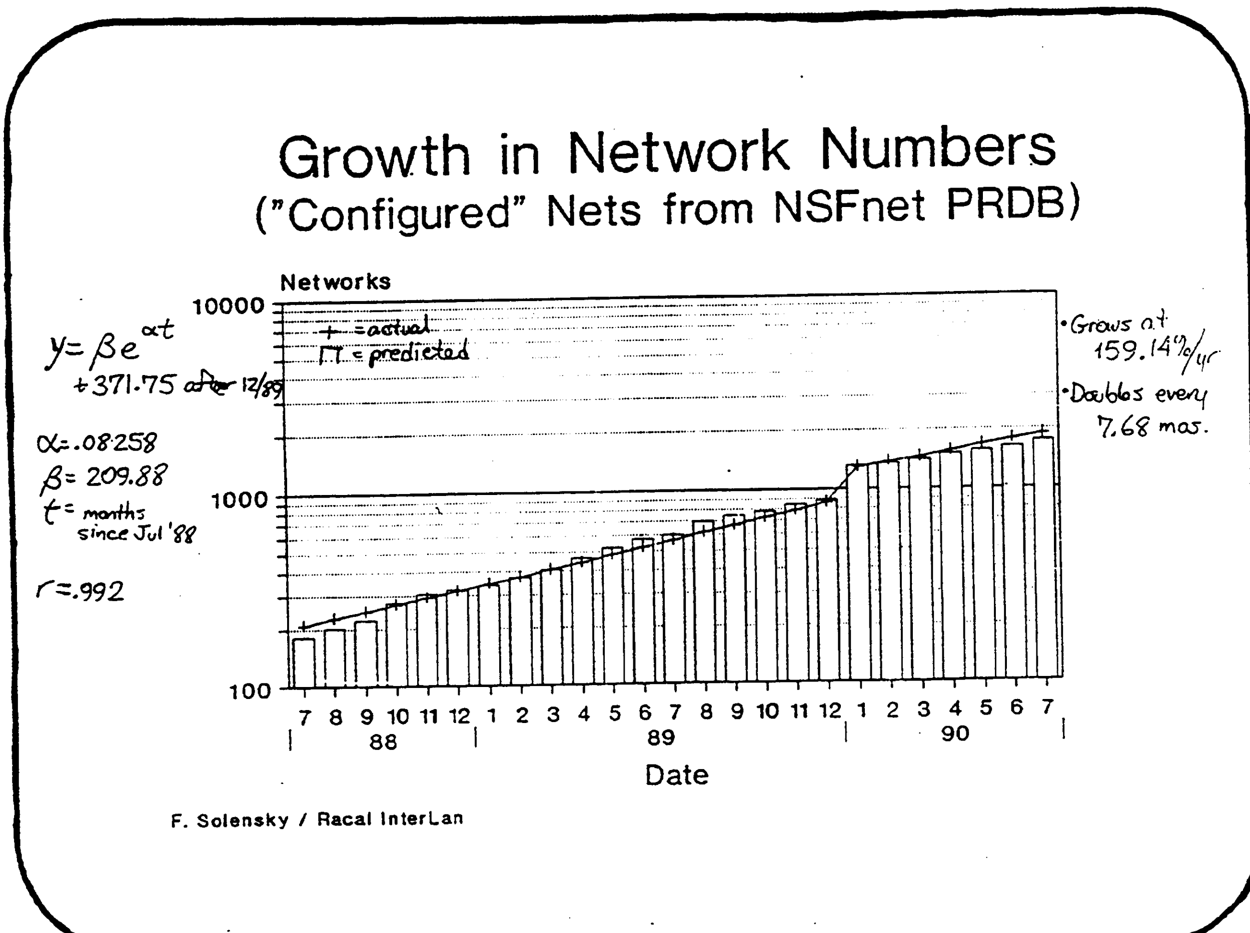
Growth in Network Numbers
("Configured" Nets from NSFnet PRDB)
Networks
+ = actual
⊓ = predicted
10000
1000
100
7 8 9 10 11 12 1 2 3 4 5 6 7 8 9 10 11 12 1 2 3 4 5 6 7
88
89
90
Date
$y=\beta e^{\alpha t}$
+371.75 after 12/89
α=.08258
β=209.88
t= months since Jul '88
r=.992
• Grows at 159.14%/yr
• Doubles every 7.68 mos.
F. Solensky / Racal InterLan

Depletion Dates

• Assigned Class "B" network numbers	Mar. 11, 1994
• NIC "connected" class B network numbers	Apr. 26, 1996
• NSFnet address space*	Oct. 19, 1997
• Assigned Class "A-B" network numbers	Feb 17, 1998
• NIC "connected" Class A-B network numbers	Mar. 27, 2000
• BBN snapshots*	May 4, 2002

* all types: may be earlier if network class address consumption is not equal.

Zaw-Sing Su
NIC
ASSIGNED NUMBERS
STATISTICS: 1983 TO 1990

MONTH/YEAR		CLASS A	CLASS B	CLASS C
Jul. 1990	Connected	29	1209	2972
	Unconnected	34	2533	16214
Aug. 1989	Connected	33	767	2666
	Unconnected	33	1524	12613
Aug. 1988	Connected	28	314	2137
	Unconnected	28	493	8198
Nov. 1987	Connected	26	230	1821
	Unconnected	27	301	7494
Mar. 1987	Connected	26	149	962
	Unconnected	27	188	6319
Nov. 1986	Connected	26	130	922
	Unconnected	27	154	5800
Dec. 1985	Connected	17	81	1510
	Unconnected	18	91	4721
Apr. 1985	Connected	17	74	1457
	Unconnected	17	77	3551
Oct. 1984	Connected	13	68	1382
	Unconnected	13	69	3209
Jun. 1984	Connected	13	52	837
	Unconnected	13	53	2398
Oct. 1983	Connected	14	37	1066
	Unconnected	14	37	1834
Jan. 1983 (Note: No distinction was made between connected or unconnected status)		31	24	1042

Sep. 1981
(Note: Class A, B, and C are established but no tallies have been done)

Jan. 1981 - Nov. 1977
(Note: The Class A, B, and C system has not been established and each Assigned Network Number is listed separately)

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director: Russ Hobby/UC Davis

Working Groups Meeting at UBC

Network FAX - The primary work that went on was to determine the type of functions the Working Group wanted to define. Two types of FAX transport were thought to be desirable.

- A protocol needs to be defined to allow FAX machines directly connected to the network to communicate. This would be similar to the way FAXs are used today on phones, only the transport is over the TCP/IP network.
- There needs to be defined a method to send FAXs in a "store and forward" type of environment. This could be as simple as using email with FAX body parts.

Network Printing Protocol - There was a change in the Chair of this Working Group. The new Chair, Glenn Trewitt of DEC, will now be chairing the Working Group. The Working Group came up with three things that need definition.

- A "wire protocol" that will allow sending a data stream over the network and out a particular wire to a device.
- A printer access protocol to connect to printers that are directly connected to the network and have an address.
- A job submission/spooler protocol to control printers and queues.

TELNET - The Linemode and Environment Options documents were submitted to become RFCs. The work at this meeting covered the issues of authentication and encryption. There was also discussion of the "wire protocol" needed by the Network Printing group. Half of the meeting was a joint meeting with the Authentication Working Group discussing common issues

Working Groups Not Meeting at UBC

Domain Name System - This Working Group is still on hold, pending issues.

Network SQL - There seems to be a lack of expertise on SQL among the current IETF population. The Chair, Cliff Lynch/UCOP has approached vendors and some SQL groups to gather the appropriate people to work on the problem. The combination of SQL experts and network expertise already in the IETF will be necessary to solve these problems.

Other Applications in Need of Work

There are still applications for the Internet that need some standards definitions.

- FTP needs some upgrading. Functions such as transferring file attributes need to be added. It would also be nice to have archiving/backup and compression functions built into FTP.
- Applications such as Telnet and email need to be updated to handle alternate character sets now that the Internet is an international network.
- Directory services are becoming increasingly important. Is X.500 going to meet our needs soon enough?
- People are asking for calendar/scheduling over the network. A protocol is needed to define calendar/schedule operations.

3.1.1 Domain Name System (dns)

Charter

Chair(s):
Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:
General Discussion: namedroppers@nic.ddn.mil
To Subscribe: namedropped-request@nic.ddn.mil

Description of Working Group:

No description available

Goals and Milestones:

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

3.1.2 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 facsimile across TCP/IP networks and to develop recommended standards for facsimile transmission across the Internet. The group is also intended to serve as a coordinating forum for people doing experimentation in this area to attempt to 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:

Aug 1990	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

- Meeting convened 9:30 AM on 8/1/90 chaired by Mark Needleman
- Attendees briefly introduced themselves and explained their interests.
- University of California and Ohio State University gave short presentations on work in progress or planned.
- Draft charter for group accepted by default in that no one raised any objections to it.
- There was a discussion of what transfer formats needed to be supported. The consensus was that there did need to be a protocol developed to support transmission of fax across the internet. At the same time support for fax using RFC822 and X.400 mail was needed. There was a sense that this 2nd issue might only need to be a reiteration of what is already stated in RFC1154(?)
- An attempt was made to define a list of requirements for the new fax transmission protocol. These were:
  - Protocol should have mechanisms for determining destination node and route to destination node.
  - Protocol should work with both G3 and G4
  - There should be translation functions for standard document formats including, at minimum, ascii to fax and postscript to fax
  - The protocol should support sending fax machine to machine via internet, computer system to fax machine via internet, and fax machine to computer system via internet
  - Protocol should support receipt information to sendors.
  - Protocol should support explosion to multiple destinations.
  - Protocol should provide interfaces to other standard fax services.
- Action items from meeting:
  - Marshall Rose (mrose@psi.com) will draft first cut at a proposed architecture based on requirements listed.
  - Karl Auerbach (karl@asylum.sf.ca.us) will draft description of what his proposed fax interface to mail looks like.
  - University of California and Ohio State University will write and distribute short descriptions of the projects they have underway or planned.

Attendees

Philip Budne phil@shiva.com

Jeffrey Case	case@utkux1.utk.edu
Andrew Cherenson	arc@sgi.com
John Cook	cook@chipcom.com
James Davin	jrd@ptt.lcs.mit.edu
Alf Farnham	carolf@mcescher.unl.edu
Michael Fidler	ts0026@ohstvma.ircc.ohio-state.edu
Robert Gilligan	gilligan@sun.com
Russell Hobby	rdhobby@ucdavis.edu
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Sam Lam	
Clifford Lynch	lynch@postgres.berkeley.edu
Stuart Lynne	sl@wimsey.bc.ca
Donald Morris	morris@ucar.edu
Mark Needleman	mhn@stubbs.ucop.edu
Cecilia Preston	ceal@asylum.sf.ca.us
Marshall Rose	mrose@psi.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.3 Network Printing Protocol (npp)

Charter

Chair(s):
Glenn Trewitt, trewitt@nsl.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.
Apr 1990	Write draft LPR specification.
May 1990	Discuss and review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.
May 1990	Submit final LPR specification including changes suggested at the May IETF. Discuss document on mailing list.
Jun 1990	Submit LPR specification as an RFC and standard.

Jul 1990	Write description of the Palladium printing protocol (2.0) in RFC format.
Aug 1990	Discuss and review the draft Palladium RFC.

CURRENT MEETING REPORT

Reported by Glenn Trewitt/DEC

NPP Minutes

We reviewed the goals of the Working Group (problems to be solved) and looked at how the problem could be partitioned into smaller subproblems. Three subproblems stood out:

1. Wire Protocol
   There is an immediate need among some vendors (notably those who make terminal servers) to have a standard that they can implement that provides the capability to create a TCP connection to one (of many) hardware "byte-stream" interfaces (either serial or parallel). The path must be capable of being 8-bit clean. It would be a good thing for the protocol to provide a mechanism for supporting "rotary groups" for groups of printers.
2. Printer Access
   Communication to a printer that has some sort of direct network connection. Presumably, the printer has its own IP address. The source of the print job is unspecified – it might just be someone's PC, or a full-blown print manager and spooler. There are many requirements here: authentication, accounting, capability negotiation (what page description languages are supported, paper sizes, special features, etc.), etc.
3. Job Submission Communication from someone who has a document to be printed to a printing manager/spooler. The current popular example of this is the lpr/lpd protocol, which most people seem to feel is inadaquate for more complex printing environments.

Decisions and Action Items

We quickly decided that problem (1), the wire protocol, deserved a general solution in a broader context, since there are a number of other applications that require it, such as data collection, modem pools, and "milking-machine" concentrators. Russ Hobby agreed, and suggested that we take that protocol project to the Telnet Working Group. Bill Westfield (cisco) agreed to do this.

In response to problem (2), Ajay Kachrani (DEC) proposed that we adopt the "Printer Access Protocol". This is the protocol used to communicate with Digital's networked PostScript printers. He handed out a description of it, which I will make available via anonymous FTP.

Richard Hart (DEC) has proposed that we address problem (3) by adopting the Palladium printing archictecture, developed at MIT's project Athena.

I suggested that we contact Adobe Systems, to see what input they could provide about printing architectures. I have since talked to Carl Orthlieb and Sherri Nichols at Adobe. They will be providing four documents about Adobe's model of printing architecture, and will participate in our activities as well.

I will shortly be setting up a directory (accessible via anonymous FTP) containing the Adobe documents, some relevant DEC Western Research Lab technical notes, as well as the documents about the two proposals (PAP and Palladium).

Other Activities

Leo McLaughlin and Robert Knight finished RFC 1179, documenting the Berkeley Line Printer Daemon Protocol.

Attendees

Philip Budne	phil@shiva.com
Anthony Chung	anthony@hls.com
Richard Hart	hart@decvax.dec.com
Russell Hobby	rdhobby@ucdavis.edu
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Stuart Lynne	sl@wimsey.bc.ca
Robert Morgan	morgan@jessica.stanford.edu
Glenn Trewitt	trewitt@nsl.dec.com
Bill Westfield	billw@cisco.com

3.1.4 TELNET (telnet)

Charter

Chair(s):

Dave Borman, dab@opus.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

We met as a group for a short period of time, and then adjourned to take part in the Authentication/Privacy and Security Research Group (PSRG) joint meeting.

In our meeting we discussed:

Problems/questions about the 4.3Reno version of telnet/telnetd. One problem is that the latest BSD release does not send telnet "Synch" commands. The code is there, but it is commented out. This is because some clients do not handle a "synch" properly, and get stuck in a loop. It was decided that it would be better to release the code with this feature turned on, with an option to turn it off if some site has a problem with it.

The "Don't Telnet" option was discussed again. There has been some, though not a great deal of interest in it. Bill Westfield said that he would send it out again with a drop dead date. If people can't agree that it is a good idea by the December IETF we will not have any future discussion on it.

There was discussion of using telnet for running printers. It was suggested that you should be able to connect to a box and say which terminal port you want to connect to, and you should be able to send information across about how to set up the hardware on the remote side. The discussion revolved around whether or not this sort of stuff belonged in telnet, or in a higher level protocol. If it belonged in telnet, how much could be done with existing options (like ENVIRON)? It was decided that before a decision can be made whether or not it belongs in telnet, someone needs to write up a list of what functionality is needed to do printers over telnet, and then look at that list and see how it maps into the current telnet spec, and if it does, will any new telnet options be needed.

There was some discussion about international character sets. At this point, we don't really know what to do about them. Should telnet know about them? How is CR/NL mapping handled? Nothing was decided, this issue will have to be pursued at a later date.

It was decided that we would like to start looking at reviewing/editing/re-writing the base Telnet RFC. Joyce Reynolds said that Jon Postel has a whole bunch of notes collected over the years that she will make available in some form. Part of the next Telnet Working Group meeting will be used to get the editing of the Telnet RFC started.

It was requested that the tn3270 mailing list be put in the minutes of the meeting. The list is at WG3270-L@UMDD.UMD.EDU. To subscribe send mail to BRUCE@UMDD.UMD.EDU.

The Working Group then adjourned, and went to meet with the Authentication/PSRG joint meeting. When we arrived, they were deep in discussion about SNMP authentication. When that was wrapped up, David Borman gave a brief overview to the group about the proposed Telnet Authentication option, and what was hoped to be accomplished with it. There was then discussion. The basic question that we wanted answered was "Should we have a Telnet Authentication option, or should the authentication be done outside the realm of Telnet, such as in a connection initiation protocol, which would happen before telnet started up?" This question was not answered, there were arguments for both sides.

Without going into all the pros and cons of what was discussed, there were two main things that came out of the discussion:

1. There needs to be a clear written description of the uses of the Telnet Authentication. The "motivation" section of the draft RFC is a bit terse, and should be expanded.
2. The Telnet Working Group will continue to develop the authentication option. Experimental implementations are being started. Both having and experiencing an implementation will be help in answering the question "Do we need it?".

The Telnet Working Group will meet next at the December IETF meeting in Colorado.

Attendees

Richard Basch	probe@mit.edu
Dave Borman	dab@opus.cray.com
Philip Budne	phil@shiva.com
Anthony Chung	anthony@hls.com
George Conant	geconant@eng.zyplex.com
Mark Crispin	mrc@cac.washington.edu
Kevin Fall	kfall@Berkeley.EDU
Neil Haller	nmh@bellcore.com
Russell Hobby	rdhobby@ucdavis.edu
Steven Hubert	hubert@cac.washington.edu
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Michael Karels	karels@berkeley.edu

Luping Liang	liang@cs.ubc.ca
Joyce K. Reynolds	jkrey@venera.isi.edu
Dana Sitzler	dds@merit.edu
Frank Solensky	solensky@interlan.interlan.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Dean Throop	throop@dg-rtp.dg.com
Bill Westfield	billw@cisco.com
Yueli Yang	yueli@bnr.ca

3.2 Host and User Services Area

Director: Craig Partridge/BBN

Host Services

Several productive Working Group meetings were held.

Special Host Requirements

This was their first meeting. The Working Group promptly found itself in a careful debate about the proper definition of a special purpose host. It was generally felt that once this definitional question was resolved, progress could be quickly made. Some definitions have been proposed and the Working Group is currently hashing out which one is best on its mailing list.

User Connectivity Problems

The group discussed a trouble-ticket scheme developed by Matt Mathis. The Working Group adopted the scheme, with some modifications, and is now working on developing a list of information that a trouble-ticket ought to contain.

Dynamic Host Configuration

This group is nearing completion of a draft of a host configuration protocol.

User Services:

Reported by Joyce Reynolds

User-Doc Working Group - Submission as Internet Draft July 3, 1990

Chaired by Karen Roubicek and Tracy LaQuey

The User-Doc Bibliography was submitted to the Internet Draft process on July 3, 1990. Final changes or amendments to the Bibliography are in process, with submission to the RFC Editor ASAP after the IETF meeting in UBC.

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

NISI - Discussion on the "Cooperation of NICs"

Chaired by Dana Sitzler

At this NISI session it was decided that instead of one major document, NISI could better serve if it put out various documentation to help regional NICs. An 11-12 bullet outline on "Recommendations to NICs" was develped at this meeting. It is an "ethics" type of document to NICs. For example, some of the bullets pertain to how NICs should interact and cooperate with each other. This document is specifically aimed at regional NICs, not campus NICs at this point in time. NISI will have a draft document to work on at the next IETF in Colorado, after which it will be submitted as an Internet Draft.

SSPHWG - Security Area/User Services Area Combined Efforts

Chaired by J. Paul Holbrook and Joyce K. Reynolds

The first meeting of the SSPHWG (Site Security Policy Handbook Working Group) was held at the May 90 IETF in Pittsburgh and second meeting was held on June 12th, at USC/Information Sciences Institute. At this third meeting the time was fully devoted to going through the first pass rough draft of the Handbook.

The draft of the Handbook was well received, and the general concensus of attendees was to keep with the direction of the document, making one more pass at the next IETF in Colorado. Submission of the Handbook to the Internet Draft process is projected to be in mid-December, for publication as an RFC FYI at the end of 1990.

This Working Group is the first to combine the efforts of two separate IETF Areas. The response to this has been successful.

USWG - Running at its Peak

Chaired by Joyce K. Reynolds

Agenda items included: Presentation and discussion of the current USWG "priority list":

- Top Priority:
  - User-Doc RFC/FYI Publication
  - NISI
  - Site Security Policy Handbook
  - Internet Installation Checklist
  - QUAIL
- Low Priority:
  - DAWG
  - Internet Stats
  - Intro Packages
  - User-Glossary

3.2.1 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.

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chair(s):

Ralph Droms, `droms@sol.bucknell.edu`

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Ralph Droms/ Bucknell

DHC Minutes

The meeting began with a presentation from Bill Nowicki of Legato about Legato's "Network Resource Administration Platform". Bill prefaced his talk with the statement that, while our Working Group is solving the low level problems associated with dynamic naming and addressing, his talk would cover more high level issues. The platform he described is intended to ease the use and management of network facilities for a running system.

The remainder of the meeting concentrated on Jeff Mogul's "Proposal for Supporting IP Address Assignment Using Coordinated BOOTP Servers". Jeff has written up the IP address allocation and delivery mechanism as discussed at the June meeting in Palo Alto (see the PSC Proceedings for details). A copy of Jeff's proposal is available for anonymous FTP on sol.bucknell.edu in file dhcwg/mogul.prop.

The discussion raised several specific questions about the proposed protocol:

- The protocol must not require a server on each subnet.
- There should be a new response "No IP address available".
- A client can release an IP address back to the IP address server - can the client also clear ARP caches of any references to that client?
- How can a client find out that its IP address has been reallocated after a network partition heals?
- Rather than periodically pinging a host to verify that its IP address is still in use, the IP address servers should only check when IP addresses need to be reused.
- We need to write into the protocol description something about the eventual use of multicast.

Steve Deering was good enough to go through the Host Requirements RFC and generate a list of per network and per interface parameters. He wondered if we should consider any of these as part of the dynamic host configuration protocol:

Per network parameters:

Gateway forwarding	on/off
Non-local source routing	on/off
Policy filters for non-local source routing	(list)

Maximum reassembly size	integer
Default TTL	integer
PMTU aging timeout	integer

Per interface parameters:

Perform mask discovery		on/off
Be a mask supplier		on/off
Perform router discovery		on/off
Router solicitation multicast address		(multicast address)
Ignore router discovery		on/off
Default router list		(addresses and preferences)
Static routes:		
	destination	(host/subnet/net)
	mask	(subnet mask)
	type of service	integer
	first hop router	(address)
	ignore redirects	on/off
	PMTU	integer
	perform PMTU discovery	on/off

Attendees

Karl Auerbach	auerbach@csl.sri.com
Richard Basch	probe@mit.edu
Scott Bradner	sob@harvard.harvard.edu
Andrew Cherenson	arc@sgi.com
Steve Deering	deering@pescadero.stanford.edu
Tom Evans	wcc@cup.portal.com
Karen Frisa	karen@kinetics.com
Robert Gilligan	gilligan@sun.com
Yong Guo	guo@cs.ubc.ca
Steven Hubert	hubert@cac.washington.edu
Holly Knight	holly@apple.com
Gregory Lauer	glauer@bbn.com
John Leong	john.leong@andrew.cmu.edu
Paul McKenney	mckenney@sri.com
Lynn Monsanto	

Robert Morgan	morgan@jessica.stanford.edu
Bill Nowicki	nowicki@sun.com
Frank Slaughter	fgs@shiva.com
John Veizades	veizades@apple.com
Walter Wimer	ww0n+@andrew.cmu.edu
Denis Yaro	DYARO@SUN.COM

3.2.3 Internet User Population (iup)

Charter

Chair(s):

Craig Partridge, craig@nnsc.nsf.net

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

3.2.4 Network Information Services Infrastructure (nisi)

Charter

Chair(s):
Dana Sitzler, `dds@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

Current Meeting Agenda:

- Recap of last meeting
- Draft document
- Discussion: Relationship between NIC's/Cooperative Agreements

Minutes:

The session began with a recap of the last NISI meeting which included discussions about where we are today in terms of network information services. A draft document summarizing these discussions will be sent to the NISI mailing list prior to the next meeting.

The session then moved into a new discussion area namely how do we propose that NICs work together to provide information services to the internet community.

Discussion:

- How should NICs work together?
- What kind of cooperative agreements can we suggest?

The idea is that NISI (or the IETF) could "encourage" NICs to "formally" work together for the benefit of network users – in much the same way people are encouraged to follow the technical specifications to 'play' in the internet. The discussion then moved onto; what are the rules that NICs should follow?

The result of this discussion was a set of suggested guidelines for NICs:

- WILL ANSWER TO BEST OF ABILITY BEFORE PASSING OFF
- NIC PHONE SHALL ALWAYS BE ANSWERED
- WILL PROVIDE AN INDEX OF SERVICES
- WILL PROVIDE REFERRAL FILES
- NIC'S SHOULD KEEP OTHER NIC'S INFORMED OF "SPECIALTIES"
- TRY TO PUT INFORMATION ON-LINE AND PROVIDE;
  - ACCESS
  - UPDATE

  - ANNOUNCEMENTS
- PARTICIPATE IN NIC FORUM AND PLAY BY RULES
- MAKE INFO AS WIDELY AVAILABLE AS POSSIBLE (COPYRIGHT TO MAINTAIN INTEGRITY OF INFORMATION, BUT STRIVE TO ENCOURAGE DISTRIBUTION – NOT NECESSARILY INTENDED TO IMPLY "FREE")
- ACKNOWLEDGE INFO SOURCE (THEREFORE, NOT LIABLE)
- DON'T MAKE UP ANSWERS
- OBLIGATED TO ANSWER? (IN SOME WAY)

While discussing these guidelines, other ideas were generated about ways of implementing them and general suggestions for improving the ability of NICs to know about each other and to work together. These ideas included:

- List of NIC's
- Index of Services
- Referral Contacts
- NIC Forum (mailing list)
- Interface to encourage electronic use, in place of paper
- Database – pull out info, put in another database for your system
- Some kind of way to do referrals
- Programs to access NIC's
- NICs should implement a common address name such as (NIC@DOMAIN) to give users a place to start

The group decided to publish the suggested guidelines as an FYI-RFC. The guidelines will be further developed, put into RFC format, and a draft will be distributed to the NISI mailing list before the next IETF meeting.

Follow-up actions:

Dana	Create a NIC Forum Mailing List – Send draft document (of last meeting) – Send recommended guidelines list
Gary	Create a Template
Mike	Bitnic Listserv
Joyce	Edit/provide input for submission of RFC
Others	Send "other projects" (Building blocks)

Attendees

Carol Farnham	carolf@mcescher.unl.edu
Russell Hobby	rdhobby@ucdavis.edu
Michael Hrybyk	mwh@educom.edu
Erik Huizer	huizer@surfnet.nl
Clifford Lynch	lynch@postgres.berkeley.edu
Gary Malkin	gmalkin@ftp.com
Berlin Moore	prepnet@andrew.cmu.edu
Fred Ostapik	ostapik@nisc.sri.com
Craig Partridge	craig@nnsc.nsf.net
Cecilia Preston	ceal@asylum.sf.ca.us
Joyce K. Reynolds	jkrey@venera.isi.edu
Mark Seger	seger@mjs1/ogo.dec.com
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov

3.2.5 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:

Jun 1990	Mailing list discussion of charter and collection of concerns.
Aug 1990	First IETF Meeting: discussion and final approval of charter; discussion and agreement on approach, including models, format, level and type of detail. Make writing assignments.
Oct 1990	First draft document.
Nov 1990	Second IETF Meeting: review first draft document, determine necessary revisions. Follow up discussion on mailing list.
Jan 1990	Revised document.
Feb 1990	Third IETF Meeting: make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail.
Apr 1991	Final draft document.
May 1991	Fourth IETF meeting: review final draft and if OK, give to IESG for publication as RFC.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

SHR Minutes

Agenda

- Is this trip really necessary?
- Principles of operation?
  - Architectural purity.
  - Interoperability.
  - Cost/benefit.
- Definition of Special-purpose Host?
  - By overall function?
    - Terminal server.
    - File server.
    - Toaster.
  - By subfunctions?
    - Network self load.
    - Programming interface.
    - Terminals (character devices).
    - Files (FTP, NFS, etc.).
    - Network management client.
    - Network management agent (e.g., bridge, router).
- RFC Format and organization?
- Specific issues?
  - IP fragment reassembly from ¡ 576 byte fragments.
  - TCP efficiency (e.g., Jacobson retransmission in a ROM).
  - Source routing.
- Contributors?
  - Analyses from vendors of example systems.
  - RFC section authors.

The Agenda's first question was "Is this trip really necessary?" The consensus was affirmative. We need some clarifications, the contention was over how far they should go and what form they should take.

On the question "Principles of operation?", we generally agreed that interoperability is the primary goal. George Conant of Xyplex suggested that our first concern should be maintaining the strength of requirements whose intention is to protect the network from misbehaving hosts. Bound by this principle, we can then apply some cost/benefit analysis to "musts" required for architectural purity or use by wizards under unusual conditions.

Considerable discussion and disagreement did not result in an answer to "Definition of Special-purpose Host?" Although subject to debate, the majority seemed to think along the lines of recognizing optional functional areas, such as an open programming interface, limited application protocols (such as just Telnet), and so on. This points to the consideration that "special-purpose host" simply means a host that isn't general purpose.

The question of "RFC Format and organization?" should have included "Degree of Specificity?". Stev Knowles of FTP Software led the charge for stating principles and omitting specifics. He was not alone (as if that matters), but the majority believed we must be more specific. David Jordan of Emulex spoke for organization by system type (for example, terminal server, file server). The strongest consensus was to organize around the RFC 1022/1023 "musts", examining each in the light of hosts with useful application subsets.

On "Specific issues?", the consensus on IP reassembly was "Shut up and do it." Source routing was less clear. The idea of keeping the requirements to forward and to respond on the reverse path but weaken the requirement to originate a source route had noticeable support. Issues such as TCP efficiency appear subject to the rule of not hurting the network while allowing space for knowing exactly how your limited TCP user (such as a ROM) will use TCP.

The answer to "Contributors?" was:

- Stev Knowles will supply a statement of principle. We will then judge whether we are done.
- Bill Westfield of cisco and Robert Elz of the University of Melbourne will each supply an analysis of the "musts" which might be subject to weakening in special cases.
- David Jordan will propose an all-inclusive list of special host types.

Discussion was lively and varied, with many valued participants other than the few mentioned above. Discussion will continue on the mailing list. According to the (unchanged by the way) charter, the next milestone is a draft document by the end of October, for review at the December IETF meeting. The above contributions will provide the text of that document.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Larry Brandt	lbrandt@sparta.com

Asheem Chandna	ac0@mtuxo.att.com
Anthony Chung	anthony@hls.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
George Conant	geconant@eng.zyplex.com
Robert Elz	kre@munnari.oz.au
Richard Fox	sytek!rfox@sun.com
Karen Frisa	karen@kinetics.com
Martin Gross	gross@polaris.dca.mil
Peter Hayden	hayden@levers.enet.dec.com
Ruei-Hsin Hsiao	nac::hsiao
Ole Jacobsen	ole@csli.stanford.edu
David Jordan	...jordan@emulex.com
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Michael Karels	karels@berkeley.edu
Frank Kastenholz	kasten@europa.interlan.com
Stev Knowles	stev@ftp.com
Sam Lam	
Paul Langille	quiver::langille@decwrl.dec.com
John LoVerso	john@loverso.leom.ma.us
Yoni Malachi	malachi@polya.stanford.edu
Keith McCloghrie	kzm@his.com
David Miller	dtm@ulana.mitre.org
Craig Partridge	craig@nnsc.nsf.net
Stephanie Price	cmcvax!price@hub.ucsb.edu
Michael Reilly	reilly@nsl.dec.com
Craig Smelser	
Bob Stewart	rlstewart@eng.xyplex.com
Bill Townsend	townsend@xylogics.com
Justin Walker	justin@apple.com
Jonathan Wenocur	jhw@shiva.com
Bill Westfield	billw@cisco.com

3.2.6 User Connectivity (ucp)

Charter

Chair(s):

Dan Long, long@bbn.com

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Dan Long/BBN

UCP Minutes

Agenda

- Introduction to UCP Working Group:
  - What is it? What's been done so far?
  - Discussion of Matt Mathis' National Trouble Ticket Tracking writeup.
  - Discussion of some operational issues by MERIT.
  - What's Next?

Dan Long (Chair) presented a brief history of the UCP Working Group:

- FSU IETF: Initial discussion
  - Structural proposals presented
  - Refine goals/scope
  - Writeups by Craig, Elise, & Martyne
- PSC IETF: Definition of terms:
  - NSC (Network Service Center)
  - P1 (user<->NSC communication protocol)
  - P2 (NSC<->NSC communication protocol)
  - Writeup by Matt

Matt Mathis (PSC) reviewed his description of a National Trouble Ticket Tracking system. A lively discussion ensued about various aspects of the proposal including:

- How do you define "closure with the user" (as in "a ticket is a contract to obtain closure with a user")?
  - What do you do about uncooperative NOC's?
  - What do you do when you cannot satisfy the user due to funding/engineering constraints?
  - Transfer of a ticket is a mechanism for obtaining closure and resolving the problem. We should acknowledge that certain problems can't be closed in a technical sense. This may be sufficient for closure with the user.
- What are the organizational implications of declaring a ticket to be a "contract"?
  - Does that mean the NSC must respond to any old barage of (nuisance) questions?
  - Can an organization commit to adhering to this system without knowing the expected demand?

- How are NSC's "certified" (as in "NSC's must be certified at least as far as adherence to the rules described in this document")?
  - We don't want to be (or can't be) coercive.
  - Needs some element of informal (polite) coercion rather than legal coercion. The problem is to get somebody to start owning the problem and a way of recording where the problem lies.
  - Makes more sense to have the system be so useful that everyone will want to join and conform.
  - Certification should only be that the NSC's adhere to the ticket hand-off protocols. Details of P2 protocol need to be fleshed out by the person who sets up the TTC.
- What about peer-bashing (i.e., pointing fingers, blaming,...)?
  - It's self-regulating (...glass houses...stones...).
  - Would a national ombudsman be reasonable?
- What about lots of users complaining about the same problem?
  - Have multiple user dialogs cross referenced with a single "problem" which has the other dialogs.
  - Closure should be obtained with each user.
  - We do want to track each caller so we know how many complaints there are.
- What about privacy of ticket information?
  - Tickets should be readable only by the owner and the ticket arbitrage center (TAC).
- What do you do with the Engineering Dialog results?
  - If the Engineering Dialog results in suggested improvements, how do those get handled?
  - Does everyone who hears about the suggestions understand the possible implementation obstacles?

Dale Johnson (Merit) led a discussion on some aspects of this system not covered in the document:

- Any national Ticket Tracking system will have to be used in conjunction with local systems. For large sites which have elaborate highly customized systems of their own, this might require software to automatically copy tickets between the local and national system. Making the national system available for all networks' local tickets could simplify operations for many NOCs, although this could result in an extremely expensive national system. If the national system was freeware or was reasonably available, then NOCs could at least use the same software for both their local and national tickets.

- NSC's still need the tools to do the diagnosis. Especially important is contact information for different network entities. The NNSC Phone Book may help solve this problem. Contact information should be both published and online.
- The NJM Working Group has started discussing common data formats and access mechanisms for the routine (SNMP and other) data that NOCs collect. Access to this kind of data from other networks could become very useful when a NOC tries to debug a complex problem outside of its own jurisdiction, or when another entity wants to aggregate or contrast data from different NOCs. NJM will continue with this project, but noted that this might also be interesting to the UCP group because it is a form of inter-NOC communication.
- How can we alert network users about outages, both planned and unplanned? How about an X.500-based (or DNS-based) posting system that people (and network utilities?) can query to determine the operational status of various network components? There was a fair amount of discussion about a low-tech short-term solution involving a standard format for problem reports posted to the NSR mailing list. The thought was that these standard reports could then be automatically collected for occasional perusal/reference by NSC staff.

Action Items

- Matt - will redraft with the suggested changes from the discussion:
  - No compulsion; be neutral
  - Privacy; tickets readable only by owner and TAC
  - TAC will mention the ombudsman role
  - Omit details of ticket format (for now)
  - Need requirements for TTC
  - It's ok for 1 ticket to have multiple user dialogs
- Dan/Craig - will clean up draft & submit into the FYI RFC pipeline
  - Check FYI RFC standards to be sure that the "2 voice" format is acceptable
  - Provide copy of draft to FARNET's September meeting

Timetable Through 1990

August	Matt will present revised draft; UCP group to comment
September	Dan/Craig will incorporate comments, and prepare draft for presentation to FARNET and submission to FYI RFC pipeline
October/November	Collect comments and refine proposal.
December	At IETF meeting, discuss deployment/future plans

Attendees

Stephen Adams	decwrl::''adams@zeppo''
Eric Carroll	eric@utcs.utoronto.ca
Carol Farnham	carolf@mcescher.unl.edu
Dale Finkelson	dmf@westie.unl.edu
Vince Fuller	fuller@jessica.stanford.edu
Steven Hubert	hubert@cac.washington.edu
Dale Johnson	dsj@merit.edu
Ken Jones	uunet!konkord!ksj
Dan Long	long@bbn.com
Matt Mathis	mathis@pele.psc.edu
Berlin Moore	prepnet@andrew.cmu.edu
Donald Morris	morris@ucar.edu
Craig Partridge	craig@nnsc.nsf.net
Dana Sitzler	dds@merit.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Carol Ward	cward@spot.colorado.edu
Robert Woodburn	woody@saic.com

3.2.7 User Services (uswg)

Charter

Chair(s):

Joyce K. Reynolds, jkrey@venera.isi.edu

Mailing Lists:

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

Description of Working Group:

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

- User-Doc Bibliography - Internet Draft, July 3, 1990
- SSPHWG - Met Wednesday morning, August 1, 1990
- NISI - Met Wednesday afternoon, August 1, 1990

Presentation and Discussion of the current USWG "priority list":

- Top Priority:
  - User-Doc RFC/FYI Publication
  - NISI
  - Site Security Policy Handbook
  - Internet Installation Checklist
  - QUAIL
- Low Priority:
  - DAWG
  - Internet Stats
  - Intro Packages
  - User-Glossary

Discussions/Reports:

QUAIL - Presented by Gary Malkin

The revamped QUAIL document was well received and will be published in mid-August as an RFC FYI. This RFC FYI is the first in a collection of FYI's called, "Questions and Answers" (Q/A) produced by the USWG. The goal of this series is to document the most commonly asked questions and answers in the Internet.

Future updates of this memo will be produced as USWG members become aware of additional questions that should be included, and of deficiencies or inaccuracies that should be amended in this document. Additional FYI Q/A's will be published which will deal with intermediate and advanced Q/A topics.

Internet Installation Checklist - Presented by Bob Enger

An installation checklist for the Internet is currently in draft stage and is intended to be of use to people of all levels; new, intermediate and advanced. It is general in nature for new and intermediate users, yet advanced users should find it an effective compilation of important information for the Internet community.

An outline and sketchy rough draft was presented by Bob Enger at the UBC IETF, with discussions and suggestions for the checklist noted. Writing will continue and the next pass draft of the checklist will be presented at the December IETF in Colorado.

Attendees

Eric Carroll	eric@utcs.utoronto.ca
Robert Enger	enger@sccgate.scc.com
Carol Farnham	carolf@mcescher.unl.edu
Phill Gross	pgross@nri.reston.va.us
Michael Hrybyk	mwh@educom.edu
Erik Huizer	huizer@surfnet.nl
Holly Knight	holly@apple.com
Gary Malkin	gmalkin@ftp.com
Marilyn Martin	martin@cdnnet.ca
Berlin Moore	prepnet@andrew.cmu.edu
Donald Morris	morris@ucar.edu
Craig Partridge	craig@nnsc.nsf.net
Joyce Reynolds	jkrey@venera.isi.edu
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Zaw-Sing Su	zsu@tsca.istc.sri.com

3.3 Internet Area

Reported by Noel Chiappa

The Connection-Oriented IP Working Group is proceeding along two tracks. The first track consists of work on the interim ST-2 protocol. A draft of that standard is in progress; this is currently at the "last but one" stage. Implementations of that spec are underway by BBN and UWashington for the SPARC and the BBN router, and a host implementation will be done by Steve Casner of ISI.

The second track consists of longer research work, suitable for use in follow-on architectures. Examples of the types of topics being addressed are resource management and connection management. Current work revolves around the construction of test bench software for experimentation and metering.

The Router Requirements Working Group has determined that three (or possibly more) documents are in fact needed. First, a Router Requirements document basically along the lines of the Host Requirements (although there is some concern at the way that document mixed fixes to previous protocols with actual requirements) which covers IP and higher levels; second, a similar document which deals with media issues (and which will be shared between the Router and future versions of the Host documents); and third, a new revision of the IP and ICMP specs (which are woefully out of date and modified by many later RFC's). Additionally, it was suggested that a document(s) to guide procurement writers in the use of the Router Requirements would be useful, but the group has decided not to tackle that windmill.

At the IETF, the Working Group had the bulk of a first draft of the first document, and went through that draft in some detail, holding discussions of a few technical topics as required.

The PPP Extensions Working Group had final drafts of some documents, and first drafts of a number of others. One final draft concerns an option to allow use of a 32 bit CRC instead of the 16 bit, another remote bridging, and another an MIB for PPP links. First drafts covered DECNet Phase IV, Appletalk, OSI, and SNAP. Note tha proposed drafts for use of XNS and IPX did not appear.

Since it is anticipated that all of these documents will be completed prior to tl next IETF meeting, this Working Group will be going dormant; i.e., it has no furth plans to meet or produce anything. The mailing list, etc., will remain however, sho someone produce additional PPP extensions.

The Router Discovery Working Group held their final meeting, and looked over w is hoped to be the "last but one" draft. Some minor fine technical points

CURRENT MEETING REPORT

Reported by Claudio Topolcic/BBN

CIP Minutes

Agenda

- ST-II specification
  - Identify remaining issues
  - Discuss remaining issues
  - Resolve remaining issues
  - Assign writing tasks
- Connection oriented protocol research collaboration
  - Discuss possible collaboration efforts

The CIP Working Group met during all five Working Group sessions. Our primary ;oal for this meeting was to resolve the remaining open issues in the ST-II protocol)ecification; three sessions were dedicated to this effort. In the other two sessions discussed collaborative experiments on connection-oriented internet protocols.

raft of the ST-II specification was distributed and discussed at the previous IETF. ral issues were resolved then and new ones uncovered. Prior to the current 'ng, Charlie Lynn distributed an updated draft incorporating the results of the us meeting plus ensuing teleconferences and email. We discussed the changes resolved issues as follows:

- ;cedence is a per-connection characteristic, and is negotiated in the flowspec. re is a separate priority on each data packet to allow for layered coding nes within one stream.
- greed that all header and option chunks should have 32-bit alignment, ing 32-bit entities within chunks, to efficiently accommodate machine ar- ures with that constraint.
- FUSE and REROUTE negative response messages will be combined and the receiver will use the reason code to determine what action is te.
- of packet rate and size are offered, agents along separate branches tion might choose incompatible combinations each of which meets m product requirement. Intermediate agents must keep track of ng each branch, so resolution can be left to the application.
- t a CHANGE won't cause the existing connection to break, the s must include the existing settings.

- It is not considered an error if the next hop on a path is out the same interface as the previous hop, to allow relay multicasting.
- The current specification does not allow a uniform way on all control messages to determine the intended client. New fields are to be added to the control message header to allow this.
- A mechanism for grouping streams is provided, but their use is not yet well enough understood, and will therefore, be left to experiment.
- To make use of IP encapsulation paths between ST agents not directly connected, the ST routing table must be extended.
- The current flowspec definition does not allow specifying a variable-rate requirement nor discrete steps in place of a range. There are provisions to define new flowspec versions as we learn what is needed through experiments.

Writing assignments were also issued for sections of the document that are incomplete but not controversial. The draft is to be ready for submission as an Internet Draft in two weeks, followed by submission as an RFC after a comment period. The protocol will have "Not-Recommended Experimental" status while the CIP Group and others conduct experiments.

Collaboration

On Wednesday, we heard status reports on experimenters' plans. Allison Mankin described her work to implement Lixia Zhang's Flow Protocol algorithms within the framework of the BSD OSI TP2 protocol. She is now implementing the virtual clock mechanism in the BSD network drivers. Allison will test the protocol in the MITRE-DCA Testbed Network; she invites others to use the testbed, too.

Charlie Lynn described the collaboration of BBN and Washington University in St. Louis to develop the "COIP-kernel" – basically a new protocol family added into the BSD socket interface around which a variety of connection-oriented protocols could be implemented. The kernel is to be done by the end of August, then during September BBN will develop a set of modules around the kernel to implement ST-II.

Paul McKenney told us about the traffic generators he is developing so that DARTnet experimenters can conduct repeatable experiments. They run in user space and can be synchronized at multiple sites, injecting packets at the NIT, RAW_IP or transport level. Measures are defined for both "best effort" and "resource reservation" types of protocols.

Finally, we discussed how members of the group might collaborate. Allison expressed interest in using the COIP-kernel to extend Flow Protocol testing to the DARTnet Sparcstation environment. Paul's traffic generators may also be usable in the network testbed. Conversely, Paul might be able to incorporate Allison's DEC-bit code into the stochastic fair queuing algorithm.

Meeting action list

Casner	Rewrite sect 2 (& 2.1?) in about 3 pages (may be ok now).
Everyone	Comment on whether sections 2.3 through 2.7 are complete.
Casner	Update old encapsulation text of sect 3.7.3.
Topolcic	Edit or rewrite section 3.7.5 on Robustness.
Lynn	Edit sect 3.7.6. on Routing to simply list the things we expect from the routing function, but state that routing is not addressed here.
Topolcic	Edit or rewrite section 3.7.8. on Groups of Streams to state that groups are a way of associating streams and to just list some possible uses of such associated groups.
Lynn	Produce text for section 3.7.9. on the Source Route Option.
Lynn	Write a section in 4.3.1, FlowSpec that addresses the Burstiness parameter.
Lynn	Edit the paragraphs in section 4.3.1. that describe LimitOnCost and LimitOnDelay to specify the units.
Topolcic	Rewrite section 4.3.5.3. on Group Parameter to simply provide suggestions for the uses of Groups.
Lynn	Expand sect 4.4.14 on use of STATUS command for failure detection.
Everyone	Help find all the constants for inclusion in section 4.5, Suggested Protocol Constants, and should suggest values.
Everyone	Help write section 6, Areas Not Addressed, and specifically to help draw up a list.
Everyone	Help identify subsets everywhere.
Schroder	Provide protocol exchange diagrams.
Everyone	Think of good way to simplify protocol demultiplexing; consider origin & target(s) of stream on same host.

Attendees

Stephen Casner	casner@venera.isi.edu
Steve Deering	deering@pescadero.stanford.edu
Kevin Fall	kfall@Berkeley.EDU
Kathleen Huber	khuber@bbn.com
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Charles Lynn	clynn@bbn.com
Allison Mankin	mankin@gateway.mitre.org
Paul McKenney	mckenney@sri.com
K.K. Ramakrishnan	rama%erlang.dec.com@decwrl.dec.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
Claudio Topolcic	topolcic@bbn.com
Sijiam Zhang	szhang@cs.ubc.ca

3.3.2 IP MTU Discovery (mtudisc)

Charter

Chair(s):

Jeff Mogul, mogul@decwrl.dec.com

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done	Decide if the proposal in RFC 1063 is sufficient, or if there are flaws to be corrected, or possible improvements to be made. Or, decide that it is unwise to create an official standard.
May 1990	Unless the proposal in RFC 1063 is acceptable, write a new RFC describing a different approach.
Ongoing	Encourage the participation of gateway implementors, since the MTU discovery process affects the design and performance of IP gateways.
Done	Encourage sample implementations of end-host and gateway portions of MTU Discovery for popular software (BSD-derived kernels, primarily). Encourage rapid implementation by major gateway vendors, since this option is relatively useless without widespread support.

3.3.3 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:

Feb 1991	Describe, in an RFC, the current set of protocols used to connect Macintoshes to IP internets.
Feb 1991	Define a MIB for the management of DDP/IP gateways.

CURRENT MEETING REPORT

Reported by Bob Morgan/Stanford

APPLEIP Minutes

IP-in-DDP

John Veizades led a discussion of his draft RFC for IP-in-DDP. These issues were discussed:

- The use of the name "MacIP" for this protocol was criticized. People are encouraged to think of a new name.
- There was agreement that gateways should never do proxy ARP replies to NBP ARPs. In fact, clients are discouraged from doing NBP ARPs at all unless they have reason to believe that the destination is on the same AppleTalk internet. Clever clients can do NBP ARPs to optimize communication in this admittedly rare case. The user will probably have to specify the zone in which to do the NBP lookup in this case.
- Clients must be prepared to get responses from multiple gateways.
- The dotted decimal format for IP addresses used in NBP lookups must be better specified. Text might be borrowed from an existing RFC.
- Gateways currently send regular NBP confirms to their IP clients to determine whether the IP address is still in use. Gateways should try to minimize the bandwidth used for this, perhaps by only doing confirms when they are running short of IP client addresses.
- It was proposed that gateways should be able to be configured with a list of acceptable zones in which to do NBP ARPs. This should help to prevent duplicate IP address assignment, and let gateways and users search the entire "subnet" more easily when necessary.
- There could be a CEASE ATP message from gateway to client to tell the client to stop using an IP address (useful in case of duplicate assignment). There could also be a REDIRECT message from gateway to client, similar to ICMP redirects.
- It was suggested that gateways should have throttles on the rate at which they forward NBP lookups, to prevent clients from flooding internetworks with broadcasts. LBL has a working implementation. Apple suggested that System 7.0 will improve Macintosh client behavior in this area.
- The gateway's ATP response to a client ASSIGN request should be able to contain more information. It was proposed to define or redefine some of the response fields. The new format will be distinguished by putting a version number in the first 16 bits of the ATP User Data area. The second 16 bits must be zero. The first version to be defined will be version 1. New field uses:

The "Other #1" field is redefined to be the subnet mask.
The "Other #2" field is defined as a time server address.

Some implementors are already using some of the "Other" fields for their own purposes. They will report on these to the RFC author.

- Gateway implementors should report any error codes that they send in ATP responses to the RFC author, who will compile a generic list.
- A new ATP REGISTER STATIC request should be defined to allow clients with static IP addresses to register them with the server and get any useful response information. The client will put the static address into the "Assigned IP address" field. Gateways should do a sanity check on the address and send an error response if necessary.
- Several changes were suggested to the draft RFC. Among them:
  - Drop references to Macintosh.
  - Drop AARP definition.
  - Drop the line "The IP address used by a gateway with multiple IP addresses is the address that is responded to using the NBP ARP."
  - Hosts do not use ATP XO requests, but ATP ALO.
  - The line "There is no response to a RELEASE packet" should be "The ATP response to a RELEASE request is empty".
  - Drop the suggestion to limit IP-in-DDP datagrams to 576 octets.
  - Drop Step 3 in the sample transaction stream.

MIB

A draft MIB, written by Steve Waldbusser of CMU, was distributed. People found it generally acceptable. There was concern that it be clearly labelled as an "AppleTalk-IP gateway MIB" and not an "AppleTalk MIB".

It was noted that there is no AppleTalk-in-PPP MIB. Frank Slaughter from Shiva , who is working on AppleTalk-in-PPP, and Steve Waldbusser will work together on this.

It was suggested that the rtmpNextHop variable be extended with a Type string to distinguish between different protocol transports such as IP, DECnet, OSI, etc.

AppleTalk-in-UDP

Allan Oppenheimer from Apple led a discussion of wide-area networking using AppleTalk encapsulated in UDP/IP. The general idea is to connect existing AppleTalk internets via the IP Internet. There are a number of issues:

- Can/should a world-wide AppleTalk Internet be created using the facilities of the existing IP Internet?

- How much administration within a site is necessary/acceptable? How much coordination between pairs of sites, or between all sites, is necessary/acceptable?
- Is administrative control of routing necessary for security purposes, or is plug-and-play more crucial?
- Can the existing DDP-in-UDP encapsulation meet the need, or are changes required?
- Can all AppleTalk-based applications be supported? Is a subset such as LaserWriter printing and AppleShare file service acceptable/easier?
- Are there solutions to network number scaling and clashes? Are there solutions to zone name scaling and clashes?
- Is it important that hosts be able to communicate directly in this internet using the standard encapsulation, or is communication through routers sufficient?

Van Jacobson from LBL described a scheme that addresses some of these issues. He has implemented this method on software running on FastPaths at LBL and some other sites.

In Jacobson's scheme, each site maintains a table with one entry for each external AppleTalk network with which it wishes to communicate. Each entry in the table contains three fields. The first is the real 16-bit AppleTalk network number of an AppleTalk network at a remote site. The second is a 24-bit IP network number that is associated one-to-one with the previous AppleTalk network number. The third is a 16-bit AppleTalk network number which is used to identify the remote network within the local AppleTalk internet. The first two numbers form a pair that a site can give to any other site with which it wishes to communicate.

The table is distributed to some number of routers in the local AppleTalk internet that are running software that understands this scheme. Not all routers in the local internet are required to run this software.

When a participating router receives a datagram to be forwarded, it looks up the destination network number in its mapping table. If the number matches an entry (using the third field as described above), the router proceeds to encapsulate the datagram in the standard DDP-in-UDP encapsulation used by KIP and CAP for transmission across the IP Internet. The router forms the destination IP address by using the IP network number from the table entry and the 8-bit DDP node number. The router also inserts the "real" AppleTalk network number from the table into the destination network field in the DDP datagram. It then transmits the IP datagram.

The datagram proceeds across the IP Internet to a router at the remote site. This router has been advertised as a router for the IP network which is associated with the destination AppleTalk network, so the datagram goes to it. Somehow this router inserts the appropriate AppleTalk network number into the source network part of the

DDP header [I DON'T KNOW HOW IT DOES THIS] and forwards the datagram to the destination AppleTalk network through the local internet.

This scheme has these advantages:

- It uses the existing DDP-in-UDP encapsulation.
- In order for two sites to communicate, each site has to manually enter the other's networks of interest into its mapping table. This provides desirable administrative control.
- By inspecting source IP addresses, a host using DDP-in-UDP (eg CAP) can communicate directly with another DDP-in-UDP host, without requiring routers, after the first few datagrams.
- Each site can have up to 64K (minus the number of internal AppleTalk networks) remote networks with which it can communicate. Since communities of interest will vary, the entire meta-internet can have many more than 64K networks.
- There is a working implementation.

People thought that Jacobson's scheme was very interesting and deserving of more study.

After this discussion, Phil Budne of Shiva volunteered to write a draft RFC describing the current practice of DDP-in-UDP encapsulation.

KIP and Phase II

Karen Frisa from Novell sent to the Apple-IP mailing list a draft proposal for extending the KIP routing and zone information protocols to handle AppleTalk Phase II. There wasn't time to discuss this proposal at this meeting .

Next Meeting

John Veizades proposed that this Working Group have another meeting before the December IETF plenary. A time in the vicinity of the October INTEROP conference was suggested.

Attendees

Philip Budne	phil@shiva.com
Cyrus Chow	cchow@orion.arc.nasa.go
Steve Deering	deering@pescadero.stanford.edu
Robert Elz	kre@munnari.oz.au
Tom Evans	wcc@cup.portal.com

Alf Farnham	carolf@mcescher.unl.edu
Karen Frisa	karen@kinetics.com
Peter Harrison	harrison@miden.ucs.unimelb.edu.au
Van Jacobson	van@helios.ee.lbl.gov
Holly Knight	holly@apple.com
Sam Lam	
Olivier Martin	martin@cearn.cern.ch
Milo Medin	medin@nsipo.nasa.gov
Robert Morgan	morgan@jessica.stanford.edu
Rebecca Nitzan	nitzan@nsipo.nasa.gov
Zbigniew Opalka	zopalka@bbn.com
Allan Oppenheimer	
Brad Parker	brad@cayman.com
Michael Roberts	roberts@educom.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Steve Waldbusser	sw01+@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com
Steve Willis	swillis@wellfleet.com
Allan Young	rcoay@possum.ecg.rmit.oz.au

3.3.4 IP over FDDI (fddi)

Charter

Chair(s):

Dave Katz, dkatz@merit.edu

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Richard Fox/Synoptics

FDDI Minutes

The meeting was solely comprised of a presentation by Caralyn Brown and Doug Bagnall called, "ARP extensions for Dual Mac Stations".

Currently ARP supports a 1-1 mapping of IP addresses to MAC addresses.

FDDI supports the notion of 1-2 mapping of IP addresses to MAC addresses.

Our goal is not to have a TCP connection break when a wrap happens. To meet this objective it was suggested that an extension to the current ARP protocol is needed, where the new ARP protocol supplies more than a 1-1 mapping but a 1-many mapping. An example of this is:

ARP response= ¡ip¿¡mac1,ring1¿¡mac2,ring2¿

One step identified in achieving this is to add a new SNAP value.

At this point 2 approaches were presented and compared.

Solution 1: Hybrid approach

Have a parameter that says that no backward compatibility is to be maintained. Thus, send old style ARP but encode stuff in target fields.

Advantages: only need to send 1 ARP for all cases. Disadvantages: encoding may break some implementations and this solution doesn't scale very well.

Some people said that this method is better solved at layer 3; reply to this was to rewrite layer 3; thus this solution is less radical than rewriting layer 3.

Solution 2: Extended ARP

This solution requires that a new ARP packet be sent out each interface (this packet is called an EARP and is slightly different than the normal ARP packet). After an EARP is sent the station must set a timer and wait for a response. If no response is received then the station must assume that the receiver of the ARP doesn't understand EARPs and so it must send out a normal ARP.

Advantages: backwards compatibility. Disadvantages: may need to send out 2 ARP requests before an answer is received.

Other issues that came up with this solution are:

- When ring wraps/unwraps stations should send ARP to itself to update everybody's ARP table – do this only after a settling period. Some people felt that the SRF frame takes care of this, others not convinced, no resolution.
  At this time we listed advantages of allowing stations to have 2 macs. The 3 identified reasons are:
  - Load balancing (transparent).
  - Transparent error recovery.
  - Dual mac in wrap: you don't know where response came from.
- Need EARP since non-wrapped stations can use wrong ring when a station is wrapped. EARPs keeps effect to wrapped stations only.(??) At this point we got into varied discussions on how wrapped rings and IP do not get along. Some people want to force all single MAC stations to be connected to the primary ring only (or at least on the same ring), others feels that this rule breaks the concept of FDDI.
- It was suggested that we continue to use RFC 1122 for ARP cache handling.

Attendees

Douglas Bagnall	bagnall_d@apollo.hp.com
Alison Brown	alison@maverick@osc.edu
Caralyn Brown	cbrown@ENR.Prime.com
Cho Chang	chang_c@apollo.hp.com
Andrew Cherenson	arc@sgi.com
Cyrus Chow	cchow@orion.arc.nasa.go
Paul Ciarfella	ciarfella@levers.enet.dec.com
Nadya El-Afandi	nadya@network.com
Richard Fox	sytek!rfox@sun.com
Michael Grobe	grobe@kuhub.cc.ukans.edu
Susan Hares	skh@merit.edu
Peter Hayden	hayden@levers.enet.dec.com
Ajay Kachrani	kachrani%regent.dec@decwrl.dec.com
Jay Kadambi	jayk@iwlcs.att.com
John LoVerso	loverso@xylogics.com
Rebecca Nitzan	nitzan@nsipo.nasa.gov
James Reeves	jreeves@synoptics.com
Bill Townsend	townsend@xylogics.com
Bert Williams	bert.synernetics@mailgate.synnet.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
Sijiam Zhang	szhang@cs.ubc.ca

3.3.5 IP over Switched Megabit Data Service (smds)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

SMDS Minutes

Review of Draft Document

The IP over Switched Multi-megabit Data Service (SMDS) Working Group met for three half-day sessions. The majority of the time was spent reviewing the text of a draft document, A Proposed Standard for the Transmission of IP Datagrams over SMDS, written by Dave Piscitello and Joe Lawrence. The configuration assumed in the document was that of a Logical IP Subnet (dubbed an LIS), in which a virtual private network supported by SMDS was treated as an IP network/subnet. The following are the requirements for an LIS configuration:

- All members have the same IP network/subnetwork number.
- All stations within an LIS are accessed directly over SMDS.
- All stations outside of the LIS are accessed via a router.
- For each LIS, a single SMDS group address (`smds$ip_ga`) has been configured that identifies all members of the LIS.

The protocol stack is assumed to be that depicted below in figure 1.

```
--------------------------------------------
|                 IP/ARP                   |
--------------------------------------------
|      Subnetwork Access Protocol (SNAP)   |
--------------------------------------------
|           IEEE 802.2 LLC Type 1          |
--------------------------------------------
|SMDS Interface Protocol (SIP) Level 3 (MAC)|
--------------------------------------------
|               SIP Level 2                |
--------------------------------------------
|               SIP Level 1                |
--------------------------------------------
```

Figure 1

In addition to the SMDS individual address associated with the Subscriber Network Interface (SNI), and to the SMDS group address associated with the LIS, the document referred to a third SMDS group address, the SMDS ARP Request Address (**smds$arp_req**). This group address is set to smds **$ip_ga**, but latter implementations may set the address to a subset of the addresses in the LIS to deal with scaling issues.

The dynamic mapping of 32 bit Internet addresses to 60 bit SMDS addresses is done via Address Resolution Protocol (ARP). ARP requests will be multicast to the **smds$arp_req** address. The ARP parameters which require specification are the following:

ar$hrd	16 bits hardware type code	<to be determined>
ar$pro	16 bits protocol type code	decimal 2048 for IP
ar$hln	8 bits octets in hardware address	decimal 8 for 64 bits
ar$pln	8 bits octets in protocol address	decimal 4 for 32 bits
ar$op	16 bits operation code	1: request 2: reply

Dave Piscitello volunteered to contact Joyce Reynolds to obtain a value for the hardware type code.

An issue arose during the discussion of ARP over SMDS concerning the encoding of the SMDS address in the ARP reply message. Following the precedence of the IP over FDDI Working Group, the document specified that the SMDS address will be carried in "canonical" format, which is the format specified in the IEEE P802.1A/D10 draft standard, in which the least significant bit of the most significant octet is transmitted first. The encoding of the 60 bit address within the SIP L_3 PDU does not conform to the canonical format, and the bits of each octet would have to be reversed. The use of the canonical format is important in transparent bridging, when LANs of a similar address space but of dissimilar address encoding schemes may be bridged. However, the group questioned the utility of transparent bridging between 802 LANs with a 48 bit address space and SMDS with a 60 bit address space. This questionable utility was compared with the potential for confusion caused by the reversal of bits in the SMDS address. In the end, the group decided not to use the canonical format, but instead to use the format specified for the SMDS "MAC" header.

No unresolved issues remained with the document and the group asked Joe Lawrence to incorporate the suggested modifications and to release the document to the email group for confirmation. Joe indicated that he might be able to release the document by mid-August.

Public Connectivity

It was felt that the draft document was adequate to define the operation of IP over small virtual private networks supported by SMDS. Discussion then turned to the issue of "public connectivity," in which an SMDS device may communicate directly with any other SMDS device. The question was asked of this model "What breaks?", and the following items were listed:

- ARP
- Routing: cost, traffic volume, table sizes
- Address management

The group was then asked whether there was any interest in pursuing this problem, and discussion led to an offer by Manoel Rodrigues and George Clapp to draft an "issues" document to attempt to clarify the issues left unresolved by the draft document.

Support of Other Protocols

Vicki Ralls pointed out that other protocols such as DECNET and XNS also need a specification to operate over SMDS, and asked whether this was of interest to the group. The group felt that IP was the appropriate topic for their work and suggested that Bellcore might be approached concerning these other protocols.

Network Management

Dave Piscitello distributed copies of three papers on network management relevant to SMDS.

- Experimental Definitions of Managed Objects for the SMDS Interface Protocol (sip) Interface Type, Kaj Tesink
- Experimental Definitions of Managed Objects for the t3-carrier Interface Type, Tracy Cox, Kaj Tesink
- Internet Draft of T1-Carrier objects, Kaj Tesink, Tracy Cox

These documents were distributed to the Working Group on an informational basis to the. The first two documents had been submitted for consideration by the TransMIB Working Group; the third had not been submitted since the points raised in the document had already been addressed by the TransMIB group.

Future Work

The work remaining for the group will be to review and possibly approve the draft document. The group may be able to approve the document at the upcoming meeting in December and, if possible, begin the process of submitting the document to become an RFC. At the same meeting, the group may review the document to be written by

Manoel Rodrigues and George Clapp.

During the IETF Plenary of Friday morning, August 3rd, Bob Hinden announced the formation of a new Working Group within the routing area, Address Resolution and Routing over SMDS and X.25 Public Data Networks. This group will be chaired by George Clapp and may investigate some of the issues left unresolved by the IP over SMDS Working Group.

Attendees

Douglas Bagnall	bagnall_d@apollo.hp.com
Chet Birger	cbirger@bbn.com
Roger Boehner	Roger.Boehner@StPaul.NCR.COM
Caralyn Brown	cbrown@ENR.Prime.com
Asheem Chandna	ac0@mtuxo.att.com
George Clapp	meritec!clapp@bellcore.bellcore.com
Tracy Cox	tacox@sabre.bellcore.com
Caroline Cranfill	rcc@bss.com
Kevin Fall	kfall@Berkeley.EDU
Michael Fidler	ts0026@ohstvma.ircc.ohio-state.edu
James Forster	forster@cisco.com
Craig Fox	foxcj@nsco.network.com
Eugene Geer	bcr!nvmxr!ewg
Neil Haller	nmh@bellcore.com
Dave Kaufman	dek@proteon.com
Alex Koifman	akoifman@bbn.com
Joseph Lawrence	jcl@sabre.bellcore.com
Walter Lazear	lazear@gateway.mitre.org
Alan Menezes	afm@cup.portal.com
David Piscitello	dave@sabre.bellcore.com
Vicki Ralls	ralls@cisco.com
Michael Reilly	reilly@nsl.dec.com
Ron Roberts	roberts@jessica.stanford.edu
Manuel Rodrigues	
Jim Showalter	gamma@mintaka.dca.mil
Frank Slaughter	fgs@shiva.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Chris Weider	clw@merit.edu
Steve Willis	swillis@wellfleet.com

ARP Extensions for
Single IP Subnet FDDI LANs

Caralyn Brown
(cbrown@enr.prime.com)
Prime Computer
July 31, 1990

Presentation Objectives

Discuss ARP extension alternatives as they have been developed so far.

Collect feedback on these alternatives to develop a clear standard.

Gather support for the suggested ARP extensions to allow a full range of FDDI services for dual MAC, single IP FDDI stations.

Problem to Solve

Provide a new ARP method that would not supersede the existing one-to-one method, but would complement it.

Provide extensions that are flexible enough for one-to-many IP to MAC addressing.

Allow ARP cache to maintain correct mappings during network transitions.

Solution

Extend the ARP protocol to include the bindings between MAC addresses and the externally visible interfaces; i.e. the interfaces to the primary and secondary rings in an FDDI LAN.

Request and reply messages contain the following information
<IP addr X> <MACx1, I1> ... <MACxn, In>

Where the pairings show the MAC address associated with their respective interface attachments. Each interface attachment would have an entry in the request and or reply message.

EARP Request/Response Format

New SNAP value

Additional fields for attachment information

Protocol version number (as suggested at previous IETF meeting)

Station state information

First hardware address corresponds to the interface sending the request.

```
16 bits    Protocol version number
16 bits    Hardware type code
16 bits    Protocol type code
 8 bits    Byte length of hardware
             address (n)
 8 bits    Byte length of protocol
             address (m)
16 bits    Opcode -- request/response
 8 bits    Station state; encoding TBD
 m bytes   Protocol address of sender
16 bits    Count of hardware addresses to
           follow

  for each hardware address
    n bytes    Hardware address of sender
    8 bits     Corresponding path interface

m bytes    Protocol address of target
n bytes    Hardware address of target
```

Backward Compatibility
Approach One - Hybrid Request

Provide tunable parameter - pure EARP network or mixed network. Default to mixed.

For initiating station:
 Is station in EARP-only network?
  YES
   send EARP request and receive EARP response.

  NO
   send hybrid EARP request and accept either ARP or EARP response.

For responding stations:
 Did the request have state and interface information?
  YES
   send EARP response

  NO
   send standard ARP response

Backward Compatibility
Approach One - Hybrid
Continued

```
16 bits    Hardware type code
16 bits    Protocol type code
 8 bits    Byte length of hardware
             address (n)
 8 bits    Byte length of protocol
             address (m)
16 bits    Opcode -- request/response
 n bytes   Hardware address of sender
 m bytes   Protocol address of sender
 n bytes   *** encoding of station state
            and path interface
mbytes     Protocol address of target
```

Path interface and station state in target hardware address space of standard ARP request.

Requesting station formats one for each interface.

Advantage: send only one message type to all stations.

Disadvantage: must receive all to load balance.

Backward Compatibility
Approach Two

Format EARP requests for each interface.

For initiating station:
 Send EARP request.
 Did we receive a response within timeout?
  YES
   record information in EARP cache

  NO
   send standard ARP request and wait for standard ARP reply.

For responding station:
 Does station support EARP?
  YES
   reply with EARP response

  NO
   drop request; unrecognized SNAP value.

Advantage: no intervention to tune system.

Disadvantage: necessary to send two requests for non-EARP stations.

Boundary Condition Processing
Station Transitions - Thru to Wrap

Transitioning station may reduce connectivity. One or more MACs may no longer be accessible.

EARP supporting stations will provide notification of wrapped state via EARP request message.

Only send transition message after station has "settled" into a state; avoid storm due to rapid transitions.

Requested IP address will be that of the wrapping station.

Station state information useful when station is in wrap.

Only end points of network wrap are affected. All other stations are undisturbed.

Boundary Condition Processing
Station Transitions - Wrap to Thru

Transitioning station may increase connectivity.

Table entries for non-EARP stations may no longer be valid.

All EARP stations will continue to communicate across wrap-thru transitions.

Transitiong stations will provide notification of thru state via EARP request message.

Only send transition message after station has "settled" into a state; avoid storm due to rapid transitions.

Requested IP address will be that of the transitioning station.

Entries obtained during wrap state for non-EARP stations are questionable. Remove or timeout.

Boundary Condition Processing
Additional Considerations

Suggestions from RFC 1122 (Requirements for Internet Hosts) under ARP Cache Validation.

1. Timeout cache entry even if it is in use.
2. Unicast poll; clear cache entry if no reply in N successive polls
3. Link layer advice; clear cache entry if link layer detects problem; e.g. no longer setting "A" indicator.
4. Higher layer advice; clear entry if higher layer indicates delivery problem.
5. Notification via EARP requests.
6. Serial MACs.
7. Aliasing; each MAC recognizes the other's address.

Properties of Extended ARP (EARP) Solution

Nodes receiving EARP requests and replies will have sufficient information stored in their EARP caches such that:

There are no false positives; nodes that intend to communicate will use the proper MAC addresses.

Resending of requests and replies are normally not needed as the network transitions.

Dual MAC (or multi MAC in the general case) stations will have sufficient information to perform load balancing.

Summary

Extensions to ARP provide

Method for load balancing over both FDDI rings

Generalized case for multi rail expansion

Method for ARP cache updating after station transitions.

Further comments may be directed to

Caralyn Brown (cbrown@enr.prime.com)
Doug Bagnall (bagnall_d@apollo.hp.com)
Doug Hunt (dhunt@enr.prime.com)
Mary Jane Strohl (strohl@apollo.hp.com)

3.3.6 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:

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

CURRENT MEETING REPORT

Reported by Fred Baker/Vitalink

PPPEXT Minutes

Point to Point MIB:

Discussion ensued on statistics per protocol, per interface. Is there duplication of objects, or a breakage of precedent? The general feeling was that there is need for counts by protocol, precedent or not, and that only some protocols are duplicated elsewhere. Therefore the MIB should contain counts by protocol.

AppleTalk:

Frank Slaughter and Steve Senum have differing approaches to Appletalk. Frank's includes a reduced overhead routing information transfer protocol. They are to coalesce their documents and put them up to the list.

Decnet IV:

At first blush, it would appear that Art Harvey and Steve Senum have dueling documents; however, this appears to be related to several miscommunications. Art is willing to see Steve's document, given certain modifications, as the standard. There are a number of problems with the use of timers in the protocol, resulting primarily from Digital's assumption that a reliable protocol such as LAPB is in use on the line. This will cause problems on unreliable links.

A General Note:

Large interest is reported for Point-to-Point Host to Router implementations over a dial up interface. This, according to Vicki Ralls, is most of the interest cisco has seen in the protocol.

The general feeling toward Art Harvey's proposal for SNAP over Point-to-Point, is that there is no overriding reason to stop the document. We should therefore, let it become a standard for generalized use of the link.

Bridge Protocol

Fred Baker submitted an alternative approach to bridge use of the link. This was generally considered superior to the approach requested by the Pittsburg IETF attendees. A document will be published.

Attendees

Fred Baker	baker@vitalink.com
Tracy Cox	tacox@sabre.bellcore.com
Dino Farinacci	dino@buckeye.esd.3com.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Craig Fox	foxcj@nsco.network.com
Russell Hobby	rdhobby@ucdavis.edu
Ken Jones	uunet!konkord!ksj
David Kashtan	
Frank Kastenholz	kasten@europa.interlan.com
Kathy Kerby	kkerby@bbn.com
Tony Lauck	lauck@tl.dec.com
Solomon Liou	solomon%penril@uunet.uu.net
John LoVerso	loverso@xylogics.com
Alan Menezes	afm@cup.portal.com
Vicki Ralls	ralls@cisco.com
Michael Reilly	reilly@nsl.dec.com
Steve Senum	sjs@network.com
Frank Slaughter	fgs@shiva.com
Glenn Trewitt	trewitt@nsl.dec.com

3.3.7 Router Discovery (rdisc)

Charter

Chair(s):

Steve Deering, deering@pescadero.stanford.edu

Mailing Lists:

General Discussion: gw-discovery@gregorio.stanford.edu
To Subscribe: gw-discovery-request@gregorio.stanford.edu

Description of Working Group:

The Router Discovery Working Group is chartered to adopt or develop a protocol that Internet hosts may use to dynamically discover the addresses of operational neighboring gateways. The group is expected to propose its chosen protocol as a standard for gateway discovery in the Internet.

The work of this group is distinguished from that of the Host Configuration Working Group in that this group is concerned with the dynamic tracking of router availability by hosts rather 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.

CURRENT MEETING REPORT

Reported by Steve Deering/Stanford

RDISC Minutes

Agenda

- Draft Specification
  - comments?
  - disposition?
- Implementations
- Black-Hole Detection

This was the fourth meeting of the Router Discovery Working Group.

The first and dominant item on the agenda was a discussion of the (late) July draft of the ICMP router discovery specification. The following improvements and changes were agreed upon:

- Add a few sentences emphasizing that this is NOT a routing protocol – hosts are expected to rely on Redirects for finding the "best" first-hop router for any given destination.
- Make it even clearer than it already is that hosts must NOT continuously send solicitations.
- Add a note explaining that, even though the timing values are defined or configured in units of seconds, randomized intervals should be computed at the best available resolution of the system's interval timer.
- Fill in the missing ICMP Type values with officially-allocated numbers.
- Change `MAX_RESPONSE_DELAY` from 5 seconds to 2 seconds.
- Change the upper bound on MaxAdvertisementInterval from (2^16 - 1) to 1800 seconds (30 minutes).
- Even when a router is configured to use multicast instead of broadcast, it may respond to a broadcast solicitation with a broadcast advertisement (if not a unicast advertisement).
- When a router performs a graceful shutdown, it should send out advertisements with a lifetime of 0, to flush its addresses from the hosts' router lists.

There was also discussion of adding an authentication field to the Router Advertisement message. Deering argued that such a field could be appended to the existing message format if and when a non-null authentication type is defined for router discovery (i.e., the absence of an authentication field indicates "null" authentication.) Noel Chiappa was not very happy with this proposal, but said he would check it out with the security gurus [which he subsequently did; apparently, Deering's proposed scheme will be acceptable].

The group then agreed that, with the above modifications, the draft specification was ready to enter into the Internet standardization track. Chiappa explained the necessary steps, as follows:

- Update the specification to incorporate the agreed changes and make it available as an Internet Draft as soon as possible.
- After a one month comment period as an Internet Draft, if no significant problems are uncovered, submit it to the IESG with the group's recommendation that it be published as a Proposed Standard.
- Operational experience with multiple, independently-developed implementations is generally required for advancement beyond Proposed Standard status. The decision to advance to the next stage (Draft Standard) is up to the IAB, with advice from the IESG.

That led to the next topic on the agenda: implementations. Andy Cherenson and Deering confirmed their previous commitment to generate an implementation of the protocol to run in user space on 4.3BSD and derived systems, perhaps starting from the source code for cisco's GDP demon; the implementation will include both the host and the router parts of the protocol. John Veizades volunteered to do a Macintosh implementation of the host part of the protocol, and said he had an environment for testing the protocol's behavior under the simultaneous startup scenario (a rack of Macs on a single power circuit). Implementations for other platforms, and at the kernel level in BSD, were solicited, but no promises were made. The importance of getting the major router vendors to implement the router part of the protocol and make it available for user testing was recognized; group members were encouraged to make that desire known to their favorite router vendors.

We then concluded that no further meetings of the Router Discovery Working Group would be necessary, if all goes according to plan. (Yah!!) We discussed the possibility of transforming into a "Black Hole Detection" Working Group, and decided not to do so. A document addressing the wider issue of host routing, of which black hole detection is a part, would be very valuable, but there was little enthusiasm for forming a new Working Group for that purpose; it might be taken up by the next incarnation of the Host Requirements Working Group, or perhaps some individual(s) will generate a document recommending (but not standardizing) good host routing strategies.

ACTION ITEMS

- Deering: Ask the Internet Assigned Numbers Authority for two new ICMP Types.
- Deering: Revise the specification as agreed at this meeting and submit it as

an Internet Draft. If no substantive, negative comments are received during a one month comment period, recommend the specification to the IESG as a Proposed Standard.

- Deering and Cherenson: Implement both the host and router parts of the protocol as a user-level demon for 4.3BSD-derived systems, and make it available to the Working Group and the wider internet community for testing and validation of the protocol.
- Veizades: Implement the host part of the protocol for Macintosh and test it in an environment with many hosts on the same subnet (especially under the simultaneous startup scenario).
- Everyone: Encourage your favorite router vendor to do a prototype implementation of the protocol, for in-house and customer- site testing.

Attendees

Zorica Avramovic	zorica@sparta.com
Art Berggreen	art@opal.acc.com
Larry Brandt	lbrandt@sparta.com
Eric Brunner	brunner@monet.berkeley.edu
Andrew Cherenson	arc@sgi.com
Steve Deering	deering@pescadero.stanford.edu
Robert Elz	kre@munnari.oz.au
Karen Frisa	karen@kinetics.com
Robert Gilligan	gilligan@sun.com
Tony Hain	alh@eagle.es.net
Steven Hubert	hubert@cac.washington.edu
Ole Jacobsen	ole@csli.stanford.edu
Ken Jones	uunet!konkord!ksj
Michael Karels	karels@berkeley.edu
Stev Knowles	stev@ftp.com
Alex Koifman	akoifman@bbn.com
Sam Lam	
Gregory Lauer	glauer@bbn.com
John Lekashman	lekash@orville.nas.nasa.gov
Solomon Liou	solomon%penril@uunet.uu.net
Yoni Malachi	malachi@polya.stanford.edu
Tony Mason	mason@transarc.com
Paul McKenney	mckenney@sri.com
John Moy	jmoy@proteon.com
John Mullen	
Stephanie Price	cmcvax!price@hub.ucsb.edu

Tim Seaver	tas@mcnc.org
Deepinder Sidhu	sidhu@umbc3.umbc.edu
Craig Smelser	
Frank Solensky	solensky@interlan.interlan.com
Martha Steenstrup	msteenst@bbn.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
John Veizades	veizades@apple.com
Carol Ward	cward@spot.colorado.edu
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	ww0n+@andrew.cmu.edu

3.3.8 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:

Aug 1990	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 both volumes.
Feb 1991	Third Internet Draft version of both volumes.

CURRENT MEETING REPORT

Reported by Jim Forster/cisco

RREQ Minutes

The Router Requirements Working Group split their work into two documents; Link Requirements and Router Requirements. The group further considered rewriting and consolidating the IP and ICMP specifications. There was consensus that this was a good thing, but there were doubts as to whether there was time or energy to do it.

The router requirements document was edited on a comprehensive full pass. Issues discussed included:

- Whether a router should support a public SNMP session. It was decided to pass this to the Interconnectivity Working Group.
- The concept of minimum configuration was rejected for this document. The vendor must make sure that when a router comes on line, it does not begin a routing function, without being correctly configured. There needs to be a sanity check on certain parameters.
- There was a discussion of routing preference order between routing protocols, such as IS-IS, and OSPF. The larger question was: Should we specify a routing entry preference? How should a router use a forwarding table and how should it be ordered? There was no consensus on this point. The only thing nearly everyone agreed on was that internal routes ought be preferred over external routes, and that RIP is unsatisfactory. Further, there must be a switch to determine which routing protocol is in charge.
- Congestion control. Note that choosing to drop the last packet on the queue is the worst possible choice via three different experiments. The document will recommend that a router ought to have 2 * bandwidth delay product buffer space in every router along a path. We need to note that a lot of the problem of congestion is poor site engineering. Note that queue length ought not be allowed to rise too long. One needs to go into congestion avoidance if this is occurring. How one then throttles a host is still a problem. We note that there is no benefit to source quench.

Attendees

Stephen Adams	decwrl::"adams@zeppo"
Nick Alfano	nick@gandalf.ca
Cathy Aronson	cja@marmot.nersc.gov
Art Berggreen	art@opal.acc.com

Chet Birger	cbirger@bbn.com
Larry Brandt	lbrandt@sparta.com
Eric Brunner	brunner@monet.berkeley.edu
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Isaac Chan	isaac@gui.consumers.bc.ca
Farokh Deboo	fjd@interlink.com
Steve Deering	deering@pescadero.stanford.edu
Robert Elz	kre@munnari.oz.au
Tom Evans	wcc@cup.portal.com
Jeffrey Fitzgerald	jjf@fibercom.com
James Forster	forster@cisco.com
Karen Frisa	karen@kinetics.com
Vince Fuller	fuller@jessica.stanford.edu
Martin Gross	gross@polaris.dca.mil
Chris Gunner	gunner@osicwg.enet.dec.com
Peter Harrison	harrison@miden.ucs.unimelb.edu.au
Charles Hedrick	hedrick@aramis.rutgers.edu
Ruei-Hsin Hsiao	nac::hsiao
Ole Jacobsen	ole@csli.stanford.edu
Van Jacobson	van@helios.ee.lbl.gov
Michael Karels	karels@berkeley.edu
Paulina Knibbe	knibbe@cisco.com
Stev Knowles	stev@ftp.com
Alex Koifman	akoifman@bbn.com
Sam Lam	
Gregory Lauer	glauer@bbn.com
John Lekashman	lekash@orville.nas.nasa.gov
Solomon Liou	solomon%penril@uunet.uu.net
Yoni Malachi	malachi@polya.stanford.edu
Gary Malkin	gmalkin@ftp.com
Milo Medin	medin@nsipo.nasa.gov
Judy Messing	messing@gateway.mitre.org
David Miller	dtm@ulana.mitre.org
Brad Parker	brad@cayman.com
Craig Partridge	craig@nnsc.nsf.net
Stephanie Price	cmcvax!price@hub.ucsb.edu
Michael Reilly	reilly@nsl.dec.com
Steve Senum	sjs@network.com
Jim Sheridan	jsherida@ibm.com
Deepinder Sidhu	sidhu@umbc3.umbc.edu
Frank Slaughter	fgs@shiva.com

Frank Solensky	solensky@interlan.interlan.com
Tony Staw	staw@marvin.enet.dec.com
Ken Stetten	kstetten@nrao.edu
Bob Stewart	rlstewart@eng.xyplex.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	ww0n+@andrew.cmu.edu
Robert Woodburn	woody@saic.com
Allan Young	rcoay@possum.ecg.rmit.oz.au
Jessica Yu	jyy@merit.edu

3.4 Network Management Area

Director: Dave Crocker/DEC

Area Summary
Reported by Greg Vaudreuil /CNRI

The Network Management area currently has 10 active working groups. Of those groups the Alert Management, Decnet Phase IV MIB, SNMP, FDDI MIB, Transmission MIB, Bridge MIB, Call Accounting, Management Services Interface, Remote Lan Monitoring, Lan Manager, and the OSI Internet Management Working Groups met.

The CMOT document was submitted to the IESG for consideration as a Draft Standard. After discussion, the IESG recommended to the IAB that CMOT be published as a proposed standard. Action by the IAB is still pending.

The SNMP Authentication document was reviewed by the Privacy and Security Task Force. Several problem areas were identified, and work is continuing.

The Alert Management working group submitted their document to the IESG for consideration as a proposed standard. No action has been taken at this time.

3.4.1 Alert Management (alertman)

Charter

Chair(s):

Louis Steinberg, louiss@ibm.com

Mailing Lists:

General Discussion: alert-man@merit.edu
To Subscribe: alert-man-request@merit.edu

Description of Working Group:

The Alert Management Working Group is chartered with defining and developing techniques to manage the flow of asynchronously generated information between a manager (NOC) and its remote managed entities. The output of this group should be fully compatible with the letter and spirit of SNMP (RFC 1067) and CMOT (RFC 1095).

Goals and Milestones:

Done	Develop, implement, and test protocols and mechanisms to prevent a managed entity from burdening a manager with an unreasonable amount of unexpected network management information. This will focus on controlling mechanisms once the information has been generated by a remote device.
Done	Write an RFC detailing the above, including examples of its conforment use with both SNMP traps and CMOT events.
May 1990	Develop, implement, and test mechanisms to prevent a managed entity from generating locally an excess of alerts to be controlled. This system will focus on how a protocol or MIB object might internally prevent itself from generating an unreasonable amount of information.
Dec 1990	Write an RFC detailing the above. Since the implementation of these mechanisms is protocol dependent, the goal of this RFC would be to offer guidance only. It would request a status of “optional”.

3.4.2 Bridge MIB (bridge)

Charter

Chair(s):

Fred Baker, baker@vitalink.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:

May 1990	Publish initial proposal
Nov 1990	Submit an Internet Draft
Feb 1991	Submit draft for RFC publication

CURRENT MEETING REPORT

Reported by Fred Baker/Vitalink

BRIDGE Minutes

The SNMP, Bridge MIB, and Transmission MIB Working Groups each met during a single Working Group session. During the short bridge MIB meeting, Paul Langille presented his work on the X.25 Bridge Entity Model.

Attendees

See the SNMP Minutes

3.4.3 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:

Jul 1990	Mailing list discussion of charter and collection of concerns.
Aug 1990	Discussion and final approval of charter; discussion and agreement on models and terminology. Make writing assignments.

Nov 1990	First draft document, discussion, additional drafts, special meeting?
Dec 1990	Review latest draft and if OK, give to IESG for publication as RFC.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

CHARMIB Minutes

Agenda

- Do we have the right starting organization?
  - Working group position in IETF hierarchy.
  - Chairman.
  - Participants.
  - Editor/author.
- Is this the right problem?
  - Character stream devices, not just terminals. That means modems, printers, RS-232, 3270, virtual ports, etc.
  - All systems, not just terminal servers. That means general- purpose hosts, bridges with a single console port, etc.
- Existing work to consider?
  - Draft standard MIBs.
  - Private MIBs?
- Technical issues?
  - List of interesting, common, reasonable information.
  - Relationship to Interface Group, considering SLIP.

To the questions "Do we have the right starting organization?" and "Is this the right problem?", the answer (by lack of disagreement) was yes. Similarly, the charter was accepted unchanged. The consensus was that this is useful, important work, and we can quickly come to a useful agreement.

The request for "Existing work to consider?" brought useful contributions from those in attendance, particularly from Bill Westfield of cisco who provided their private terminal MIB. The consensus was that the various existing private MIBs are quite similar, with most differences considered as desirable additions.

The "Technical issues?" topic resulted in sufficient conclusions for the following first-draft MIB model. Character devices are a separate group, analagous to the Interface Group. The group contains physical and logical ports in one table, indexed by sequential integers, with their real identification and type as objects in the table. Each table entry contains such objects as counters for characters in and out, parity errors, and framing errors. It has configuration information such as parity, speed, and bits per character. It also has status information, such as the state of modem control signals. The Character Group also contains session information for each session on each port.

Character devices that support SLIP have a corresponding entry in the Interface Group, which uses the MIB-II object ifSpecific to point to the corresponding character MIB entry. When SLIP is active, the Interface Group entry has an ifOperStatus value of "up". When SLIP is inactive, the Interface Group status is "down".

The group agreed to have a working meeting at the INTEROP conference.

Those who have private terminal MIBs that have not been submitted to the group are to do so as quickly as possible. If their company requires confidentiality, such submissions can be made through Marshall Rose, who will preserve anonymity. Bob Stewart is to provide a statement of the working model, as outlined above.

The next milestone in the charter is a first draft by November. Given the above-mentioned submissions, I will attempt to prepare the draft by the beginning of October, so it can be reviewed at our INTEROP meeting.

Attendees

Anthony Chung	anthony@hls.com
George Conant	geconant@eng.zyplex.com
John Cook	cook@chipcom.com
James Davin	jrd@ptt.lcs.mit.edu
David Jordan	...jordan@emulex.com
Satish Joshi	sjoshi@mvis1.synoptics.com
Frank Kastenholz	kasten@europa.interlan.com
John LoVerso	loverso@xylogics.com
Keith McCloghrie	kzm@his.com
Donald Merritt	don@brl.mil
David Perkins	dave_perkins@3com.com
Marshall Rose	mrose@psi.com
Bob Stewart	rlstewart@eng.xyplex.com
Bill Townsend	townsend@xylogics.com
Bill Westfield	billw@cisco.com

3.4.4 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):

Jonathan Saperia, `saperia%tcpjon@decwrl.dec.com`

Mailing Lists:

General Discussion: `phiv-mib@jove.pa.dec.com`
To Subscribe: `phiv-mib-request@jove.pa.dec.com`

Description of Working Group:

The DECNet Phase IV MIB Working Group will define MIB elements in the experimental portion of the MIB which correspond to standard DECNet Phase IV objects. The group will also define the access mechanisms for collecting the data and transforming it into the proper ASN.1 structures to be stored in the MIB.

In accomplishing our goals, several areas will be addressed. These include: Identification of the DECNet objects to place in the MIB, identification of the tree stucture and corresponding Object ID's for the MIB elements, Generation of the ASN.1 for these new elements, development of a proxy for non-decnet based management platforms, and a test implementation.

Goals and Milestones:

Done	Review and approve the charter and description of the Working Group, making any necessary changes. At that meeting, the scope of the work will be defined and individual working assignments will be made.
Sep 1991	Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. If possible, prototype implementation to begin after revisions have been made.
Dec 1990	Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.
Mar 1990	Review final draft and if OK, give to IESG for publication as RFC.

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

CURRENT MEETING REPORT

Reported by Jon Saperia/DEC

DECNETIV Minutes

1. An early draft with 28 groups was distributed for discussion purposes, so that we could begin the process of removing redundant or unnecessary variables.
2. It was agreed that we would reorganize the MIB into groups that correspond to the various layers of software found in DECNet Phase 4. For example, the X.25, Network, Session, Routing, Data Link, and End Communication Layer Groups. This will also make it easier to use the same approach to optional and mandatory variables that is used for the Internet Standard MIB. For example, X.25 and all variables in that branch of the tree will be mandatory in implementations that support X.25 and not required for those implementations which do not provide X.25 service. More work is needed in this area and I will attempt to recast what we have defined into these groups.
3. Several people expressed the desire to keep the total number of variables down to less than 80. We will attempt this, however; since a prime purpose of the MIB is to allow DECNet Phase IV objects (including end systems) to be managed via SNMP, more DECNet variables will have to be implemented for the MIB than are currently found in some of the implementations in router products.
4. Each branch of the tree will be further devided into three sub-groups, these will be the parameters, counters and events sub-groups. In order to support the events sub-groups we will be defining DECNet Phase IV traps. Steve Willis will be writing up something to cover experimental trap id's.
5. For the sake of consistency each variable will have deciv prepended to it.
6. There will be a Working Group meeting before the October INTEROP time-frame so that these changes can be reviewed. Since a number of vendors have already implemented some portion of a DECNet MIB in their proprietary MIBs this will be an opportunity to merge them.
7. Where information is available in other MIBs, we will not include that as part of the DECNet phase IV mib. An example of this is the new ethernet MIB.
8. After the meeting, it was suggested that we may want to consider publishing the MIB in portions such as the Network Layer or DECNet Phase IV Routing MIB rather than waiting to do the entire piece at once. Comments on this appoach would be appreciated.
9. Members of this list will be contacted separately to set up the September Meeting.

Attendees

Chris Chiotasso	chris@sparta.com
Farokh Deboo	fjd@interlink.com
Nadya El-Afandi	nadya@network.com
Stanley Froyd	sfroyd@salt.acc.com
Charles Hedrick	hedrick@aramis.rutgers.edu
Steven Hunter	hunter@ccc.mfecc.arpa
David Perkins	dave_perkins@3com.com
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Steve Willis	swillis@wellfleet.com

3.4.5 FDDI MIB (fddimib)

Charter

Chair(s):

Jeffrey Case, case@utkux1.utk.edu

Mailing Lists:

General Discussion:
To Subscribe:

Description of Working Group:

No description available

Goals and Milestones:

none specified

3.4.6 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:

May 1990	Policy Models Examined.
Aug 1990	Meter Working Draft Written.
Nov 1990	Collection Protocols Working Papers Written.
Feb 1991	Meter Final Draft Submitted.
Feb 1991	Collection Protocol Working Papers Reviewed.
May 1991	Collection Protocol Recommendation.

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

ACCT Minutes

Agenda:

Wednesday	Reports and Presentations.
	Review of Document Outlines.
Thursday	Review of Meter Services Draft.

Summary:

Don Hirsh reported on findings for developing a LAN accounting resource and conducted a review of existing network accounting systems. (Slides attached.) The Internet Accounting Working Group reviewed the first draft of the Meter Services document and proposed some modifications. Group members received copies of the ISO accounting meter function and accounting document drafts.

Action Items during Next Period (ending Dec 1, 1990):

- Meter Services: C.Mills
  - Revise Internet Accounting Background Draft
  - Revise Internet Accounting Architecture Draft
  - Revise Meter Services Draft
- Collection Protocol: M.Dubetz
  - Write first draft

Attendees

Dave Crocker	dcrocker@nsl.dec.com
Martin Dubetz	dubetz@wugate.wustl.edu
Tony Hain	alh@eagle.es.net
Neil Haller	nmh@bellcore.com
Brian Handspicker	bd@vines.enet.dec.com
Don Hirsh	hirsh@magic.meridianpc.com
Joel Jacobs	jdj@mitre.org
Ken Jones	uunet!konkord!ksj

Keith McCloghrie	kzm@his.com
Cyndi Mills	cmills@bbn.com
Mark Seger	seger@mjs1/ogo.dec.com
Martha Steenstrup	msteenst@bbn.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
Jessica Yu	jyy@merit.edu

Q: What is existential hell for a "Live Free or Die" kind of a computer scientist?

A Promiscuous LAN Accounting Resource

Your congenial presentor:
Don Hirsh, hirsh@meridiantc.com

Meridian
TECHNOLOGY CORPORATION

This work has been funded by the National Science Foundation, ISI-8960397

Outline of Presentation

- Current Network Accounting Practices
  FARNet Study
  WU
  BBN/Milnet
- Review NSF Sponsored Research on LAN Accounting

Meridian
TECHNOLOGY CORPORATION

Accounting Practice in the Field

Have found 3 sources of USEFUL information.

- FARNet Study
- Washington U. Billing Practices
- BBN/SRI Milnet Experiences

Meridian
TECHNOLOGY CORPORATION

Summary of FARNet Study

General Summary:
- Operations - we have 'em!
- Commercial and Public Sector relations - we're for 'em!
- Cost recovery - we're agin' it, but we've published acceptable use statements that conform to NSF recommendations.

Informal survey of regional network operation, policy, management - conducted in 10-11/1988.

Questions in 4 areas:
- Operations
- Public Sector Relations
- Commercial Relations
- Usage/Charges

Meridian
TECHNOLOGY CORPORATION

WU Billing in a Nutshell

Stand up and talk, Martin Dubetz.

Meridian
TECHNOLOGY CORPORATION

BBN/SRI Experiences

Zaw-Sing Su and Cindy Mills,
come on down!

Meridian
TECHNOLOGY CORPORATION

NSF Proposal

"N-Level Protocol Parsing in Real-Time:
A Framework for Local Network Resource Accounting"

Design a system capable of decoding LAN frames all the way up to the application "on the fly." Think of it as a successor to Braden's statspy.

We think such a system will be a useful, configurable tool for IAWG experiments.

Meridian
TECHNOLOGY CORPORATION

What We Did

System Architecture
Network Characterization
Protocol Characterization
Processing Characterization
Monitor Design

Meridian
TECHNOLOGY CORPORATION

Why Do Accounting?

Perspectives:

- Providers - Recover Costs Equitably
- Users - Maximize Utility/Minimize Expense
- CPAs - Because it's There

Meridian
TECHNOLOGY CORPORATION

What Does One Account For?

Device	the behavior of a data-link peer
Network	the behavior of a particular network address associated with a particular device or interface
User	a process owned by a unique user-id at a unique network or device address

Network and Device accounting are often isometric, but not always.

Meridian
TECHNOLOGY CORPORATION

What Should One Account For?

Device	frames and their attributes
Network	packets and their attributes transactions and their attributes
User	transactions and their attributes

Meridian
TECHNOLOGY CORPORATION

What Can One Account For?

Accounting Feature	TELNET	RLOGIN	CTERM	LAT	HPS (XNS ts	OSI NVT
A. Device accounting						
Source Device Identification	Y	Y	Y	Y	Y	Y
Dest. Device (datalink addr	Y	Y	Y	Y	Y	Y
B. Network accounting						
Source Network Address	Y	Y	Y	N	Y	?(Y)
Dest. Network Address	Y	Y	Y	N	Y	?(Y)
C. Process accounting						
Source Process (port,etc.)	Y	?(Y)	Y	?	Y	?(Y)
Dest. Process	Y	?(Y)	Y	?	Y	?(Y)
D. User accounting						
Source User Identification	N	?	Y (multiuser)	?	?	?
Dest. User Identification	Y (very hard)	?	?	?	?	?
Connection-based	Y	Y	Y	Y	Y	?(Y)
Static protocol identification	Y	?(Y)	N	?(Y)	Y	?

Meridian
TECHNOLOGY CORPORATION

Protocol Characterization

The Essential Problem for General Purpose N-Level Decoding:

oftenthereisn'tanyfielddemultiplexor

Meridian
TECHNOLOGY CORPORATION

What Can One Account For?

Accounting Feature	TELNET	RLOGIN	CTERM	LAT	HPS (XNS ts)	OSI NVT
A. Device accounting						
Source Device Identificatior	Y	Y	Y	Y	Y	Y
Dest. Device (datalink addr)	Y	Y	Y	Y	Y	Y
B. Network accounting						
Source Network Address	Y	Y	Y	N	Y	?(Y)
Dest. Network Address	Y	Y	Y	N	Y	?(Y)
C. Process accounting						
Source Process (port,etc.)	Y	?(Y)	Y	?	Y	?(Y)
Dest. Process	Y	?(Y)	Y	?	Y	?(Y)
D. User accounting						
Source User Identification	N	?	Y (multiuser)	?	?	?
Dest. User Identification	Y (very hard)	?	?	?	?	?
Connection-based	Y	Y	Y	Y	Y	?(Y)
Static protocol identification	Y	?(Y)	N	?(Y)	Y	?

Meridian
TECHNOLOGY CORPORATION

A Parse Tree

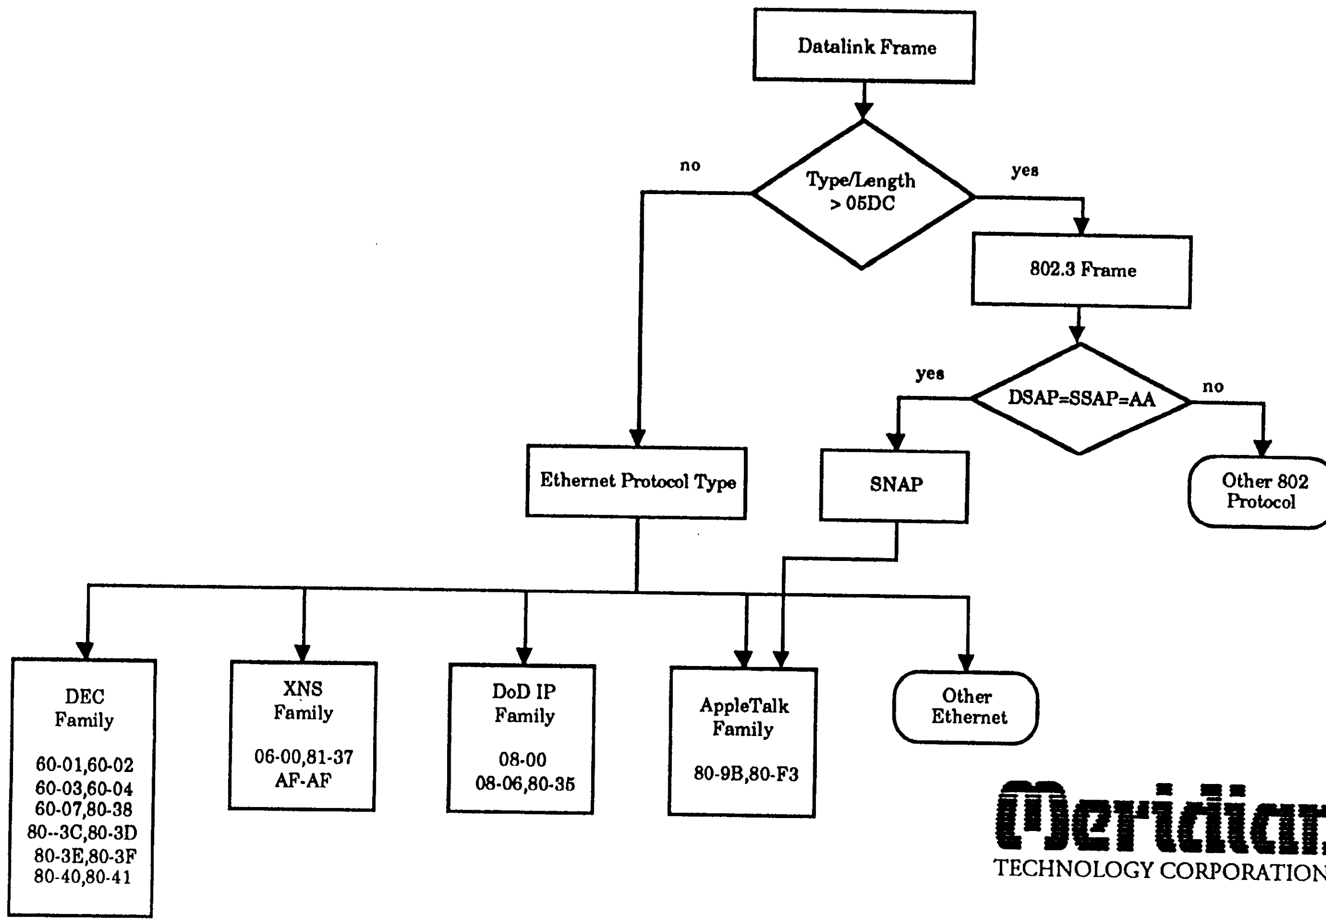

Meridian TECHNOLOGY CORPORATION

Protocol Frames

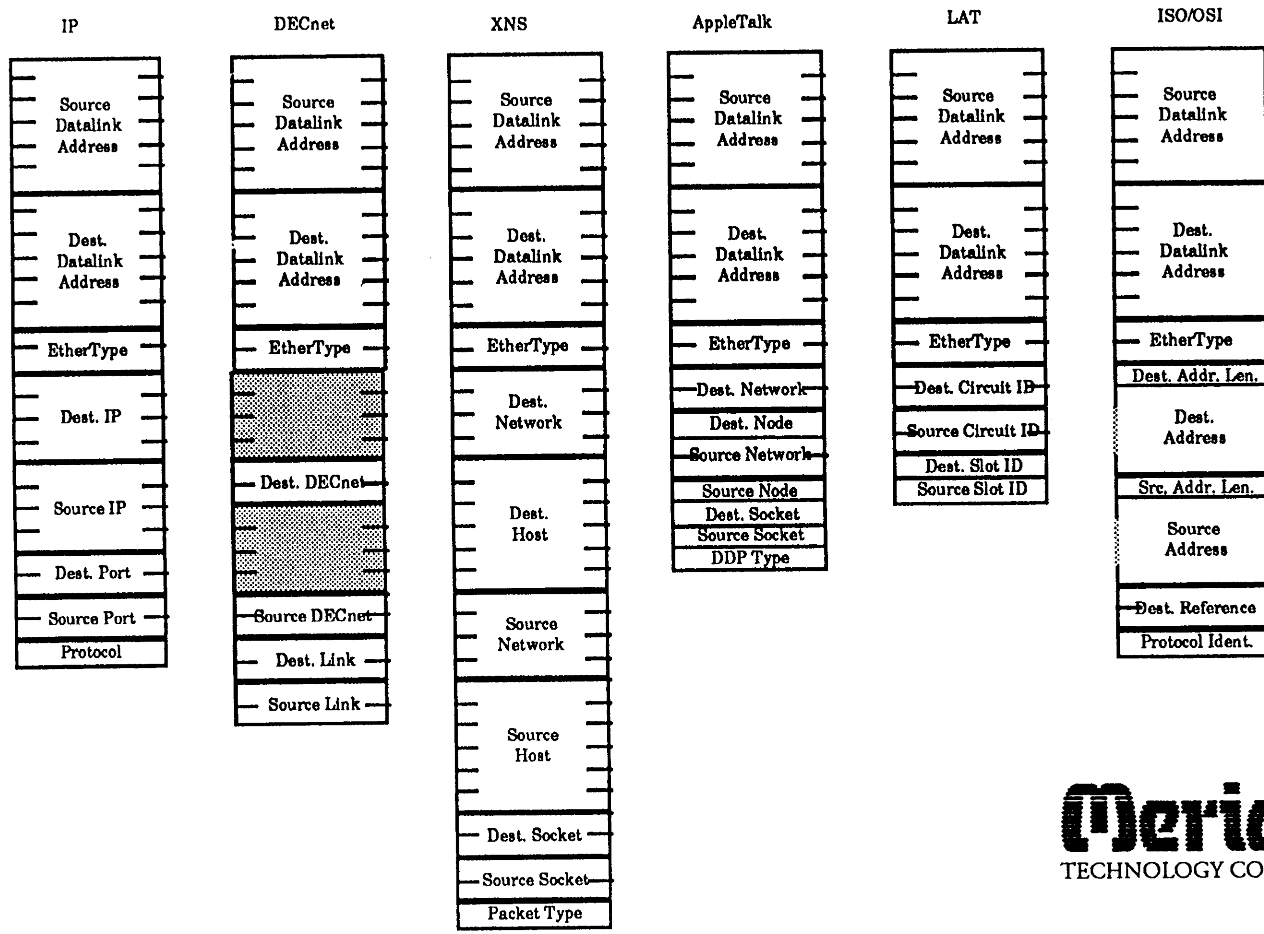

Meridian TECHNOLOGY CORPORATION

Processing Characterization

Stateless parsing: equivalent to traversing a tree from root to leaf, with demultiplexors as branches

Stateful parsing: same notion of traversing a tree but there is memory at different places (leaf, vertex...) requires more complex structures and computations.

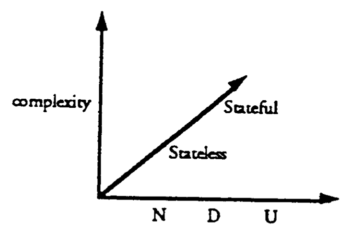

Meridian TECHNOLOGY CORPORATION

System Architecture

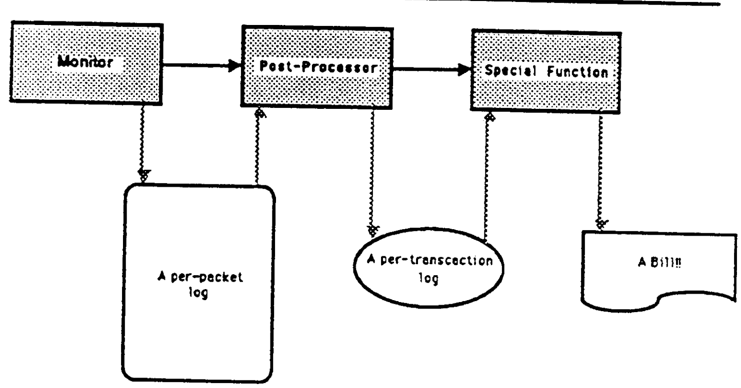

Meridian TECHNOLOGY CORPORATION

Network Characterization

LAN studies w/distributed monitors and standard protocol analyzers. 2 important observations:

1) ratios of multicast to "unicast" is large
2) size distribution is bi-modal

Meridian
TECHNOLOGY CORPORATION

3.4.7 LAN Manager (lanman)

Charter

Chair(s):

David Perkins, dave_perkins@3com.com

Mailing Lists:

General Discussion: lanmanwg@cnd.hp.com
To Subscribe: lanmanwg-request@cnd.hp.com

Description of Working Group:

This Working Group is chartered to define and maintain the MIB and relevant related mechanisms needed to allow management overlap between the workgroup environment (LAN Manager based) and the enterprise environment (based on TCP/IP management).

This translates into three basic objectives:

- Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
- Develop requirements for additional network management information, as needed, and work to extend the LAN Manager interfaces to support such information.

Goals and Milestones:

TBD	Define a minimal set of MIB objects using the existing LAN Manager APIs and file system APIs for LAN manager version 1.x. Start MIB in standards track.
TBD	Define an upwards compatible MIB for LAN Manager version 2.x.
TBD	Work to influence Microsoft, the developer of LAN Manager, to add/change APIs so that MIB developed can be consistant in style and information content with MIBs developed by other MIB Working Groups.

none specified

CURRENT MEETING REPORT

Reported by Dave Perkins/3Com

LANMAN Minutes

New Chair:

Jim Greuel from Hewlett Packard, the previous Chair, was unable to attend. In mail messages he indicated that he would no longer be able to participate and nominated Dave Perkins as a replacement. This nomination was approved by the Working Group.

Lan Manager I MIB:

The current MIB was posted in Internet-Drafts. There are two parts. The names of the two documents are:

- draft-ietf-lanman-mib-00.txt
- draft-ietf-lanman-alerts-00.txt

These MIBs were briefly reviewed and appeared to be in great shape. The next step is to encourage more implementations of these MIBs and start them in the standards track.

Lan Manager II MIB:

Eric Peterson from Microsoft posted a proposal for LAN Manager II MIB before the meeting for review. Most of the time spent in the Working Group was spent in reviewing his proposal. Eric had taken the LAN Manager I MIBs and combined them together and added some new information that is available in LAN Manager version 2.x. The selection rules that he used were:

- Keep the total number of MIB variables below 200
- Define primarily read-only objects
- Add the objects that "real" network managers use in day to day operations

The feedback on the proposal was the following:

- The new MIB must be upwards compatible with the first version.
- Add table of currently logged on users.
- Add a group that contained information about the current domain. Include in it a table that has the list of all servers in a domain.
- Variables should be added so that the information in TRAPs can be determined via polling.

- Check on adding tables that list 1) all the Users at a server, 2) all the USEs at workstations, 3) all the SHAREs at a server,

Eric will post an updated proposal by September 15th so that a meeting, if necessary, can be scheduled during the INTEROP show in October.

Next Meeting:

Depending on demand, a meeting will be held during the INTEROP show (Oct 8-12) or at the next IETF meeting in Colorado (Dec 3-7).

Attendees

Jonathan Biggar	jon@netlabs.com
Theodore Brunner	tob@thumper.bellcore.com
Chris Chiotasso	chris@sparta.com
Cyrus Chow	cchow@orion.arc.nasa.go
Dave Crocker	dcrocker@nsl.dec.com
Dwaine Kinghorn	microsoft!dwaink
David Perkins	dave_perkins@3com.com
Eric Peterson	microsoft!ericpe
Jim Reinstedler	jimr@ub.com
Marshall Rose	mrose@psi.com
Mark Sleeper	mws@sparta.com
Mark Wood	mark1@iw/cs.att.com

3.4.8 Management Services Interface (msi)

Charter

Chair(s):

Oscar Newkerk, newkerk@decwet.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 Oscar Newkerk/DEC

MSI Minutes

The Management Services Interface Working Group met to discuss the latest revision of the draft API document. The following actions were taken.

- The section on authentication information (Section B.4) was modified to indicate that the authentication information should be passed as an AVL instead of the previously indicated C structure.
- The rough outline of an interoperability statement. This statement will document the requirement that implementations of MSI behave the same regardless of the underlying protocol.
- The Interoperability statement effort produced an issue that must be addressed in order for an implementation of the MSI API to function. In order for an implementation of MSI to behave the same regardless of the protocol, the MIB for new classes must be defined in both SNMP terms and OIM terms. This requires that objects and events be documented in the current template format as well as the ISO GDMO format. Without this information, it is impossible for an MSI implementation to 'translate' a management request into both an SNMP PDU and a CMOT PDU. There was no resolution of this issue, but it was agreed that it should also be raised in the OIM Working Group meeting the next day and raised to the Network Management Area Chair.

Comments on the API draft from the UBC meeting will be incorporated in the next revision of the draft.

Attendees

Stephen Adams	decwrl::"adams@zeppo"	
Amatzia Ben Artzi		
Roger Boehner	Roger.Boehner@StPaul.NCR.COM	
Jeffrey Buffum	jbuffum@apollo.hp.com	
Stanley Froyd	sfroyd@salt.acc.com	
Satish Joshi	sjoshi@mvis1.synoptics.com	
Jay Kadambi	jayk@iwlcs.att.com	
Lynn Monsanto		
Oscar Newkerk	newkerk@decwet.dec.com Bill Nowicki	\> \verb nowicki
James Reeves	jreeves@synoptics.com	
Jim Reinstedler	jimr@ub.com	

Raphael Renous	
Jim Sheridan	jsherida@ibm.com
Cheng Song	song@ibm.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Denis Yaro	DYARO@SUN.COM

3.4.9 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.

CURRENT MEETING REPORT

Reported by Brian Handspicker/Digital

OIM Minutes

Agenda

- OIM-MIB-II
- Security Association Policy (CMOT/CMIP)
- MIB/Interoperability Issues
- CMOT Revision/Interoperability Test Presentation

OIM-MIB-II

System Title:

Sysname should be GET-REPLACE but should have a strong health warning about the inadvisability of changing sysname.

Note, there was some discussion of the ISO/ANSI position that the system title should be a Distinguished Name (DN). We agreed that once the Internet defines a containment tree root for Director Services, the OIM group will define a new attribute for the system object called System ID (?), which will be a DN. This DN may include as one of its components, sysname.

One member of the AT&T Bell Labs group raised the concern that NMF uses OID for system title. This did not get much support.

CREATION/DELETION for Objects Defining Containment Hierarchy:

All okay except TCPConn Entry and UDP Entry need DELETE, but do not need CREATE.

Progression of OIM-MIB-II to RFC: Questions but no objections.

Proposing OIM-MIB-II as Proposed Standard: No objections.

Add in statement in status: "This RFC obsoletes MIB definitions included in RFC 1095 (CMOT)".

Security:

Need to restrict Masquerade, Modification, Disclosure ANSI X3T5.7—ISO SC21/WG1—(NIST)OIWNMSIG/SECURITY SIG—O

X.500 Hash Function, Public Key Encryption

```
                v---- ROS/CMIP -----------v
+---------------------------------------+
| hash    |          |          |       |
+---------------------------------------+
```

Need coordination in future with NMF. Looking for proposed solutions/implementations for experimentation.

Association Policy:

Proposed replace existing CMOT/CMIP ACN's with ACN's defined in ISO DIS 10040 (SMO). This means the RFC1095 ACN is no longer defined in the CMOT revision (though still defined in RFC1095). Also the original new 4 ACN has been reduced to 3 ACN's: agent, manager and agent-manager. NIST OIW Association Policy likely to move to Stable Agreement in January 1991.

ACTION: BDH to revise Association Policy

MIB/Interoperability Issues:

MSI requires MIBS to be defined such that OSI SMI and IETF mappings, attributes and objects (?) are defined. This places a requirement either on all the MIB definition groups or on OIM for these mappings.

Fiction: BDH, Lee, etc., to provide the How To Write A MIB document by the next meeting.

ACTION: Lee to define appendices to FDDI MIB, etc., with OSI SMI Mappings for MIBS.

CMOT Revision Interoperability Testing Presentation

Repeated presentation of the plenary presentation. Announced intent to hold another round of Interoperability testing. U.C. London suggested as another potential participant. Four of the participants supporting RFC1095 Revision replacement as Draft Standard. No one objected to its replacement as Draft Standard. CMOT/CMIP 1095 Revision based on NIST OIW IA's.

GDMO Templates

Need tools for converting IETF SMI MO definitions to ISO SMI and vice versa. Jeff Case suggested that his research project may be able to provide public domain tools for this.

System Management Functions

Lee to distribute NIST OIW proposed implementors agreements for System Management Functions. Proposed OIM based SMF IA's on these OIW IA's.

Attendees

Jonathan Biggar	jon@netlabs.com
Yvonne Biggar	yvonne@cam.unisys.com
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Buffum	jbuffum@apollo.hp.com
Jeffrey Case	case@utkux1.utk.edu
Asheem Chandna	ac0@mtuxo.att.com
Mike Erlinger	mike@mti.com
Chris Gunner	gunner@osicwg.enet.dec.com
Brian Handspicker	bd@vines.enet.dec.com
Alex Koifman	akoifman@bbn.com
Lee LaBarre	cel@mbunix.mitre.org
Don McWilliam	
Alan Menezes	afm@cup.portal.com
Lynn Monsanto	
Oscar Newkerk	newkerk@decwet.dec.com
Jim Reinstedler	jimr@ub.com
Harvey Shapiro	shapira@wnyosi2.arpa
Jim Sheridan	jsherida@ibm.com
Mark Sleeper	mws@sparta.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Justin Walker	justin@apple.com
Denis Yaro	DYARO@SUN.COM

3.4.10 Remote LAN Monitoring (rlanmib)

Charter

Chair(s):
Mike Erlinger, `mike@mti.com`

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

Description of Working Group:

The LAN Monitoring MIB Working Group is chartered to define an experimental MIB for monitoring LANs.

The Working Group must first decide what it covers and what terminology to use. The initial thought was to investigate the characteristics of some of the currently available products (Novell's LANtern, HP's LanProbe, and Network General's Watch Dog). From this investigation MIB variables will be defined. In accomplishing our goals several areas will be addressed. These include: identification of the objects to place in the MIB, identification of the tree structure and corresponding Object ID's for the MIB elements, generation of the ASN.1 for these new elements, and a test implementation.

Goals and Milestones:

Jul 1990	Mailing list discussion of charter and collection of concerns.
Aug 1990	Discussion and final approval of charter; discussion and agreement on models and terminology. Make writing assignments.
Dec 1990	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1990	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Michael Erlinger/Micro Technology

RLANMIB Minutes

This was the first meeting of this Working Group. The activities centered on getting to know one another and brainstorming on the concept of remote LAN monitoring and associated problems. The following lists some of the major discussion topics:

- The features of two monitoring boxes, HP LAN Probe and Novell LANtern, were presented by members of the audience as representative of the marketplace.
- The concepts of LAN monitoring and packet capture with packet analysis (e.g., SNIFFER) were discussed. Packet capture and analysis by remote LAN probes seemed to stretch the data transfer capabilities of SNMP.
- The concept of filtering was discussed in great detail. In particular, the approaches to filtering by various manufacturers and the different approaches to combining filters and traps.
- Relationships to other IETF Working Groups were discussed and the need for close interaction was noted by all.
- Finally, there was much discussion about remote LAN monitoring and SNMP. In particular: do smart agents violate the spirit of SNMP? How does a large amount of agent-captured data move to an NMS? and how can various probes and NMSs be synchronized within a particular LAN?

There was no attempt to reach any consensus on these issues, but the group did come up with the following action items:

1. The list of attendees would be sent to rlanmib-request for addition to the mailing list. Done.
2. The group would like to change its name to the Remote LAN Management Working Group and make appropriate charter changes. Mike Erlinger took on this assignment.
3. Steve Waldbusser will attempt to generate a review document on the various network probes. He will try to discern the common features of these devices.
4. Mike Erlinger will attempt to generate a review document on the concept of filters as they apply to remote probes/agents.

Attendees

Scott Bradner	sob@harvard.harvard.edu
Phil Budne	phil@shiva.com

Martin Dubetz	dubetz@wugate.wustl.edu
Mike Erlinger	mike@mti.com
Michael Fidler	ts0026@ohstvma.ircc.ohio-state.edu
Olafur Gundmundsson	ogud@cs.umd.edu
Ken Jones	uunet!konkord!ksj
Kathy Kerby	kkerby@bbn.com
Jim Kinder	jdk@fibercom.com
Paul Langille	quiver::langille@decwrl.dec.com
Cyndi Mills	cmills@bbn.com
Robert Morgan	morgan@jessica.stanford.edu
Robert Pinna	bwp!hpctlb.hp.com
K.K. Ramakrishnan	rama%erlang.dec.com@decwrl.dec.com
Marshall Rose	mrose@psi.com
Glenn Trewitt	trewitt@nsl.dec.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Mark Wood	mark1@iw/cs.att.com

3.4.11 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 Marshall Rose/PSI

SNMP Minutes

Met jointly with Transmission Working Group and (newly formed) Bridge Working Group.

A draft "SNMP Implementation Profile Questionnaire" was handed out for comment.

Experimental MIB's discussed:

- SNMP Views
- Generic Interface Extensions
- T1-Carrier
- Ethernet
- Token Ring
- Token Bus

Actions:

- All above Experimental MIB's are to receive final editing and then be submitted to Internet-Drafts.
- Tracy Cox of Bellcore will submit drafts of Experimental MIBs for T3 and IP over SMDS.
- Rich Fox of Synoptics will submit a draft of a proxy-by-community proposal.

Attendees

Stephen Adams	`decwrl::"adams@zeppo"`
William Anderson	wda@mitre-bedford.org
Amatzia Ben-Artzi	amatzia@synoptics.com
Jonathan Biggar	jon@netlabs.com
Chet Birger	cbirger@bbn.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Asheem Chandna	ac0@mtuxo.att.com
Chris Chiotasso	chris@sparta.com
Cyrus Chow	cchow@orion.arc.nasa.go
Paul Ciarfella	ciarfella@levers.enet.dec.com
Rob Coltun	rcoltun@trantor.umd.edu

George Conant	geconant@eng.zyplex.com
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Nadya El-Afandi	nadya@network.com
Mike Erlinger	mike@mti.com
Jeffrey Fitzgerald	jjf@fibercom.com
Richard Fox	sytek!rfox@sun.com
Stanley Froyd	sfroyd@salt.acc.com
James Galvin	galvin@tis.com
Jim Goetz	goetz@hav2d.att.com
Hellmut Golde	golde@june.cs.washington.edu
Chris Gunner	gunner@osicwg.enet.dec.com
Peter Harrison	harrison@miden.ucs.unimelb.edu.au
Peter Hayden	hayden@levers.enet.dec.com
Ruei-Hsin Hsiao	nac::hsiao
Steven Hunter	hunter@ccc.mfecc.arpa
Satish Joshi	sjoshi@mvis1.synoptics.com
Jay Kadambi	jayk@iwlcs.att.com
Frank Kastenholz	kasten@europa.interlan.com
Kathy Kerby	kkerby@bbn.com
Jim Kinder	jdk@fibercom.com
Paul Langille	quiver::langille@decwrl.dec.com
John LoVerso	john@loverso.leom.ma.us
Keith McCloghrie	kzm@his.com
Don McWilliam	
Donald Merritt	don@brl.mil
Paul Mockapetris	pvm@isi.edu
David Perkins	dave_perkins@3com.com
Robert Pinna	bwp!hpctlb.hp.com
David Piscitello	dave@sabre.bellcore.com
James Reeves	jreeves@synoptics.com
Jim Reinstedler	jimr@ub.com
Raphael Renous	
Robert Reschly	reschly@brl.mil
Marshall Rose	mrose@psi.com
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Harvey Shapiro	shapira@wnyosi2.arpa
Cheng Song	song@ibm.com
Bob Stewart	rlstewart@eng.xyplex.com
Bill Townsend	townsend@xylogics.com
Steve Waldbusser	sw01+@andrew.cmu.edu

Bert Williams	bert.synernetics@mailgate.synnet.com
Mark Wood	mark1@iw/cs.att.com
Jean Wu	eskovgaa@uvcw.uvic.ca
Allan Young	rcoay@possum.ecg.rmit.oz.au

3.4.12 Transmission Mib (transmib)

Charter

Chair(s):

John Cook, cook@chipcom.com

Mailing Lists:

General Discussion: unknown
To Subscribe: unknown

Description of Working Group:

The objective of the Transmission Architecture Working Group is to drive the development, documentation and testing of MIB objects for the physical and data-link layers of the OSI model. The Working Group attempts to consolidate redundant MIB variables from new specifications into a universal structure.

Goals and Milestones:

Ongoing	Provide a forum for vendors and users of MAC layer communications equipment.
Ongoing	Form sub-Working Groups of experts to define object for the following at the data-link layer: X.25, Ethernet, Token, FDDI and T1.
Done	Form a core group to evaluate the work of the sub-Working Groups.
Ongoing	Act as a liaison between sub-Working Groups and the network management protocol Working Groups, including SNMP, OIM, IEEE 802.1, etc.

CURRENT MEETING REPORT

Reported by John Cook/Vitalink

TRANSMIB Minutes

The SNMP, Bridge MIB, and Transmission MIB Working Groups each met during a single Working Group session. During the short Transmission MIB meeting, John Cook presented his work on the revised T1 mib document.

Attendees

See the SNMP Minutes

3.5 OSI Integration Area

Directors: Ross Callon/DEC and Rob Hagens/University of Wisconsin

OSI Integration Area Report

The OSI General Working Group discussed the draft document that has been produced by the FNC OSI Planning Group (FOPG). This document "OSI Integration Coexistence and Interoperability Issues" is available as an Internet Draft. The purpose of this document is to provide 1) a snapshot of where we are in the process of OSI Integration, 2) a record of issues that have been resolved, and 3) a list of issues that have not yet been resolved and require funding.

The X.400 Working Group discussed a proposal to use the Domain Name System to aid in the operation of RFC 987/RFC 1148 mail gateways. These gateways join together the Internet Standard mail system based upon RFC 822/SMTP and the OSI Message Handling System (X.400). The working group also discussed the format of X.400 addresses that will be used by the Internet Pilot X.400 project.

The NSAP Guidelines Working Group has produced a guidelines document that should be available as an Internet Draft by the next IETF meeting. They have also produced a short RFC which obsoletes RFC 1069 and indicates that the recommended NSAP address structure is that defined by GOSIP, version 2. Finally, they have begun work on a new paper: "A proposal for administration of NSAP allocations".

The X.500 Working Group did not meet. However, they will be meeting later in the Fall at the INTEROP 90 conference.

3.5.1 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:

Dec 1990	Produce a paper describing guidelines for the acquisition and administration of NSAP addresses in the Internet.
Dec 1990	Have the paper published as an RFC.
Dec 1990	Have the paper incorporated, in whole or in part, into the "GOSIP User Guide" and the FNC OSI Planning Group document.

CURRENT MEETING REPORT

Reported by Jim Showalter/DCA

OSINSAP Minutes

The meeting was chaired by Richard Colella (NIST).

Agenda

- Recording of Minutes
- Status of the NSAP RFC
- Status of the NSAP Guidelines Paper
- Proposed NSAP Administration Paper
- Address Transition Issues

Status of the NSAP RFC

Ross Callon (OSI Area Co-director/DEC) gave a brief status of the NSAP RFC. The RFC, which supersedes RFC 1069, is a recommended structure for OSI NSAPs for use in the Internet. At present it is an Internet Draft out for comment. Ross proposed that the group recommend to the IESG that the draft be progressed as an RFC. Although unrelated to the actual status report the door was opened for discussion of whether other addresses could be used and still be GOSIP V.2 compliant. The answer was yes. Essentially, GOSIP does not preclude any NSAP structure. If IS-IS is to be used efficiently, however, the NSAP must carry a 6 octet System ID field and a 1 octet network selector field in the last 7 octets of the DSP.

There was also some discussion on who or what organization has responsibility for assigning addresses. This was prompted by the fact that the NSAP RFC simply points to GOSIP V.2 for NSAP format structure rather than specifying the structure in the RFC. The reason is that the Internet (thus far) is recommending use of the GOSIP format. If the format should change, then the RFC will not have to be republished. In the unlikely event that the GOSIP format should change to such a degree that the Internet experts are uncomfortable with it then the NSAP RFC could be modified to reflect the required format rather than point to GOSIP. Following the discussion a vote was taken on whether or not to recommend to the IESG to advance the NSAP Internet Draft to RFC status. The vote was 17 for and 0 against.

NSAP Guidelines Status

Not much was done since the last meeting. After some discussion it was agreed by consensus that the NSAP Guidelines paper would be updated. All editors' comments would be resolved and the paper would be mailed out for review by the end of August.

A Working Group meeting is tentatively planned to be held at INTEROP in October to review the document prior to the December IETF meeting.

NSAP Administration Proposal

Richard noted that, under current GSA guidelines for administration of GOSIP NSAPs, GSA will entertain proposals from any organization wishing to be assigned AA values under ICD 0005. He recommended that the Working Group develop such a proposal, which would be the administrative counterpart to the NSAP Guidelines paper. The proposal would request one or more AA values from GSA and elaborate on how these would be administered. An organization that is willing to provide the administrative support should be identified to submit the proposal to GSA. NSF was suggested as a possible candidate, and there may be others.

Sue Hares (Merit) volunteered to begin drafting the administration document. If you would like to contribute she can be reached at skh@merit.edu.

4. Address Transition

This subject had arisen on the Working Group mailing list and Richard wanted to ensure that there was no disagreement before updating the Guidelines paper. Subsequent to the explanation of the issue, which is detailed below, there was no significant discussion and no disagreement.

Address transition has to do with the interaction between hierarchical address assignment and the way IS-IS routers handle areas that move from one routing domain to another. For example, assume an area, represented by the area address ABC (i.e., a prefix), moves to another routing domain and retains its area address. If the area address is allocated from the (shorter) prefix of the original routing domain, AB (i.e., hierarchical address assignment), two problems are created. First, in the source routing domain, the ISs must advertise externally to other routing domains that they can reach all addresses that start with AB *except* the addresses that start with ABC (i.e., the recently-moved area). Second, in the destination routing domain, the ISs must advertise externally to other routing domains that they can reach all those addresses that they could reach before, e.g., those that begin with prefix XY, but *also* the area address of the newly-acquired area, ABC.

If there is no address reclamation, over time this will lead to "address entropy", or flat addressing. Any gains in address collapse from originally allocating addresses hierarchically will eventually disappear. It is, therefore, necessary that the area eventually relinquish its old area address to the original routing domain.

Attendees

Nick Alfano	nick@gandalf.ca
Colin Amor	uunet!rti.rti.org!bnrunix!cja
Ross Callon	callon@bigfut.enet.dec.com
C. Allan Cargille	cargille@cs.wisc.edu
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@nhis.navy.mil
Nick Di Iorio	nicola@napoli.att.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Ella Gardner	epg@gateway.mitre.org
Michael Grobe	grobe@kuhub.cc.ukans.edu
Robert Hagens	hagens@cs.wisc.edu
Susan Hares	skh@merit.edu
Ken Jones	uunet!konkord!ksj
Paulina Knibbe	knibbe@cisco.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Chuck Martin	
David Miller	dtm@ulana.mitre.org
Cyndi Mills	cmills@bbn.com
Douglas Montgomery	dougm@osi3.ncsl.nist.gov
Mark Needleman	mhn@stubbs.ucop.edu
Rebecca Nitzan	nitzan@nsipo.nasa.gov
Yakov Rekhter	yakov@ibm.com
Jim Sheridan	jsherida@ibm.com
Jim Showalter	gamma@mintaka.dca.mil
Keith Sklower	sklower@okeeffe.berkeley.edu
Erik Skovgaard	eskovgaa@uvcw.uvic.ca
Zaw-Sing Su	zsu@tsca.istc.sri.com
Justin Walker	justin@apple.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
Jean Wu	eskovgaa@uvcw.uvic.ca

3.5.2 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.

CURRENT MEETING REPORT

Reported by Rob Hagens/University of Wisconsin

OSIGEN Minutes

Agenda

Review and Discuss the Internet Draft "OSI Integration Coexistence and Interoperability Issues".

The meeting was convened by co-Chairs Ross Callon and Robert Hagens. This entire meeting was spent reviewing the draft document titled "OSI Integration Coexistence and Interoperability Issues". This document is available as an Internet Draft.

The meeting was very successful. A summary of the sections of the paper that need revision is presented below.

Volunteers For Text Modifications (that know about it...):

Martin Gross
Sue Hares
Judy Messing
Mark Needleman
Erik Skovgaard
Mark Sleeper

Action Items:

1. Sec 3, pg.4: Modify paragraph explaining regional network intentions. Get survey results regarding regional networks routing ISO 8473 - Rob Hagens
2. Sec 5.2, pg 6, Data Link Layer: Rewrite. Include current status of PPP - Eric Skovgaard
3. Sec 5.2.1, pg 6: RFC needs to be written on IP over HDLC - need volunteer
4. Sec 5.2.7, pg 7: Talk to Dave Crocker about network management tools (trace route as well) - Sue Hares and Ross Callon
5. Sec 5.2.8, pg 7: Modify (including: gtw requirements for dual stacks, congestion bit, and layer 3 requirements) - Ross Callon
6. Sec 5.6, pg 7-8: Modify (including: human-friendly X.400 addresses, and add paragraph on content type) - Rob Hagens and Erik Skovgaard
7. Sec 5.7, pg 11: Review by a VTP expert - need volunteer
8. Sec 5.8, pg 12: Modify (including FTP-¿FTAM appl gtw etc) - Martin Gross
9. Sec 5.?, pg 12: Add section on X Windows over OSI - Mark Needleman

10. Sec 5.?, pg 12: Add additional info on applications over OSI - Judy Messing
11. Sec 6.1, pg 12-13: Contact Steve Kille to review Directory Service issues - Ross Callon
12. Sec 6.2, pg 13: Needs review (John Lynn, Steve Crocker) - Need volunteer
13. Sec 6.3, pg 13: Network Management Modifications - Mark Sleeper, Richard Colella (Sue Hares will review)
14. Sec 7.3.2, pg 17: Modifications to encapsulation - Steve Willis? Keith Sklower?
15. Sec 7.?: Additional section on bridge/packet size etc. Link level issue - need volunteer

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Cathy Aronson	cja@marmot.nersc.gov
Ross Callon	callon@bigfut.enet.dec.com
Isaac Chan	isaac@gui.consumers.bc.ca
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@nhis.navy.mil
Mark Crispin	mrc@cac.washington.edu
Farokh Deboo	fjd@interlink.com
Dale Finkelson	dmf@westie.unl.edu
Ella Gardner	epg@gateway.mitre.org
Michael Grobe	grobe@kuhub.cc.ukans.edu
Martin Gross	gross@polaris.dca.mil
Robert Hagens	hagens@cs.wisc.edu
Susan Hares	skh@merit.edu
Ken Jones	uunet!konkord!ksj
Paulina Knibbe	knibbe@cisco.com
Judy Messing	messing@gateway.mitre.org
David Miller	dtm@ulana.mitre.org
Cyndi Mills	cmills@bbn.com
Douglas Montgomery	dougm@osi3.ncsl.nist.gov
Mark Needleman	mhn@stubbs.ucop.edu
Rebecca Nitzan	nitzan@nsipo.nasa.gov
Mark Seger	seger@mjs1/ogo.dec.com
Steve Senum	sjs@network.com
Keith Sklower	sklower@okeeffe.berkeley.edu
Erik Skovgaard	eskovgaa@uvcw.uvic.ca
Mark Sleeper	mws@sparta.com
Tony Staw	staw@marvin.enet.dec.com
Ed Stern	els@proteon.com

Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Zaw-Sing Su	zsu@tsca.istc.sri.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Justin Walker	justin@apple.com
John Wieronski	john@osc.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@igloo.osc.edu
Jean Wu	eskovgaa@uvcw.uvic.ca

3.5.3 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.

CURRENT MEETING REPORT

Reported by Robert Hagens/University of Wisconsin

OSI X.400 Minutes

Agenda

- Review of the Draft Proposal for the use of the Internet DNS to maintain RFC 987/RFC 1148 Address Mapping Tables.
- Discussion of the structure of O/R Addresses used by the Wisconsin Pilot X.400 project.
- Address mechanisms that allow non-X.400 users (i.e., RFC 822 mail users) to address X.400 users.

The meeting was convened by Chair Robert Hagens. An attendance list will be published with the Proceedings of the IETF. The meeting had several attendees from the NIST/OSI workshop, X.400 SIG.

A proposal has been circulated on several mailing lists; "Draft Proposal for the use of the Internet DNS to maintain RFC 987/RFC 1148 Address Mapping Tables" (by Cole and Hagens) which describes how the DNS could be used to store, retrieve, and maintain the mappings between RFC 822 domain names and X.400 O/R addresses.

Implementations of RFC987 gateways require that a database store address mapping information for X.400 and RFC822. This information must be disseminated to all RFC987 gateways. In the internet community, the DNS has proven to be a practical means for providing a distributed nameservice. Advantages of using a DNS based system over a table based approach for mapping between O/R addresses and domain names are:

1. It avoids fetching and storing of entire mapping tables by every host that wishes to implement RFC987.
2. Modifications to the DNS based mapping information can be made available in a more timely manner than with a table driven approach.
3. Table management not necessarily required for DNS sites.
4. One can determine the mappings in use by a remote gateway by querying the DNS (remote debugging).

The proposal was discussed. A scenario was presented which demonstrated an example lookup:

Given O/R Address:
"/c=us/admd= /prmd=nren/o=uw-madison/ou=cs/ou=dip/s=hagens"

and DNS record

"*.cs.uw-madison.nren. .us.x400" IN TO-822 6 cs.wisc.edu

1. O/R Address is rewritten as a domain name with attribute values used as domain components: dip.cs.uw-madison.nren. .us
2. Lookup domain name within X.400 top-level domain: lookup(dip.cs.uw-madison.nren. .us.x400) returns cs.wisc.edu (count = 6)
3. Since the count indicates that only 6 of the 7 attributes were matched, any unmatched components must be prepended. In this case, prepend "dip".
4. Result: dip.cs.wisc.edu

The proposal received general acceptance. Several changes to the approach have been suggested which differ from that specified in the proposal. These changes are summarized below:

1. DNS representation of O/R address to use O/R attribute values directly, not appendix F notation.
2. The new tree of X.400-¿RFC 822 resource records should be placed within a new top level domain (the name of this top level domain is undecided).
3. Generation of table information from DNS is performed via recursive zone transfers of the x.400 tree (instead of an automated submittal process). This is probably the biggest issue to be resolved. It is vital that the process of extracting the mappings from the DNS be given a thorough analysis so as to insure that it is feasible.
4. Wildcard count field can be changed so that it is statically entered in authoritative input data, instead of computed by authoritative servers.
5. Discard preference field in proposed resource records.

A portion of the X.400 session was spent discussing X.400 naming and in particular the construction of RFC822 addresses to reach users who are really using X.400. This discussion was led by Allan Cargille, University of Wisconsin

I. Naming Choices.
When determining initial X.400 O/R Names, one can either derive the new X.400 names from existing RFC822 addresses, or can start afresh with new names that take advantage of the semantics of the O/R Name structure. In particular, one can select X.400 Organization and Organizational Unit names that are more suitable for database lookup. For example, at the University of Wisconsin-Madison, they have existing addresses of the form user@cs.wisc.edu. Constructing the X.400 O/R Name from the existing RFC822 name could yield something like:

c=us; admd= ; prmd=xnren; o=wisc+edu; ou=cs; s=user while starting afresh could yield names like:

c=us; admd= ; prmd=xnren; o=uw-madison; ou=cs; s=user

So far in the NSF X.400 project they have taken the second approach, that of constructing new O/R Names instead of deriving them from existing domain names.

Group opinion was that sites should have the freedom to select whatever O/R Name they felt would be most helpful, either derived from an existing domain name, or newly selected.

II. Addressing X.400 Users From The RFC822 World.
There are several approaches that can be taken. All have technical advantages and disadvantages – it is not obvious that any choice would be "right" or "wrong". Assume that there are people in the U.S. Internet that are using X.400 as their email service. Users in the RFC822 world need to be able to address these X.400 users. It is assumed that part of the user population at a site may move to X.400, while the remainder of the users continue to use RFC822 mail.

A. Default solution as per RFC987. Mail would be explicitly sent to an RFC987 gateway, with the X.400 address on the left hand side of the "@" and the gateway address on the right hand side. This would look like

"c=us;admd= ;prmd=xnren;o=uw-madison;ou=cs;s=user"@x400.gateway.us.

This scheme does not require any special mapping records in the RFC987 gateway.

B. RFC987 Regular Mapping Rule. This solution has been adopted by some European countries. The RFC822 address for an X.400 user is composed by using concatenating values of the X.400 address. For example, a user with the X.400 address

c=us;admd= ;prmd=xnren;o=uw-madison;ou=cs;s=user

would be addressed as "user@cs.uw-madison.xnren.us" (or something similar). This looks much like an existing Internet address. One would also register MX records to direct mail for xnren.us or organization.xnren.us to an RFC987 gateway.

One complication of this scheme is that it requires a REGULAR rule for constructing the RFC822-style address from the X.400 address. This could be problematic in the U.S. in large. For example, some government sites will be using a value in the ADMD field, whereas other sites will only use a blank in that field.

This scheme requires placing records in the global RFC987 mapping tables but only a few, because general mapping rules are being used.

This scheme creates a new address space inside the U.S. Internet in parallel to existing addresses.

For a user who switched from RFC822 to X.400, mail to the that user's "old" Internet address would still work due to the use of a system alias or .forward file to forward the mail to the new address (and thus to the RFC987 gateway).

C. Mapping to Existing Names. This solution would keep the names used to reach X.400 users consistent with the existing domain names. Each site would register a local MX record in their existing domain name space that points to an RFC987 gateway. This would look very much like just another hostname. Mail to the X.400 users would be sent to this new MX record and be forwarded to a gateway. For example, in the University of Wisconsin Computer Science Department, addresses look like user@cs.wisc.edu. Several people are starting to use X.400, and RFC822 mail was directed to them as:

Last@x400.cs.wisc.edu, or First.Last@x400.cs.wisc.edu

This scheme requires entering a mapping record for every organization into the global RFC987 mapping tables.

Discussion. The Working Group recommended solution C above because it is most consistent with existing domain names, and does not require the creation of any new high-level domains. The Working Group expressed concern at the "x400" string being used as part of a user address (even though this is really just part of an MX record name) because in general we do not want to encourage people to externalize the kind of email end-system inside the email address. Based on this input, the Wisconsin NSF X.400 project has changed to Internet-style addresses of the form:

Last@pilot.cs.wisc.edu, or First.Last@pilot.cs.wisc.edu

Action Items:

Prepare a new version of the DNS proposal. Complete by next IETF meeting.

Attendees

Dave Borman	dab@opus.cray.com
David Brent	brent@staff.ucs.ubc.ca
C. Allan Cargille	cargille@cs.wisc.edu
Isaac Chan	isaac@gui.consumers.bc.ca
Andrew Cherenson	arc@sgi.com
Richard Colella	colella@osi3.ncsl.nist.gov

Curtis Cox	zk0001@nhis.navy.mil
Mark Crispin	mrc@cac.washington.edu
Ella Gardner	epg@gateway.mitre.org
Martin Gross	gross@polaris.dca.mil
Robert Hagens	hagens@cs.wisc.edu
Erik Huizer	huizer@surfnet.nl
Jim Knowles	jknowles@trident.arc.nasa.gov
Neil Koorland	nkoo@cs.ubc.ca
Walter Lazear	lazear@gateway.mitre.org
Judy Messing	messing@gateway.mitre.org
Paul Mockapetris	pvm@isi.edu
Jim Reinstedler	jimr@ub.com
Erik Skovgaard	eskovgaa@uvcw.uvic.ca
Einar Stefferud	EStefferud@ECL
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Peter Vanderbilt	pv@sun.com
Chris Weider	clw@merit.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Jean Wu	eskovgaa@uvcw.uvic.ca

3.5.4 OSI X.500 (osix500)

<u>Charter</u>

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:

This document suggests an initial scope for the IETF OSI Directory Services Working Group (OSI-DS). Brief summary of group: to be supplied after detailed suggestions have been discussed. Timeframe: need to add some timeframes and tighten objectives. Most of this is appropriate for the first meeting.

Goals and Milestones:

TBD	X.500 does not have sufficient functionality for full deployment on the Internet. This group should identify areas where extensions are required.
TBD	The directory can be used to support a wide range of applications. It is necessary to evaluate which are important for the Internet, and what level of priority they should be given within the community. White Pages type of application is likely to be given a high priority.
TBD	A Schema (Naming Architecture) should be defined for the Internet. A requirement for a schema should be defined, and inputs evaluated. Various approaches to specification of Schema from a user and system standpoint should be considered, including update mechanisms.
TBD	There is a requirement for representation of Directory Names, as these will need to be communicated "out of band". An Internet approach to this should be defined.

Ongoing	Liaisons should be established as appropriate. In particular: RARE WG3, to harmonize work with European activities, NIST, to coordinate with the Directory SIG.

3.6 Operation Area

Interim Director: Phill Gross/CNRI

At the Vancouver meeting of the Network Joint Monitoring Working Group, we spent some time discussing the possible organization of the Operations Area. We formed the notion of a "technical board" or "Directorate" to support the Area Director. Some very early thoughts about the charter and mission of such a Directorate are included below.

The IESG Operations Area Directorate

The Operations Area Directorate of the IETF would be a board of advisors comprised of national and international network operators. The Chair of the Operations Area Directorate would serve as the Operations Area Director on the IESG. Some responsibilities of the Operations Area Directorate might include:

- Guidance to other IETF technical development efforts.

  The IETF was formed as a technical development body in support of operational networks. Current IETF activities are still motivated by the goal of improving the operations of real networks. The Operations Area Directorate would help define operational requirments and set priorities for development in other IESG technical areas.

- Development of operations methods, practices, and policies.

  The Operations Area would take an active role in developing guidelines and practices for internet operations, management, and interconnection. This could include attempting to reach consensus upon common joint management policies for common links. It could include specifying common managment tools, common minimum collection metrics, common data storage formats for interchange of information, common display and reporting formats (e.g., performance data or topology maps). These consensus guidelines would be applied in the next two bullets.

- Coordination between operational groups.

  The Internet is now an international communications inter-network. There are many hundreds of administrative domains, thousands of networks, and hundreds of thousands of end systems. It is no longer possible for a single group to act as the main focus for operations of this global enterprise. However, coordination and liaison are possible and crucial. It would be the goal of the Operations Area to encourage coordination and liaison between the various national and international operational groups, and to encourage the usage of commonly agreed methods and practices.

- Coordination between network planners.

  The goals in this bullet are similar to the previous bullet. However, in this case, we distinguish between existing operational networks and those networks in earlier stages of planning. Existing networks may always have aspects of "grandfathered" policies, whereas newly planned networks have an opportunity to follow new practices and guidelines established by consensus.

To help bring a broader operations perspective to the IESG, it may make sense to institutionalize the notion of co-Area Directors on the IESG (perhaps serving as the Chair and Vice-Chair of the Directorate).

As with the newly announced Network Management Directorate (see Chair's message in these Proceedings), we are still at an early stage. However, we hope to be able to announce more concrete results at the Boulder IETF meeting.

3.6.1 Benchmarking Methodology (bmwg)

Charter

Chair(s):

Scott Bradner, sob@harvard.harvard.edu

Mailing Lists:

General Discussion: bmwg@harvisr.harvard.edu

To Subscribe: bmwg-request@harvisr.harvard.edu

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices 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:

Dec 1989	Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.
Feb 1989	The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring

concentrators as well as detail the relative importance of various performance criteria within each class.

TBD Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

BMWG Minutes

The draft version of the terminology memo was reviewed. A number of changes were agreed on and will be made.

Work was started on the methodology memo. A video conference will be set up for sometime in September to continue this work.

A number of attendees expressed a desire that the Working Group quickly start concerning itself with the performance of host implementations of TCP/IP and other protocols.

Attendees

Arthur Berggren	
Chet Birger	cbirger@bbn.com
Scott Bradner	sob@harvard.harvard.edu
Michael Grobe	grobe@kuhub.cc.ukans.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Ruei-Hsin Hsiao	nac::hsiao
Michael Karels	karels@berkeley.edu
David Kaufman	
Joseph Lawrence	jcl@sabre.bellcore.com
John Lekashman	lekash@orville.nas.nasa.gov
Yoni Malachi	malachi@polya.stanford.edu
Gary Malkin	gmalkin@ftp.com
John Mullen	
Bill Nowicki	nowicki@sun.com
K.K. Ramakrishnan	rama%erlang.dec.com@decwrl.dec.com
Ron Roberts	roberts@jessica.stanford.edu
Manuel Rodrigues	
Dean Throop	throop@dg-rtp.dg.com
John Wieronski	john@osc.edu
Walter Wimer	ww0n+@andrew.cmu.edu

3.6.2 DDN Interconnectivity (ddniwg)

Charter

Chair(s):
Kathleen Huber, khuber@bbn.com

Mailing Lists:
General Discussion: unknown
To Subscribe: unknown

Description of Working Group:

No description available

Goals and Milestones:

none specified

CURRENT MEETING REPORT

Reported by Kathy Huber/BBN

DDNIWG Minutes

The first meeting of the DDN Interconnectivity Working Group was at UBC. The group discussed a charter and plan of action. The purpose of the meeting was to determine interest in issues pertaining to:

- Internet Routing
- Connectivity Protocols
- Policy and Procedures
- Monitoring Heterogeneous Systems in the DDN

A mailing list will be set up of attendees and DCA to create a charter and a forum for discussion.

The most critical items discussed were:

- Router requirement concerns with regard to DDN performance
- Inter-agency monitoring and control issues
- Issues pertaining to interconnectivity with the rest of the Internet
- Access control
- The effect of internet growth on DDN users

Attendees

Zorica Avramovic	zorica@sparta.com
Mark Crispin	mrc@cac.washington.edu
Robert Enger	enger@sccgate.scc.com
Kathleen Huber	khuber@bbn.com
Kathy Kerby	kkerby@bbn.com
Walter Lazear	lazear@gateway.mitre.org
Donald Merritt	don@brl.mil
Paul Mockapetris	pvm@isi.edu
Zbigniew Opalka	zopalka@bbn.com
Robert Reschly	reschly@brl.mil
Harvey Shapiro	shapira@wnyosi2.arpa
Thomas Von Deak	tvondeak@nasamail.nasa.gov

3.6.3 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 Phill Gross/CNRI

NJM Minutes

Agenda

- The role of the NJM WG
- Request for statistics from researchers
- Maps
- Operations Area

The role of the NJM WG.

Historically the National Nets have been represented. Matt Mathis wants to hear national net information at NJM meetings. For example, SNMP Session names, and information on T3 migration. However, we agreed this should be more than just an NSFnet group.

Although there is some overlap with TEWG, TEWG concentrates on topology while NJM concentrates on management and monitoring.

Request for statistics by researchers

Phill Gross has received several requests for stats from bona fide researchers. The issue is to get real data for models and simulations, and to do performance analysis. Gross asked how many regional networks kept stats. About 10, the majority of those present, kept stas. Mathis: PSC collects about 20 mbyte per month. Most goes to tape. PSC tends to get their needs met without requiring much analysis.

Guy Almes listed 3 kinds of stats: 1) Reliability, 2) Character of usage (i.e., NNstat, end points pairs, protocol types), and 3) Performance, congestion.

Finkelson: Most analysis programs are local with different formats. He uses NYSER package to collect SNMP data and uses graduate students for data reduction.

That seemed typical. Data tends to be regional specific, and format is special, often reduced. Also, commercial clients of the nets may object to net management data being freely given out to researchers.

Almes: "Character of usage" data is sensitive to some users. Perhaps, IRTF could characterize kinds of data they want and we could work toward provding it. Real data is best. Using real data is good for us too.

Dan Wiverhan (OSU): OSU has tool that monitors multiple ethernets; runs on PC; tells protocols and end-points; Based on KA9Q; Supports SNMP.

Long discussion ensued on various types of date collection. Gross: It either MIB or NNstat-based.

Apparently, some vendors are reluctant because makes their boxes slower. Perhaps a smaller set of common stats would meet better acceptance from vendors. We tried to draw up a list of "low impact, high yield" metrics. Perhaps we should develop a subset of MIB that everyone should archive in standard format, with standardized time granularity.

Maps

Questions: Should USWG catalog all on-line MAP sites? Should format be standardized (e.g., Dated to show currency, All line speeds shown)? Who is working on mapping techniques (MERIT and Bellcore)?

Are maps really used to debug a problem to another site? They are often too far out of date. Many folks said they used maps to debug. Also useful in topology planning. Counter by Mathis: Maps often don't show interesting (i.e., surprising) links.

Ted Brunner (Bellcore) is working on auto map generation. [He gave a demo that afternoon.] Uses MIB. For better display, will probably need more MIB objects than currently available. He has extended MIB to hold extra information. This became an interesting topic. Gross was asked to make sure this was on the agenda for the next meeting.

Other efforts? Gross and Enger mentioned the Contel "net-feeb" program. People are interested. Will contel make available?

Bottom line of this topic - Maps are potentially very useful, but would be much better with better methods (e.g., common formats, auto generation, up-to-date).

Operations Area

Gross: This may not be in NJM Charter, but this group can give important feedback/advice. Should IETF Operations Area be pro-active or re-active? Formal or informal? Liason only? Should IETF propose a set of guidelines for Internet operations? What other groups should be involved? Should we reach to local managers directly, or through Farnet? Should we identify one Area Director or "Board of Directors"?

Almes: More operations folks need to attend the standards sessions. Gross: Differences between ANSI and IETF standards setting process. More user and operations input at IETF.

Bottom line – An active operations area is important. Interaction with protocol development is very important. See the operations area report for more detailed description of the group consensus about the operations area. The IETF Chair gives his thanks to the NJM Working Group for helping to formulate the direction for the IETF Operations Area.

3.6.4 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.

CURRENT MEETING REPORT

Reported by Guy Almes/Rice

TEWG Minutes

The TEWG met for a single session on Wednesday morning, August 1st. Scott Brim, Chair of TEWG, was unable to attend and asked Guy Almes to chair the session in his place.

The session focused on sharing information about three increasingly important areas of Internet topology: Europe, the Pacific, and the new Army Supercomputer Network (ASnet).

Rudiger Volk, of the University of Dortmund and a participant in RIPE, led a presentation and discussion of connectivity both between Europe and North America and within Europe. Olivier Martin of CERN also contributed to the discussion.

Rudiger first focused on trans-Atlantic connectivity. Among the most important links are the following:

- The 64kb/s line from EUnet at CWI in Amsterdam to the UUnet site in Virginia. This serves the EUnet community directly and serves as a backup for other nets, e.g., NORDUnet.
- The 64kb/s line from NORDUnet in Stockholm to the NSFnet site at JvNC. This serves the NORDUnet community directly.
- The T1 line from CERN in Geneva to the NSFnet site at Cornell University. This serves the EASInet community directly.
- The 56kb/s line from INRIA near Nice to Princeton University. This serves users within France directly.
- The 56kb/s satellite line from DFN/WIN in Garching to the ESnet site at Fermilab. This serves the DFN/WIN community within Germany directly.
- The 9.6kb/s line from Karlsruhe to NYSERnet serves another community within Germany.

Further, there is a planned upgrade of the DFN-to-ESnet line to use one of the two 'fat pipes'.

Rudiger and Olivier mentioned two problems that lead to asymmetric and sub-optimal routes to Europe:

- The heavy use of default routes within some parts of Europe often result in asymmetric routes in which packets go from North America to Europe via some explicit route, while return packets use a default path. Increased deployment

of dynamic routing within Europe should improve this situation.
- The use of MX records for some European sites cause very suboptimal routes to be taken in some cases.

RIPE is working with others to help solve these problems.

There was a brief discussion of the situation in Britain. As a general rule, IP traffic from the outside world enters the UK via an application-level gateway in London, and is transmitted via JAnet using the Coloured Book protocols to individual campuses. There are several exceptions to this that we discussed. First, the University of Kent at Canterbury is on EUnet, and thus connects to Europe and thence to North America via CWI in Amsterdam. Also Milo Medin reported that, as part of the 'fat pipe' to London, some British sites will be served by JAnet using 'IP-over-X.25' techniques; this should be an improvement over the current use of the JAnet application-level gateway.

Rudiger closed with some thoughts on a possible outline for an intra-European backbone. He noted that currently, the three most important trans-Atlantic lines are those at:

- CWI in Amsterdam,
- The NORDUnet hub in Stockholm, and
- CERN in Geneva.

There are plans underway to upgrade the bandwidth of lines from Stockholm to Amsterdam and from Amsterdam to Geneva.

Milo Medin, of NASA, reported on a recent meeting of PACCOM, which coordinates the Internet within the Pacific Rim. At the present, there is a 512kb/s terrestrial line from NASA/Ames to Hawaii, which serves Hawaii and the following other sites:

- Japan via four 64kb/s terrestrial circuits. There is some work to combine these to a single 256kb/s circuit.
- Australia via a 56kb/s satellite circuit. There is some work on increasing the bandwidth of this circuit within the year. Unfortunately, it will be quite some time until the circuit can be converted from satellite to terrestrial.
- New Zealand via a 14kb/s analog circuit. There is some work on using better modems, and possibly real-time compression boxes, to increase the effective bandwidth of this line.
- Korea via a 56kb/s circuit.

Among the coming developments are the following:

- The possibility of a 64kb/s line from Japan to Europe. This would complicate routing within the Pacific. The effective use of the current low-speed lines is eased by the ability to use default routing heavily from Pacific Rim countries to Hawaii.

- Discussions of adding Singapore and Taiwan.

In response to a question about networking to sites in Antarctica, Milo expressed regret over current technical problems that prevent the placement of a geostationary satellite there. More seriously, he mentioned that work is being done on networking to Antarctica.

Bob Reschly, of ASnet and BRL, reported on the ongoing deployment of ASnet, which serves the Army supercomputer centers and other Army labs. The initial topology is a mixed T1/56kb/s topology centered at BRL. ASnet is 138.18.

Connectivity to NSFnet is primarily through the ASnet site at the Minnesota Supercomputer Center via MRnet and CICnet. A secondary connection from the ASnet site in Vicksburg to the SURAnet site at Jackson, Mississippi and through SURAnet is planned.

Several ASnet sites are also on MILnet, and a subset of these will be used to route traffic between ASnet and MILnet. An ASnet router at FIX-Ease would improve connectivity both to MILnet, to NSFnet, and to other parts of the Internet.

One interesting technical aspect of ASnet is its planned use of crypto equipment on all serial lines.

ASnet is openly connected to the rest of the Internet, and is to be used only for science/research uses within the Army.

Attendees

Guy Almes	almes@rice.edu
William Anderson	wda@mitre-bedford.org
Jeffrey Burgan	jeff@nsipo.nasa.gov
Eric Carroll	eric@utcs.utoronto.ca
Rob Coltun	rcoltun@trantor.umd.edu
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Dale Finkelson	dmf@westie.unl.edu
Vince Fuller	fuller@jessica.stanford.edu
Hellmut Golde	golde@june.cs.washington.edu
Michael Grobe	grobe@kuhub.cc.ukans.edu
Phill Gross	pgross@nri.reston.va.us
Michael Hrybyk	mwh@educom.edu

Steven Hubert	hubert@cac.washington.edu
Dan Jordt	danj@cac.washington.edu
Kathy Kerby	kkerby@bbn.com
Dan Long	long@bbn.com
Olivier Martin	martin@cearn.cern.ch
Matt Mathis	mathis@pele.psc.edu
Milo Medin	medin@nsipo.nasa.gov
Paul Mockapetris	pvm@isi.edu
Philippe Park	ppark@bbn.com
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Ken Stetten	kstetten@nrao.edu
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Rudiger Volk	rv@informatik.uni-dortmund.de
Tom VonDeak	tvondeak@nasamail.nasa.gov
Carol Ward	cward@spot.colorado.edu
John Wieronski	john@osc.edu
Dan Wintringham	danw@igloo.osc.edu
Robert Woodburn	woody@saic.com

3.7 Routing Area

Director: Robert Hinden/BBN

Area Summary

Interconnectivty Working Group (Guy Almes)

The Border Gateway Protocol (BGP) was made a Proposed Standard. It is described in:

RFC 1163: A Border Gateway Protocol (BGP)
RFC 1164: Application of the Border Gateway Protocol in the Internet

The majority of the meeting was spent discussing a MIB for the management of agents that speak BGP. A second draft was provided in advance by Steve Willis and served as the reference document for the discussion. Among the key points of discussion:

- Actions that should be taken by an agent upon state transitions.
- Variables in the MIB that could be eliminated or streamlined in the interests of simplicity of definition and implementation.
- Contents of tables that describe attribute lists of routes.

Steve Willis took these decisions and will use them to produce a revised document. This MIB will be used provisionally in our early use of BGP Version 2.

The remainder of the meeting was spent discussing early experience implementing and using BGP-2. Dennis Ferguson and Yakov Rekhter were among the early implementors present, and Dennis Ferguson and Jessica Yu were among the early users present.

PDN Routing (Carl-H. Rokitansky)

The Working Group discussed general usage of running IP over X.25 public networks. Topics discuss included address resolution on public X.25 networks and reverse charging mechanisms. It was agreed that the current ARP protocol could be used with a server on a public X.25 network to perform X.25 to IP address mapping. Also discussed was Carl-Herbert Rokitansky's clustering techniques.

Multicast OSPF (Steve Deering)

A rough draft document written by John Moy describing the Multicast Extension to the Open SPF Protocol (OSPF) was circulated and discussed.

The remainder of the meeting was spent discussing OSPF Version 2 Specification which is available as an Internet Draft.

IS-IS Routing (Ross Callon)

The latest version of the Internet Draft for the Integrated IS-IS specification was reviewed. Topics discussed included:

- Authentication mechanisms
- Inter-Domain TAG information encoding
- Amount of Inter-Domain routing information data carried
- Number of addresses per interface

as well as a number of small clarifications in the document.

The Working Group agreed that after the changes were made the spec would be ready to be published as an Internet Draft, and submitted as an RFC.

Routing Working Group Changes

The PDN Routing Working Group was retired. I would like to thank Roki (Carl-H. Rokitansky) for his management of this group and his many trips from Europe to attend the IETF meetings.

A new routing Working Group was formed. It will be called The Routing and Address Resolution over SMDS and X.25 Public Data Networks working group. The chair will be George Clapp of Ameritech. The group will address routing issues and algorithms necessary to run Internet protocols on large public networks.

3.7.1 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.

CURRENT MEETING REPORT

Reported by Ross Callon/DEC

IS-IS Minutes

The IS-IS Working Group met the morning of August 1, 1990, at the IETF meeting in Vancouver, BC. We reviewed the most current Integrated IS-IS specification.

The greatest amount of discussion was on the authentication field. Several problems with the current text in the spec were pointed out. Also, whatever we do will probably conflict with whatever the authentication folks eventually tell us to do. One option was therefore to go back to what was originally in the spec, which is to leave the contents of the authentication field unspecified. However, there is an urgent need for the most basic form of error supression. For example, it is very useful to provide a simple mechanism for preventing mis-configuration of a single link from causing two large routing domains to inadvertantly merge into one domain.

After a great deal of discussion, it was agreed that we would like to do just about the same thing that OSPF already does: provide a simple password mechanism with an escape to allow future identification of other mechanisms. Ross Callon (as editor for the IS-IS specification) was instructed to remove the details of the authentication field from the main body of the spec, specifying the contents of the field as "to be determined", and to provide an annex to the spec specifying how to use the authentication field for carrying a simple password. Also, we agreed to use the same value for the authentication type field as used by OSPF, in the off-chance that future assignments between authentication type fields could be kept in alignment.

It was pointed out that the current definition of the manner of carrying TAG information in the "interdomain routing protocol information field" was difficult to process (in particular, it required that before processing an "IP External Reachability Information" field, the implementation would first have to check what the following field is, and if it is an "Interdomain Routing Protocol Information" field, then process the two fields in parallel). After discussion, an alternate encoding was agreed upon.

There was a discussion of the possibility that the amount of information carried in the Inter-Domain Routing Protocol Information field may be large, and that in some cases the bulk of level 2 routers (those that don't do inter-domain routing directly) would therefore be required to store information that they don't have any use for. This would appear to mean that folks determining how to use this field need to give careful consideration to what inter-domain routing information should be put into this field, and what should be carried by other means. Ross agreed to add a note to the spec describing this issue.

The limit on the maximum number of addresses that can be assigned to a single interface was discussed. There was general agreement that multiple IP addresses per interface was useful in some cases (particularly for transition), but there was no obvious reason to limit a router to two addresses per interface (as in the current spec). It was agreed that a better limit was whatever number of addresses could fit into one occurrence of the "IP Interface Address" field in IS-IS Hello packets, which implies a maximum of 63 IP addresses per interface. It was agreed that this limit was plenty big enough, also that there was no need to pick a smaller limit.

Rob Hagens pointed out that the use of the term "segmentation" in section 3.6 was inconsistent with the terminology used in the OSI spec (the meaning was consistent, just the terminology was different). Ross agreed to fix this.

It was agreed that after these changes were made, the spec was ready to be published as an Internet Draft, and submitted as an RFC. Ross agreed to send the draft spec to the Working Group first in case anyone could find any nits.

A few other minor editorial nits were also transmitted to Ross during side discussions.

Attendees

Karl Auerbach	auerbach@csl.sri.com
Fred Baker	baker@vitalink.com
Art Berggreen	art@opal.acc.com
Chet Birger	cbirger@bbn.com
Ross Callon	callon@bigfut.enet.dec.com
C. Allan Cargille	cargille@cs.wisc.edu
Curtis Cox	zk0001@nhis.navy.mil
Farokh Deboo	fjd@interlink.com
Dino Farinacci	dino@buckeye.esd.3com.com
Jeffrey Fitzgerald	jjf@fibercom.com
Chris Gunner	gunner@osicwg.enet.dec.com
Yong Guo	guo@cs.ubc.ca
Robert Hagens	hagens@cs.wisc.edu
Tony Hain	alh@eagle.es.net
Susan Hares	skh@merit.edu
Peter Harrison	harrison@miden.ucs.unimelb.edu.au
Kathleen Huber	khuber@bbn.com
Paulina Knibbe	knibbe@cisco.com
Holly Knight	holly@apple.com
Alex Koifman	akoifman@bbn.com
Gregory Lauer	glauer@bbn.com

Walter Lazear	lazear@gateway.mitre.org
Solomon Liou	solomon%penril@uunet.uu.net
Yoni Malachi	malachi@polya.stanford.edu
Douglas Montgomery	dougm@osi3.ncsl.nist.gov
Rebecca Nitzan	nitzan@nsipo.nasa.gov
Zbigniew Opalka	zopalka@bbn.com
Brad Parker	brad@cayman.com
Michael Reilly	reilly@nsl.dec.com
Jim Reinstedler	jimr@ub.com
Jim Showalter	gamma@mintaka.dca.mil
Keith Sklower	sklower@okeeffe.berkeley.edu
Frank Solensky	solensky@interlan.interlan.com
John Veizades	veizades@apple.com
Chris Weider	clw@merit.edu
Steve Willis	swillis@wellfleet.com
Walter Wimer	ww0n+@andrew.cmu.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Allan Young	rcoay@possum.ecg.rmit.oz.au

3.7.2 Interconnectivity (iwg)

Charter

Chair(s):
Guy Almes, almes@rice.edu

Mailing Lists:
General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done	Complete development of version 2 of the Border Gateway Protocol (BGP).
Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
May 1990	Develop a MIB for BGP.
Jun 1990	Work with the Security Area to enhance the provision for security in BGP.
Jul 1990	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

CURRENT MEETING REPORT

Reported by Guy Almes/Rice

IWG Minutes

The most important agenda item was the review and approval of a MIB for the management of agents that speak BGP. A second draft was provided in advance by Steve Willis and served as the reference document for the discussion. Among the key points of discussion:

- What action should be taken by an agent upon state transitions (as defined by the finite automaton in the BGP protocol document)? We agreed that SNMP traps would be defined for a subset of these transitions and we agreed on the information to be provided upon each such trap.
- What variables in the MIB could be eliminated or streamlined in the interests of simplicity of definition and implementation? The final MIB will reflect a significant reduction in the total number of variables defined in the second draft.
- There were two tables in the second draft that describe the attribute lists of routes. One table describes all received routes, and the other describes those actually in use. We tightened the description of just when entries in these tables existed and what they would contain.

Steve took these decisions and used them to produce a revised document Tuesday evening. This MIB will be used provisionally in our early use of BGP version 2, and will be the MIB submitted when we propose advancement of BGP to 'Draft Internet Standard' status.

The rest of our time was spent discussing early experience implementing and using BGP-2. Dennis Ferguson and Yakov Rekhter were among the early implementors present, and Dennis Ferguson and Jessica Yu were among the early users present. Among the items discussed were:

- Since the BGP-2 header is an odd number of bytes, implementors should be careful of the C-language size of operator.
- In view of the overhead of processing the message and update headers and the attribute lists of each BGP update message, the inclusion of many routes per update message is an extremely important efficiency concern.
- In BGP-3 we should seriously consider letting the 'next hop' attribute of an update message default to the IP address of the speaker. This would not only simplify the implementation, but would allow an identical update message to be sent to several peers in even more cases than at present.

- Dennis reports a problem with the FSM in the case when two peers try to connect to one another at the same time. This causes a 'BGP Transport connection open' event in the OpenSent state, which causes both ends to disconnect and return to the Idle state, all with no particular reason to think it won't happen again. An improved FSM would fix this.
- Dennis reports the need for a default inter-AS metric attribute. Without one, it is not clear how to compare an advertisement from one peer with an explicit metric with an advertisement from another peer with no metric.
- There was great appreciation for the lack of split horizon in BGP-2. Since each update message contains a complete AS-level path, there is no need for split horizon. Further, by having speaker A advertise to speaker B the nets it gets to via speaker B in a safe way, two significant advantages arise:
  - assembly of update messages is considerably simplified by not having the identity of the peer influence the update message. For example, when A assembles update messages for B and C, it can use the same update for both despite the fact that some of the routes it is advertising may have been derived from B. In many cases, particularly with IBGP, identical update messages can be sent to several peers.
  - the use of BGP-2 for monitoring inter-AS routing is considerably improved, since a speaker learns more fully what routes its peer uses. For example, when A advertises to B even the routes A has derived from B, B learns that A is actually using the advertised routes. This will allow useful sanity checks.
- Similarly, the lack of need for having a Holddown period, as in BGP-1, is taken by the implementors as a major improvement.

In view of the mild nature of the 'problems' encountered by early implementors, continued deployment of BGP-2 throughout the Internet appears likely.

Due to a very strong overlap of IWG and NJM, we decided to cancel the afternoon session which had been planned. We agreed that gaining experience with the implementation and use of BGP-2 during the next several months will be an important task for the IWG. At the Boulder IETF meeting, we will need to review this experience with a view toward moving BGP, with possible revisions, to the Draft Internet Standard level.

Attendees

Guy Almes	`almes@rice.edu`
Jeffrey Burgan	`jeff@nsipo.nasa.gov`
Dino Farinacci	`dino@buckeye.esd.3com.com`

Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Vince Fuller	fuller@jessica.stanford.edu
Robert Hinden	hinden@bbn.com
Dan Jordt	danj@cac.washington.edu
Alex Koifman	akoifman@bbn.com
Solomon Liou	solomon%penril@uunet.uu.net
Dan Long	long@bbn.com
Olivier Martin	martin@cearn.cern.ch
Matt Mathis	mathis@pele.psc.edu
Milo Medin	medin@nsipo.nasa.gov
Philippe Park	ppark@bbn.com
Yakov Rekhter	yakov@ibm.com
Martha Steenstrup	msteenst@bbn.com
Rudiger Volk	rv@informatik.uni-dortmund.de
Steve Willis	swillis@wellfleet.com
Robert Woodburn	woody@saic.com

3.7.3 Multicast Extentions to OSPF (mospf)

Charter

Chair(s):
Steve Deering, deering@pescadero.stanford.edu

Mailing Lists:
General Discussion: mospf@devvax.tn.cornell.edu
To Subscribe: mospf-request@devvax.tn.cornell.edu

Description of Working Group:

This Working Group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done	Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.
Done	Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.
Aug 1990	We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using the existing BSD OSPF code, written by Rob Coltun of the University of Maryland, as an example.
Dec 1990	Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations. The specification should now be ready to submit as an RFC.

CURRENT MEETING REPORT

Reported by Steve Deering/Stanford

MOSPF Minutes

The agenda for this meeting was to discuss the draft OSPF Multicast extensions. Unfortunately, John Moy was unable to attend the meeting and had not yet completed the draft specification. He did send along a partial draft and outline, but since no one had had a chance to read it, we decided to adjourn the meeting and to organize a later meeting or videoconference sometime before the next IETF plenary.

Attendees

Fred Baker	baker@vitalink.com
Rob Coltun	rcoltun@trantor.umd.edu
Steve Deering	deering@pescadero.stanford.edu
Dino Farinacci	dino@buckeye.esd.3com.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Michael Fidler	ts0026@ohstvma.ircc.ohio-state.edu
Stanley Froyd	sfroyd@salt.acc.com
Vince Fuller	fuller@jessica.stanford.edu
Robert Gilligan	gilligan@sun.com
Robert Hinden	hinden@bbn.com
Gregory Lauer	glauer@bbn.com
Luping Liang	liang@cs.ubc.ca
Deepinder Sidhu	sidhu@umbc3.umbc.edu
Frank Solensky	solensky@interlan.interlan.com
Martha Steenstrup	msteenst@bbn.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Steve Willis	swillis@wellfleet.com
Robert Woodburn	woody@saic.com
Yueli Yang	yueli@bnr.ca

3.7.4 Open Systems Routing (orwg)

<u>Charter</u>

Chair(s):

Martha Steenstrup, `msteenst@bbn.com`

Mailing Lists:

General Discussion: `open-rout-interest@bbn.com`
To Subscribe: `open-rout-request@bbn.com`

Description of Working Group:

The Open Systems Routing Working Group is chartered to develop a policy-based AS-AS routing protocol that will accommodate large size and general topology.

Goals and Milestones:

Done	Write an architecture document.
TBD	Draft Protocol Specification of key elements of the protocol.

3.7.5 Private Data Network Routing (pdnrout)

Charter

Chair(s):
CH Rokitansky, `roki@isi.edu`

Mailing Lists:
General Discussion: `pdn-wg@bbn.com`
To Subscribe: `pdn-request@bbn.com`

Description of Working Group:

The DoD INTERNET TCP/IP protocol suite has developed into a de facto industry standard for heterogenous packet switching computer networks. In the US, several hundreds of INTERNET networks are connected together; however the situation is completely different in Europe.

The only network which could be used as a backbone to allow interoperation between the many local area networks in Europe, now subscribing to the DoD INTERNET TCP/IP protocol suite, would be the system of Public Data Networks (PDN). However, so far, no algorithms have been provided to dynamically route INTERNET datagrams through X.25 public data networks. Therefore, the goals of the Public Data Network Routing Working Group are the development, definition and specification of required routing and gateway algorithms for an improved routing of INTERNET datagrams through the system of X.25 Public Data Networks (PDN) to allow worldwide interoperation between TCP/IP networks in various countries. In addition, the application and/or modification of the developed algorithms to interconnect local TCP/IP networks via ISDN (Integrated Services Digital Network) will be considered.

Goals and Milestones:

Done	Application of the INTERNET Cluster Addressing Scheme to Public Data Networks.
Done	Development of hierarchical VAN-gateway algorithms for worldwide INTERNET network reachability information exchange between VAN-gateways.

Done	Assignment of INTERNET/PDN-cluster network numbers to national public data networks. (Mapping between INTERNET network numbers and X.121 Data Network Identification Codes (DNICs)).
Done	Assignment of INTERNET/PDN-cluster addresses to PDN-hosts and VAN-gateways according to the developed hierarchical VAN-gateway algorithms.
Done	Definition of the PDN-cluster addressing scheme as an Internet standard.
Done	Specification of an X.121 Address resolution protocol.
Oct 1989	Specification of an X.25 Call Setup and Charging Determination Protocol.
Oct 1989	Specification of an X.25 Access and Forwarding Control Scheme.
Oct 1989	Specification of routing metrics taking X.25 charges into account.
TBD	Delayed TCP/IP header compression by VAN-gateways and PDN-hosts.
TBD	Provide a testbed for worldwide interoperability between local TCP/IP networks via the system of X.25 public data networks (PDN).
TBD	Implementation of the required algorithms and protocols in a VAN-Box.
TBD	Interoperability between ISO/OSI hosts on TCP/IP networks through PDN.
TBD	Consideration of INTERNET Route Servers.
TBD	Interoperability between local TCP/IP networks via ISDN.
TBD	Development of Internetwork Management Protocols for worldwide cooperation and coordination of network control and network information centers.

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

PDNROUT Minutes

This was the last meeting of the PDN Routing Working Group. Topics discussed included address resolution on public X.25 networks and reverse charging mechanisms. It was agreed that the current ARP protocol could be used with a server on a public X.25 network to perform X.25 to IP address mapping. Also discussed was Carl-Herbert Rokitansky's clustering techniques.

Attendees

Larry Brandt	lbrandt@sparta.com
George Clapp	meritec!clapp@bellcore.bellcore.com
Jeffrey Fitzgerald	jjf@fibercom.com
Yong Guo	guo@cs.ubc.ca
Robert Hinden	hinden@bbn.com
Philippe Park	ppark@bbn.com
Stephanie Price	cmcvax!price@hub.ucsb.edu
Carl-Herbert Rokitansky	roki@isi.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Rudiger Volk	rv@informatik.uni-dortmund.de

3.8 Security Area

Director: Steve Crocker/TIS

Introduction

The work in the security area is carried out both in Working Groups specific to the security area and in cooperation with Working Groups in other areas. Related work also takes place in the Privacy and Security Research Group (PSRG), which is part of the Internet Research Task Force (IRTF). Working Groups active in the security are the Security Policy Working Group (SPWG), Site Security Policy Handbook Working Group (SSPHWG) and the SNMP Authentication Working Group.

Interaction with the PSRG

The PSRG arranged to have its meeting in conjunction with the IETF meeting this time, and that provided an opportunity for the IETF and PSRG members to interact on a variety of topics. PSRG members got a chance to see the workings of the IETF, which led to some discussion of PSRG members also attempting to attend IETF meetings on a regular basis.

A current focus of the PSRG is the development of privacy enhanced mail (PEM). RFCs 1113, 1114 and 1115 specify a PEM protocol, and multiple implementations are underway. TIS demonstrated a version at the Vancouver meeting, and other implementations are known to exist at DEC, MIT, RSADSI and in Germany. The TIS version is undergoing beta testing and will be released to the Internet community later this year.

A handful of technical and policy issues related to PEM are still open. The technical issues include various nits about certificate format, the interaction of mail transport systems with the new focus on trying to send messages which authenticated but not encrypted in a form which can be read with existing mail systems.

An important technical issue is the choice of the message digest function. In the RFCs, MD2 is specified as the primary choice for a message digest function. A similar function is needed in the SNMP Authentication protocol, and there was an exploration of choices of algorithms which might satisfy both requirements. MD4 was chosen as the best choice, because it is much faster than other known algorithms and because it has not been shown to be defective.

The last criterion, viz that it hasn't been broken, is uncomfortably weak. It would be far better to have algorithms that are known to be cryptographically strong. Unfortunately, there is no rigorous way to guarantee that property. A useful alternative is to subject an algorithm to strong review by experts and to expose it to use over a long period of time. MD4 has not been around long, and hence there has only been a modest level of review and moderate use. Nonetheless, it seems far better than any other choice.

In response to concerns voiced over the lack of knowledge of the strength of MD4, a panel of cryptography experts is being formed to review MD4. This panel should be formed during the fall, and results will be announced as they become available. In the meantime MD4 remains the algorithm of choice although all protocols which use it are structured to permit the introduction of alternative algorithms in the future.

The Security Area Working Groups

The Security Policy Working Group met and focused on finding a small set of basic principles around which to build a policy statement. Previously the group had concentrated on describing the full space of issues and formulating a basis for organizing all of the diverse issues. A fuller description of the work on the principles is described in the Working Group's report, and a complete report is expected in October.

The Site Security Policy Handbook Working Group is busy working on a handbook. From the minutes of its last meeting:

> The first pass draft of the Handbook was well received, and the general consensus of attendees is to keep with the direction of the document, with one more pass at the next IETF in Colorado. Submission of the Handbook to the Internet Draft process is projected to be in mid-December, for publication as an RFC FYI at the end of 1990.

The SNMP Authentication Working Group has a new draft of the protocol in three documents in the Internet Drafts directory for review. A fourth document is being prepared which provides an overview of the protocol and its use. During the Vancouver meeting the PSRG and the SNMP Authentication Working Groups met and discussed various aspects of the protocol. Most of the issues involved clarification of the protocol. however, one technical issue that has emerged is how the authentication extensions interact with the addressing structure. In particular some implementations of SNMP impose more meaning on the addressing structure than was originally intended, and there is some concern that such implementations may not extend easily to include authentication. This aspect of the protocol design is under review.

Other topics

A number of other topics have received less attention but are on the agenda for increased activity.

Telnet authentication: The Telnet Working Group desires to add a form of authentication to the Telnet protocol. Ideally, there should be a general form of authentication that applies to multiple protocols rather than having a separate design for each protocol. Both the PSRG and IETF members are thinking about this general issue, but no plan exists yet.

IP Security Option: Work on an IP Security Option existed a few years ago and has been pursued outside the IETF structure. The existing work will be reviewed within the IETF to see if action is needed to standardize the protocol.

Improved login security: DARPA/ISTO has expressed concern over the most obvious weakness in Internet systems, viz the security of the login process. Two aspects are of concern, the quality of the passwords that people use and the fact that passwords are transmitted in the clear. An ad hoc group consisting of Steve Kent, Steve Crocker, and members of the CERT met to discuss what could be done. A Working Group needs to be formed, but it remains dormant for lack of available people.

3.8.1 IP Authentication (ipauth)

Charter

Chair(s):

Jeffrey Schiller, `jis@bitsy.mit.edu`

Mailing Lists:

General Discussion: `awg@bitsy.mit.edu`
To Subscribe: `awg-request@bitsy.mit.edu`

Description of Working Group:

To brainstorm issues related to providing for the security and integrity of information on the Internet, with emphasis on those protocols used to operate and control the network. To propose open standard solutions to problems in network authentication.

Goals and Milestones:

TBD	RFC specifying an authentication format which supports multiple authentication systems.
TBD	Document discussing the cost/benefit tradeoffs of various generic approaches to solving the authentication problem in the Internet context.
TBD	Document to act as a protocol designers guide to authentication.
TBD	RFC proposing A Key Distribution System (emphasis on "A" as opposed to "THE"). MIT's Kerberos seems the most likely candidate here.

3.8.2 Internet Security Policy (spwg)

Charter

Chair(s):

Richard Pethia, rdp@sei.cmu.edu

Mailing Lists:

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

Description of Working Group:

The Security Policy Working Group is chartered to create a proposed Internet Security Policy for review, possible modification, and possible adoption by the Internet Activities Board. The SPWG will focus on both technical and administrative issues related to security, including integrity, authentication and confidentiality controls, and the administration of hosts and networks.

Among the issues to be considered in this Working Group are:

- Responsibilities and obligations of users, 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.
May 1990	Revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.

Jul 1990	Prepare first draft policy recommendation for Working Group review and modification.
Sep 1990	Finalize draft policy and initiate review following standard RFC procedure.

CURRENT MEETING REPORT

Reported by Steve Crocker/TIS, Richard Pethia/CERT, J. Paul Holbrook/CERT

SPWG Minutes

The Security Policy Working Group (SPWG) met in Vancouver. The Chair, Richard Pethia, was unable to attend, and the meeting was co-Chaired by Paul Holbrook and Steve Crocker.

Background

Prior meetings had opened up a range of topics including whether there should be a security policy for the Internet, what aspects of security were important, who should implement the policy, and what means should be used. A three dimensional framework had been proposed to help categorize the issues. The three dimensions are:

Security services, including:

- Protection of information from unauthorized disclosure
- Protection of information from unauthorized modification
- Protection from denial of service
- Protection from unauthorized use of facilities

Who is affected

- Users
- Host operators
- Local network operators
- Regional and Backbone network operators
- Host operating system vendors
- Network component suppliers, e.g., router vendors

Means to implement

- Administrative
- Technical
- Legal and Legislative

The Vancouveri Meeting

At the Vancouver meeting, we shifted focus and attempted to find a consensus on what the central elements of an Internet policy might be.

The group engaged in an experiment in which each participant attempted to write a set of principles. This exercise worked very well, and the responses from the group showed a surprising amount of agreement. Joel Jacobs from Mitre took the task of trying to synthesize the writings of the group into a single strawman security policy. A summary (and interpretation) of some of the thoughts of the group is included at the end of these minutes.

A fuller summary of the exercise conducted at the Vancouver meeting will be coming out in October. Some points emerged fairly clearly. There is a common understanding that sites are fundamentally reponsible for their own security and that in a community as large as the Internet there are some individuals who will attempt to violate the security of systems. Against this backdrop, two ideas emerged fairly clearly as principles to build into the policy.

1. Users have a positive obligation to respect the security of the systems on the Internet. This includes not attempting to penetrate systems they don't have access to and not exceeding the authorized use of the systems they have access to. As simple as this statement seems to be, it establishes the idea that security in the Internet is not a game. Without a clear statement along these lines, it might be considered fair game to try to break into systems just to see if it can be done.
2. Sites and network operators should cooperate with each other on security matters. Again, this statement seems simple on its face, but it establishes the idea that sites, local nets, etc., have an obligation to assist each other instead of leaving each site strictly to its own defense.

These ideas and others will be elaborated upon in the next few months.

Selected Observations

What follows are some of the themes the group seems to agree upon coupled with explanatory paragraphs in which I (Paul Holbrook) try to interpret the thinking of the group.

A caveat: the information in this document has been filtered several times. Steve Crocker provided the original bullets, and thus provided his own view of what the group said. The paragraph after each bullet is my interpretation of what the group was thinking about. In particular, where the explanation says people 'should' do something, that does not mean that everyone agreed to propose this, just that this is one interpretation of where the group was going. The result is that the people who were at the meeting may not agree with what follows.

Internet, regionals/backbones, sites, hosts – all should have security policies.

Security policies and procedures are needed at all levels of the Internet. The policies will be different for different groups, and the general level of security expected may be different. For example, the policy may encourage regional networks to protect the network infrastructure such as the routers and other network equipment, but may put the burden of privacy on hosts. Thus, a regional would make it's best effort to protect the network, but would not provide a guarantee of privacy for the hosts that use it.

Emphasis on user responsibility, identification, and accountability.

The policy should state clearly that users are responsible for their own actions regardless of the level of security a site maintains. By analogy, even if you leave your front door unlocked, that doesn't give someone else permission to enter your house.

Sites should also have policies that support identifying and (if necessary) accounting for individual users. If your site is used to break into another site, that other site may ask for your help in tracking down the problem. It should be possible for you to figure out what user's account at your site was used. This requires that all users be individually identified, and that enough accounting records be kept to identify when users were on systems. (On Unix systems, the normal login accounting may well be sufficient for what we're after here.)

This last requirement is likely to be controversial. There are sites that keep guest or group accounts for their own convenience, terminal servers that allow access out to the Internet without logging into a local system, and so forth. There was some irony in this proposal, since we all enjoyed this kind of open access out to the Internet at UBC, yet this was the very kind of access we were proposing limiting.

Emphasis on mutual assistance

- Preference for investigation
- Concern for privacy

Where possible, sites should assist each other in investigating security incidents. Sites should provide contact points to help facilitate communication about security problems.

When a security incident occurs, a site has two main choices:

- Try to watch or trace the intruder(s) in an effort to see how widespread the problem is and hopefully identify who is responsible;
- Identify the vulnerabilities or lapses that led to the incident, clean up the systems and lock the intruder(s) out.

Some people leaned towards encouraging sites to investigate problems. In many cases, locking an intruder out will force them to find another site to use, but will not stop them from breaking into systems. The decision about what to do about an intrusion will always be up to the site, but the community standard should be to try to solve the problem. This does not necessarily advocate prosecution or law enforcement involvement. Once an intruder is identified, there are many possible courses of action.

Encouragement to use good security controls

Policies and procedures are not a substitute for putting good security controls in place and making sure systems are securely configured. The policy should encourage sites to put useful security controls in place.

The Need for Unforgeable User Identification

Vint Cerf/CNRI

FIRST DRAFT

Summary

This brief memorandum motivates the need for Internet mechanisms and facilities for authenticating user identification and for assuring that such identification cannot be forged.

Introduction

The Internet has reached a point in its evolution where some of the services accessible require compensation from the using parties (or an entity which accepts responsibility for paying for services rendered).

At the application level, such compensation is required for use of information services such as bibliographic databases (National Library of Medicine MEDLARS; Research Libraries Information Network, etc.)

Commercial electronic messaging providers (e.g. MCI Mail, Compuserve, ATT Mail, Sprint Mail, BT Dialcom, QUIK-COMM, etc.) normally charge for their services. Some, such as Compuserve and MCI Mail provide access to commercial information services (e.g., Dow Jones News & Retrieval). Under the present terms and conditions, commercial email services do not charge Internet users for delivering email sent from Internet sites to commercial email boxes. Even if this provision remains in place, there are other services such as fax and hardcopy delivery, bulletin boards and information services which, at present, are not accessible to Internet users because there is no secure way to identify a billable account to which to charge these special services.

Passwords carried in plaintext form across the Internet, whether in a Telnet session or via email, are not sufficiently protected to make the risk of compromise acceptable. Moreover, there is no currently standardized means of authenticating whether the use of a particular billable account is legitimate (once a password is compromised, it can be used at will, for simple, password-based account identification methods.)

Example Requirements

At least two applications need reliable, secure account authentication capability:

- Remote login
- Email store and forward services

In the first instance, it is required that the user/account identification provided to the server be protected from capture and re-use by hostile third parties and that the serving site can verify that the identification has not been forged.

In the second case, it is required that at an email relay, an arriving message to be passed into the next email system can be reliably and authentically associated with an account in the next email system, if necessary, for purposes of accounting and validation that the message originator is authorized to use the services requested.

For example, it should be possible for an Internet user to send email to fax recipients by way of ATT Mail and for ATT Mail to correctly account and bill for this usage. This means that the originator must supply information associated with a message which identifies account information needed to complete processing of the message at the Internet/commercial email interface. The provision of this account identifying information needs protection from compromise and validation that its use is legitimate.

Questions

1. Can the same techniques work for remote login and store-and-forward services?
2. Even if a "password" can be encrypted for confidentiality and signed for authenticity, how can the recipient be sure that the encrypted and signed object has not been hijacked by an abusing third party? (i.e. "stealing and reuse")
3. Given that there must be some kind of authenticated exchange between user and server just to set up an account, can we take advantage of this to carry out any additional exchanges needed to support the confidentiality and authenticity required for these account validation applications?

Scope of the Internet Security Policy

J. Paul Holbrook/CERT/SEI/CMU:

This proposal deals with two areas that the Internet Security Policy is concerned with: the scope of the Internet Policy, and lines of authority or responsibility at a site. These are separate issues, so I'll treat them that way.

Scope of the Policy

The Internet Security Policy should not mandate security policies for sites beyond what is necessary for maintaining the security of the Internet. The policy should not mandate the form of a site's internal response to security problems. However, it should require that a site have policies in place which meet a minimum set of requirements to allow effective prevention of and response to Internet security problems. Helping a site develop a more complete set of security policies and procedures is the goal of the the Site Security Policy Handbook.

The goal of the policy is to ensure that each site responsibly protects and audits access to the Internet, and maintains a point of contact so that each site can get information about security problems and also assist others in dealing with security problems that involve their site.

The policy covers all "network-capable" devices that may affect the Internet. Thus, in addition to hosts, terminal servers, routers, and other network management devices are covered. Other machines that may indirectly allow unaudited access to the Internet are also covered. For example, if a host that has access to the Internet also trusts other hosts on a site's local network, the policy covers those other machines as well. As an example, if an Internet host trusts a local PC via some mechanism such as rlogin or special trusted accounts, a user might be able to use the PC to gain access to the Internet without proper auditing. In this case, the PC is covered under the policy. (If the Internet host does not trust the PC, the PC does not come under the policy.)

Site Authority

In this proposal, I use the term 'site' to mean every resource-owning organization, including regional networks and other entities. I've used the terms 'MUST' and 'SHOULD' in capitals to help point out suggested policy directions.

> [Comments in brackets are notes to help explain the reasoning behind some of the statements. These comments would not appear as part of a policy, though they might appear as a commentary that goes along with the policy.]

Site Security Contact

Every site MUST have a site security contact. This may or may not be the same as the normal site contact or network manager. A site security contact can be an individual or an organization. The site security contact SHOULD be familiar with the technology and security of all systems at that site. If that is not possible, the security contact MUST be able to get in touch with the people that have this knowledge 24 hours a day.

> [At the CERT we've been in touch with sites only to find out that they have no idea who is responsible for security or how to get in touch with them.]
>
> [A point of terminology: in his 'responsibility' writeup, James VanBokkelen refers to 'network managers' and 'host managers'. The site security contact is a peer to the network manager; it might even be the same person. Others in the Internet community have used the term 'site contact', which I've used because it helps to emphasize that a site security contact may have to deal with both network and host issues. Certainly a regional network or other network provider can (and should) have a 'site security contact.' However, the terminology is certainly open to change.]

Security Contact Availability

The site security contact MUST provide other designated organizations in the Internet with a 24 hour point of contact. At a minimum, this should be a phone number which is answered during 'business hours' 5 days a week, and equipped with an answering machine that is checked at least once every day (including weekends) to cover off hours. Sites SHOULD consider providing 'real time' response: e.g., home phone numbers, pager numbers, or other means of contacting people. However, being able to get directly in touch with the security contact at any time is not required.

> [This is a compromise statement; it's hard to require a site to provide around-the-clock response without proof that it would be worth the cost. At the CERT we've found almost all problems can be dealt with by having a contact who is available during business hours. However, large sites or sites that care about the availability and security of their systems will probably want to provide 24 hour access to their security contact.]

Sites MUST ensure that some backup security contact can be reached if the primary security contact is unavailable. This can take the form of a secondary contact person or organization. If outside organizations must use some different procedure to get to the backup security contact, sites MUST ensure that these procedures are communicated to the outside organizations.

The 'designated' or 'outside' organizations have this contact information might be a local Network Control Center or Network Information Center, or might be security response centers such as CERT. Since security organizations might need access to this information anytime, organizations that keep this information MUST make it available 24 hours a day.

> [The User Connectivity Problem (UCP) Working Group is working on the problem of how to get site contact information propagated around so that network problems can be dealt with. We should consider using whatever means they come up with for distributing this kind of information. In any case, the specifics of how this works are an operational matter that doesn't belong in a policy.]

Security Policy Issues

Although the initial response to a security incident is often a technical one, policy issues also need to be dealt with. Should an intruder be shut out or watched? Should law enforcement be involved? Should a site disconnect itself from the network to avoid a worm or intruders? These decisions are not strictly technical; they may affect many people. Sites MUST ensure that people with the authority to decide these kinds of issues are available in the event of a serious security problem.

If the site security contact does not have the authority to make these kinds of decisions, sites are encouraged to have a 24 hour administrative contact. (This administrative contact does not need to be visible to people outside the site.) Sites SHOULD also have policies that state who has the authority to make decisions and take actions in response to security problems, and under what circumstances administrators or decision makers should be brought in on an active security incident. The goal should be that a site security contact can quickly (i.e., in a few hours) take action to deal with a security problem, if necessary getting in touch with someone who can authorize their actions.

At some sites, policy makers could give advance authorization to the site security contact and other system managers. For example, the site may give their technical people the authority and license to make their best efforts to deal with security problems. In this case, the policy also protects the technical people from 'retribution' from policy makers after the fact.

> [The motivation here is that policy makers should be involved early on if a serious security incident is underway. Policy makers may have little to do with the day-to-day operation of systems, but they will be concerned if a serious security incident has serious impact on a site and it's operation.

Among other things, if decision makers are not involved and understand the nature of security problems, they might impose policies after the fact to 'deal with the security problem.' For example, the CERT has heard of sites where the local policy maker's response to a security incident was to advocate permanently disconnecting from the Internet.

However, since this issue is mostly a matter of site internal policies, the Internet Security Policy should not mandate an administrative contact. The Site Security Policy Handbook will help flesh out this area by going into detail about how site policy makers should be involved in setting security policy and procedures.]

Attendees

Alison Brown	alison@maverick@osc.edu
Steve Crocker	crocker@tis.com
Terry Gray	gray@cac.washingtom.edu
J. Paul Holbrook	ph@sei.cmu.edu
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Joel Jacobs	jdj@mitre.org
David Jordan	...jordan@emulex.com
Tim Seaver	tas@mcnc.org
Mark Stein	marks@eng.sun.com
Dale Walters	
John Wieronski	john@osc.edu
C. Philip Wood	cpw@lanl.gov

3.8.3 SNMP Authentication (snmpauth)

Charter

Chair(s):

Jeffrey Schiller, jis@bitsy.mit.edu

Mailing Lists:

General Discussion: awg@bitsy.mit.edu
To Subscribe: awg-request@bitsy.mit.edu

Description of Working Group:

To define a standard mechanism for authentication within the SNMP.

Goals and Milestones:

May 1990	Write an RFC specifying procedures and formats for providing standardized authentication within the SNMP.

3.8.4 Site Security Policy Handbook (ssphwg)

<u>Charter</u>

Chair(s):

J. Paul Holbrook, ph@sei.cmu.edu
Joyce K. Reynolds, jkrey@venera.isi.edu

Mailing Lists:

General Discussion: ssphwg@cert.sei.cmu.edu
To Subscribe: ssphwg-request@cert.sei.cmu.edu

Description of Working Group:

The Site Security Policy Handbook Working Group is chartered to create a handbook that will help sites develop their own site-specific policies and procedures to deal with computer security problems and their prevention.

Among the issues to be considered in this group are:

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:

May 1990	Review, amend, and approve the charter as necessary. Examine the partcular customer needs for a handbook and define the scope. Continue wok on an outline for the handbook. Set up a SSPHWG "editorial board"for future writing assignments for the first draft of document.
Jun 1990	Finalize outline and organization of handbook. Partition out pieces to interested parties and SSPHWG editorial board members.
Aug 1990	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 forpublication.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI and J. Paul Holbrook/CERT

SSPHWG Minutes

The first pass draft of the Handbook was well received, and the general consensus of attendees is to keep with the direction of the document with one more pass at the next IETF in Colorado. Submission of the Handbook to the Internet Draft process is projected to be in mid-December, for publication as an RFC FYI at the end of 1990.

Attendees

L Allyson Brown	allyson@umd5.umd.edu
Richard Colella	colella@osi3.ncsl.nist.gov
Mark Crispin	mrc@cac.washington.edu
Carol Farnham	carolf@mcescher.unl.edu
J. Paul Holbrook	ph@sei.cmu.edu
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Joel Jacobs	jdj@mitre.org
Dale Johnson	dsj@merit.edu
Gary Malkin	gmalkin@ftp.com
Berlin Moore	prepnet@andrew.cmu.edu
Gerard Newman	gkn@sds.sdsc.edu
Fred Ostapik	ostapik@nisc.sri.com
Tim Seaver	tas@mcnc.org
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
John Wieronski	john@osc.edu
C. Philip Wood	cpw@lanl.gov

Chapter 4

Network Status Briefings

4.1 Mailbridge Report

Reported by Kathleen Huber/BBN and Zbigniew Opalka/BBN

MAILBRIDGES

Kathleen Huber, Zbigniew Opalka

July, 31 1990

BBN Communications
A Division of Bolt, Beranek and Newman, Inc.

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH SUMMARY

- 1432 Networks Advertised
- 2995 Networks Registered

NUMBER OF NETWORKS
LINEAR
DECEMBER 1983-JULY 1991

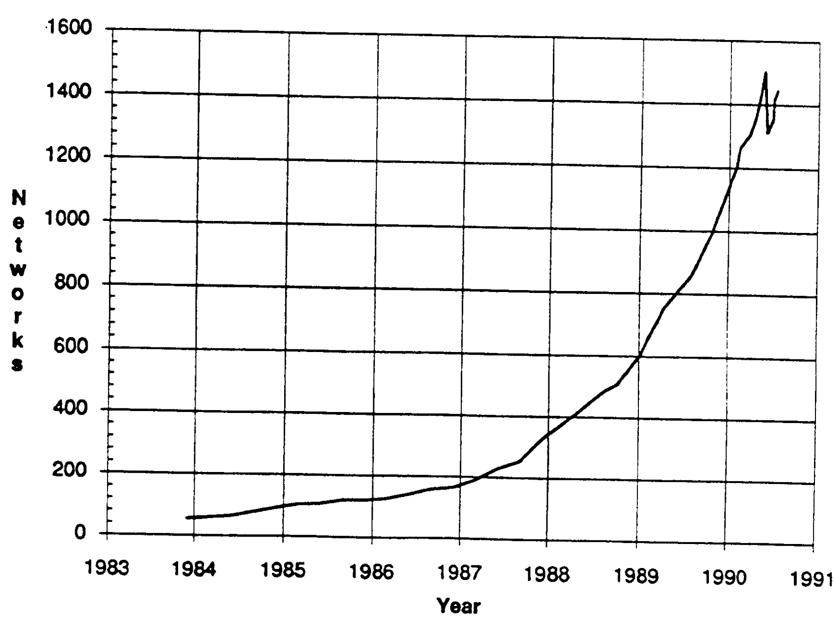

NUMBER OF NETWORKS
LOGARITHMIC
DECEMBER 1983-JULY 1991

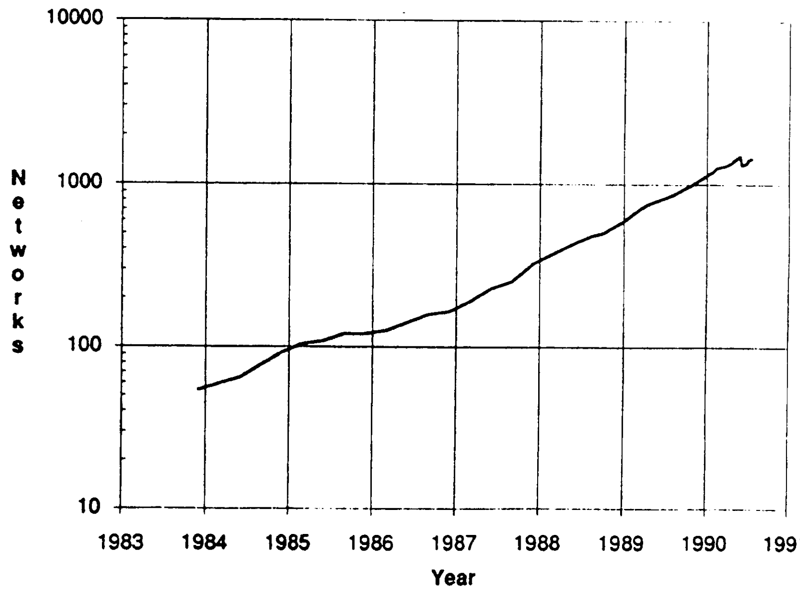

DDN MAILBRIDGES

CURRENT STATUS

- Six DDN Butterfly Mailbridges Operational
  - Two Mailbridges interface the NSFNET
- Seventh Mailbridge
  - August 10
  - Randolph Air Force Base, Texas
- Increase EGP Poll Interval
  - Negotiable up to 60 minutes
  - Testing with Air Force concentrators
- Update Configuration of FIX-EAST

FIX-EAST
(PRESENT)

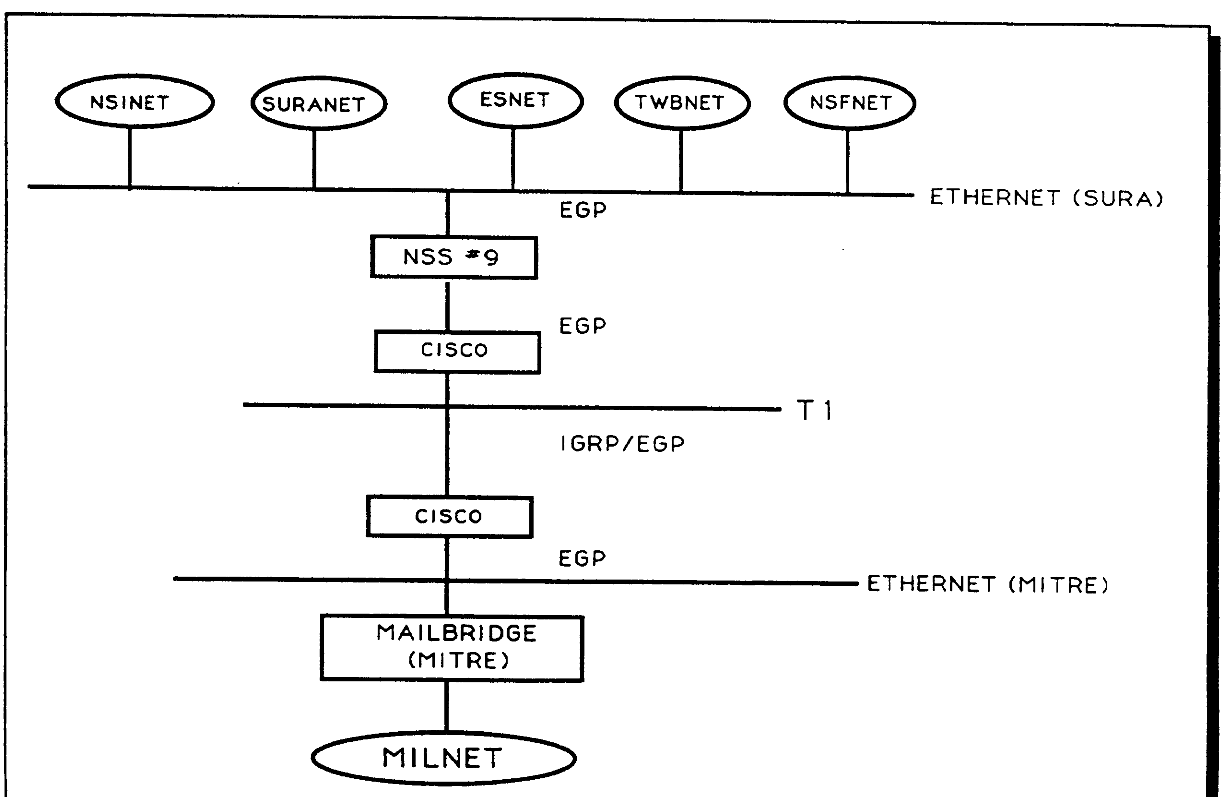

FIX-EAST
(PLANNED)

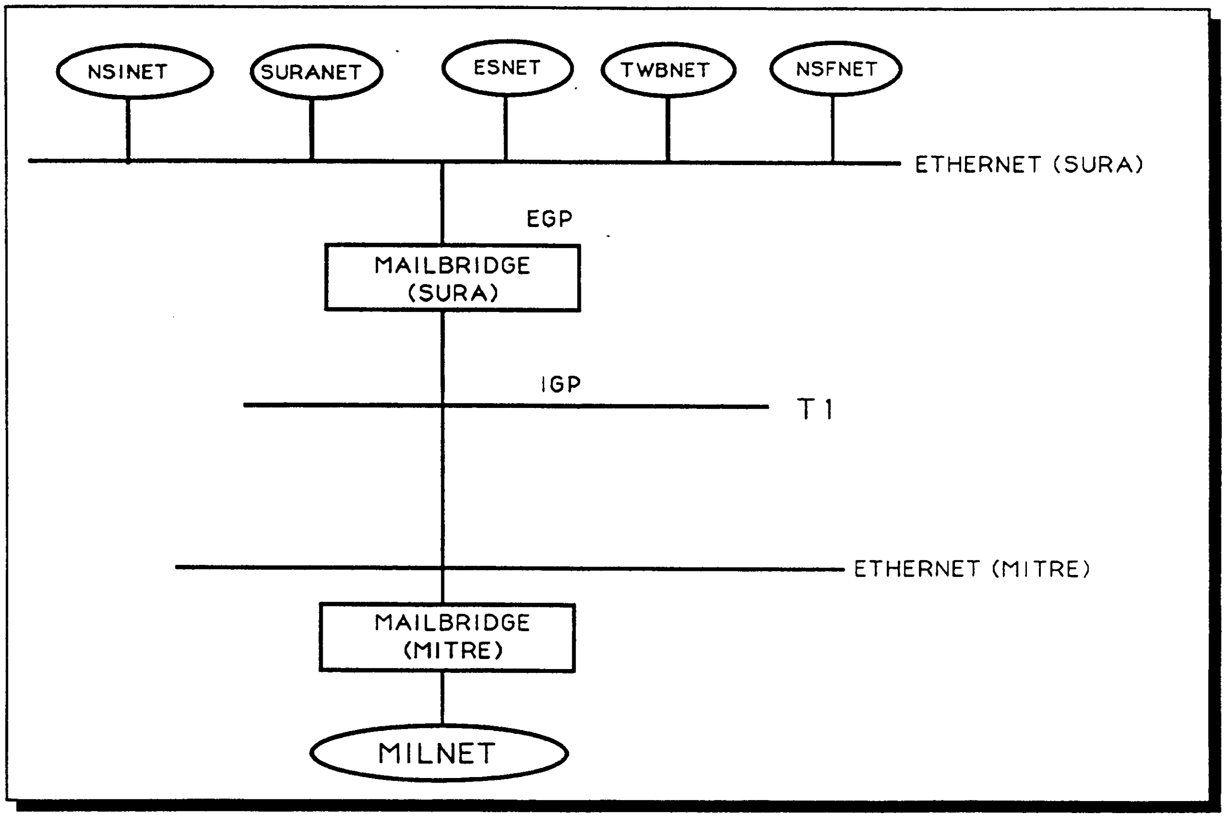

CURRENT STATUS

- ARPANET Termination - June 1
  - 30% decrease in Total Traffic
    - fewer routing loops
    - less routing updates
- Heaviest Traffic Originators
  - SIMTEL20 - 20% To-BMILAMES traffic
    - 12% To-BMILMITRE traffic
  - Aberdeen (UMd) - 24% To-BMILAMES traffic
    - 11% To-BMILMITRE traffic

MILNET
GATEWAY AND INTERNET TRAFFIC
HOMING ANNOUNCEMENT

- Balance Mailbridge User Demand For:
  - Internet traffic service
  - EGP service

- Implement Changes Only Between:
  - 2100 hours Friday, August 10 and,
  - 0600 hours Monday, August 13

ASSIGNMENTS

- Assign EGP Servers
  - BMILAMES
  - BMILMTR
  - BMILBBN
  - BMILISI
  - BMILRAN
- Assign AF Concentrators
  - BMILDCEC
  - BMILLBL
- Assign Internet Traffic Servers
  - ARPANET Termination June 1
  - NSFNET interfaces
    - BMILAMES (FIX-WEST) 192.52.195
    - BMILMTR (FIX-EAST) 192.52.194

EGP NEIGHBOR COMPARISON

	DIRECT NEIGHBORS April	DIRECT NEIGHBORS June/July
BMILAMES	90	68
BMILBBN	141	111
BMILDCEC	112	99
BMILISI	69	62
BMILLBL	43	63
BMILMTR	105	76

TRAFFIC SUMMARY COMPARISON

	Avg. Pkts/Day Forwarded		Avg. Bytes/Pkts.		Avg. Pkts Dropped	
	Jan-April	May-July	Jan-April	May-July	Jan-April	May-July
BMILAMES	4,460,790	3,983,027	144	161	2.1%	0.7%
BMILBBN	2,539,730	1,251,380	131	212	3.2%	5.5%
BMILDCEC	2,648,190	1,251,969	138	204	2.7%	4.4%
BMILISI	1,552,510	523,932	227	253	0.1%	0.2%
BMILLBL	224,139	430,421	397	277	0.0%	0.1%
BMILMTR	3,581,250	2,982,371	149	149	0.9%	0.8%

TRAFFIC TO BMILAMES

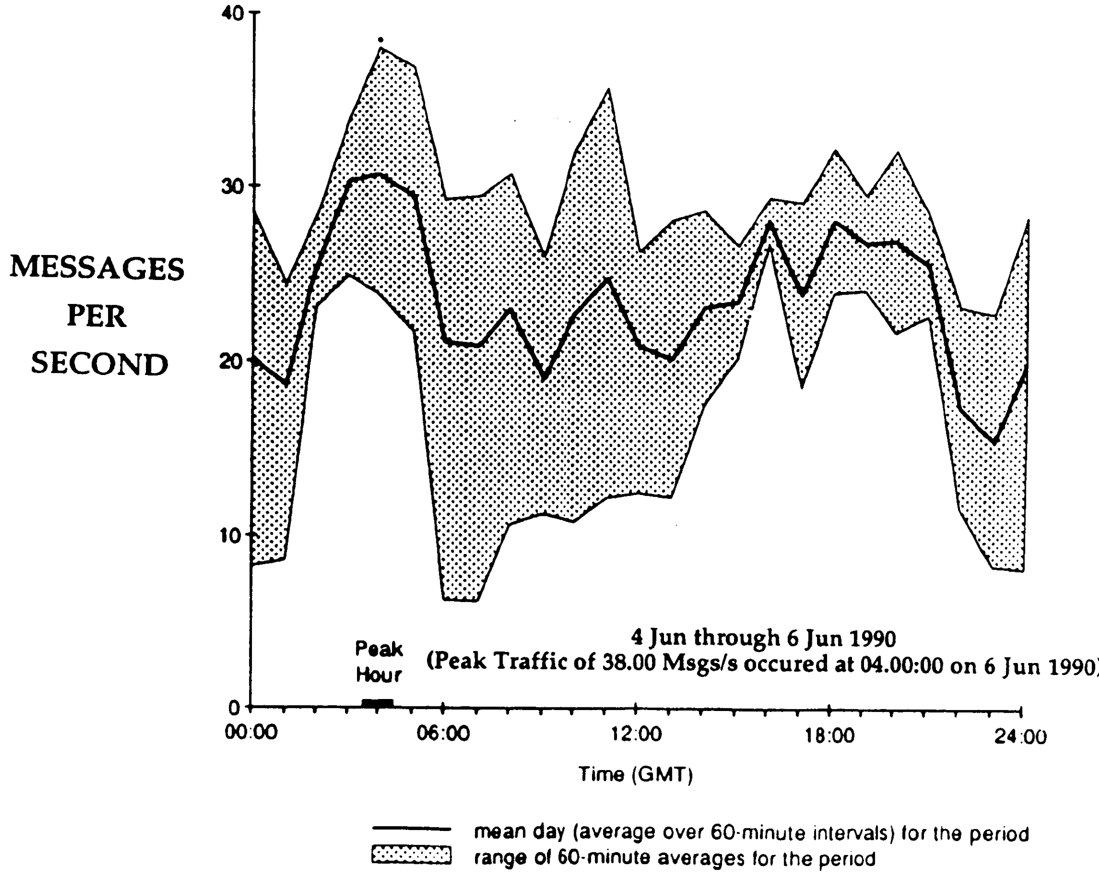

TRAFFIC FROM BMILAMES

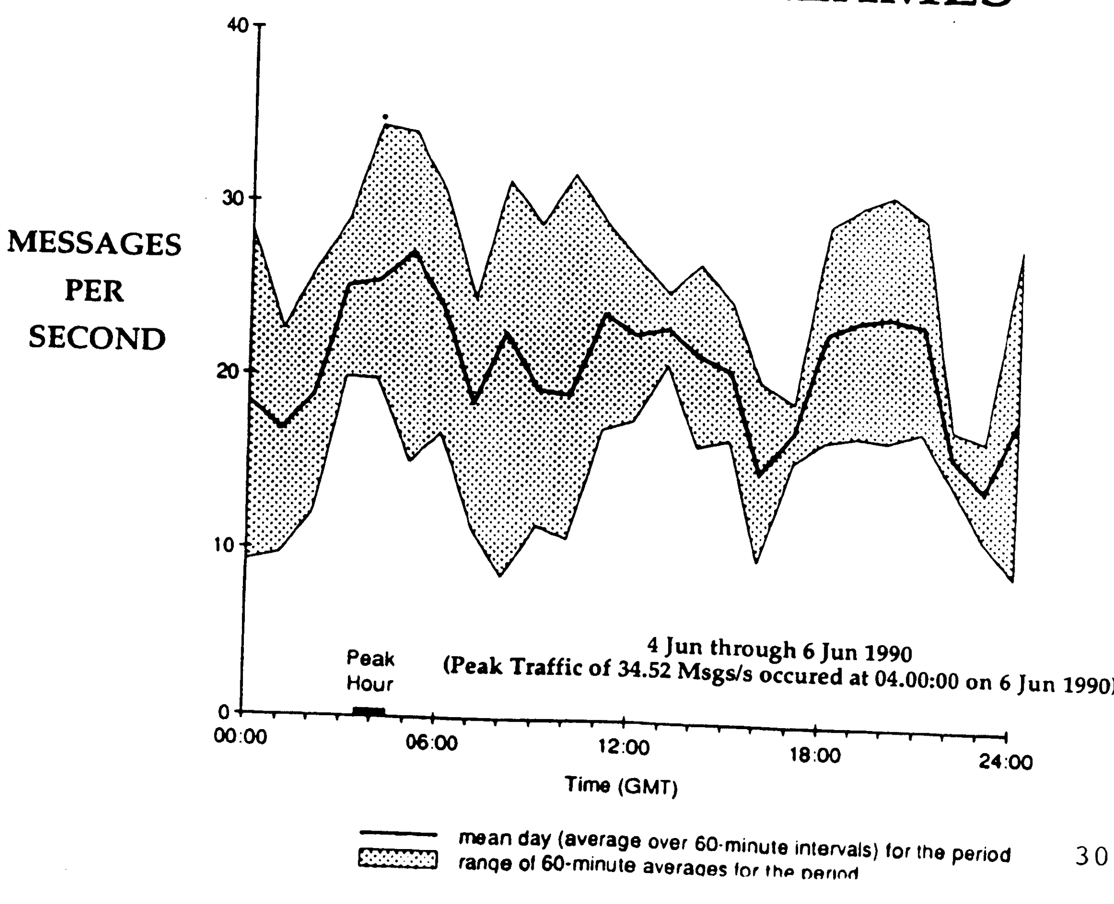

THROUGHPUT WEEKLY AVERAGES

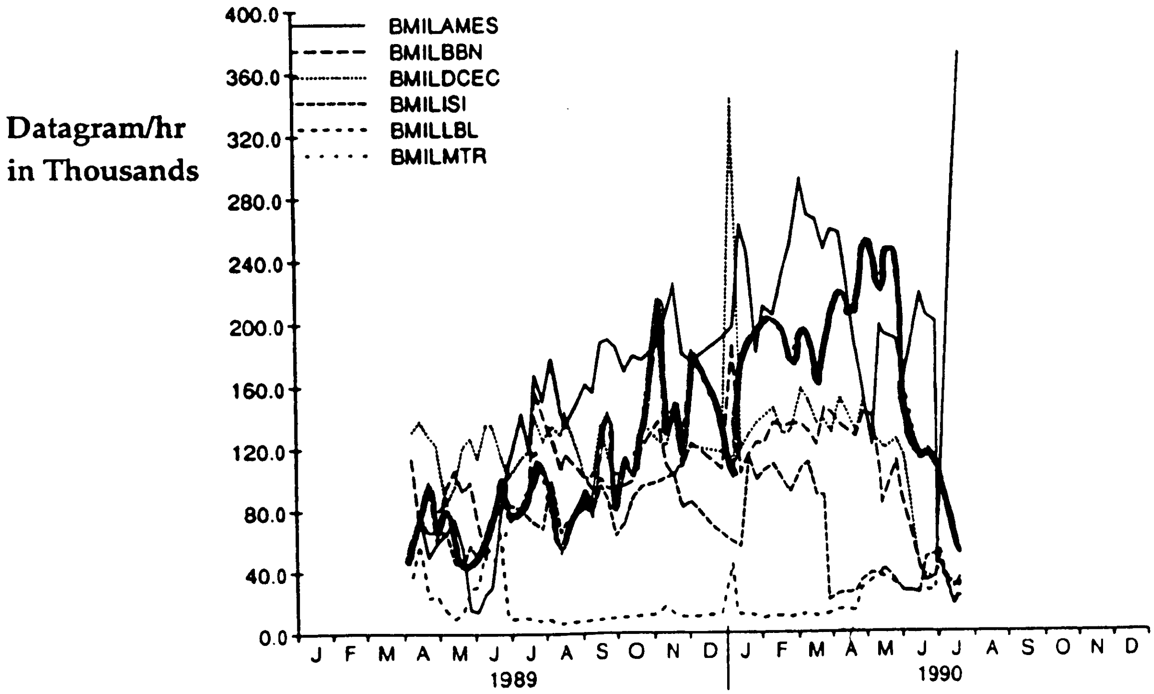

TO-GATEWAY TRAFFIC WEEKLY AVERAGES

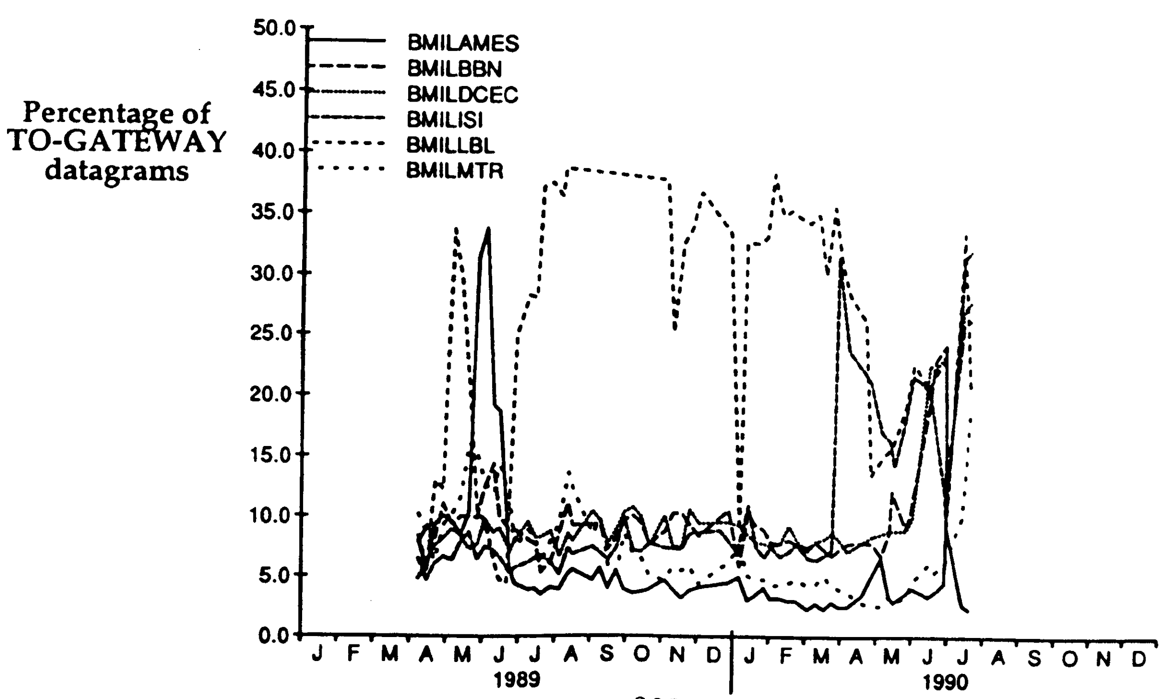

DAILY THROUGHPUT FOR BMILAMES

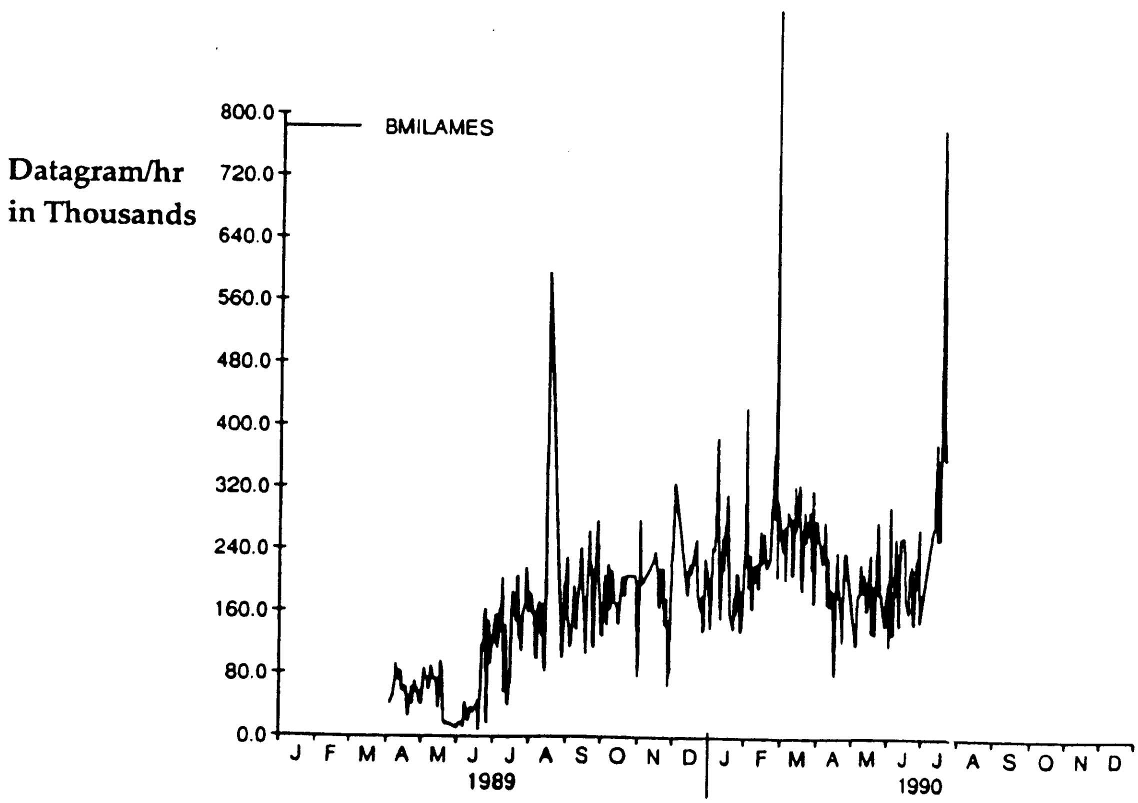

DAILY THROUGHPUT FOR BMILMTR

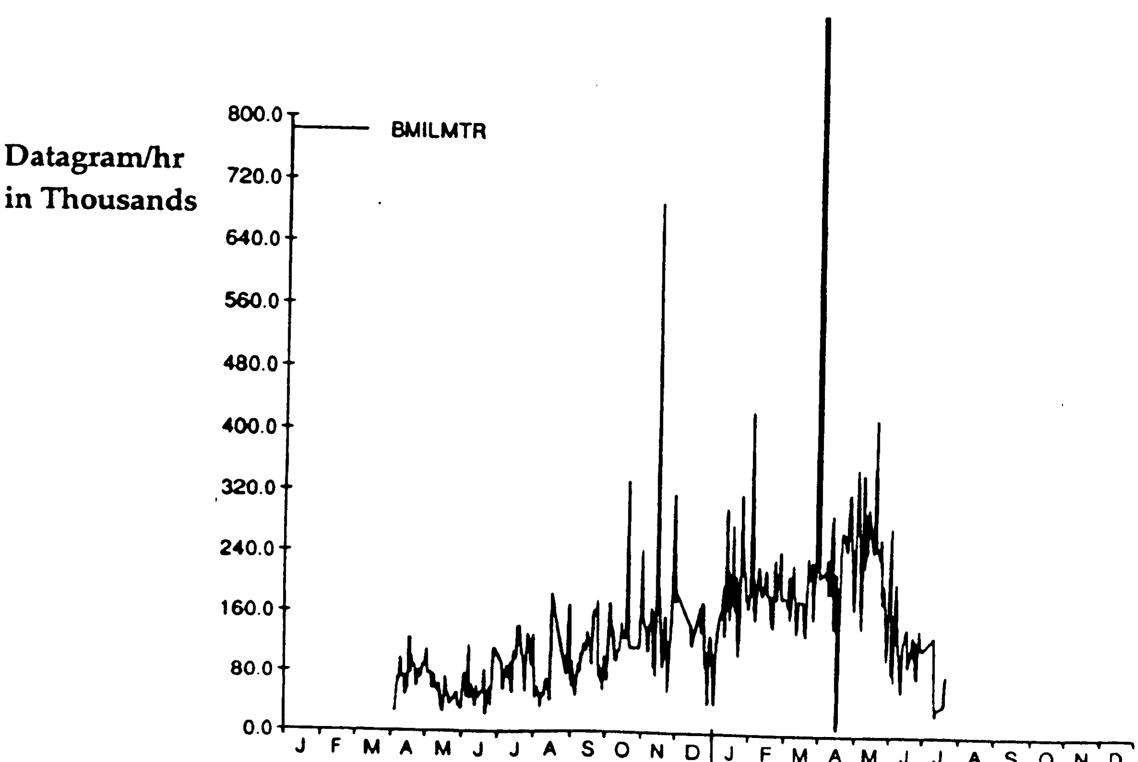

DAILY THROUGHPUT FOR BMILBBN

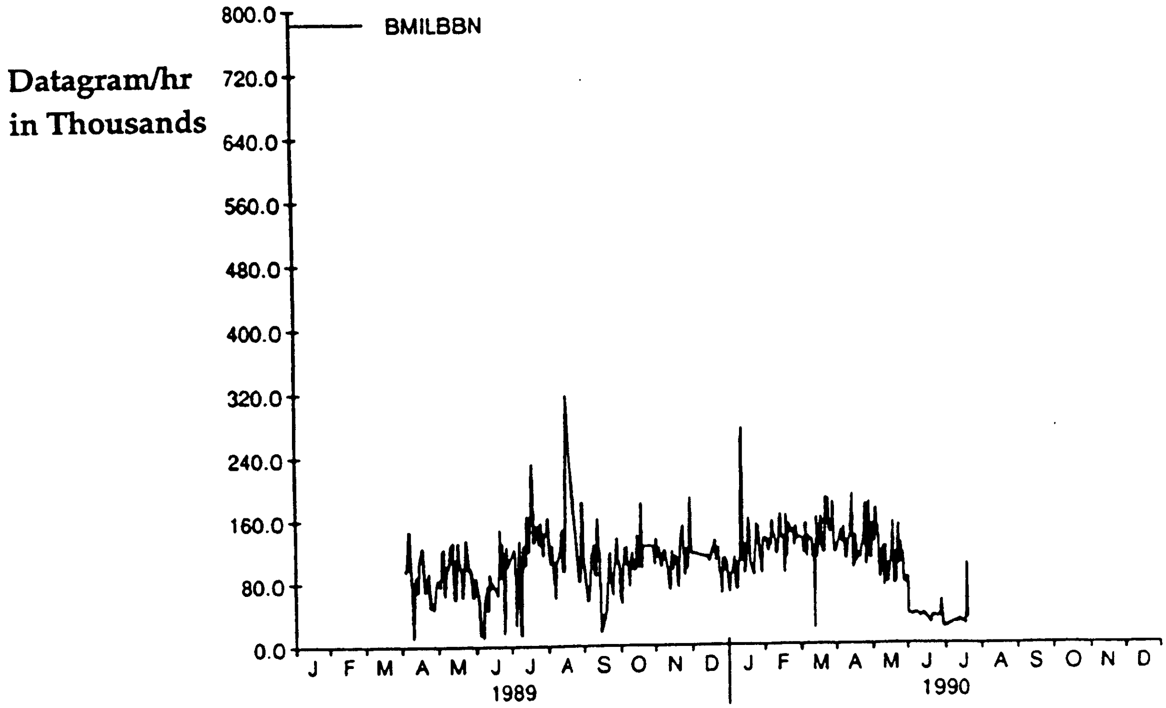

PACKETS DROPPED WEEKLY AVERAGES
(PARTIAL SCALE)

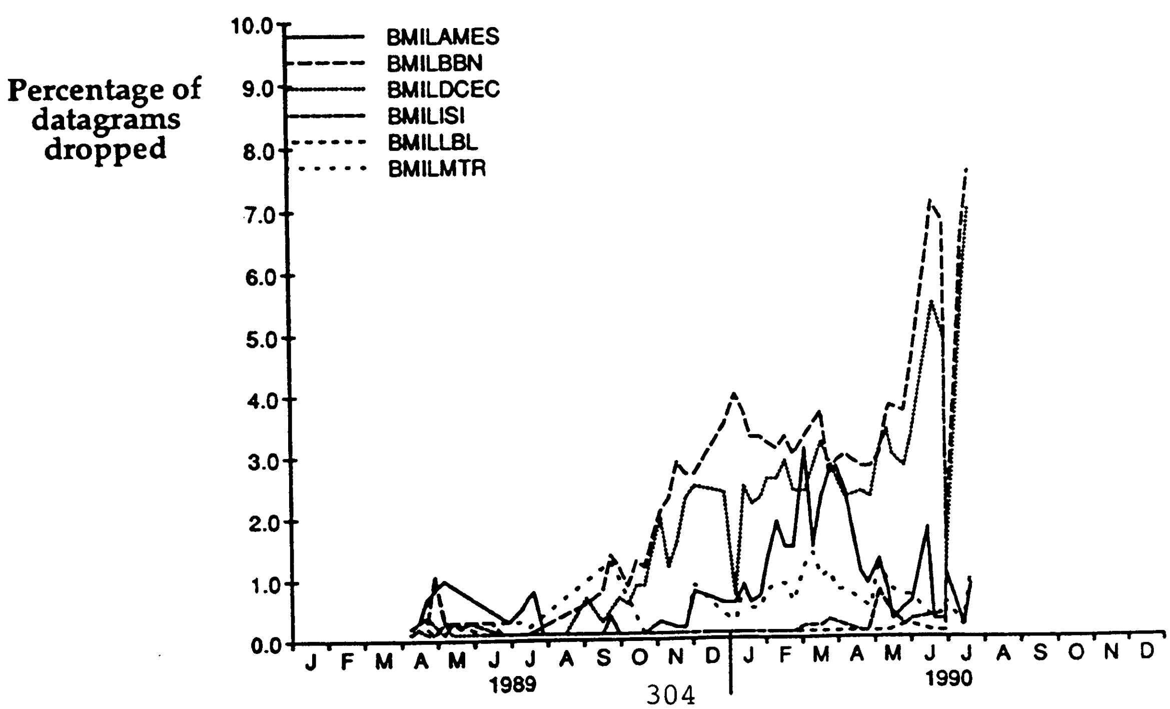

PROBLEMS AND ISSUES
Routing Difficulties

- EGP Server Overload
- "Routing Loops" - no longer see
- "Net Unreachables" - still a problem
  - Outdated static routing tables
    - referencing ARPANET
    - referencing non-existent paths
  - Routing update loss due to EGP server overload

SUMMARY

- Current Actions
  - ARPANET termination
  - Balance EGP service
  - Re-assign AF concentrators
  - Deploy 7th mailbridge
  - Increase EGP Poll Interval
- Future Possibilities
  - Further improve EGP performance
    - distribute EGP processing among multiple processors
    - upgrade hardware to Butterfly Plus Platform
    - EGP replacement

4.2 ESnet

reported by Tony Hain

ESnet

ESNET - STATUS REPORT

JULY 1990

TONY HAIN

ESNET BACKBONE - 1990

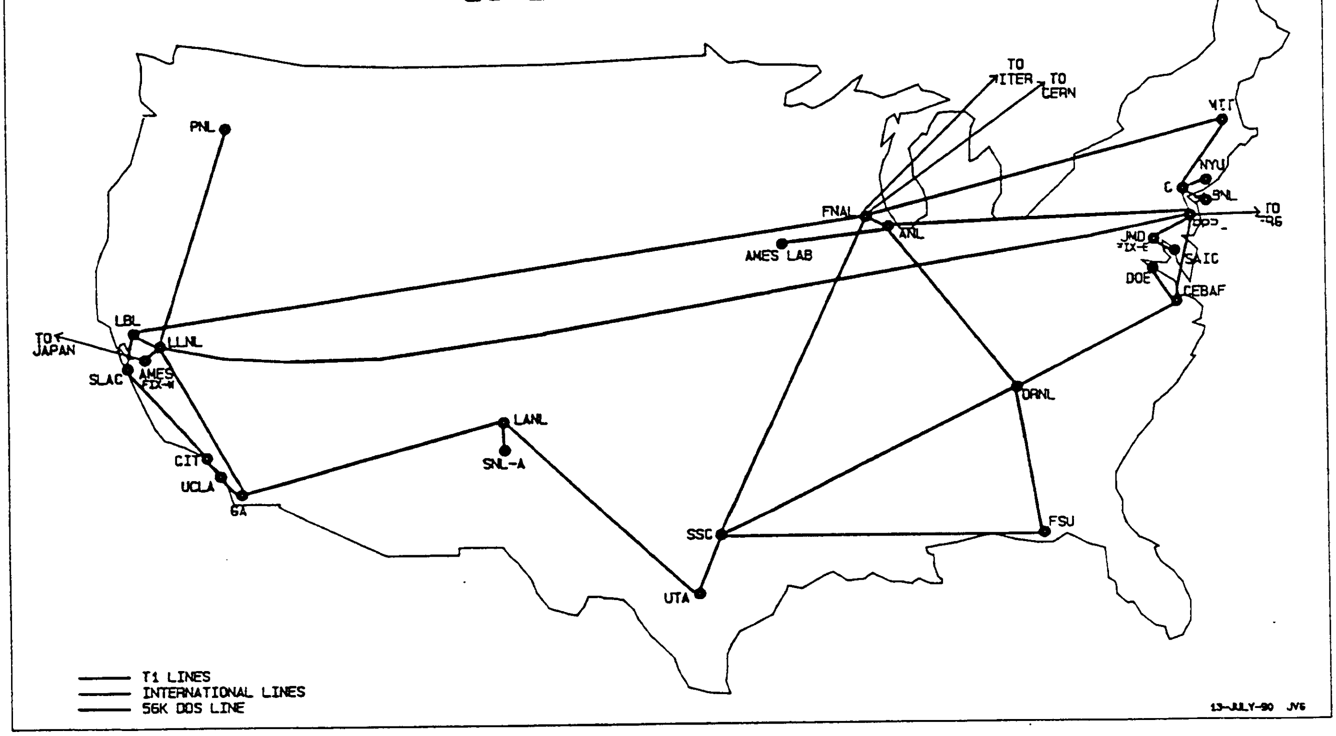

ESnet July '90

PAST ACTIVITIES:

CONTINUED WORKING TROUBLE WITH NNT PROVIDED CIRCUITS

REFINING PEERING DETAILS WITH SOME SITES

FSU & ANL/ORNL T1 LINES INSTALLED

INITIAL IP ROUTING THROUGH GARCHING TO FRG

BEGAN UPGRADE TO CISCO CSC-3 PROCESSORS

DISTRIBUTED NETMGR RELEASE USING MFENET II AS DEFAULT

ESnet July '90

PLANED ACTIVITIES:

COMPLETE CSC-3 UPGRADE

ADD CIRCUITS FOR AMES-IOWA, SAIC, & DOE

DEPLOY CISCO X.25 SWITCHING

SHUTDOWN MFENET I & INTERIM X.25 BACKBONE BY OCT. 1

INSTALL NEW CIRCUIT PPPL/FRG

ESnet

July '90

STATS:

23 Routers Installed

44 Directly Connected Networks

404 Regional Connected Networks via 6 Regionals

.381G Packets Received

65% IP / 35% DECnet

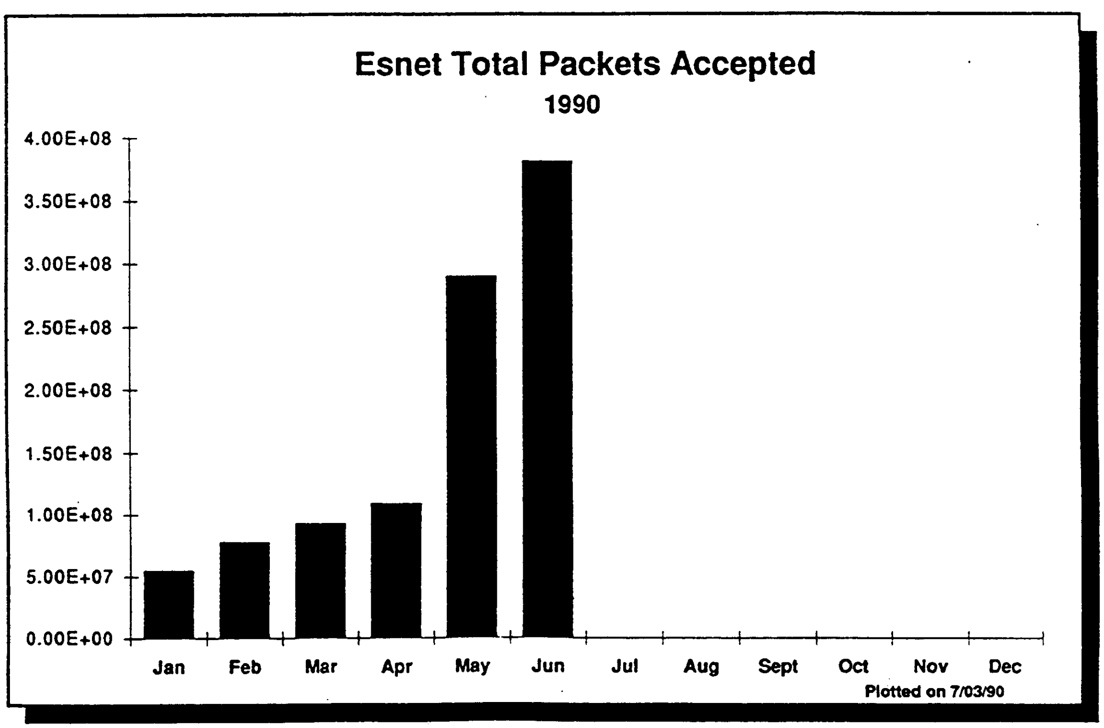

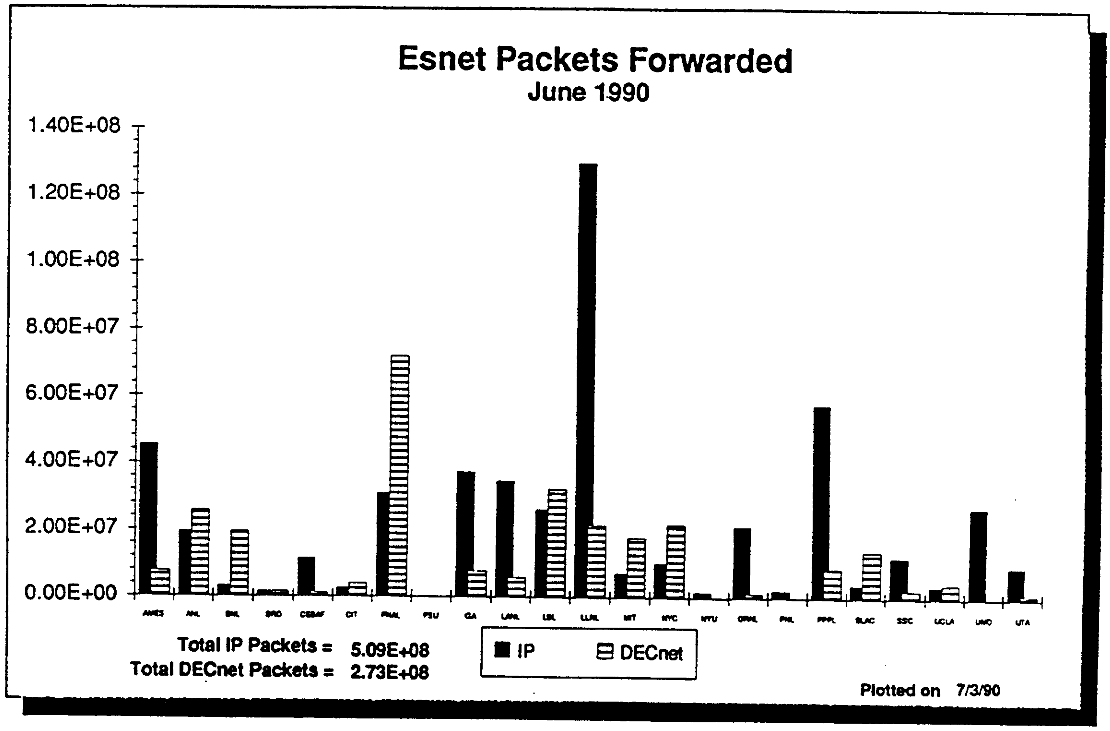

Esnet Packets Forwarded
June 1990
1.40E+08
1.20E+08
1.00E+08
8.00E+07
6.00E+07
4.00E+07
2.00E+07
0.00E+00
Total IP Packets = 5.09E+08
Total DECnet Packets = 2.73E+08
IP
DECnet
Plotted on 7/3/90

4.3 NASA Sciences Internet

Reported by Jeffery G. Burgan

NASA Science Internet
Status Report

July 1990

Jeffrey G. Burgan
Sterling Software
NASA Science Internet Project Office
NASA Ames Research Center

NASA Science Network

- Multiprotocol Routers (50 Proteons)
  - TCP/IP
  - DECnet Phase IV
- Approximately 290 nets
- Direct connections to other National Backbones
  - NSFnet
  - ESnet
  - MILnet
  - TWBnet

- **International Links**
  - Paccom
  - UK "fat pipe"

- **New Connections**
  - Cerro Tololo Inter-American Observatory (CTIO)
  - Fairbanks, Alaska
    - Gilmore Creek
    - Alaska SAR Facility
  - Rice University

NASA Science Network

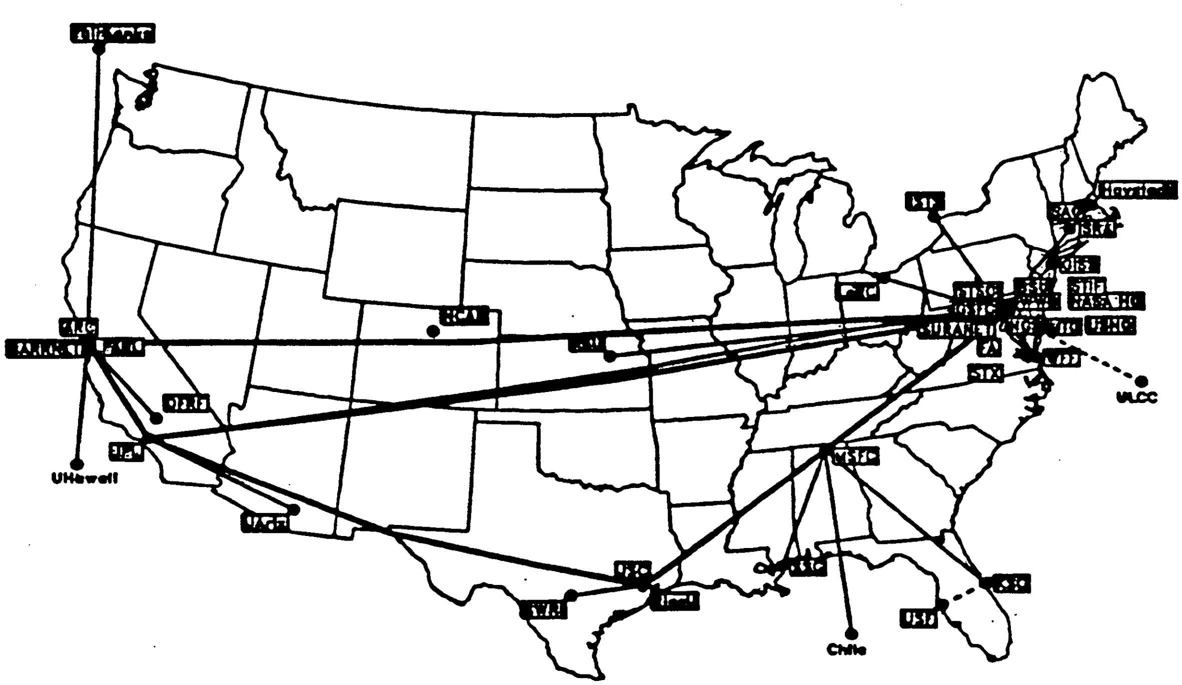

Prepared for: NASA Science Internet Project Office
by: Sterling Software, NAS2-13210
July 20, 1990

OSPF Update

- 11 Routers running OSPF
- Routes
  - 20 Internal SPF
  - Default External Type 1
  - 355 External Type 2
  - 1280 EGP
- OSPF Version 2
  - Stub Area support
  - Forwarding address in External LS Advertisements

NSN OSPF System

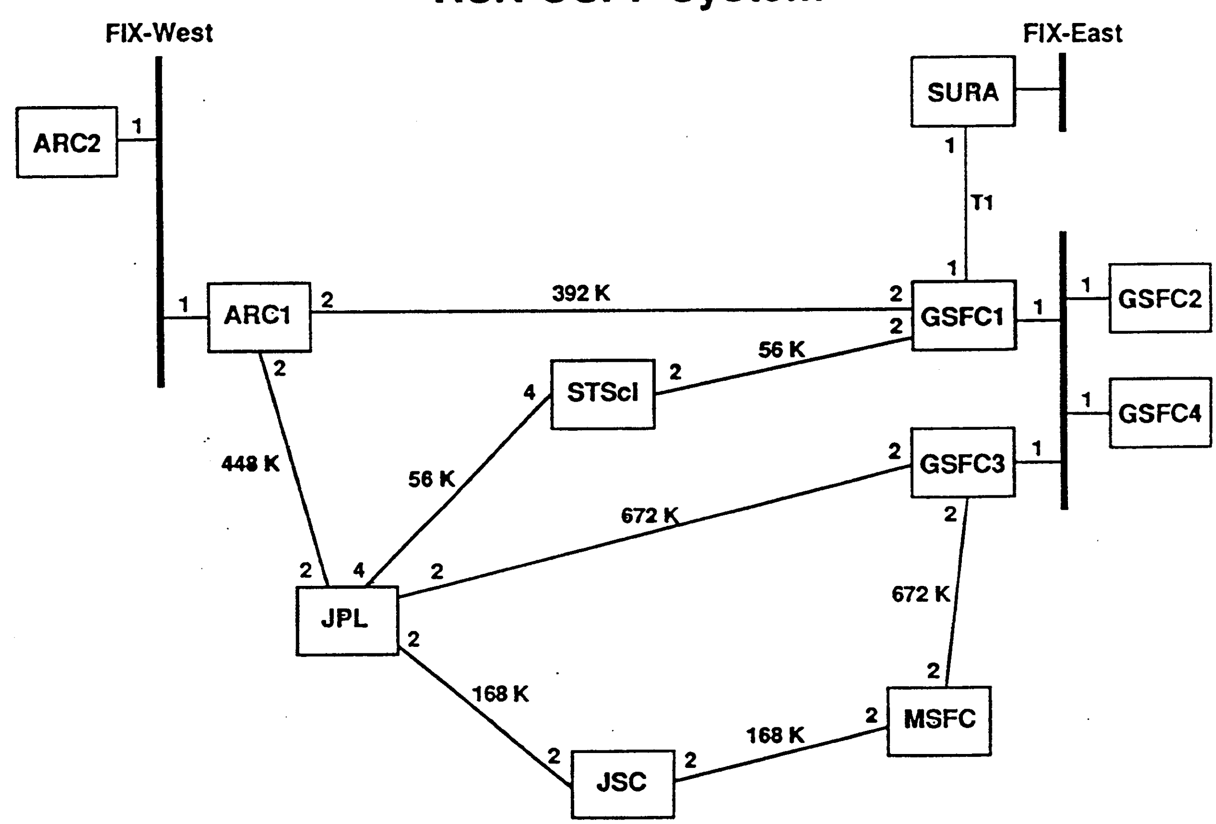

Routing Overhead (RIP vs. OSPF)

RIP (293 nets)

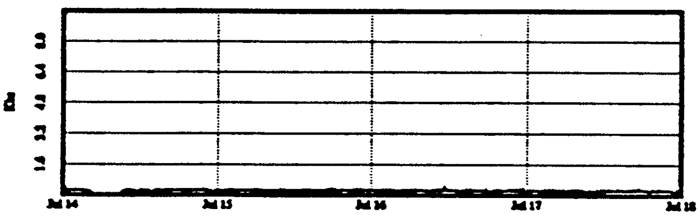

OSPF (394 nets)

4.4 NSFnet

reported by Dale Johnson

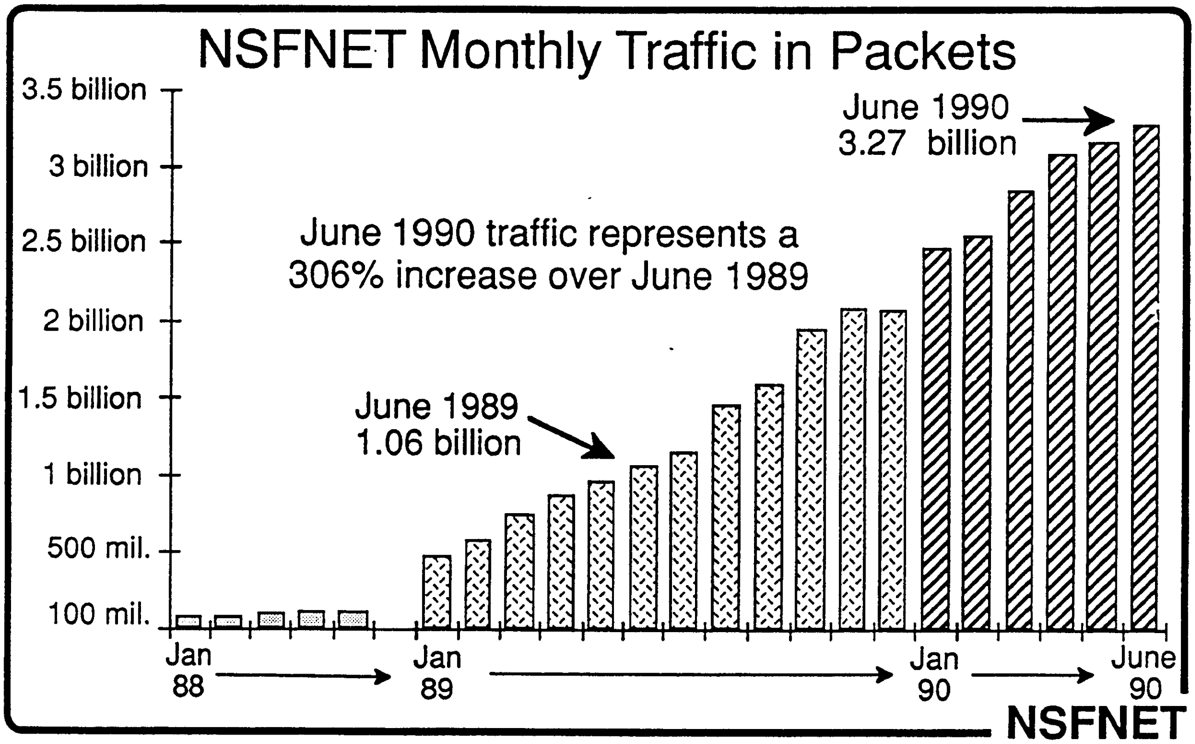
NSFNET Monthly Traffic in Packets
3.5 billion
3 billion
2.5 billion
2 billion
1.5 billion
1 billion
500 mil.
100 mil.
June 1990
3.27 billion
June 1990 traffic represents a
306% increase over June 1989
June 1989
1.06 billion
Jan 88
Jan 89
Jan 90
June 90
NSFNET

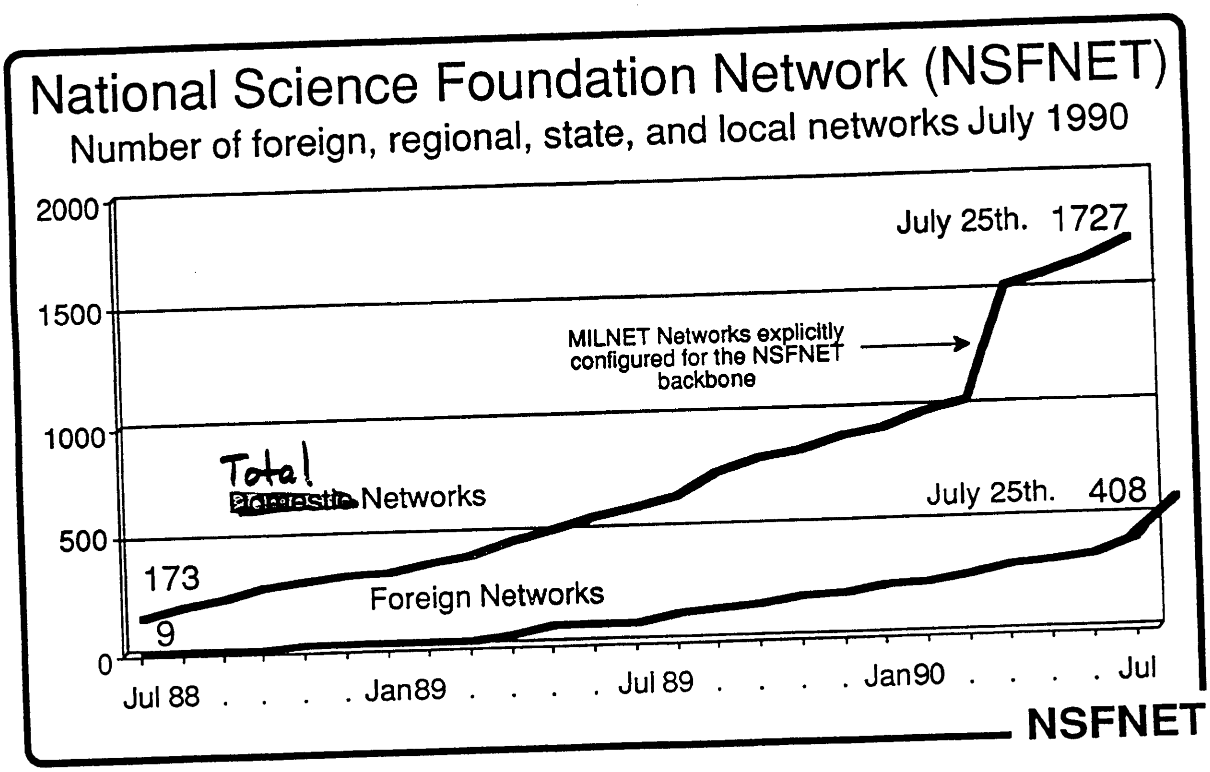
National Science Foundation Network (NSFNET)
Number of foreign, regional, state, and local networks July 1990
2000
1500
1000
500
0
July 25th. 1727
MILNET Networks explicitly
configured for the NSFNET
backbone
Total ~~Domestic~~ Networks
July 25th. 408
173
9
Foreign Networks
Jul 88 . . . . Jan89 . . . . Jul 89 . . . . Jan90 . . . . Jul
NSFNET

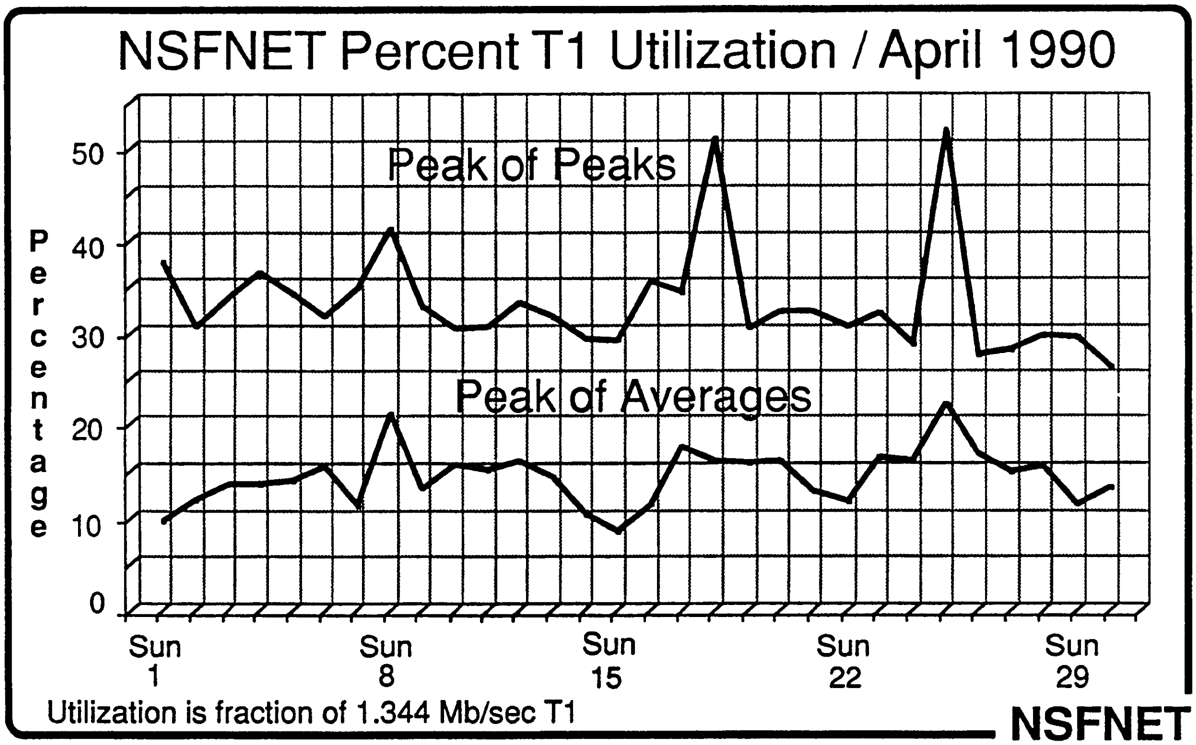

NSFNET Percent T1 Utilization / April 1990
Peak of Peaks
Peak of Averages
Percentage
0
10
20
30
40
50
Sun 1
Sun 8
Sun 15
Sun 22
Sun 29
Utilization is fraction of 1.344 Mb/sec T1
NSFNET

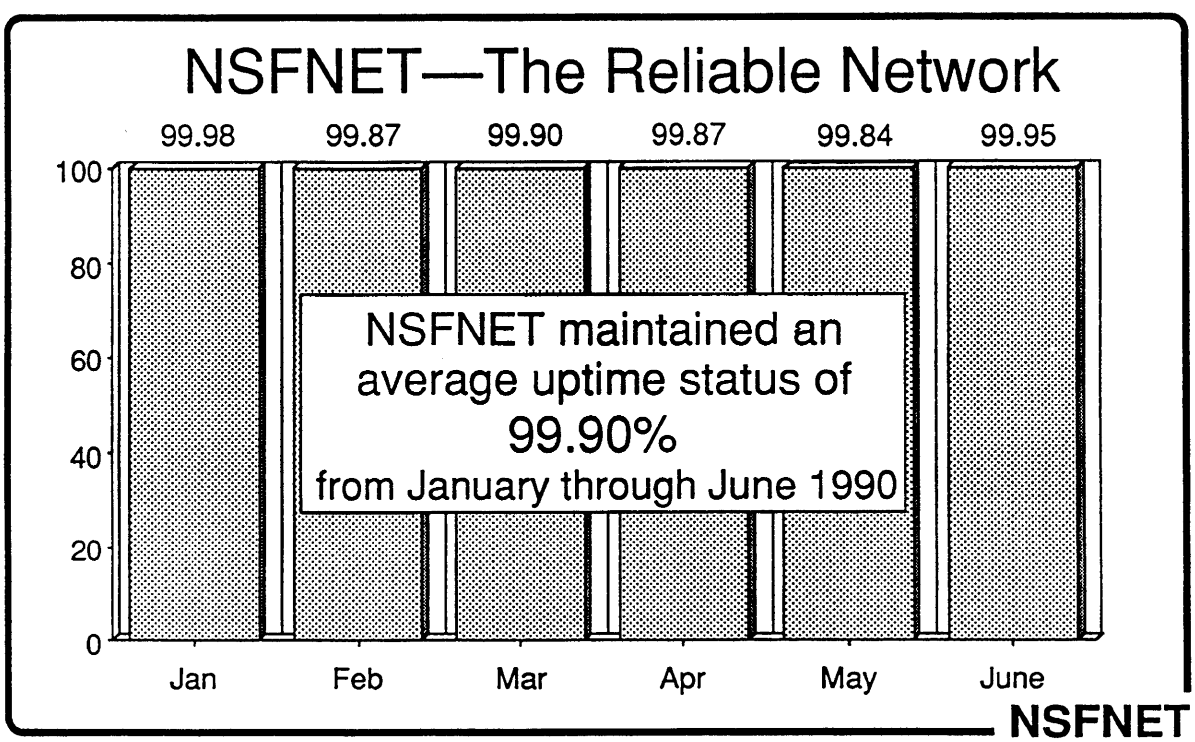

NSFNET—The Reliable Network
99.98
99.87
99.90
99.87
99.84
99.95
0
20
40
60
80
100
Jan
Feb
Mar
Apr
May
June
NSFNET maintained an average uptime status of 99.90% from January through June 1990
NSFNET

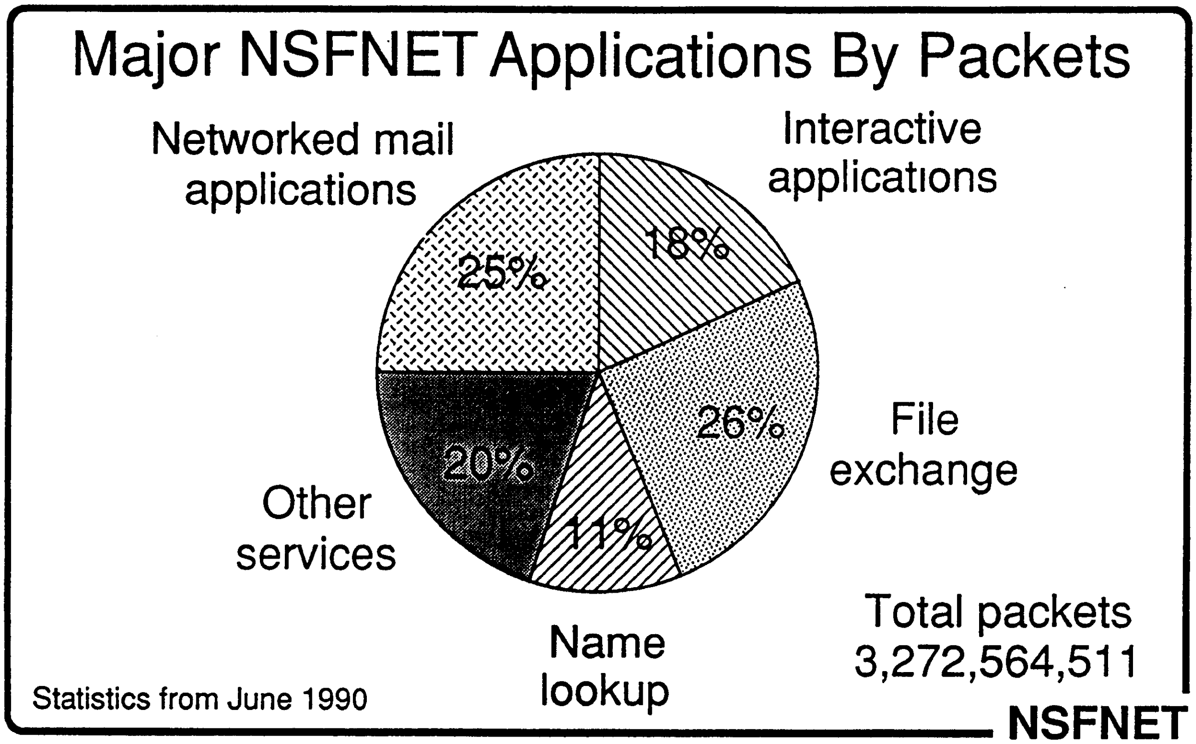
Major NSFNET Applications By Packets
Networked mail applications
25%
Interactive applications
18%
File exchange
26%
Name lookup
11%
Other services
20%
Total packets 3,272,564,511
Statistics from June 1990
NSFNET

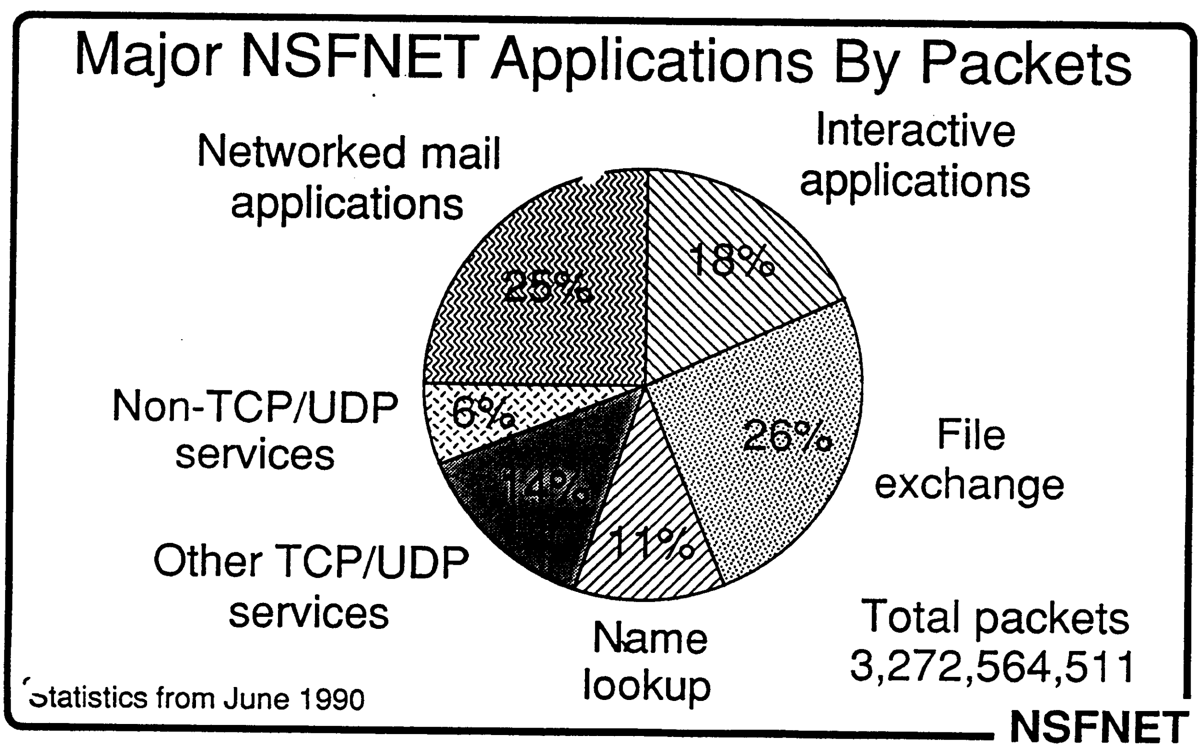
Major NSFNET Applications By Packets
Networked mail applications
25%
Interactive applications
18%
File exchange
26%
Name lookup
11%
Other TCP/UDP services
14%
Non-TCP/UDP services
6%
Total packets 3,272,564,511
Statistics from June 1990
NSFNET

Improvement in Round Trip Time
With New NSS Configuration

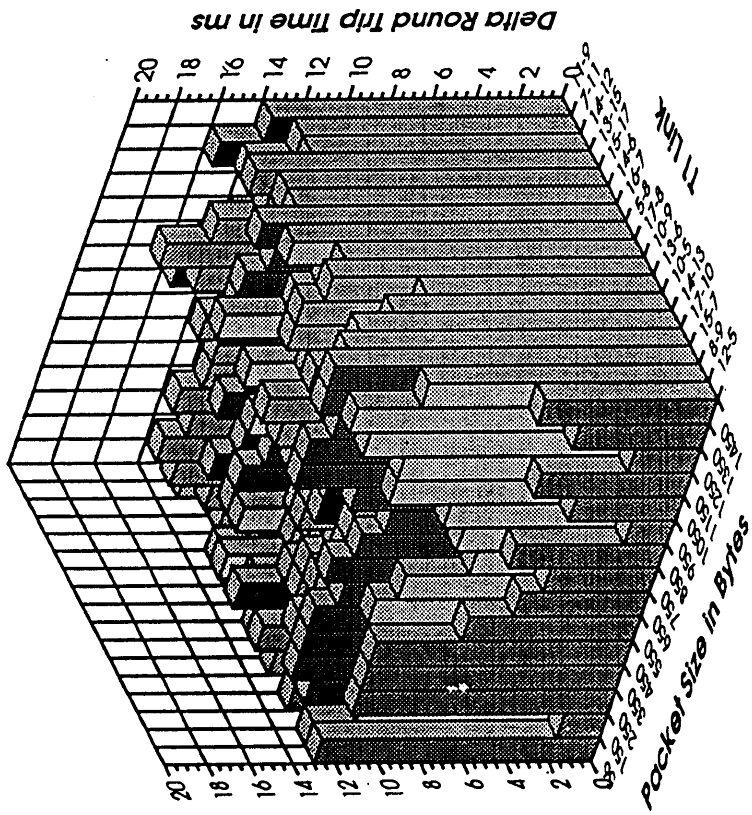

OSI

- Code on the Research Network
- Want to test applications?

  See Sue Hares (skh@merit.edu)

NSFNET

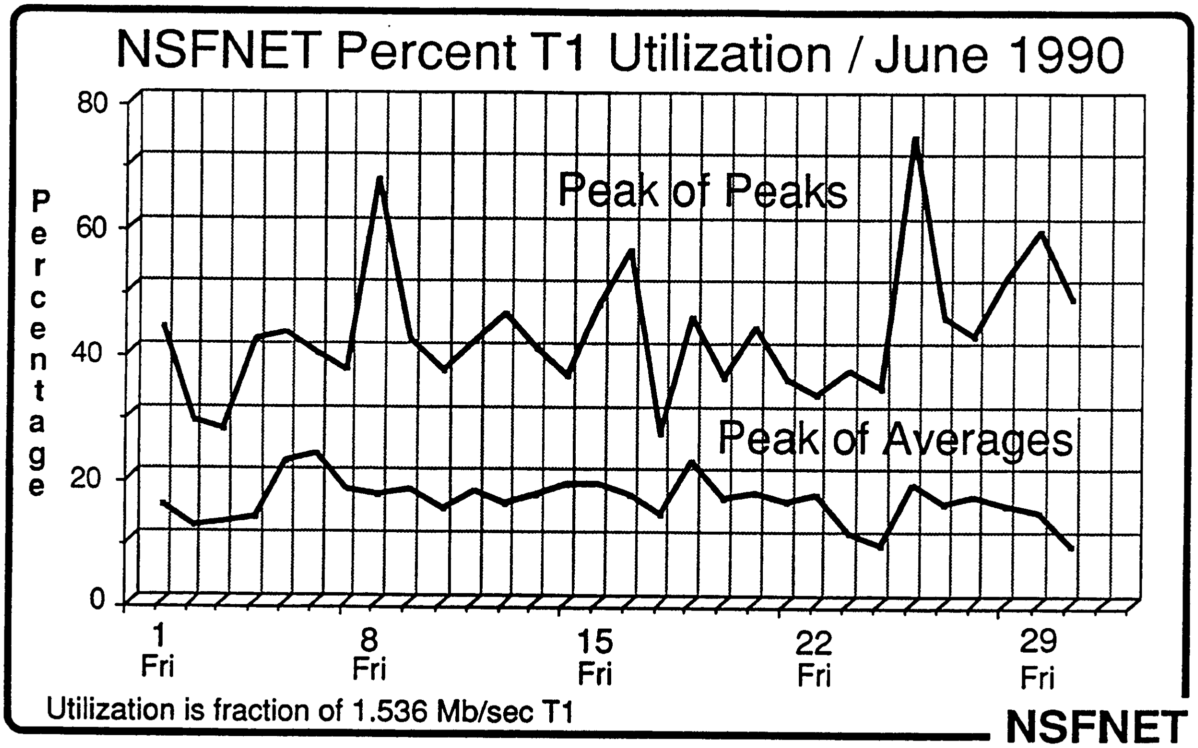
NSFNET Percent T1 Utilization / June 1990
Percentage
80
60
40
20
0
Peak of Peaks
Peak of Averages
1 Fri
8 Fri
15 Fri
22 Fri
29 Fri
Utilization is fraction of 1.536 Mb/sec T1
NSFNET

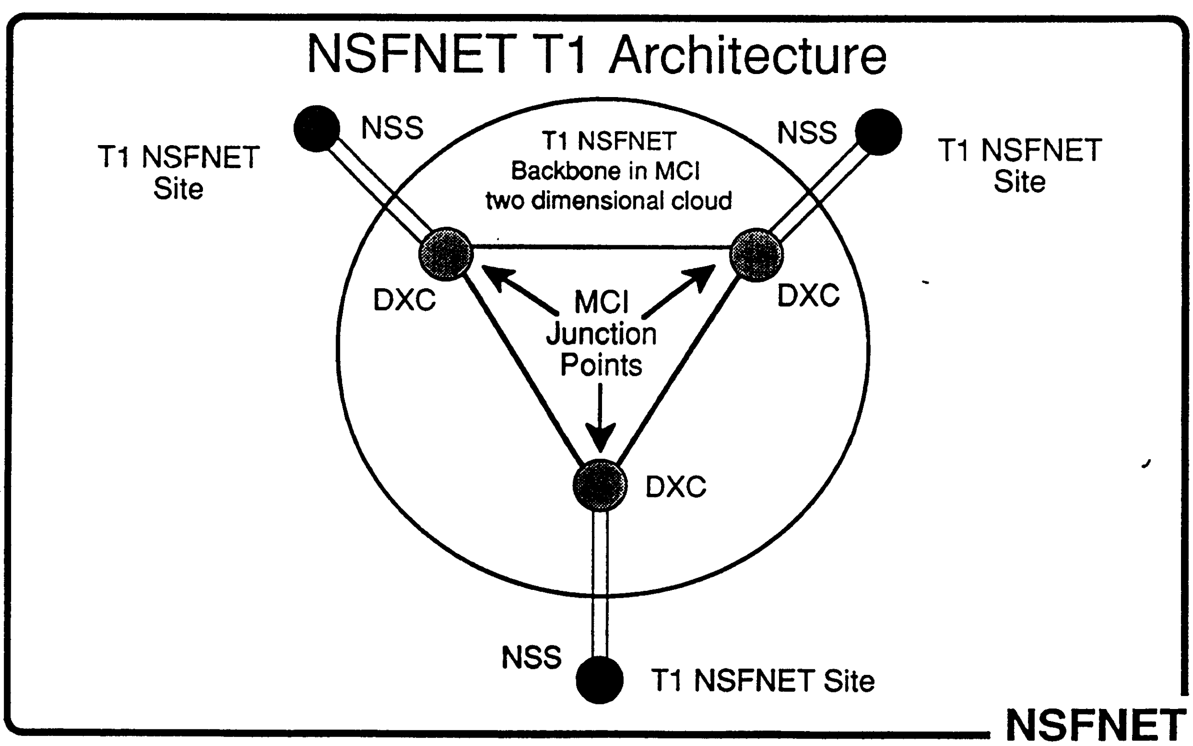
NSFNET T1 Architecture
T1 NSFNET Site
NSS
T1 NSFNET
Backbone in MCI
two dimensional cloud
NSS
T1 NSFNET Site
DXC
MCI
Junction
Points
DXC
DXC
NSS
T1 NSFNET Site
NSFNET

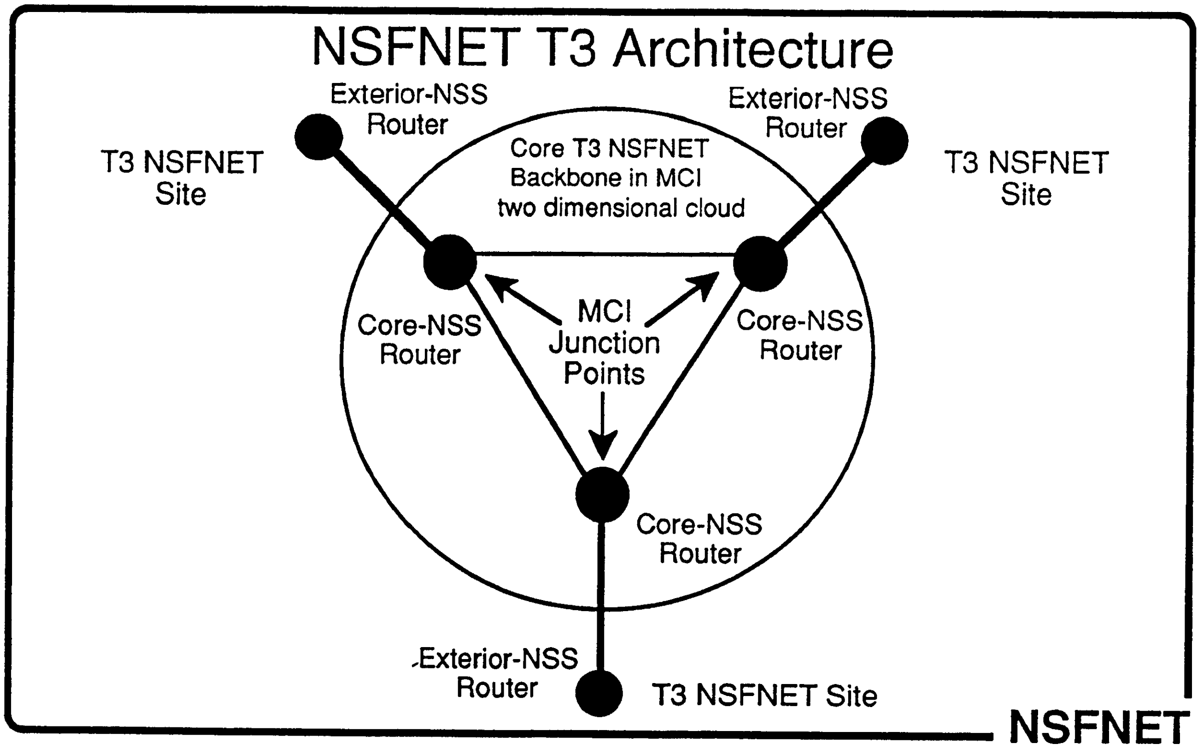
NSFNET T3 Architecture
Exterior-NSS Router
Exterior-NSS Router
T3 NSFNET Site
T3 NSFNET Site
Core T3 NSFNET Backbone in MCI two dimensional cloud
Core-NSS Router
MCI Junction Points
Core-NSS Router
Core-NSS Router
Exterior-NSS Router
T3 NSFNET Site
NSFNET

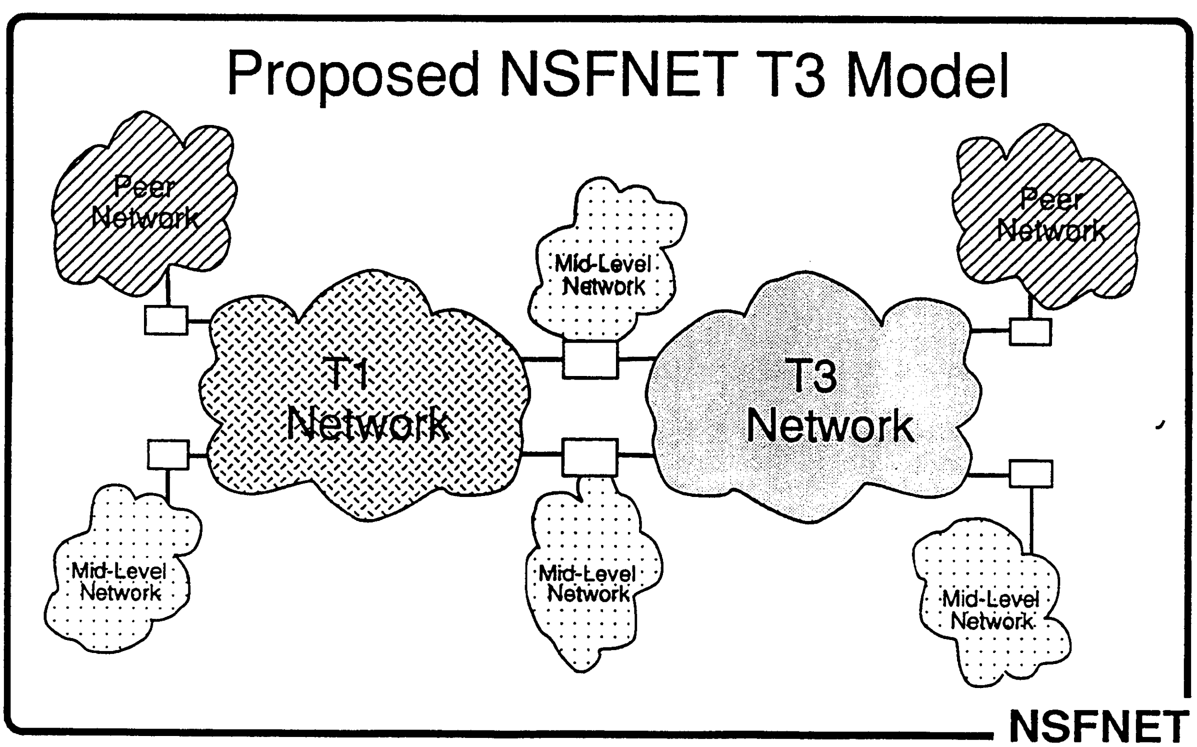
Proposed NSFNET T3 Model
Peer Network
Mid-Level Network
Peer Network
T1 Network
T3 Network
Mid-Level Network
Mid-Level Network
Mid-Level Network
NSFNET

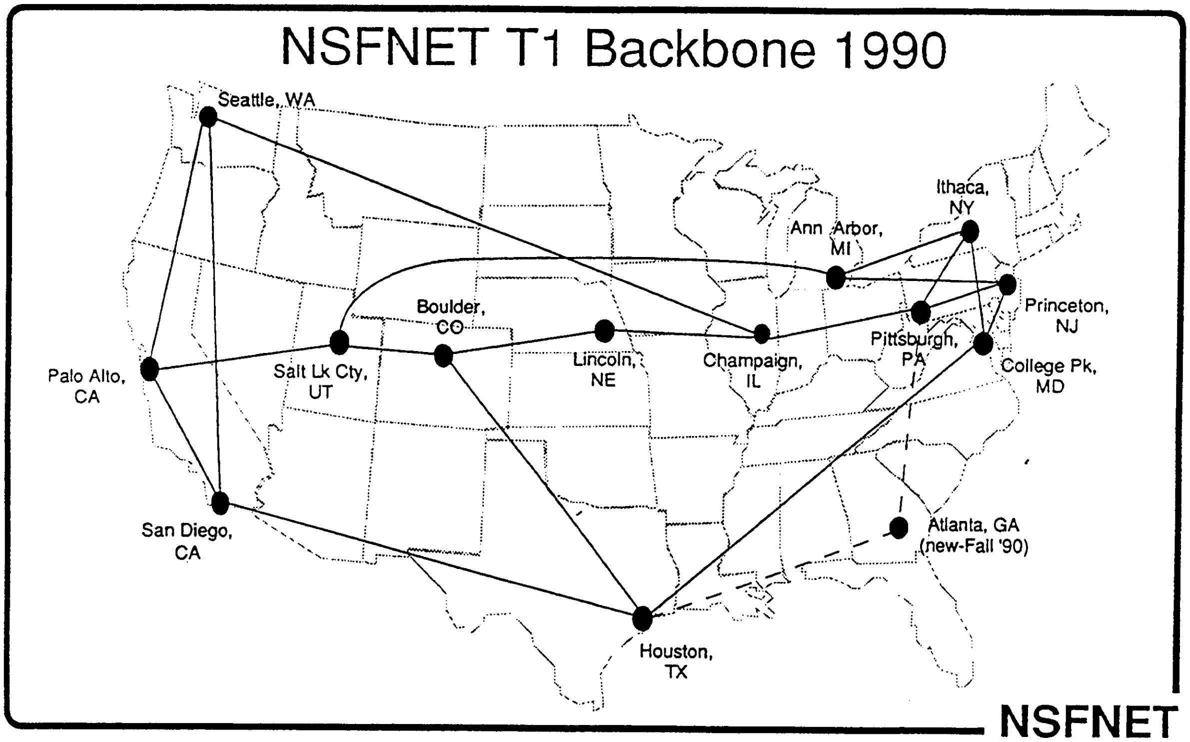
NSFNET T1 Backbone 1990
Seattle, WA
Ithaca, NY
Ann Arbor, MI
Princeton, NJ
Boulder, CO
Pittsburgh, PA
Palo Alto, CA
Salt Lk Cty, UT
Lincoln, NE
Champaign, IL
College Pk, MD
San Diego, CA
Atlanta, GA (new-Fall '90)
Houston, TX
NSFNET

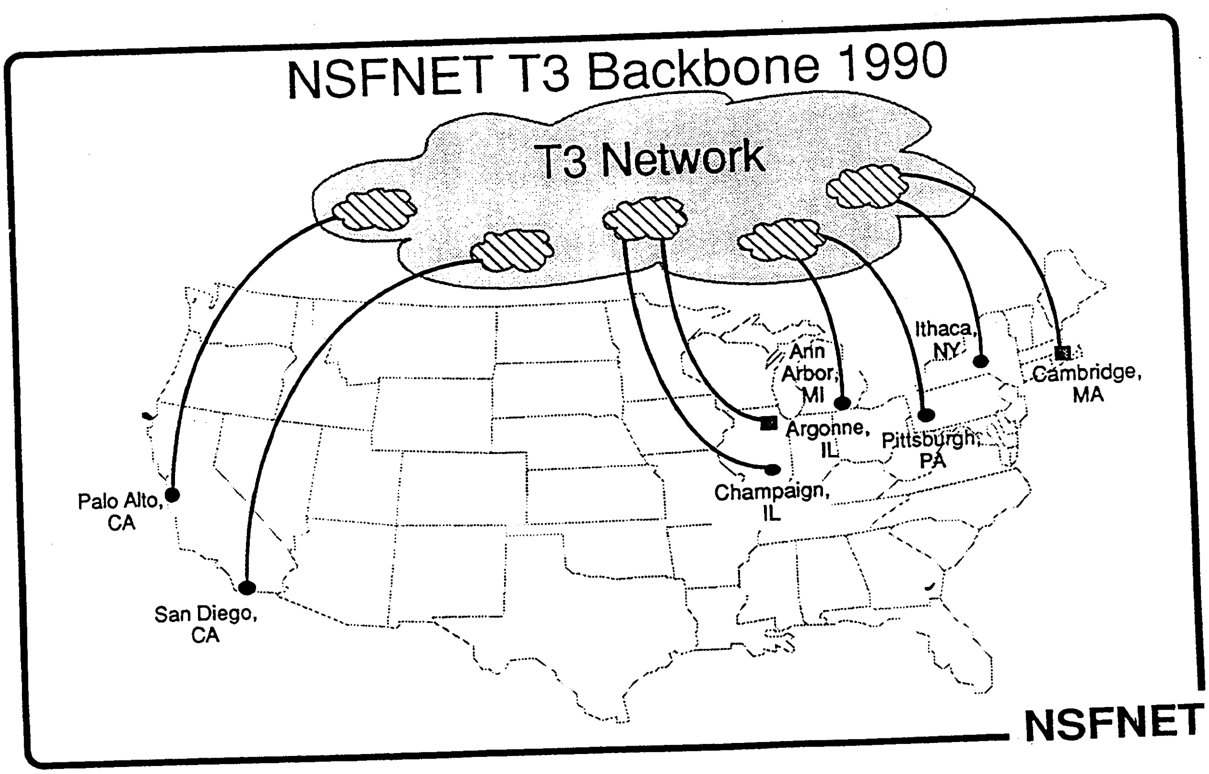
NSFNET T3 Backbone 1990
T3 Network
Ithaca, NY
Ann Arbor, MI
Cambridge, MA
Argonne, IL
Pittsburgh, PA
Palo Alto, CA
Champaign, IL
San Diego, CA
NSFNET

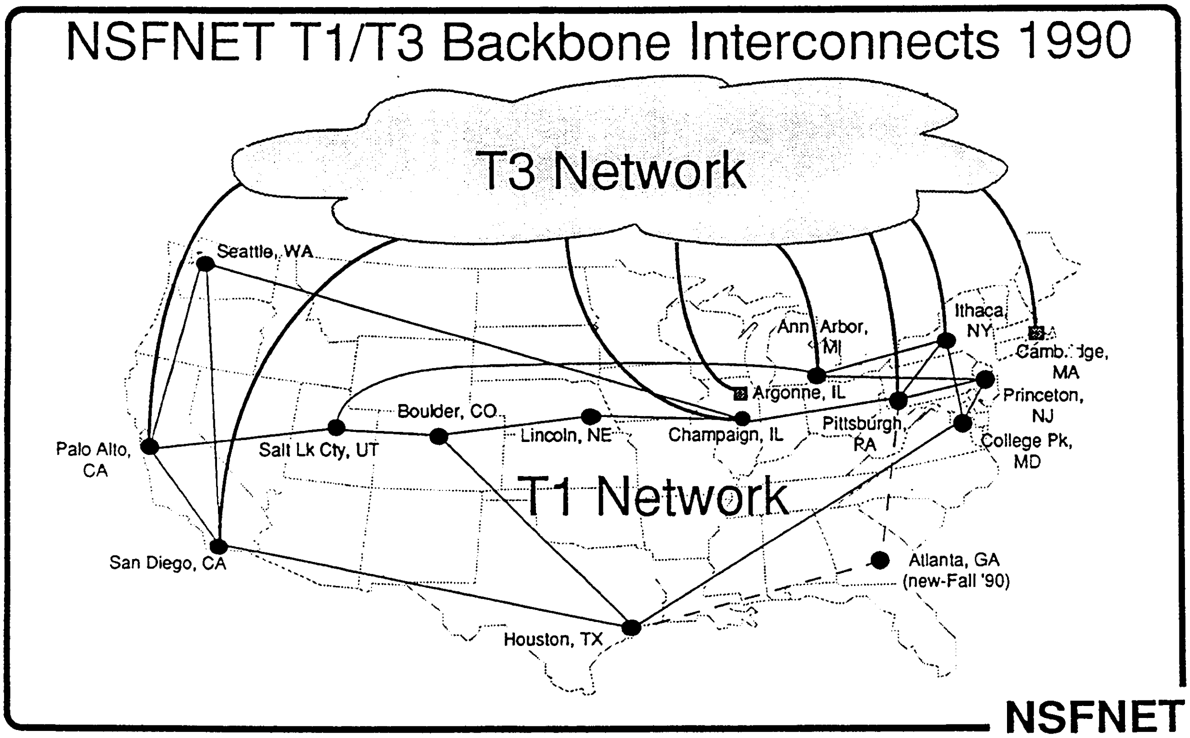
NSFNET T1/T3 Backbone Interconnects 1990
T3 Network
Seattle, WA
Ann Arbor, MI
Ithaca, NY
Camb.·.dge, MA
Argonne, IL
Princeton, NJ
Boulder, CO
Lincoln, NE
Champaign, IL
Pittsburgh, PA
College Pk, MD
Palo Alto, CA
Salt Lk Cty, UT
T1 Network
San Diego, CA
Atlanta, GA (new-Fall '90)
Houston, TX
NSFNET

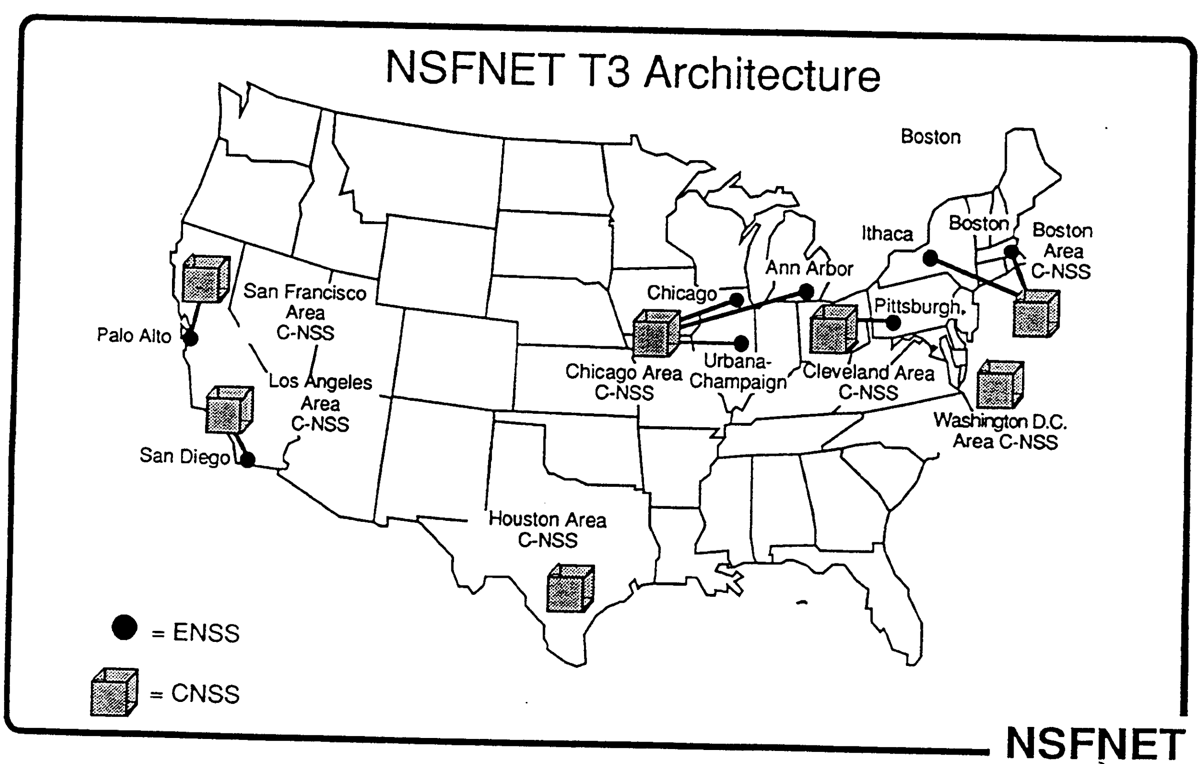
NSFNET T3 Architecture
Boston
Ithaca
Boston
Boston Area C-NSS
Ann Arbor
Chicago
Pittsburgh
San Francisco Area C-NSS
Palo Alto
Urbana-Champaign
Chicago Area C-NSS
Cleveland Area C-NSS
Los Angeles Area C-NSS
Washington D.C. Area C-NSS
San Diego
Houston Area C-NSS
= ENSS
= CNSS
NSFNET

Chapter 5

IETF Protocol Presentations

5.1 CMIP over TCP/IP

Presentation by Brian Handspicker/DEC

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

This memo is a revision of RFC 1095 - "The Common Management Information Services and Protocol over TCP/IP" [27]. It defines a network management architecture that uses the International Organization for Standardization's (ISO) Common Management Information Services/Common Management Information Protocol (CMIS/CMIP) in the Internet. This architecture provides a means by which control and monitoring information can be exchanged between a manager and a remote network element. In particular, this memo defines the means for implementing the International Standard (IS) version of CMIS/CMIP on top of both IP-based and OSI-based Internet transport protocols for the purpose of carrying management information defined in the Internet-standard management information base. Together with the relevant ISO standards and the companion RFCs that describe the initial structure of management information and management information base, these documents provide the basis for a comprehensive architecture and system for managing both IP-based and OSI-based internets, and in particular the Internet.

In creating this revision of RFC 1095, the following technical and editorial changes were made:

- The tutorial section on OSI Management included in RFC 1095 has been removed from this document. After some revisions, the tutorial material may be reintroduced as an Internet FYI Draft.
- The sections in RFC 1095 which discussed the semantics of how to interpret requests in the context of Internet MIBs has been removed from this protocol document. This topic is now discussed in the OIM-MIB-II document. This protocol should be useable with MIB-I or MIB-II. But, it will also be able to

exploit the new features of the OIM-MIB-II.

- This document is based on the final International Standards for CMIS/CMIP (ISO 9595/9596) rather than the Draft International Standards.
- Many of the original agreements defined in RFC 1095 have been accepted and included in the OIW NMSIG implementors agreements. Rather than duplicating these agreements, they have been removed from this revision of RFC 1095. This document should be read in conjunction with ISO 9595/9596 (CMIS/CMIP) and the OIW Stable Agreements document.
- The Association Negotiation describe in RFC 1095 has been changed to align with current international and national agreements. But, it has retained backwards compatibility with the assignment of an Application Context Name for 1095CMOT which is identical to the Application Context Name specified in RFC 1095.

CMIP over TCP/IP

Brian D. Handspicker

OSI Internet Management CoChair

bd@vines.enet.dec.com

508-486-7894

CMIP over TCP/IP

Revision of RFC 1095

- Protocol based on IS CMIP
- Application Context Negotiation based on DIS SMO
- Based on NIST OIW Stable Implementers Agreements
- 1095 Tutorial Removed
- 1095 MIB moved to OIM-MIB-II
- Submitted to Internet Drafts

CMIP over TCP/IP

Diffs between DIS & IS CMIP

- Use ACSE Association Services directly
- Application Context & Functional Units Changed
- Scoping extended
- Superior Class parameter added to CREATE
- Mandatory parameter now optional (Class for Action)
- Optional return parameter now Mandatory (Invoke Id)
- Additional Error Codes
- Semantic Clarifications
- Editorial Changes

Add/Remove Addendum: Add/Remove support ~~for SET~~

CMIP over TCP/IP

Backward Compatibility (IS → DIS)

MOOT: Vendors in OIM committed to products based on IS CMIP

CMIP over TCP/IP

CMOT & Internet MIBs

Support for IETF SMI (RFC 1065)
Support for OSI SMI (DIS 10165-4)

100% Support for IETF SMI-based MIBs

OIM extensions to MIB-II ease naming

CMIP over TCP/IP

IOP Testing

DIS & IS CMOT

OpenLab July 23-26

GET, SET, CREATE, DELETE tested

Problems w/ Interpretation & version

Corrected on-the-fly w/ Patches

Participants:

3 Com	DIS
Digital	DIS, IS
HP	IS
Netlabs	DIS
	DIS, IS

Independent:

Mitre	DIS, IS
UW	DIS, IS
3 Com	DIS

CMIP over TCP/IP

Conclusion

6 organizations testing implementations
12 organizations planning products
Interop Product demonstrations
Continued Government & European Interest

Revise RFC 1095 as Draft Standard

Chapter 6

Technical Presentations

6.1 IMAP Services

Presentation by Mark Crispin

The intent of the Interactive Mail Access Protocol, Version 2 (IMAP2) is to allow a workstation, personal computer, or similar small machine to access electronic mail from a mailbox server. Since the distinction between personal computers and workstations is blurring over time, it is desirable to have a single solution that addresses the need in a general fashion. IMAP2 is the "glue" of a distributed electronic mail system consisting of a family of client and server implementations on a wide variety of platforms, from small single- tasking personal computing engines to complex multi-user timesharing systems.

Although different in many ways from the Post Office Protocols (POP2 and POP3, hereafter referred to collectively as "POP") described in RFC 937 and RFC 1081, IMAP2 may be thought of as a functional superset of these. RFC 937 was used as a model for this RFC. There was a cognizant reason for this; POP deals with a similar problem, albeit with a less comprehensive solution, and it was desirable to offer a basis for comparison.

Like POP, IMAP2 specifies a means of accessing stored mail and not of posting mail; this function is handled by a mail transfer protocol such as SMTP (RFC 821).

This protocol assumes a reliable data stream such as provided by TCP or any similar protocol. When TCP is used, the IMAP2 server listens on port 143.

IMAP SERVICES

Mark Crispin
University of Washington
mrc@CAC.WASHINGTON.EDU
August 2, 1990 2:15 PM

What is IMAP?

- a protocol to access electronic mailboxes from a remote site
- an operating-system independent representation of electronic mail
- a pre-processor for mail

IMAP PROTOCOL

- Text-based. No binary
- Structured
- Allows division of labor between client and server
- Allows server to manipulate client state without explicit client action

```
S:   * OK SUMEX-AIM.Stanford.EDU Interim Mail Access Protocol II
       Service 6.1(349) at Thu, 9 Jun 88 14:58:30 PDT
U:   a001 login crispin secret
S:   a002 OK User CRISPIN logged in at Thu, 9 Jun 88 14:58:42 PDT, job 76
U:   a002 select inbox
S:   * FLAGS (Bugs SF Party Skating Meeting Flames Request AI Question
       Note \\XXXX \\YYYY \\Answered \\Flagged \\Deleted \\Seen)
S:   * 16 EXISTS
S:   * 0 RECENT
S:   a002 OK Select complete
U:   a003 fetch 16 all
S:   * 16 Fetch (Flags (\\Seen) InternalDate " 9-Jun-88 12:55:44 PDT"
       RFC822.Size 637 Envelope ("Sat, 4 Jun 88 13:27:11 PDT"
       "INFO-MAC Mail Message" (("Larry Fagan" NIL "FAGAN"
       "SUMEX-AIM.Stanford.EDU")) (("Larry Fagan" NIL "FAGAN"
       "SUMEX-AIM.Stanford.EDU")) (("Larry Fagan" NIL "FAGAN"
       "SUMEX-AIM.Stanford.EDU")) ((NIL NIL "rindflEISCH"
       "SUMEX-AIM.Stanford.EDU")) NIL NIL NIL
       "<12403828905.13.FAGAN@SUMEX-AIM.Stanford.EDU>"))
S:   a003 OK Fetch completed
U:   a004 fetch 16 rfc822
S:   * 16 Fetch (RFC822 {637}
S:   Mail-From: RINDFLEISCH created at  9-Jun-88 12:55:43
S:   Mail-From: FAGAN created at  4-Jun-88 13:27:12
S:   Date: Sat, 4 Jun 88 13:27:11 PDT
S:   From: Larry Fagan  <FAGAN@SUMEX-AIM.Stanford.EDU>
S:   To: rindflEISCH@SUMEX-AIM.Stanford.EDU
S:   Subject: INFO-MAC Mail Message
S:   Message-ID: <12403828905.13.FAGAN@SUMEX-AIM.Stanford.EDU>
S:   ReSent-Date: Thu, 9 Jun 88 12:55:43 PDT
S:   ReSent-From: TC Rindfleisch <Rindfleisch@SUMEX-AIM.Stanford.EDU>
S:   ReSent-To: Yeager@SUMEX-AIM.Stanford.EDU, Crispin@SUMEX-AIM.Stanford.EDU
S:   ReSent-Message-ID: <12405133897.80.RINDFLEISCH@SUMEX-AIM.Stanford.EDU>
S:
S:   The file is <info-mac>usenetv4-55.arc  ...
S:   Larry
S:   -------
S:   )
S:   a004 OK Fetch completed
U:   a005 logout
S:   * BYE DEC-20 IMAP II server terminating connection
S:   a005 OK SUMEX-AIM.Stanford.EDU Interim Mail Access Protocol Service logout
```

IMAP History

1985 Early experiments

1986 DEC-20 server, Xerox Lisp client. First IMAP protocol

1987 IMAP2 developed. DEC-20 server and Xerox Lisp client updated. Unix server started.

1988 IMAP2 published. Mac client started. Unix client.

1989 NeXT client

1990 IMAP for the world

PROTOCOL CONCEPTS

. Must be implementable on a small machine
. Minimize network traffic
. Make it simple for small clients to do complex things - envelope parsing and searching in the server
. have a featureful state on the server mailbox
. COMPATABILITY!!!

PROTOCOL DETAILS:

. tagged requests and responses
. untagged responses used to transmit data to client
. protocol requests:
  . Login / logout
  . Fetch data
  . Store status
  . Search
  . Expunge mailbox
  . Copy messages to folder
  . Select new folder

SOFTWARE STATUS

SERVERS:
  . BSD UNIX
  . DEC-20 - no further development

CLIENTS:
  . Xerox Lisp
  . TI Explorer / Common Lisp
  . C-client:
    . Macintosh
    . NeXT (2 clients)
    . UNIX
    . PC under development

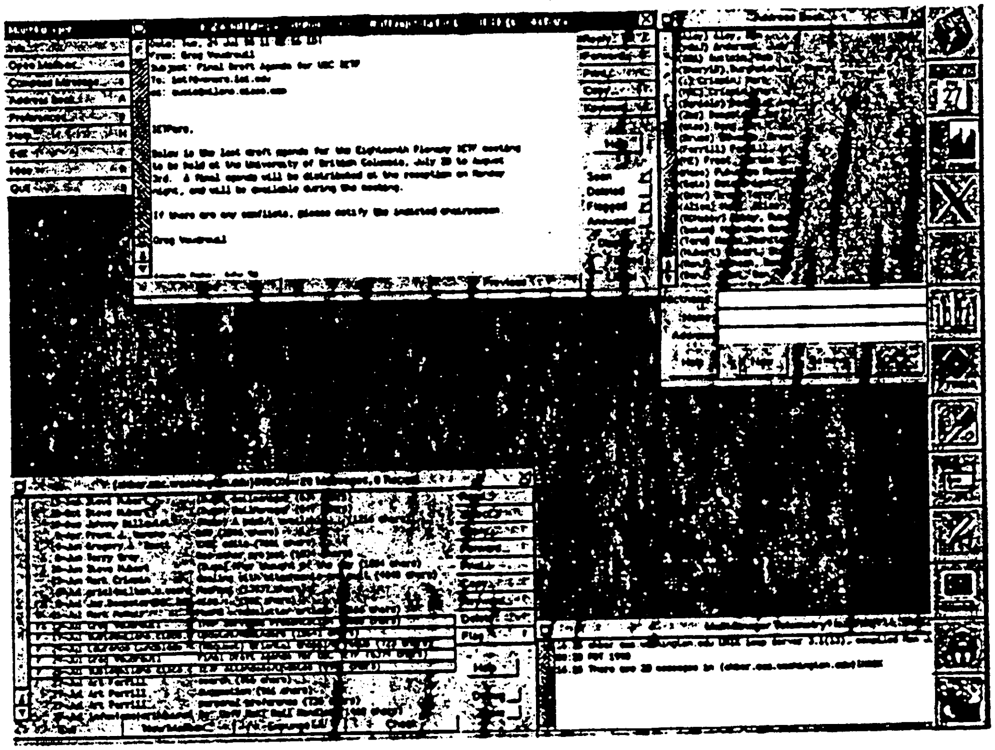

C - Client

A library of C functions to deal with IMAP and local mailboxes in a transparent manner.

FTPHOST.CAC.WASHINGTON.EDU
128.95.112.1
pub/imap.tar.Z

6.2 CA*net

Presentation by Dennis Ferguson

CA*net Status Report

Dennis Ferguson

University of Toronto

The Internet in Canada Prior to CA*net

- Five regional (read: provincial) networks
  - BCnet (British Columbia), with a connection to the NSFnet at Seattle
  - ARnet (Alberta Research network), connected to BCnet
  - Onet (Ontario), with a connection to the NSFnet at Ithaca
  - RISQ (Recherche Interordinateurs Scientifique Quebecois), a NYSERnet member
  - NSTN (Nova Scotia Technical Network), connected to RISQ
- A few miscellaneous "mission oriented" connections
  - Defence Research Establishment, a NYSERnet member
  - An NSI connection to ISTS
  - One or two MILNET sites

CA*net's Reason for Being

- A very long time ago the National Research Council (not quite the equivalent of the NSF) began to talk about an IP network
- The Internet is a desirable research resource. There was considerable demand
- Most places in Canada are closer to the US than they are to each other. Proliferation of southbound links was a real possibility
- The above is illegal
- Regional disparity. Some places are neither close to each other nor the US
- Perhaps a sense that it is appropriate to pay one's own way

The Establishment of CA*net

- The NRC published an RFI for parties interested in the establishment of CA*net in December, 1988, and an RFP in July, 1989
- A proposal submitted jointly by the University of Toronto, IBM Canada and Integrated Network Services Inc (INSINC) was successful
- The University of Toronto agreed to install and operate the network
- IBM Canada provided routers at an advantageous price
- INSINC, a telecommunications reseller, provided the circuits at close to their cost

CA*net Overview

- A backbone network serving "regionals"
- Ten nodes (there are ten provinces)
- Nine internal links (the minimum required to connect ten nodes...)
- All internal links are 56 kbps
- Three links to the NSFnet, Vancouver-Seattle, Toronto-Ithaca and Montreal-Princeton
- NSFnet links are small fractions of T1

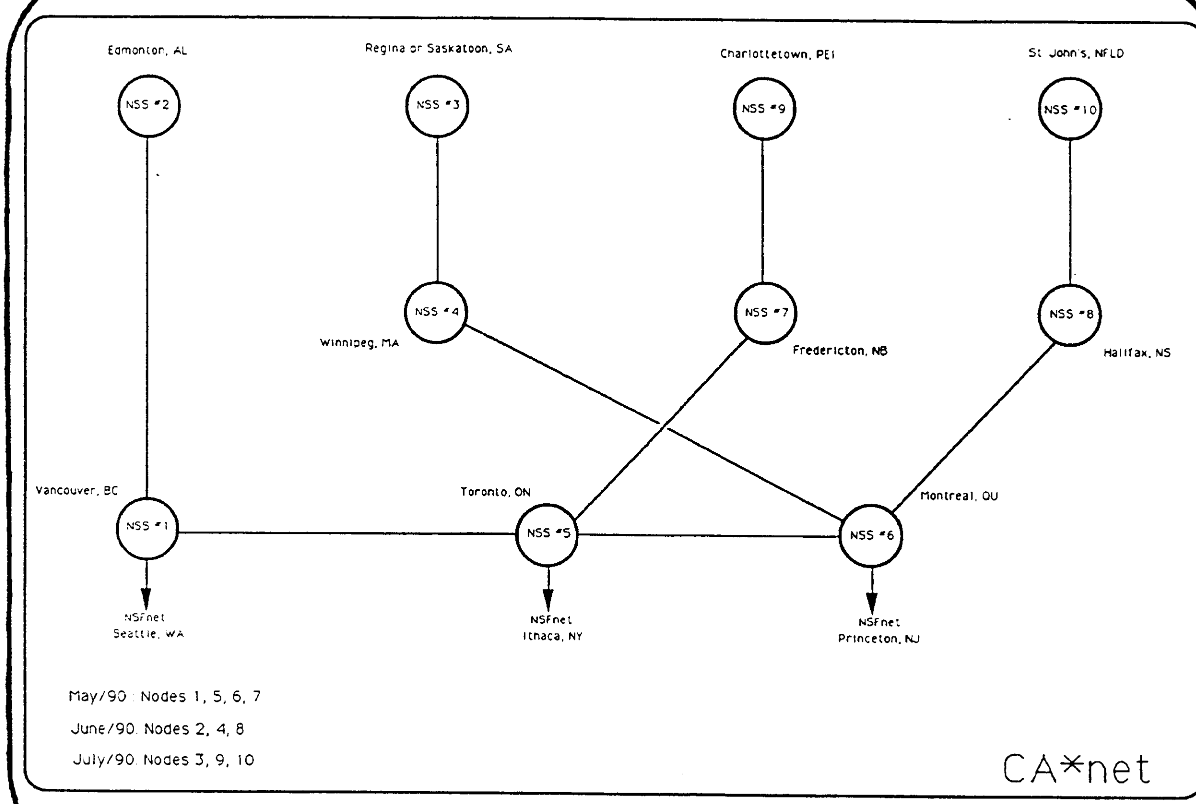

Figure 1

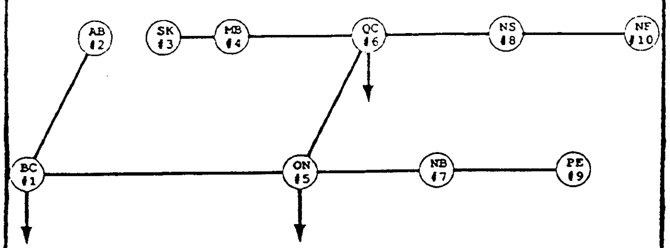

Who Pays, and for What?

- The network's furthest eastern point is about 3800 miles (4 1/2 hours worth of time zones) from its furthest western point
- The network consists of over 7000 miles of DS0
- The cost of digital circuits here is about 8 to 10 times the cost in the US
- There aren't a lot of people to pay for this
- The NRC contribution provided about 60% of the funding for three years of operation
- The other 40% will be paid for by the regional networks
- The NSF paid for the US portions of the three links to the NSFnet, a contribution for which we are very grateful
- The IBM Canada pricing amounted to a considerable donation
- INSINC provided circuits at cost, for other considerations
- This all was still painful. The network in the RFC was not the network that was built

Installation

- A tight schedule. The equipment was delivered in the third week of April for preparation. As was the money to pay for anything. The first four nodes were scheduled for installation on May 1
- INSINC cooperated by slipping on the delivery dates of the links, some worse than others
- The first three routers were installed by the middle of May, Vancouver-Toronto-Montreal. The NSFnet links all slipped badly, but we inherited a circuit from Toronto to Ithaca which began to work towards the end of May
- Machines were sent configured. Installation and cabling of the equipment was done by local site people.
- Debugging on the fly. The first real traffic moved across the network during the second week of June.
- Have now installed eight of the ten routers
- We await the NSFnet links

CA*net Routers and Miscellany

- A reading of the response to RFP would lead one to believe that the routers are not RT's. They are, at least for the moment
- The software is essentially the NSFnet software, collapsed so you can put all the serial cards in one box
- In addition to the single-RT routers, we have additional machines to do traffic measurements
- Routing exchange with our clients is via EGP, or BGP
- The local touchdown arrangement varies from the NSFnet. We supply the ethernet and ask our clients to attach to it. This has the benefit of keeping things simple
- The NSFnet links are handled by a separate machine, to allow them to get faster and to provide a buffer AS between the NSFnet and CA*net

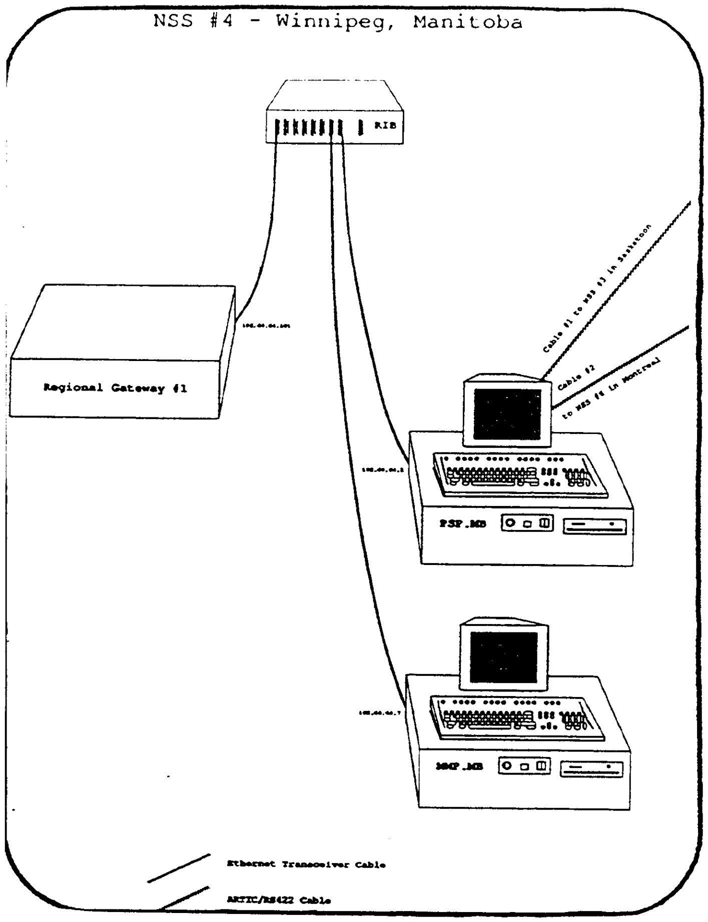

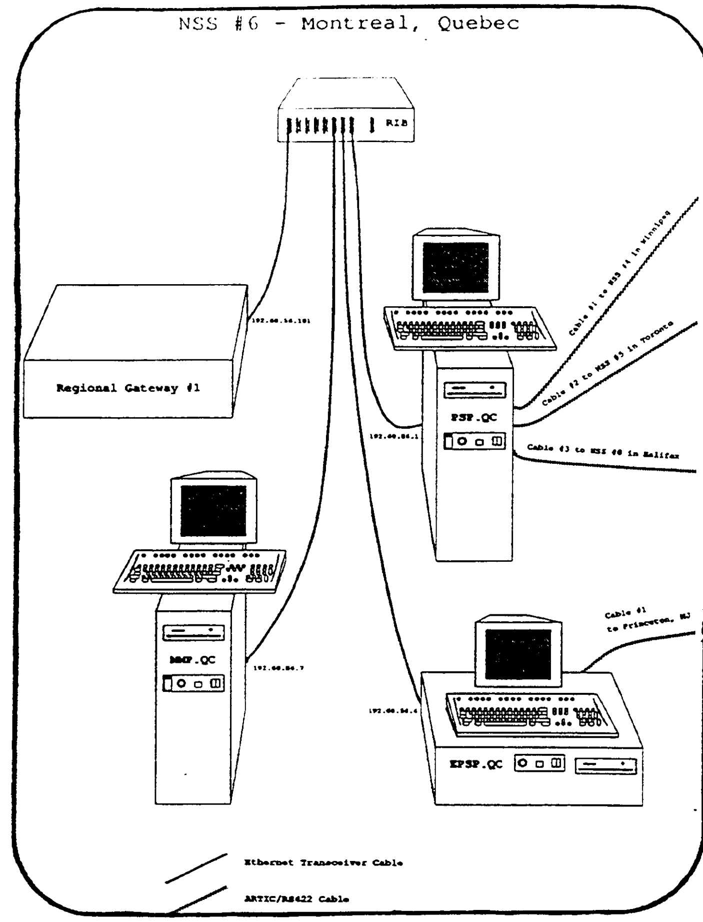

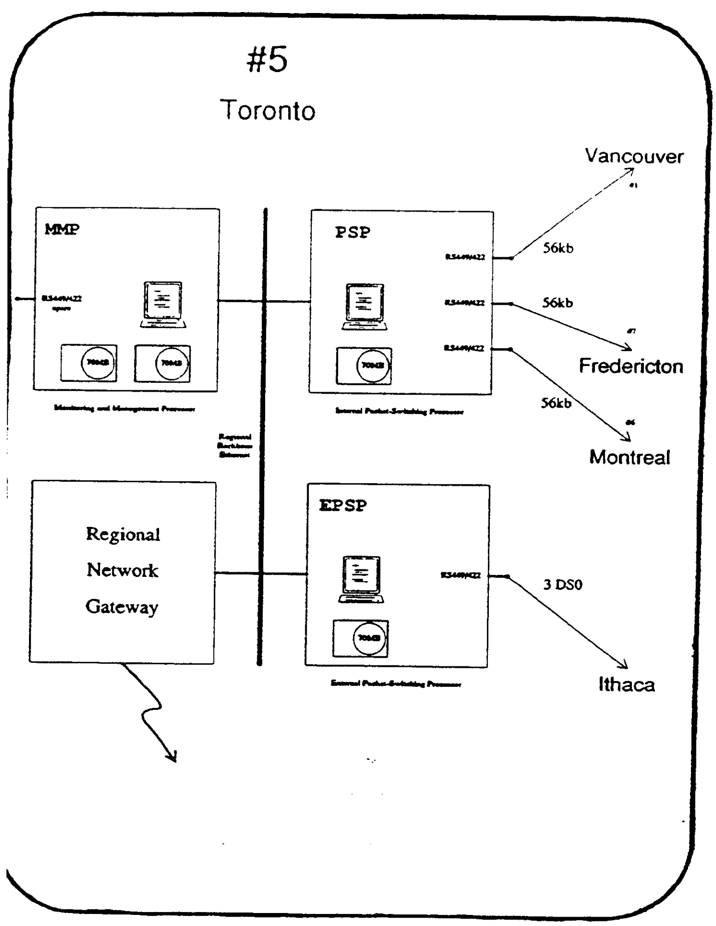

Routing and Routing Policy

- We fully configure routing for our clients, like the NSFnet, with the additional complication of having to also deal with the configuration of the three NSFnet gateways
- Configuration files are generated from an AS/network database, using homegrown software. Regionals tell us what networks they wish to announce, and from where
- Policy is that Canadian traffic stays in Canada whenever possible, while other traffic is routed to the NSFnet via the shortest path available. This policy is enforced by the routers at the locations with NSFnet links
- Networks are added to the NSFnet database with the appropriate link preferences so that routing is usually symmetric
- Fallback routing via the NSFnet is done by the regionals at locations with an NSFnet link. Healing backbone partitions via the NSFnet requires a BGP protocol violation, but may be done

Problems Downscaling to DS0

- The NSFnet software had only ever run in one environment
- The priority queuing done by the kernel was inappropriate at 56 kbps. Changes were made to ensure that at least some user traffic was moved
- The NSFnet IGP includes internal and external links in a single link state update, which is flooded in its entirety whenever something changes. The link state update from the node at the NSFnet gateway was large, and was sent frequently. This was "fixed" by moving the external networks via internal BGP instead
- The kernel was modified to allow some experimentation with congestion avoidance and control strategies

CA*net Futures

- NetNorth traffic. Many (most) of the existing NetNorth circuits are targetted for deletion in favour of VMNET
- Closing the loops to minimize problems resulting from circuit failures
- If the tariffs drop, more bandwidth where justified. Not likely T1, though, except maybe to the US
- Routing coordination. Now things are simple, but this is guaranteed not to be so in the future
- OSI is in your future. Development as necessary

Other Futures

- The ISTC (Industry, Science and Technology Canada, a federal government department) has produced a report recommending the establishment of a much more ambitious network (March, 1990)
- DS1/DS3 bandwidth, with a substantial amount devoted to infrastructure in the regionals
- Rumoured target date is 1991
- Pointedly divorced from CA*net (whew!)

Other Futures

- The ISTC (Industry, Science and Technology Canada, a federal government department) has produced a report recommending the establishment of a much more ambitious network (March, 1990)
- DS1/DS3 bandwidth, with a substantial amount devoted to infrastructure in the regionals
- Rumoured target date is 1991
- Pointedly divorced from CA*net (whew!)

6.3 Engineering the CREN

Presentation by Michael Roberts and Michael Hrybyk/EDUCOM

The Corporation for Research and Educational Networking (CREN) is the result of the 1989 merger of CSNET and BITNET. These two academic networks had their origins in the early Eighties in university computer science departments and university computer centers. In recent years, both have grown and have seen the interests of their users grow much closer together. CREN currently has more than six hundred members. The Corporation is governed by a twelve person board elected by its members. Its services, amounting to approximately $2.6 million a year, are provided by a contract between CREN and EDUCOM. EDUCOM in turn has a contract with BBN for services connected with CSNET. There is no longer any distinction in CREN membership between use of CSNET services and BITNET services. A member may use either or both types of service. CREN is part of the family of NSFNET networks and is connected to NSFNET in Cambridge and San Diego.

CREN has been active in two areas with regard to its BITNET operations. First, BITNET's topology has undergone major changes of late. Second, BITNET has begun to address inter-operability issues, especially relating to IP-based networks.

The topology of BITNET has changed rather radically over the past 18 months. BITNET was initially designed as a strict tree, with each node having only one parent. Routing decisions were simple, since there was only one path from one node to any other. Sites began to install extra lines, disturbing this delicate arrangement, resulting in sub-optimal routes or, worse, the possibility of loops.

Princeton University developed a TCP encapsulation scheme for NJE records, which form the basis of the protocol used throughout BITNET. The Princeton VMNET product for VM/CMS implements this scheme, and other providers (Joiner Associates for VMS, Pennsylvania State University for UNIX) have added support to their offerings. This capability allows two consenting BITNET nodes to establish a virtual connection using the Internet.

The testing phase of VMNET has recently ended, and major sites are now moving traffic from slow, 9600 baud, leased lines to TCP links. A group of BITNET hub nodes are now all directly interconnected using the TCP encapsulation scheme. This has been dubbed the "BITNET Core." The throughput has increased dramatically due to this radical change.

BITNET no longer resembles a convenient tree. Peter Honeyman's pathalias has been adapted for use within BITNET in order to generate routing tables for each node based on assigned link weights. The use of VMNET and newer routing tools has served to reduce the width of the network drastically, and as a result, has increased response times.

Interoperability at the application layer has become increasingly important to CREN/BITNET members. CREN will begin to require that all BITNET nodes make use of RFC822/821 compliant mailers. Vendors (especially IBM) have been petitioned to bundle such a mailer as part of their basic program product. There is also a push to allow more than 8 characters for a node name in NJE records, as well as expanding the range of values allowed.

CREN has supported the integration of BITNET with the Internet community. BITNET can be seen as a set of services (mail, file and information servers, interactive messaging facilities, etc.) provided to member nodes. Those services should be independent of transport, and can be implemented over top of existing IP network facilities. BITNET services can be provided across the Internet, and CREN has begun to facilitate that goal.

CREN has been active on the CSNET side of its operation as well. The west coast cluster is now operational. Plans for integration of dial-up IP services have been drawn. The latter is seen as a low cost alternative for connection to the Internet, especially for members of smaller institutions.

CREN Technical Activities

- Conditions of membership outlined.
- BITNET Topology changes.
- BITNET/Internet Interoperability activities.
- CSNET activities.

1

CREN Membership Terms and Conditions - General

- Follow acceptable use guidelines.
- Abide by technical standards.
- Appoint member representative.
- Pay dues.
- Accept and forward traffic from members.
- Provide up to date member and node information.

2

CREN Membership Terms and Conditions - BITNET

- Guarantee bandwidth of 9600 on principal link.
- Offer at least one more port/connection to another member to establish connectivity.
- Offer 20hr/day, 7 days per week availability of principal connection.
- Install routing tables on a timely basis.
- Strongly recommend use of rfc821/822 compliant mailer.

3

CREN Membership Terms and Conditions - CSNET

- Sign proper software license agreement.
- Have a registered Internet Domain Name.

4

The Old BITNET Topology

- Network constrained to be a tree.
- Tree was rooted at CUNY.
- Only a single path from node to node supported.
- Dependent on order imposed by members.

5

Current BITNET Topology

- Use of the Internet to carry traffic.
- Core sites fully connected using virtual links.
- Use of NJE encapsulation scheme within TCP packets.
- Relaxation of leased line requirement.
- Implementations:
  1. VMNET from Princeton U. for VM/CMS on IBM 370 hardware.
  2. Harvard VMNET driver for UREP and Unix.
  3. Joiner Assoc. VMNET driver for VAX/VMS.

6

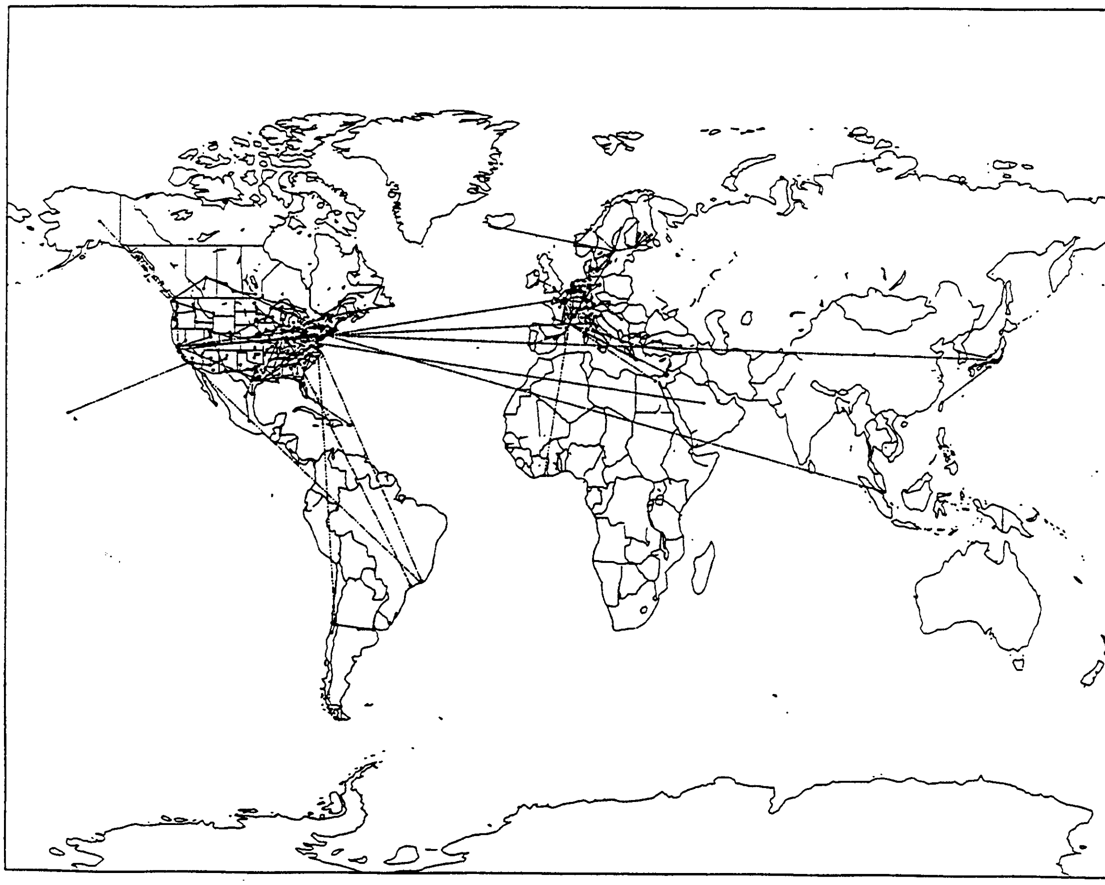

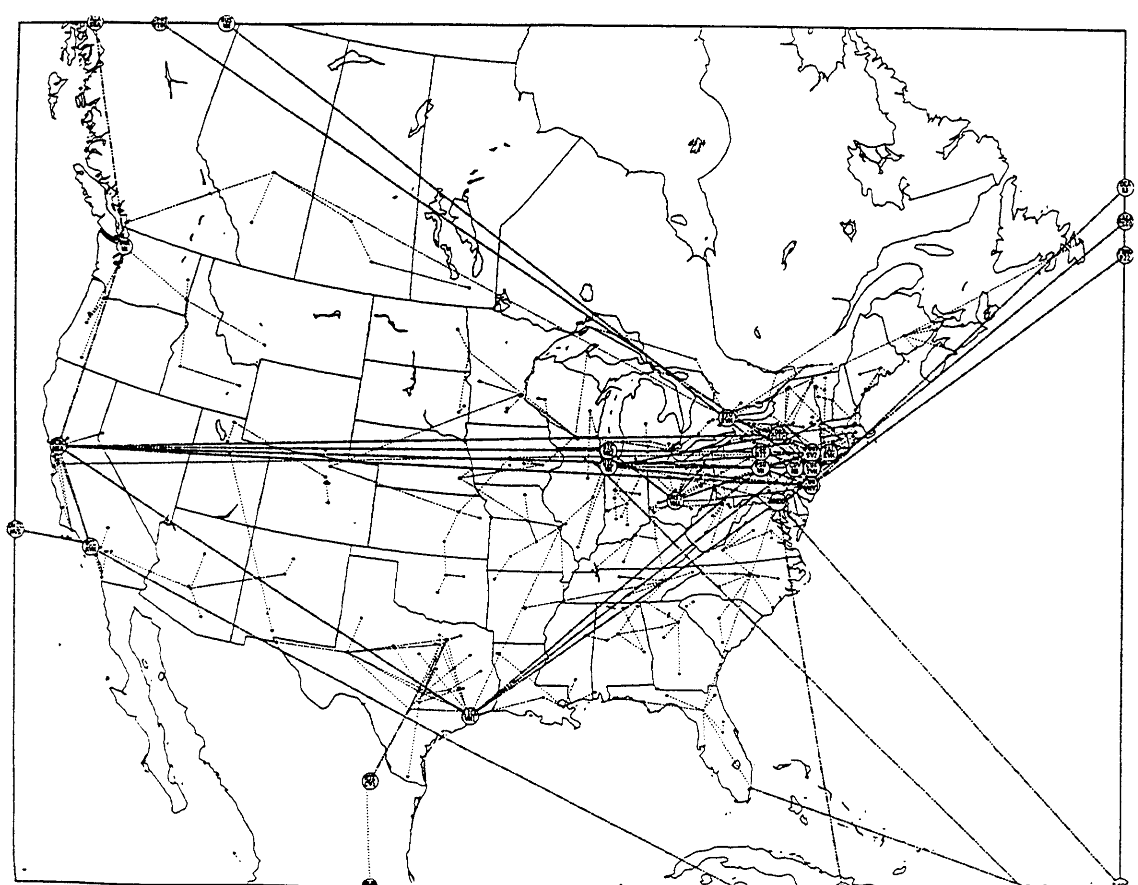

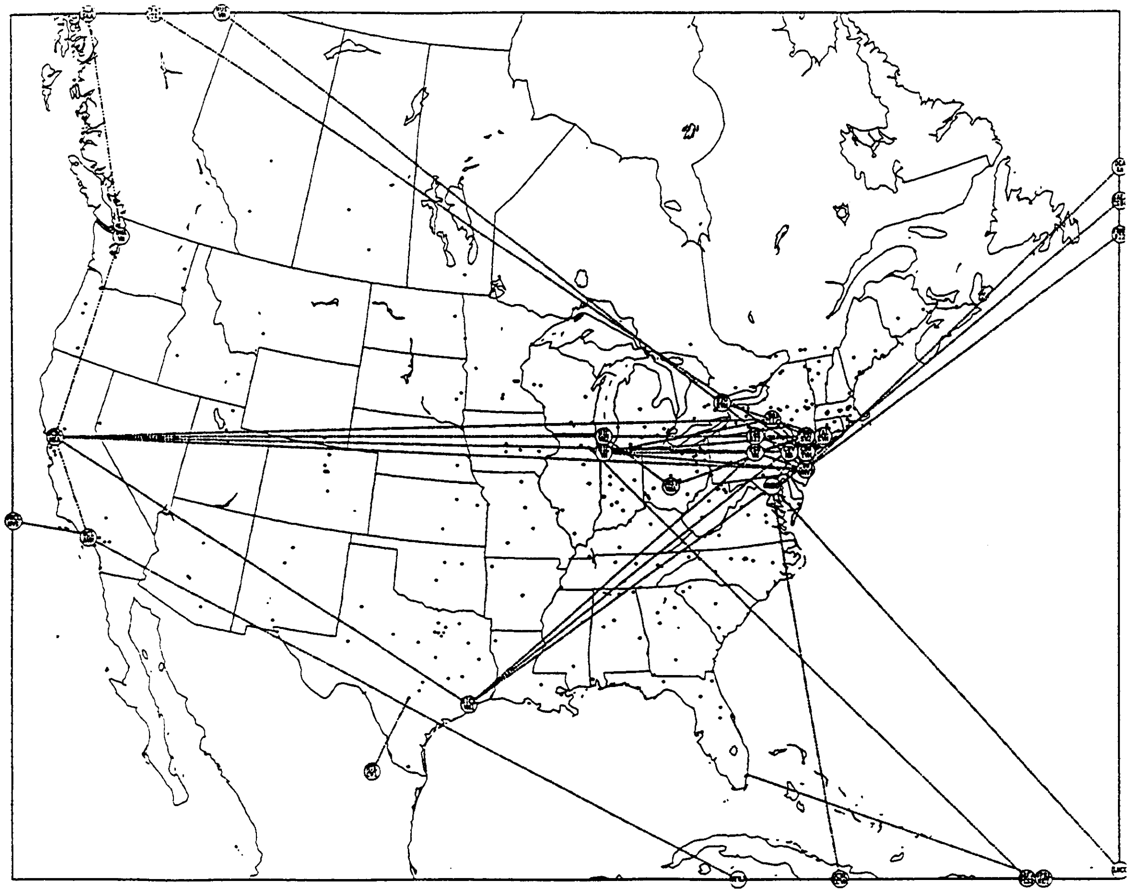

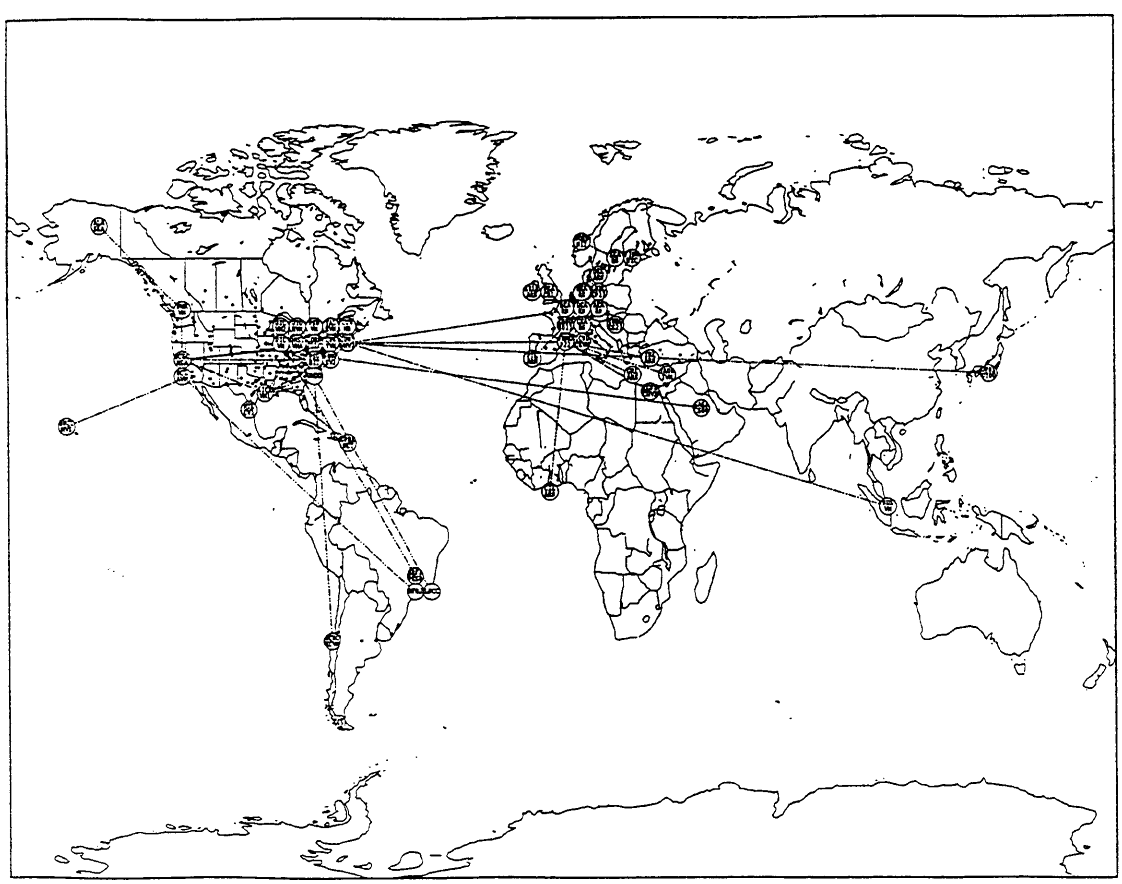

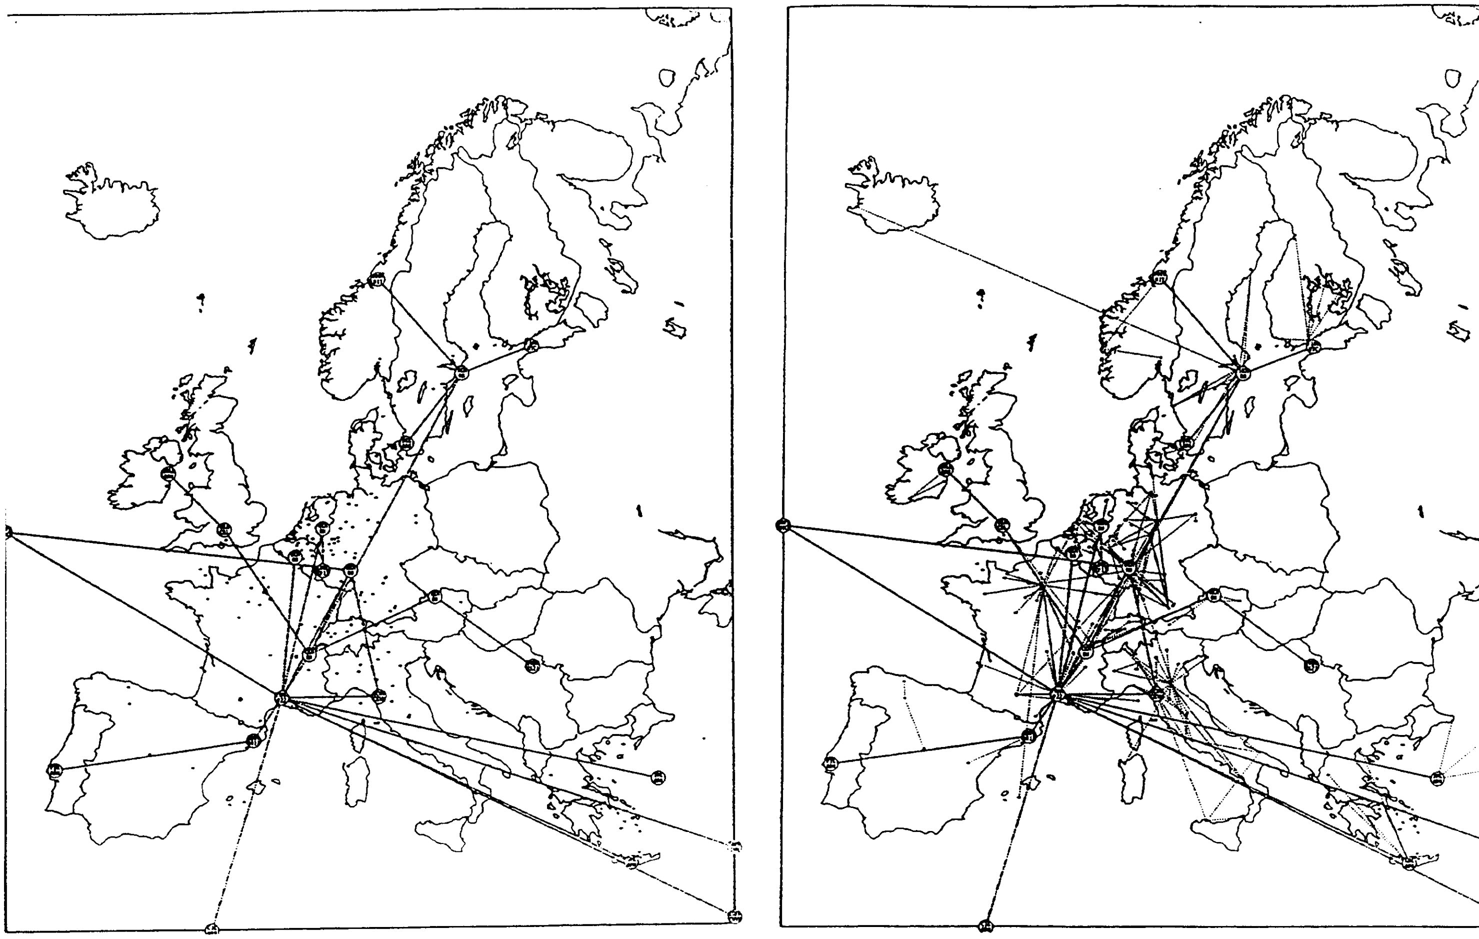

BITNET Routing

- Still uses static routing tables, updated every 30 days.
- Shortest path now calculated (using a modified version of Honeyman's pathalias).
- Link weights assigned by line speed (but can accomodate other variables).
- Use of dynamic routing being investigated.

7

Interoperability

- How to integrate store-and-forward with packet-switched networks?
- Maturation of s-and-f services has occured:
  - Sender-initiated file transfer.
  - Interactive messaging systems.
  - Mailing list and bulletin board services.
- Use of protocol encapsulation to preserve applications.

8

BITNET Interoperability Goals

- Use of rfc821/22 compliant mailers.
- Domain registration for all BITNET members nodes (and use of appropriate MX records).
- Use of a native RSCS IP line driver for VM systems.
- Lippke's FRED project.
- Message routing problems.

9

CSNET Developments

- West Coast Cluster.
- Plan for use of dial-up IP service.

10

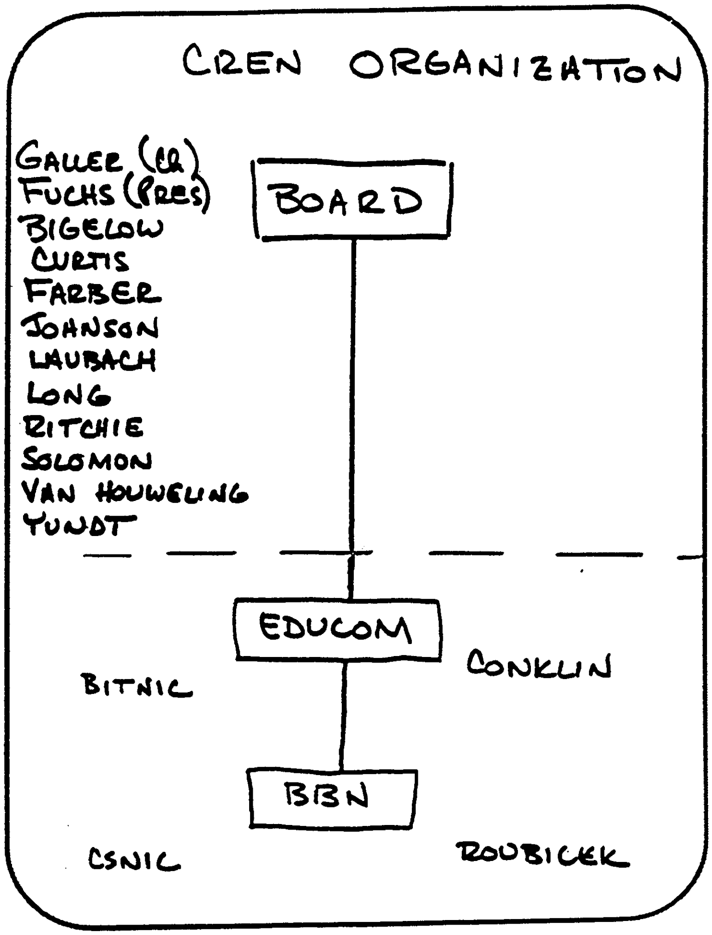

INTERNET / NSFNET / CREN

CONNECTIVITY IN U.S. HIGHER ED.

TODAY:	≅ 400	NSFNET
	≅ 500	CREN
	100	OTHER ?
TARGET:	100	RESEARCH UNIV
	300	OTHER UNIV
	1000	OTHER 4 YR
	1000	2 YR
	1000	OTHER SPECIALIZED
	3400 ±	

1989 FRICC NREN PROGRAM PLAN:

"1300 BY 1995"

≅ 6 MILLION STUDENTS
1 MILLION FACULTY + RESCH STAFF

6.4 Perspectives on Research Networks in Europe

Presentation by Eric Huizer and Rudiger Volk / Rare-Ripe

PERSPECTIVES ON RESEARCH NETWORKS IN EUROPE

IETF
Vancouver 2 - 8 - 1990

Erik Huizer
SURFnet, The Netherlands

huizer@SURFnet.nl

Multinational Networks

HEPnet
SPAN
EUnet
EARN
NORDUnet
EASInet
etc.

National Networks / Networking Organisations

Janet (uk, gb)
DFN (de)
SURFnet (nl)
SWITCH (ch)
GARR (it)
etc.

Regional Networks

RARE
National members (e.g. SURFnet)
International members (e.g. EUnet)
Liason members (e.g. EWOS)
Associate members (e.g. Korea)

some coordination bodies:
EURO-CCIRN
EEPG
IXI-CC
RIPE
WG's

Task-Forces / Working Groups

Ripe task-forces (mostly IP)
Rare WG's (mostly OSI)

RARE Working groups

WG1:	MHS	X.400, RFC-987/1148
WG2:	File Transfer	FTAM, FTAM-FTP GWY
WG3:	Directory Services	X.500, wpp
	Information Services	EIS
	User Services	User support
WG4:	Transport services	TP0/TP4, CONS/CLNS, NSAP
WG5:	Virtual Terminal	VT, X11
WG6:	High Speed	
WG8:	Network Management	SNMP, CMIP
	Security	

COSINE

Financing from combined European governments for creation of a European network infrastructure for research.

Without interfering with autonomy of existing networks!

IXI

X.25 (1984) 64kb/s
Free till mid 1991

What then???
Too expensive?
Too slow?
Single protocol?

Or
developing into an affordable multi-protocol backbone?

C***OSI***NE

Pan-European Backbone

Multi-protocol
High speed
"Fixes" for intercontinental links

Don't forget Eastern Europe

The European Backbone

Inventory:
Maps (RIPE, HEPnet, EEPG)
Usage inventory (EEPG, Rare WG6)

Intermediate solutions:
Line sharing, multi-lateral agreements
e.g. :
Stockholm
|
Amsterdam
|
Geneva (CERN)
|
Bologna
|
Rome

Structural solutions

Intercontinental lines

Coordination
CCIRN

Technical Advice to RARE
EEPG & RIPE

1 "fat pipe" (>128kb/s) in place
Cornell - CERN (EASInet)

2nd f.p. needed for
load-balancing
back-up

ORGANISATION EUROPEENNE POUR LA RECHERCHE NUCLEAIRE
EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH
Laboratoire Européen pour la Physique des Particules
European Laboratory for Particle Physics

Networking MAPs

To understand the world around us, we maintain a series of MAPS of networks to which our organization is connected.

They are prepared for CERN internal use from available information, and may contain errors or omissions.

Those interested may obtain copies from, or mention corrections to:
F.G.de Bilio, CERN CN/CS, e-mail:bilio@cernvm.cern.ch

F.Fluckiger
CERN/CN/CS
June 1990

MAIN International Academic and Research Leased Lines in EUROPE (*)

(*): - incl. ordered lines and agreed groupings
- CERN considered as a country

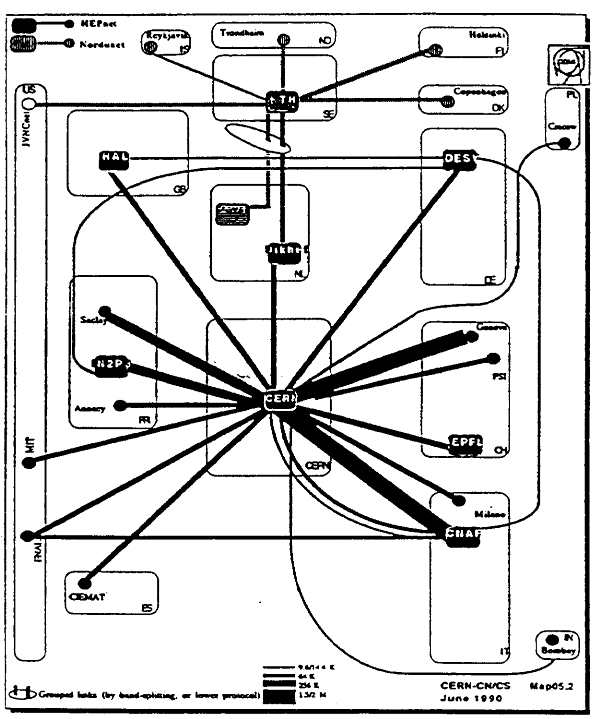

(Prepared for CERN internal use from available information. Can be obtained from F.G.de Bilio CERN-CN/CS)

MAIN International Academic and Research Leased Lines in EUROPE (*)

(*): - incl. ordered lines and agreed groupings
- CERN considered as a country

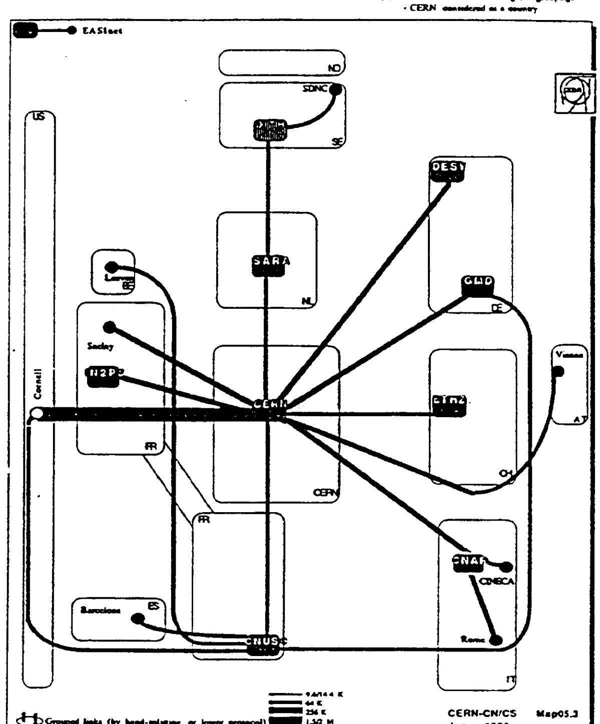

(Prepared for CERN internal use from available information. Can be obtained from F.G.de Bilio CERN-CN/CS)

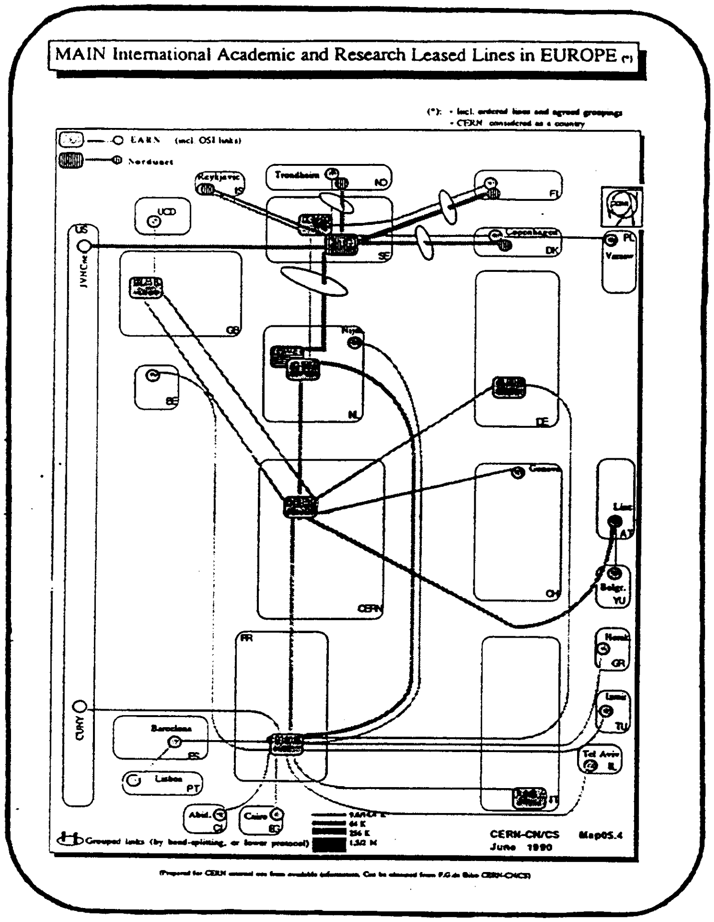
MAIN International Academic and Research Leased Lines in EUROPE (*)
(*): - incl. ordered lines and agreed groupings
- CERN considered as a country
EARN (incl. OSI links)
Nordunet
Reykjavik IS
Trondheim NO
FI
UCD
SE
Copenhagen DK
PL Varsaw
US
JVNCnet
GB
Nijm.
NL
BE
DE
Geneva
CERN
Linz AT
CH
Belgr. YU
FR
Nemk. GR
CUNY
Barcelona
ES
Lisbon PT
Izmir TU
Tel Aviv IL
IT
Abid. CI
Cairo EG
9.6/14.4 K
64 K
256 K
1.5/2 M
Grouped links (by band-splitting, or lower protocol)
CERN-CN/CS
June 1990
Map05.4

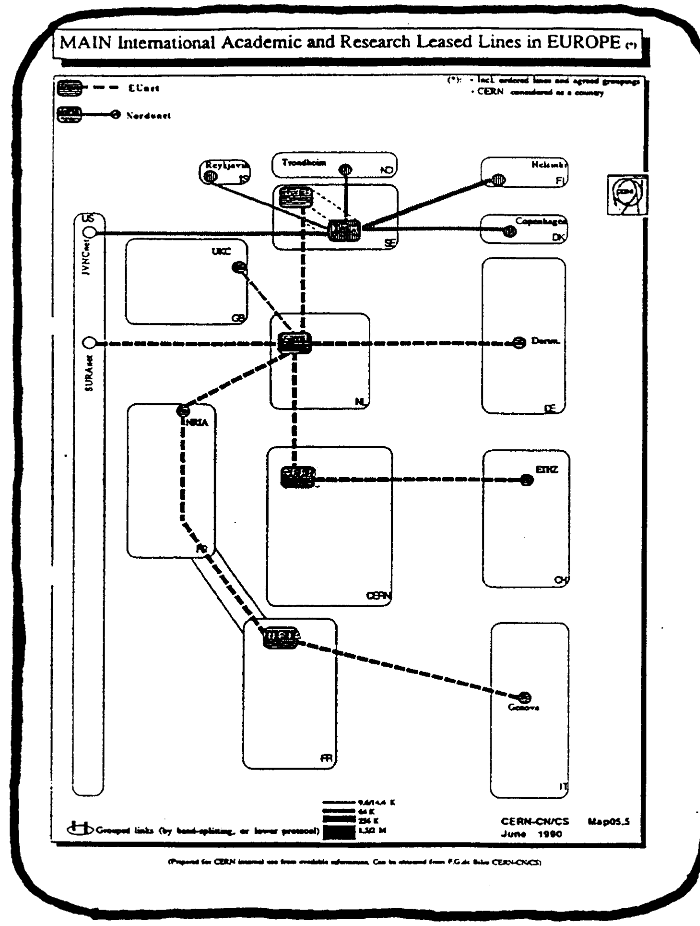
MAIN International Academic and Research Leased Lines in EUROPE (*)
(*): - incl. ordered lines and agreed groupings
- CERN considered as a country
EUnet
Nordunet
Reykjavik IS
Trondheim NO
Helsinki FI
SE
Copenhagen DK
US
JVNCnet
UKC
GB
SURAnet
NL
Dortm.
DE
INRIA
ETHZ
FR
CERN
CH
Genova
FR
IT
9.6/14.4 K
64 K
256 K
1.5/2 M
Grouped links (by band-splitting, or lower protocol)
CERN-CN/CS
June 1990
Map05.5

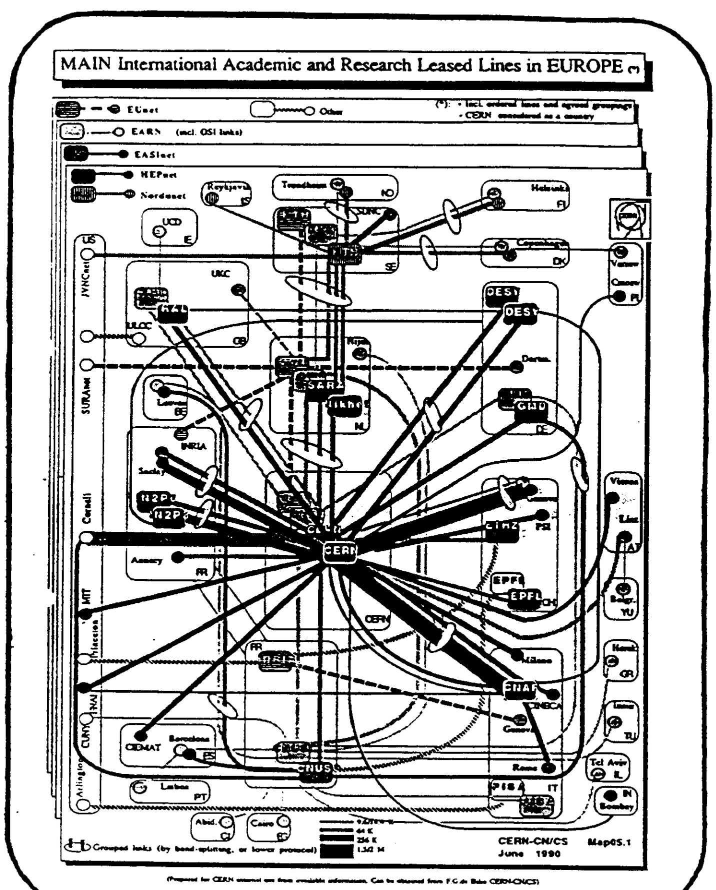
MAIN International Academic and Research Leased Lines in EUROPE (*)
(*): - incl. ordered lines and agreed groupings
- CERN considered as a country
EUnet
Other
EARN (incl. OSI links)
EASInet
HEPnet
Nordunet
Reykjavik IS
Trondheim NO
Helsinki FI
UCD IE
Copenhagen DK
SE
US
JVNCnet
UKC
RAL
ULCC
GB
DESY
Dortm.
SURAnet
SARA
Nikhef
NL
GMD
DE
BE
INRIA
Saclay
IN2P3
Cornell
Amsey
FR
CERN
ETHZ
PSI
EPFL
CH
MIT
Princeton
FNAL
CUNY
Arlington
CIEMAT
Barcelona
ES
Lisbon PT
CNUS
CNAF
CINECA
Genova
Roma
PISA
IT
Milano
Vienna
Linz AT
Belgr. YU
Nemk. GR
Izmir TU
Tel Aviv IL
IN Bombay
Abid. CI
Cairo EG
9.6/14.4 K
64 K
256 K
1.5/2 M
Grouped links (by band-splitting, or lower protocol)
CERN-CN/CS
June 1990
Map05.1

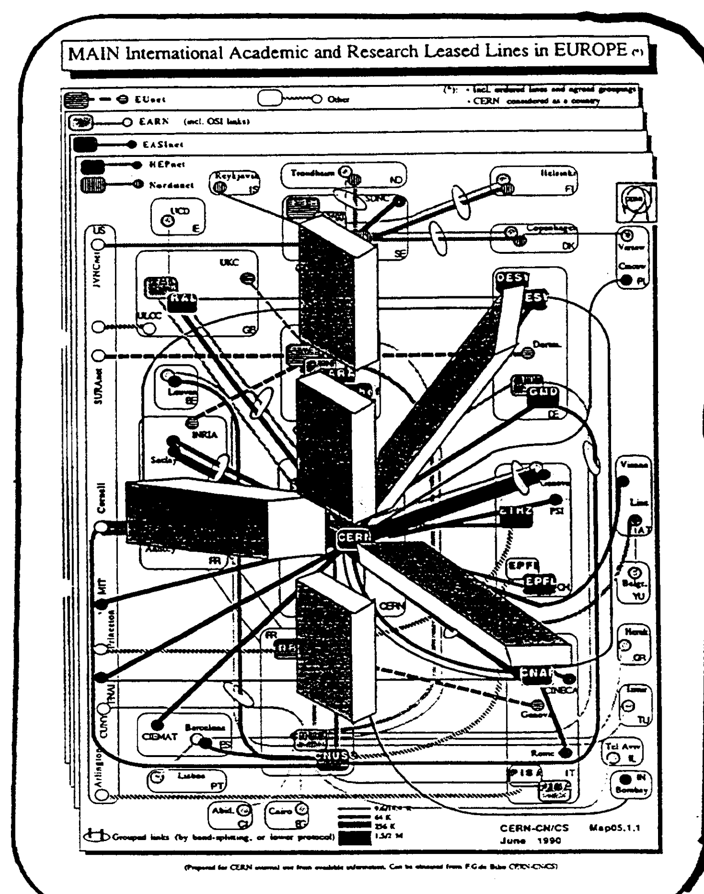
MAIN International Academic and Research Leased Lines in EUROPE (*)
(*): - incl. ordered lines and agreed groupings
- CERN considered as a country
EUnet
Other
EARN (incl. OSI links)
EASInet
HEPnet
Nordunet
Reykjavik IS
Trondheim NO
Helsinki FI
UCD E
Copenhagen DK
SE
US
JVNCnet
UKC
RAL
ULCC
GB
DESY
Dortm.
SURAnet
GMD
DE
BE
INRIA
Saclay
Cornell
FR
CERN
ETHZ
PSI
EPFL
CH
MIT
Princeton
FNAL
CUNY
Arlington
CIEMAT
Barcelona
ES
Lisbon PT
CNUS
CNAF
CINECA
Genova
Roma
PISA
IT
Vienna
Linz AT
Belgr. YU
GR
TU
Tel Aviv IL
IN Bombay
Abid. CI
Cairo EG
9.6/14.4 K
64 K
256 K
1.5/2 M
Grouped links (by band-splitting, or lower protocol)
CERN-CN/CS
June 1990
Map05.1.1

Main International Leased Lines using IP in Europe (incl. ordered lines)

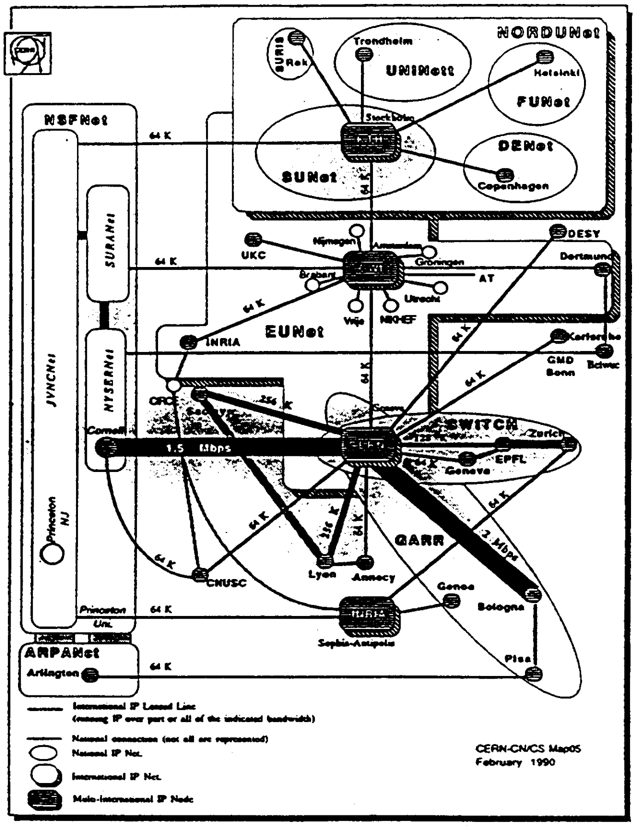

(Prepared for CERN internal use from available information. Can be obtained from F.G. de Bilio CERN-CN/CS)

6.5 Berkeley TCP Evolution from 4.3-Tahoe to 4.3-Reno

Presentation by Van Jacobson/LBL

There is now a new mbuf structure in 4.3-reno requiring a new code be written following a sockets-streams model called b-streams. A new driver calling sequence is partly in the way of the new b-streams architecture. They want to put a lot more caching of information into 4.4, which requires both letting the drivers have a place to cache and sending the drivers a convenient hole as a root. For right now we have added a new parameter to all driver calls (which are roots) and pretty soon some of the old parameters will go away. An example would be the current interface pointer which happens to be redundant because you can find it by looking inside the root. If you have written drivers for Berkeley Unix, you might consider rewriting them so they will work under 4.4.

Provision has been made for having higher level protocols lay out data in the right way to go out on a wire. This should leave space for lower level protocol headers including link layer headers so as drivers are coming up, they leave information laying around that higher level protocols can use (like what is the largest size link header this driver might require) and you can use that up at the socket level to leave a hole to put that link header. We are doing this to get around the model that was in 4.3 and 4.2 where all network interfaces were considered to be ethernets. It may be things with larger or smaller than 14 byte link headers.

There is a brand new routing algorithm which approximates Patricia, done by Keith Scaller (sp?). As a functional equivalent, it uses a radix algorithm but carries a per-entry mask and carries the same hierarchy the Patricia algorithm allows. You can have subnets carry their own private masks and different width masks on different networks. Routes now carry full path characteristics inside the route, including time to live, pipe size, the socket buffer to be used, the mtu, minimum rtt and every other thing we could possible think to throw into it. We are trying to teach transport protocols to utilize this information. Right now tcp is the main thing that makes use of this information.

I use that path characteristic all the time. It is really nice for slip links because they are wierd and have a tiny bandwidth delay product. Ideally, you would like them to have a small mtu so there is not a lot of interference between interactive and background traffic like ftp's. If you have a slip link at home, all you need to do to set it up is to add a default route to your slip link and add these new parameters, such as mtu = 296, 1k pipesize, rtt ¿ 2sec. Put that entry in the slip line start up, and from then on, every TCP uses the correct parameters for that slip link.

TCP changes fall into two categories, 1) speed ups, alot of which is the header prediction work, though not all, because alot of it relys on the not-yet available b-streams changes. Easy things were 1) TCP ports coming in had to be turned into your local datastructures, the state structure of TCP. It turns out that if you cache just the last PCB lookup, and do a quick compare to see if these ports match the last lookup, which is a real cheap 96 byte compare, you get about a 90% hit rate on any host. We added a 1 BAT (?) cache which saves thousands of instructions with of an IMPC (?) lookup. There was an unfortunate need to do D to M memory to memory copies to generate occasional error messages. These copies under m-mpullups were a major performance loss. Most of the code has been changed so they are no longer there. You should almost never see 4.3 reno doing an m-mpullup or memory to memory copy when a packet comes in. All the data stays in a buffer until you get it to either a user of NFS or it's final destination.

The Protocol is smart on the way out, in that it builds packets that look just as they will going out on the wire. The chunks are copies in from the user in 1MSS size units and you leave enough space in the front for the TCP, IP, and link level header so all you have to do is stick the header on the front of the chunk and blast it out on the wire. You never need to repacketize data. There is at most one memory to memory copy. In some virtual memory schemes, such as the one in 4.4, there may not even be one memory to memory copy. The one copy is feeding the data to the higher level applications socket.

About 1.2 of the header prediction code is in 4.3 reno. The input side is in the release, but the similar code for the outbound side is not yet working. You keep a few bits around that summarize some of the past actions like whether an ack packet showed up, or whether two ack or whether the user used the data as it arrived. There are five bits of data available, you can switch on the bits, in some cases determining which packets were generated, and in others determining the packet by carefully inspecting your state.

Something modivated by some test results on the NSFnet backbone sending packets near the bandwidth delay product (30-40 kbytes), the silly window code began sending fragments. Particularly when copying data in from the user, you're laying it down in units that look just like the packets that are being layed down on the wire. So, you are copying data from the user in 460 byte chunks so the silly window code sends 3 bytes of that chunk, and now you are in exactly the wrong place in every chunk in the buffer. Rather than getting the best possible performance, you get the worst possible performance. This somewhat mitigated the advantages in the new code. Make sure you never get out of line in the chunk with the buffers set up for you. In the process, the data looks a lot better on the wire. You should never see 4.3 reno laying little bitty packets.

There are a whole set of changes for low speed links. The route path characteristics are a real win for SLIP. They allow you to correctly configure the link. One nice thing about TCP is that if you set one side right, the other side is automatically configured correctly. The max segment size gets minimized between the two offers. SLIP nodes is usually a client leaf node, setting up default characteristics is all you ever need to do. There were many things that were giving spurious retransmitts. Because header compression gives up on retransmitted packets, you really notice a retransmit over a 2400 baud line. Everytime one one of those happened I dug into the networking code and figured out why it happened and fixed it.

The major problem turned out to be tied to the slow start code. When slow starting, we start out with a real small window, and as acks are received, we open up the window. You will be sitting there in a telnet session with a one packet window. As you login, the stram of characters fully opens the transmit window. If you then VI a file, you dump 2K worth of characters on the line to repaint the screen. The slow start is now a no-op. The RTT suddenly goes up by factor of 2000. The rtt counter does not deal with a factor of 2000. The idea behind slow start is to get this ack clock going. You are trying to get told when to drop a packet into the line. The problem is that the connection has gone idle, and there is not much data, a sudden lump of data will not have a string of acks returning. After 1 RTT there are no more acks returning. If there were any acks, they would return in 1 RTT. To solve this problem we put in a simple test: If the connection has been idle for more that 1 RTT, then slow start. This got rid of most spurious retransmits. This did more than help interactive traffic. This also helps SMTP and NNTP which have short dialogue sessions which tend to open the window followed by a large blast of data which tends to overwhelm the line.

There was another set of changes for the case where the line is truly bandwidth limited, so as the packet size changes, the RTT changes a lot. The RTT estimator was too heavily filtered to be effective. This is a case where the estimator was just too aggressive. This was the old algorithm ner RTT=1RTT+2*est(variance). We changed the formula to new RTT=1RTT+4*est(variance) and almost all the spurious retransmitts disappeared, including all the bandwith related retransmits. There were one or two left which were solved by punting. We said, this line has got 1k of buffer and is a 9600 baud. We can see excursions in RTT up to 1 second as this buffer fills and emoties, so don't let the RTT go under a second. So, we put in a minimum path RTT of 1 second.

The enhancements work better over fast links. Because we have per path characteristics, we can use the right mtu, not just 576 bytes for each path when the NSFnet can handle 1500 bytes. We can now use the right pipe size. You are not constrained to using 4K worth of buffer when all you need is 3 packets worth of buffer on an ethernet, and you really need 40 packets worth of buffer on the NSFnet.

With the new timeout when you are feeding into a long fat pipe, a pipe requiring a 40k window, the transients from cross traffic if everyone else is using 40k windows, can be pretty extraordinary and with the new timer, where we put more weight on the variance, you are much less likely to do a false retransmit. I did a bunch of NSFnet throughout tests, and throughout those tests, I did not see any spurious retransmits with the new code. If you are running an nntp for example, it has the property where it sends these short control messages and then blasts out a large article or lump of data and that would full the heck out of slow start. The control messages open the window, and the blast drives the timers through the roof. The slow start on idle code has solved this problem for fast links as well. Retransmits have gone down to essentially zero.

Tutorial ——

Berkeley TCP Evolution from 4.3-tahoe to 4.3-reno

Van Jacobson
Lawrence Berkeley Laboratory

18th Internet Engineering Task Force meeting
University of British Columbia
Vancouver, BC, Canada
July 31–August 3, 1990

General network changes in 4.3-reno

- New mbuf structure.
- New driver calling sequence.
- Drivers tell higher level protocols how much space needed for headers (max_linkhdr).
- New routing algorithm (≈ PATRICIA — carries per-entry mask).
- Routes carry 'path characteristics': TTL, MTU, pipesize, minimum RTT, estimated RTT, estimated RTT var., pipelimit, 'keep-alive' enable, etc.

General network changes in 4.3-reno (cont.)

(Path characteristics very nice for SLIP links. E.g., at client end of link just do

```
route add default slip-link 1 mtu 296 pipesize 1024 minrtt 2
```

and all connections will use good window, mtu and timer values.)

TCP changes in 4.3-reno

It's faster:

- Last PCB lookup cached.
- No m_pullup's.
- Space left for IP and link-level headers.
- ≈ 50% of Header Prediction added.
- Sender silly-window code more effective.
- Socket buffers rounded up to integral multiple of MSS.

TCP changes in 4.3-reno (cont.)

It works better over slow links:

- Route path characteristics allow individual connection configuration (smaller MTU and pipesize).
- Slowstarts if connection idle for $\geq$ 1 RTT. (Prevents most spurious retransmissions on screen repaints, SMTP, NNTP, etc).
- Retransmit timeout changed from $rtt + 2 * var$ to $rtt + 4 * var$ (catches most remaining spurious retransmits).
- Per-path minimum RTT (prevents rest of spurious retransmits).

TCP changes in 4.3-reno (cont.)

It works better over long, high-speed links:

- Route path characteristics allow individual connection configuration (bigger MTU and pipesize).
- New timeout better for loaded, fat pipes.
- Slowstart-on-idle much better for interactive and NNTP traffic.
- New "Fast Recovery" algorithm.

NSFNet tests — trace startup detail

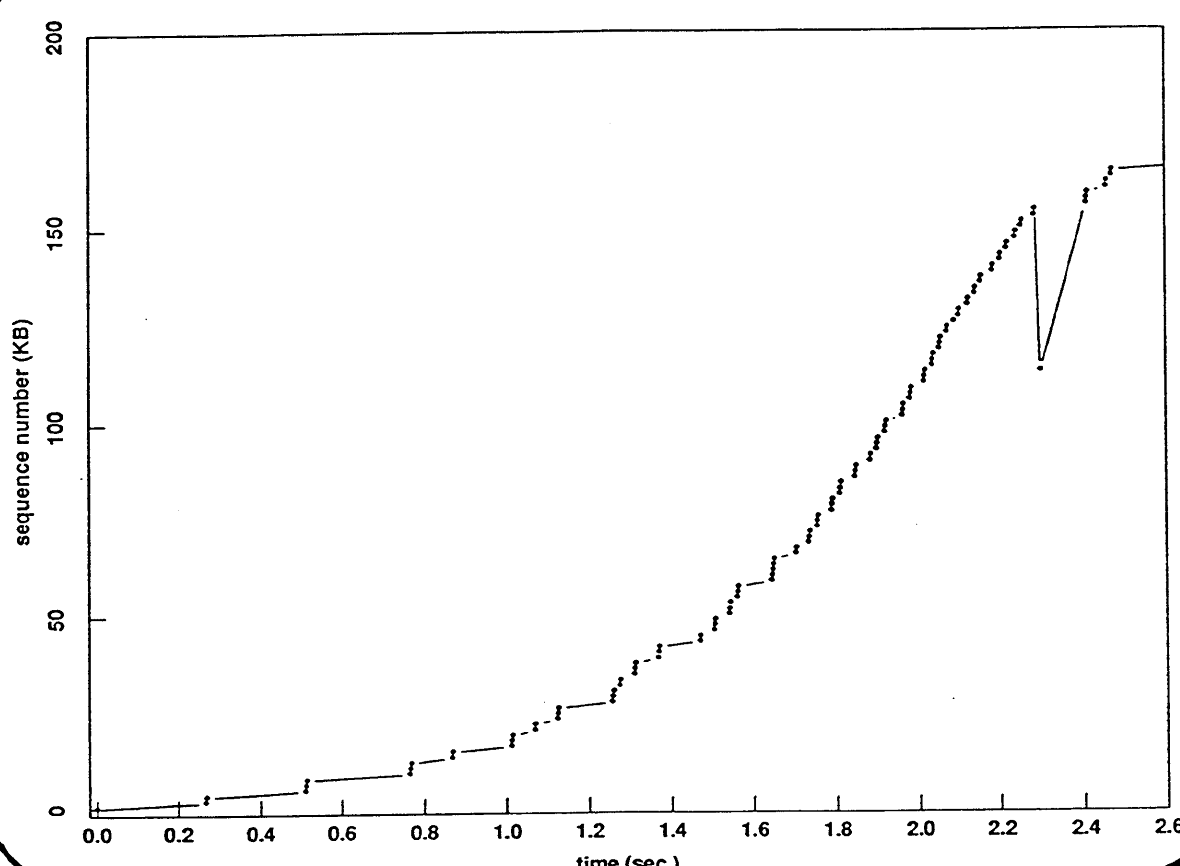

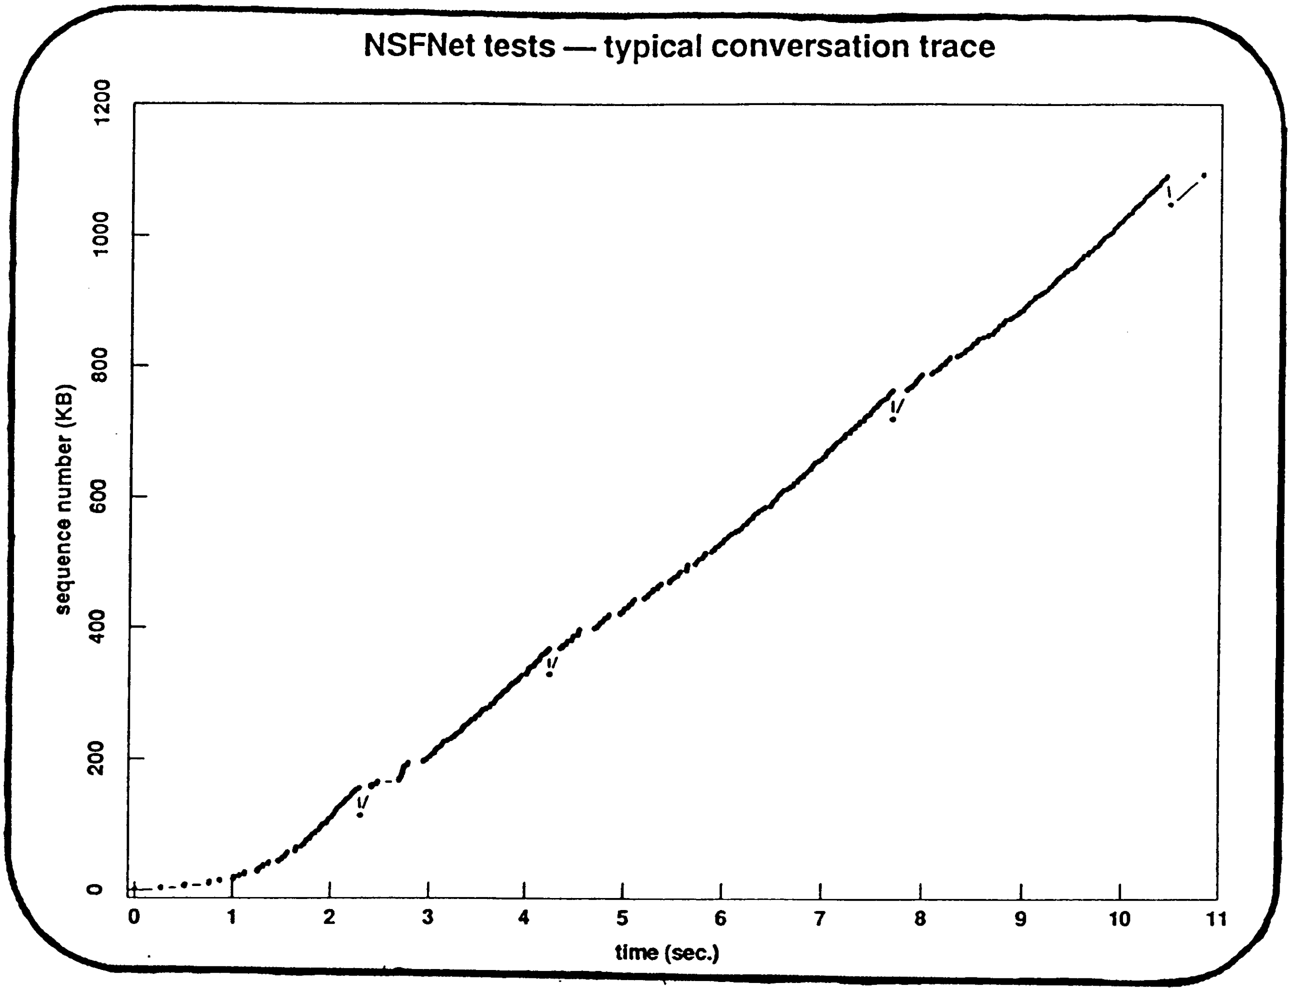
NSFNet tests — typical conversation trace
sequence number (KB)
time (sec.)
0
200
400
600
800
1000
1200
0
1
2
3
4
5
6
7
8
9
10
11

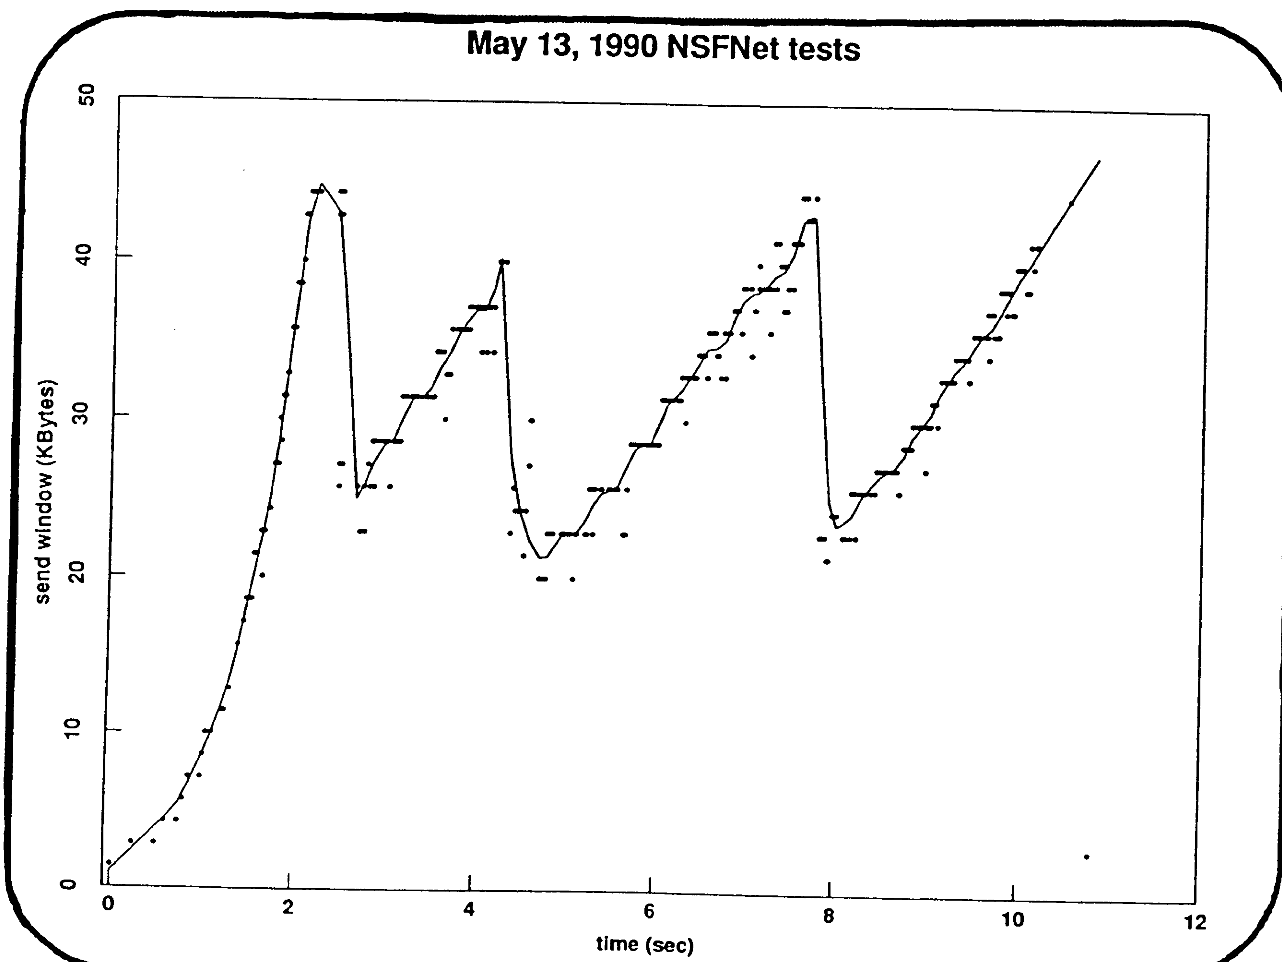
May 13, 1990 NSFNet tests
send window (KBytes)
time (sec)
0
10
20
30
40
50
0
2
4
6
8
10
12

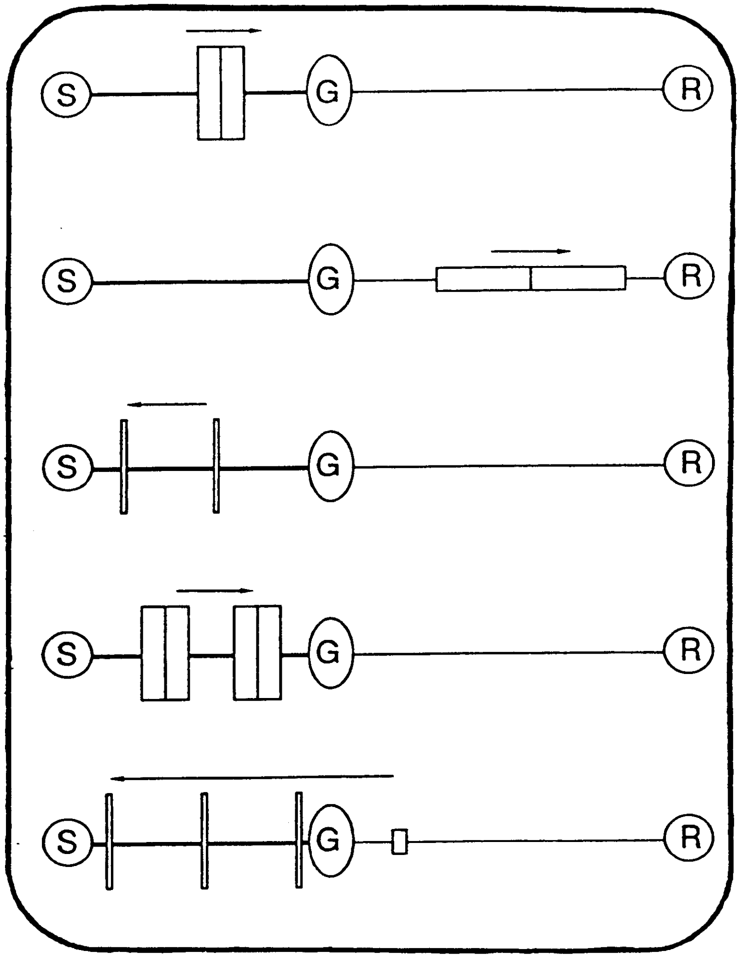
S
G
R
S
G
R
S
G
R
S
G
R
S
G
R

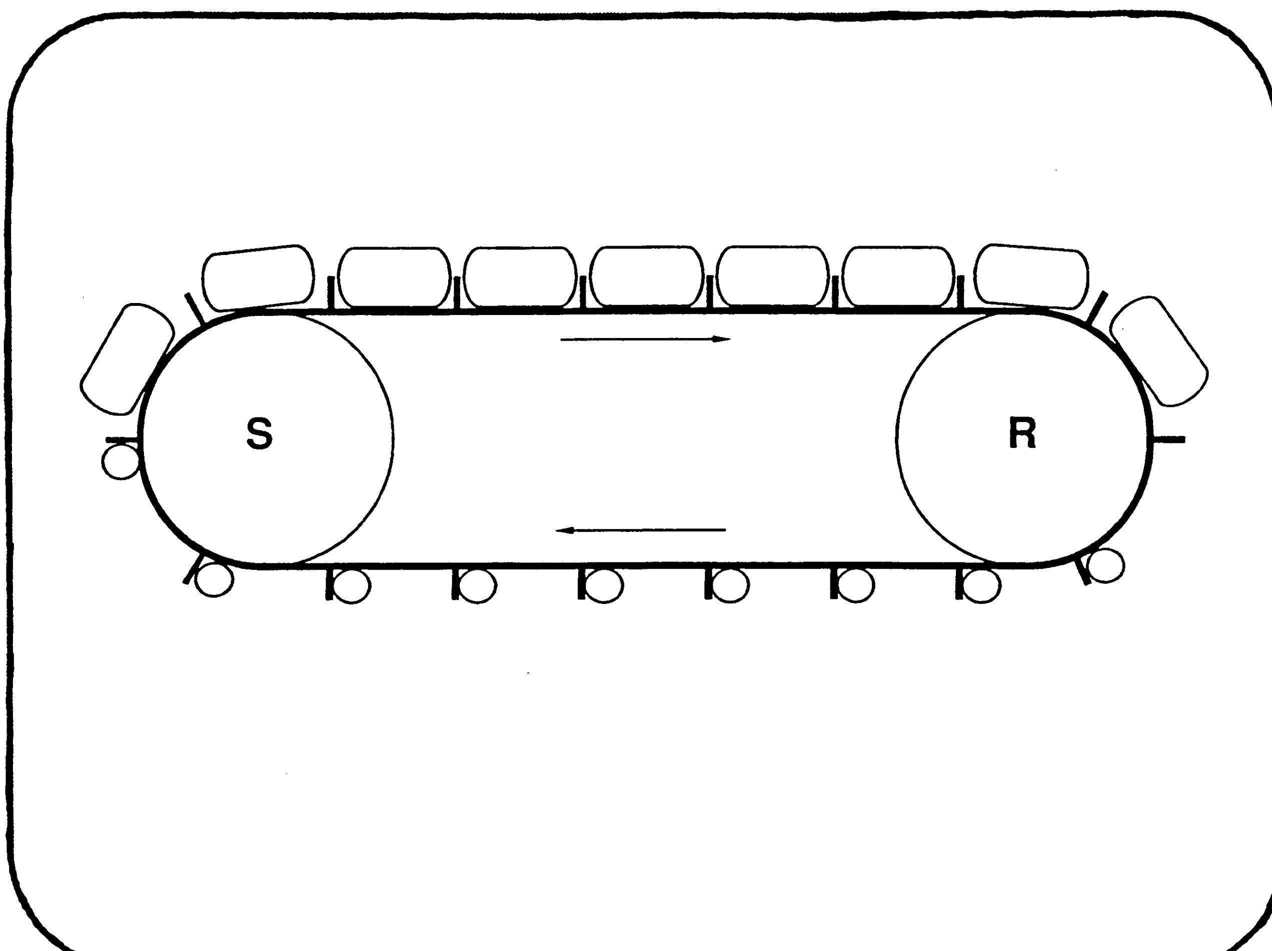
S
R

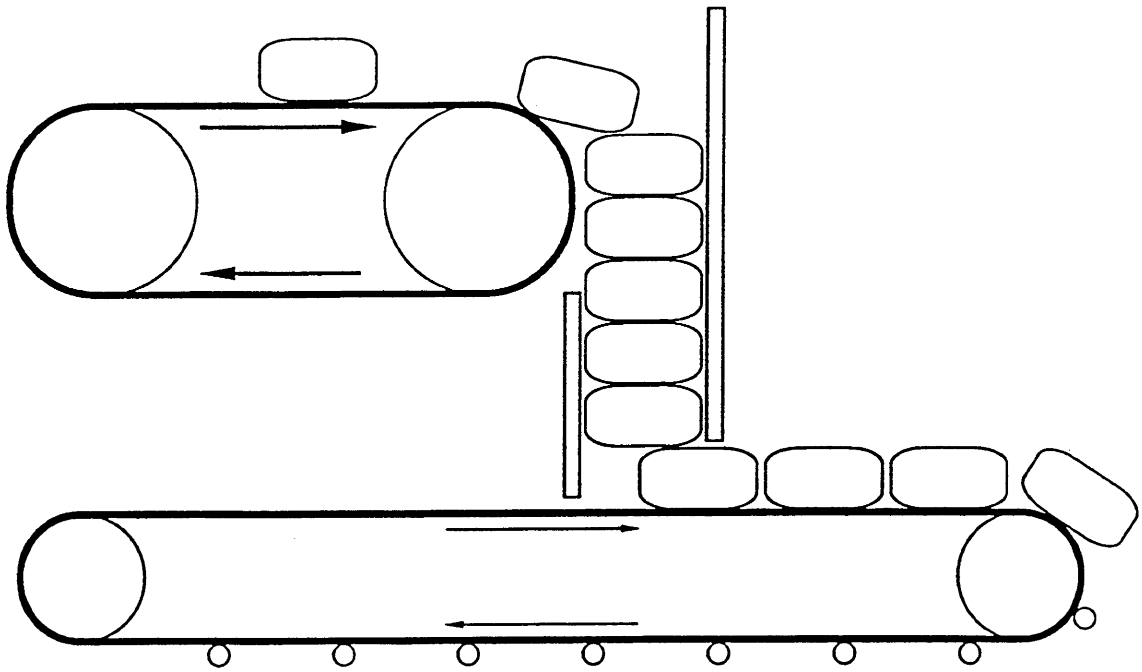

Time x — Packet damaged or dropped.

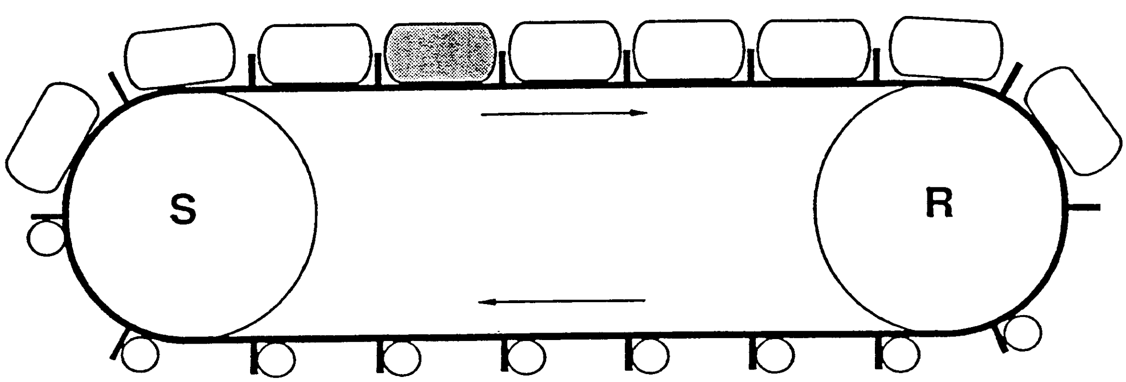

Time 0^- — Damaged / missing packet arrives at receiver.

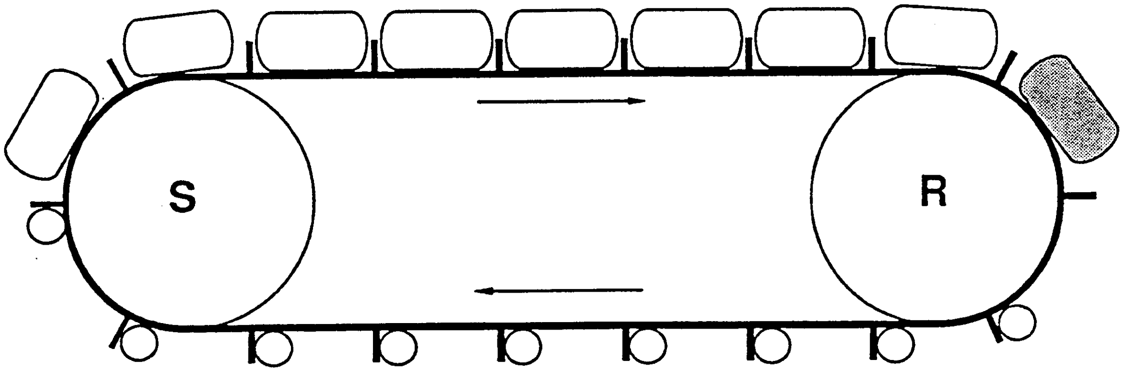

Time 0^+ — Hole in seq. space – send duplicate ack.

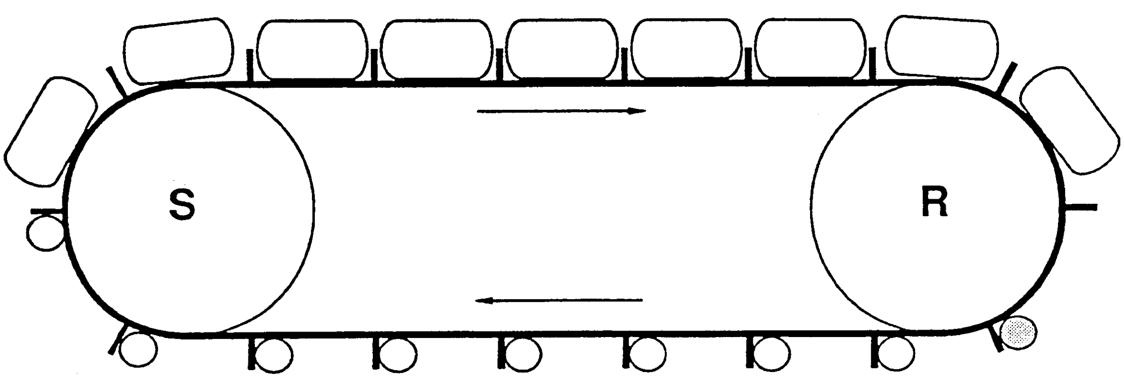

Time 0.5 R⁻ — First duplicate ack arrives.

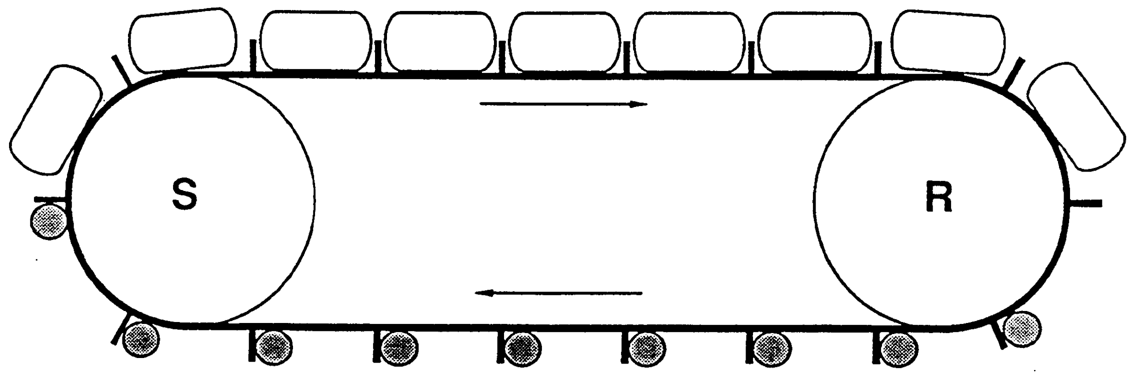

Time 0.5 R⁺ — Missing packet retransmitted.

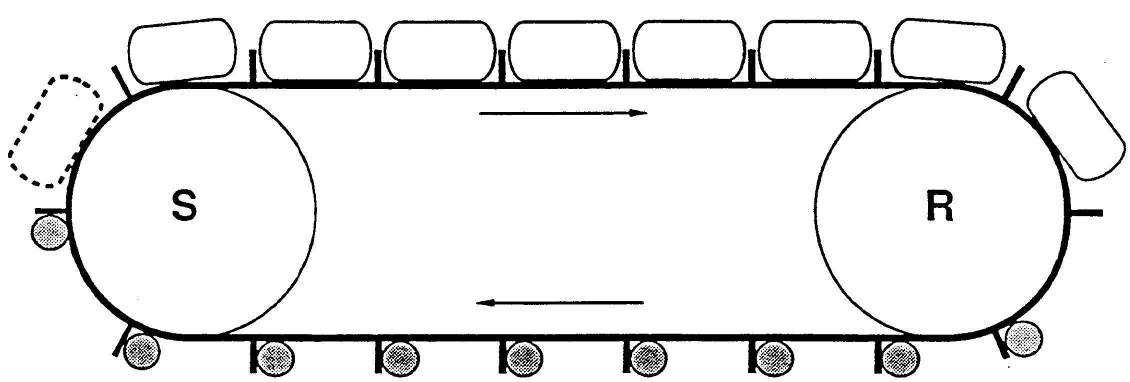

Time 1 R⁻ — Retransmit arrives at receiver.

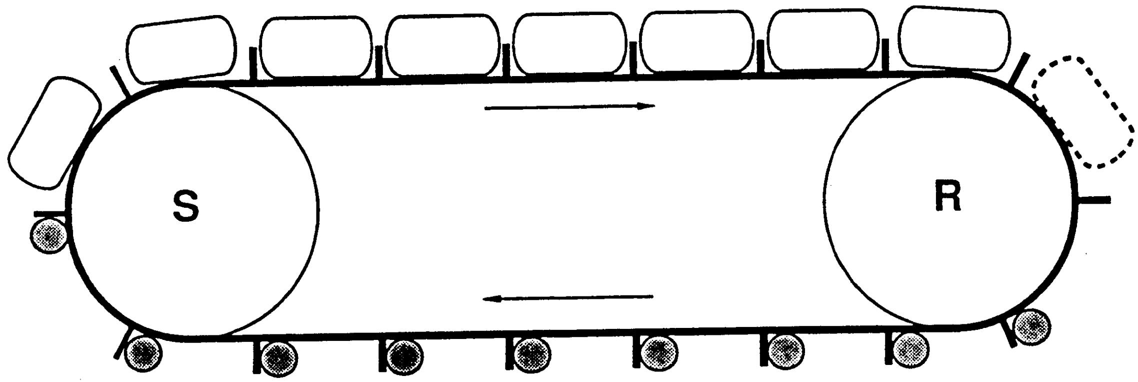

Time 1 R+ — Seq space hole filled.

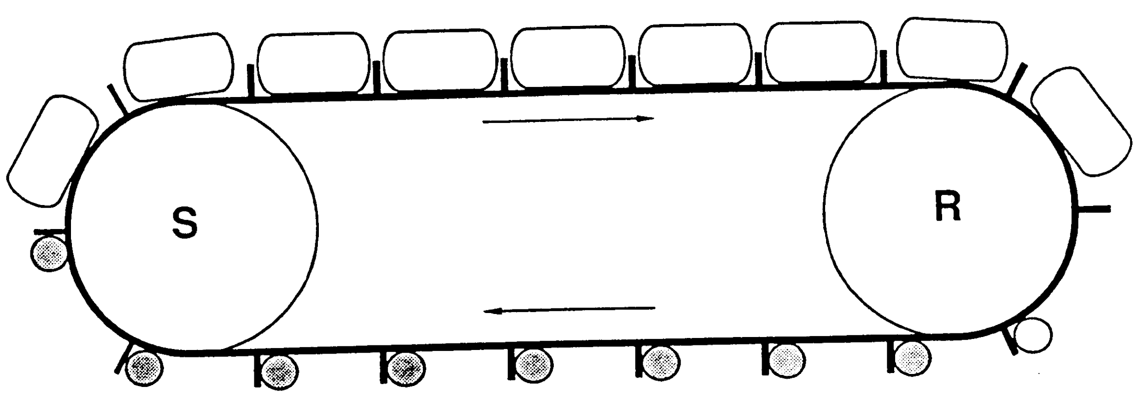

Time 1.5 R^- — Almost back to normal.

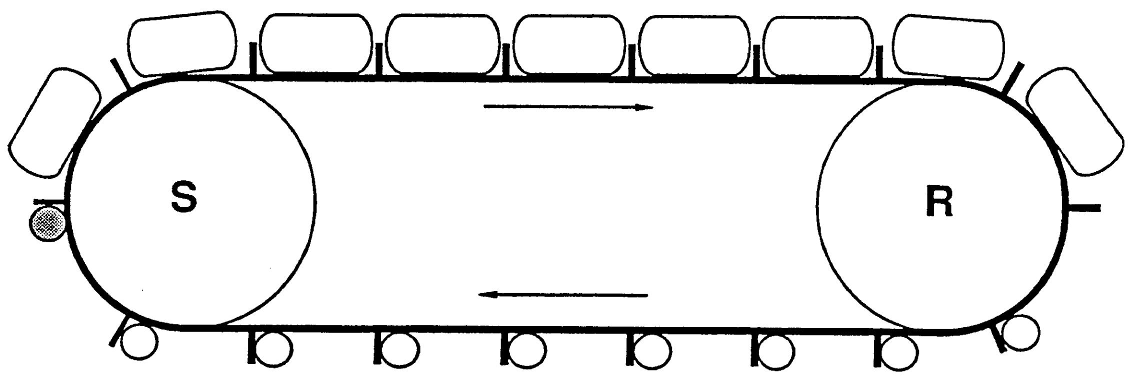


```
if (packet is duplicate ack) {
        if (++tp->dupacks == tcprexmtthresh) {
                tp->snd_ssthresh = max(tp->snd_cwnd / 2, tp->maxseg);
                tp->snd_nxt = tp->snd_una;
                tp->snd_cwnd = tp->maxseg;
                tcp_output(tp);
                tp->snd_nxt = tp->snd_max;
                tp->snd_cwnd = tp->snd_ssthresh + tp->dupacks * tp->maxseg;
        } else if (tp->dupacks > tcprexmtthresh) {
                tp->snd_cwnd += tp->maxseg;
                tcp_output(tp);
        }
        goto drop;
}
if (tp->dupacks >= tcprexmtthresh  &&  tp->snd_cwnd > tp->snd_ssthresh)
        tp->snd_cwnd = tp->snd_ssthresh;
tp->dupacks = 0;
```

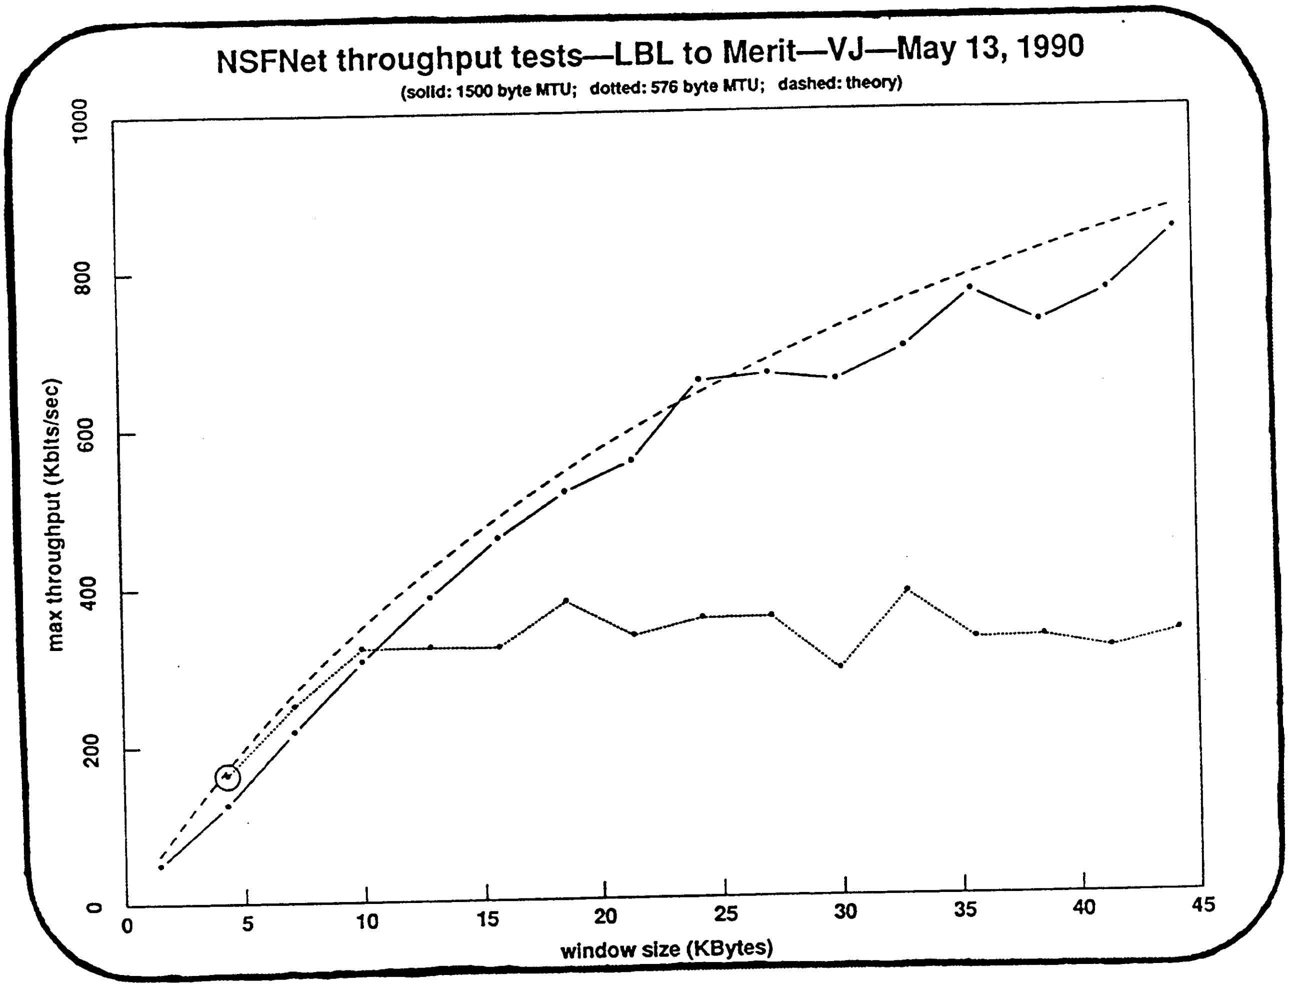
NSFNet throughput tests—LBL to Merit—VJ—May 13, 1990
(solid: 1500 byte MTU; dotted: 576 byte MTU; dashed: theory)
max throughput (Kbits/sec)
0
200
400
600
800
1000
0
5
10
15
20
25
30
35
40
45
window size (KBytes)

6.6 Scaling and Policy in the Internet

Presentation by Paul Tsuchiya

SCALING AND POLICY IN THE INTERNET

PAUL F. TSUCHIYA

Problems

ROUTING DOESN'T SCALE

ADDRESS DEPLETION

NO POLICY ROUTING

THIS TALK:

- Scaling and Policy Routing
  - Using Multiple Hierarchical Addresses (Polly)
- Efficient and Flexible Hierarchical Address Assignment
  - Using hierarchical, non-contiguous masks (Kampai Addressing)

New Terms

Polly

- What I call the technique of using multiple hierarchical addresses for scaling and policy

Kampai Addressing

- What I call the technique of assigning addresses from the bottom-up using non-contiguous masking

Rumperephobia

- Just wait and see

Scaling: How To

Tricks Exist:

- Default Route
- Localize Routing Table
  - Tunnel through AS between Border Gateways
  - Remote query to Route Server to find route (ORWG technique)

But only one fundamental scaling technique:

- HIERARCHICAL ADDRESSES

Current and Future "Internet" Structure

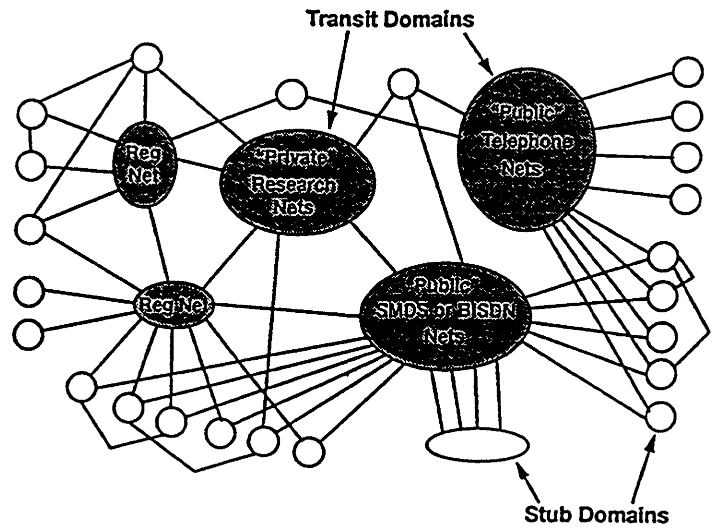

Problems with Hierarchical Addressing

We have ~~two~~ three problems using hierarchical addressing:

1. IP addresses aren't hierarchical (or at least, not very)
2. Hierarchy and policy tend to be contradictory goals
3. Hierarchical Addresses couple Domains to backbones
   - Addresses change when backbone changes
   - Multiple backbone connections

Effect of Hierarchical Addresses
Scaling, but Little Policy

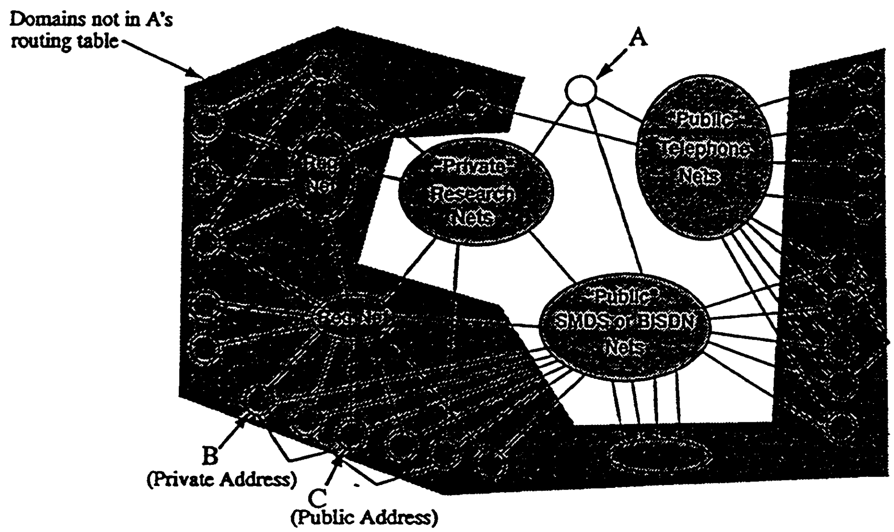

Effect of Ignoring Hierarchical Addresses
Policy, but Little Scaling

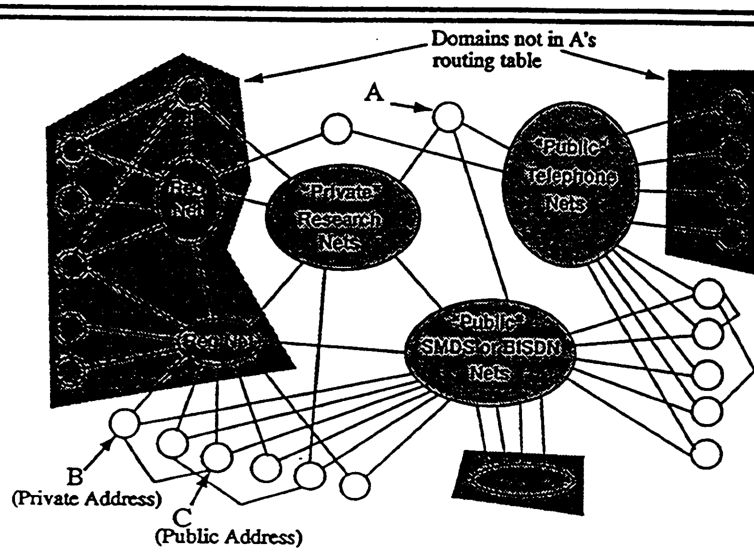

Solution: Multiple Hierarchical Addresses Scaling and Policy

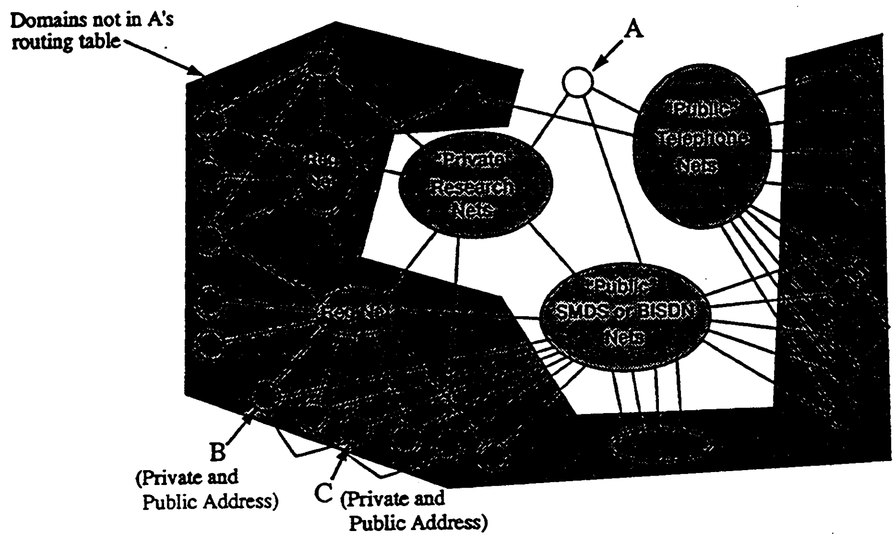

Basic Idea Behind Multiple Hierarchical Addresses ("Polly")

Before connection establishment, Directory Service, or Domain Name System (DNS), returns *K* addresses, one for each backbone path

"User" chooses the appropriate address as a policy decision

Routing efficiently routes packets, because address is hierarchical

BUT............ doesn't this do nothing more than solve the problem in one area (routing) by creating one in another (naming)?????

EXACTLY!!!
Impact of Polly

Directory Service must return multiple addresses

- But this only <u>increases</u> load by a <u>small factor K</u>
  - K = number of addresses
  - Think of returning multiple addresses as returning one big one

User must pick an address

- But this just part of user's normal policy decision

Intra-domain routing load also <u>increased</u> by a <u>small factor K</u>

Inter-domain routing load <u>decreased</u> from roughly **N^2** to **NlogN**

<u>Routing load decrease outweighs Directory Service load increase</u>

Policy Routing Background

Hop-by-hop (BGP)

- Routers calculate multiple paths to destinations
  - Either link-state or distance-vector
  - Use TOS field in header or equivalent
- Routers must understand host policies

Source Routing (ORWG)

- Each "host" calculates multiple paths to destination
- Path is setup in advance of data packets
  - Use some kind of Path ID field in header
  - Usually link-state
- Rely on caching in routers

Using Polly for Policy: Yet Another Tool

Polly enhances, NOT replaces, other techniques

Hop-by-hop

- Relieves some router burden of knowing host policy

Source Routing

- Provides easier scaling
  - Path setup only to backbone network

Can *engineer* use of multiple addresses, in conjunction with other techniques

Address Explosion?

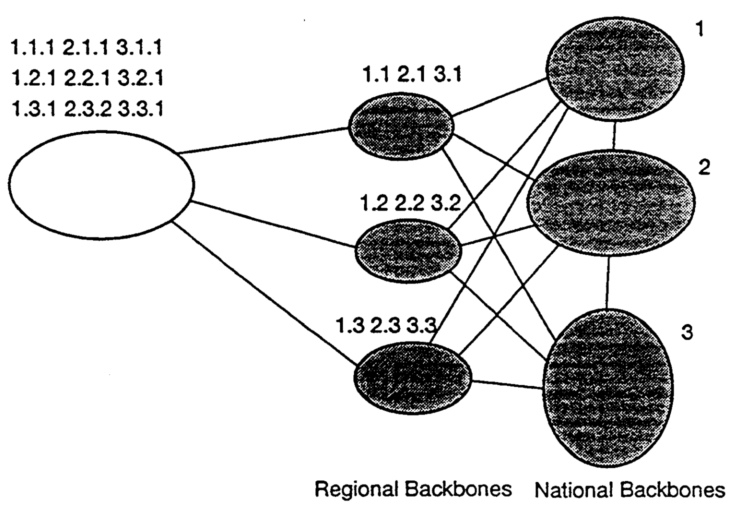

NO!
Engineer Addresses as Appropriate

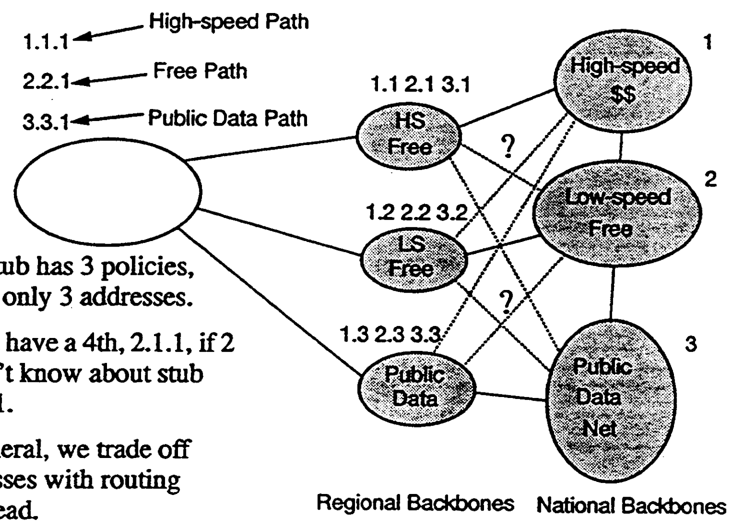

The stub has 3 policies, needs only 3 addresses.

Could have a 4th, 2.1.1, if 2 doesn't know about stub via 2.1.

In general, we trade off addresses with routing overhead.

Address Depletion Problem
Options:

<u>Make IP Address bigger</u>

- Variable-length address
- New option field with additional address space

- Requires changes to *hosts and routers*
- Why not just use ISO IP in that case?

<u>Use existing IP Address more efficiently</u>

- Current IP Address assignment techniques:
  - Scales poorly (not enough hierarchy)
  - Uses space inefficiently
- Requires only changes to *routers*

Use Existing Space More Efficiently:

Use variable-length, non-contiguous address fields

Assign from the bottom up as needed

Distinguish address fields using masks

- Routing algorithm carries masks

I call this Kampai Addressing

Rumperephobia

Fear or intense dislike, usually irrational, of non-contiguous bit masks

- Often associated with the more common isophobia

Based on recognition that:

- if two different masks have the same number of bits, one cannot determine which is "more general"
- one can derive addresses and associated non-contiguous masks that freak-out patricia trees

If two different masks have the same number of bits, then neither is more general

It is easy to avoid masks that freak-out patricia trees

Basic Procedure: Kampai Addressing

Initially assign just enough address space to cover needs

- Large mask (many "ones")

When more addresses needed, double allocated space

- Change a "1" to a "0" in mask

Simple algorithm insures that:

- Old addresses still valid under new mask
- Other addresses still valid under their masks

Execute at each level of (topological) hierarchy

- Hosts requests from subnet, subnets request from network, network from regional, regional from national, national from root (Postel)

Kampai Addressing Example (for one level of address hierarchy)

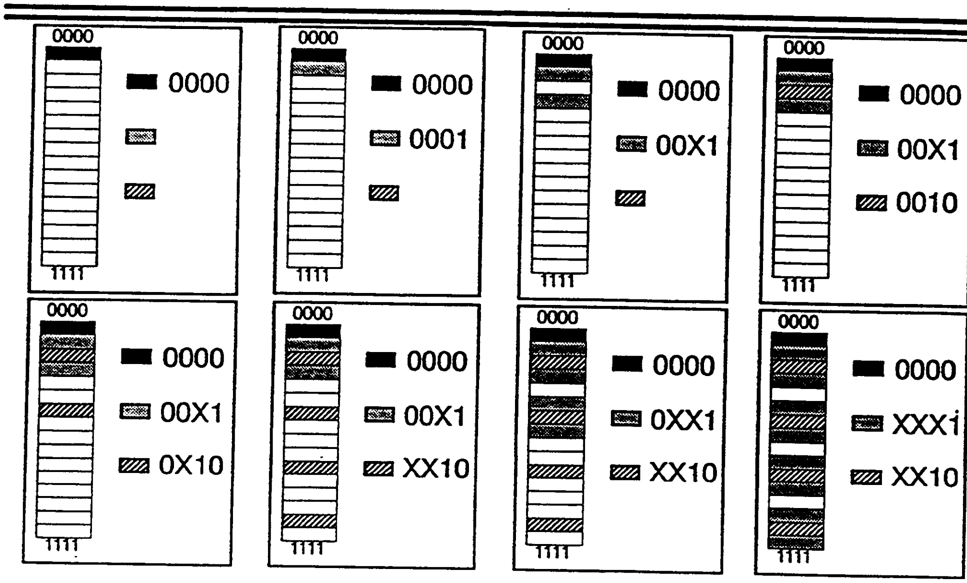

Kampai Addressing Efficiency

Around 50% address utilization at each level of assignment

- 45% utilization after request, 90% utilization before request
- Average around 65%
- Drop to 50% to include small gaps, shrinkage
- Example: if 5 levels of hierarchy, we get *6% efficiency*

Patricia lookup faster with kampai addresses than with traditional

- Because meaningful information encoded in fewer bits
- Patricia looks at fewer bits before making decision

Now Hard Part
Transition to Polly and Kampai

No changes to hosts necessary

- Although doing so gives policy control

Should make changes to intra-domain routing

- But possible to do without—less efficient

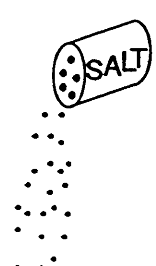

Need changes to inter-domain routing

- Modify BGP

Domain-name must return multiple addresses

- Within spec, but how about (host and DNS) implementations?

Description of Polly (complete version)
Before Connection

Determine appropriate addresses for domain

- According to connectivity, policies, routing TOS, etc.

Assign (multiple) addresses to hosts

- SNMP

Update DNS with multiple addresses

Intra-domain routing keeps track of multiple addresses

- Per host, per subnet, per area, etc.

Inter-domain routing keeps track of multiple addresses

- Per domain

Description of Polly
Full Host Version

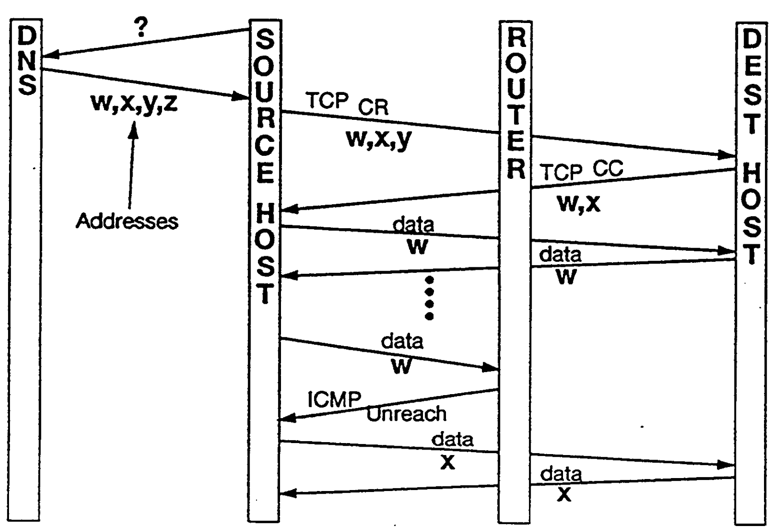

Description of Polly (complete version) During Connection

User starts connection using host name (foo.bar.xx)

Source host gets multiple addresses from DNS

According to policy, source host:

- Throws away unwanted addresses
- Rank orders remaining addresses

Source host sends TCP connect request (using first address)

- Identifies rank order of addresses in TCP option

According to policy, destination host throws away unwanted addresses

Description of Polly (complete version) During Connection (cont)

Destination host sends TCP connect accept

- Identifies remaining acceptable addresses in TCP option

Source host removes addresses unacceptable to destination host

Source and destination hosts establish connection

- Both can identify connection with any of the chosen addresses
  - No loss of performance

Assume destination host becomes unreachable during connection

Router sends ICMP destination unreachable to source host

Source host tries next available address (IP level function)

Polly Options: No Host Changes

Host cannot choose between multiple destination addresses

- Or, have no policy server, and let host just pick first address

Host doesn't recognize own multiple addresses

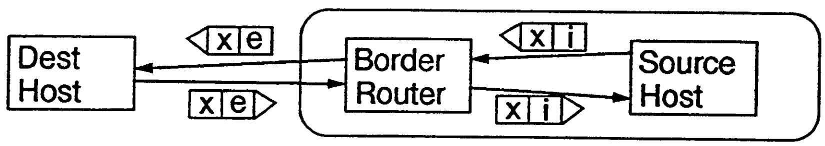

- Simple algorithmic manipulation, since "host" part of address unchanged

Intra-domain Routing Changes

Must recognize variable-length, non-contiguous masks

- OSPF already does this

Should recognize multiple addresses per host

- Can avoid this if Border Routers do address translation

Inter-domain Routing Changes

Must recognize variable-length, non-contiguous masks

Must recognize multiple addresses per domain

- BGP already does this

Address Assignment

Until all intra-domain routing protocols recognize non-contiguous masks, IP Addresses must follow old style (Class A/B/C)

- Let's consider this the transition period

During transition, let major backbones assign addresses from Class A spaces

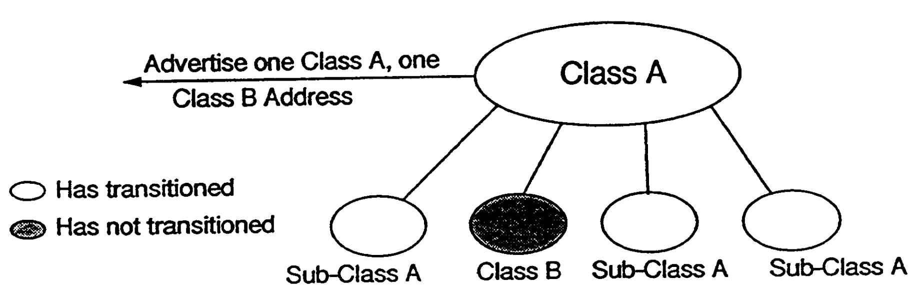

Address Assignment (cont)

After transition, backbones can obtain address outside of Class A space

- According to rules of Kampai Addressing

Each domain should always have its own, non-backbone oriented address space

- Not advertised outside domain
- Allows for stable addressing internally

Summary

Kampai Addressing solves scaling and address depletion problems

- To the extent possible, given only 32 bits of address space

Polly provides policy routing control

Transition seems doable

- Phased transition possible
- Hosts can avoid changes in perpetuity (although not recommended)
- Routing changes not overly complex

<u>OR: Just wait for OSI—they may have Polly and Kampai anyway</u>

6.7 NASA ACTS Satellite

Presentation by Tom VonDeak/NASA Lewis

The NASA ACTS Program was initiated to advance the state-of-the- art in satellite communications. It however has proved capable of advancing, in general, the state-of-the-art in communications. In the early 1980's, members of the ACTS Program identified key, high risk technologies for advanced development. These technologies have been combined to form a platform which the ACTS Program will use to verify and demonstrate ACTS developed technologies. The ACTS satellite is scheduled to be launched into a geosynchronous orbit by the Shuttle in May 1992.

A primary function of the ACTS satellite is its availability for the development of applications. Some of the key application areas that are being currently developed within our program are narrowband ISDN, HDTV, and Gigabit/second communications. The network services provided to the science community are in the initial stages of development. Even at this stage of development it is recognized that in order to support the science community to the fullest extent possible it is necessary that the ACTS satellite network provide connections into existing data networks. Identifying the segment of the science community best able to make use of the ACTS Program and defining their data networking requirements has become a top priority of the ACTS Program. Among the science oriented applications under study are the placement of 46.828 MBit/sec Earth Stations at Palmer Station, Antarctica, astronomer access and control of remote observatory instruments, database file transfer, and researcher access and control of unattended telemetry instruments.

Among the enabling ACTS technologies are: Ka-band Multi-beam antenna, a mechanically steerable antenna, On-board N x 64 KB/S circuit switching, Real time circuit allocation, and three active transponders each having 900 MHz of bandwidth .

The Multi-beam antenna system developed for the ACTS Program accommodates the formation of 1/2 degree (about 150 miles diameter) beams at Ka-band using antenna sizes of 2.2 and 3.3 meters, respectively. In addition, the ACTS satellite has a 1 degree beam that can be moved to any location in the Western Hemisphere ranging from within the Antarctic/Arctic Circles to the Azores and Midway Island. The narrow beamwidth concentrates the power of satellite transmissions and increases the dB/K of the satellite receiver. The increased antenna gain at Ka-Band along with the improved transmission/reception characteristics of the ACTS satellite result in Earth Stations with small diameter antennae and low power transmitters being able to communicate with other Earth Stations at high data rates. ACTS Earth Stations, under current program development, will have throughputs of 1.792 Mps and 46.828

Mps. ACTS Earth Stations capable of transmitting at BISDN rates of 150 Mps to 933 Mps are under initial stages of definition and development.

On-board N x 64 KB/S circuit switching supports communications between Earth Stations residing in separate beams of the ACTS satellite. The majority of the ACTS Ka-band beams are separated into two groups referred to as West and East families. In addition, there are three separate beams on Cleveland, Tampa, and Atlanta. These groups are serviced by two baseband processors each capable of routing 1,200 individual 64 KB/S circuits. The baseband processor multiplexes the outgoing circuits from uplink beams into a serial data stream for the destination beam. This stream is transformed into a downlink TDMA burst from the satellite to Earth Stations residing in that beam.

Real time circuit allocation is carried out using a low data rate communication path between Earth Stations and the Master Control Station. This communication path is integrated into the TDMA bursts and allows the earth station and master control station, located at Lewis Research Center, to exchange information and commands such as circuit setup and disconnect messages. The master control station creates a routing table based on the circuit commands exchanged with the earth stations. This routing table is transmitted to the satellite where it is implemented to establish n x 64 Kbps channels between earth stations. Existing channels between earth stations remain intact while new circuits are established. The time between a circuit request and its actual creation varies according to the number of outstanding requests and network traffic but is estimated to be less than 4.5 seconds under all but the most extreme circumstances.

The 900 MHz wide transponders of the ACTS satellite offer a unique opportunity to conduct experiments with Broadband ISDN (B-ISDN) rates ranging from 150 MB/S to 933 MB/S. The 900 MHz bandwidth capacity of the three Ka-Band ACTS transponders is accessed by the Master Control Station issuing appropriate commands to the satellite. While in this mode of operation, beam locations are interconnected via an analog Microwave Switch Matrix rather than the digitally oriented baseband processor. The Microwave Switch Matrix has bandwidth characteristics matching those of the transponders. This bandwidth coupled with the high G/T and dB/K of the ACTS satellite have made it possible to pursue the development of Earth Stations with gigabit/sec transmission capabilities.

The ACTS satellite is proving itself to be a flexible platform rich with features capable of supporting a wide range of interests. The requirements imposed by applications and uses of ACTS will determine the earth station interfaces and the direction of any developmental efforts.

Although the on-orbit phase of the ACTS Program is 22 months in the future, the planning phase of the Experiments Office is at the height of its activity. Earth Stations are currently being defined as to the interfaces and protocols that they will support. When that definition is complete, it will (for the most part) dictate what can and cannot be accomplished in the ACTS Demonstration Network. One area under study is the integration of packet- switching into the ACTS Demonstration Network. The variable on- demand bandwidth in increments of 64 KBPS, the software flexible circuit set-up interface, and the real-time switching operation of the satellite switch can be developed into a very efficient satellite-based packet switching network. One of the key issues in this area is the identification of the network equipment configuration that would best serve ACTS goals in implementing a packet-switching network.

It would be inaccurate to say that the applications being developed through the ACTS Program are satellite oriented. The capabilities of the ACTS satellite enables the development of advanced applications that can be demonstrated before the matching terrestrial capabilities will be made available. The next generation of communications begins operation with the ACTS satellite.

NASA Advanced Communications Technology Satellite

Internet Engineering Task Force

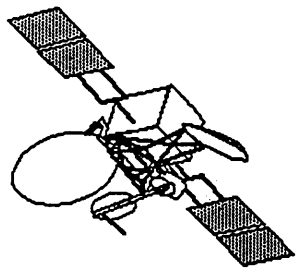

Thomas vonDeak
NASA Lewis Research Center
ACTS Experiments Office
Mail Stop: 54-6
Cleveland, Ohio 44135

AGENDA

- PROGRAM OVERVIEW
- SYSTEM OVERVIEW
- EARTH STATIONS
- EXPERIMENTS PROGRAM

ACTS

NASA

PROGRAM OVERVIEW

ACTS NASA

NASA COMMUNICATIONS PROGRAM—SPINOFFS

PROGRAMS	LAUNCHED	TECHNOLOGY SPINOFF	APPLICATIONS SPINOFF
ECHO I,II	1960, 1964	30m INFLATABLE MYLAR ANTENNA (PASSIVE REFLECTOR) SATELLITE GROUND TERMINALS	1st TRANSOCEANIC 2-WAY VOICE VIA SATELLITE
RELAY I, II	1962, 1964	10 WATT TRANSPONDERS	LIVE TV WORLDWIDE
SYNCOM II, III	1963, 1964	1st GEOSTATIONARY SATELLITE STATION KEEPING AND TRACKING TECHNIQUES	PRECURSOR TO DOMESTIC SATELLITES AND INTELSAT DEVELOPED INTEGRATED LAUNCH VEHICLE, SPACECRAFT AND COMMUNICATIONS TECHNOLOGIES FOR COMMERCIAL USE
ATS–1,3,5	1966, 67, 69	ELECTRONICALLY AND MECHANICALLY DESPUN ANTENNAS MULTIPLE ACCESS TECHNIQUES	LAND, AERONAUTICAL AND MARITIME MOBILE SPACE PHOTOGRAPHY POSITION LOCATION DISASTER COMMUNICATIONS DATA COLLECTION EMERGENCY MEDICAL RURAL HEALTH CARE PUBLIC SAFETY TRANSPORTATION

NOTE: SCIENCE SPINOFFS HAVE ALSO BEEN EXTENSIVE BUT ARE NOT INCLUDED HERE. EXAMPLES INCLUDE:
SPIN SCAN CLOUD CAMERA, HIGH RESOLUTION RADIOMETERS, CODING FOR DEEP SPACE COMMUNICATIONS, ELIMINATION OF RADIO INTERFERENCE PROBLEMS, SCINTILLATION MEASUREMENTS, COMMUNICATION LINK CHARACTERIZATION FOR DEEP SPACE DATA ACQUISITION.

ACTS CD-45456 NASA

NASA COMMUNICATIONS PROGRAM—SPINOFFS (CONT'D)

PROGRAMS	LAUNCHED	TECHNOLOGY SPINOFF	APPLICATIONS SPINOFF
ATS–6	1974	SPOT BEAMS 10 METER UNFURABLE ANTENNA NETWORKING HAND HELD GROUND TERMINALS FOR MOBILE COMMUNICATIONS VIA SATELLITE MULTIFREQUENCY-MULTIFUNCTIONAL TECHNOLOGY	DIRECT BROADCAST RURAL EDUCATION AND HEALTH CARE VSAT DEVELOPMENT INTERACTIVE VIDEO TELECONFERENCING AND TEACHING PUBLIC SAFETY ALASKAN TV MOBILE AND PERSONAL COMMUNICATION PRECURSOR TO INDIA'S SPACE PROGRAM THROUGH SITE EXPT. OPENED UP NEW SATELLITE FREQUENCY BANDS: "L"-BAND 800 MHz AND S-BAND
CTS (NASA/ CANADA)	1976	HIGH POWER 12 GHz 200 W TRANSMITTER LARGE LIGHTWEIGHT FOLDED SOLAR ARRAYS ADVANCED SPACECRAFT STABILIZATION TECHNIQUE	DIRECT BROADCAST AND FIXED SATELLITE THIN ROUTE COMMUNICATIONS DISASTER COMMUNICATIONS VSAT DEVELOPMENT RURAL COMMERCIAL, HEALTH, EDUCATION AND LIBRARY SERVICES OPENED UP 12 GHz BAND TO COMMERCIAL, DOMESTIC, AND INTERNATIONAL OPERATIONS INTERACTIVE VIDEO LONG LIFE-TIME SATELLITES (10 YR VERSUS 1–3 YRS)

CD-45457

ACTS NASA

NASA COMMUNICATIONS PROGRAM—SPINOFFS

PROGRAMS	LAUNCHED	TECHNOLOGY SPINOFF	APPLICATIONS SPINOFF
MOBILE SATELLITE (MSAT-X)	1992/3 (COMMERCIAL)	LARGE SPACECRAFT ANTENNA NETWORK CONTROL SMALL, LOW COST PHASED ARRAY AND OMNI MOBILE GROUND ANTENNA PROPAGATION EFFECTS	PUBLIC SAFETY DISASTER COMMUNICATIONS TRANSPORTATION COMMERCE RURAL TELEPHONE CELLULAR INTERCONNECT DATA ALL LAND, MARITIME AND AERONAUTICAL MOBILE VOICE AND DATA SERVICES
ACTS	1992	"HOPPING" SPOT BEAMS ONBOARD SWITCHING 30/20 GHz S/C AND GROUND TERMINAL TECHNOLOGY VERY HIGH CHANNEL CAPACITY, ON DEMAND DYNAMIC RAIN FADE COMPENSATION	SUPER COMPUTER NETWORKING SCIENCE DATA NETWORKING FIBER OPTIC INTERCONNECT RURAL COMMERCIAL, HEALTH, EDUCATION AND LIBRARY SERVICES DISASTER COMMUNICATIONS MOBILE COMMUNICATIONS HDTV (HIGH DEFINITION TV) INTEGRATED SERVICES DIGITAL NETWORK SHOULD OPEN UP 20/30 GHz BAND FOR COMMERCIAL DOMESTIC AND INTERNATIONAL OPERATIONS WILL RELIEVE SPECTRUM AND ORBIT CONGESTION

CD-45459

ACTS NASA

WHAT IS ACTS?

- ADVANCED COMMUNICATIONS TECHNOLOGY SATELLITE (ACTS)
- AN EXPERIMENTAL SATELLITE SPONSORED BY NASA TO PAVE THE WAY FOR NEXT GENERATION COMMUNICATION SATELLITE
- GOVERNMENT-INDUSTRY PARTNERSHIP
- A NATIONAL RESEARCH FACILITY INCORPORATING ADVANCED CONCEPTS
- REDUCE RISK SUFFICIENTLY TO STIMULATE COMMERCIAL USE OF TECHNOLOGIES
- LAUNCH DATE: MAY 1992 (STS/TOS)
- MISSION LIFE
  - 2 YEAR EXPERIMENT PERIOD
  - 4 YEAR STATIONKEEPING FUEL

ACTS PROGRAM GOALS

- DEVELOP THE HIGH-RISK ADVANCED COMMUNICATIONS TECHNOLOGY USABLE IN MULTIPLE FREQUENCY BANDS TO SUPPORT A WIDE RANGE OF FUTURE COMMUNICATIONS SYSTEMS, FOR NASA, OTHER GOVERNMENTAL AGENCIES, AND INDUSTRY
  - ENABLE GROWTH IN THE CAPACITY AND EFFECTIVE UTILIZATION OF THE FREQUENCY SPECTRUM
  - ENABLE CONTINUED U.S PREEMINENCE IN SATELLITE COMMUNICATIONS

OBJECTIVES/APPROACH OF ACTS PROGRAM

	GOAL	APPROACH
ACTS FLIGHT PROGRAM	SUPPORT CONTINUED U.S. INDUSTRY LEADERSHIP IN THE WORLD COMMUNICATIONS SATELLITE MARKET	FLIGHT TEST HIGH RISK TECHNOLOGIES WHICH FALL OUTSIDE SPONSORSHIP CAPABILITY OF PRIVATE SECTOR USE TECHNOLOGIES IN AN EXPERIMENT PROGRAM
ACTS EXPERIMENT PROGRAM	STIMULATE COMMERCIAL USE OF ACTS TECHNOLOGIES	DEMONSTRATE TECHNICAL FEASIBILITY THROUGH <u>TECHNICAL PERFORMANCE EVALUATION EXPERIMENT</u> DEMONSTRATE APPLICATIONS THROUGH <u>APPLICATIONS EXPERIMENT</u>

ACTS NASA

ACTS SPACECRAFT CHARACTERISTICS

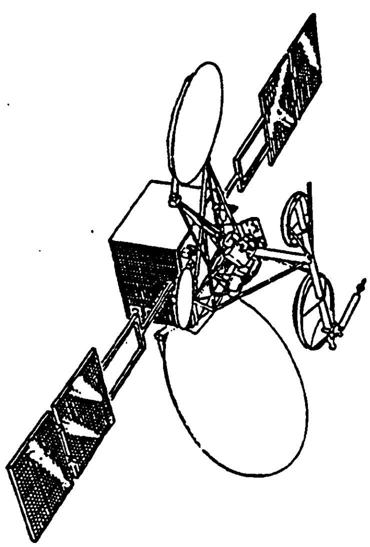

WEIGHT:	3400 lbs (ON-ORBIT)
POWER:	1770 W. BOL FOUR PANEL SOLAR ARRAY (134.5 $ft.^2$)
FREQUENCY BANDS:	Ka-BAND (30/20 GHz)
PAYLOAD:	MULTIBEAM ANTENNA, ON-BOARD PROCESSING AND ROUTING
SPACECRAFT POINTING ACCURACY:	$\pm 0.025°$
LAUNCH DATE:	MAY 1992 (STS/TOS)
MISSION LIFE:	2 yr EXPERIMENT PERIOD 4 yr STATION KEEPING FUEL

CD-45463

ACTS NASA

KEY ACTS TECHNOLOGIES

HIGH EIRP FAST HOPPING SPOT BEAMS

- SPECTRAL REUSE THROUGH SPATIAL DIVERSITY
- HIGHER THROUGHPUT VSATs (T1 RATE)
- SMALLER GROUND TERMINALS
- EFFICIENT CAPACITY ASSIGNMENT TO GEOGRAPHICALLY NON-UNIFORM DEMAND
- STEERABLE BEAM

ON-BOARD PROCESSING

- SWITCHING & ROUTING ON-BOARD AT INDIVIDUAL VOICE CIRCUIT LEVEL
- SINGLE HOP MESH VOICE NETWORK
- IMPROVED SIGNAL-TO-NOISE RATIO

Ka-BAND

- OPENING A NEW BAND
- 2.5 GHz BANDWIDTH
- DYNAMIC RAIN FADE COMPENSATION

CD-45611

COMPARISON OF ADVANCED SATELLITE TECHNOLOGY PROGRAMS

TECHNOLOGY	JAPAN	U.S.	EUROPE
L-Band (1.5 GHz)	ETS-V, VI; CS-4	Industry	SAT-2; Industry
S-Band (2.5 GHz)	ETS-VI; CS-4	Industry; TDRS	SAT-2
C-Band (6/4 GHz)	Industry; CS-4	Industry	Industry
Ku-Band (14/12 GHz)	Industry; CS-4	Industry	Industry
Ka-Band (30/20 GHz)	CS-3,4; ETS-VI; Industry	ACTS	SAT-2;OLYMPUS; ITALSAT
Optical Inter-Satellite Links	ETS-VI	LCT	SAT-2
Large, High Power Bus	CS-4	Industry	OLYMPUS
Onboard Processing	---	ACTS	ITALSAT; SAT-2 (?) OLYMPUS II (?)
Hopping Spot Beams	---	ACTS	---
Dynamic Fade Compensation	---	ACTS	OLYMPUS

ACTS SCHEDULE

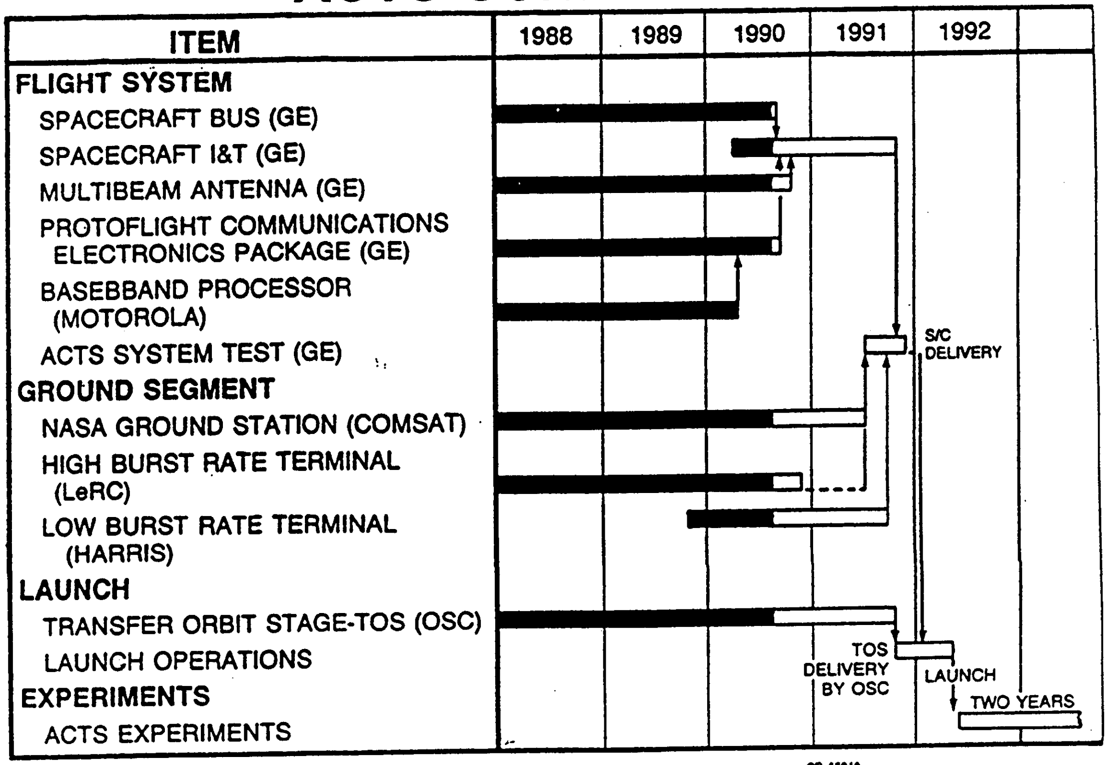

CD-40919

SYSTEM OVERVIEW

Spacecraft Configuration

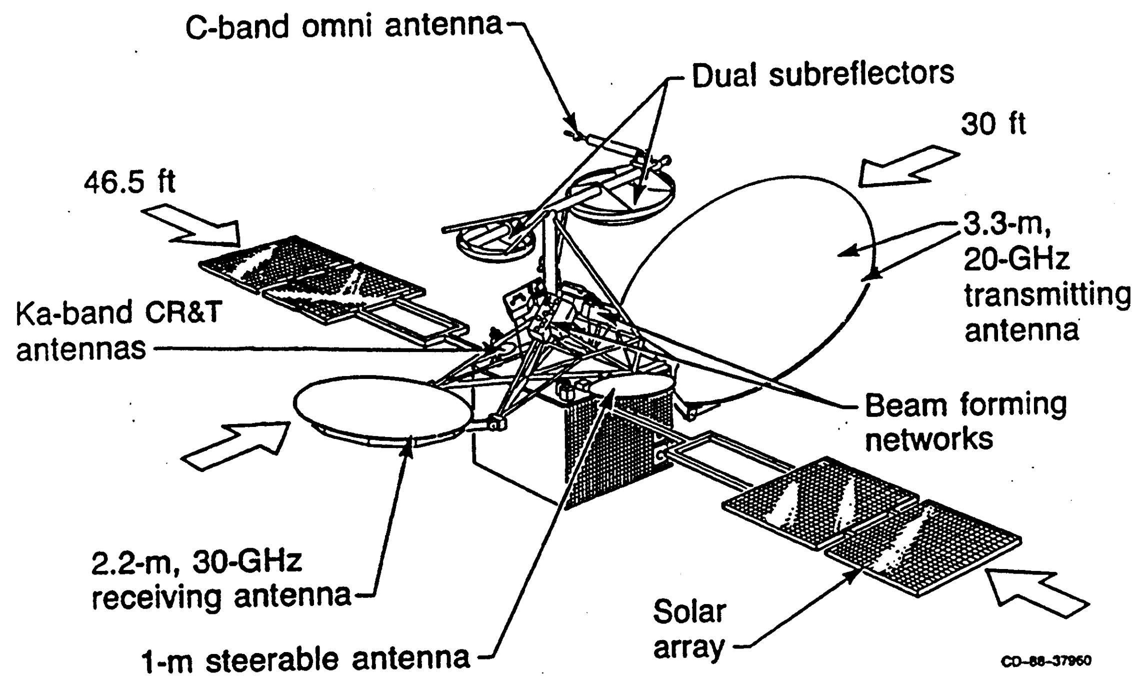

ACTS NASA

ACTS SWITCHING AND PROCESSING TECHNOLOGY

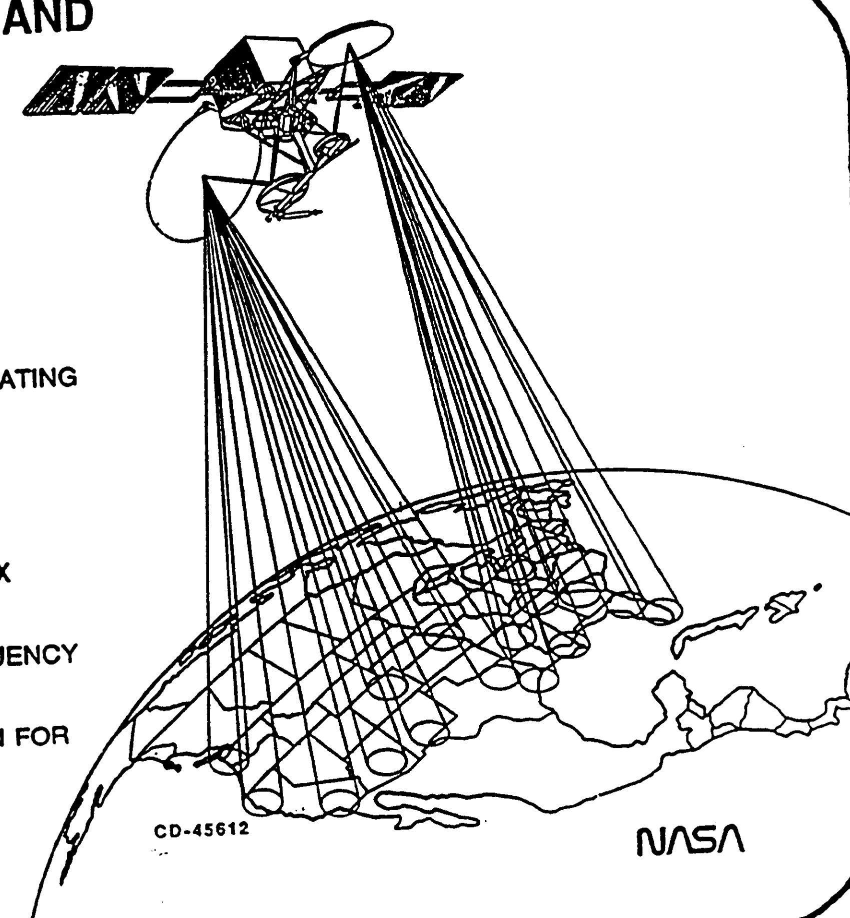

BASEBAND PROCESSOR
- DEMODULATING/REMODULATING
- DECODING/ENCODING
- ROUTING
- CIRCUIT SWITCHING
- ONBOARD MEMORY

MICROWAVE SWITCH MATRIX
- DYNAMIC "BENT PIPE" BEAM-TO-BEAM ROUTING
- UPLINK/DOWNLINK FREQUENCY TRANSLATION
- NO ONBOARD MEMORY
- STATIC-MODE OPERATION FOR CONTINUOUS CARRIERS

ACTS NASA

ACTS Multibeam Antenna Coverage

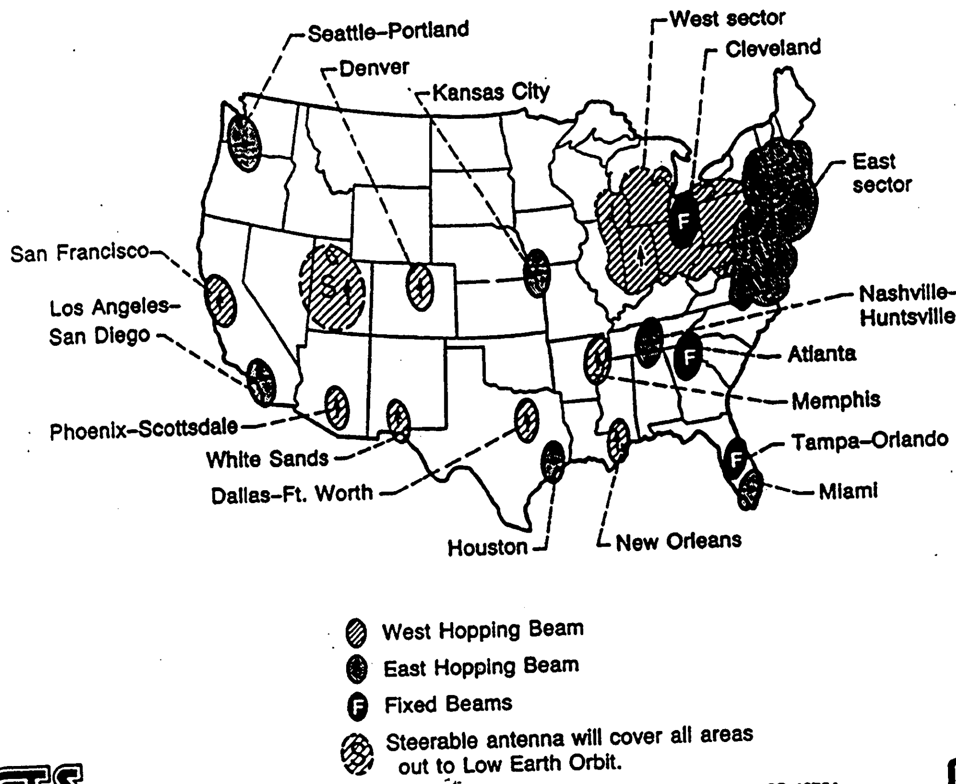

CD-48754

ACTS MECHANICALLY STEERABLE BEAM ANTENNA COVERAGE

- STEERABLE ANTENNA PRODUCES A SINGLE 1° BEAM
- TO EXTEND THE COVERAGE
- MECHANICALLY STEERED TO ANYWHERE WITHIN DISK OF EARTH AS SEEN FROM 100° WEST LONGITUDE
- STEERING RATE IS APPROXIMATELY 1 DEGREE/MINUTE

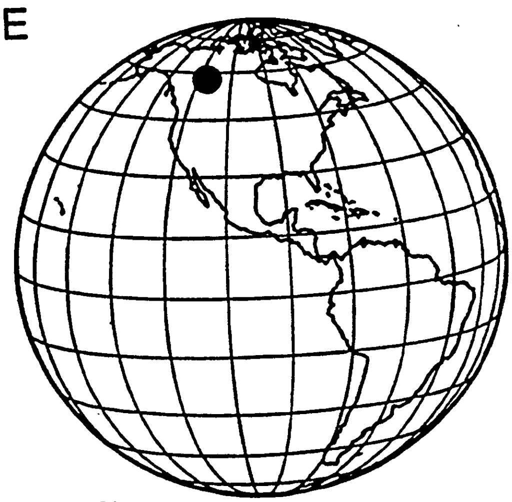

CD-43931

MODES OF OPERATION

ACTS CAN BE CONFIGURED NOMINALLY TO OPERATE IN ONE OF TWO MODES

MODE 1—BASEBAND PROCESSOR MODE (TDMA)

TWO ACTIVE BEAMS
USE THE TWO SCANNING BEAMS
ON-BOARD BASEBAND PROCESSOR DEMOD., REMOD.
UPLINK: TWO 27.5 MSPS OR ONE 110 MSPS
DOWNLINK: ONE 110 MSPS

MODE 2—MICROWAVE SWITCH MATRIX MODE

THREE ACTIVE BEAMS
ANY THREE OF THE FIVE POSSIBLE BEAMS
NO ON-BOARD BASEBAND PROCESSING
4 × 4 ON-BOARD IF SWITCH
4 FOR 3 REDUNDANCY
900 MHz BANDWIDTH

CD-44887

ACTS COMMUNICATIONS SYSTEM

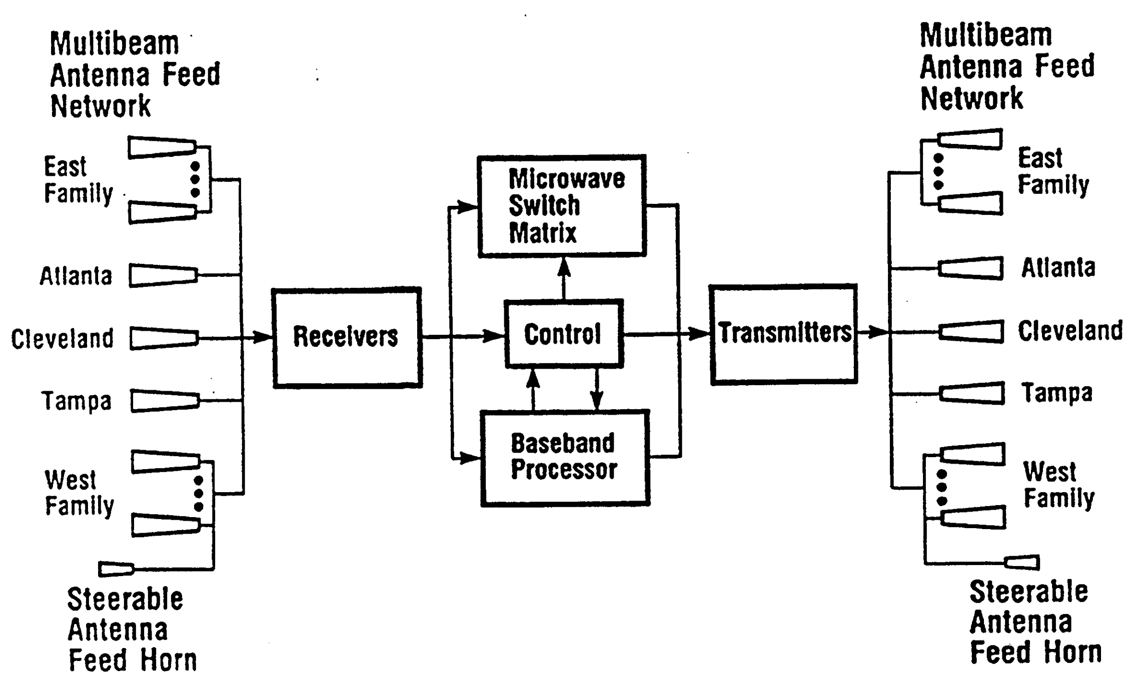

CD-43929

MICROWAVE SWITCH MATRIX MODE

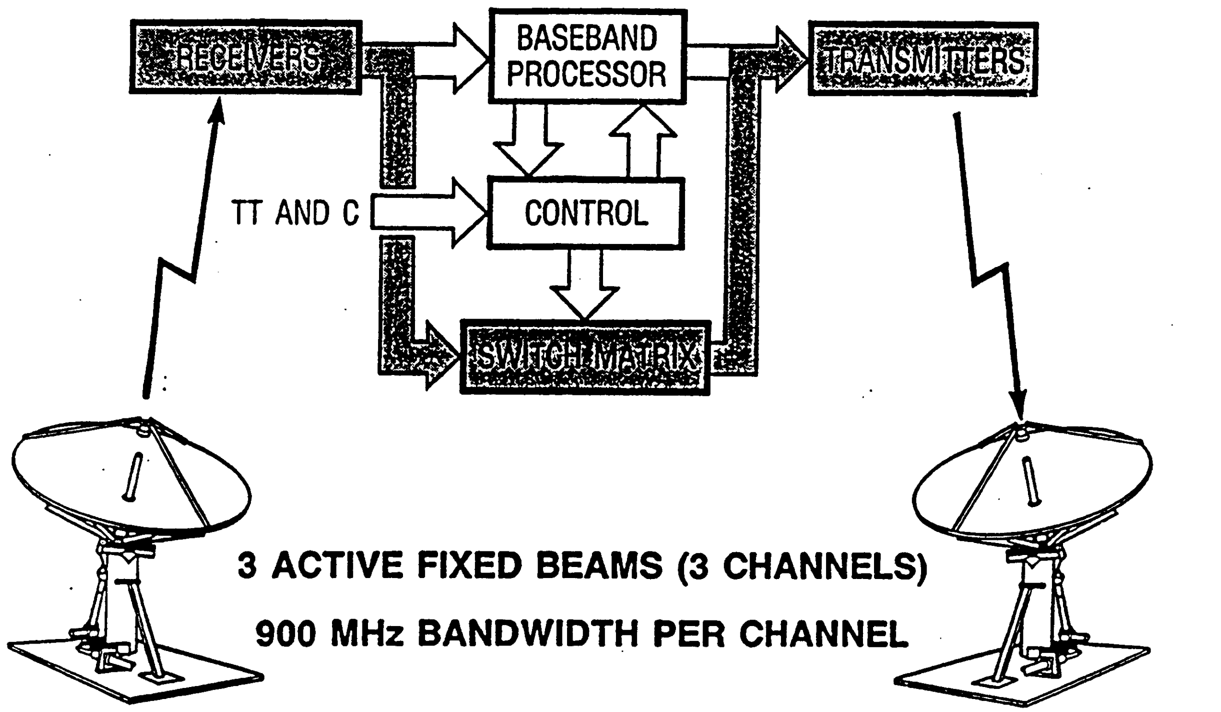

3 ACTIVE FIXED BEAMS (3 CHANNELS)

900 MHz BANDWIDTH PER CHANNEL

CD-44780

ACTS NASA

EARTH STATIONS

ACTS NASA

ACTS EXPERIMENTER TERMINAL CATEGORIES

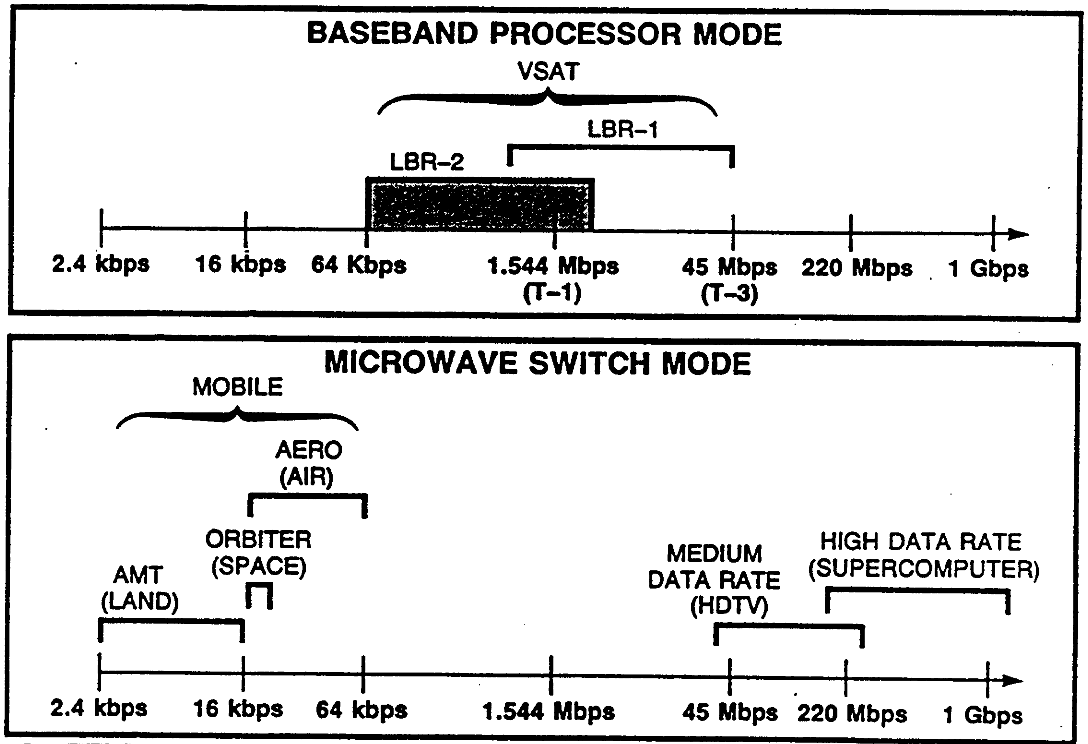

PROTOTYPE CURRENTLY FUNDED

CD-47620

BBP TYPE OF ACTS EARTH STATIONS

	NASA Ground Station	LBR-2	LBR-1
EIRP, dBW	74/68	66/60*	77/72*
G/T, dB/ºK	>27	~22/16*	22
Antenna diameter, m	~5.0	2.4/1.2*	2.4
High-power-amplifier power, W	54/14	~16	~240/60*
Uplink burst rate, Mbps	LBR-1/LBR-2	27.5/13.8**	110
Downlink burst rate, Mbps	110/55	110/55**	110

*Values shown are for baseline/option/LBR-2/LBR-1 design, Earth Station (Requirements depend on location within ACTS antenna Coverage, ninety percent (90%) of expected experimenter locations can use the lower EIRP or G/T values.

**Lower value is reduced burst rate to compensate for rain fade throughput, however, is not reduced since dwell time increases proportionately.

CD-48052

LBR-1 USER INTERFACE CONCEPT

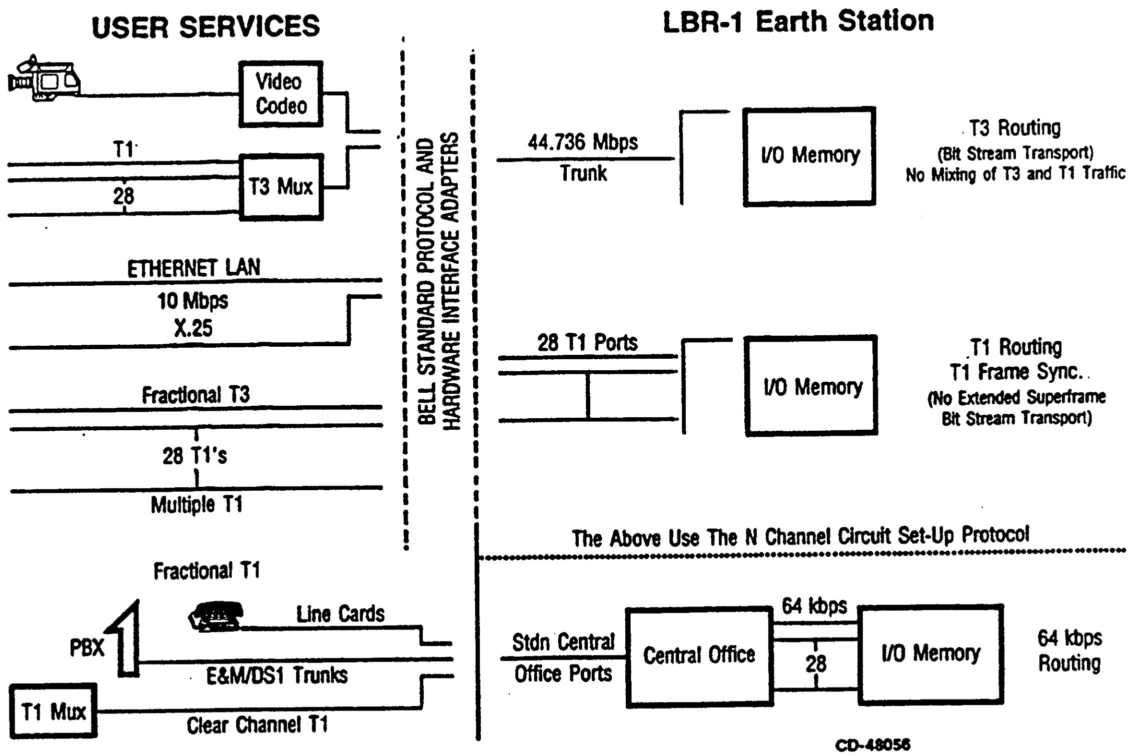

THE LBR-2 EARTH STATION PROVIDES STANDARD TELEPHONE COMMUNICATIONS INTERFACES

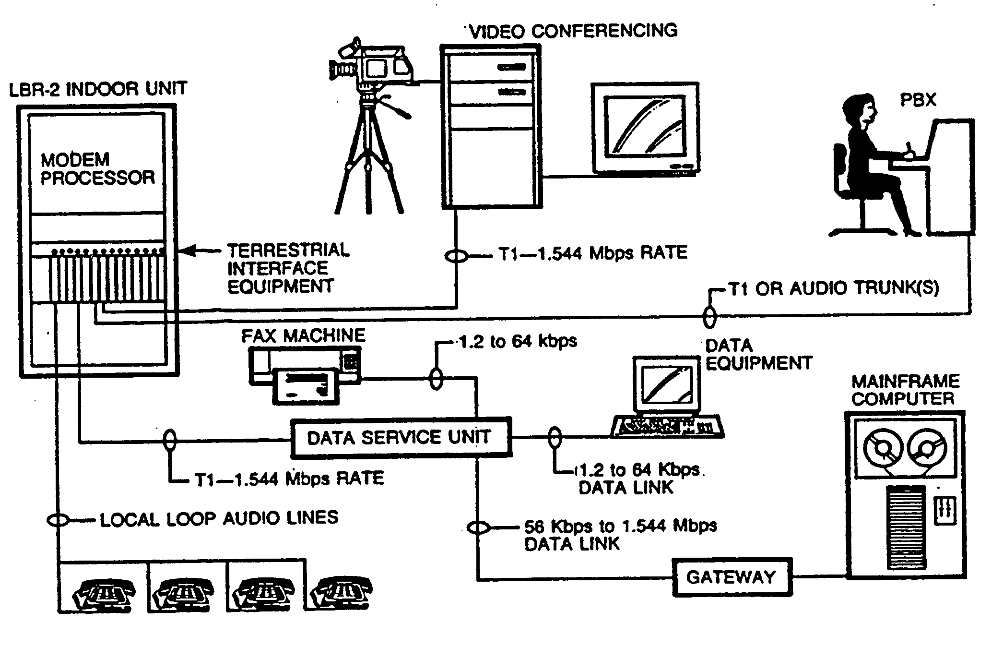

ACTS VSAT EARTH STATION
(LBR-2)

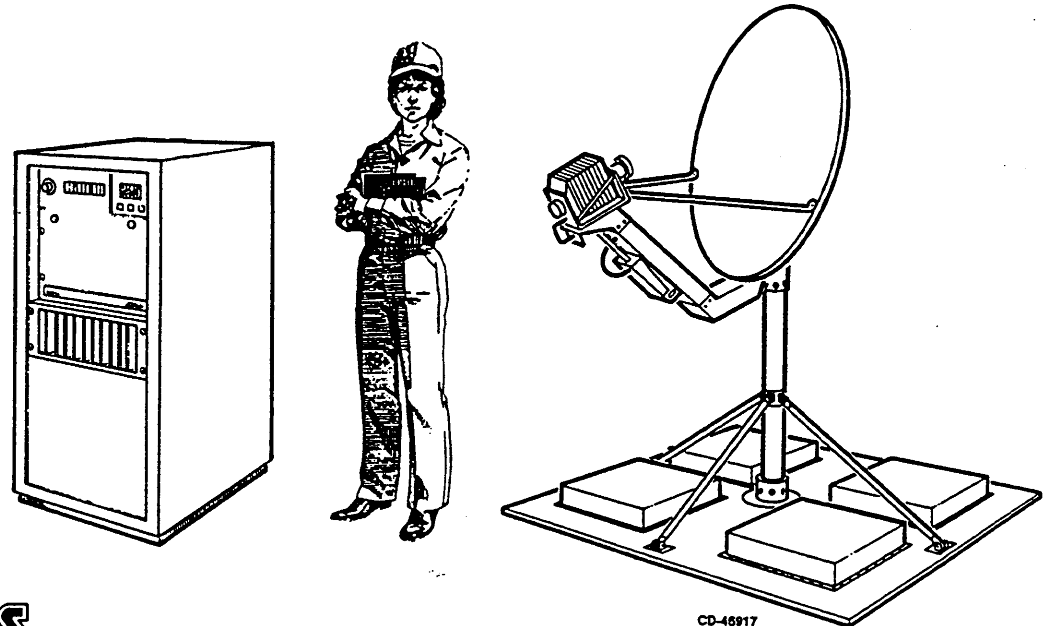

ACTS

NASA

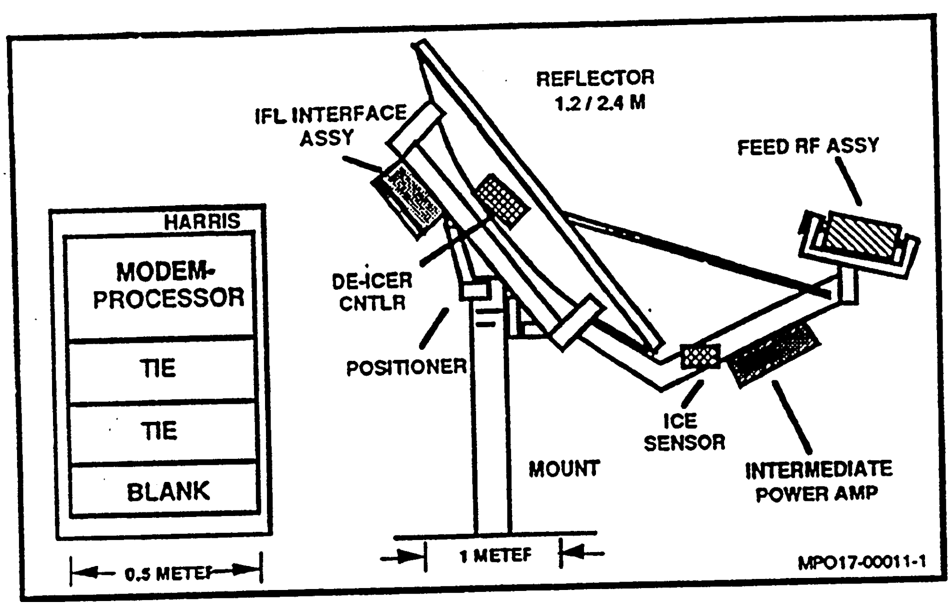

Physical Layout of ACTS Earth Station

ACTS

NASA Lewis Research Center

MSM TYPE OF ACTS EARTH STATIONS

	NASA Link Evaluation Terminal	Medium Data Rate Earth Stations	High Data Rate Earth Stations	Very Low Data Rate Earth Stations	Propagation Earth Stations
EIRP, dBW	76/68	>80	>84	TBD	——
G/T, dB/K	>27	27/21	27	TBD	15 to 18
Antenna Diameter, m	4.77	5.0	5.0	TBD	1.2
High-power-amplifier power, W	60	150	300	<1	——
Uplink burst rate Mbps	220 or 110	220	>900	<16 KBPS	——
Downlink burst rate, Mbps	220 or 110	220	>900	16 KBPS	——

CD-48054

ACTS NASA

AVAILABLE ACTS TEST BED FOR TECHNICAL INVESTIGATIONS

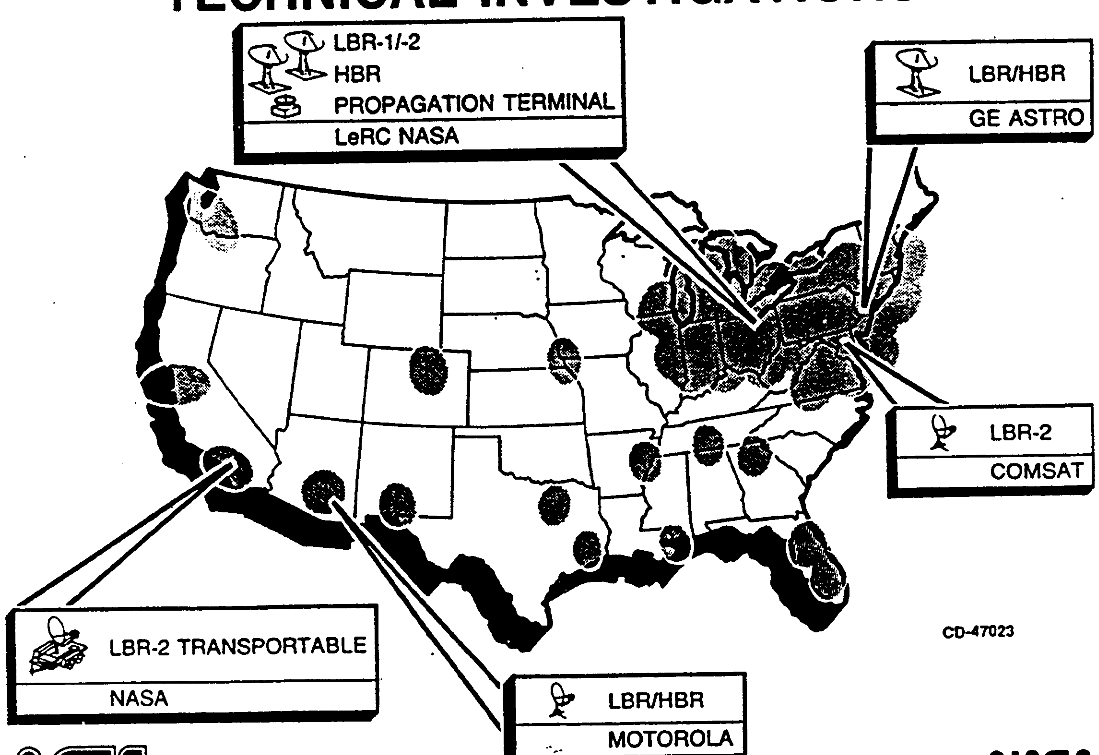

CD-47023

ACTS NASA

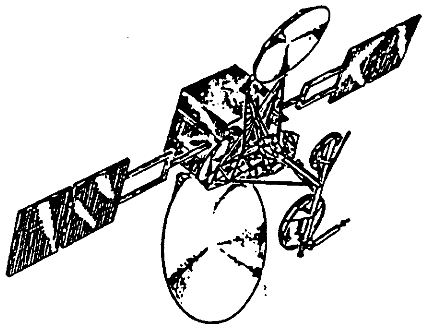

NASA

ACTS

Advanced Communications Technology Satellite

Experiments Program

CD-47728

EXPERIMENT PROGRAM OBJECTIVES

THE OBJECTIVES OF THE ACTS EXPERIMENTS PROGRAM ARE TO:

1. CONDUCT A COMPLETE SET OF TECHNOLOGY VERIFICATION EXPERIMENTS;

2. CONDUCT OF A BALANCED SET OF EXPERIMENTS WHICH EVALUATE THE POTENTIAL APPLICATIONS OF THE ACTS TECHNOLOGIES, WITH THE GOAL OF AT LEAST ONE SIGNIFICANT EXPERIMENT IN EACH OF THE APPLICATION AREAS.

QUASI-OPERATIONAL SATELLITE COMMUNICATIONS SYSTEM EXPERIMENTS MAY BE PERFORMED, AS TARGETS OF OPPORTUNITY, PROVIDED THEY DO NOT INTERFERE WITH ACCOMPLISHMENT OF THE TECHNOLOGY VERIFICATION AND APPLICATIONS OBJECTIVES.

NASA C-90-05118

THE SPECTRUM OF EXPERIMENTS

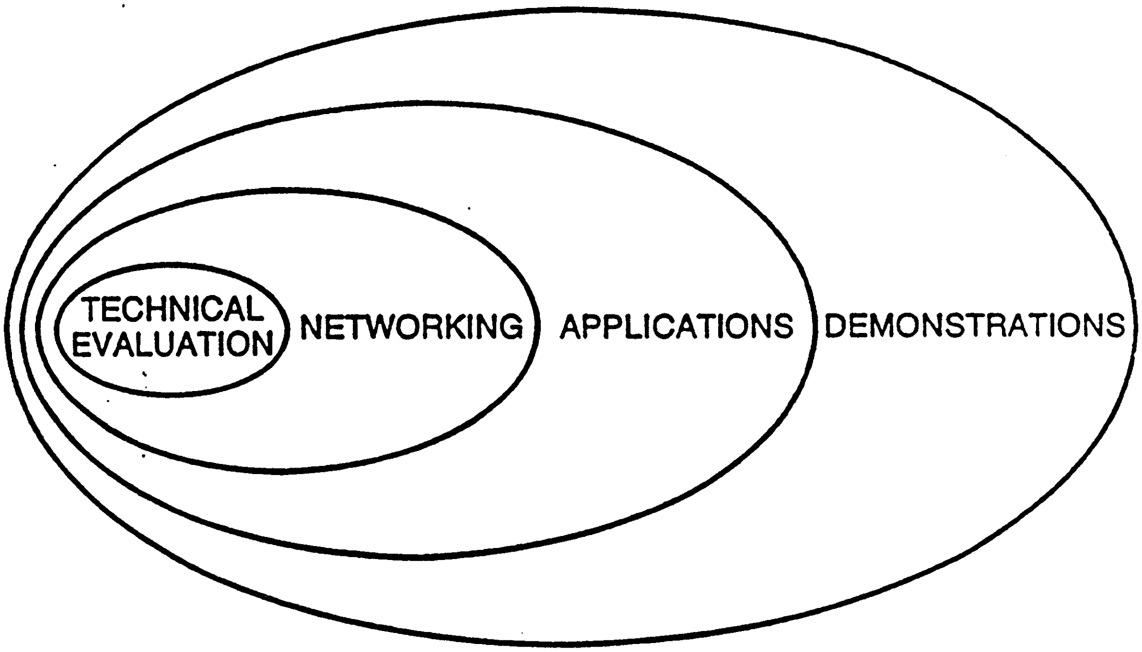

CD-46606

ACTS NASA

T1 (1.544 MBPS) VSAT NETWORK

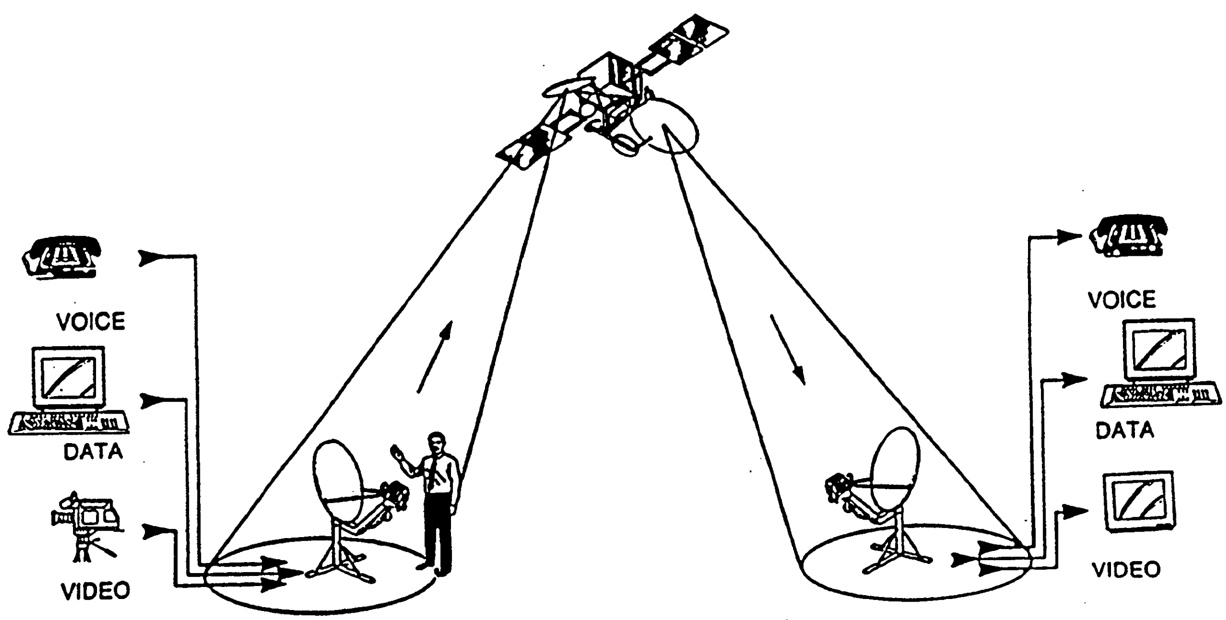

CD-44470

ACTS NASA

LAN INTERCONNECTIVITY

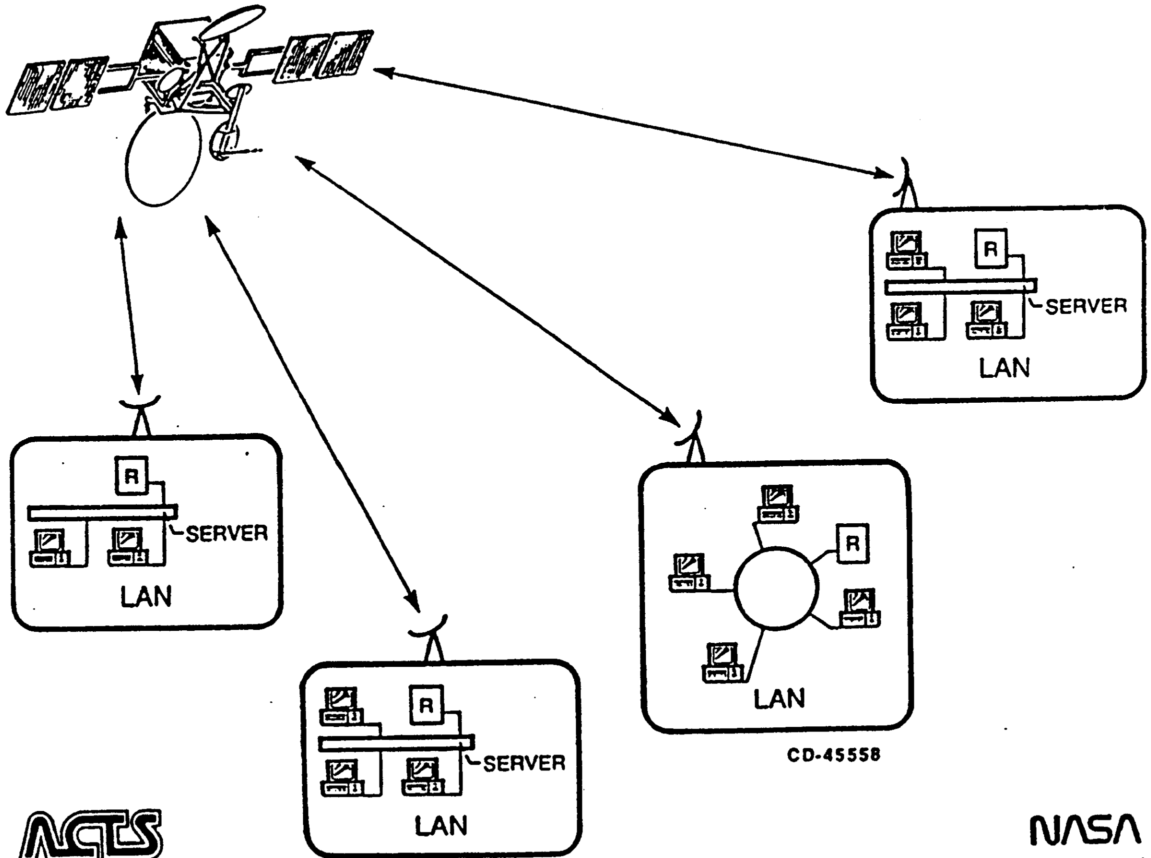

ACTS/Packet: Experimenter Support

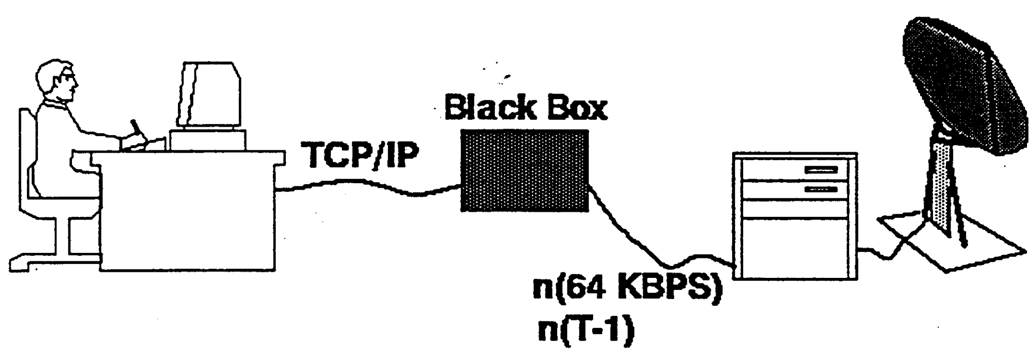

USER	INTERFACE	ACTS
Experimenter	To be found	LBR1
- Astrophysics	or	LBR2
- Astronomer	To be developed	HBR
- Antarctica	- satellite compatible -	

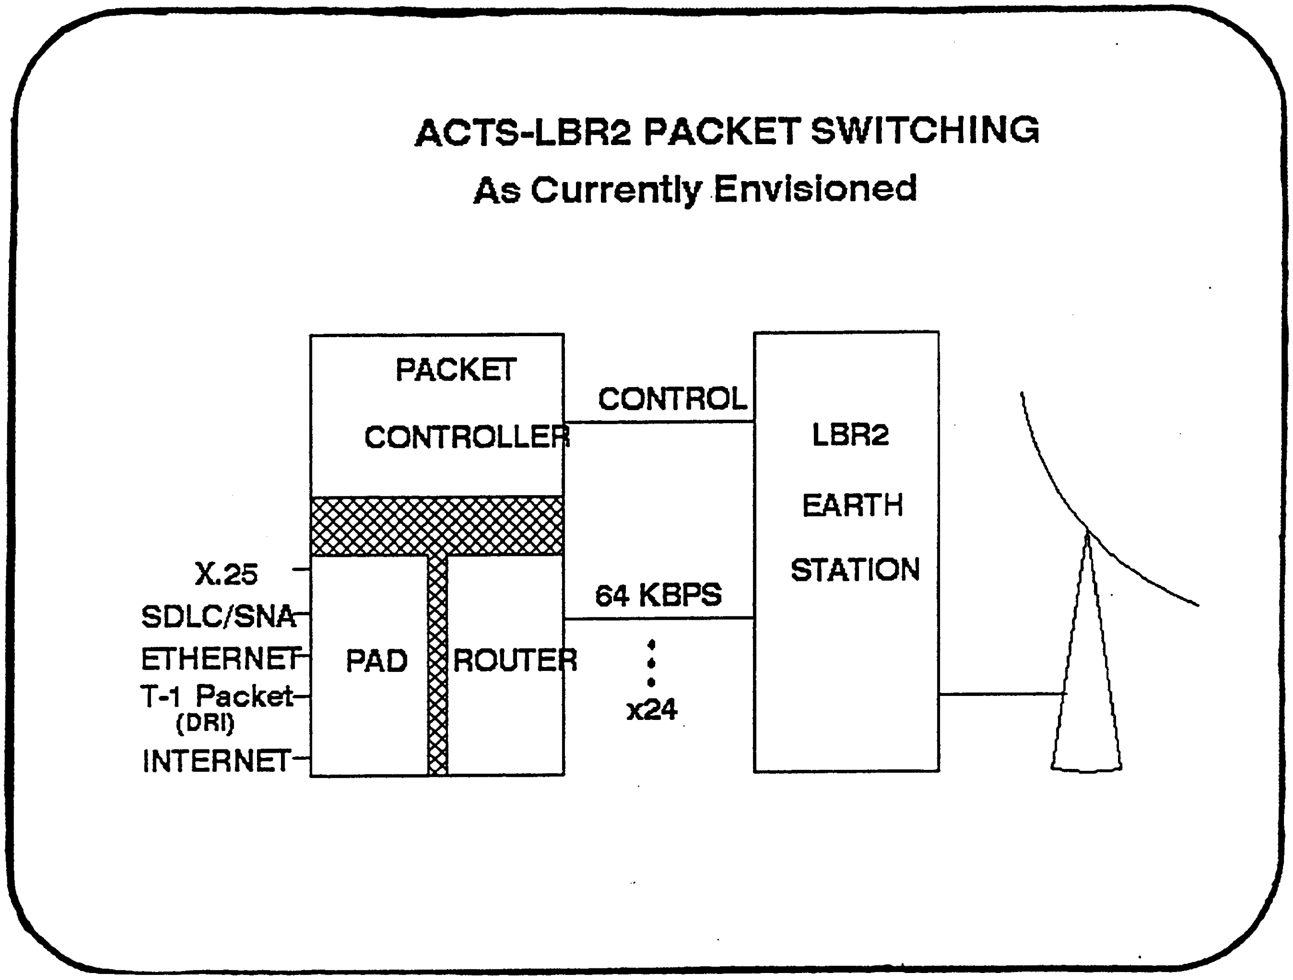
ACTS-LBR2 PACKET SWITCHING
As Currently Envisioned
PACKET
CONTROLLER
CONTROL
LBR2
EARTH
STATION
X.25
SDLC/SNA
ETHERNET
T-1 Packet
(DRI)
INTERNET
PAD
ROUTER
64 KBPS
x24

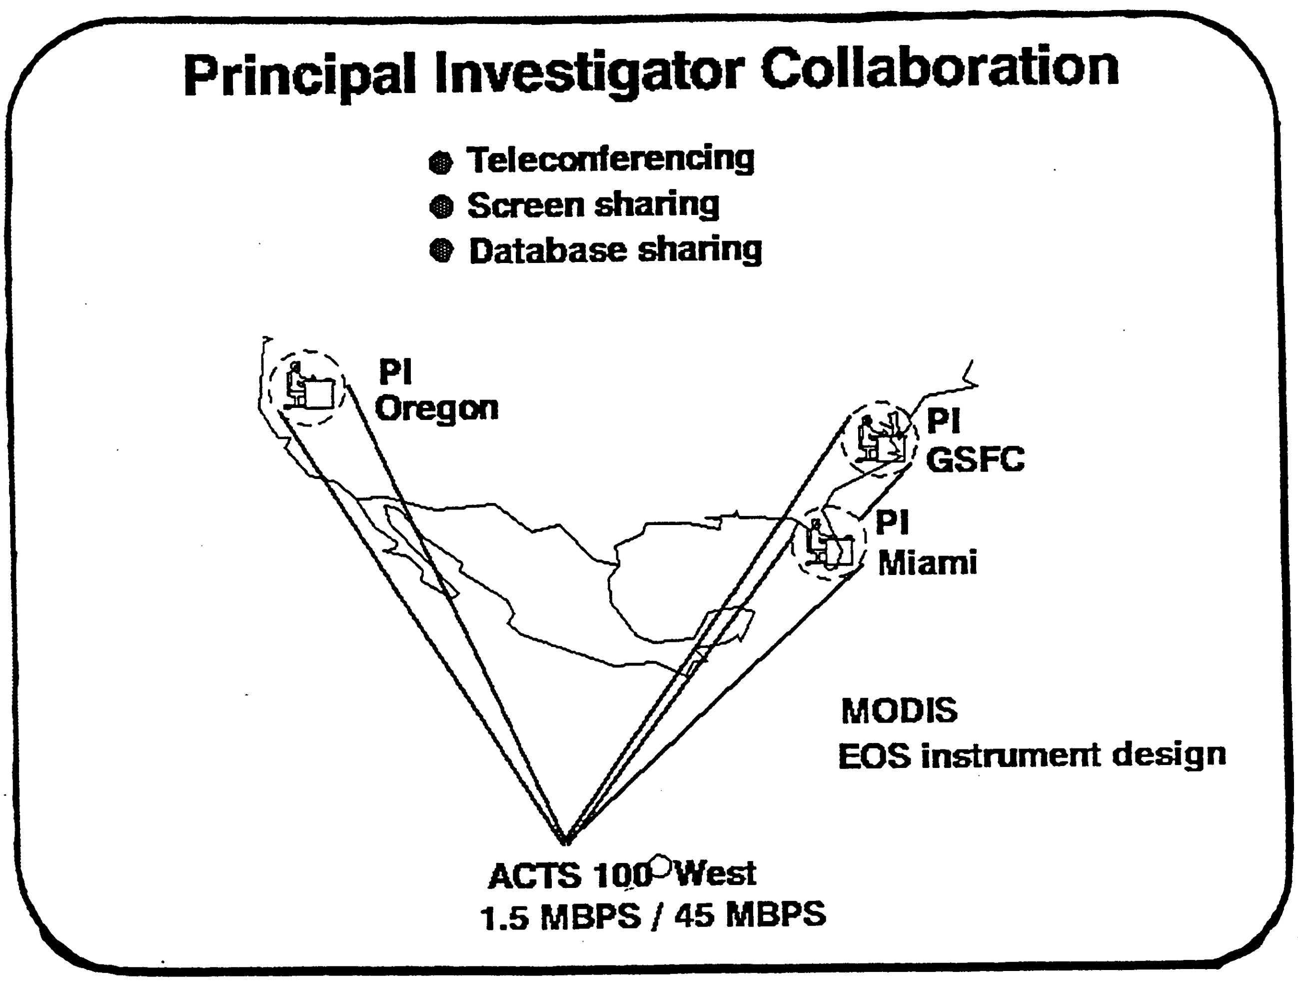
Principal Investigator Collaboration
Teleconferencing
Screen sharing
Database sharing
PI
Oregon
PI
GSFC
PI
Miami
MODIS
EOS instrument design
ACTS 100° West
1.5 MBPS / 45 MBPS

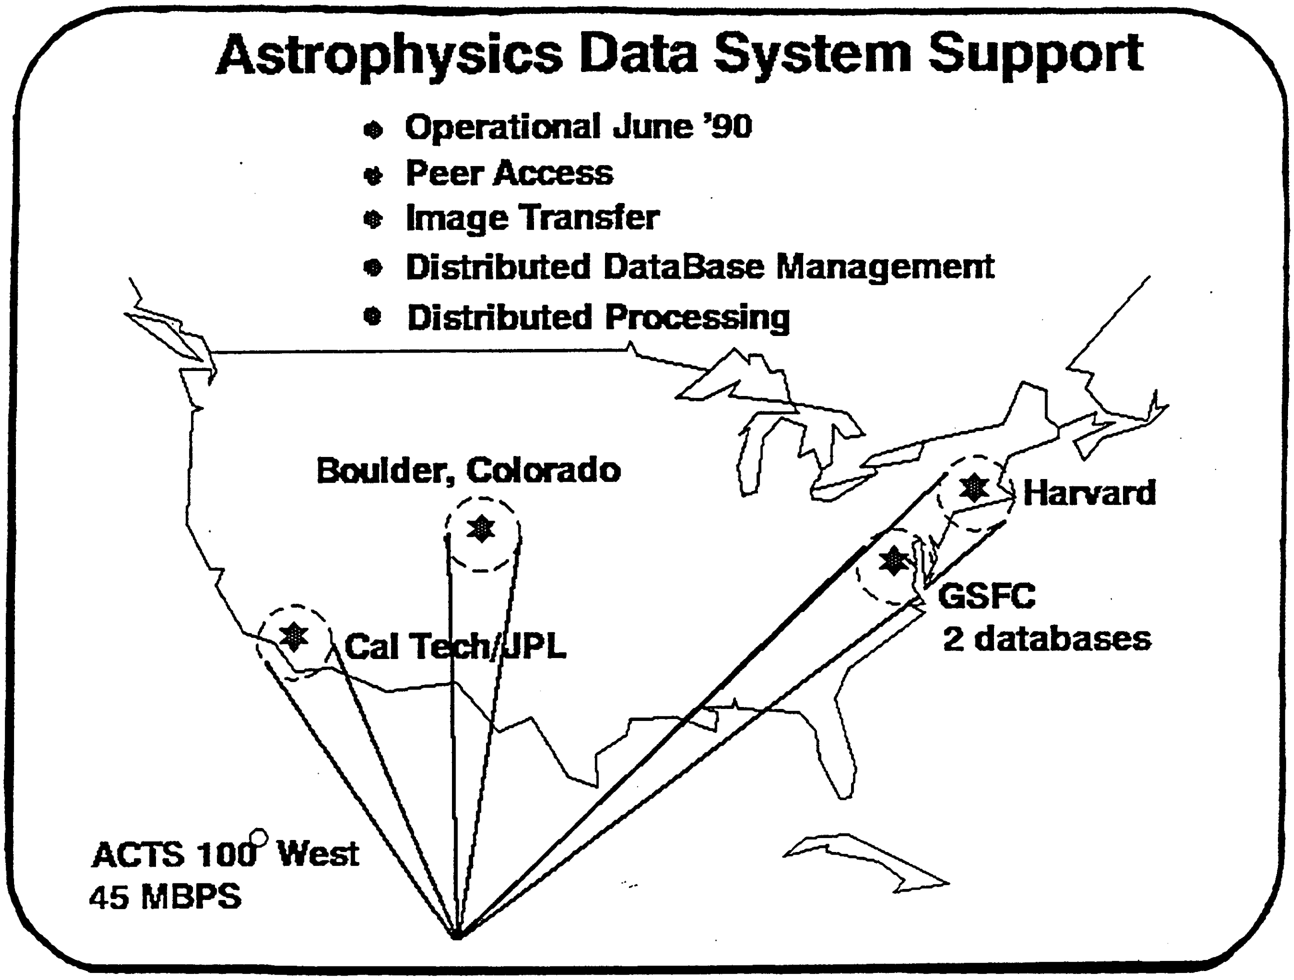
Astrophysics Data System Support
Operational June '90
Peer Access
Image Transfer
Distributed DataBase Management
Distributed Processing
Boulder, Colorado
Harvard
GSFC
2 databases
Cal Tech/JPL
ACTS 100° West
45 MBPS

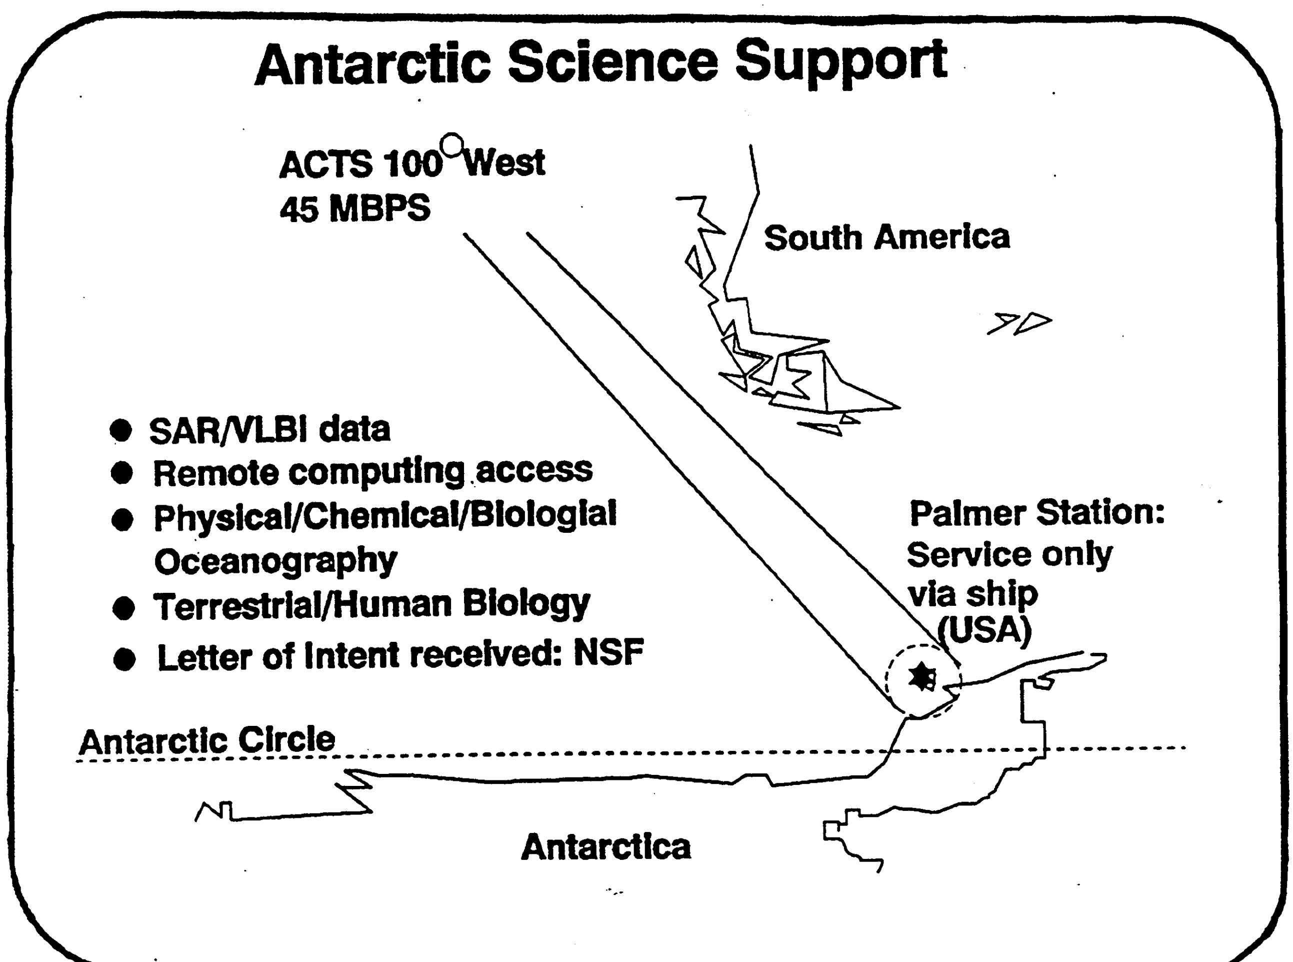
Antarctic Science Support
ACTS 100° West
45 MBPS
South America
SAR/VLBI data
Remote computing access
Physical/Chemical/Biologial Oceanography
Terrestrial/Human Biology
Letter of Intent received: NSF
Palmer Station:
Service only
via ship
(USA)
Antarctic Circle
Antarctica

Science Disciplines in Antarctica

- Upper Atmoshpere Physics
  Sensor data, real time computing
- Solid Earth Geophysics
  NSF Continental Lithosphere Program
- Oceanography
  Voice communications outside Antarctica
- Meteorology
  Image Data Access
- Glaciology
  Data transmission
- Astronomy & Astrophysics
  Infrared and microwave radio astronomy
- Geodesy
  -

Buoy Data Collection

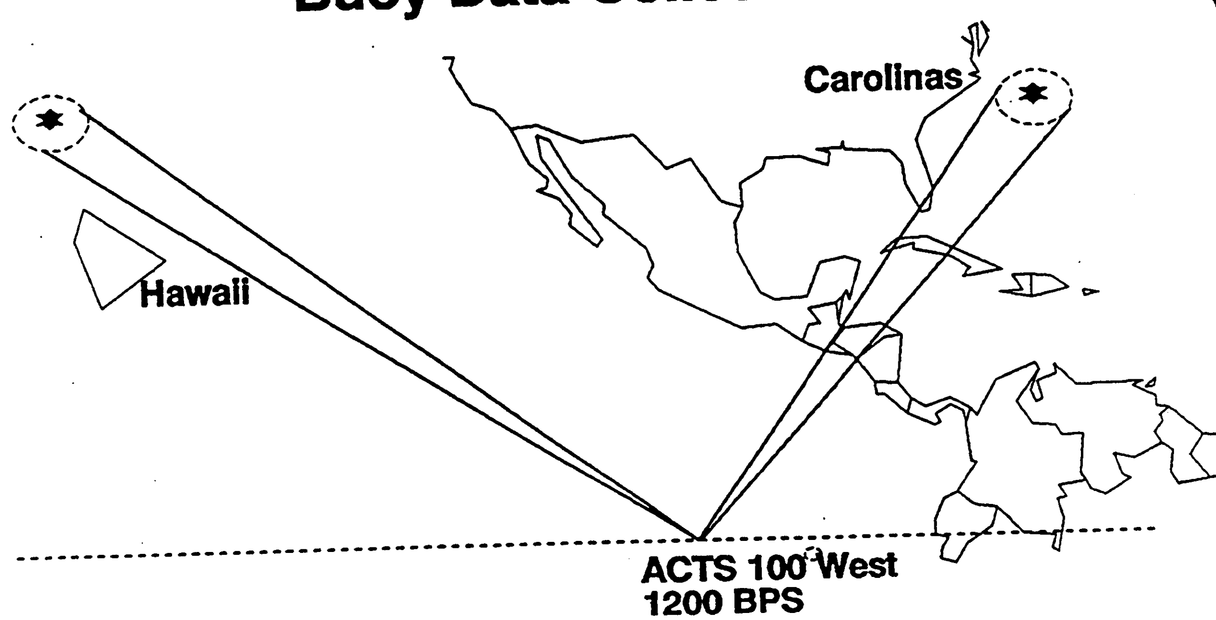

- 1 MBit/Buoy/Day (Approximately)
- SeaWIFS Calibration/Validation
- Unattended Operation
- Extended buoy life
- Proof of Concept

Real-Time Distribution

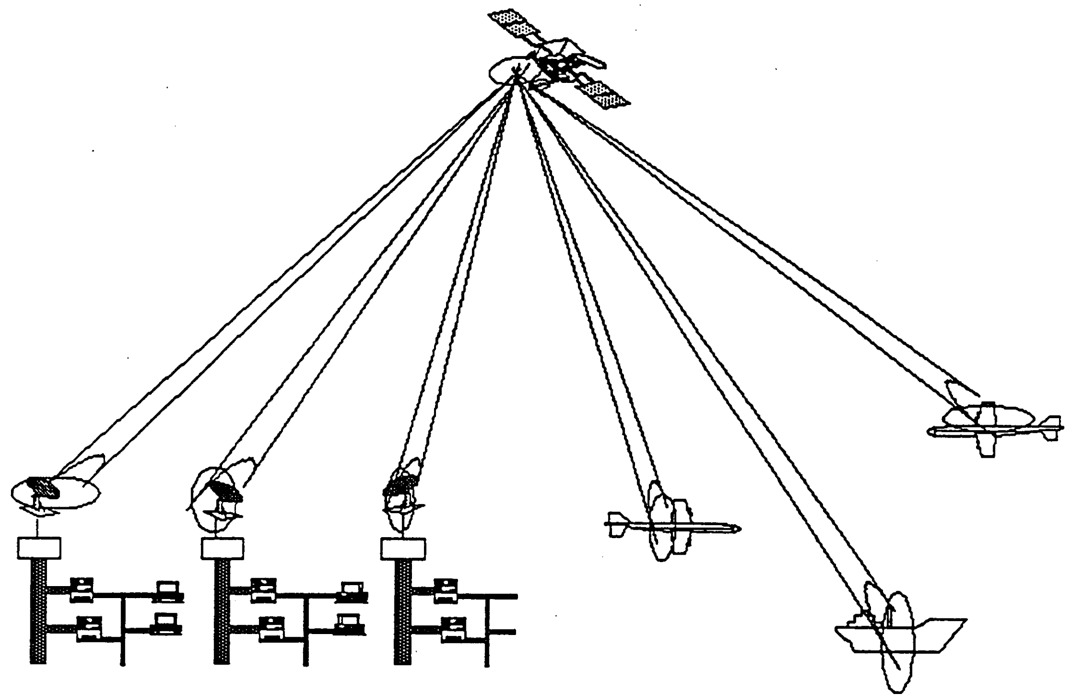

Measurement Validation Support

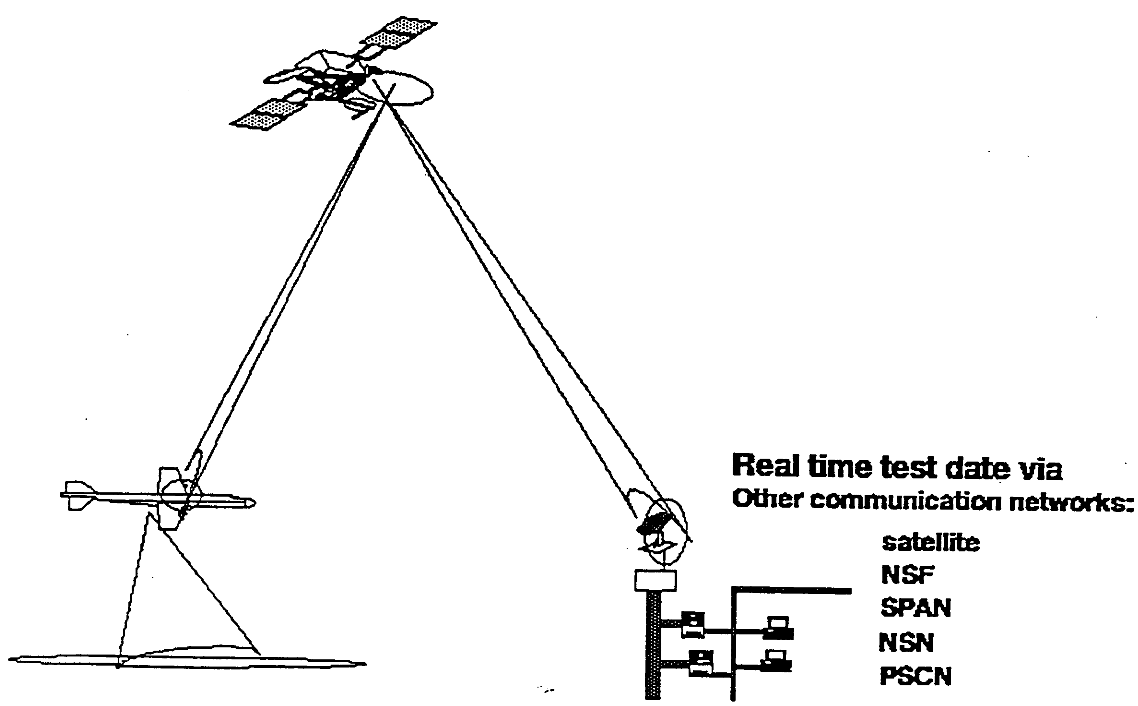

ACTS HDTV APPLICATION AREAS

NETWORK/CABLE Television

Network Feeds
Studio Links
Special Events/HDTV Theatre
Broadcast to Homes
Regional Broadcast Using Spot Beams

COMMERCIAL/INDUSTRIAL

Remote Monitoring
Remote Surveillance
Training/Seminars
Video Conference
CAD/CAM

GRAPHIC ARTS APPLICATIONS

POP Advertising
CAD

MEDICAL

Remote Diagnosis
Surgical Video
Training/Seminars
Video Conference
Remote Diagnosis
Radiological Imaging (PACS)
Remote Picture Storage (PACS)

SCIENTIFIC

Remote Monitoring
Video Conference
Training/Seminars
Image Analysis
High Speed Video

ACTS

5/8/90
RAB APPLHD2

WHY HDTV OVER ACTS?

CONGRESSIONAL DIRECTION

- ACTS VIEWED AS AN IDEAL TESTBED FOR HDTV DISTRIBUTION MEDIUM

WIDE BANDWIDTH AT NEXT HIGHER FREQUENCY BAND

- T1 AND T3 CAPABILITY WITH BASEBAND PROCESSOR
- MICROWAVE SWITCH MATRIX OFFERS 900 MHz CLEAR CHANNEL
- DIGITAL OR ANALOG
- POSSIBLE FUTURE DBS ALLOCATION

TESTBED FOR FUTURE DISTRIBUTION ALTERNATIVE

- SPOTBEAMS MAY MODEL A REGIONAL BEAM
- TEST POSSIBILITY OF VERY FINE SPOTBEAMS FOR EXCLUSIVE DISTRIBUTION
- DIGITAL TRANSMISSION MAY BE EXPLORED

ACTS

NASA

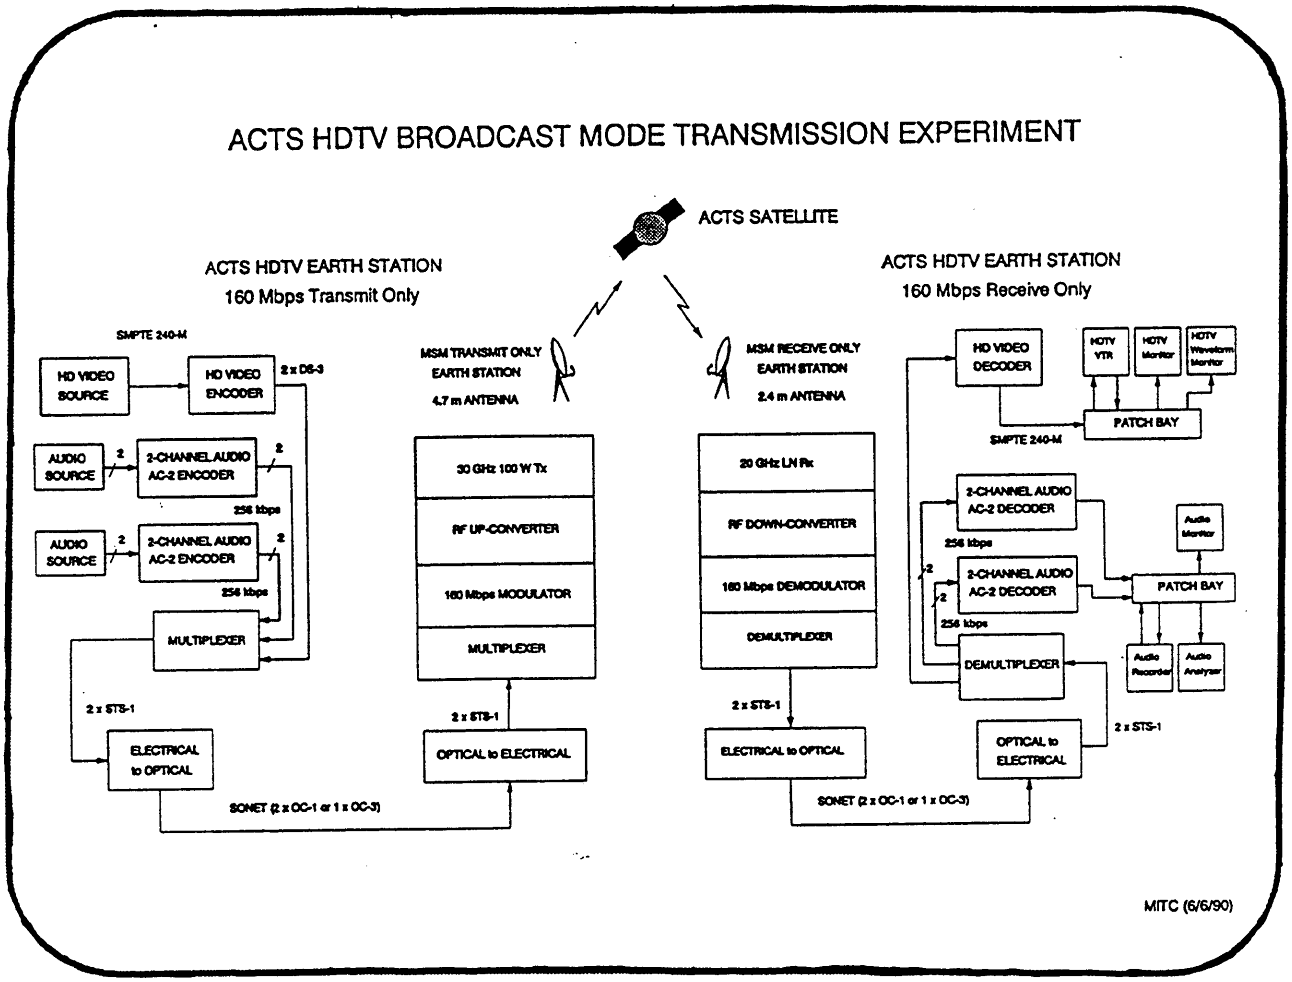

ACTS HDTV BROADCAST MODE TRANSMISSION EXPERIMENT
ACTS SATELLITE
ACTS HDTV EARTH STATION
160 Mbps Transmit Only
SMPTE 240-M
HD VIDEO SOURCE
HD VIDEO ENCODER
2 x DS-3
AUDIO SOURCE
2
2-CHANNEL AUDIO AC-2 ENCODER
2
256 kbps
AUDIO SOURCE
2
2-CHANNEL AUDIO AC-2 ENCODER
2
256 kbps
MULTIPLEXER
2 x STS-1
ELECTRICAL to OPTICAL
SONET (2 x OC-1 or 1 x OC-3)
MSM TRANSMIT ONLY EARTH STATION
4.7 m ANTENNA
30 GHz 100 W Tx
RF UP-CONVERTER
160 Mbps MODULATOR
MULTIPLEXER
2 x STS-1
OPTICAL to ELECTRICAL
MSM RECEIVE ONLY EARTH STATION
2.4 m ANTENNA
20 GHz LN Rx
RF DOWN-CONVERTER
160 Mbps DEMODULATOR
DEMULTIPLEXER
2 x STS-1
ELECTRICAL to OPTICAL
SONET (2 x OC-1 or 1 x OC-3)
ACTS HDTV EARTH STATION
160 Mbps Receive Only
HD VIDEO DECODER
HDTV VTR
HDTV Monitor
HDTV Waveform Monitor
PATCH BAY
SMPTE 240-M
2-CHANNEL AUDIO AC-2 DECODER
256 kbps
2
2-CHANNEL AUDIO AC-2 DECODER
2
256 kbps
DEMULTIPLEXER
PATCH BAY
Audio Monitor
Audio Recorder
Audio Analyzer
OPTICAL to ELECTRICAL
2 x STS-1
MITC (6/6/90)

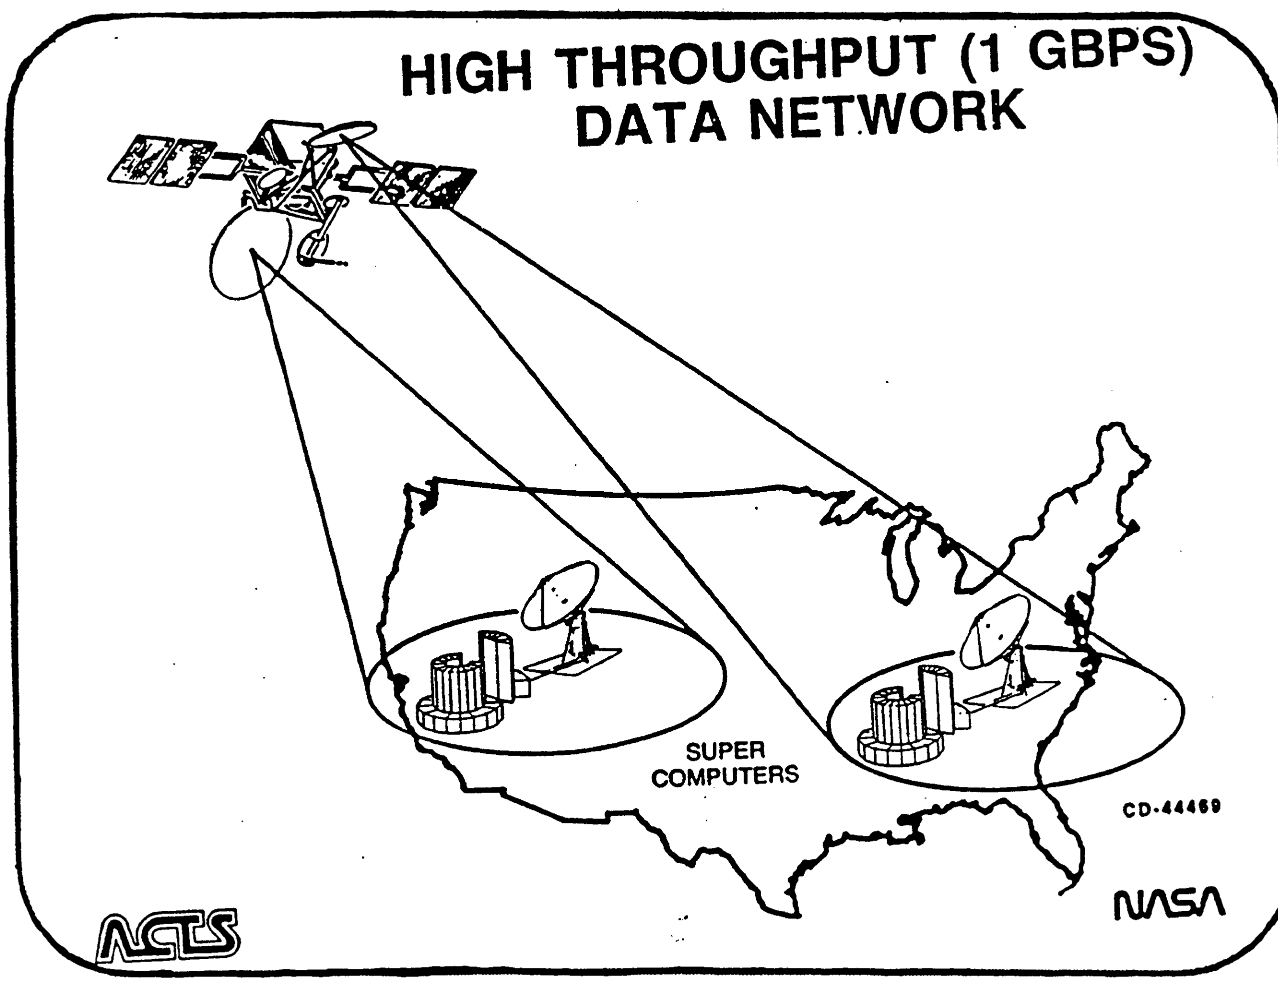

HIGH THROUGHPUT (1 GBPS)
DATA NETWORK
SUPER COMPUTERS
CD-44469
ACTS
NASA

SUPERCOMPUTING APPLICATIONS

ADVANTAGES OF ACTS:

- CAPACITY
  - 900 MHZ CHANNELS ALLOWS HIGH DATA RATE (GIGABIT PER SECOND) COMMUNICATIONS THROUGHPUT
- SWITCHING
  - MICROWAVE SWITCH MATRIX ALLOWS SATELLITE TO INTERCONNECT MULTIPLE NODES WITH HIGH DATA RATE THROUGHPUT
- CONNECTIVITY
  - SATELLITE PROVIDES ACCESS TO USERS ANYWHERE WITHIN THE SATELLITE ANTENNA COVERAGE FOOTPRINT (ESPECIALLY TO "REMOTE" AREAS OFF FIBER BACKBONE)
- COST
  - COST BILLED ON USAGE, USAGE ON DEMAND. OPERATIONAL SATELLITE SYSTEM STUDIES FOR GIGABIT PER SECOND LINKS SHOW COST ADVANTAGES.

ACTS NASA

DARPA BROAD AGENCY ANNOUNCEMENT (BAA)

- Objective: Solicit proposals for various aspects of networking in support of DARPA Basic Research in Experiments with Very High Speed Gigabit and Terabit Networks Including:
  - Innovative Networking Experiments for Using the High Capacity Channels of ACTS
  - Short-term Design Studies of Ground Terminal Technology to Support ACTS Experiments
- Projected Funding: $2.5 Million
- Schedule:

CBD Notice (Constitutes BAA)	May 15, 1990 (C)
Abstracts Due	June 15, 1990 (C)
Proposals Due	August 17, 1990
DARPA Responses	Within 30 days of receipt

EXPERIMENT DEVELOPMENT

ACADEMIA

- FUNDING: GOVERNMENT & INDUSTRY
- PURPOSE: GENERAL COMMUNICATIONS R&D
- COMMITMENT: NRA PROPOSAL

INDUSTRY

- FUNDING: INDUSTRY
- PURPOSE: EVALUATE TECHNOLOGY FOR FUTURE OPERATIONAL SYSTEMS
- COMMITMENT: EOA PROPOSAL

GOVERNMENT

- FUNDING: AGENCY
- PURPOSE: QUASI-OPERATIONAL EVALUATION, R&D, EVALUATE TECHNOLOGY FOR FUTURE OPERATIONAL SYSTEMS
- COMMITMENT: EOA PROPOSAL
  DEVELOP MOU'S

CD-45619

ACTS EXPERIMENTS PROGRAM

- **NASA RESEARCH ANNOUNCEMENT (NRA)**

  FUNDING PROVIDED BY NASA TO UNIVERSITIES AND OTHER ORGANIZATIONS WITH LIMITED DISCRETIONARY MONIES.

  RELEASE DATE: MARCH 5, 1990
  PROPOSALS DUE: JUNE 4, 1990

  EVALUATION AND SELECTION IN PROCESS

- **EXPERIMENT OPPORTUNITY ANNOUNCEMENT (EOA)**

  FUNDING PROVIDED BY SUBMITTING ORGANIZATION (INDUSTRY, MILITARY AND GOVERNMENT ORGANIZATIONS, INCLUDING NASA)

  RELEASE DATE: FALL, 1990
  PROPOSALS DUE: EARLY, 1991

ACTS EXPERIMENTS PROGRAM

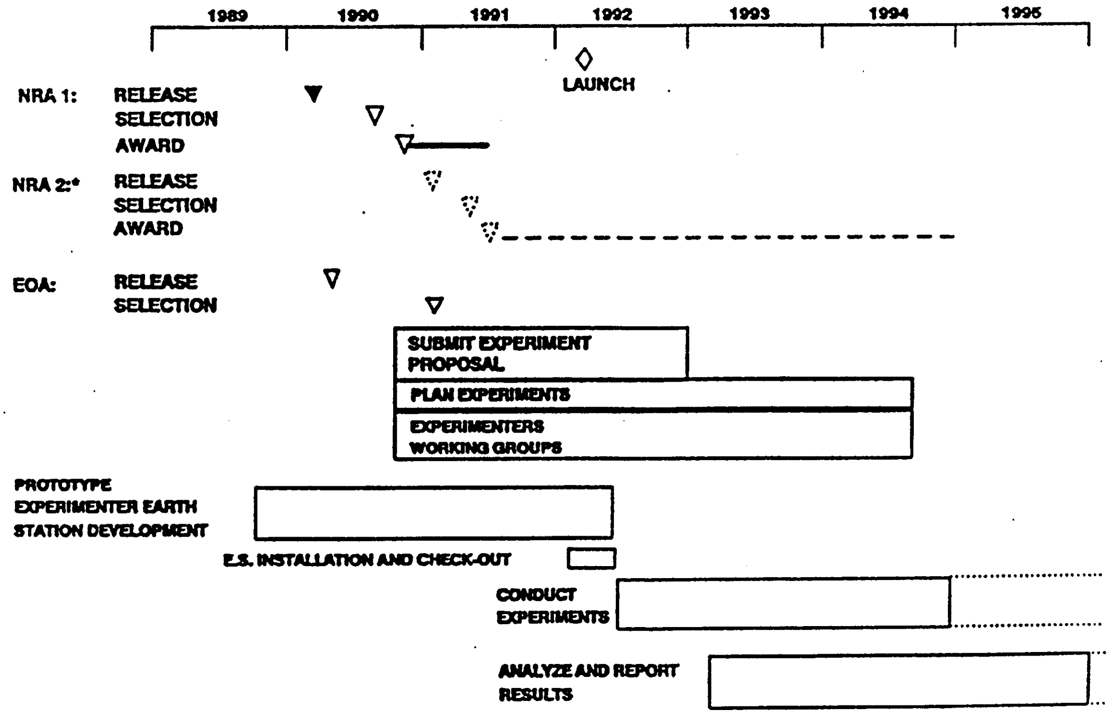

NRA-NASA RESEARCH ANNOUNCEMENT
EOA-EXPERIMENT OPPORTUNITY ANNOUNCEMENT

ACTS

* PLANNED (DEPENDENT ON EXPT. FUNDING AUGMENTATION)

EXPERIMENT OPPORTUNITY ANNOUNCEMENT (EOA)

☐ PURPOSE: TO OBTAIN FROM POTENTIAL EXPERIMENTERS A PROPOSAL TO PERFORM AN ACTS EXPERIMENT

☐ KEY ELEMENTS OF EOA SUBMITTAL:

- COVER LETTER–CERTIFICATION OF ORGANIZATIONAL SUPPORT AND SPONSORSHIP
- IDENTIFYING INFORMATION–TITLE, INVESTIGATORS, ORGANIZATION
- SUMMARY OF EXPERIMENT
- EXPERIMENT OBJECTIVES
- APPROACH
- EXPERIMENT PLAN–INCLUDING NUMBER AND LOCATION OF GROUND TERMINALS;
  - –PARAMETERS TO BE MEASURED;
  - –DATA REQUIRED FROM NASA;
  - –ANTICIPATED SCHEDULE;
  - –EXPERIMENT DURATION
- RESOURCES–PROVIDED BY EXPERIMENTER
  - –REQUIRED FROM NASA
- ANTICIPATED RESULTS

CD-41314

ACTS NASA

ACTS EXPERIMENT PROGRAM

NASA WILL PROVIDE:
- SPACECRAFT TIME DURING EXPERIMENT PERIOD
- MASTER CONTROL STATION OPERATIONS
- EXPERIMENT PROGRAM MANAGEMENT
- NASA GROUND STATION FOR EXPERIMENTER USE
- DATA MEASUREMENT ABOARD SPACECRAFT AND AT MASTER CONTROL STATION

NASA WILL ASSIST EXPERIMENTERS IN:
- EXPERIMENTERS PLANNING AND DESIGN
- DEVELOPING OR PURCHASING THEIR OWN GROUND TERMINALS
- UTILIZING NASA GROUND STATION OR OTHER ORGANIZATIONS' GROUND TERMINALS

EXPERIMENTER PROVIDES:
- EXPERIMENT PLAN
- RESOURCES TO CONDUCT EXPERIMENTS (INCLUDING THAT FOR EARTH STATIONS)
- EXPERIMENTER WORKING GROUP SUPPORT
- CONDUCT OF EXPERIMENTAL INVESTIGATION
- ANALYSIS OF EXPERIMENT RESULT

ACTS

NASA

FOR INFORMATION CONTACT:

RONALD J. SCHERTLER
ACTS PROJECT EXPERIMENTS MANAGER
NASA LEWIS RESEARCH CENTER
21000 BROOKPARK RD., MS 54-6
CLEVELAND, OH 44135
(216) 433-3527

TOM VONDEAK
NASA LEWIS RESEARCH CENTER
21000 BROOKPARK RD., MS 54-6
CLEVELAND, OH 44135
(216) 433-3277
EMAIL: TVONDEAK@NASAMAIL.NASA.GOV

ACTS

NASA

Appendix A

Attendees

Stephen Adams
Digital Equipment Corporation
550 King Street
M/S LKG2-1/N2
Littleton, MA 01460
(W)(508) 486-7521
decwrl::"adams@zeppo"

Nick Alfano
Gandalf Data Limited
130 Colonnade Road S
Ottawa, ON K2E 7M4
Canada
(W)(613) 723-6500
nick@gandalf.ca

Guy Almes
Rice University
Department of Computer Science
P.O. Box 1892
Houston, TX 77251-1892
(W)(713)527-6038
almes@rice.edu

Philip Almquist
Consultant
214 Cole Street
San Francisco, CA 94117
(W)415-752-2427
almquist@jessica.stanford.edu

William Anderson
Mitre
Burlington Road
M/S EO66
Bedford, MA 01730
(W)(617)271-3388
wda@mitre-bedford.org

Cathy Aronson
Lawrence Livermore National Laboratory
East Avenue L 561
P.O. Box 808
Livermore, CA 94550
(W)(415)422-4016
cja@marmot.nersc.gov

Karl Auerbach
Epilogue Technology Corporation
P.O. Box 217
Capitola, CA 95010-0217
(W)(408) 426-8786
auerbach@csl.sri.com

Zorica Avramovic
SAIC
8619 Westwood Center Drive
Vienna, VA 22180
(W)(703) 448-1683
zorica@sparta.com

Douglas Bagnall
Hewlett-Packard
Apollo Division
330 Billerica Road
M/S CHR 03 DC
Chelmsford, MA 01824
(W)(508)256-6600
bagnall_d@apollo.hp.com

Fred Baker
Advanced Computer Communications
6607 Kaiser Drive
Fremont, CA 94555
(W)415-794-1100
fbaker@acc.com

David Balenson
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)(301) 854-5358
balenson@tis.com

Richard Basch
Massachusetts Institute of Technology
One Amherst Street
M/S E40-342C
Cambridge, MA 02139
(W)617-253-0100
probe@mit.edu

Amatzia Ben-Artzi
SynOptics Communications
501 E. Middlefield Road
M/S B3
Mountain View, CA 94043
(W)415-960-1100
amatzia@synoptics.com

Art Berggreen
Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
art@opal.acc.com

Jonathan Biggar
NetLabs
11693 San Vicente Boulevard
Suite 348
Los Angeles, CA 90049
(W)(213) 824-2500
jon@netlabs.com

Yvonne Biggar
Unisys
5151 Camino Ruiz
M/S 02-C210
Camarillo, CA 93065
(W)(805) 526-1606
yvonne@cam.unisys.com

Chet Birger
BBN Communications
150 Cambridge Park Drive
Cambridge, MA 02140
(W)(617)873-2676
cbirger@bbn.com

Roger Boehner
NCR NPD
2700 Snelling Avenue North
M/S 010
St. Paul, MN 55126
(W)(612) 638-7274
Roger.Boehner@StPaul.NCR.COM

David Borman
Cray Research
1440 Northland Drive
Mendota Heights, MN 55120
(W)(612)681-3398
dab@opus.cray.com

Scott Bradner
Harvard University
William James Hall 1232
33 Kirkland Street
Cambridge, MA 02138
(W)(617)495-3864
sob@harvard.harvard.edu

Larry Brandt
Fibronics International
Spartacus Group
1 Lowell Research Center
Lowell, MA 01852
(W)(508) 937-1600
lbrandt@sparta.com

David Brent
University of British Columbia
Computing Centre
Vancouver, BC V6T 1W5
Canada
(W)(604) 228-5736
brent@staff.ucs.ubc.ca

Caralyn Brown
Prime Computer
500 Old Connecticut Path
M/S 10-10
Framingham, MA
(W)(508) 879-2960
cbrown@ENR.Prime.com

Jack Brown
US Army
ASQB-SEP-C, Building 31043
CDR USAISEC
Fort Huachuca, AZ 85613
(W)(602) 533-2873
jbrown@huachuca-emh8.army.mil

L Allyson Brown
U.S. Department of Defense
9800 Savage Road
M/S C322
Ft. George G. Meade, MD 20755
(W)(301) 684-7302
allyson@umd5.umd.edu

Eric Brunner
Tule Network Services
801 Church Street
Suite 1327
Mountain View, CA 94041
(W)(415) 965-7710
brunner@monet.berkeley.edu

Theodore Brunner
Bellcore
445 South Street
M/S MRE 2P252
Morristown, NJ 07960
(W)201-829-4678
tob@thumper.bellcore.com

Philip Budne
Shiva
155 Second Street
Cambridge, MA 02141
(W)(617) 384-8500
phil@shiva.com

Jeffrey Buffum
Apollo Computer
330 Apollo Drive
M/S CHD 02 DN
Chelmsford, MA 01824
(W)(508) 256-6600
jbuffum@apollo.hp.com

Jeffrey Burgan
NASA Ames
M/S 233-8
Moffett Field, CA 94035
(W)415-604-5705
jeff@nsipo.nasa.gov

Ross Callon
Digital Equipment Corporation
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-486-5009
callon@bigfut.enet.dec.com

C. Allan Cargille
University of Wisconsin
Computer Sciences Department
1210 W. Dayton Street
Madison, WI 53706
(W)615-974-0822
cargille@cs.wisc.edu

Eric Carroll
University of Toronto
Computing Services
4 Bancroft Avenue, Room 102
Toronto, Ontario M5S 1C1
Canada
(W)(416) 978 6134
eric@utcs.utoronto.ca

Jeffrey Case
University of Tennessee
Department of Computer Science
107 Ayres Hall
Knoxville, TN 37996
(W)615-974-8992
case@utkux1.utk.edu

Stephen Casner
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
(W)(213) 822-1511
casner@venera.isi.edu

Vinton Cerf
Corp. for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)(703) 620-8990
vcerf@NRI.Reston.VA.US

Isaac Chan
Consumers Software
73 Water Street
Suite 603
Vancouver, BC V6B 1A1
Canada
(W)(604)688-4548
isaac@gui.consumers.bc.ca

Stephen Chan
SFU LCCR
San Francisco University, Burnaby
Burnaby, BC V5A 1S6
Canada
(W)(604)291-4399
chan@cs.sfu.ca

Asheem Chandna
AT&T Bell Laboratories
480 Red Hill Road
Room 2K058
Middletown, NJ 07748
(W)(201)615-5819
ac0@mtuxo.att.com

Cho Chang
Hewlett-Packard/Apollo
250 Apollo Drive
Chelmsford, MA 01824
(W)(508)256-6600
chang_c@apollo.hp.com

Samuel Chanson
University of British Columbia
Department of Computer Science
Vancouver, BC V6T 1W5
Canada
(W)604/ 228-6667
chanson@cs.ubc.ca

A. Lyman Chapin
BBN Communications Corporation
150 Cambridge Park Drive
Cambridge, MA 02140
(W)(617) 873-3133
Lyman@merit.edu

Brett Chappell
NSWC
Code N 35
Dahlgran, VA 22448-5000
(W)(703)663-1571
behappe@relay.nswc.navy.mil

Andrew Cherenson
Silicon Graphics
2011 N. Shoreline Boulevard
PO Box 7311
M/S 9U-510
Mountain View, CA 94039-7311
(W)415-962-3486
arc@sgi.com

J. Noel Chiappa
Consultant
708 E. Woodland
Grafton, VA 23692
(H)(804) 898-8183
jnc@ptt.lcs.mit.edu

Chris Chiotasso
Spartacus/A Fibronics Co.
One Lowell Research Center
847 Rogers Street
Lowell, MA 01852
(W)(508)937-1600
chris@sparta.com

Cyrus Chow
NASA AMES
M/S 233 18
Moffett Field, CA 94035
(W)415/604-6843
cchow@orion.arc.nasa.go

Anthony Chung
Hughes LAN
1225 Charleston Road
M/S 7
Mountian View, CA 94043
(W)(415)966-7430
anthony@hls.com

Paul Ciarfella
Digital Equipment Corporation
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)(508)952-3548
ciarfella@levers.enet.dec.com

George Clapp
Ameritech Services
Gould Center, Building 40
2850 Golf Road
Rolling Meadows, IL 60008
(W)708-806-8318
meritec!clapp@bellcore.bellcore.com

Graeme Clark
University of British Columbia
Department of Computer Science
Vancouver, BC V6W 1Z7
Canada
clark@cs.ubc.ca

Richard Colella
National Institute of Standards and Technology
Building 225
Room B217
Gaithersburg, MD 20899
(W)301-975-3627
colella@osi3.ncsl.nist.gov

Rob Coltun
Wellfleet
510 New Mark Esplanade
Rockville, MD 20850
(W)301-454-0863
rcoltun@trantor.umd.edu

George Conant
Xyplex
Engineering
330 Codman Hill Road
Boxborough, MA 01719
(W)508/ 264-9900
geconant@eng.zyplex.com

John Cook
Chipcom
Southborough Office Park
118 Turnpike Road
Southborough, MA 01772
(W)508-460-8900
cook@chipcom.com

Curtis Cox
Department of Defense
NARDAC
Code 3031, Bldg 196
Washington Navy Yard
Washington, DC 20374-1435
(W)202-433-4026
zk0001@nhis.navy.mil

Tracy Cox
Bellcore
331 Newman Springs Rd.
Redbank, NJ 07701
(W)201/ 758-2107
tacox@sabre.bellcore.com

Caroline Cranfill
Bell South Services
675 W. Peachtree St.
M/S 41B50 SBC
Atlanta, GA 30375
(W)404-420-8432
rcc@bss.com

James Davin
Massachusetts Institute of Technology
Computer Science Lab, NE43-507
545 Technology Square
Cambridge, MA 02139
(W)617-253-6020
jrd@ptt.lcs.mit.edu

Mark Crispin
6158 Lariat Loop NE
Bainbridge Island, WA 988110-2020
(W)206/842-2385
mrc@cac.washington.edu

Farokh Deboo
Interlink Computer Science
47370 Fremont Boulevard
Fremont, CA 94538
(W)415-657-9800
fjd@interlink.com

Dave Crocker
Digital Equipment Corporation
Network Systems Lab
505 Hamilton Avenue
Palo Alto, CA 94301
(W)415-688-1320
dcrocker@nsl.dec.com

Steve Deering
Xerox PARC
1017 Mallet Court
Menlo Park, CA 94025
(W)415-321-0224
deering@pescadero.stanford.edu

Steve Crocker
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
crocker@tis.com

John Demco
University of British Columbia
6356 Agricultural Road
Vancouver, BC V6T 1W5
Canada
(W)604-228-6537
demco@ean.ubc.ca

Megan Davies
Corp. for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)(703) 620-8990
mdavies@nri.reston.va.us

Nick Di Iorio
AT&T Bell labs
Crowfords Corner Road
Room 4K-325
Holmdel, NJ 07733
(W)201-949-8741
nicola@mapoli.att.com

Andris Dindzans
Newbridge Networks
600 March Road
P.O. Box 13600
Kananta, ON K2K 2E6
Canada
(W)613-591-3600

Ralph Droms
Bucknell University
Computer Science Department
323 Dana Engineering
Lewisburg, PA 17837
(W)717-524-1145
droms@sol.bucknell.edu

Martin Dubetz
Washington University
Campus Box 1048
One Brookings Drive
St. Louis, MO 63130
(W)314-726-7389
dubetz@wugate.wustl.edu

Nadya El-Afandi
Network Systems
7600 Boone Avenue, N.
Brooklyn Park, NJ 55428
(W)612-424-4888
nadya@network.com

Robert Elz
University of Melbourne
Computer Science
Parkville, VIC 3161 AUSTRALIA
(W)+6133445225
kre@munnari.oz.au

Robert Enger
Contel Corporation
P.O. Box 10814
Chantilly, VA 22021-0814
(W)703-818-5555
enger@sccgate.scc.com

Mike Erlinger
Micro Technology
5065 East Hunter Ave
Anaheim, CA 92807
(W)714-970-0300
mike@mti.com

Tom Evans
Webster Computer Corporation
2109 O'Toole Avenue, Suite J
San Jose, CA 95131-1303
(W)408-954-8054
wcc@cup.portal.com

Kevin Fall
University of California, Berkeley
571 Evans Hall
Berkeley, CA 94720
(W)415-642-3979
kfall@Berkeley.EDU

Dino Farinacci
3Com Corporation
2081 N. Shoreline Boulevard
Mountain View, CA 94043
(W)415-940-7661
dino@buckeye.esd.3com.com

Alf Farnham
Boeing Computer Services
P.O. Box 24346
M/S 6H-22
Seattle, WA 98124-0346
(W)206-234-1382
carolf@mcescher.unl.edu

Carol Farnham
University of Nebraska
Lincoln, NE 68588-0496
(W)402-472-5878
carolf@mcescher.unl.edu

Dennis Ferguson
University of Toronto
5 King's College Road
Toronto, ONTARIO M5S 1A4
Canada
(W)416-978-2455
dennis@gw.ccie.utoronto.ca

Michael Fidler
IRCC, Ohio State University
1971 Neil Avenue, Rm 406
Columbus, OH 43210-1210
(W)(612) 292-4843
ts0026@ohstvma.ircc.ohio-state.edu

Dale Finkelson
University of Nebraska-Lincoln
29 Walter Scott Engr Ctr
MIDnet
Lincoln, NE 68588-0534
(W)402-472-5032
dmf@westie.unl.edu

Jeffrey Fitzgerald
Fibercom, Inc.
P.O. Box 11966
Roanoke, VA 24022
(W)703-342-6700
jjf@fibercom.com

James Forster
cisco Systems
1525 O'Brian Drive
Menlo Park, CA 94025
forster@cisco.com

Craig Fox
Network Systems
7600 Boone Avenue North
Minneapolis, MN 55428
(W)612 424-4888
foxcj@nsco.network.com

Richard Fox
SynOptics Communications
950 Linden Ave. #208
Sunnyvale, CA 94086
(W)408-691-7272
sytek!rfox@sun.com

Karen Frisa
Novell
1340 Treat Boulevard, Suite 500
Walnut Creek, CA 94596
(W)415-947-0998
karen@kinetics.com

Stanley Froyd
Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
sfroyd@salt.acc.com

Robert Gilligan
Sun Microsystems
2550 Garcia Avenue
Mountain View, CA 94043
(W)415 336-1012
gilligan@sun.com

Vince Fuller
Stanford University
Networking & Communication
115 Pine Hall
Stanford, CA 94305-4122
(W)415-723-6860
fuller@jessica.stanford.edu

Jim Goetz
AT&T Bell Laboratories
5 Lowell Court
M/S HRIF213
Freehold, NJ 07728
(W)201-615-5598
goetz@hav2d.att.com

James Galvin
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
galvin@tis.com

Hellmut Golde
University of Washington
Dept of CS and Engineering
M/S FR-35
Seattle, WA 98195
(W)(206)543-6175
golde@june.cs.washington.edu

Ella Gardner
Mitre Corporation
7525 Colshire Drive
M/S W425
McLean, VA 22102-3481
(W)703-883-5826
epg@gateway.mitre.org

Terry Gray
University of Washington
3737 Brooklyn Avenue NE
M/S HG-45
Seattle, WA 98105
(W)(206) 543-3880
gray@cac.washingtom.edu

Eugene Geer
Bellcore
331 Newman Springs Road
M/S NUC1H310
Red Bank, NJ 07701
(W)201-758-2149
bcr!nvmxr!ewg

Larry Green
Protocol Engines
1900 State Street
Santa Barbara, CA 93101
(W)805-965-0825
green@pei.com

Michael Grobe
University of Kansas
Academic Computing Services
Lawrence, KS 66045
(W)(913) 864-0100
grobe@kuhub.cc.ukans.edu

Martin Gross
Defense Communications Agency
1860 Wiehle Avenue
Code DRFE
Reston, VA 22090-5500
(W)703-437-2165
gross@polaris.dca.mil

Phill Gross
Corp. for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)703-620-8990
pgross@nri.reston.va.us

Olafur Gudmundsson
University of Maryland
Dept. of Computer Science
College Park, MD 20742
(W)(301) 454-6497
ogud@cs.umd.edu

Chris Gunner
Digital Equipment Corporation
P.O. Box 121
Worton Grange, Imperial Way
Reading, RG2 0TU
GB
(W)44-734-853516
gunner@osicwg.enet.dec.com

Yong Guo
University of British Columbia
Computer Science Department
Vancouver, BC V6T 1W5
Canada
(W)(604) 228-5401
guo@cs.ubc.ca

Robert Hagens
University of Wisconsin
Computer Science Dept.
1210 West Dayton Street
Madison, WI 53706
(W)608-262-1017
hagens@cs.wisc.edu

Tony Hain
Lawrence Livermore National Laboratory
PO Box 5509
Livermore, CA 94550
(W)415-422-4017
alh@eagle.es.net

Neil Haller
Bellcore
445 South Street
Morristown, NJ 07960-1910
(W)201-829-4478
nmh@bellcore.com

Brian Handspicker
Digital Equipment Corporation
550 King Street
LKG1-2/E19
Littleton, MA 01460
(W)508-486-7894
bd@vines.enet.dec.com

Susan Hares
Merit Computer Network
1075 Beal Ave, NDSB
NDSB
Ann Arbor, MI 48109-2112
(W)313 936 2095
skh@merit.edu

Peter Harrison
MIDEN Pacific
16 National Circuit
Barton, ACT 2600
ARLA
(W)61-6-271-2222
harrison@miden.ucs.unimelb.edu.au

Richard Hart
Digital Equipment Corporation
110 Spit Brook Raod
ZKO3-3/W20
Nashua, NH 03062-2698
(W)603-881-0418
hart@decvax.dec.com

Peter Hayden
Digital Equipment Corporation
153 Taylor Street
Littleton, MA 1460-1407
(W)(508) 952-3547
hayden@levers.enet.dec.com

Charles Hedrick
Rutgers University
Hill Center, POB 879
CCIS
Piscataway, NJ 08854
(W)(201) 932-3088
hedrick@aramis.rutgers.edu

William Hiles
NSWC
Code N-35
Dahlgran, VA 22448-5000
(W)(703) 663-1571
whiles@realy.nswc.navy.mil

Robert Hinden
BBN Communications
50 Moulton Street
Cambridge, MA 02138
(W)617-873-3757
hinden@bbn.com

Don Hirsh
Meridian Technology
15965 Manchester Rd.
P.O. Box 2006
St. Louis, MO 63011
(W)314-394-1600
hirsh@magic.meridianpc.com

Russell Hobby
University of California, Davis
Computing Services
Surge II - Room 1400
Davis, CA 95616
(W)916-752-0236
rdhobby@ucdavis.edu

J. Paul Holbrook
Carnegie Mellon University
CERT-SEI
4500 Fifth Ave.
Pittsburgh, PA 15213-3890
(W)412-268-7720
ph@sei.cmu.edu

Greg Hollingsworth
Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20723
(W)301-953-6065
gregh@mailer.jhuapl.edu

Terrence Holm
MPR Teltech
8999 Nelson Way
M/S 23A
Burnaby, BC V5A 4B5
Canada
(W)(604) 293-6079

Michael Hrybyk
EDUCOM
1112 16th Street NW
Washington, DC 20036
(W)(202) 872-4200
mwh@educom.edu

Ruei-Hsin Hsiao
Digital Equipment Corporation
550 King Street
LKG2-2/P5
Littleton, MA 01460-1289
(W)(508) 486-7628
nac::hsiao

Kathleen Huber
BBN Communications
50 Moulton Street
Cambridge, MA 02138
(W)617-873-2520
khuber@bbn.com

Steven Hubert
University of Washington
Networks & Dist Computing
M/S HG-45
Seattle, WA 98195
(W)206-543-6384
hubert@cac.washington.edu

Erik Huizer
SURFnet
P.O. Box 19035
3501 DA Utrecht
NETHERLANDS
(W)+31 30 310290
huizer@surfnet.nl

Steven Hunter
Lawrence Livermore National Laboratory
PO Box 808
M/S L-561
Livermore, CA 94550
(W)415-423-2219
hunter@ccc.mfecc.arpa

Joel Jacobs
MITRE Corporation
Burlington Road
M/S K302
Bedford, MA 01730
(W)(617) 271-7373
jdj@mitre.org

Ole Jacobsen
Interop, Inc.
806 Coleman Avenue, #9
Menlo Park, CA 94025
(W)415-325-9542
ole@csli.stanford.edu

Van Jacobson
Lawrence Berkeley Laboratory
One Cycloctron Road
M/S 46A
Berkeley, CA 94720
(W)415-486-6411
van@helios.ll.lbl.gov

Dan Jordt
University of Washington
170 Academic Computer Center
3737 Brooklyn Avenue NE
Seattle, WA 98105
(W)206-543-7352
danj@cac.washington.edu

Dale Johnson
Merit NOC
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)(313) 936-2270
dsj@merit.edu

Satish Joshi
SynOptics Communications
501 East Middlefield Road
Mountian View, CA 94043
(W)415 960 1100
sjoshi@mvis1.synoptics.com

Bruce Jolliffe
University of British Columbia
Computing Center
Vancouver, BC V6T 1W4
Canada
(W)(604) 228-6478
bruce_joliffe@mtsg.ubc.ca

Ajay Kachrani
Digital Equipment Corporation
4 Technology Drive
M/S DSG1-1/F8
Westford, MA 01886-4196
(W)(508) 635-8614
kachrani%regent.dec@decwrl.dec.com

Ken Jones
Concord Communications
753 Forest Street
Marlborough, MA 01752
(W)(508) 460-4646
uunet!konkord!ksj

Jay Kadambi
AT&T Bell Laboratories
110 E. Warrenville Road
M/S 1Z-319
Naperville, IL 60566
(W)(708) 979-8490
jayk@iwlcs.att.com

David Jordan
Emulex
2880 Zanker Road, Suite 204
San Jose, CA 95134
(W)(408) 452-4780
...jordan@emulex.com

Michael Karels
University of California
CSRG Computer Science Div. EECS
457 Evans Hall
Berkeley, CA 94720
(W)415-642-4948
karels@berkeley.edu

Susie Karlson
Interop, Inc.
480 San Antonio Road, Suite 100
Mountian View, CA 94040
(W)415 941-3399
susie@milano.cisco.com

Frank Kastenholz
Racal Interlan
155 Swanson Road
Boxborough, MA 01719
(W)508-263-9929
kasten@europa.interlan.com

David Kaufman
Proteon
Two Technology Drive
Westborough, MA 01581-5008
(W)508-898-2800
dek@proteon.com

Stephen Kent
BBN Communications
150 Cambridge Park Drive
Cambridge, MA 02140
(W)(617) 873-3988
kent@bbn.com

Kathy Kerby
BBN Communications
50 Moulton Street
Cambridge, MA 02138
(W)617 862 8784
kkerby@bbn.com

Jim Kinder
FiberCom, Inc.
P.O. Box 11966
Roanoke, VA 24022
(W)(703) 342-6700
jdk@fibercom.com

Dwaine Kinghorn
Microsoft
1 Microsoft Way
Redmond, WA 98052
(W)(206) 882-8080
microsoft!dwaink

Fred Klein
Bellcore
331 Newman Springs Road
M/S 2X-353
Red Bank, NJ 07701-7030
(W)(201) 758-5089

Paulina Knibbe
cisco Systems, Inc.
1525 O'Brian Boulevard, #13
Menlo Park, CA 94025
(W)(415) 326-1941
knibbe@cisco.com

Holly Knight
Apple Computer
10440 Bubb Road
M/S 58A
Cupertino, CA 95014
(W)(408) 974-5110
holly@apple.com

Jim Knowles
Sterling Software
NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035
(W)(415) 604-5707
jknowles@trident.arc.nasa.gov

Stev Knowles
FTP Software
26 Princess Street
Wakefield, MA 01880-3004
(W)617-246-0900
stev@ftp.com

Alex Koifman
BBN Communications
50 Moulton Street
M/S 031
Cambridge, MA 02138
(W)617-873-8610
akoifman@bbn.com

Lee LaBarre
Mitre Corporation
Burlington Road
M/S EO66
Bedford, MA 01730
(W)617-271-8507
cel@mbunix.mitre.org

Mary Louise Laier
Apple Computer
1675 Kennard Way
Sunnyvale, CA 94087
(W)(408) 974-6034

Sam Lam
Wimsey Associates
15585 M.P.O.
Vancouver, BC V6B 5B3
Canada
(W)(604) 325-4125

Paul Langille
Digital Equipment Corporation
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)(508) 952-3500
quiver::langille@decwrl.dec.com

Tony Lauck
Digital Equipment Corporation
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-496-7644
lauck@tl.dec.com

Gregory Lauer
BBN Communications
10 Moulton Street
Caimbridge, MA 02138
(W)(617) 873-3562
glauer@bbn.com

Joseph Lawrence
Bellcore
331 Newman Springs Rd.
M/S NVC-1C202
Redbank, NJ 07701
(W)201-758-4146
jcl@sabre.bellcore.com

Walter Lazear
Mitre Corporation
7525 Colshire Drive
M/S W4ZZ
McLean, VA 22102
(W)703-883-6515
lazear@gateway.mitre.org

Luping Liang
University of British Columbia
Computer Science Department
Vancouver, BC V6T 1W5
Canada
(W)(604) 228-3061
liang@cs.ubc.ca

James Leighton
Lawrence Livermore National Laboratory
P.O. Box 808
700 East Avenue
Livermore, CA 94551
(W)415-422-4025
jfl@nersc.gov

Fernando Liello
RARE
Via Valerio
2 Trieste
Italy
(W)39-40-5603392
liello@trieste.infn.it

John Lekashman
NASA Ames
M/S 258-6
Moffett Field, CA 94035
(W)415 694-4359
lekash@orville.nas.nasa.gov

John Linn
Digital Equipment Corporation
85 Swanson Road
M/S BXB1-2/D04
Boxborough, MA 01719-1326
(W)(508) 264-5491
linn@ultra.enet.com

Mark Leon
NASA Ames
Moffett Field
M/S 233 8
Mountain View, CA 94035
(W)415-604-6498
leon@nsipo.arc.nasa.gov

Solomon Liou
Penril Datacomm
1300 Quince Orchard Boulevard
Gaithersburg,, MD 20878-4106
(W)(301) 921-8600
solomon%penril@uunet.uu.net

John Leong
Carnegie Mellon University
4910 Forbes Avenue
M/S UCC124
Pittsburgh, PA 15213
(W)412-268-6722
john.leong@andrew.cmu.edu

John LoVerso
Concurrent Computer
One Technology Way
Westford, MA 01886
(W)(508)692-6200
john@loverso.leom.ma.us

Dan Long
BBN Communications
Systems and Technologies; NEARNet
10 Moulton Street
Cambridge, MA 02138
(W)617-873-2766
long@bbn.com

Clifford Lynch
University of California Berkley
300 Lakeside Drive
8th Floor
Oakland, CA 94612-3550
(W)(415) 987-0522
lynch@postgres.berkeley.edu

Charles Lynn
BBN Communications
Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3367
clynn@bbn.com

Stuart Lynne
Unifax Communications
225B Evergreen Drive
Port Moody, BC V3H 1S1
Canada
(W)(604) 937-7532
sl@wimsey.bc.ca

Yoni Malachi
Chipcom
188 Turnpike Road
Southborough, MA 01772-1886
(W)972 8 460068
malachi@polya.stanford.edu

Gary Malkin
FTP Software
26 Princess Street
Wakefield, MA 01880
(W)617-246-0900
gmalkin@ftp.com

Allison Mankin
Mitre
7525 Colshire Drive
M/S W425
McLean, VA 22102
(W)(703) 883-7907
mankin@gateway.mitre.org

Marilyn Martin
CDNnet
CDNnet Headquarters, UBC
403 - 6356 Agricultural Road
Vancouver, BC V6T 1W5
Canada
(W)604-228-6537
martin@cdnnet.ca

Olivier Martin
CERN CN Division
CH-1211
Geneva, 23
Switzerland
(W)41 22 7674894
martin@cearn.cern.ch

Tony Mason
Transarc
The Gulf Tower
707 Grant Street
Pittsburgh, PA 15219
(W)412-338-4400
mason@transarc.com

Matt Mathis
Pittsburgh Supercomputer Center
4400 Fifth Avenue
Pittsburgh, PA 15213
(W)412-268-3319
mathis@pele.psc.edu

Keith McCloghrie
Hughes LAN Systems
1225 Charleston Road
Mountain View, CA 94043
(W)415-966-7934
kzm@his.com

Paul McKenney
SRI International
333 Ravenswood Avenue
EK 360
Menlo Park, CA 94025
(W)(415) 859-4910
mckenney@sri.com

Don McWilliam
University of British Columbia
6356 Agricultural Road
Vancouver, BC
Canada

Milo Medin
Sterling Software
Ames Research Center
Science Internet Project Office
M/S 233-18
Moffett Field, CA 94035
(W)415-490-9157
medin@nsipo.nasa.gov

Alan Menezes
T3 Plus Networking
2840 San Tomas Expressway
Santa Clara, CA 95051
(W)(408) 727-4545
afm@cup.portal.com

Donald Merritt
Balistic Research Lab
Attn: AMXBR-SECAD
Aberdeen Proving Ground, MD 21005-5066

(W)301-278-6808
don@brl.mil

Judy Messing
Mitre Corporation
7525 Colshire Drive
McLean, VA 22102-3481
(W)(703) 883-6670
messing@gateway.mitre.org

David Miller
Mitre Corporation
Burlington Road
D115
M/S EO66
Bedford, MA 01730
(W)617-271-3993
dtm@ulana.mitre.org

Cyndi Mills
BBN Communications
Communications
150 CambridgePark Drive
M/S 20/608
Cambridge, MA 02140
(W)617-873-4143
cmills@bbn.com

Paul Mockapetris
DARPA/ISTO
1400 Wilson Boulevard
Arlington, VA 22209-2308
(W)(213) 822-1511
pvm@isi.edu

Lynn Monsanto
Sun Microsystems
2550 Garcia Avenue
Mountain View, CA 94043
(W)(415) 691-4458

Douglas Montgomery
National Institute of Standards and Technology
Building 225, B-217
Gaithersburg, MD 20899
(W)301 975-3630
dougm@osi3.ncsl.nist.gov

Berlin Moore
Bell of Pennsylvania
PrepNet
530 N. Nevill Street
3rd Floor
Pittsburgh, PA 15213
(W)412-268-7873
prepnet@andrew.cmu.edu

Robert Morgan
Stanford University
115 Pine Hall
Stanford, CA 94305-4122
(W)415 723-9711
morgan@jessica.stanford.edu

Donald Morris
NCAR
Scientific Computing Division
P.O. Box 3000
Boulder, CO 80307-3000
(W)303-497-1282
morris@ucar.edu

James Moulton
Open Network Solutions, Inc.
P.O. Box 669
Sterling, VA 22170
(W)(703) 430-2668
Cobra!jmoul@bellcore.com

John Mullen
767 Juanita Ave
Santa Barbara, CA 93109
(W)(805)966-5286

Mark Needleman
University of California
Division of Library Automation
Kaiser Center Room 854
300 Lakeside Drive
Oakland, CA 94612-3550
(W)(415) 987-0530
mhn@stubbs.ucop.edu

Daniel Nessett
Lawrence Livermore National Laboratory
P.O. Box 808
M/S L-60
Livermore, CA 94551
(W)(415) 422-4033
nessett@nmfecc.arpa

Oscar Newkerk
Digital Equipment Corporation
14475 NE 24th St.
Bellevue, WA 98007
(W)206-865-8913
newkerk@decwet.dec.com

Gerard Newman
San Diego Supercomputer Center
10100 John Jay Hopkins Drive
PO Box 85608
San Diego, CA 92138
(W)619-534-5076
gkn@sds.sdsc.edu

Rebecca Nitzan
NASA Headquarters
600 Maryland Avenue, SW
Suite 200 East
Washington,, DC 20024
(W)(202) 488-1116
nitzan@nsipo.nasa.gov

Bill Nowicki
Sun Microsystems
2550 Garcia Avenue
M/S 14-49
Mountain View, CA 94043
(W)415-690-1300
nowicki@sun.com

Mike Oki
Newbridge Networks
600 March Road
P.O. Box 13600
Kanata, ON K2K 2E6
Canada
(W)(613) 591-3600

Zbigniew Opalka
BBN Communications
150 CambridgePark Drive
Cambridge, MA 02140
(W)617-873-2888
zopalka@bbn.com

Allan Oppenheimer
Apple Computer
19925 Stevens Creek Blvd.
M/S 35K
Cupertino, CA 95014
(W)(408)974-4744

Fred Ostapik
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
(W)(415) 859-6111
ostapik@nisc.sri.com

Philippe Park
BBN Communications
Systems and Technologies
10 Moulton Street
Cambridge, MA 02238
(W)617-873-2892
ppark@bbn.com

Brad Parker
Cayman Systems
26 Landsdowne Street
Cambridge, MA 02139
(W)617-494-1999
brad@cayman.com

Craig Partridge
BBN Communications
Systems and Technologies
10 Moulton Street
M/S 6/5A
Cambridge, MA 02138
(W)(617) 873-2459
craig@nnsc.nsf.net

David Perkins
3Com
2081 N. Shoreline Boulevard
Mountain View, CA 94043
(W)415-694-2808
dave_perkins@3com.com

Eric Peterson
Microsoft
1 Microsoft Way
Redmond, WA 98052
(W)(206) 882-8080
microsoft!ericpe

Robert Pinna
Hewlett-Packard
P.O. Box 7050
5070 Centennial Boulevard
Colorodo Springs, CO 80933-7050
(W)719-531-4370
bwp!hpctlb.hp.com

David Piscitello
Bellcore
331 Newman Springs Road
M/S NVC 1C-236
Red Bank, NJ 07701
(W)201-758-2286
dave@sabre.bellcore.com

Mary Ann Potts
University of British Columbia
Physics Department
6224 Agricultural Road
Vancouver, BC V6T 2A6
Canada
(W)(604) 228-6913
map@physics.ubc.ca

Cecilia Preston
SCR
2140 Shattuck Avenue, #2170
Emeryville, CA 94704
(W)(415) 644-4216
ceal@asylum.sf.ca.us

Stephanie Price
CMC
125 Cremona
Santa Barbara, CA 93117
(W)805-968-4262
cmcvax!price@hub.ucsb.edu

Vicki Ralls
cisco Systems
1525 O'Brien Drive
Menlo Park, CA 94025
(W)415 688-8245
ralls@cisco.com

K.K. Ramakrishnan
Digital Equipment Corporation
LKG 1-2/A19
550 King Street
Littleton, MA 01460-1289
(W)508-486-7267
rama%erlang.dec.com@decwrl.dec.com

James Reeves
SynOptics Communication
501 E. Middlefield Rd.
Mountain View, CA 94043
(W)(415)691-7240
jreeves@synoptics.com
(W)(415)691-7240
jreeves@synoptics.com

Michael Reilly
Digital Equipment Corporation
Network Systems Lab
505 Hamilton Ave.
Palo Alto, CA 94301
(W)415-853-6593
reilly@nsl.dec.com

Jim Reinstedler
Ungermann-Bass
2560 Mission College Blvd.
Santa Clara, CA 95054
(W)408 562-5636
jimr@ub.com

Yakov Rekhter
IBM Corporation
TJ Watson Research
Route 134, PO Box 218
Yorktown Heights, NY 10598
(W)914-945-3896
yakov@ibm.com

Raphael Renous
Fibronics, Advanced Technology Center
Haifa, 31905
ISRAEL
(W)972-4-313604

Robert Reschly
U.S. Army
SLCBR-SE-A (Reschley)
Ballistic Research Laboratory
Aberdeen Proving Ground, MD 21005-5066

(W)301-278-6808
reschly@brl.mil

Joyce K. Reynolds
USC Information Sciences Institute
4676 Admiralty Wy #1001
Marina del Rey, CA 90292-6695
(W)213-822-1511
jkrey@venera.isi.edu

Michael Roberts
EDUCOM
1112 16th Street, NW
Suite 600
Washington, DC 20036
(W)(202) 872-4200
roberts@educom.edu

Ron Roberts
Stanford University
BARRNet
Pine Hall, Room 175-B
Stanford, CA 94305-4122
(W)415-723-3352
roberts@jessica.stanford.edu

Manuel Rodrigues
AT&T Bell Laboratories
Crawfords Corner Road
Room 1F-412
Holmdel, NJ 07733
(W)(201)949-4655

Marshall Rose
Performance Systems International
California Office
5210 Great American Parkway
Suite 3106
Santa Clara,, CA 95054
(W)(415) 961-3380
mrose@psi.com

Kenneth Rossen
BBN Communications
150 Cambridge Park Drive
M/S 20/10B
Cambridge, MA 02140
(W)(617)873-3140
kenr@bbn.com

Jonathan Saperia
Digital Equipment Corporation
550 King Street
LKG 1-2/B13
Littleton, MA 01460-1289
saperia%tcpjon@decwrl.dec.com

Jeffrey Schiller
Massachusetts Institute of Technology
1 Amherst Street
Room E40-311
Cambridge, MA 02139
(W)617-253-8400
jis@bitsy.mit.edu

Tim Seaver
Microelectronics Center of North Carolina

P.O. Box 12889
Research Triangle Park, NC 27709
(W)919-248-1973
tas@mcnc.org

Mark Seger
Digital Equipment Corporation
550 King Street
M/S LKG1-2/87
Littleton, MA 01460
(W)(508)486-5530
seger@mjs1/ogo.dec.com

Steve Senum
Network Systems
7600 Boone Avenue North
Minneapolis, MN 55428
(W)612-424-4888
sjs@network.com

Harvey Shapiro
NARDAC
Code 4211 Bldg 143 4th Floor
Washington Navy Yard
Washington, DC 20374
(W)(202)433-5422
shapira@wnyosi2.arpa

Jim Sheridan
IBM
166 East Shore Drive
PO Box 334
Whitmore Lake, MI 48189
(W)313-393-6537
jsherida@ibm.com

Robert Shirey
Mitre Corporation
7525 Colshire Dr.
M/S Z286
McLean, VA 22102-3481
(W)703-883-7210
shirey@mitre.org

Jim Showalter
Defense Communications Agency
1860 Wiehle Avenue
Reston, VA 22090-5500
(W)703-437-2580
gamma@mintaka.dca.mil

Deepinder Sidhu
University of Maryland - BC
Baltimore, MD 21228
(W)(301) 455-3028
sidhu@umbc3.umbc.edu

Dana Sitzler
Merit Computer Network
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-3000
dds@merit.edu

Keith Sklower
University of California Berkeley
Computer Science Dept.
570 Evans Hall
Berkeley, CA 94720
(W)415-642-9587
sklower@okeeffe.berkeley.edu

Erik Skovgaard
PSC
745 Clark Drive
Vancouver, BC V5L 3J3
Canada
(W)(604)253-0766
eskovgaa@uvcw.uvic.ca

Frank Slaughter
Shiva
One Cambridge Center
Cambridge, MA 02141
(W)617-864-8500
fgs@shiva.com

Mark Sleeper
Sparta
7926 Jones Branch Dr.
McLean, VA 22102
(W)703-448-0210
mws@sparta.com

Craig Smelser
Digital Equipment Corporation
110 Spit Brook Road
M/S ZK03-3/U14
Nashua, NH 03062

Miles Smid
National Institute of Standards Technology
A-216 Technology
Gaithersburg, MD 20899
(W)(301)975-2938
smid@st1.ncsl.nist.gov

Brock Smith
University of British Columbia
Networking & Communications
6356 Agricultural Road
Computer Science Building
Vancouver, BC V6T 1W5
Canada
(W)(604)228-6062
smith@netcom.ubc.ca

Frank Solensky
Racal Interlan
155 Swanson Road
Boxborough, MA 01719
(W)508-263-9929
solensky@interlan.interlan.com

Cheng Song
IBM Corporation
472 Wheelers Farms Road
M/S 91
Milford, CT 06460
(W)(203)783-7039
song@ibm.com

Charlie Stark
NCR
2700 Snelling Ave. N.
M/S S 075
St. Paul, MN 55113
(W)(612) 638-7218

Tony Staw
Digital Equipment Corporation
P.O. Box 121
Worton Grange, Imperial Way
M/S REO-2 G/G9
Reading, Berkshire RG2 0TU
GB
(W)0734 868711
staw@marvin.enet.dec.com

Martha Steenstrup
BBN Communications
150 Cambridge Park Drive
Room 20/665
Cambridge, MA 02140
(W)(617)873-3192
msteenst@bbn.com

Einar Stefferud
Network Management Associates, Inc.
17301 Drey Lane
Huntington Beach, CA 92647
(W)(714) 842-3711
EStefferud@ECL

Mark Stein
Sun Microsystems
2550 Garcia Avenue
M/S 14-40
Mountain View, CA 94043
(W)(415)336-2435
marks@eng.sun.com

Ed Stern
Proteon
Two Technology Drive
Westborough, MA 01581
(W)(508)898-2800
els@proteon.com

Ken Stetten
National Radio Astronomy Observatory
Box O
Socorro, NM 87801
(W)(505) 835-7234
kstetten@nrao.edu

Bob Stewart
Xyplex
330 Codman Hill Road
Boxborough, MA 01719
(W)(508)264-9900
rlstewart@eng.xyplex.com

Roxanne Streeter
Sterling Software NASA Ames
Science Internet Project Office
M/S 233-18
Moffett Field, CA 94035
(W)415-694-4845
streeter@nsipo.arc.nasa.gov

Allen Sturtevant
Lawrence Livermore National Laboratory
PO Box 808
M/S L-561
Livermore, CA 94550
(W)415-422-8266
sturtevant@ccc.nmfecc.gov

Zaw-Sing Su
SRI International
333 Ravenswood Avenue
EJ280
Menlo Park, CA 94025
(W)(415) 859-4576
zsu@tsca.istc.sri.com

Sally Tarquinio
Mitre Corporation
7525 Coldshire Drive
M/S W425
McLean, VA 22102
(W)(703)883-7987
sally@gateway@mitre.org

Dean Throop
Data General Corporation
62 Alexander Dr.
Research Triangle Park, NC 27709
(W)919-549-8421
throop@dg-rtp.dg.com

Claudio Topolcic
BBN Communications
Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3874
topolcic@bbn.com

Bill Townsend
Xylogics
53 Third Avenue
Burlington, MA 01803
(W)(617)272-8140
townsend@xylogics.com

Glenn Trewitt
Digital Equipment Corporation
505 Hamilton Avenue
M/S UCH-1
Palo Alto, CA 94301
(W)415-688-1323
trewitt@nsl.dec.com

Paul Tsuchiya
Bellcore
435 South Street
M/S MRE 2P-251
Morristown, NJ 07960
(W)(201) 829-4484
tsuchiya@thumper.bellcore.com

Gregory Vaudreuil
Corp. for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)(703) 620-8990
gvaudre@nri.reston.va.us

John Veizades
Apple Computer
20525 Mariani Avenue
M/S 67-B
Cupertino, CA 95014
(W)408-974-2672
veizades@apple.com

Sudhanshu Verma
Hewlett-Packard
19420 Homestead Road
Cupertino, CA 95014
(W)408-447-3417
verma@hpindbu.cup.hp.com

Rudiger Volk
Universitaet Dortmund
Informatik IRB
Postfach 500 500
D-4600 Dortmund 50
Dortmund,
GERM
(W)+49-2317554760
rv@informatik.uni-dortmund.de

Tom VonDeak
NASA Lewis
21000 Brookfield Road
Cleveland, OH 44135
(W)(216)433-3277
tvondeak@nasamail.nasa.gov

Steve Waldbusser
Carnegie Mellon University
4910 Forbes Avenue
UCC 130
Pittsburgh, PA 15213
(W)412-268-6628
sw01+@andrew.cmu.edu

Justin Walker
Apple Computer
10440 Bubb Road
M/S 58A
Cupertino, CA 95014
(W)(408)974-4427
justin@apple.com

Dale Walters
National Institute of Systems Technology
270 Quince Orchard Road
Building 225 B217
Gaithersburg, MD 20899
(W)(301)975-3641

Carol Ward
University of Colorado
Westnet
3645 Marine Street
Boulder, CO 80309-0455
(W)303-492-5860
cward@spot.colorado.edu

Chris Weider
Merit
1075 Beal Ave
Ann Arbor, MI 48104
(W)(313)936-3000
clw@merit.edu

Jonathan Wenocur
Shiva
One Cambridge Ctr
Cambridge, MA 02141
(W)617-864-8500
jhw@shiva.com

Bill Westfield
cisco Systems
1525 O'Brien Drive
Menlo Park, CA 94025
(W)415 326-1941
billw@cisco.com

John Wieronski
Ohio State University
Ohio Supercomputer Center
1224 Kinnear Road
Columbus,, OH 43212
(W)(614)292-9290
john@osc.edu

Steve Wilbur
University College London
Department of Computer Science
Gower Street
London, WC1E 6BT
ENGLAND
(W)44 71 387 7050
wilbur@uk.ac.ucl.cs

Bert Williams
Synernetics
85 Rangeway Road
North Billerica, MA 01862
(W)(508)670-9009
bert.synernetics@mailgate.synnet.com

Steve Willis
Wellfleet Communications
12 DeAngelo Drive
Bedford, MA 01730
(W)617-275-2400
swillis@wellfleet.com

Hon-Man Wong
Simon Fraser University
Centre for Systems Science
Burnaby, BC V5A 1S6
Canada
(W)(604)291-4711
honman@cs.sfu.ca

Walter Wimer
Carnegie Mellon University
Networking & Communications
4910 Forbes Avenue
M/S UCC 132
Pittsburgh, PA 15213-3890
(W)412-268-6252
ww0n+@andrew.cmu.edu

C. Philip Wood
Los Alamos National Laboratory
Network Engineering
M/S B255
Los Alamos, NM 87545
(W)(505)667-2598
cpw@lanl.gov

Linda Winkler
Argonne National Laboratory
9700 S. Cass Avenue
Building 221, B-251
Argonne, IL 60439
(W)708-972-7236
b32357@anlvm.ctd.anl.gov

Mark Wood
AT&T Computer Systems
1100 E. Warrenville Road
M/S 1A 314
Naperville, IL 60566
(W)(708)979-7521
mark1@iw/cs.att.com

Dan Wintringham
Ohio State University
Ohio Supercomputer Center
1224 Kinnear Road
Columbus, OH 43212
(W)614-292-0901
danw@igloo.osc.edu

Robert Woodburn
SAIC
CSEIC
8619 Westwood Center Drive
Vienna, VA 22182
(W)703-734-9000
woody@saic.com

Jean Wu
PSC
745 Clark Drive
Vancouver, BC V5L 3J3
Canada
(W)(604)253-0766
eskovgaa@uvcw.uvic.ca

Yueli Yang
Bell Northern Research
P.O. Box 3511
Station C
M/S 7B31
Ottawa, ON K1Y 4H7
Canada
(W)(613)763-4243
yueli@bnr.ca

Denis Yaro
Sun Microsystems
2550 Garcia Avenue, 14-49
Mountain View, CA 94043
(W)(415)336-5435
DYARO@SUN.COM

Allan Young
Royal Melbourne Institute of Technology
Electronic Communications Group
124 Latrobe Street
Melbourne, 3000 AUSTRALIA
(W)011 6136602799
rcoay@possum.ecg.rmit.oz.au

Jessica Yu
Merit Computer Network
1075 Beal Avenue
Ann Arbor, MI 48104
(W)(313)936-2655
jyy@merit.edu

Sijiam Zhang
University of British Columbia
327 E. 46th Avenue
Vancouver, BC V6W 1Z7
Canada
(W)(604)322-0687
szhang@cs.ubc.ca