Proceedings of the Twenty-Second Internet Engineering Task Force

Los Alamos National Laboratory
Santa Fe, New Mexico
November 18 - 22, 1991

Corporation for National Research Initiatives

PROCEEDINGS OF THE TWENTY-SECOND INTERNET ENGINEERING TASK FORCE

LOS ALAMOS NATIONAL LABORATORY
SANTA FE, NEW MEXICO
November 18 - 22, 1991

Compiled and Edited by
Megan M. Davies
Cynthia Clark
Debra Legare

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 Twenty-Second Internet Engineering Task Force was held in Santa Fe, New Mexico, during the week of November 18th. Our Local Host, Los Alamos National Laboratory provided excellent Internet Connectivity. Specifically, we would like to express our appreciation to Dale Land and Peter Ford, who, along with several others, including Randy Hoebelheinrich, Glenn Michel, C. Philip Wood, Steve Howard, and Paul Norris, ensured the smooth operation of the Terminal Room. We wish to recognize as well, the following organizations for their generous contributions:

US West	T1 Circuit
Digital Equipment	(3) DS5000, (1) VS3100, (2) VT1200, (2) VT220
Sun Microsystems	Color Sparc2, Sparc ELC, Sparc Printer
Hewlett Packard	(2) HP 720 Color Workstations
BFA Sales	Datability Terminal Server
LANL	Mac IIfx, IBM PS2/60, Mac IIci, Apple Laserwriter, cisco CGS Router, (6) VT200 Terminals, LAN

We also wish to acknowledge those individuals who contributed to the high quality of the Technical Presentation Sessions. Philip Almquist, Jordan Becker, Laura Breeden, George Clapp, Steve Deering, Susan Estrada, Deborah Estrin, Tate Jennings, Mark Knopper, Tom Lyon, John Morrison, Radia Perlman, Martha Steenstrup, and Scott Williamson.

Finally, a big thank you to Debra Legare for her efficient management of the registration process as well as to Paramjit (Pam) Bhamrah and Sarah Hebert (LANL) for their invaluable assistance. Thanks also to Cynthia Clark and Greg Vaudreuil for their help with these Proceedings.

Megan Davies/CNRI

Contents

Chair's Message

The IETF met in Santa Fe on November 18-22, 1991. The meeting was hosted by Los Alamos National Laboratory (LANL). Many thanks are due to Dale Land, John Morrison, C. Philip Wood, Peter Ford, and many others at LANL for the amazing amount of work that went into hosting this meeting. The facilities were outstanding and the location was beautiful. Numerous folks mentioned to me that this was a very productive IETF meeting. I think we can thank LANL (and perhaps, the clear mountain air?) for helping to make this such a productive meeting).

The meeting was attended by approximately 350 people. It was quite productive with 46 working groups and 11 BOF's meeting in over 80 separate sessions. Three IETF Area "advisory" groups met – the Security Area Advisory Group (SAAG), the Operational Requirements Area Directorate (ORAD), and the User Services Area Council (USAC).

We are very pleased that FARNET chose to meet in Santa Fe during the same week, so that there was quite a bit of interaction between IETF and FARNET interests during that time. In particular, ORAD met jointly with FARNET, and had a very productive session pursuing FARNET's topic for the week – "Hardening the Mid-level Networks".

The Internet Activites Board (IAB) also took this opportunity to meet in Santa Fe. It was quite helpful to have IAB members in attendance at the IETF, and this helped increase the communication and positive interaction between the IAB, the IESG, and the IETF. I feel that the IETF benefits greatly from the direct participation of IAB members in the various working group activities. I hope we will continue to see this close interworking between the IAB and IETF.

There were 14 technical presentations during the week. As it turns out, there was an increased interest in ATM at this meeting, with three separate presentations on the basic technical details of ATM and an interesting approach to using ATM in local area networks. There was also a BOF on "IP over ATM", which will become a working group at the next meeting.

There was an important focus on routing at this IETF. Martha Steenstrup (BBN) presented a status report on Inter-Domain Policy Routing (IDPR), and how IDPR might interwork with BGP (or other inter-domain routing protocols). Deborah Estrin (USC) presented a proposal, co-authored with Yakov Rekhter (IBM) for a "Unified" Inter-Domain Routing Protocol. Noel Chiappa ran a BOF on his proposal for a new routing and addressing architecture. Noel's BOF was based on his presentation at the July IETF meeting in Atlanta, and will likely evolve into a working group effort. The BGP Working Group had several very important sessions. During one

BGP Working Group session, Jessica Yu (Merit) introduced a new working group effort under the Operational Requirements Area to concentrate on the operational deployment of BGP.

In another BGP Working Group session, Phill Gross (ANS) led a discussion on introducing address masks into BGP, including the notion of "supernet masks" to condense information in routing tables. The discussion soon expanded to encompass the related problem of IP address depletion. As a result, the assembled group, along with the IAB and IESG, organized the Routing and Addressing Working Group (ROAD).

The goal of the Working Group will be to propose methods to deal with the related problems of routing table scaling and IP address depletion. The ROAD Working Group will hold its first meeting at the IETF meeting in San Diego (March 16-20, 1992). The IETF effort dovetailed very nicely with the results of the IAB/IESG Architecture Retreat in June (reported at the July IETF meeting), which recommended (in part) that an IETF working group be formed to pursue this crucial matter. In an attempt to help focus the activities of this important Group several members of the IAB retreat have joined some participants from the BGP Working Group session to set the agenda for the ROAD Working Group in March, and explore some of the various alternatives.

IESG and IAB Reporting of Internet Standardization

The procedures for reporting and tracking Internet standardization activities have grown in an ad hoc fashion over the last several years as the IETF standardization activities have expanded.

In Santa Fe, the IAB and IESG wrote down the following sequence of procedures for reporting and tracking Internet standarization actions to the IETF and the wider Internet community:

Announce WG Progress	I-D announcements, To: IETF
Announce WG Completion	i.e., "Last Call", From: IESG, To: IETF, IAB
Announce IESG Recommendation	From: IESG, To: IAB, cc: IETF
Announce IAB Outcome	From: IAB, To: IETF, IESG
RFC Published	RFC List

Essentially the same procedure is followed for standards actions at any of the three levels of Internet standardization – Proposed, Draft, Internet Standard. Note that the second step ("last call") is new. It was added to assure that interested parties will have additional notification and time to make comments on upcoming standards actions.

Upcoming IETF Meetings

The next IETF meeting will be hosted by San Diego Supercomputer Center on March 16-20, 1992. E. Paul Love, Jr. and Hans-Werner Braun will act as local hosts. Reservation material will be sent to the IETF mailing list in January 1992. Note that this is the same week of the America's Cup, so San Diego will be VERY crowded. Please try to make your reservations as early as possible.

We are now working very hard to schedule IETF meetings further into the future. Our goal is to schedule meetings at least one year in advance.

Please note that we are now planning to hold our first IETF meeting outside North America in the Fall of 1993 in Europe. This is a natural step, with the Internet Society beginning operation in 1992, and with the IETF finding itself increasingly involved in international issues. More information on this important development will be made available as the plans become firm.

IETF Report in the Internet Society Quarterly Newsletter

The Internet Society will be publishing a newsletter on a quarterly basis. Activities in the IETF will be reported regularly in this newsletter.

Phill Gross
IETF Chair

Agenda of the Twenty-Second IETF

(November 18-22, 1991)

MONDAY, November 18, 1991

8:00-9:00 am	IETF Registration and Continental Breakfast	
9:00-9:30 am	Introductions and Local Arrangements • Local Host Orientation (Dale Land/LANL) • FARNET Overview (Laura Breeden/FARNET)	
9:30-12:00 noon	Morning Sessions	
	APP	Internet Message Extensions WG (Greg Vaudreuil/CNRI)
	INT	IP over Appletalk WG (John Veizades/Apple)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	MGT	X.25 Management Information Base WG (Dean Throop/Data General)
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	RTG	Inter-Domain Policy Routing WG (Martha Steenstrup/BBN)
	SEC	Security Area Advisory Group (Stephen Crocker/TIS)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	SNMPMIB Compiler BOF (Dave Perkins/SynOptics)
	APP	Internet Message Extensions WG (Greg Vaudreuil/CNRI)
	INT	IP over Appletalk WG (John Veizades/Apple)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	SEC	Common Authentication Technology WG (John Linn/DEC)
	USV	User Services WG (Joyce Reynolds/ISI)
3:30-4:00 pm	Break (Refreshments provided)	

4:00-6:00 pm	Afternoon Sessions II	
	APP	Internet Message Extensions WG (Greg Vaudreuil/CNRI)
	INT	IP over Appletalk WG (John Veizades/Apple)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	RTG	Border Gateway Protocol WG (Yakov Rekhter/IBM)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	SEC	Privacy-Enhanced Electronic Mail WG (Steve Kent/BBN)
	USV	Directory Information Services Infrastructure WG (Chris Weider/Merit)

TUESDAY, November 19, 1991

8:30-9:00 am	Continental Breakfast
9:00-9:30 am	Technical Presentations • "Overview of ATM" (George Clapp/Ameritech)
9:30-12:00 noon	Morning Sessions

APP	Internet SMTP Extensions WG (Greg Vaudreuil/CNRI)
APP	Network Database WG (Daisy Shen/IBM)
APP	Telnet WG (Steve Alexander/INTERACTIVE Systems)
OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB and Rob Hagens/UWisc)
RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
RTG	Border Gateway Protocol WG (Yakov Rekhter/IBM)
RTG	Inter-Domain Policy Routing WG (Martha Steenstrup/BBN)
TSV	Service Location Protocol WG (John Veizades/Apple)
USV	Network Information Services Infrastructure WG (April Marine/SRI and Pat Smith/Merit)

Breaks	Coffee available throughout morning.
1:30-3:30 pm	Afternoon Sessions I

APP	Internet SMTP Extensions WG (Greg Vaudreuil/CNRI)
APP	Network Database WG (Daisy Shen/IBM)
MGT	IEEE 802.3 Hub MIB WG (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
OPS	Operational Requirements Area Directorate (Susan Estrada, CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB and Rob Hagens/UWisc)
RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
RTG	Multicast Extensions to OSPF WG (Steve Deering/Xerox PARC)
TSV	Domain Name System WG (Mike Reilly/DEC)

	SEC	Common Authentication Technology WG (John Linn/DEC)
	USV	Internet School Networking WG (John Clement/EDUCOM, Connie Stout/TheNet and Art St. George/UNM)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	INT	Router Requirements WG (Philip Almquist/Consultant)*
	MGT	IEEE 802.3 Hub MIB WG (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
	OPS	User Connectivity Problems WG (Dan Long/BBN)
	OSI	Office Document Architecture WG (Peter Kirstein/UCL)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	Inter-Domain Policy Routing WG (Martha Steenstrup/BBN)*
	RTG	Open Shortest Path First IGP WG (John Moy/Proteon)
	TSV	Distributed File Systems WG (Peter Honeyman/UMich)
	USV	Internet Anonymous FTP Archives WG (Peter Deutsch/McGill and Alan Emtage/McGill)
7:00-10:00 pm	Evening Sessions	
	BOF	Authentication, Authorization, and Accounting Issues for Terminal/Network Servers BOF (Larry Blunk/Merit)
	BOF	Dynamic Creation of Network Links BOF (Andy Nicholson/Cray Research)
	BOF	Teleconferencing BOF (Russ Hobby/UCDavis)
	BOF	RIP V2 Internet Draft BOF (Gary Malkin/FTP)
	BOF	Router Requirements Checklist BOF (Susan Estrada/CERFnet)
	BOF	SNMP Device Discovery BOF (Fred Baker/ACC)
	OSI	Network OSI Operations WG (Sue Hares/Merit)

*IDPR and RREQ will meet jointly

WEDNESDAY, November 20, 1991

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	• "ATM in LANs" (Tom Lyon/Sun Microsystems)	
9:30-12:00 noon	Morning Sessions	
	APP	Network Fax WG (Mark Needleman/UC)
	BOF	IP Over ATM BOF (Bob Hinden/BBN)
	INT	Point-to-Point Protocol Extensions WG (Brian Lloyd/Consultant)
	MGT	Internet Accounting WG (Cyndi Mills/BBN and Gregory Ruth/BBN)
	OPS	Benchmarking WG (Scott Bradner/Harvard)
	OPS	Network Status Reports (Phill Gross/ANS)
	OPS	User Connectivity Problems WG (Dan Long/BBN)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB and Rob Hagens/UWisc)
	SEC	Privacy-Enhanced Electronic Mail WG (Steve Kent/BBN)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	New Technology TCP BOF (Sean O'Malley/UAZ)
	APP	Telnet WG (Steve Alexander/INTERACTIVE Systems)
	INT	Multi Media Bridging WG (Jeff Fitzgerald/Fibercom)
	INT	Point-to-Point Protocol Extensions WG (Brian Lloyd/Consultant)
	MGT	Remote LAN Monitoring WG (Mike Erlinger/Micro Technology)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	ISIS for IP Internets WG (Ross Callon/DEC)
	OPS	Network Joint Management WG (Gene Hastings/PSC)
	OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB and Rob Hagens/UWisc)

	SEC	Privacy-Enhanced Electronic Mail WG (Steve Kent/BBN)
	USV	Internet User Glossary WG (Tracy LaQuey Parker/UTexas and Gary Malkin/FTP Software)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Technical Presentations	

- "A Unified Approach to Inter-Domain Routing" (Deborah Estrin/USC and Yakov Rekhter/IBM)
- "Practical Method of Foiling the Denial of Service Threat in the Network" (Radia Perlman/DEC)
- "Development of ATM and the ATM Standards" (George Clapp/Ameritech)

7:00-10:00pm	Evening Session	
	BOF	Developing Operational Measurement Criteria BOF (Susan Estrada/CERFnet and Tim Salo/MSC)
	BOF	IP Routing for Wireless/Mobile Hosts BOF (Steve Deering/Xerox)
	BOF	NOOP OSI Network Tools BOF (Sue Hares/Merit)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	RTG	BGP Deployment and Application BOF Jessica Yu/Merit)
	SEC	SNMP Security WG (James Galvin/TIS and Keith McCloghrie/Hughes)
	USV	NOC-Tools Catalogue Revisions WG (Robert Enger/Contel and Darren Kinley/RISQ)

THURSDAY, November 21, 1991

8:30-9:00 am	Continental Breakfast	
9:00-12:00 noon	Morning Sessions	
	BOF	Audio/Video Transport BOF (Steve Casner/ISI)
	BOF	New Internet Routing and Addressing Architecture BOF (Noel Chiappa)
	APP	Internet SMTP Extensions WG (Greg Vaudreuil/CNRI)
	INT	Multi Media Bridging WG (Jeff Fitzgerald/Fibercom)
	MGT	Internet Accounting WG (Cyndi Mills/BBN and Gregory Ruth/BBN)
	MGT	Ethernet MIB WG (Frank Kastenholz/Clearpoint)*
	MGT	Simple Network Management Protocol WG (Marshall Rose/DBC)*
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	SEC	Security Area Advisory Group (Stephen Crocker/TIS)
	TSV	TCP Large Windows WG (Dave Borman/Cray Research)
Breaks	Coffee available throughout the morning.	
1:30-3:30 pm	Technical Presentations • "Advanced Network Activities at Los Alamos" (John Morrison/LANL) • "Inter-Domain Policy Routing" (Martha Steenstrup/BBN) • "Router Requirements" (Philip Almquist/Consultant)	
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Open Plenary and IESG • Protocol Standards Actions	

*SNMP and Ethermib will meet jointly.

FRIDAY, November 22, 1991

8:30-9:00 am	Continental Breakfast	
9:00-12:00pm	Technical Presentations	

- "FARNET and the IETF" (Susan Estrada/CERFnet)
- "Router Discovery" (Steve Deering/Xerox)
- "SMDS Showcase and Applicable RFC's" (Tate Jennings/Bellcore)
- "NSFnet Update" (Mark Knopper/Merit and Jordan Becker/ANS)
- "NIC Services" (Scott Williamson/Network Solutions)

1:30-3:30 Summary Reports

APP	Applications Area (Russ Hobby/UC Davis)
INT	Internet Area (Noel Chiappa and Philip Almquist/Consultant)
MGT	Network Management Area (James Davin/MIT)
OPS	Operational Requirements Area (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
OSI	OSI Integration Area (Ross Callon/DEC)
RTG	Routing Area (Bob Hinden/BBN)
SEC	Security Area (Steve Crocker/TIS)
TSV	Transport and Services Area (Dave Borman/Cray Research)
USV	User Services Area (Joyce K. Reynolds/ISI)

3:30-4:00 pm Concluding Remarks (Phill Gross/ANS)

Key to Abbreviations

APP	Applications Area
BOF	Birds of a Feather Session
INT	Internet Area
MGT	Network Management Area
OSI	OSI Integration Area
OPS	Operational Requirements Area
RTG	Routing Area
SEC	Security Area
TSV	Transport and Services Area
USV	User Services Area

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) is the protocol engineering, development, and standardization arm of the Internet Architecture Board (IAB). The IETF began in January 1986 as a forum for technical coordination by contractors for the U.S. Defense Advanced Projects Agency (DARPA), working on the ARPANET, U.S. Defense Data Network (DDN), and the Internet core gateway system. Since that time, the IETF has grown into a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet protocol architecture and the smooth operation of the Internet.

The IETF mission includes:

1. Identifying and proposing solutions to pressing operational and technical problems in the Internet,

2. Specifying the development (or usage) of protocols and the near-term architecture to solve such technical problems for the Internet,

3. Making recommendations to the IAB regarding standardization of protocols and protocol usage in the Internet,

4. Facilitating technology transfer from the Internet Research Task Force (IRTF) to the wider Internet community, and

5. 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 nine technical areas. Each is led by an Area Director who has primary responsibility for that one area of IETF activity.

Together with the Chair of the IETF, these nine technical Directors (plus, a Director for Standards Procedures) compose the Internet Engineering Steering Group (IESG).

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

IETF and IESG Chair:	Phill Gross/ANS
Applications:	Russ Hobby/UC-Davis
Internet:	Noel Chiappa
	Philip Almquist/Consultant
Network Management:	James Davin/ MIT
OSI Integration:	Dave Piscitello/Bellcore
	Erik Huizer/SURFnet
Operational Requirements:	Phill Gross/ANS
	Bernhard Stockman/NORDUnet
	Susan Estrada/CERFnet
Routing:	Robert Hinden/BBN
Security:	Steve Crocker/TIS
Transport and Services	Dave Borman/Cray Research
User Services	Joyce K. Reynolds/ISI
Standards Management:	Dave Crocker/TBO

The IETF has a Secretariat, headquartered at the Corporation for National Research Initiatives in Reston, Virginia, with the following staff:

IETF Executive Director:	Steve Coya
IESG Secretary:	Greg Vaudreuil
IETF Coordinator:	Megan Davies
Administrative Support:	Debra Legare
	Cynthia Clark

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 4.5 day plenary sessions three times a year. These plenary sessions are composed of Working Group Sessions, Technical Presentations, Network Status Reports, working group reporting, and an open IESG meeting. A Proceedings of each IETF plenary is published, which includes reports from each Area, each working group, and each Technical Presentation. The Proceedings includes a summary of all current standardization activities.

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, requests should be sent to `ietf-info@nri.reston.va.us`. An archive of mail sent to the IETF list is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 Future IETF Meeting Sites

Spring 1992

San Diego Supercomputer Center
Host: E. Paul Love, Jr. and Hans-Werner Braun
March 16-20, 1992

Summer 1992

Massachusettes Institute of Technology
Host(s): Dave Clark and James Davin
July 13-17, 1992

Fall 1992

U.S. Sprint
Host: Robert Collet
November 16-20, 1992 (tentative)

Spring 1993

OARnet and Ohio State University
Host: Kannan Varadhan
March 1993 (tentative)

Summer 1993

CRIM
Host: Darren Kinley
TBD

1.2 On Line IETF Information

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

Directory Locations

Information pertaining to the IETF, its working groups and Internet Drafts can be found in either the "IETF" Directory or the "Internet-Drafts" Directory. (For a more detailed description of these Directories, please see Section 1.2.1 and 1.2.2). To retrieve this information via FTP, establish a connection, then Login with username "anonymous" and the password requested by the system. This password will either be your login name or "guest". When logged in, change to the directory of your choice with the following commands:

cd ietf
cd internet-drafts

Individual files can then be retrieved using the GET command:

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

East Coast (US) Address: nnsc.nsf.net (128.89.1.178)

West Coast (US) Address: ftp.nisc.sri.com (192.33.33.22)

Internet Drafts are available by mail server from this machine. To retreive a file mail a request:

To: mail-server@nisc.sri.com
Subject: Anything you want

In the body put a command of the form:
send internet-drafts/1id-abstracts.txt or
send ietf/1wg-summary.txt

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

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

Europe Address: nic.nordu.net (192.36.148.17)

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

1.2.1 The IETF Directory

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

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

FILE NAME	
0mtg-agenda	The current Agenda for the upcoming IETF plenary, containing scheduled Working Groups meetings, Technical Presentations and Network Status Reports.
0mtg-at-a-glance	The announcement for the upcoming IETF plenary, containing specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and meeting costs.
0mtg-rsvp	A standardized RSVP form to notify the secretariat of your plans to attend the upcoming IETF meeting.
0mtg-sites	Current and future meeting dates and sites for IETF plenaries.
1id-abstracts	The Internet Drafts currently on-line in the Internet-Drafts Directory.
1id-guidelines	Instructions for authors of Internet Drafts.
1ietf-description	A short description of the IETF, the IESG and how to participate.
1wg-summary	A listing of all current working groups, the working group Chairs and their email addresses, working group mailing list addresses, and where applicable, documentation produced. This file also contains the standard acronym for the working groups by which the IETF and Internet-Drafts Directories are keyed.

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

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

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

1.2.2 The Internet-Drafts Directory

The Internet-Drafts Directory has been installed to make available, for review and comment, draft documents that will be submitted ultimately to the IAB and the RFC Editor to be considered for publishing as RFC's. These documents are indexed in the file 1id-abstracts.txt in the Internet-Drafts Directory. Comments are welcome and should be addressed to the responsible person whose name and email addresses are listed on the first page of the respective draft.

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

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

where <std wg abrev> is the working group acronym, <docname> is an abbreviated version of the document title, and <rev> is the revision number.

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

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

where <author> is the author's name.

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

1.3 Guidelines to Authors of Internet Drafts

The Internet-Drafts Directories are available to provide authors with the ability to distribute and solicit comments on documents they plan to submit as a Request for Comments (RFC). Submissions to the Directories should be sent to "**internet-drafts@nri.reston.va.us**".

Internet Drafts are not an archival document series. These documents should not be cited or quoted from in any formal document. Unrevised documents placed in the Internet-Drafts Directories 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 Directories with an announcement to that effect for an additional six months.

Following the practice of the RFCs, submissions are to be sent in ASCII. Postscript is also acceptable, however, we still require the submission of a matching ASCII version (even if figures must be deleted) for readers without postscript printers and for on-line searches.

Internet Drafts are generally in the format of an RFC. 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 document should have an abstract section, containing a two-to-three paragraph description suitable for referencing, archiving, and announcing the document. This abstract will be used in the id-abstracts index and in the announcement of the Draft. The abstract should follow the "Status of this Memo" section.

The Internet Draft should neither state nor imply that it is a Proposed Standard. To do so conflicts with the role of the IAB, the RFC Editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the Internet Draft. These are common words in the "Status of the Memo" section and may cause confusion if placed in the title. If the Internet 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 Internet Activities Board as a standards document. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

2. This document will be submitted to the Internet Activities Board as a Proposed Standard. This document defines an experimental extension to the SNMP MIB. Upon publication as a proposed standard, a new MIB number will be assigned. This

is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

3. This draft document will be submitted to the RFC editor as an informational document. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

4. This draft document will be submitted to the RFC editor as an experimental protocol. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

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

Chapter 2

Steering Group Report

2.1 Standards Progress Report

Between the July meeting hosted by BellSouth in Atlanta and the November meeting hosted by Los Alamos National Laboratory, there have been many IETF originating protocols and informational documents published as RFC's.

In preparation for the many upcoming routing protocol documents, the IESG published a checklist for advancing routing protocols.

RFC 1264	Internet Routing Protocol Standardization Criteria

The Open Shortest Path First Version 2 was elevated to Draft Standard status.

RFC 1245	OSPF Protocol Analysis
RFC 1246	Experience with the OSPF Protocol
RFC 1247	OSPF Version 2
RFC 1253	OSPF Version 2 Management Information Base Several versions of the MIB were released prior to RFC 1253.

The Border Gateway Protocol Version 3 was elevated to Draft Standard status.

RFC 1265	BGP Protocol Analysis
RFC 1266	Experience with the BGP Protocol
RFC 1267	A Border Gateway Protocol 3 (BGP-3)
RFC 1268	Application of the Border Gateway Protocol in the Internet
RFC 1269	Definitions of Managed Objects for the Border Gateway Protocol

Network Management

RFC1270	SNMP Communications Services
RFC1271	Remote Network Monitoring Management Information Base
RFC1272	Internet Accounting: Background

Internet Control

RFC1254	Gateway Congestion Control Survey
RFC1256	ICMP Router Discovery Messages

2.2 Minutes of the Open Plenary and IESG

Agenda:

- IETF Protocol Actions
- IAB Meeting Report
- Open Plenary

IETF Protocol Actions

X.400

Two X.400 documents were presented for Proposed Standard status. These were "Mapping between X.400(1988)/ISO 10021 and RFC 822" and "X.400 1988 to 1984 Downgrading". The IESG had discussed these documents and found them to be necessary pieces to deploying an X.400 infrastructure in the Internet. These documents were discussed within the Rare community and were found to be reasonable. No objections were raise in the Open Plenary session.

Router Requirements

The Router Requirements documents were not ready for "Prime Time". A presentation was made by Philip Almquist earlier in the week. The intention was to issue the set of documents one last time after the November 1991 IETF, and submit them to the IAB before the March 1992 IETF.

IAB Report

The Internet Activities Board (IAB), the parent organization to the IETF, held a meeting during the IETF Plenary. Vint Cerf gave a report of that meeting to the Plenary. In brief, the IAB approved the IESG recommendation that OSPF be designated as the "common" Internet Interior Gateway Protocol. The IAB also approved the long awaited RFC 1108 DOD IP Security Option as a Proposed Standard.

Digital Equipment offered the SPX authentication technology to the IETF. A letter was written by DEC granting the IAB change control of this protocol. At this weeks meeting the IAB accepted this generous offer.

With the observation that the DNS is still not being used by several communities who continue to rely on the hosts.txt file provided by the nic.ddn.mil, the IAB agreed to encourage transition away from the static tables and towards the more modern DNS.

Open Plenary

RFC Copyright Issues

The IETF mailing list prior to the Plenary meeting became quite active on the subject of assigning copyright of the RFC's. Vint Cerf made a brief presentation explaining that the IAB and the now forming Internet Society are working on protecting the spirit of the RFC's. No final decision has been made as to the legal status of the RFC's. Several goals were identified and discussed. Among the goals of this effort are:

1. Preserve the freely available nature of the RFC's. Once published they should remain freely reproducible. Current copyright law can be read to require that the author be contacted for permission to reproduce the document. This was seen as burdensome.

2. Protect against someone writing a "bogus" document and calling it an Internet RFC.

3. Protect against actions which would limit access to RFCs.

IETF/Internet Society Relationship

The Plenary discussed the relationship between the IETF, ISOC and INET '92. While the formal arrangements have not yet been completed, ISOC is intended to be the umbrella organization under which the IAB and IETF will make standards. Membership in ISOC will not be required to participate in the IETF standards making activities. The INET '92 conference is a technical conference, not a plenary meeting of the Internet Society proper. Internet Society business will be conducted on-line, except for balloting which requires authenticated mail. The INET conference does not directly impact the standards making activities of the IETF.

The Standards Process

Bill Simpson began a discussion of the current standards process. He expressed concern over the multiple levels of approval required for standards. He proposed a radical restructuring to include:

1. Elimination of the IAB as an active approver of standards to be replaced by an ISOC rubber stamp,

2. Shortening the "last call" process by overlapping the "last call" and two week Internet Draft posting period, and

3. Eliminating the Draft Standard Stage for protocols in favor of a simple proposed to full standard progression.

These ideas were not adopted but led to further discussion of the current standards process. Among the many ideas discussed was the relationship between the Internet Drafts and the various levels of standardization to the implementation and product development cycles. The only concrete advice offered in the Plenary was to not ship products claiming to be a standard based on an Internet Draft, since an Internet Draft is not a standard. Other than that, the implementation of Proposed and Draft Standards is at the discretion of the company.

The utility of Draft Standards was reiterated by many attendees who saw a Draft Standard as a signal to implement. A Draft Standard has a demonstrated high level of confidence while still not having met the full requirements of a Full Standard.

The popular misperception that all RFCs are Standards was discussed at length. Again, attendees expressed their frustrations with this association, particularly on the part of marketing folks and writers of Request-for-Proposals. Toward the goal of eliminating this association, a new RFC sub-series, called an STD has been proposed. If adopted, all standards will be given an STD number, and that number will remain constant through the various versions as the protocols travel through the standards process.

The requirements for Proposed Standard were discussed. There is a wide mis-understanding and a lot of folklore surrounding the implementation requirements needed for an Internet Draft to become a Proposed Standard. This confusion is beginning to affect the timeliness of documents. A Proposed Standard is required only to have a credible specification, and to have demonstrated a significant constituency. Implementation is simply the best, but not the only, means to demonstrate a credible specification. What this requirement means varies from protocol to protocol and area to area. Routing protocols have a well specified set of criterion, which in recognizing the complexity of routing protocols requires an implementation to become a Proposed Standard. The Network Management Area follows a similar principle for MIB's, but it is less formal and not documented. Other protocol areas are more ad-hoc, but in general, implementations are not required.

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UCDavis

Area Overview

The Applications Area of the IETF is moving to bring multimedia capabilities to the Internet. One Working Group in particular, The Internet Message Format Extensions Working Group (822ext), has made great strides in this direction. This Working Group is finishing the specifications to allow email to have multiple parts to the message where each part may be text, image, audio, video or other types of information to be presented to the end user. The Network News Transport Protocol Working Group (nntp) is working closely with the new message format to bring these capabilities to the network news world. The Teleconferencing BOF explored the idea of desktop video conferencing. The general goal of the area is to define the protocols to create an interoperable multimedia distributed computing environment for the Internet.

Internet Message Format Extensions

The Working Group is finishing the document on multi-part mail messages that will replace RFC 822, and plans to submit the document as a Proposed Standard in early January. This will complete the work of the group.

Internet Mail Extensions

The Working Group has a new Chair, John Klensin (MIT). The Group had to decide if progress could be made towards a method to allow eight bit characters in SMTP. The Group decided to define a means for negotiating the transport of eight bit characters. It was thought that the method could also be useful for negotiation of other items, such as allowed message size.

Network News Transport Protocol

This Working Group did not meet in Santa Fe, but has been making good progress on the mailing list and has a document about ready to be issued as an Internet Draft.

Automated Internet Mailing List Services

Unfortunately the Chair of this Working Group, David Lippke, had to resign due to a reassignment of work duties. The Group will be on hold until a new Chair is found.

Network Fax

The Working Group finished work on the image format to be used for transporting FAX on the Internet. The Document will be available as an Internet Draft soon.

Network Database

The Working Group continued work on the definition of SQL transactions over TCP/IP networks. The Group is small and there needs to be involvement from other SQL implementors.

TELNET

The Working Group made further progress on authentication and encryption for TELNET sessions. It was decided that authentication and encryption need to be closely tied together in operation.

Teleconferencing

At this BOF several individuals presented work being done on teleconferencing over the Internet. After the presentations there was discussion on how the problem can be broken in work that can be done by various working groups. One working group was created to define methods for real-time transport of audio and video.

CURRENT MEETING REPORT

Reported by Russ Hobby/UC Davis

Minutes of the Teleconferencing BOF (TELECONF)

There has been considerable discussion about using the Internet in support of remote conferencing. Ideas range from using the network to provide a shared workspace, such as common document viewing and editing, to full motion video conferencing. The BOF in Santa Fe was intended to bring together people that are doing current research in the area and to see if a common direction can be found. Several researchers presented their work.

Steve Casner/ISI presented an overview of multimedia conferencing including work that has been done on the Terrestrial Wideband Network and DARTnet. Paul Milazzo/BBN talked about their efforts with workstation based video windows. Hans Eriksson/SICS told the Group about their efforts with the MultiG - research program which is working toward a Collaborative Desktop and Telepresence. Sze-Ying Wuu/Bellcore explained the workings of their Touring Machine, as a system for the management of mixed multimedia. Yee- Hsiang Chang/MCNC told the Group about their project to use packet video for use on the Concert network. Peter Kiestein/UCL provided information on the meeting on International Multimedia Conferencing held at ISI on November 13, 1991.

After the presentations there was discussion about the various parts of multimedia conferencing and how the job of creating a working system may be broken into manageable tasks. It was suggested that there are four areas of work:

1. Shared Workspace. This includes things like shared whiteboards, editors and generalized windows for viewing other applications.

2. Conference Management. This concerns conference setup, connection management and coordination for the various parts of the overall conference.

3. Transport Formats. This will define the formats of media (i.e., video, audio, image) and how they are to be transported over the network.

4. Data Delivery. This area needs to address the service guarantees needed by real-time data and to provide reliable multicast capabilities.

A new working group, chaired by Steve Casner, was created to coordinate work being done on the transport formats of video and audio over UDP. It was recognized that UDP does not provide the necessary service guarantees for real-time data. However, it was viewed that useful work could be done over UDP on lightly loaded networks until a better means of data delivery is made available.

It was felt by the Group that the IETF should continue work in these areas and should be coordinated with other standards groups.

Attendees

Harald Alvestrand	herald.alvestrand@delab.sintef.no
James Beers	beers@nr-tech.cit.cornell.edu
David Borman	dab@cray.com
Robert Braden	braden@isi.edu
Scott Brim	swb@nr-tech.cit.cornell.edu
Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@concert.net
Richard Cogger	rhx@cornellc.cit.cornell.edu
Steve Deering	deering@xerox.com
Barbara Denny	denny@sri.com
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Hans Eriksson	hans@sics.se
Farrell Gerbode	farrell@rice.edu
Ittai Hershman	ittai@nis.ans.net
Russ Hobby	rdhobby@ucdavis.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Cheryl Krupczak	cheryl@cc.gatech.edu
Ruth Lang	rlang@nisc.sri.com
Eliot Lear	lear@sgi.com
Brian Lloyd	brian@ray.lloyd.com
E. Paul Love	loveep@sdsc.edu
Paul Milazzo	milazzo@bbn.com
Greg Minshall	minshall@wc.novell.com
David O'Leary	oleary@sura.net
Ari Ollikainen	ari@es.net
Joe Ragland	jrr@concert.net
Daisy Shen	daisy@watson.ibm.com
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
Andrew Veitch	aveitch@bbn.com
Sze-Ying Wuu	syww@thumper.bellcore.com

A Roadmap to Internet Multimedia Conferencing

Where are we now?
Where do we want to go?
How do we get there?

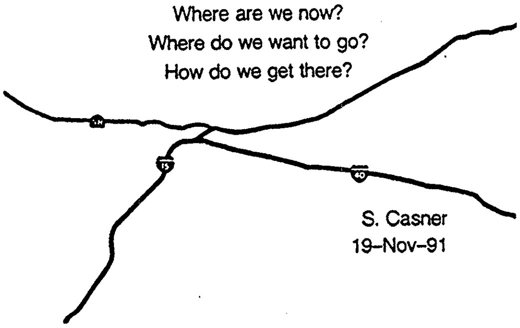

S. Casner
19-Nov-91

Where are we now?

- "Experimentally operational" conferencing on TWBnet
- DARTnet packet audio & video, resource mgmt experiments
- SICS MultiG project: packet audio & video, shared workspace
- UPenn has reported packet audio across Internet
- InterOp demo of remote radio from Australia
- There must be lots of LAN hacking with SPARCs, NeXTs ...
  - need to organize/standardize some of this activity

Terrestrial Wideband Net and Conference Sites

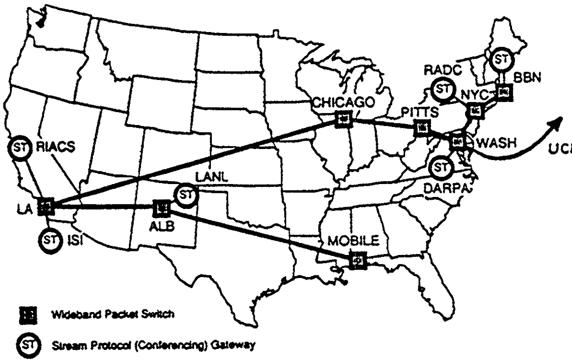

The DARPA Research Testbed Network (DARTnet)

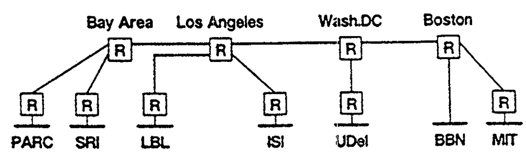

T1 (1.5Mb/s) lines: Cross-country spine + tail circuits

Routers: Sun SPARCstation 1+

Packet Voice and Video Hardware

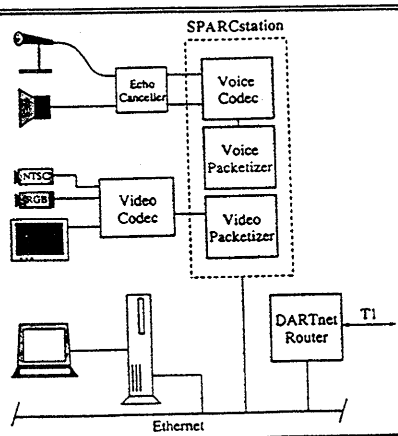

Conferencing Experiments on DARTnet

- Weekly audio conferences using UDP and IP multicast
- DES encryption added for privacy with IP multicast
- Packet audio using ST-II also tested on all nodes of DARTnet
- Packet video between ISI and BBN using ST-II on SPARCs
- Packet video between MIT and PARC using UDP on 386 PC

Where we are

- High cost -> few sites on TWBnet -> shared rooms
- "Personal Conferencing" on workstations is being developed
- Experiments are planned on DARTnet in several areas

Where do we want to go?

Dave Farber: *"Maybe a good activity for the IETF (the community) would be to make the technology capable of supporting such a meeting [IETF]."*

Jack Haverty: *"How can Internet technologies, in the fuzziest definition, be used to support the processes involved in the IETF activities?"*

▸ Widespread personal teleconferencing over Internet

How do we get there?

E-mail discussion showed 3 areas of interest for possible WGs:

- Enable distributed IETF meetings for reduced travel and increased international participation — may use existing technology
- Shared workspace, plus telephone initially
- Develop real-time packet audio and video over Internet

- Connection (or session) management

Enabling Distributed IETF Meetings

- Explore the use of commercial services, covering the costs (using the Internet would be great, but is not necessary)
  - Start with speakerphones to bring in non-attendees?
- Can we find a set of sites capable of multiple simultaneous WG teleconferences? Is that the right model?
- How to manage time zone differences: Async conferencing?
- Some human factors work (I/O devices, mics, echo cancellers) may apply to IETF meetings and personal conferencing, too

Shared Workspaces

Some existing systems may be ready for deployment:

- MMConf (used in TWBnet system) — replicated architecture
- Shared X (HP, DEC, Bellcore, others) — centralized arch.
- SICS Mdraw — open floor, ISIS for consistency
- Commercial products (PC links, MacMICA, Aspects)

Plus group activity applications to be built upon these platforms.

▸ Tablets or touch screens may be important

Real-Time Packet Audio and Video

Requirements in several areas:

- End-system hardware, especially video codecs
- Development and deployment of high-BW nets (T3, ATM)
- Network protocols for real-time services (at IP level):
  - Resource mgmt. in hosts, on LANs
  - Resource mgmt. in routers, intra-domain and inter-domain
- Audio/video transport/application protocols
- Operating System scheduling for real-time processes

Video Codecs are a Roadblock!

- Box-level codecs (e.g., PictureTel) are too expensive
- Concept is good for experiments but is a dead-end product
- ISI has two DIME board prototypes from Sun (DVI chipset)
- JPEG chips are in NeXT, Parallax but don't compress enough
- H.261 and MPEG coding standards are coming, then real products built into workstations
- Can do low-cost, low-performance video with frame grabbers (e.g., VideoPix) and simple software compression initially

Need Cooperation Among Researchers and Vendors

- We can't do it all...
  - Need workstation mfrs to build in audio and video compression plus camera, video windows, mic, speaker
- Workstation vendors can't do it all either...
  - Need widespread, standardized network protocols to allow interoperation and achieve critical mass
  - Not motivated to make heterogeneous configurations work

Protocols for Real-Time Services: Under Construction

- Soft state and hard state schemes for resource management
  - Van Jacobson and Dave Clark (soft state)
  - ST-II (hard state) plus resource management mechanism
  - Lixia Zhang's Flow Protocol is an example mechanism
  - Berkeley (Anderson) Session Reservation Protocol (SRP)

Must scale:

- To wide deployment
- To gigabit speeds

Take a Shortcut: Start with UDP

- UDP will work fine in some places at some times: T3 backbone, lightly-loaded T1's and Ethernets
- Use simple TOS-based priority and IP multicast where available
- Buffers can accommodate seconds of delay; people might not! (8 KB per second for PCM audio)
- Congestion discard is a bigger problem, need loss < 1%

Need a Highway Patrol!

- How do we control usage with no flow control mechanism?
- Well-behaved TCPs will pull aside for the UDP road hogs
- Usage seems likely to grow until quality degrades, leaving a service that's not useful

Audio / Video Transport / Application Protocols

- Several UDP-based implementations already exist, using incompatible header formats
  - First step is to design a common header format
- Is there a light-weight, real-time transport protocol hiding here?
- Integration of synchronization mechanism for inter-media and/or inter-site synchronization (e.g., BBN's Synchronization Protocol)

Next Stop: Connection (Session) Management

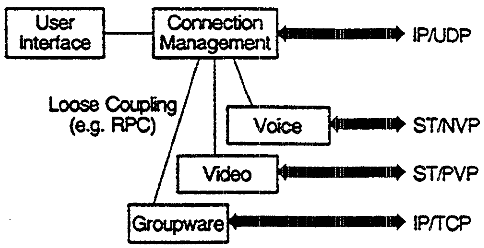

The Roadmap

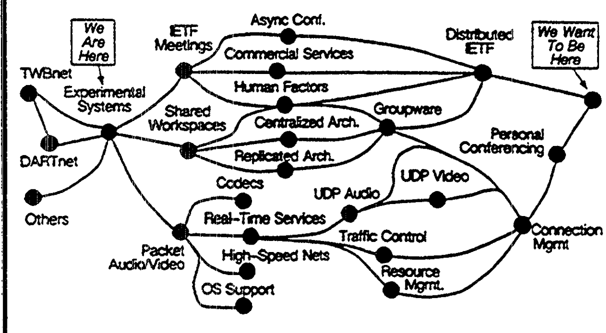

Summary

Enabling technology is both hardware and software:

- DVI, H.261, MPEG video codec chips and boards
- T3 and ATM to provide high-bandwidth network infrastructure, but we need resource management protocols, too
- Shared Workspace development and standardization
- Connection management protocols and user interfaces to make conferencing easy to use
- Operating system support to integrate all of this

Hans Eriksson
Swedish Institute of Computer Science

Videoconferencing in MultiG

- MultiG — research programme
- Current implementations
- Pilots
- Different Scenarios – different compromises
- Research interest

MultiG — research programme

- Collaborative
  - SICS, KTH, Ericsson, Televerket, ...
- High-Speed stuff
  - n*Gb/s
- from the fiber to the application
  - host interfaces
  - network/transport protokols
  - distributed multimedia support
  - multimedia applications
    - 2D: Collaborative Desktop
    - 3D: Telepresence (virtual reality)

Current implementations

- Sharing workspace
  - Mdraw
    - N-way – X11 – ISIS
- Personal Communication
  - PicturePhoneTalk
    - 2-way – video/audio – TCP
  - Teleconference (Dec 16)
    - N-way – video/audio – UDP
- CoDe
  - protype of some tools for CSCW
    - audio, group-map, answering machine
- Combo
  - PPTalk
  - MDraw (telepointer)
- Later
  - Telepresence conference

Pilots

- SICS — KTH
  - Ether – FDDI – Ether
- SICS — Ericsson — Telia
  - local FDDIs interconnected with 34Mb/s
- SICS — Oslo Univeristy
  - via NORDUnet (256kb/s)
- UCL, ISI, DARTnet ...

Different Scenarios – different compromises

- Catching TV broadcast
  - real-time – only one chance
- Conference
  - Brainstorming
    - high interaction -> low latency
  - Collegues
    - *what* they say
  - Negotiations
    - *how* they say it
- Lecture
  - low interaction -> latency not important
  - pupils-teacher -> video important -> high bandwidth
- Meeting
  - a mix of the above
  - a sequence of mini-lectures

Research interest (HE)

- dynamic adaptability
  - varying loads (network, cpu, ...)
- multimedia synchronisation
  - lip–synch
- Striking balances
  - brainstorming
  - conferences...
- Heterogenous networks
  - mix of FDDIs, ethers, 64s...

CSCW UI
operating support
ST-2 impl
3D...

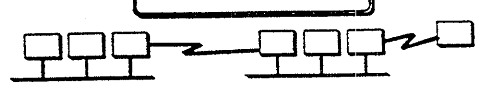

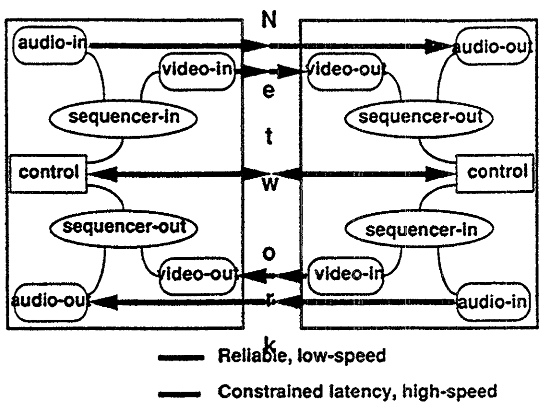

TOURING MACHINE:

Distributed Systems for Multimedia Communication

Mauricio Arango, Peter Bates, Jane Cameron, Brian Coan, Gita Gopal, Nancy Griffeth, Gary Herman, Takako Hickey, Will Leland, Victor Mak, Lillian Ruston, Mark Segal, John Unger, Mario Vecchi, Abel Weinrib, Doris Woods, Sze-Ying Wuu

Bellcore

New Touring Machine Features

Applications Programming Interface (API)
- "language" for writing multimedia communications applications
- reflects separation of application *policy* from network *mechanism*

Separate control of media
- audio
- video
- data

Fully integrated name server
- name and access transient objects (e.g., communications sessions)

Rich network infrastructure
- multiple switches (routing, etc.)
- allocation of specialized hardware (e.g., bridges)

New Applications
- Multimedia Telecommunication Service (MTS)
- CRUISER™ service (2146)
- RENDEZVOUS™ system (2146)
- Touring Messaging
- Match Maker

AW 11/15/91

Software Architecture

<u>Long-lived objects</u>

Station Manager (optional)
- implements resource-sharing policies among clients

Station Object
- provides interface to Touring Machine
- manages station ports

Resource Manager
- allocates physical resources

Resource objects
- control physical resources

Name Server
- repository for static and dynamic system information

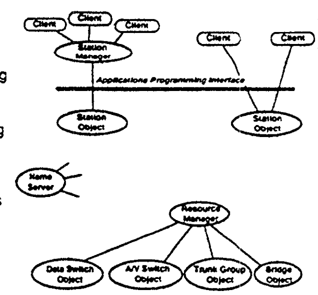

AW 11/15/91

Software Architecture

<u>Transient Objects</u>

Session Object
- site for negotiation between clients
- maintains logical state of session

Transport Object
- maintains logical-to-physical mapping for session

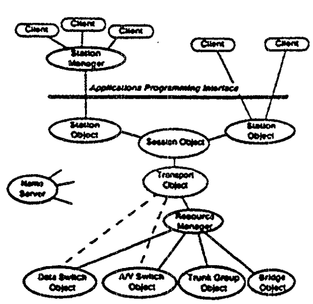

AW 11/15/91

Application Programming Interface

Client registration

```
(registerClient <token> <clientName> <regAction>+)
(registerChange...)
```

- initiate and authorize client interaction with Touring Machine
- register endpoints (audio, video, data)

Session establishment and modification

```
(sessionCreate <token> <sessionName> <clientID>
<sessionAction>+)
(sessionChange...)
```

- associate clients, negotiate "call"
- establish connectors between sources and sinks

Local resource control

```
(endpointMap, endpointUnmap, portCreate,
endpointAssign)
```

- map and unmap endpoint to assigned port
- create port (data)
- assign endpoints to ports

Name server queries `(nsQuery... <keys> <attributes>)`

Inter-client message forwarding `(messageSend)`

Error notification `(errorNotify...)`

AW 11/15/91

Packet Video Project

Packet Video for Videoconference Project at MCNC

Yee-Hsiang Chang, Ph.D.
Resident Scientist, Communications Research
MCNC Center for Communications

email: yhc@concert.net

MCNC

Packet Video Project

Private non-profit corp
- To promote the growth of education and research in North Carolina institutions
- To foster the economic development of North Carolina

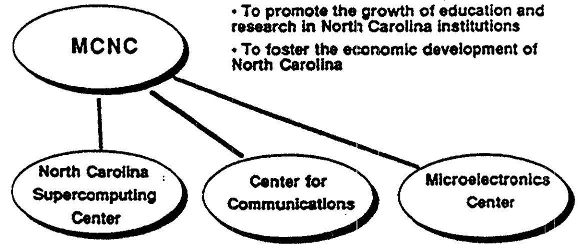

supercomputer
- operations
- research
- education

data and video network
- operations
- research

CMOS VLSI
- fabrication
- process
- design tools
- packaging

CONCERT

Packet Video Project

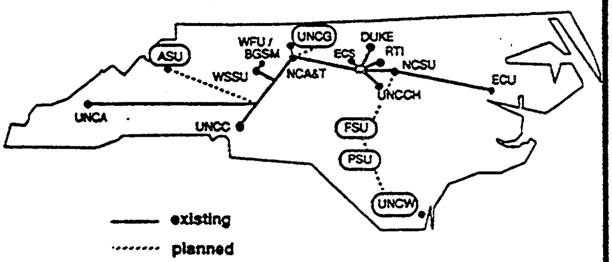

Communications Research Programs

Packet Video Project

- VISTAnet, gigabit network testbed
- Shared workstation, X windows shared space
- Packet Video, remote scientific visualization
- ATM, supercomputing gateways

Virtual proximity
Networked access to supercomputing

MCNC Center for Communications: CONCERT Data Network

Packet Video Project

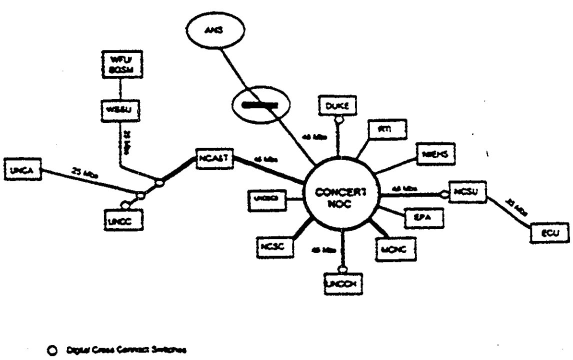

MCNC Center for Communications: *CONCERT* Video Network

Packet Video Project

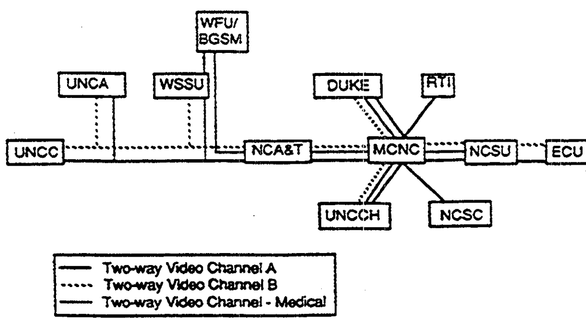

Multimedia Conferencing Using Packet Video over the Internet

Packet Video Project

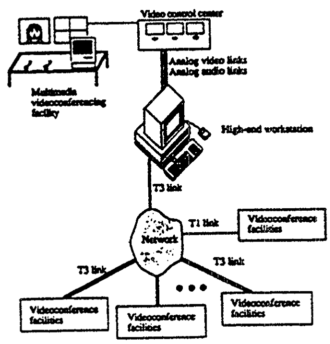

Objectives

Packet Video Project

- Deliver current CONCERT analog videoconference capability in digital Internet T3 environment, using standard CODEC equipment.
  - Investigate subset capabilities for T1 and sub-T1
  - Investigate workstations as multimedia devices
  - Accomplish in five phases over a two year period
  - Start in January, 1992
- Perform work through a consortium of industry and university members.
- Utilize emerging video standards and demonstrate multi-vendor compatibility/interoperability.
- Stimulate development of video application platforms and applications.
- Influence the development of commercial workstations as effective multimedia devices.

Schedule

Packet Video Project

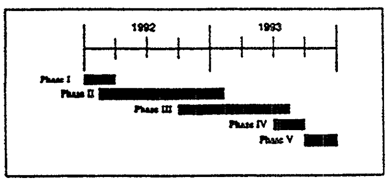

Phase I:
- Equipment and laboratory setup. It also includes the task of writing software that will serve as programmable impairment inducer.
- Performance measurement for improvements.
- Traffic pattern measurement for different compression schemes.

Schedule (Continued)

Packet Video Project

Phase II:
- Resource allocation mechanisms, latency control, and scheduling algorithms at routers.
- Session layer software.
- Presentation layer software.
- Error control at transport layer.

Phase III:
- Reliable and unreliable IP multicast.
- IP multicast routing
- Resource allocation and latency control (continued from Phase II for further refinements).
- Presentation layer software (continued from Phase II for further refinements).

Phase IV:
- Transport protocol comparison.
- Application software.

Phase V:
- Hardware specifications of the new multimedia workstation.
- Software specifications (e.g., operating systems and X window) for the multimedia workstation.

Phase I

Packet Video Project

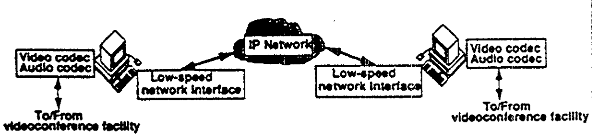

- Evaluate available codecs based on the perceived picture quality, delay, and the traffic pattern.
- Identify the areas for improvement.

Phase II

Packet Video Project

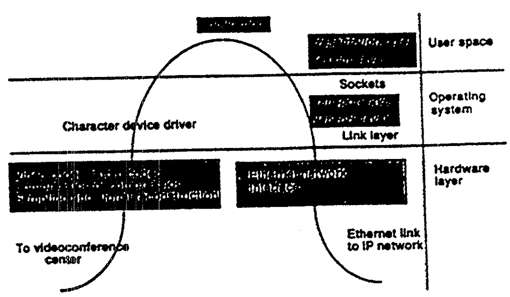

Phase II (Continued)

Packet Video Project

- Application layer: simple program to read/write data from character device driver/network interface.
- Presentation layer: data buffering, smoothing, and synchronization for the received data, and data conversion to/from external data representation for the sent and received data .
- Session layer: videoconferencing session initiating, terminating, joining, leaving and provisioning. The provisioning function determines the type of equipment at all sites and the associated network bandwidth for the quality of service.
- The session and presentation layer software will be built into multimedia application libraries.
- Transport layer: error control for multimedia application (e.g., forward error correction).
- Network layer: resource allocation, and latency control mechanisms. We will investigate and implement the scheduling algorithm at routers.

Resource Allocation and Latency Control

Packet Video Project

- Build on the already installed base of IP
  - The connection-oriented vs. connectionless issue is irrelevant to the resource management function (DARTnet).
- Delay tolerance
  - 400 - 600 ms round trip delay
  - Four sources
    - coder delay: time to code/compress or decode/decompress
    - signal propagation: time to travel down the wire
    - switch buffers: delay through the router queues
    - jitter compensation: max time each packet must wait in a buffer to maintain isochronous delivery to the application

Resource Allocation and Latency Control (Continued)

Packet Video Project

Delay Prediction: A single queue to represent T3 network

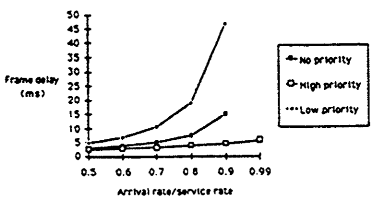

- Adding a priority mechanism greatly helps the resource availability.
- Delay will be limited if the real-time traffic is limited to a portion of the total traffic.

Phase III

Packet Video Project

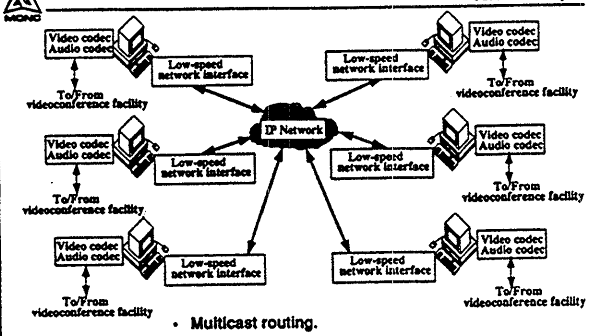

- Multicast routing.
- Reliable and unreliable multicast.
- The impact of congestion on multicast.

Phase V

Packet Video Project

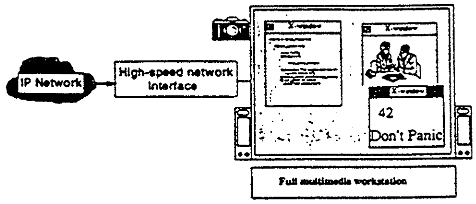

- Hardware specifications of the new multimedia workstation.
- Software specifications (e.g., operating systems and X window) for the multimedia workstation.

Phase IV

Packet Video Project

- We expect better codecs are coming out and we will try them.
- Evaluate other transport protocols such as VMTP and XTP on top of IP.
- Application program for network management and control: this software will enable the network operator to control bandwidth, routing, and video resolution at multiple site.

3.1.1 Distributed Scheduling Protocol (chronos)

Charter

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

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

Description of Working Group:

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

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

Goals and Milestones:

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

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

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

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

3.1.2 Internet Mail Extensions (smtpext)

Charter

Chair(s):
John Klensin, `klensin@infoods.mit.edu`

Mailing Lists:
General Discussion: `ietf-smtp@dimacs.rutgers.edu`
To Subscribe: `ietf-smtp-request@dimacs.rutgers.edu`
Archive: `~ftp/pub/ietf-smtp-archive:dimacs.rutgers.edu`

Description of Working Group:

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

Goals and Milestones:

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

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

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

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

Mar 1992 — Review and finalize the SMTP Extensions document.

Mar 1992 — Submit the SMTP Extensions document as a Proposed Standard.

Internet Drafts:

"SMTP Extensions for Transport of Enhanced Text-Based Messages", 07/10/1991, John Klensin <draft-ietf-smtpext-8bittransport-02.txt>

CURRENT MEETING REPORT

Reported by John Klensin/MIT

Minutes of the Internet Mail Extensions Working Group (SMTPEXT)

The meeting on the 19th was long, and very intense. The result was a narrowing of the focus on the transport extensions. The meeting began with Greg Vaudreuil introducing John Klensin and handing over the Chair. Klensin then announced that the Meta-Agenda for the day was to either focus sufficiently so that a clean plan and schedule could emerge or be able to report that the Working Group was going nowhere and should be abandoned.

Klensin then introduced a decision model for the major options facing the Working Group. That model was refined somewhat in Group discussion and appeared as follows:

```
           Negotiate
      -----------------------------------------------------
       |              |                      |             |
     Yes             No              Can't decide     Decide not to
       |          ----|-----------      -----|             |
     --|----     ?*    Prime     |      |                  |
     |     |         ---|--      |      |             ?    |
 RFCZZZZ  ?*       Prime   |     |      |<-----------------|
   bis               bis   |     |------|                  |
                      |    |         |                     |
                      v    v    Let 1000 flowers           v
                      @    @       bloom            New protocol ?*
                                     @
```

Notation:

?* -> Need to make a plan

@ -> Need to consider--or abdicate--damage control for ''old'' systems, especially wrt blowups and information loss.

''Prime'' refers both to the specific proposal from Robert Ullman and David Robinson to the class of proposals who share the concept that existing ''8 bit clean/8 bit transmitting'' implementations are acceptable, should be encouraged, and should be joined by others.

Considerable discussion ensued. There was no sympathy expressed for the sending of unnegotiated 8bit data as a long-term strategy, but a general understanding that it was undesirable to leave existing implementations that do that without plausible transition paths.

The Working Group concluded that the receipt of data with the 8th bit set but without negotiation was an error, and proceeded to analyse the error states. The conclusion was that an originating SMTP client was non-conforming if it transmitted any data with the 8th bit set without prior negotiation and agreement. A destination server receiving such a message could respond in one of three ways and be conforming:

1. Reject the message as an invalid transport case, presumably using a 520 error code.

2. Deliver the message in 8bit form. This option requires that the MTA "know" that such delivery can be accomplished accurately (i.e., without loss of information). This would normally be the case when both delivery MTA and UA were in a "8bit clean" environment.

3. If sufficient information is available, downgrade the message to 7bit RFC-XXXX. Since the Working Group did not consider it acceptable to "guess" at what the character set might be, or to make an assumption based on, e.g., the sending or receiving country, the "sufficient information" condition will in general be met only if the incoming message is already in valid RFC-XXXX format.

If a message with leading bits set arrives at a relay host without prior negotiation, the relay has the additional option of transparently forwarding that message. The destination host is no worse off in this case than it would be had the message been sent without the relay. In other words, the Working Group agreed that there was no significant benefit in imposing additional requirements on relays for policing protocol conformance. Relays would, of course, retain the options of rejecting or downgrading, as provided in (1) and (3) above.

There was then general agreement that "doing nothing" was undesirable. For some people, the above analysis was acceptable only if the Working Group proceeded to define and agree upon a negotiation model; others were convinced that the analysis and agreement was useful in itself.

The various large scale options of RFC-ZZZZ (November 6th draft) were then reviewed, with backward references to the pre-St. Louis version of that document. The options of "new protocol" and "move more rapidly toward X.400" were raised as alternatives, but quickly dismissed in the context of the current charge of the Working Group, since they do not address the very real issues of existing 8bit transport over existing ports and protocols.

The session on November 21st began with a review of an intermediate draft of RFC-ZZZZ which Klensin had prepared to incorporate the changes agreed to on the 19th. The meeting then went through an interim "outstanding issues" list, eliminating many of the issues and deferring others. As one might expect, some issues were controversial, others were not. A review of the interactions between SIZE and the capabilities concept introduced on 19th led to the partial restoration of the former while retaining the latter.

The morning's greatest controversy was over exactly what requirement to impose for the capability for conversion from 8bit to 7bit transport forms in mail relays. The issue is complex because it is seen by some as an issue of keeping mail relays simple and, in particular, not requiring that each one have gateway capability, and by others as an issue of increased mail interoperability (or of avoiding decreased interoperabiliy). After lengthy and sometimes heated discussion, it was agreed to adopt a rule designed to reduce as much as possible the chance of deferred rejection of 8bit mail as a result of encountering an 8->7 boundary. A host accepting 8bit mail is not permitted to have the mail later rejected as a result of a conversion requirement. This means, in essence, that any host accepting 8bit mail must either be able to guarantee (through out-of-band information) that it can make final 8bit delivery to the addresses in the message, or must be prepared to arrange for conversion to seven-bit form. The Working Group understands that the conditions for guaranteeing an unobstructed 8bit path can rarely be met in practice and that this requirement means that a mechanism for conversion to 7bit forms is therefore essentially a requirement of a host that is implementing server support for the SMTP "emal" verb. Probably the only exception that does not depend on considerable out-of-band information and very early verification of addresses would be for a server that supported only local delivery, with no capability for relaying, automatic forwarding, or providing mail exchanger services for other hosts.

There was then a discussion of newly-written "packetized data stream" and "binary" proposals by Neil Katin. The discussion of the former was carried far enough to reach general agreement on a model: sending and acknowledgement (in a request-and-wait mode, paralleling DATA) of a "packet mode" command. If that command is accepted, the sender can send packetized streams of data using an introducing "packet N" command followed by N octet of data without regard to line lengths or delimiters. Each packet would be acknowledged by the server, but the model is designed so that these acknowledgements can be handled asynchronously by the client (permitting batching). After each such packet, the server would expect to receive either another "packet N" command; the "packet 0" command, indicating end-of-data; or RSET or QUIT. Lengths of packets would be as chosen by the sender. The question of need for a receiver-imposed maximum packet length was discussed. It was finally concluded that such sizes were not an issue given TCP buffering capability; the issue will be revisited if anyone can identify a case in which server-imposed restrictions are actually needed.

Agreement was reached in principle on incorporating packetized data stream (as described above) and binary mail. Joint work with the 822-Extensions Group was done to provide additional specifications for the handling of error messages that must be mailed back to the sender (rather than reported as part of the SMTP transaction). These efforts will be incorporated into RFC-ZZZZ if they converge rapidly enough and are appropriate; otherwise they will be handled as separate documents.

Specific conclusions about RFC-ZZZZ were:

1. There is, at present, no real demand for transport forms wider than 8 bits or for addressing the issues such transport would cause. The question will be revisited when and if there is a requirement for such transport.

2. It is important to clarify and establish an extension model for RFC821 now, even if no substantive changes were incorporated into that extension model.

3. There is no demand for 8-bit versions of SOML and SAML, since it is unlikely that anyone would really want RFC-XXXX messages delivered directly to their screens. An 8-bit version of SEND FROM is problematic, since such messages are typically transported without headers, leaving ambiguitites about the character set in use as soon as the characters are not clearly ASCII. If and when there is demand and a definition for an enhanced SEND, an extension can be proposed and considered.

4. As a result of (1) and (3), the marginal "cost" of a new transport variation (e.g., binary or ESND) becomes one verb, not four verbs. And, since there is willingness to defer extended-width (past 8) entirely and predict that it will not be needed, the complexities and additional states associated with the TYPE verb can be eliminated by getting rid of that verb. This, of course, implies that EMAL limits the message being transported to being one in which ASCII (with a leading zero bit) can be successfully used in trace fields. That does not appear to be a severe restriction in practice, regardless of the theoretical possibilities.

5. While the concept of a SIZE inquiry is desirable, it was felt that several other inquiries may be useful also and that it was not desirable to worsen the query-and-wait transaction model. Consequently SIZE (as an inquiry) is to be removed and replaced by a capability inquiry (CPBL) to which a server would return such information as what size messages were normally acceptable and what other options were supported in a canonical way. The format of the canonical response awaits further definition, although there was sympathy for something of the attribute=value character. There was also discussion about the implications of denial of the availability of a service without general agreement other than a client should not "try anyway" if some capability were explicitly denied. There was also a discussion of the fact that some hosts might wish to avoid giving out capability information as a security measure in order to avoid disclosing operating system or similar information. This may imply that hosts should be able to respond to a capability request by explicitly asserting certain services, by explicitly denying them, or by providing no information (in which case the client would normally behave as if the inquiry had not been made).

6. The SIZE verb, used to alert the server of the approximate size of a file that is about to be transmitted, is retained. This verb serves two main purposes: early rejection of large messages, rather than having to transmit them first and providing receivers

some ability to prepare for large messages. The latter may actually permit larger messages to be delivered.

7. Additional text should be put into the document that explicitly identifies the results of experiments with existing servers relative to handling of unknown verbs and recommending behavior if commands are refused with syntax errors.

8. The explanatory/discussion sections should be retained, although we may wish to start identifying those that are intended for a final document separately from those which are to be retained only during discussion.

9. Support of EVFY is required of any server that supports EMAL and support of CPBL is required of any server that supports any enhanced capability (beyond those of SMTP). For the latter, "support" is defined as the ability to return useful information on which the client is expected to take action. Mechanisms for CPBL responses that do not reveal information will be considered only if an explicit request or requirement is received from the security area.

10. While enhanced trace field capabilities and requirements are needed if enhanced mail features are not going to make it appreciably harder to identify and fix problems (it is already bad enough), that material will be removed to a separate document if agreement cannot be reached quickly enough. The Working Group identified one specific concern, which is the need to bind conversion-tracing fields to RFC-XXXX body parts, not whole messages, since some conversions will be performed one body part at a time. The requirement for this body part header has been brought to the attention of the RFC-XXXX authors.

11. The material on RSET and defining new FROM verbs is useful and should be retained. Some textual improvements are needed.

12. CPBL does not accept an argument; the use of one is a syntax error.

    The following issue is considered resolved unless new issues and alternatives are raised. It differs from the above because, rather than being discussed at length, there has apparently been no interest in taking issue with it since the first version appeared in the first Internet Draft version of RFC-ZZZZ.

13. The model for which error/response codes are used in various situations. The placeholder for this has been changed to a "tentative agreement" paragraph.

Summary, Schedule, and Plan.

After discussion with the Applications Area Director and the IETF Chair, we should plan on requesting that RFC-ZZZZ be promoted to Proposed Standard status not later than the end of the March 1992 IETF meeting. It appears after the November 1991 Working Group meetings that there are no "show stopper" issues remaining. There are several issues for which options, details, or explicit text still need to be worked out. Any of those that cannot be worked out and agreed upon by March will be removed from RFC-ZZZZ and handled separately.

A new version of RFC-ZZZZ has been prepared and is being submitted for publication as an Internet Draft. Note that this version supercedes the one announced on the list and circulated to the November 21st Working Group meeting.

Attendees

Mary Artibee	artibee@sgi.com
Nathaniel Borenstein	nsb@thumper.bellcore.com
James Conklin	conklin@bitnic.educom.edu
Mark Crispin	mrc@cac.washington.edu
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Erik Fair	fair@apple.com
Ned Freed	ned@innosoft.com
Olafur Gudmundsson	ogud@cs.umd.edu
Russ Hobby	rdhobby@ucdavis.edu
Christian Huitema	christian.huitema@sophia.inria.fr
Neil Katin	katin@eng.sun.com
John Klensin	klensin@infoods.mit.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@sgi.com
Gordon Lee	gordon@ftp.com
Jack Liu	liu@koala.enet.dec.com
Paul Milazzo	milazzo@bbn.com
Keith Moore	moore@cs.utk.edu
Robert Morgan	morgan@jessica.stanford.edu
Chris Myers	chris@wugate.wustl.edu
Mark Needleman	mhn@stubbs.ucop.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Daniel Newman	dan@innosoft.com
Michael Patton	map@lcs.mit.edu
Robert Purvy	bpurvy@us.oracle.com
Robert Shirey	shirey@mitre.org
Keld Simonsen	keld@dkuug.dk
Ursula Sinkewicz	sinkewic@decvax.dec.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.3 Internet Message Extensions (822ext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done Review the Charter, and refine the Group's focus. Decide whether this is a worthwhile effort.

Done Discuss, debate, and choose a framework for the solution. Assign writing assignments, and identify issues to be resolved.

Done Review exiting writing, resolve outstanding issues, identify new work, and work toward a complete document.

Done Post a first Internet Draft.

Nov 1991 Review and finalize the draft document.

Dec 1991 Submit the document as a Proposed Standard.

Internet Drafts:

"MIME (Multipurpose Internet Mail Extensions): Mechanisms for Specifying and Describing the Format of Internet Message Bodies", 06/18/1991, Nathaniel Borenstein, Ned Freed <draft-ietf-822ext-messagebodies-03.txt, .ps>

"A User Agent Configuration Mechanism For Multimedia Mail Format Information", 06/18/1991, Nathaniel Borenstein <draft-ietf-borenstein-configmech-03.txt, .ps>

"Mnemonic Text Format", 07/08/1991, Philippe-Andre Prindeville, Keld Simonsen <draft-ietf-822ext-qreadable-02.txt>

"Character Mnemonics and Character Sets", 07/08/1991, Keld Simonsen <draft-ietf-822ext-charsets-02.txt>

"Representation of Non-ASCII Text in Internet Message Headers", 11/14/1991, Keith Moore <draft-ietf-822ext-msghead-01.txt>

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

Minutes of the Internet Message Extensions Working Group (822EXT)

Agenda

- Discuss and resolve outstanding issues in Quad-x.
- Discuss and complete the header character set proposal.

Resolve outstanding issues in Quad-X

A list of outstanding issues was reviewed and amended. Note, the term Quad-x was coined for RFCXXXX at this meeting, and is used throughout these Minutes.

1. Audio Format

   The Working Group was presented with two proposals for the format of audio/basic. Both proposals were based on the NeXT audio formats, one had attributes in the content-type headers and the other had the attributes in the file header in the body. After discussion, the Working Group concluded that it had no basis for choosing a standard **# extensible #** audio format and left the work for a future group. The NeXT format was seen by many to be too machine dependent, and had too many options, even as profiled by Marshall Rose.

   A simple format was agreed to for audio/basic which has no options and is not extensible. This definition for audio basic was defined as u- law, 1-channel, 8 khz. The data in the bodypart is straight u-law.

2. Message Integrity Check

   The Working Group expressed a strong need to define a message integrity check for message bodies. This was felt to be more general than would be available by adding a checksum to the base 64 encoding. No clear specification was available at this meeting. In the interests of making forward progress, the Working Group agreed that the absence of a MIC was not a "show stopper", and if a solid proposal is ready, and can be approved by the list by December 16th, it would be included in the document.

   ACTION: Ned Freed and Jim Galvin – Write a MIC proposal to include the preferred MIC as suggested by the Security Area Advisory Group.

3. Multipart/Alternative

   Multipart alternative was enthusiastically endorsed as a transition mechanism to encourage the sending of richer formats than may otherwise be used. By allowing a sender to send both a richly formatted document and include in a systematic way a simpler version, one which may be "cat'ed" to the screen, concern for the lowest common denominator will not have to be a restriction on the use of new features.

4. Character Set Issues

   The Working Group specified the definition of a character set for the purposes of Quad-x to be a unique mapping of a byte stream to glyphs, a mapping which does not require external profiling information.

   (a) ISO 2022-jp

       ISO 2022 is not strictly speaking a character set. It is a switching mechanism which requires an external profile to be useful. The Japanese have defined such a profile, and that profile will be documented and considered a character set for the purposes of Quad-x.

   (b) Mnemonic

       Keld Simonsen's mnemonic proposal as currently written requires the external specification of a character set and an escape character. As such, it does not fit the general requirements of a character set. A lunch sub-group defined a profile for mnemonic, with a lead-in character of "&" (ASCII 38) and ASCII as the default character set. With the profile, the Working Group accepted mnemonic as an acceptable character set for Quad-x.

5. Application Specifications

   The Working Group agreed upon several criterion for the specification of new application subtypes to be defined in the Quad-x proposal. A new application must include in attribute-value pairs, the profile, macro packages used, and any external pre-processors needed to use the included data. The security implications of using the particular applications data without authentication must also be discussed.

   (a) PostScript

       Adobe has defined Postscript in such a way that it does not require profiling information. A security considerations section was written by Ned Freed, pointing out the nature of the risk associated with file operations, and recommending that they be disabled. Macintosh postscript files, which require laserprep header, as well as other postscript files generated by programs such as FrameMaker which call external libraries, must be sent with all such libraries prepended the mailed postscript to avoid the need to externally specify profiling information.

(b) .nroff and TeX

No person in the Working Group felt comfortable writing a complete profile for the use of either TeX or .nroff. The specification of these popular applications was left as a future effort.

6. Alphabet for Boundary Markers

The current alphabet for boundary markers makes it difficult to construct markers which are compatible with RFC934 and existing digesting software. The addition of space as a valid character would satisfy this need. Further discussion resulted in the adoption of a more general alphabet, to include the invariant set of characters defined for the use of Base-64 to be used in boundary markers. Trailing spaces are not permitted. When spaces are used in a marker, the entire marker will have to be quoted in the header.

7. Binary Type Definition

An unscheduled discussion on the need for the Binary type was held. With the clarification of the Applications type, and the difficulty of specifying exactly what initial content-types Binary should have, the Working Group, without objection, decided to drop it in favor of Application/Octet-Stream.

This was a natural progression from the realignment of content-types in terms of system resources begun before the Atlanta meeting. Application and Binary both require the ability to handle arbitrary Binary data, and require external programs to use the information.

8. Application/External-Reference

External Reference was seen by the Working Group to be a very useful feature, but inadequately defined in Quad-x. The current syntax provides no mechanism for multiple simultaneous retrieval mechanisms, the specification of syntax for mail-servers, or prioritizing the retrieval order. The use of specific Application/FTP and Application/NFS when used with Multi part/Alternative seems to be a reasonable approach, and was to be written up Borenstein.

As with the MIC, the absence of this feature was not seen to be a show-stopper. A new proposal will be submitted to the mailing list and if acceptable will be included in the document.

ACTION: Nathaniel Borenstine – Write up and submit to the mailing list a new proposal for application/external reference.

9. Use of Defaults

   The current Quad-x document specifies defaults for only selected content-types. In the case where defaults are not specified, and when the specified default may cease to be useful, possible ambiguity results. A strong view expressed before this meeting by Dave Crocker was supported by most attendees that defaults should be prohibited and that the subtype should always be specified. For broken mail which is sent with incomplete content-types, behavior of the reader is left up to the implementor and user. It was felt that because the message was already "broken" any uniform assumption could not be reliable.

10. Portable End-Of-Line Markers in Base 64

    The Working Group deleted end of line markers in Base 64, leaving it to the specific content-type to define the semantics of end of record. This decision has the advantage of restoring symmetry and transport independence between Base 64 and Quoted-Printable

11. Compression

    Compression was raised in the context of the Binary content type. Participants have expressed a desire, and the pragmatic realization that the use of "compressed, uuencoded, tar" files will continue to be sent and need to be indicated in the message. The Working Group previously stated it's preferences and rationale for not supporting uuencode, but has never clearly expressed it's position on compression. The issue was tabled pending a proposal to be sent to the mailing list. Again, if the proposal is acceptable it will be included, and it's absence will not be a show-stopper.

    ACTION: Neil Katin – Draft a proposal for the use of the compress algorithm in the Quad-X proposal.

    (a) Internal Reference in Richtext

        A proposal was made at this meeting to expand the richtext definition by including an internal-reference token. It was envisioned that this token would allow the insertion of objects in other parts of the message into the richtext stream. While many people supported this idea, no concrete proposal was submitted. If a proposal is approved by the mailing list, it will be included in the document.

        ACTION: Harri Salminen – Draft a proposal for Internal reference in the richtext content subtype.

With the conclusion of the meeting, five issues were left open. A new version of Quad-x, along with the proposals for the open issues, is due on December 6th. A new Internet Draft is expected at that time. The final comment period will end with the posting of a final version of Quad-x in the first week of January when the Working Group will submit the document to the IESG for Proposed Standard Status.

Header character set proposal

The Working Group began a review of the proposal submitted by Keith Moore to include character set identification and encoding information in the headers of a document.

The discussion was unstructured and resulted in a productive stream-of-consciousness review. The Working Group approved of the general approach and with the changes discussed, approved the proposal. Below are the main issues discussed and their resolution.

1. Multiple Encoded Words

   The Working Group felt that it should be acceptable to use multiple encoded words. Furthermore, the Working Group agreed that the length of encoded words should not be limited by this document, but rather by implementors of software in consideration of the pragmatic guidelines in the Quad-x document.

2. Character Set Names

   The Working Group committed to aligning the character set names between the header document, Quad-x and Simonsen's charset document. The use of the numeric identify was dropped, both as a result of allowing longer lines by specifying multiple encoded words, and out of consideration in making the encoded word more user-readable with old software.

Timetable for completion

This document will be aligned with Quad-x, and a new version will be submitted to the Internet-Drafts Directory by December 6th. At that time, the Working Group may decide to combine the two documents, or progress them jointly as a single standard. In any event, the Working Group committed to the submission of the header document and Quad-x as a bound set.

Attendees

Harald Alvestrand	`herald.alvestrand@delab.sintef.no`
Mary Artibee	`artibee@sgi.com`
Nathaniel Borenstein	`nsb@thumper.bellcore.com`
Ronald Broersma	`ron@nosc.mil`
Cyrus Chow	`cchow@ames.arc.nasa.gov`
James Conklin	`conklin@bitnic.educom.edu`
Robert Cooney	`cooney@wnyose.nctsw.navy.mil`
Mark Crispin	`mrc@cac.washington.edu`

Erik Fair	fair@apple.com
Ned Freed	ned@innosoft.com
James Galvin	galvin@tis.com
Jisoo Geiter	geiter@gateway.mitre.org
Russ Hobby	rdhobby@ucdavis.edu
William Jackson	jackson@manta.nosc.mil
Borka Jerman-Blazic	jerman-blazic@ijs.ac.mail.yu
William Jolitz	william@okeeffe.cs.berkeley.edu
Neil Katin	katin@eng.sun.com
Tom Kessler	kessler@sun.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@sgi.com
Gordon Lee	gordon@ftp.com
Jack Liu	liu@koala.enet.dec.com
Paul Milazzo	milazzo@bbn.com
Keith Moore	moore@cs.utk.edu
Mark Needleman	mhn@stubbs.ucop.edu
Daniel Newman	dan@innosoft.com
Michael Patton	map@lcs.mit.edu
Harri Salminen	hks@funet.fi
Keld Simonsen	keld@dkuug.dk
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.4 Network Database (netdata)

Charter

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

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Review and approve the Charter, making any changes necessary. Examine needs, resources for this network database protocol and define the scope of work. Begin work on a framework for the solution. Assign writing assignments for first draft of the document.
Done	First draft to be completed.
Aug 1991	Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting.
Dec 1991	Continue revisions based on comments received at meeting and e-mail. Start making document an Internet Draft.

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

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

Internet Drafts:

"Network Database Protocol", 06/26/1991, Daisy Shen <draft-ietf-netdata-netdata-01.txt>

"Network Database Implementation Information", 12/16/1991, Daisy Shen <draft-ietf-netdata-implement-00.txt>

CURRENT MEETING REPORT

Reported by Daisy Shen/IBM

Minutes of the Network Database Working Group (NETDATA)

This is the third meeting of the Working Group chaired by Daisy Shen. The meeting Agenda is shown below:

- Review the Charter.
- Review the Draft.
- Discuss the Draft and problems that are related to the subject.
- Report the status of the first implementation.
- Look for the second implementation.
- Discuss the effort of other vendors and OSF related to the subject.
- Future work.

Review the Charter

Most of the attendees were new to the Working Group; therefore, we reviewed the Charter and agreed that we followed the Charter and met the milestones on schedule by December 1991. Although the Group was small, a lot of valuable discussions were held.

Review the Draft

The original draft was separated into two documents. One was the protocol itself, and the other was implementation information. We reviewed both documents. The biggest mistake was a typo of ASN. Members suggested that the draft should define the requirements more clearly rather than explain one of the alternatives. Once the requirements are defined, it is up to the implementer which alternative to choose to implement the Network Database System. The draft will be revised according to the suggestion.

Discuss the draft and problems that are related to the subject

We discussed some problems and resolutions during the last meeting at Atlanta. We continued to discuss more issues, and resolved the following issues:

1. Multiple threads between clients and a server
   - Unit of Work
2. Multiple servers on a host
   - Program number and Port number
3. Data Buffers
   - Database <--> DB Utility <--> Server <--> Client
4. Data Conversion
   - Character strings with ASN.1 and BER(Basic Encoding Rules)

5. Security
   - Security is required, but the protocol should not be limited to Kerberos only. Kerberos is suggested, but not required. The implementer can choose a means that suits his/her own system.

Report the status of the first implementation

A version of Sun Microsystems <--> IBM VM is implemented. It follows the draft except for data conversion. It does not use ASN.1 but a somewhat similar method to manage data conversion.

Look for the second implementation

We are looking for volunteers to do the second implementation.

Discuss the effort of other vendors related to the subject

We would like to know more about the work that the SQL access group has done, but have not yet been able to.

Future Work

- Update the protocol draft and provide clear requirements.
- Discuss Error Recovery.
- Compare the performance of data conversion between the first implementation and using ASN.1.
- Give a presentation on ASN.1.
- Give a presentation on ISO standard.
- Run a demo.
- Contact the Operational Statistics Working Group.
- Find volunteers to do the second version of the implementation.

Attendees

L. Dain Gary	ldg@cert.sei.cmu.edu
William Jackson	jackson@manta.nosc.mil
Dale Johnson	dsj@merit.edu
Bill Melohn	melohn@auspex.com
Mark Needleman	mhn@stubbs.ucop.edu
Robert Purvy	bpurvy@us.oracle.com
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Daisy Shen	daisy@watson.ibm.com

3.1.5 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
Archive: /pub/netfax@stubbs.ucop.edu

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Review and approve Charter making any changes deemed necessary. Refine definition of scope of work to be accomplished and initial set of RFC's to be developed. Begin working on framework for solution.
Done	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 submit 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.

Internet Drafts:

"A File Format for the Exchange of Images in the Internet", 10/15/1991, Alan Katz, Danny Cohen <draft-ietf-netfax-netimage-02.txt>

CURRENT MEETING REPORT

Reported by Mark Needleman/U California

Minutes of the Network Fax Working Group (NETFAX)

The NETFAX Working Group met at the IETF meeting in Santa Fe. The main goal of the meeting was to go over the Internet Draft on a file format for transferring bitmapped images in the Internet.

- The Internet Draft was discussed and revised at the meeting. The proposed changes plus others that had been discussed on the list will be incorporated into the document and it will be put out for a short review again as an Internet Draft.

- Consensus was reached at the meeting that, provided there are no technical objections to the new version of the document, after the review period of a couple of weeks ends, the document will be progressed to proposed RFC status.

- A discussion was held as to what should be the future work of the Group, if anything, now that the document on file formats was nearing completion. It was brought up that at earlier meetings there had never been any consensus achieved on how to go forward on any of the other ideas the Group had discussed and that maybe the Group should not attempt to pursue anything further for awhile until everyone had a clear idea of what work was needed and how to do it. However it was decided that the Group should at least make one more attempt to define some of the issues and problems in things like addressing and routing that had been discussed at previous meetings. Dan Newman agreed to take some work he had previously posted to the list on this and expand it and repost it. It is hoped that this will become the basis for something that could be turned into an RFC discussing these issues and proposing solutions.

- A discussion was held on building interoperable implementations based on the file format now that the document has become standardized. Mark Needleman mentioned that the University of California under the auspices of the Coalition for Networked Information will move forward with plans to get organizations that have already done work in the area of networked fax to convert their projects to the standard file format and to get those projects to interoperate with each other. This will serve both to test out the proposed RFC and will also provide the requisite number of implementations that are required before a document can become a full RFC. Other participants in the meeting also indicated they would begin working on building implementations.

- Some discussion was held on the issue of testing and building conformance suites. It was agreed that some test files would be made available that could be used to test an implementation. Mark Needleman and Carl Malamud agreed to discuss between themselves where to locate these files. The idea being that they would either reside on stubbs.ucop.edu or on a host that Carl has access to.

Attendees

Jeffrey Buffum	buffum@vos.stratus.com
Ned Freed	ned@innosoft.com
Russ Hobby	rdhobby@ucdavis.edu
Ole Jacobsen	ole@csli.stanford.edu
Scott Kaplan	scott@ftp.com
Neil Katin	katin@eng.sun.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Peter Liebscher	plieb@sura.net
Carl Malamud	carl@malamud.com
Mark Needleman	mhn@stubbs.ucop.edu
Daniel Newman	dan@innosoft.com
Jon Postel	postel@isi.edu
Daisy Shen	daisy@watson.ibm.com
Andrew Veitch	aveitch@bbn.com
William Yundt	gd.why@forsythe.stanford.edu

3.1.6 Network News Transport Protocol (nntp)

Charter

Chair(s):
Eliot Lear, `lear@sgi.com`

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

Description of Working Group:

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

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

o User level and site designated authentication methods;

o Binary transfer capability;

o Minimization of line turnaround; and

o Stronger article selection capability.

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

Goals and Milestones:

Done	Define scope of work.
Jun 1991	Submit Internet Draft for review and comment.
Jun 1991	Possibly meet at USENIX for further comment.
Jul 1991	Meet at IETF for further comment.
Aug 1991	Submit RFC to IESG.

Internet Drafts:

"Network News Transfer Protocol Version 2: A Protocol for the Stream-Based Transmission of News", 09/30/1991, Eliot Lear <draft-ietf-nntp-news-00.txt, .ps>

3.1.7 Network Printing Protocol (npp)

Charter

Chair(s):
Glenn Trewitt, trewitt@pa.dec.com

Mailing Lists:
General Discussion: print-wg@pa.dec.com
To Subscribe: print-wg-request@pa.dec.com
Archive:

Description of Working Group:

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

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

Goals and Milestones:

Done	Review and approve the Charter, making any changes deemed necessary. Review the problems of printing in the Internet.
Done	Write draft LPR specification.
Done	Discuss and review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.
Done	Submit final LPR specification including changes suggested at the May IETF. Discuss document on mailing list.
Done	Submit LPR specification as an RFC and standard.
Jul 1990	Write description of the Palladium printing protocol (2.0) in RFC format.
Aug 1990	Discuss and review the draft Palladium RFC.

3.1.8 TELNET (telnet)

Charter

Chair(s):
Steve Alexander, stevea@i88.isc.com

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

Description of Working Group:

The TELNET Working Group will examine RFC 854, "Telnet Protocol Specification", in light of the last six 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

Internet Drafts:

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

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

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

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

Request For Comments:

RFC 1116 "Telnet Linemode option"

RFC 1184 "Telnet Linemode Option"

CURRENT MEETING REPORT

Reported by Steve Alexander/INTERACTIVE Systems Corporation

Minutes of the TELNET Working Group (TELNET)

An initial Agenda of possible topics included:

- Administration
- Authentication Option
- Environment Option
- Encryption Option
- Future of the Working Group

Steve Alexander replaced Dave Borman as Working Group Chair and introduced himself to the Group. Some members raised concerns about the functionality of the mailing list and Dave said he would look into it.

We then discussed the Authentication Option and what needed to be done to publish it as an Experimental RFC:

- Dave Borman will incorporate Jeff Schiller's security considerations.
- Kerberos IV and V documents will be split out.
- Ted T'so will verify the Kerberos IV draft.
- Dave Borman will talk to John Cole about Kerberos V.
- Four separate drafts will be issued for a two week comment period.
  - Basic option
  - Kerberos IV
  - Kerberos V
  - SPX

The next item was the Environment Option. This was fairly non-controversial.

- Dave Borman will split well-defined and arbitrary variables.
- There was discussion about passing some sort of OSTYPE variable based on the Assigned Numbers list.
- Dave will transfer editing of this option to Steve.

Both of these are expected to be completed by the end of 1991.

The next item was the encryption option and there was much spirited discussion around a security vs. performance balance.

- The current encryption option cannot stop an active attacker.
- Having encryption on might be too slow on some PCs.

It was decided to send the current authentication option forward to Experimental and then work on tying authentication and encryption together.

- Jeff Schiller will work with Steve Crocker to get expertise in this area.

Attendees

Steve Alexander	stevea@i88.isc.com
Robert Austein	sra@asylum.sf.ca.us
James Barnes	barnes@xylogics.com
Larry Blunk	ljb@merit.edu
David Borman	dab@cray.com
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Neil Haller	nmh@bellcore.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Mark Lewis	mlewis@telebit.com
John Linn	linn@zendia.enet.dec.com
Glenn McGregor	ghm@merit.edu
Jeffrey Schiller	jis@mit.edu
Richard Smith	smiddy@pluto.dss.com
Bruce Taber	taber@interlan.com
Theodore Tso	tytso@mit.edu
Kathleen Wilde	wilde@decvax.dec.com

3.2 Internet Area

Director(s):

- Philip Almquist: almquist@jessica.stanford.edu
- Noel Chiappa: jnc@ptt.lcs.mit.edu

Area Summary reported by Philip Almquist and Noel Chiappa

Four Internet Area working groups met in Santa Fe. The Internet Area also hosted two Birds-of-a-Feather (BOF) sessions.

IP over Appletalk

The Apple-IP Working Group revised their AURP (IP over Appletalk) and MacIP (Appletalk over IP) drafts. The Group expects that both of these documents are now finished, but will allow a final comment period before submitting them for standardization. SNMP over Appletalk is ready to be submitted for standardization. Work continues on Appletalk over PPP. Appletalk MIB enhancements are on hold pending further implementation experience.

Multi-Media Bridging

The Multi-Media Bridge Working Group has been working on a replacement for RFC1042 (IP over 802). This work is intended to better handle the peculiarities of 802.5 yet remain backwards-compatible with RFC1042. The Group also continues to consider the problems of bridging dissimilar networks.

Router Requirements

The Router Requirements Working Group revised and approved a Forwarding Table MIB document and made some minor revisions to the Router Requirements draft. The Group's Chair gave a plenary presentation on the Router Requirements draft in anticipation of its imminent completion. The Group held a joint session with the IDPR Working Group to ensure that the output of the two Groups will be consistent. For similar reasons, some members of the Group attended the BGP Working Group's discussions of route leaking between OSPF and BGP.

Point-to-Point Protocol Extensions

The Working Group decided, based on implementation experience, that some changes to the protocol were needed. In particular, they revised the definitions of the Link Quality Monitoring and IP Address Negotiation facilities. The Group also did some work on their PPP Authentication draft.

IP over Asynchronous Transfer Mode BOF

A BOF Chaired by Bob Hinden met to determine whether sufficient interest in ATM networks existed to justify the formation of an IP over ATM Working Group. The answer seemed to be a resounding yes.

Dynamic Creation of Network Links BOF

Another BOF, Chaired by Andy Nicholson, met to discuss experiments at Cray Research in "Dynamic Creation of Network Links" (basically, using switched T3 services to add Internet paths on demand). This BOF has met before, and will probably become a working group.

CURRENT MEETING REPORT

Reported by Andy Nicholson/CRAY

Minutes of the Dynamic Creation of Network Links BOF (DCNL)

BOFs were held on this subject at the 20th and 21st IETF's, under the name "Conditioning of By-Request Network Resources". This is a continuation of that interchange. The change of name was suggested by Noel Chiappa. This meeting attracted much more interest than the previous meetings. Attendees generally suggested developing a Charter for the possibility of starting an IETF Working Group.

While working with circuit-switched T3 networks, developers at Cray Research, Inc., determined that there would be advantages to defining a standard way to control certain classes of network resources through the Internet. In the case of a circuit-switched T3 line, the line should be switched on only when there are active transport connections which can fully utilize the service. Due to the high cost of the resource, underutilization would be particularly undesirable. The developers believe that this capability might have other applications in the Internet and that an effort should be made to define a standard protocol.

Minutes:

The meeting began with a presentation by Andy Nicholson regarding the work done at Cray Research with circuit-switched T3 networks. This was a review of the Internet Draft draft-nicholson-conditioning-00.txt.

This was followed by a short discussion of the Link Control Protocol used by the Cray Research demonstration software. This protocol is mentioned in the existing draft and will be fully described in an upcoming Internet Draft.

The attendees discussed different methods of supporting this service and how it might fit into the infrastructure of the Internet. One possibility is that rather than transport providers deciding when to activate and deactivate links, intermediate routers in the network may decide to perform this function when there is a need. In this way the network could automatically adapt to changing network load and delay conditions.

Some of the attendees were suggesting other possible uses of this capability and some discussion ensued. The most likely additionally use of this type of support for circuit-switched links would be for planned capacity management where an administrator might bring extra capacity on-line in the network during peak load times. Other possible uses are for backups to existing primary links or for bypass links when there is sufficient traffic between distant endpoints to avoid multiple hops between those endpoints.

Bill Jolitz suggested considering the management of the dynamic links. Rather than developing a new protocol (LCP) to create and manage links, SNMP could be used if a proper MIB was developed. He went on to suggest that this could be used as a methodology of defining the work to be done by making a first cut at a MIB for this facility. A working group could then iteratively improve the MIB while refining the functionality of dynamically created links.

Also discussed were various issues presented in a slide, and the attendees suggested other issues requiring consideration. One issue is that when bringing extra bandwidth into the network to alleviate congestion, the relief may only be temporary. As senders discover the extra bandwidth, they may increase their output to use it up.

Another possibility is that the routing protocols may dynamically change the network topology to suit the changing demands, and this would add new complexity to routing and routing protocols.

Ken Hayward was concerned whether this service would have a useful lifetime, considering that there does not seem to be, at present, an analog to switched T3 in the ATM world, and that ATM might address the issues presented here. There was general agreement that this was a good point, but that we could not predict the future. It was further noted that some networks might wish to have dynamic control of slower links, such as in the case of bringing backup links on-line when a primary link fails.

The attendees generally agreed that this is an interesting topic of discussion and expressed a desire to see a concrete description of the problems to be addressed. Andy Nicholson agreed to develop a Charter which addresses these concerns. He also agreed to install a mailing list for discussion of this topic. If there is sufficient interest then he will present the Charter to the IESG for working group status.

At the meeting Nicholson described a paper published by CICNet in July, 1991, which mentioned their interest in circuit-switched T3 services. He promised to include instructions on how to get this paper. That information is presented here.

A report titled "High Performance Applications on CICNet: Impact on Design and Capacity" is available from CICNet via anonymous FTP.

ftp: NIC.CIC.NET

cd: /pub/reports

get: ds3-report.[ps or txt]

ABSTRACT: This twenty-three page report summarizes available network technologies, reports on a survey of the needs of researchers and faculty at CIC institutions, and provides detailed studies of network requirements in four areas of contemporary, scientific research.

The needs of these four areas of research are then summarized in terms of network requirements, and specific recommendations are presented by the Working Group to CICNet, Inc. The report was authored by the CICNet DS-3 Working Group, which was chaired by Mike Enyeart of Indiana University.

Final Note: A mailing list for this work has been set up, its address is dcnl-ietf@cray.com.

Attendees

Jordan Becker	becker@nis.ans.net
Richard Cherry	rcherry@wc.novell.com
Erik Fair	fair@apple.com
Craig Fox	foxcj@network.com
John Gong	jgong@us.oracle.com
Kenneth Hayward	crm57d@bnr.ca
Greg Hollingsworth	gregh@mailer.jhuapl.edu
William Jolitz	william@okeeffe.cs.berkeley.edu
Andy Nicholson	droid@cray.com
Timothy Salo	tjs@msc.edu
Stephen Shew	sdshew@bnr.ca
Martha Steenstrup	msteenst@bbn.com
Brad Steinka	brad@python.eng.microcom.com
Preston Wilson	preston@i88.isc.com
Robert Woodburn	woody@cseic.saic.com
Richard Woundy	rwoundy@ibm.com

CURRENT MEETING REPORT

Reported by Peter Honeyman/UMich

Minutes of the IP Routing for Wireless/Mobile Hosts BOF (MOBLHOST)

The IP Routing for Wireless/Mobile Hosts BOF met on Wednesday evening during the Santa Fe IETF. Notes for that meeting were taken by Peter Honeyman and were edited by Steve Deering.

Agenda

- Report on IEEE 802.11
- IP over Ham Radio
- Duchamp Report
- Karlberg Presentation
- Working Group Status

Report on IEEE 802.11. (Deering)

The IEEE 802.11 committee is working on a wireless LAN standard compatible with other protocols in 802.x suite. It is not clear what compatible means beyond 48 bit addresses. The standard is intended to support packet voice as well. The scope of the work includes:

- Physical layer: identifying possible physical layers.
  - Radio frequency (single-multi-channel, spread-spectrum, frequency hopping, ...)
  - Infra-red
- MAC layer: Narrowed down to four candidates that differ with respect to:
  - Random access vs. TDMA
  - Single- vs. multi-channel
  - Need for wired infrastructure or distinguished node (base station).
- Architecture: How (whether) to tie cells together.
  - Basic service area vs. extended service area, likely to use 802-style bridging, contributions welcome.

The following were comments in response to questions:

- Concerned with both stationary wireless and mobile/ephemeral wireless.
- No apparent bias towards radio frequency over infra-red.
- Architecture tending towards base-station/mobile model, rather than more general architecture.

- Data rates of interest are roughly 1 to 20 Mbps.
- Max packet size not yet pinned down.
- The IEEE 802.11 Committee has met for just over one year; with their first meeting taking place in September 1990. Standard voting rights rules (attend some number of consecutive meetings.)

IP over Ham Radio

Question: Interested in MAC-layer issues, but not tackling mobile?
Answer: Not at all! Not too uncommon to have 10 meter packet radio attached to laptop, relay back to home station. Most is AX.25. Some is TCP/IP. Mobile includes moving among variety of nets (ISDN, wireless LAN, Ethernet)...

Deering's classification of host mobility:

1. "Permanent" relocation (weeks to years)
   - Manual configuration, BOOTP, ...
2. "Temporary" relocation (hours to weeks)
   - Originator only - Dynamic Host Configuration Protocol (DHCP)
   - Originator + target
   - DHCP + DNS updates
   - Or keep old address
3. Roaming
   - Keep same address, to keep connections alive.

Report from Dan Duchamp, Columbia

See SIGCOMM paper or notes from last mobile BOF. The basic problem is to keep sessions alive in a mobile environment.

Goal	Hop networks "seamlessly" No change to hosts not involved with mobile computers (as distinguished from the Sony Labs effort)
Setting	IP suite, level three (*not* the link layer)
Gist	Put mobile hosts on "virtual network" Maintain a distributed database (DDB) mapping hosts on virtual network to a location on a physical network. (The DDB serves the functions of router configuration and administration.)
Terms	To talk to mobile hosts (MHs), network has "mobile support stations" (MSS). Traffic to/from MHs is gatewayed through MSS via IP-IP encapsulation protocol. A virtual subnet is a subnet on which MHs live.

Intra-campus routing	MH must have MSS as its gateway (implies every cell must contain an MSS). First hop of stationary host route to a virtual subnet is via its MSS.

```
Host  MH
 |  |
     IP |   | IP
 |     IPIP   |
MSS ----------- MSS
```

(Imagine a picture here with a cloud in the middle, a host attached to the cloud, a couple cells hanging off the cloud, and an MH communicating with an MSS:

```
(^^^^^^^^^)----- Host
(   )
(          )
vwwwwwwwwwv
  |  |
  |  |
    ------+--    --+------
     |       |
     |       |
      ( (( MSS )) )     ( (( MSS )) ) ) ) ) MH
```

Q: Why not host routes?
A: Fast topology changes, highly dynamic, much state.
IPIP description: Packets sent through two MSSs using IPIP tunneling. e.g., send from stationary host to MH:

```
+---------------------------------+
| Link: Dest = local MSS          |
|    +-------------------------+  |
|    | IPIP: Dest = MH's MSS   |  |
|    |      +----------------+ |  |
|    |      | IP: Dest = MH  | |  |
|    |      +----------------+ |  |
|    +-------------------------+  |
+---------------------------------+
```

Q: MH has one address?
A: Yes! Fixed, forever, in its virtual subnet.

Q: What about broadcast, multi-cast?
A: Evil but necessary. We punt.

Flies in the ointment:

1. Heuristic for beacon accept/reject.
2. No IP broadcast.
3. Security? What's that?
4. Need careful placement of MSSs – enough of them, powerful enough, ...
5. Is inter-campus case adequate?
6. Need (worldwide) common beacon conventions.
7. Plenty of knobs to turn. e.g., beacon aging.
8. No relationship with: IDPR, DHC, 1241 encapsulation, Chiappa, XXX discovery, for all XXX.

An RFC is in the works.

Presentation by Ken Carlberg, MCIC

1. Primary Goals
   - Focus work on additions to existing/proposed ISO standards.
   - Minimize responsibility/workload of end systems.
     - End system transparency concerning mobile end system (MES) movement.
2. Two-tiered design
   - Augment intra IS-IS PDUs.
   - Inclusion of Directory Service for IDRP.

Use X.500 to discover proper domain, IDRP thereafter. Re-register at the domain level.

MES has permanent logical address. When enter area, router dynamically assigns router address. When move, routing address changes.

Working Group Status

A vote was taken to determine whether the BOF should become a Working Group. The Group approved and there was discussion of the implications and expectations.

- Proposed Charter:
  - Develop/adopt architecture and protocols to support mobile hosts in the Internet.

  - Convey Internet mobility concerns and ideas to relevant working groups and standards bodies.

- Scope of work:
  - Issues above media access: addressing, naming, routing, bridging.
  - Issues beyond DHC: roamers, temporarily relocated hosts.
  - Mobile hosts, networks, collections of networks.
  - Mobility across multiple link layers (wired and wireless).
  - Multi-protocol as well as IP-only.
  - Impact on higher layers, e.g., transport layer.
  - Accommodation of sleeping hosts and off-line hosts.
  - Cellular topology and general topology, with and without wired infrastructure.
- Outside of Scope:
  - Solutions that do not interoperate with existing IP hosts and routers.
  - Issues of delay/jitter-sensitive traffic, TOS queuing/routing.
  - Congestion avoidance and control.
  - Compression.
  - Privacy-not at this layer. (This does not exclude authentication.)
  - No MIB. (Applause.) Although it may be a late addition.

Deering volunteered to chair the Group. No other volunteers.

Attendees

Mary Artibee	artibee@sgi.com
Karl Auerbach	karl@eng.sun.com
William Barns	barns@gateway.mitre.org
James Beers	beers@nr-tech.cit.cornell.edu
Ken Carlberg	carlberg@cseic.saic.com
Steve Deering	deering@xerox.com
John Demco	demco@cs.ubc.ca
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Hans Eriksson	hans@sics.se
Karen Frisa	karen.frisa@andrew.cmu.edu
Jisoo Geiter	geiter@gateway.mitre.org
Jim Ghadbane	jimgh@newbridge.com
Kenneth Hayward	crm57d@bnr.ca
Peter Honeyman	honey@citi.umich.edu
William Jolitz	william@okeeffe.cs.berkeley.edu
Neil Katin	katin@eng.sun.com
Tom Kessler	kessler@sun.com
Vincent Lau	vincent.lau@eng.sun.com
Gordon Lee	gordon@ftp.com
Mark Lewis	mlewis@telebit.com
Bill Melohn	melohn@auspex.com
Paul Milazzo	milazzo@bbn.com

Greg Minshall	minshall@wc.novell.com
Zbigniew Opalka	zopalka@roswell.spartacus.com
J. Bradford Parker	brad@cayman.com
Allan Rubens	acr@merit.edu
Eric Smith	eric@telebit.com
Martha Steenstrup	msteenst@bbn.com
Osamu Takada	takada@sdl.hitachi.co.jp
Andrew Veitch	aveitch@bbn.com
John Veizades	veizades@apple.com
Walter Wimer	walter.wimer@andrew.cmu.edu

3.2.1 Connection IP (cip)

Charter

Chair(s):

Claudio Topolcic, `topolcic@nri.reston.va.us`

Mailing Lists:

General Discussion: `cip@bbn.com`
To Subscribe: `cip-request@bbn.com`
Archive:

Description of Working Group:

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

Goals and Milestones:

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

Request For Comments:

RFC 1190	"Experimental Internet Stream Protocol, Version 2 (ST-II)"

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chair(s):
Ralph Droms, droms@bucknell.edu

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

Description of Working Group:

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

Goals and Milestones:

Done — We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to: Exchange Internet packets with other hosts, Obtain packet routing information, Access the Domain Name System, and Access other local and remote services.

Done — We will summarize those mechanisms already in place for managing the information identified by Objective 1.

Jan 1991 — We will suggest new mechanisms to manage the information identified by Objective 1.

Jan 1991 — Having established what information and mechanisms are required for host operation, we will examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.

TBD — Write a bootp extensions document

Internet Drafts:

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

"Dynamic Host Configuration Protocol", 07/09/1991, R. Droms <draft-ietf-dhc-protocol-01.txt, .ps>

3.2.3 IP over AppleTalk (appleip)

Charter

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

Mailing Lists:
General Discussion: apple-ip@apple.com
To Subscribe: apple-ip-request@apple.com
Archive:

Description of Working Group:

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

Goals and Milestones:

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

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

Internet Drafts:

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

"Tunnelling AppleTalk through IP", 10/30/1991, Alan Oppenheimer <draft-ietf-appleip-aurp-02.txt, .ps>

"SNMP over AppleTalk", 12/23/1991, G. Minshall, M. Ritter <draft-ietf-appleip-snmp-appletalk-00.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

INTERIM MEETING REPORT

Reported by John Veizades/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)
October, INTEROP '91

The meeting started with the following Agenda:

- MacIP
- AURP
- AppleTalk and SNMP
- AppleTalk and PPP

The consensus of the Group was that the MacIP document should be supplemented with some of the comments made by Tom Evans in an appendix to the document as a set of implementors notes. The Group also felt that the specification is riddled with problems and that this specification should be worked on to provide a long-term solution. The author was willing to do this if he is not chartered to work on the long-term direction of this specification.

The MacIP document will be moved to the draft stage as soon as these changes are made.

The AURP discussion then began Alan Oppenheimer of Apple led this discussion. The discussion started with a description of the protocol changes between this document and the previous version of the document. The changes are as follows:

- Addition of Open-Req and Open-Rsp packets
- Changes to zone-based packets
  - Move into connection stream
  - Combined into one packet type
- Packet header changes
- Changes to details of connection teardown
  - Keep the one-way connections independent
  - Only reset tickle timer on AURP packets
- Calling out of transport as a separate layer "AURP-Tr"

On the last subject there was heated debate on whether this was the correct way to design the protocol and it was stated that this may tie the protocol to IP and not make it portable to other protocols. The following points were thought to be essential to allowing the protocol to be accepted as an internet standard. The type of service required by the underlying media should be called out specifically. The transport should be separated from the rest of the protocol so that other transports could be used. A state diagram should be included in the document.

Some discussion on the ability of AURP to solve the problems of the mythical customer were held. Some real customers (Boeing and NASA) expressed their doubts as to whether there was a need for this functionality in their networks.

The state diagram of the AURP protocol will be specified and then the document will be put in the Internet-Drafts Directory.

The next topic was AppleTalk and PPP. A preliminary document specifying the use of SNMP over DDP was presented by Mike Ritter from Apple. This document is available on apple.com in the directory /pub/apple-ip. The AppleTalk MIB II document is also available there. These documents are also available on AppleLink in the Developer Support folder.

The final topic that was discussed was the start of work by this Group on the AppleTalk over PPP specification. Brad Parker has been working with others on a preliminary specification for this functionality and this will be available from him shortly. This specification supports much of the functionality that is available with the ARAP (Apple Remote Access Protocol) specification that is available through APDA. Brad is also working with the PPP Extensions Working Group to arrive at a specification for dial-back and security on PPP links.

The next meeting will be held in conjunction with the IETF meeting in Sante Fe, New Mexico, November 18-22, 1991.

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)

The Working Group met and discussed work in the following areas:

SNMP

Work in the SNMP area is split into three areas. The AppleTalk MIB Plus (the first version is now RFC 1243) (this MIB will no longer be called MIB 2) is now out for comment as an Internet Draft. There are implementations of the RFC 1243 MIB available on Shiva, Cayman, Farallon, 3Com and ACC Systems. Implementation and use experience has led to the following list of problems with this MIB: it is felt that there may be more variables than is needed, this MIB does not allow for the configuration of routers and there are questions on whether this MIB supports half routers well. It is felt that there are significant areas for discussion and implementation. The Group is not trying to rush the MIB Plus document and is waiting for appropriate comment. The SNMP over AppleTalk document is ready to move forth as a Proposed Standard and will be doing so shortly after comments from this meeting are incorporated in the document. Concern was raised about getting major console manufacture is to incorporate this standard into their consoles. Concern was also raised as to the ability of the MIB to be used for the global changing of a network's zone list. Test tools are available from Mike Ritter (MWRitter@Applelink.apple.com). The last item was the Macintosh system MIB which is now out for general comment.

AURP

The AURP (Apple Update Based Routing Protocol) will be progressing from Internet Draft to Proposed Standard after revising the state diagram. The completed document will be submitted as a Proposed Standard in the Internet community and will be made available as an APDA document. A vendor product "bakeoff" is scheduled for MacWorld in January, with about seven companies at various stages of implementation (Cayman, cisco, Shiva, DEC, Farallon, Compatible Systems, Novell, 3Com, Pacer and Apple). Seeding of some of these products to sites around the world is also planned in the next few months.

ABGP

A presentation was made on the possibility of introducing a BGP like protocol as a border gateway protocol for AppleTalk. Greg Bruell from Shiva made the presentation, Yakov Rekhter (IBM) and Scott Brim (Cornell) were in attendance. Why BGP? It looks a lot like AURP when you make some needed extensions to BGP to incorporate AppleTalk. Transport stays the same as BGP except that it uses a different TCP Port. Message layer stays the same as BGP, the autonomous system number maps to the domain identifier and there is a change to the network list into the network zone tuple list. Some advantages are that BGP is a well known implementation, decisions on policy are outside of the update protocol and it is easy to implement.

AppleTalk and OSPF

Greg Bruell led a discussion on using an OSPF like protocol to replace the AppleTalk IGP which is RIP-like.

PPP and AppleTalk

The document presented is close. Comments will be incorporated and reissued as an Internet Draft for comments from the AppleTalk community as well as the PPP community. Additions to the current document include calling out and describing the operation of several common cases; node to server, node to node and half routing. Comments on hop count incrementing and which options should be negotiated for each case will be added. Operation with AURP will be left to the AURP effort. The smartbuffering compression algorithms are available through APDA in the document which describes the operation of the AppleTalk Remote Access Protocol (ARAP). Implementations are in progress by cisco, Cayman, Shiva, Novell, Telebit, A/UX and Farallon. The AppleTalk over PPP work was presented to the PPP Extensions Working Group. The PPP Extensions Working Group added functionality that will allow all that is needed for call-back in the security fields of the LCP. Both Brad Parker and John Veizades presented the Apple communities' view on dial-back and security. The version of the PPP document that will contain the PPP Extensions for security will include everything needed for dial-back as presented in the ARAP specification as well as the ability for the user to specify the number string to be called back at. The security specification will also contain whatever is necessary for "secure ID" extensions.

MacIP

Three outstanding comments were brought up and will be incorporated into the current document which will be posted for final review before moving the protocol to Proposed Standard. The areas of comments were ICMP messages, out-of-zone-operation, and multiple servers in the same zone. In the area of ICMP messages it was decided that ICMP redirects will be gleaned by the macIP gateway when it is doing proxy arp for nodes in the AppleTalk network that are on the same logical subnet as the gateway. In the area of out-of-zone-operation, if two hosts use the same address in the AppleTalk internet packets it is destined that one will be reliably dropped. When two servers are in the same zone some election mechanism will be used to choose one of them as the gateway though others will be kept to use as secondaries if the first fails to provide registration or services. Two features should be added for the rebuilding of the AppleTalk address to IP address mapping on server restart, one is the Phil Koch algorithms for gleaning address mappings and the other is the ability to send NBP lookups to specified zones to rebuild the mapping table.

OAF - Open AppleTalk Federation

Discussion was held on how to continue the growth of the infrastructure related AppleTalk protocols. Most ideas evolve them by moving them into the IETF community. Work is being done on charter definition, vendors buy in and on discussing these issues with the relevant Apple people. This effort would proceed within the infrastructure of the IETF. The IETF has been approached as to the viability of this undertaking and they advise that the work could be accomplished under an AppleTalk directorate within the IETF. Concern was raised as to Apple's role in such a venture and what Apple's commitment to such a venture would be.

Attendees

James Beers	beers@nr-tech.cit.cornell.edu
Craig Brenner	Brenner2@applelink.apple.com
Gregory Bruell	gob@shiva.com
Philip Budne	phil@shiva.com
Peter Caswell	pfc%pacvax@uunet.uu.net
Richard Cherry	rcherry@wc.novell.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Dino Farinacci	dino@cisco.com
Karen Frisa	karen.frisa@andrew.cmu.edu
John Gawf	gawf@compatible.com
Bob Jeckell	robert_jeckell@nso.3com.com
Holly Knight	holly@apple.com
Louise Laier	laierl@applelink.apple.com
Joshua Littlefield	josh@cayman.com
Greg Merrell	merrell@greg.enet.dec.com
Greg Minshall	minshall@wc.novell.com
Robert Morgan	morgan@jessica.stanford.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Chandy Nilakantan	csn@3com.com
Alan Oppenheimer	oppenheimer1@applelink.apple.com
J. Bradford Parker	brad@cayman.com
Christopher Ranch	cranch@novell.com
Michael Ritter	mwritter@applelink.apple.com
Eric Smith	eric@telebit.com
Evan Solley	solley@applelink.apple.com
David S.A. Stine	dstine@cisco.com
John Veizades	veizades@apple.com
Lee Wade	wade@nsipo.arc.nasa.gov

3.2.4 IP over Asynchronous Transfer Mode (atm)

Charter

Chair(s):
Robert Hinden, hinden@bbn.com

Mailing Lists:
General Discussion: atm@bbn.com
To Subscribe: atm-request@bbn.com
Archive: Send message to atm-request@bbn.com

Description of Working Group:

The IP over ATM Working Group will focus on the issues involved in running internetworking protocols over Asynchronous Transfer Mode (ATM) networks. The final goal for the Working Group is to produce standards for the TCP/IP protocol suite and recommendations which could be used by other internetworking protocol standards (e.g., ISO CLNP and IEEE 802.2 Bridging).

The Working Group will initially develop experimental protocols for encapsulation, multicasting, addressing, address resolution, call set up, and network management to allow the operation of internetwork protocols over an ATM network. The Working Group may later submit these protocols for standardization.

The Working Group will not develop physical layer standards for ATM. These are well covered in other standard groups and do not need to be addressed in this Group.

The Working Group will develop models of ATM internetworking architectures. This will be used to guide the development of specific IP over ATM protocols.

The Working Group will also develop and maintain a list of technical unknowns that relate to internetworking over ATM. These will be used to direct future work of the Working Group or be submitted to other standard or research groups as appropriate.

The Working Group will coordinate its work with other relevant standards bodies (e.g., ANSI T1S1.5) to insure that it does not duplicate their work and that its work meshes well with other activities in this area. The Working Group will select among ATM protocol options (e.g., selection of an adaptation layer protocol) and make recommendations to the ATM standards bodies regarding the requirements for internetworking over ATM where the current ATM standards do not meet the needs of internetworking.

Goals and Milestones:

Done First Meeting. Establish detailed goals and milestones for Working

Group.

Jan 1992 Circulate drafts of IP over ATM Specifications.

Mar 1992 Review approaches to running IP over ATM.

none specified

CURRENT MEETING REPORT

Reported by Bob Hinden/BBN

Minutes of the IP over Asynchronous Transfer Mode BOF (ATM)

The meeting was organized as a BOF to determine if there was enough interest in forming an IETF working group to develop protocols to run IP over ATM networks.

The first half of the meeting was a presentation by Bob Hinden/BBN covering the reasons why the BOF was organized. The talk included:

- Motivation for the IP over ATM Group
- Why ATM is interesting
- Relationship to other standards groups
- Area for work
- Next steps

The second half of the meeting was a discussion of whether the Group should be formed and if so what area the Group should first focus on. The work areas discussed included:

- ATM Internet Architecture
- ATM and Internetwork Operation
- Adaptation Layer Selection(s)
- Flow Control/Congestion Avoidance
- Flow Setup/Connection Establishment
- IP Encapsulation
- Multicast
- Network Management
- Physical Media
- Security

Other issues discussed include whether the Group should initially develop standard protocols or experimental protocols.

After much discussion the Group decided that there was enough interest that an IETF Working Group should be formed. There was also a consensus that the Group initially develop experimental protocols.

There was agreement on the following topics:

- Approaches for Interoperability with the Internet (e.g., IP over ATM)

- Develop a list of unknowns that need to be worked on.

- Set up a separate mailing list from the IPLPDN Working Group.

- Identification of Architectural Alternatives to Internetworking over ATM.
- That the work of this Group not duplicate the work that is being done in other ATM standard groups.
- Physical Layer Standards are very well covered in other standard groups and do not need to be addressed in this Group.
- There should be communication with other ATM Standard Bodies.
- Initially to not work on Congestion Control issues.
- The initial focus of the Group should be local with wide area internetworking.

The Chair was tasked to revise the draft Charter to be consistent with these agreements.

Attendees

James Barnes	barnes@xylogics.com
Art Berggreen	art@acc.com
David Bolen	db31@nis.ans.net
Robert Braden	braden@isi.edu
Theodore Brunner	tob@thumper.bellcore.com
John Burnett	jlb@adaptive.com
Lida Carrier	lida@apple.com
Yee-Hsiang Chang	yhc@concert.net
A. Lyman Chapin	lyman@bbn.com
George Clapp	clapp@ameris.center.il.ameritech.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Danny Cohen	cohen@isi.edu
Dave Cullerot	cullerot@ctron.com
Steve Deering	deering@xerox.com
Kurt Dobbins	dobbins@ctron.com
Hans Eriksson	hans@sics.se
Erik Fair	fair@apple.com
Jim Ghadbane	jimgh@newbridge.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Herve Goguely	lircorp1!roger@uunet.uu.net
Kenneth Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Christine Hemrick	hemrick@cisco.com
Robert Hinden	hinden@bbn.com
William Jolitz	william@okeeffe.cs.berkeley.edu
Satish Joshi	sjoshi@synoptics.com

Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Manu Kaycee	kaycee@ctron.com
Yoav Kluger	ykluger@fibhaifa.com
Cheryl Krupczak	cheryl@cc.gatech.edu
Anthony Lauck	lauck@tl.enet.dec.com
Joseph Lawrence	jcl@sabre.bellcore.com
Tom Lyon	pugs@eng.sun.com
Andrew Malis	malis@bbn.com
Allison Mankin	mankin@gateway.mitre.org
David McDysan	0002806303@mcimail.com
Bill Melohn	melohn@auspex.com
Donald Merritt	don@brl.mil
Greg Minshall	minshall@wc.novell.com
Dave Monachello	dave@pluto.dss.com
Andy Nicholson	droid@cray.com
Zbigniew Opalka	zopalka@roswell.spartacus.com
Venkat Prasad	vsp@3com.com
Yakov Rekhter	yakov@watson.ibm.com
Manoel Rodrigues	manoel.rodrigues@att.com
Philippe Roger	lircorp1!roger@uunet.uu.net
Allan Rubens	acr@merit.edu
Timothy Salo	tjs@msc.edu
Sharad Sanghi	sharad@ans.net
Tim Seaver	tas@mcnc.org
John Seligson	johns@ultra.com
Stephen Shew	sdshew@bnr.ca
Mike Spengler	mks@msc.edu
Ravi Srinivasan	ravi@eng.vitalink.com
Martha Steenstrup	msteenst@bbn.com
Brad Steinka	brad@python.eng.microcom.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Sven Tafvelin	tafvelin@ce.chalmers.se
Osamu Takada	takada@sdl.hitachi.co.jp
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Lee Wade	wade@nsipo.arc.nasa.gov
Walter Wimer	walter.wimer@andrew.cmu.edu
Richard Woundy	rwoundy@ibm.com

IP over ATM

BOF

Robert Hinden

November 20, 1991

INTRODUCTION

- Motivation for IP/ATM Group
- Why ATM is Interesting
- Relationship to Other Standards Groups
- Areas for Work
- Next Steps

MOTIVATION FOR IP/ATM W.G.

- Growing Interest in ATM Cell Relay
  - Carriers
  - Network Equipment Vendors
  - Computer Manufactures
- IETF is Ideal Body to Provide Feedback to Other Standard Organizations
  - IETF best represents Internetworking Designers and Users
  - Unparalleled Understanding of Data Networking
- Now is the Ideal Time to Develop Common Protocols to allow interoperation between:
  - New ATM Infrastructure
  - Existing Internetworking Infrastructure

WHY IS THIS NOT IN IPLPDN W.G.

- Multimedia is a Key Driver for ATM
  - Requires Guarantees of Service
- ATM is Not Limited to Large Public Data Networks
  - ATM will be prevalent in LAN's well before it appears in the Public Networks
  - ATM will also be used in Private ATM Networks
- Important IP/ATM Issues are Not Routing and/or Framing
  - Route Set Up
  - Type of Service Requests and Management

WHY IS ATM INTERESTING

- Simple Fast Switching
  - VCI Based Addresses
  - Amenable to VLSI Implementations
- Small Fixed Size Packets (a.k.a. Cells)
  - Low Switching Latency
  - Bandwidth Control
- Scalable
  - Range of Speeds
  - Media Independent
- Standards
  - Public Carriers are Committed to ATM
  - International Infrastructure Likely to be Developed

SWITCHED v.s. SHARED MEDIA LANS

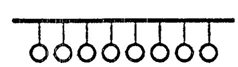

- Port Bandwidth = 1 / N
- Port Latency = N
- Port Cost = Aggregate Bandwidth
- Aggregate Bandwidth = Constant
- Technology Hard to Scale
- 5-10 Year Cycles

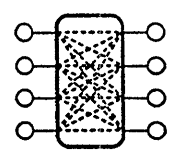

- Port Bandwidth = Constant
- Port Latency = Constant
- Port Cost = Constant
- Aggregate Bandwidth = N
- Technology Easy to Scale
- 18 - 24 Month Cycle

COMPANIES ANNOUNCING ATM PRODUCTS

- Adaptive
- BBN
- David System
- CISCO
- FORE Systems
- NET
- Newbridge
- Northern Telecom
- Stratacom
- Ungerman Bass
- Hughes LAN Systems
- Synoptics
- ATT
- Alcatel
- Thompson CSF
- Siemens
- OKI
- Timeplex

IS ATM TOO GOOD TO BE TRUE?

- Surprisingly Large Amount of Agreement over the Use of ATM
  - Carriers
  - Computer Manufacturers
  - Switch Vendors
- Could Everyone Be Right?
- More Likely that if Everyone Agrees
  - They must be Wrong!
- Good News !
  - Agreement is only at Top Level
  - Lots of Disagreement on Next Level

    AAL's, Flow Control, Setup, Framing, ...

FUTURE ATM INTERNET

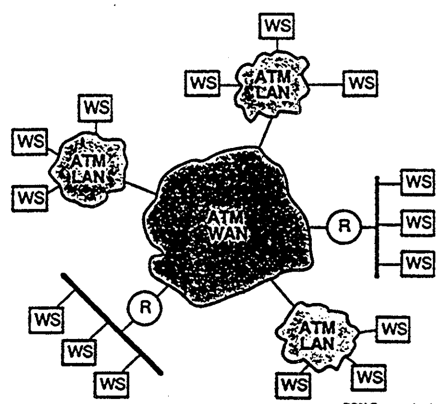

ATM INTERNET SERVERS

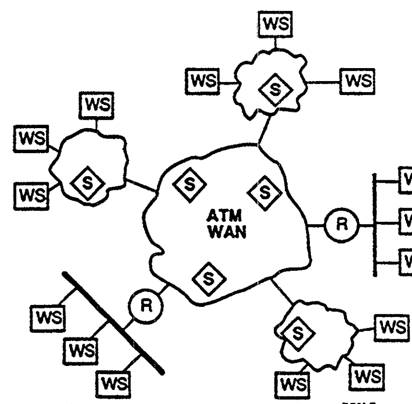

RELATIONSHIP TO OTHER STANDARDS GROUPS

- ANSI T1S1.5
  - Broadband ATM Standards
- IEEE 802.6
  - DQDB / SMDS ATM Standards
- ATM Forum
  - Vendor Forum Developing ATM Implementation Agreements

AREAS FOR WORK

- ATM Internet Architecture
- Flow Setup / Connection Establishment
  - Type of Service Selection
- Flow Control / Congestion Avoidance
- ATM and Internetwork Interoperation
- IP Encapsulation
- Adaptation Layer Selection(s)
  - AAL1, AAL2, ... AALn
- Physical Media
- Network Management
-
-
-

IP / ATM — November 20, 1991

ISSUES

- Should IP/ATM become a IETF Working Group?
- Initially Develop Standard Protocols or Experimental Protocols?
- How to Interface with Other Standard Groups?
- Which Technical Areas to Address First?

BBN Communications

IP / ATM — November 20, 1991

NEXT STEPS

- Formalize Working Group
- Revise Charter
- Develop Work Plan
  - Work Items
  - Schedule
  - Assignments
- Setup Formal Liaison with Other Standards Groups

BBN Communications

3.2.5 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
Archive:

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 liaison with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards. This Group is specifically not chartered to provide solutions to mixed media bridging problems.

Goals and Milestones:

Done — Write a document specifying the use of IP on a single MAC FDDI station.

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

Request For Comments:

RFC 1188 — "A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks"

3.2.6 Multi-Media Bridging (mmb)

Charter

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

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

Description of Working Group:

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

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Minutes of the Multi-Media Bridging Working Group (MMB)

Report not submitted. Refer to Area Report for a brief summary.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Karl Auerbach	karl@eng.sun.com
Art Berggreen	art@acc.com
David Bridgham	dab@asylum.sf.ca.us
John Burnett	jlb@adaptive.com
J. Noel Chiappa	jnc@lcs.mit.edu
Rob Coltun	rcoltun@ni.umd.edu
Kenneth Crepea	crepea@cisco.com
Dave Cullerot	cullerot@ctron.com
Jeff Fitzgerald	jjf@fibercom.com
Craig Fox	foxcj@network.com
Darrell Furlong	
Jim Ghadbane	jimgh@newbridge.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Ronald Jacoby	rj@sgi.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Yoav Kluger	ykluger@fibhaifa.com
Stev Knowles	stev@ftp.com
John Lekashman	lekash@nas.nasa.gov
Keith McCloghrie	kzm@hls.com
Bill Melohn	melohn@auspex.com
David Minnich	dwm@fibercom.com
Dean Morris	morris@marvin.dec.com
Venkat Prasad	vsp@3com.com
Anil Rijsinghani	anil@levers.enet.dec.com
Philippe Roger	lircorp1!roger@uunet.uu.net
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Ravi Srinivasan	ravi@eng.vitalink.com
Iris Tal	437-3580@mcimail.com
Yuan Wang	natadm!ycw@uunet.uu.net
John Ziegler	ziegler@artel.com

3.2.7 Point-to-Point Protocol Extensions (pppext)

<u>Charter</u>

Chair(s):
Brian Lloyd, `brian@ray.lloyd.com`

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

Description of Working Group:

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

Goals and Milestones:

None specified

Internet Drafts:

"Definitions of Managed Objects for the Point-to-Point Protocol", 09/10/1990, Frank Kastenholz <draft-ietf-pppext-pppmib-01.txt>

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

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

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

"The Point-to-Point Protocol for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links", 07/01/1991, W A Simpson <draft-ietf-pppext-lcp-02.txt>

"The PPP Internet Protocol Control Protocol (IPCP)", 07/01/1991, G McGregor <draft-ietf-pppext-ipcp-03.txt>

"Proposed Point-to-Point Procotol for AppleTalk", 07/08/1991, S. Senum, J. Muchow, F. Slaughter, B. Parker <draft-ietf-pppext-appletalk-00.txt>

"The PPP OSI Network Layer Control Protocol (OSINLCP)", 07/25/1991, D. Katz <draft-ietf-pppext-osinlcp-00.txt>

"The PPP Authentication Protocols", 07/25/1991, B. Lloyd, W.A. Simpson

<draft-ietf-pppext-authentication-02.txt>

"PPP Link Quality Monitoring", 12/30/1991, W. A. Simpson <draft-ietf-pppext-lqm-01.txt>

Request For Comments:

RFC 1220 "Point-to-Point Protocol Extensions for Bridging"

CURRENT MEETING REPORT

Reported Brian Lloyd/Telebit

Minutes of the Point-to-Point Protocol Extensions Working Group (PPPEXT)

Brian Lloyd welcomed the Group, asked for sign-in and led a short discussion on the mailing list, PPP archive availability and a history of the Working Group. Also Brian discussed current status and current implementations of PPP.

Bill Simpson reported that SAAG has reviewed the PPP authentication draft document. The result is that the message digest algorithm used in the Challenge Handshake Authentication Protocol (CHAP) may be either MD4 or MD5. The document is being changed to support this. The default algorithm will be the same as that chosen by the SNMP Working Group for SNMP authentication. [MD5 was chosen].

Brian Lloyd reported on the IPLPDN discussion on frame relay, X.25 and PPP over the same physical interface. They decided to use XID to distinguish which protocol will run on the link.

Brad Parker of Cayman gave a synopsis of the work on PPP in the AppleTalk Working Group. Apple has chosen to use PPP instead of a proprietary point-to-point protocol, thus paving the way for both IP and AppleTalk on the same serial interface. The result is a document that is ready for review by the PPPEXT Working Group. Two implementations are available. Brad has partially completed work on the drivers and an individual at the University of Michigan is planning on continuing the effort.

Philip Almquist presented the comments on the PPP requirements portion of the Router Requirements document. The members of the RREQ Working Group objected to listing line speeds above which Van Jacobson (VJ) header compression should not be used. The result was that the recommendation from the PPPEXT Working Group was changed to read that VJ header compression should be used below 20Kbps and may be used at any speed above that. The upper bound above which VJ header compression should not be used, previously set at 64Kbps, was removed.

Philip also reported that there were objections by the members of the RREQ Working Group to the requirement for Link Quality Monitoring (LQM). This led into a discussion of LQM. The issue was also raised that some of the vendors wish to do other forms of proprietary LQM.

One of the problems with the existing LQM is that it is considered to be part of the Link Control Protocol (LCP) and hence must use an Async Control Character Map (ACCM) of all 1's. This just about doubles the size of an LQM packet on an async link.

As a result, the LCP document will be modified to support a slightly different LQM negotiation that can support multiple types of LQM. If an implementation supports LQM at all, it must support the existing type of LQM so that there will be a common denominator (analogous to MIB-1 and MIB-2 of SNMP).

As a result of the LQM problem the Group decided that all LCP packet/option codes less than or equal to seven that are needed to bring the LCP to the open state must be escaped using the all-ones ACCM. After the link is open the other options, i.e., authentication, new LQM, etc., may be transmitted using the negotiated ACCM and compression options even though these packets are ostensibly LCP packets.

There is a problem that occurs when the LCP goes to the open state and a frame that has the ACCM set to zero (control characters not escaped) arrives at the receiver before the receiver has updated its ACCM and changed to the open state (this often occurs when the first Network Control Protocol (NCP) packet immediately follows the last LCP ack). The NCP frame is discarded at the receiver. There was a suggestion to insert a delay to allow the receiver to get to the open state before sending the NCP packet. It was noted that this is not a serious problem because the standard error recovery sequence properly deals with this. It was decided not to make a change in the state machine and to add an implementation note describing the problem.

There was concern about the length of time that it can take to determine that a link has failed (ten retries with three seconds between retries). The final decision was to make it clear that the three second delay may be adjusted to accommodate links with lower latency, i.e., that high speed link interfaces timeout values should be smaller. This information will be added to the LCP document and the default timeout value will become part of the PPP MIB.

Glenn McGregor presented his IPCP document and discussed the changes to the VJ header compression as used in PPP. Now, the slot number – which is used to identify a particular session being compressed – is not compressed. This greatly improves error recovery if a packet is lost or damaged in transit.

PPPEXT Minutes Evening Session

IP Address discussion continued. The Working Group decided to remove the feature for negotiating/reporting multiple IP addresses on an interface.

In addition the Working Group decided that the IP address negotiation procedure was too complicated to ensure that it worked properly. The Group decided on a much simpler scheme that retains all the features of the earlier version without the complexity. The IPCP document will contain a description of the old method along with a strong note indicating that implementations should use the new IP address negotiation procedure, and that the old IP address negotiation will be eliminated sometime in the not-too-distant future as the IPCP document proceeds down the standards track.

Bill Simpson and Brian Lloyd presented the Authentication Document. The section on management of secrets (keys) has a hole due to the lack of availability of a secure mechanism for the dissemination of the "secret". This will be gated by the work on Common Authentication Technology (CAT) and on SNMP secret dissemination technology.

Also the Challenge Handshake Authentication Protocol (CHAP) will change the way it uses MD5 to generate the authentication "signature" so as to be 100 percent compatible with SNMP. This should allow the core authentication procedures to be completely interchangeable between PPP and SNMP.

The discussion then proceeded to the call-back field of CHAP. The purpose of this field is for one end or the other to indicate to the peer that it wishes to terminate the link and call-back, primarily for purposes of reversing charges (some indicated that call-back may prove useful for enhancing security). Several people indicated that multiple call-back destinations may be desirable so a call-back address (phone number) field was defined and added.

Marty Del Vecchio from Shiva Corp presented Netware IPX Control Protocol which he has implemented. The Group suggested a number of changes and improvements. Marty will do further research and present an improved document soon.

Other documents were discussed. It was noted that 3Com has implemented stripped down versions of most of the NCPs. There was nothing to report on CLNP/OSI over PPP. AppleTalk over PPP is very close to completion. Michele Wright of Timeplex will take over the DECnet over PPP document. Several of the implementors present indicated that they are actively working on an implementation of PPP that supports DECnet.

The topic of conversation then moved on to switched circuit (dial-up, ISDN, etc.) connection techniques. A discussion then ensued about techniques for automatically starting PPP during a login process. It was noted that the first PPP frame on an async link consists of the octet sequence "7e ff 7d 03". This makes it possible for a terminal server or host to recognize that the peer wishes to run PPP and may start PPP immediately.

The discussion also went back to PPP over ISDN. The XID technique for determining which protocol would run, e.g., PPP, Frame Relay, or X.25, was discussed again.

The discussion then proceeded to the topic of inverse multiplexing, e.g., using multiple PPP links to simulate a single link/interface with greater bandwidth. There is a need to add a mechanism to indicate to the remote peer that one end or the other needs to increase capacity and will be opening an additional link. It was suggested that the new link need only open the LCP and authenticate, and there is no need to renegotiate the NCPs. The magic number that is negotiated on a link could be used as a logical connection number and can be made unique across all of the logical PPP connections, e.g., all physical connections that are part of a single logical interface will use the same magic number.

Results and Decisions

1. The Group decided to move the status of the LCP document back to "Proposed" because of the changes to LQM.

2. The Group decided to move the status of the IPCP document back to "Proposed" status because of the desired changes to the IP address negotiation.

3. The Group decided to keep the status of the Authentication document at "Proposed" status due to the changes in the CHAP.

Attendees

Steve Alexander	stevea@i88.isc.com
Fred Baker	fbaker@emerald.acc.com
Dean Cheng	dean@sunz.retix.com
Richard Cherry	rcherry@wc.novell.com
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Kenneth Crepea	crepea@cisco.com
Marty Del Vecchio	marty@shiva.com
Craig Fox	foxcj@network.com
Chris Gunner	gunner@osicwg.enet.dec.com
Bob Jeckell	robert_jeckell@nso.3com.com
William Jolitz	william@okeeffe.cs.berkeley.edu
Frank Kastenholz	kasten@europa.clearpoint.com
Tom Kessler	kessler@sun.com
Holly Knight	holly@apple.com
Gordon Lee	gordon@ftp.com
Brian Lloyd	brian@ray.lloyd.com
Glenn McGregor	ghm@merit.edu
Robert Morgan	morgan@jessica.stanford.edu
Dean Morris	morris@marvin.dec.com
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Chandy Nilakantan	csn@3com.com
J. Bradford Parker	brad@cayman.com
Miguel Sasson	sasson@xylogics.com
Mark Schaefer	schaefer@davidsys.com
William Simpson	Bill_Simpson@um.cc.umich.edu
Eric Smith	eric@telebit.com
Ravi Srinivasan	ravi@eng.vitalink.com
Bruce Taber	taber@interlan.com
Mark Therieau	markt@python.eng.microcom.com
William Townsend	townsend@xylogics.com
Maurice Turcotte	dnmrt@interlan.com
John Veizades	veizades@apple.com
Yuan Wang	natadm!ycw@uunet.uu.net
Scott Wasson	sgw@sgw.xyplex.com
Preston Wilson	preston@i88.isc.com
L. Michele Wright	uncng!michele@uunet.uu.net

3.2.8 Router Requirements (rreq)

<u>Charter</u>

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

Mailing Lists:
General Discussion: `ietf-rreq@Jessica.Stanford.edu`
To Subscribe: `ietf-rreq-request@Jessica.Stanford.edu`
Archive:

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 OSPF and BGP.

The Working Group will also instigate, review, or (if appropriate) produce additional RFCs on related topics. To date, Group members have produced draft documents discussing the operation of routers which are in multiple routing domains (3 papers), TOS, and a routing table MIB.

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 Working Group has decided that, unlike RFC-1009, the Router Requirements document should not discuss Link Layer protocols or address resolution. Instead, those topics should be covered in a separate Link Layer Requirements document, applicable to hosts as well as routers. Whether this Group will create the Link Layer Requirements is still to be determined.

Goals and Milestones:

Done	First Internet Draft version.
Done	Second Internet Draft version.
Done	Third Internet Draft version.
Sep 1991	Fourth Internet Draft version
Oct 1991	Final Internet Draft version.
Nov 1991	Submission for Proposed Standard.

Internet Drafts:

"Requirements for Internet IP Routers", 09/17/1990, Philip Almquist <draft-ietf-rreq-iprouters-03.txt>

"Ruminations on Route Leaking", 07/25/1991, Philip Almquist <draft-almquist-leak-00.ps>

"Ruminations on the Next Hop", 07/25/1991, Philip Almquist <draft-almquist-nexthop-00.ps>

"Type of Service in the Internet Protocol Suite", 07/25/1991, Philip Almquist <draft-almquist-tos-02.txt>

"Some Thoughts on Multi-Domain Routing", 07/25/1991, Ross Callon <draft-callon-routing-00.txt>

"IP Forwarding Table MIB", 08/14/1991, Fred Baker <draft-ietf-rreq-forwarding-04.txt>

CURRENT MEETING REPORT

Reported by Philip Almquist

Minutes of the Router Requirements Working Group (RREQ)

The Router Requirements Working Group met much more briefly than it had at previous IETF meetings. There were three primary activities:

1. Discussion and approval of a Forwarding Table MIB.
2. Revision of the Router Requirements Draft.
3. Coordination with other working groups.

In addition, the Chair gave a lengthy plenary presentation on Router Requirements in anticipation of its becoming a Proposed Standard before the next IETF meeting in San Diego.

Each of the three activities listed above is described in more detail below. The Chair would like to thank Frank Solensky for recording the decisions reached during the meeting.

Forwarding Table MIB

The Working Group discussed Fred Baker's Forwarding Table MIB proposal. One substantive flaw was found and fixed, and the revised version was deemed ready to be passed to the IESG as a candidate for a Proposed Standard.

The Group hotely debated the question of whether the proposed MIB ought to handle routing of IP multicasts and, if so, what modifications to the MIB would be required. The range of conflicting views on these questions suggested that multicast routing will need to be better understood before appropriate MIB support can be standardized. The Group felt that the Forwarding Table MIB was too valuable to be placed on hold indefinitely while multicast routing matures, but will revisit these issues before requesting that the MIB be advanced to Draft Standard status.

Router Requirements Document

The Working Group discussed a number of minor technical issues and requests for clarification. Most were disposed of with little debate. Probably the only notable decision was the one to lower the requirement level of MIBs other than MIB-II from MUSTs to SHOULDs, on the grounds that the number of MIBs which have been developed has reached the point where implementing all relevant MIBs may be becoming onerous.

Coordination with Other Working Groups

The Working Group held a joint meeting with the Inter-Domain Policy Routing Working Group to try to ascertain whether there were inconsistencies between the specifications the two Groups are producing. The primary focus of the discussion was the implications of the

IP-over-IP encapsulation used by IDPR (e.g., on ICMP error messages). There was also some discussion of the interactions between the IDPR protocol and other routing protocols. No particular inconsistencies between the work of the two Groups were identified, but apparently both Groups found the discussions interesting and informative.

Although the Working Group did not hold an official joint meeting with the BGP Working Group, the RREQ Chair and some other RREQ Working Group members attended the BGP sessions at which route leaking between BGP and OSPF was discussed. The goal (or at least the goal of the RREQ Chair) was to try to achieve consistency between the BGP Group's work and the parallel work on route leaking between arbitrary routing protocols that is being done in the RREQ Group.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Fred Baker	fbaker@emerald.acc.com
Atul Bansal	bansal@wile.nac.dec.com
William Barns	barns@gateway.mitre.org
Art Berggreen	art@acc.com
William Biagi	bbiagi@cos.com
Rob Coltun	rcoltun@ni.umd.edu
Dave Cullerot	cullerot@ctron.com
John Damiano	
Kurt Dobbins	dobbins@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Jim Ghadbane	jimgh@newbridge.com
Chris Gunner	gunner@osicwg.enet.dec.com
Frank Heath	heath@cmc.com
Ronald Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Michael Karels	karels@cs.berkeley.edu
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Yoav Kluger	ykluger@fibhaifa.com
Stev Knowles	stev@ftp.com
Ron Lau	rlau@synoptics.com
John Lekashman	lekash@nas.nasa.gov
Tony Li	tli@cisco.com
Gary Malkin	gmalkin@ftp.com
Mike Marcinkevicz	mdm@csu.net
April Merrill	abmerri@tycho.ncsc.mil
Donald Merritt	don@brl.mil
Dave Monachello	dave@pluto.dss.com
Dean Morris	morris@marvin.dec.com

Dennis Morris	morrisd@imo-uvax.dca.mil
John Moy	jmoy@proteon.com
Thomas Pusateri	pusateri@cs.duke.edu
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Richard Smith	smiddy@pluto.dss.com
Frank Solensky	solensky@clearpoint.com
Michael St. Johns	stjohns@umd5.umd.edu
Brad Steinka	brad@python.eng.microcom.com
Iris Tal	437-3580@mcimail.com
Sally Tarquinio	sally@gateway.mitre.org
William Townsend	townsend@xylogics.com
Yuan Wang	natadm!ycw@uunet.uu.net
Scott Wasson	sgw@sgw.xyplex.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov
Richard Woundy	rwoundy@ibm.com
John Ziegler	ziegler@artel.com

3.2.9 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
Archive:

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done Mailing list discussion of Charter and collection of concerns.

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

3.3 Network Management Area

Director(s):

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

Area Summary reported by James Davin/MIT

At the Santa Fe meeting of the IETF, six working groups of the Network Management Area held one or more sessions throughout the week. Two Birds of a Feather sessions were also held.

The SNMP Network Management Directorate reviewed six MIB specifications that had been recently reported out of working groups. Three of these were products of the Character MIB Working Group: the MIB for Character Stream Devices, the MIB for Parallel-Printer-Like Hardware Devices, the MIB for RS232-Like Devices. Also reviewed were the IP Forwarding MIB produced by the Router Requirements Working Group, the Frame Relay MIB produced by the IPLPDN Working Group, and the SMDS Interface Protocol MIB produced by the SNMP Working Group. The IESG announced its intention to consider these MIBs as candidates for Proposed Standard status after final text is available in the Internet-Drafts repository.

In addition, the Directorate discussed the problem of representing elaborate protocol stacks using the abstractions provided by the "interfaces" Group of MIB 2. The Directorate discussion was premised on the notion that is implicit in MIB 2 that an "interface" object is only used to represent protocol entities below the internetwork (e.g., IP) layer. The problem addressed has arisen in any number of Working Group discussions: although the interfaces Group in MIB 2 is a convenient abstraction for managers, it doesn't support specific transmission media or elaborate protocol stacks that may involve both downward and upward multiplexing.

The Directorate discussion came to three conclusions:

1. Every entry in a media-specific MIB table is paired one-to-one with a single entry in the interfaces table of MIB 2. The media-specific entry can be reached from the generic interfaces table entry by using information in the ifType object together with information in the ifIndex object.

2. Media-specific MIB table entries can (and often do) include "pointer" information that represents user-service relations among entities in a more or less elaborate protocol stack below the internetwork layer. This pointer information variously takes the form of OBJECT IDENTIFIER values (as in the Character MIB) or combinations of OBJECT IDENTIFIER and INTEGER values.

3. If every protocol entity below the internetwork layer is represented by an entry in the MIB 2 interfaces table, then all possible user-service relations among such entities may be concisely represented as a set of ordered pairs of ifIndex values. A simple MIB to represent such a set of ordered pairs was deemed desirable.

A document presenting these conclusions in greater detail will be prepared as a basis for broader discussion of this problem.

X.25 Management Information Base

The Working Group met to consider three documents: one that instruments X.25 link-layer functionality, one that instruments X.25 layer 3 functionality, and one that instruments convergence functions necessary to run IP over X.25.

At this meeting, the Working Group decided that the scope of instrumentation in the link-layer MIB will be confined to the LAPB protocol. The Working Group also concluded that the objects in these MIBs should be reviewed for actual usefulness in managing networks and that some pruning or alteration in conformance posture may be desirable. The Working Group noted that the IPLPDN Working Group was contemplating a revision to RFC 877 and decided to monitor that activity to determine if it may warrant revision to the IP/X.25 convergence MIB. The Group also discussed at some length the problems of representing X.25 protocol stacks in MIBs and suggested that the SNMP Directorate might pay some attention to this problem.

Remote LAN Monitoring

This Working Group met informally to discuss implementation experience with the recently published RMON MIB. At the suggestion of members who had attended the Birds of a Feather session on SNMP Device Discovery earlier in the week, the Working Group spent some time discussing ways in which RMON technology could be applied to the device discovery problem. The meeting also recommended that a new working group be formed to address extensions of the RMON MIB for Token Ring media.

IEEE 802.3 Hub MIB

This Working Group met to discuss the current draft of an SNMP MIB for 802.3 Repeater devices. The Chair reported on IEEE reaction to this first draft of the SNMP MIB. A presentation was made on ideas for a “Chassis MIB” that is useful in instrumenting communications products that encompass multiple devices. As a result of this presentation, the Working Group concluded that its repeater MIB need not accommodate multiple repeater devices as this need was better addressed by the notion of a Chassis MIB. The Working Group recommended that effort be applied to development of the Chassis MIB ideas.

Internet Accounting

This Working Group met in two sessions during the Santa Fe IETF meeting. The first session reviewed the Internet Accounting Background document (RFC 1272). Some time was spent bringing newcomers up to date with the Working Group's purpose and efforts. New attendees brought fresh perspectives and offered many comments, criticisms, and suggestions that will be incorporated into either a new version of the RFC or into follow-on documents.

The second session was spent in discussion of the latest draft of the Internet Accounting Architecture. Although this document has existed for several months now and has undergone three or four extensive revisions, it still needs work, both in form and content. The stated scope of the document was tightened. The Internet Accounting model and its difference from the OSI accounting model was more clearly defined. A decision was made to combine the metering services document (formerly to be separate) with the architecture document. A decision was made to announce the Working Group's intention to produce a draft MIB document before its work is concluded. Discussion of the architecture document will continue with a view to advancing it to the status of Internet Draft by the next IETF conference.

Simple Network Management Protocol

This Working Group met briefly in Santa Fe to conclude its business. The only item of outstanding business was the resolution of issues surrounding the Ethernet MIB. The Working Group Chair reviewed the course of action that had been previously discussed on the mailing list. With the formation of the Ethernet MIB Working Group to resolve outstanding issues, the SNMP Working Group adjourned and disbanded. The scheduled time that remained after adjournment of the SNMP Working Group was devoted to the first meeting of the new Ethernet MIB Group.

Ethernet MIB

The Ethernet MIB Working Group met for the first time in Santa Fe to begin its resolution of outstanding issues in the Ethernet MIB. The Working Group Charter was presented and interpreted by the Chair. The Working Group decided to omit from the current version of the MIB the language that dissociates conformance to the standard from actual implementation of the relevant objects. The Working Group felt that resolution of the issues required a more straightforward strategy that ties implementation requirements to particular operating environments. The Working Group also decided that distinctions between 802.3 and Ethernet environments could be a useful principle in articulating conformance requirements. The Group also agreed that distinctions between hardware and software implementations of MAC layer functions would also be an important consideration.

SNMP MIB Compiler

A Birds of a Feather session on SNMP MIB Compiler technology was conducted by Dave Perkins of SynOptics. Dave presented his recent work on MIB compiler technology and explained how it could be valuable both in syntax checking of MIB documents and as a tool to support development of SNMP agents.

SNMP Device Discovery

A Birds of a Feather session on SNMP Device Discovery was conducted by Fred Baker of ACC. Much time was spent in this session attempting unsuccessfully to formulate an adequate definition of the problem. The session articulated some ideas on how remote monitoring technology could be applied to the device discovery problem, and these were subsequently presented to the RMON MIB Working Group for its consideration.

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

Minutes of the SNMP Device Discovery BOF (DEVDISC)

Prior to meeting in Santa Fe, there was an extended discussion on the finder@emerald.acc.com and SNMP-WG@nisc.nyser.net mailing lists. It is summarized as follows, as it represents a significant context to the BOF held at the IETF meeting.

Essentially, we have two problems and at least three solutions on the table. The purpose of the BOF is exploratory - there exists a subset of individuals who feel that there is no viable problem to solve, and if there is it should not be solved; there are others who support various viewpoints. We need to put all of the issues on the table and come up with a problem statement before we can either proceed or decide not to proceed. The problems are:

1. Within a single administrative domain, it should be possible for Network Management Systems to locate all of the systems appropriate for them to manage (e.g., with SNMP) without preconfiguration. This is believed to be helpful to network managers in that they now have positive assurance that they do in fact know all of the key devices in their networks. This viewpoint has been presented by a couple of vendors, and was in fact the start of the discussion.

2. Within a single administrative domain, it is possible and probable that devices are added to the network without the knowledge of the network manager. Several network managers have indicated a desire to know literally all of the devices on their networks, and their network layer attributes.

The potential solutions may be classified as "first person", "second person", and "third person" solutions, and there are a couple of variations on each of those:

First Person:

Examples of current deployment:

- Wide area: RWHO...
- Immediate Neighbor: OSPF, ES-IS, IS-IS, DECNET, RIP, DECNET, DEC MOP, DEC LAT...

Each SNMP-manageable device on the network periodically emits a trap which announces its presence to interested parties. The trap is sent to a multicast which is received by interested parties on the extended LAN. Its contents include Object Identifiers of MIB Groups supported by the device, system.sysObjectID, and the Read-only community string/party to be used with this agent. If we presume that the probability that a multicast will reach all of its intended recipients > some value, then the probability that all of the network managers know about all of the devices they should manage within some amount of time is a function of the emission rate and the time limit.

A second version of this might use IP Multicasting to propagate information throughout the administrative domain.

Concerns:

- First approach: Impact of SNMP Security Architecture not yet analyzed. Does not propagate information beyond router.

- Second approach: Scaling, definition of administrative boundaries, some details in SNMP. Impact of SNMP Security Architecture not yet analyzed.

- Doesn't solve second problem.

Second Person:

Examples of current deployment:

- ARP
- 802.5 RIF Discovery
- DEC RBMS

Each interested party does something to elicit a response from the systems it is concerned about. This might include sweeping MIBs and then pinging new folks discovered in ARP caches, etc. Someone has suggested letter bombs - broadcast a GET system.sysDescr, and collect the responses. In the latter class of solution, there would need to be either some random "host delay" to avoid flooding the network, or an "exclusion group" to advise responders to NOT respond.

Concerns:

- Scaling, traffic level, both burst and sustained, definition of administrative boundaries.

Sweeps may solve second problem, or at least part of it, but this is not assured. broadcast "pings" only solve it for the architectures whose "ping" is used, and not all architectures define a "ping".

Third Person:

Examples of current deployment:

- RMON MIB

A subset of the systems in the network actively notify the interested NMSs of new systems that they detect. "Detection" is somewhat imprecise - one proposal defines detection to be a protocol specific neighboring relationship; another defines it as the use of a LAN source address. In the latter, the RMON MIB is proposed as a solution.

Concerns:

- With the RMON MIB, no network layer information is captured. If the network manager is not on the local wire with the system found, it has no information other than the MAC Address and the location of the monitor with which to do anything further and no protocol with which to get it.

- With the RMON MIB, only LAN systems are detected, and then only on LANs that have objects defined in RMON. As it stands today, RMON is fairly obviously targeted at Ethernet. For use on Token Ring or FDDI, there is additional work defined by the RMON Working Group. Multipoint networks such as SMDS and Frame Relay are not addressed; this may or may not be an issue - can we assume that contracts exist in the presence of these technologies? Are private networks a concern?

- With the protocol specific detection, a router or bridge could advertise the MAC and network layer information to the NMS; the fact that a TRAP is unreliable means that the NMS might nonetheless fail to learn the information. Use of a SET has been suggested, but some feel that specifying an application residing in the router or bridge is distasteful. Each NMS could also poll the subset of systems (monitors, routers, etc., a limited subset of the network) for new information.

The BOF was started with a presentation by Anil Rijsinghani of Digital, whose question on the SNMP Mailing List is what actually started the whole debate. His fundamental concern, echoed by some other vendors, was that there is today no single, reliable, way to find all of the SNMP Manageable devices in an administrative domain. As a corollary to that, there is no way to determine what MIBs any given station supports. Even a MIB walk may not return that information if a MIB is primarily composed of tables and the service is not currently configured or active. Mechanisms that are available depend on assumptions that may not hold, such as the use of the "public" community in SNMP or that SNMP capable systems periodically send SNMP messages. Other drawbacks of existing mechanisms may include: they are complex, generate excessive traffic, and require every NMS to perform its own discovery. Requirements of a solution to this problem include: it should be reliable (discover every SNMP device), be simple, use small amount of network bandwidth, require a small amount of agent effort, should work regardless of powerup sequence, impose a low load on others and convey useful standardized information.

The remainder of the BOF was given over to determining what problem the assembled company wanted to solve; this is a non-trivial problem in its own right. The discussion was wide-ranging, and a number of quite divergent opinions were presented. It was generally felt that the problems of finding all SNMP capable systems, finding all SNMP/UDP/IP capable systems, and finding all systems that use the Internet were quite distinct and call for different solutions, and that finding all equipment attached to the Internet is not a solvable problem.

After much discussion, it was concluded that the fundamental problem seeking solution borrowed components of each of these problems. Network managers do in fact need to know what equipment is attached to their networks, and are helped by products which will perform this function. Products that do this utilize the RMON MIB, proprietary MIBs and algorithms, and scan such tables as the ARP cache and Routing cache. However, the problem of device discovery does not include a number of other functions (such as drawing a picture or matrix of Internet connectivity). These are "next step" processes which follow the discovery of the systems in the network.

Given this much problem definition, the conclusion was reached that the RMON MIB could be extended to solve much of the discovery process. The reasons that it is inadequate now are:

- It is limited to finding systems attached to LANs, and
- It does not capture the protocol type or network layer protocol addresses that a device is using.

As a result, the information captured about a system found by RMON, as it stands, cannot be used to perform the next step, that of pinging the device, especially if the device is separated from the NMS by a router. Therefore, the ultimate solution reached was to recommend that the RMON MIB be extended with a table containing, at minimum, the following information:

```
deviceTable
  deviceEntry [deviceMacAddress, deviceProtocol]
    deviceMacAddress OCTET STRING
    deviceProtocol OCTET STRING or OBJECT IDENTIFIER
    deviceProtocolAddress OCTET STRING
```

There may not be a protocol address for all protocols layered onto the Data Link Layer, so the NMS must expect that deviceProtocolAddress may have a length of zero octets.

A prototype MIB will be forwarded to Mike Erlinger for consideration by the RMON Working Group.

Attendees

Miriam Amos Nihart	miriam@decwet.zso.dec.com
Robert Austein	sra@asylum.sf.ca.us
Fred Baker	fbaker@emerald.acc.com
Steve Bostock	steveb@novell.com

David Bridgham	dab@asylum.sf.ca.us
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Buffum	buffum@vos.stratus.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Richard Cherry	rcherry@wc.novell.com
James Codespote	jpcodes@tycho.ncsc.mil
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
William Jackson	jackson@manta.nosc.mil
Ron Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Scott Kaplan	scott@ftp.com
Frank Kastenholz	kasten@europa.clearpoint.com
David Kaufman	dek@proteon.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
Stev Knowles	stev@ftp.com
Deidre Kostick	dck2@sabre.bellcore.com
Ron Lau	rlau@synoptics.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
David Minnich	dwm@fibercom.com
Michael Newell	mnewell@nhqvax.hg.nasa.gov
David Perkins	dperkins@synoptics.com
Radia Perlman	perlman@radia.enet.dec.com
Robert Purvy	bpurvy@us.oracle.com
Anil Rijsinghani	anil@levers.enet.dec.com
Michael Ritter	mwritter@applelink.apple.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Timon Sloane	peernet!timon@uunet.uu.net
Ravi Srinivasan	ravi@eng.vitalink.com
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Iris Tal	437-3580@mcimail.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
John Veizades	veizades@apple.com
Sudhanshu Verma	verma@hpspdla.hp.com

Lee Wade	wade@nsipo.arc.nasa.gov
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jeremy Wilson	
Junekang Yang	natadm!yang@uunet.uu.net
John Ziegler	ziegler@artel.com

CURRENT MEETING REPORT

Reported by David Perkins/SynOptics

Minutes of the SNMPMIB Compiler BOF (MIBCOMP)

This BOF was a presentation by Dave Perkins of a unique architecture for an SNMP MIB Compiler. The important aspects of this project are that:

1. The compiler is split into a separate front-end and a replaceable back-end;
2. The front-end does extensive syntax and semantic checks much better than MOSY and includes support for traps, multiple modules, imports, and textual conventions;
3. The back-ends are easy to write and can be used for specific applications; and
4. The intent is to make the source code "public domain" so that it can be used by any interested party.

Example users of the system include the following:

1. SNMP agent developers - a back-end can be written to generate MIB data structures and dispatch tables specific to an agent implementation.
2. MIB developers - the front-end provides extensive error checking. A back-end can be written to print "reports".
3. Management station developers - a back-end can be written which merges additional fields with those from MIB objects so that database records can be generated for a generic MIB query system.
4. SNMP tool developers - a back-end can be written which formats the MIB information so that it can be used by existing tools that require MIB object information in a format other than the concise MIB format.

Dave presented an overview of the MIB compiler architecture and gave a status report of his current implementation experience. Key points included the following: the front-end is currently implemented under MS-DOS and was used on a laptop at the IETF to do instant MIB checking; much testing has been done including all the MIBs in RFCs and many of the MIBs in Internet Drafts; and an example back-end was written which demonstrated selection of MIB groups for reporting.

There were many action items to be done by the next IETF meeting. The key ones included: porting to UNIX; working out the details so the code could be made "freely available"; and setting up a directory where sources and documentation could be reached via FTP access.

Attendees

Miriam Amos Nihart	miriam@decwet.zso.dec.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
David Bridgham	dab@asylum.sf.ca.us
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Buffum	buffum@vos.stratus.com
Dave Cullerot	cullerot@ctron.com
Michael Erlinger	mike@lexcel.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Ronald Jacoby	rj@sgi.com
Manu Kaycee	kaycee@ctron.com
Yoav Kluger	ykluger@fibhaifa.com
Ron Lau	rlau@synoptics.com
Ron Mackey	rem@dsiinc.com
David Perkins	dperkins@synoptics.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
June-Kang Yang	natadm!yang@uunet.uu.net
Wengyik Yeong	yeongw@psi.com
John Ziegler	ziegler@artel.com

3.3.1 Bridge MIB (bridge)

Charter

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

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

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 draw from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

Goals and Milestones:

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

Request For Comments:

RFC 1286	"Definitions of Managed Objects for Bridges"

3.3.2 Character MIB (charmib)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

General Discussion: char-mib@decwrl.dec.com
To Subscribe: char-mib-request@decwrl.dec.com
Archive:

Description of Working Group:

The Character MIB Working Group is chartered to define a MIB for Character Stream Ports that attach to such devices as terminals and printers.

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

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

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

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

Goals and Milestones:

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

Internet Drafts:

"Definitions of Managed Objects for RS-232-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-rs232like-03.txt>

"Definitions of Managed Objects for Parallel-printer-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-parallelprinter-02.txt>

"Definitions of Managed Objects for Character Stream Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-charmib-02.txt>

3.3.3 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):
Jonathan Saperia, saperia@tcpjon.enet.dec.com

Mailing Lists:
General Discussion: phiv-mib@jove.pa.dec.com
To Subscribe: phiv-mib-request@jove.pa.dec.com
Archive:

Description of Working Group:

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

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

Goals and Milestones:

Done	Review and approve the Charter and description of the Working Group, making any necessary changes. At that meeting, the scope of the work will be defined and individual working assignments will be made.
Done	Mailing list discussion of Charter and collection of concerns.
Done	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.
Done	Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.
Done	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.

Request For Comments:

RFC 1289 "DECnet Phase IV MIB Extensions"

3.3.4 Ethernet MIB (ethermib)

Charter

Chair(s):

Frank Kastenholz, kasten@europa.clearpoint.com

Mailing Lists:

General Discussion: enet_mib@europa.clearpoint.com
To Subscribe: enet_mib-request@europa.clearpoint.com
Archive: Not available

Description of Working Group:

This Working Group is charged with resolving the outstanding conformance issues with the Ethernet MIB in preparation for its elevation from Proposed to Draft Standard status. Specifically, this Working Group shall:

(1) Develop a document explaining the rationale for assigning MANDATORY status to MIB variables which are optional in the relevant IEEE 802.3 specification (the technical basis for the Internet Ethernet MIB). This shall not be a standards-track document.

(2) Develop an implementation report on the Ethernet MIB. This report shall cover MIB variables which are implemented in both Ethernet interface chips, and in software (i.e., drivers), and discuss the issues pertaining to both. This report shall also summarize field experience with the MIB variables, especially concentrating on those variables which are in dispute. This document shall not be a standards-track document. While the Ethernet MIB is progressing through the standardization process, this document shall be periodically updated to reflect the latest implementation and operational experience.

(3) Work to reconcile the differences regarding MANDATORY and OPTIONAL MIB variables with the IEEE 802.3 Management Specification.

(4) Extend explicit invitations to the members, reviewers, and participants of the IEEE 802.3 committee to participate in the Working Group's efforts. This will ensure that as much Ethernet and IEEE 802.3 expertise as possible is available.

(5) Maintain a liaison with the IEEE 802.3 committee. All documents produced by the Working Group will be forwarded to the IEEE 802.3 committee for their consideration as contributions to their efforts.

(6) Modify the "grouping" of variables in the MIB, in the light of the implementation and operational experience gained, in order to effect the desired conformance groupings.

This Working Group is chartered to make only changes to the MIB that fall into the following categories:

(1) Division of variables into MIB groups. This may necessitate adding or deleting groups and conceptual tables and moving variables among said groups and conceptual tables. Doing so may require the addition or deletion of variables necessary to support the conceptual tables (e.g., the ...Table, ...Entry, and ...Index types of variables). These changes may be necessary to align the MIB with the work of other standards bodies, the needs of implementors, and the needs of network managers in the Internet.

(2) Changing the conformance requirements of the MIB groups in order to align the MIB with the work of other standards bodies, the needs of implementors, and the needs of network managers in the Internet.

(3) Deleting variables from the MIB on the basis of implementation and operational experience showing that the variables are either unimplementable or have little practical operational value.

The Working Group is explicitly barred from making changes to the definition or syntax of objects nor may the Working Group add objects to the MIB except as may be required by Point 1 above.

Goals and Milestones:

TBD Draft Variable Status Rationale document.

TBD Develop Implementation Report.

CURRENT MEETING REPORT

Reported by Frank Kastenholz/Clearpoint Research Corp

Minutes of the Ethernet MIB Working Group (ETHERMIB)

The first order of business was to review the administrative issues surrounding the Working Group:

1. The Charter and the "Pax Davin" were reviewed and discussed at length. All present agreed to strictly adhere to the rules of Charter and the "Pax Davin".

2. The Working Group decided to try to have the documents required by the Charter ready for publication at the same time as the MIB is put forward for Draft Standard status. This would be in about six months. These two documents are:

   (a) An explanation for the assignment of MANDATORY status to optional 802.3 variables, and

   (b) An implementation report on the MIB variables.

   (c) A mailing list will need to be created.
   [This has been done – enet_mib@europa.clearpoint.com.]

After discussing administrative issues, the Working Group turned its attention to the MIB itself. The following items were discussed (thanks to Anil Rijsinghani of DEC for his notes of the meeting). Any changes to the MIB will be made to the version to be put forward for DRAFT STANDARD status.

1. The Working Group discussed the text in the MIB which allows an implementation to return 0 for counters for which the underlying events are not counted. It was realized that this wording makes it impossible to disambiguate the two cases of not implementing a counter and and 0 occurrences of the underlying event.

   The Working Group discussed the issue and a vote was taken on it. The Group decided to remove the offending text from the document. The Working Group realized that for variables to which this text applies, there are four alternatives which should apply:

   (a) Delete the variable from the MIB entirely as its utility has not been demonstrated by wide implementation experience,

   (b) Move the variable into a separate optional MIB group,

(c) The implementor must figure out some way to support the variable, and

(d) The implementor would not implement the variable, return noSuchName errors whenever the variable is accessed and not claim compliance to the MIB.

2. The possibility of an 802.3 specific (as opposed to DIX Ethernet specific or common to both) MIB group was discussed. It was decided to continue this discussion on the mailing list.

3. The dot3StatsExcessiveDeferrals object is implemented in only one chip out of 14 studied by Anil Rijsinghani. It should either be made optional or be removed from the MIB. The other mandatory objects are implementable with commonly available chips and supporting software. This will be considered for further study on the mailing list.

4. The TDR definition in the MIB is not sufficient, given that IEEE 802.3 does not define this object. It does not describe the two conditions under which the object is defined, and how to distinguish between them (short and open cable faults). Anil Rijsinghani will contribute text to clarify this.

Attendees

Steve Alexander	stevea@i88.isc.com
Miriam Amos Nihart	miriam@decwet.zso.dec.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Philip Budne	phil@shiva.com
Jeffrey Buffum	buffum@vos.stratus.com
James Codespote	jpcodes@tycho.ncsc.mil
Tracy Cox	tacox@sabre.bellcore.com
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
William Jackson	jackson@manta.nosc.mil
Ole Jacobsen	ole@csli.stanford.edu
Ronald Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Scott Kaplan	scott@ftp.com
Frank Kastenholz	kasten@europa.clearpoint.com

Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
Ron Lau	rlau@synoptics.com
Evan McGinnis	bem@3com.com
Lynn Monsanto	monsanto@sun.com
David Perkins	dperkins@synoptics.com
Anil Rijsinghani	anil@levers.enet.dec.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@interlan.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Frank Welch	
Jeremy Wilson	
June-Kang Yang	natadm!yang@uunet.uu.net
Wengyik Yeong	yeongw@psi.com
John Ziegler	ziegler@artel.com

3.3.5 IEEE 802.3 Hub MIB (hubmib)

Charter

Chair(s):

Keith McCloghrie, kzm@hls.com
Donna McMaster, mcmaster@synoptics.com

Mailing Lists:

General Discussion: hubmib@synoptics.com
To Subscribe: hubmib-request@synoptics.com
Archive: pub/humbib:sweetwater.synoptics.com

Description of Working Group:

This Working Group will produce a document describing MIB objects for use in managing Ethernet-like hubs. A hub is defined as a multiport repeater that conforms to Section 9, "Repeater Unit for 10 Mb/s Baseband Networks" in the IEEE 802.3/ISO 8802-3 CSMA/CD standard (2nd edition, Sept. 1990). These Hub MIB objects may be used to manage non-standard repeater-like devices, but defining objects to describe vendor-specific properties of non-standard repeater-like devices are outside the scope of this Working Group. The MIB object definitions produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions.

In order to minimize the instrumentation burden on managed agents, the MIB definitions produced by the Working Group will, wherever feasible, be semantically consistent with the managed objects defined in the IEEE draft standard P802.3K, "Layer Management for Hub Devices." The Working Group will base its work on the draft that is the output of the July 1991 IEEE 802 plenary meeting. The Working Group will take special cognizance of Appendix B of that specification that sketches a possible realization of the relevant managed objects in the SNMP idiom.

Consistent with the IETF policy regarding the treatment of MIB definitions produced by other standards bodies, the Working Group may choose to consider only a subset of those objects in the IEEE specification and is under no obligation to consider (even for "Optional" status) all objects defined in the IEEE specification. Moreover, when justified by special operational needs of the community, the Working Group may choose to define additional MIB objects that are not present in the IEEE specification.

Although the definitions produced by the Working Group should be architecturally consistent with MIB-II and related MIBs wherever possible, the Charter of the Working Group does not extend to perturbing the conceptual models implicit in MIB-II or related MIBs in order to accommodate 802.3 Hubs. In particular, to the extent that the notion of a "port" in an 802.3 Hub is not consistent with the notion of a network "interface" as articulated in MIB-II, it shall be modelled independently by objects defined in the Working Group.

Because the structure of 802.3 Hub implementations varies widely, the Working Group shall take special care that its definitions reflect a generic and consistent architectural model of Hub management rather than the structure of particular Hub implementations.

The IEEE Hub Management draft allows an implementor to separate the ports in a hub into groups, if desired. (For example, a vendor might choose to represent field-replaceable units as groups of ports so that the port numbering would match a modular hardware implementation.) Because the Working Group Charter does not extend to consideration of fault- tolerant, highly-available systems in general, its treatment of these groups of ports in an 802.3 Hub (if any) shall be specific to Hub management and without impact upon other portions of the MIB.

Goals and Milestones:

Done	Distribute first draft of documents and discuss via E-mail.
Done	Working Group meeting as part of IETF to review documents.
Sep 1991	Distribute updated documents for more E-mail discussion.
Nov 1991	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Jan 1992	Documents available with specified changes incorporated.

Internet Drafts:

"Definitions of Managed Objects for IEEE 802.3 Repeater Devices", 07/23/1991, Donna McMaster, Keith McCloghrie <draft-ietf-hubmib-mib-00.txt>

CURRENT MEETING REPORT

Reported by Donna McMaster/SynOptics

Minutes of the IEEE 802.3 Hub MIB Working Group (HUBMIB)

The meeting was called to order by Co-Chairs Keith McCloghrie and Donna McMaster.

Agenda

- Introduction
- Chassis MIB presentation (Keith)
- Repeater ID discussion and resolution
- Report on IEEE 802.3 Hub Management ballot (Donna)
- Discussion of outstanding issues
    - From Section 8 of the current draft
    - From the mailing list since the Atlanta meeting

There were no changes to the draft, mailing list, or archive site since the last meeting. The current draft is still the July 22, 1991 version. The Working Group mailing list is hubmib@synoptics.com. Requests should be sent to hubmib-request@synoptics.com. Drafts and mail are archived in pub/hubmib on sweetwater.synoptics.com, and can be accessed using anonymous ftp.

Donna will add all meeting attendees to the hubmib mailing list.

Chassis MIB

There has been significant discussion about the repeater ID. Several parties have expressed the opinion that the repeater ID is not the best solution to the problem of managing multiple repeaters with a single agent, but that the problem needs to be addressed.

Keith presented an alternate proposal, dubbed a "Chassis MIB." This MIB defines objects for managing a "box" containing assorted network devices such as repeaters, bridges, routers, and/or terminal servers. Keith's slides are reproduced below.

```
CHASSIS MIB

  How to manage a box containing multiple modules.
  o  Multiple Physical Modules - slots
  o  Multiple Logical Devices - repeaters, bridges, etc.
  o  Multiple Backplane ''Wires'' - Ethernet, Token Ring, FDDI, etc.
  o  Power Supply - need separate MIB
```

PHYSICAL DEVICE TABLE

What's in the Slot ?
- o Index by slot-number
- o Board Type - an OID, common values defined for empty and unknown
- o Last change - sysUpTime at last insert/removal

LOGICAL DEVICE TABLE

- o Index by integer
- o Function - a sum of values, one value for each of repeater, bridge, router, terminalServer, management card, etc.
- o ObjectId = sysObjectId
- o Party - a SNMP party OID, or 'noParty'
- o Community - community-string or empty
- o IpAddress - IP Address for use with community

BACKPLANE WIRES TABLE

- o Indexed by integer
- o Type - an OID
- o Other ??

RELATION TABLE

Which device(s) are in which slot(s) and connected to what wires on the backplane
- o Each entry represents one relation
- o Each entry contains three pointers:
- o 1st pointer is the slot number
- o 2nd pointer is the logical device index
- o 3rd pointer is the backplane wire index

An entry means that the module in the indicated slot is (part of) the indicated logical device and is connected to the indicated backplane wire.

EXAMPLE

Slot	Device	Backplane
1	1	1
2	1	1
3	2	1
3	2	2
4	3	2

```
        5            3              2

    devices 1,3 are repeaters; device 2 is a bridge.
```

Vigorous discussion ensued. There were many questions, and general enthusiasm. Some of the issues raised:

- Questions about physical vs. logical devices
- Use netAddr instead of IP address
- Multiple addresses for the same agent (router, or OOB): could make multiple entries in the device table. Would need an additional index variable for the table.
- What community string does each device send in traps – that is, if one agent represents multiple repeaters, how does the trap receiver determine which repeater is referenced by the trap?

The enthusiasm threatened to use all the time allocated for discussion of repeater MIB issues, so a straw poll was taken to see if a new effort should be undertaken to develop a Chassis MIB. Straw poll question: Do people believe that the development of a Chassis MIB is a useful and feasible project? Strong consensus that a Chassis MIB is both useful and feasible, no opposition was expressed.

Repeater ID

Keith briefly recapped the repeater ID issue and opened the floor to debate. Several people expressed the opinion that the repeater ID is not the appropriate mechanism for handling multiple repeaters, and that energy should be directed instead toward development of a Chassis MIB.

No one was speaking in favor of the repeater ID, so a straw poll was taken. Twelve people indicated preference for dropping the repeater ID; one (Jeff Case) wanted to keep the ID. When asked for comment, Jeff explained that it was a simple solution to a current, real problem, but that he knew better than to fight overwhelming odds.

Donna presented a letter from IEEE 802.3 Hub Management members Kathy de Graaf (DAVID Systems), Steve Horowitz (Chipcom), and Jim Reinstedler (Ungermann-Bass), arguing to keep the repeater ID. Their conclusion is that the repeater ID "provides a simple, inexpensive, standard, interoperable, and useful way of allowing a single agent to address multiple repeaters." (Full text of the letter will be published in the Proceedings.) Discussion was invited; no one had changed his/her opinion. No representatives from Chipcom or Ungermann-Bass were present to comment. Mark Schaefer from DAVID Systems declined to comment on the letter, saying that he personally prefers the Chassis MIB solution.

In light of the strong consensus, the Working Group officially decided to remove the repeater ID from the MIB, effectively making the MIB definitions represent a single repeater instead of a collection of repeaters.

IEEE Report

Donna presented a summary of IEEE 802.3 Hub Management Task Force (802.3 HMTF) activities. The 802.3 HMTF circulated a draft for letter ballot in early August. The draft received 325 comments from 64 balloters, with an initial approval rate of 71 percent. All comments were addressed in meetings held at the IEEE 802 Plenary November 10-15. Enough comments were favorably resolved to raise the approval rate above the 75 percent needed to consider the ballot formally passed.

802.3 HMTF made a number of changes in their draft as a result of the comment resolution process. A new draft will be mailed out for confirmation ballot in December, closing in January. (The confirmation ballot process is intended to verify that changes address voters' concerns without creating new problems.)

The overall 802.3 Working Group also chartered new activities for defining MAU management information and for rewriting the current 802.3 layer management standard in the ISO GDMO format. The MAU management effort will include such information as media type (e.g., 10BASE-T or coax) and link status.

A summary of the major changes being made in the 802.3 Hub Management draft:

1. The term "hub" is being changed to "repeater."

2. The SNMP encodings in Annex B are being replaced with a reference to the work of the IETF Hub MIB Working Group.

   Case questioned whether IEEE was dropping the SNMP encodings because they consider SNMP to be a "substandard" management protocol. Donna stated that 802.3 uses ISO GDMO encodings because their standards are forwarded to ISO after adoption by IEEE. Removing the SNMP encodings was done to acknowledge that the IEEE does not believe it appropriate to "compete" with the IETF in developing SNMP MIBs.

   However, the 802.3 HMTF is very interested in SNMP, and most of the companies represented in that group are implementing SNMP management of their repeaters. Given that strong level of interest in SNMP, their action indicates a willingness to "trust" the IETF Hub MIB Working Group.

3. The concept of "groups" was modified in several ways. The "group" concept has always been a logical concept with references to possible physical mappings. In the new draft, all references to physical embodiments of groups are being removed, making a group a purely logical construct.

The group definition has also been changed to allow non-contiguous port numbering, and to allow ports to be added to groups or removed from groups without resetting management. Previously, a group had a fixed number of ports "N", and the ports in the group were numbered from 1 through N. To effect this change, the groupNumberOfPorts attribute was replaced with groupPortCapacity, and a groupPortMap attribute was added.

These 802.3 HMTF port/group changes generated much discussion in the IETF Hub MIB Working Group, detailed in section V below.

4. The repeater-level MJLPs counter was replaced with a per-port equivalent called "veryLongEventReceived" counter.

5. ExecuteSelfTest2 was considered to be redundant with the resetHub command, and was eliminated. ExecuteSelfTest1 was renamed to be execNonDisruptiveSelfTest.

6. One balloter suggested that hubHealthData should be left for vendor extensions, as it cannot be interpreted in a vendor-independent manner. After some discussion, 802.3 HMTF decided to keep hubHealthData as "an opportunity for implementation agreements."

7. The shortEvents and runts counter definitions were changed, and several other counter definitions were made more clear. The "runtMaxTime" number (that differentiates between a long but legal collision fragment and a late collision) was debated and left unresolved. A conference call between repeater experts is being scheduled, and 802.3 HMTF agreed to let the members of the conference call specify the value to be used.

The next questions for the Hub MIB Working Group (IETF flavor) are whether to incorporate these 802.3 HMTF changes in our draft, and if so, when the changes should be made. All agreed that technical changes to counter definitions must be reflected in the IETF MIB. We also agreed to wait until after the confirmation ballot closes so that our draft doesn't thrash unnecessarily.

When the confirmation ballot is complete, Donna will convey the ballot results to the Working Group along with a proposal for incorporating changes.

Draft Status

Jeff Case suggested the draft might be ready for forwarding to Proposed Status. There were mutterings of concern over changes that might be made in this meeting. Agreement was reached to postpone the question until later in the meeting.

We later agreed that we will not forward the document to the IESG. The editors will update the draft with changes from this meeting and from the IEEE confirmation ballot, and publish for discussion at the next IETF meeting. The goal will be approval of the Working Group and submission as recommended for Proposed Standard status.

Groups of Ports

(In reference to IEEE item 3, above.) The Hub MIB Working Group members shared a strong consensus that the reason for defining port groups is to assist the user in mapping the port numbers to the physical devices. This is in direct opposition to the IEEE's direction of stating that the group mapping is purely logical. The Working Group agreed that the draft will continue to state that implementors may assign group and port numbers as desired, but that we strongly recommend that group and port mappings match the physical manifestation of the repeater as closely as possible.

The Working Group agreed to accept the IEEE's change to allow ports within a group to come and go. Does this imply a need for portUpTime as well as for groupUpTime? This would add complexity to every implementation whereas having ports moving between groups/repeaters is expected to be the less common case. Much discussion, decided not to add portUpTime.

Discussion of portMap. The Working Group observed that this information can be deduced from other existing objects in a single powerful Get-Next PDU (though not in a single wimpy Get PDU), and also observed that this configuration information will not change frequently. The same applies to the groupMap. Both groupMap and portMap are therefore redundant, and there was a general feeling that the overhead of collecting the information does not justify the optimization of packaging the information into a bit map. We decided that groupMap will be removed, and we will not add portMap.

How to handle the table rows for groups that are removed from a repeater or ports that are removed from a group? Delete the rows? Or have a state column in the table with a "not here" value to indicate a port/group that has trotted off into the sunset? Jeff Case: in other such cases, we have left this to the discretion of the implementor. There was general agreement that the implementor should choose when it is appropriate to remove the table row and when it is appropriate to return a state indicating that the group/port is unavailable for service.

It was further observed that "not here" could mean "switched to the other repeater in this box" or it could mean that a plug-in module was removed or had failed. There was some discussion about having an operState column that could be used for various flavors of broken or "not here." This idea was greeted favorably, and discussed with other objects later in the meeting (below).

Issues from Draft Section 8

Some of the section 8 issues had been previously resolved; we covered them briefly just for completeness. Numbers below correspond to Section 8 headings.

8.1. Optional groups: agreed to keep all three groups (mandatory Basic, optional Monitor and Address Tracking).

8.2. Multiple repeaters: removing repeater ID, see II above.

8.3. System objects (rptrBasManufacturer, rptrBasProduct, rptrBasVersion): agreed to take them out.

8.4. Health information: Agreed to take out rptrHealthData. Should be in vendor-specific MIBs, since it cannot be decoded in a standard way. "If people implement this instead of something that users can understand, we've done a disservice."

8.5. Additional group information: Keith showed a matrix of administrative objects relating to repeater, groups, and ports, and the Working Group discussed which administrative objects should be included for each of the three. The resulting table is shown below. The only changes from the current draft are in the operState column. Details of the proposed changes are listed below the matrix.

	admin state	oper state	reset	self test	upTime
repeater	NO	YES (1)	YES	YES	NO
group	NO	YES (2)	NO	NO	YES
port	YES	YES (3)	NO	NO	NO

1. Rename rptrHealthState to be rptrOperState.
2. Add new groupOperState object.
3. Add new portOperState object. Some discussion about whether this should be combined with autoPartitionState. Donna disagreed, because autoPartitionState is very specifically defined for repeater hardware. Agreed to define enumerations for portOperState and see then whether combining with autoPartitionState makes sense.

8.6. Carefully-crafted counter comments: committee condemns; clearly cannot condone.

Issues from Mailing List

Keith had slides listing all issues discussed on the mailing list since the last meeting, and the Working Group addressed each of them in turn.

Broadcast, Multicast Counters: These were not included in IEEE and earlier IETF drafts because they can be collected by a promiscuous monitor anywhere in the unbridged LAN segment and mapped to senders using the packets' source addresses. After discussion, there was agreement not to add broadcast, multicast counters.

Total Counters Discussed optimizing the collection of counts for a repeater by offering repeater (or even group?) total counts. This issue is similar to the portMap/groupMap issue, but counters (esp. errors) need to be collected much more frequently in order to track the health of the network. Also, it is not unusual for a single repeater to have over 100 ports, causing high collection overhead.

After discussion, the Working Group agreed that total counts are appropriate for some set of information. Proponents of totals are asked to submit proposed sets of total counters to the mailing list for further discussion.

Suggestion from Bob Faulk regarding address search object: No one expressed interest in pursuing this proposal, and it was suggested that it was more appropriate as a vendor extension.

IEEE 802.3 Hub Management Letter

```
To:     Donna McMaster
        Keith McCloghrie
        Repeater Management MIB Working Group
        IETF

From:   Kathy de Graaf
        Steve Horowitz
        Jim Reinstedler

For over two years we, as members of the IEEE 802.3 Repeater
Management Task Force, have worked very hard to  develop a standard for
managing IEEE 802.3 repeaters.  802.3 has approved the current draft in a
letter ballot, and on November 14, 1991 affirmed this work by voting
overwhelmingly to send the current draft to a confirmation ballot.

The members of the 802.3, representing almost all the major hub vendors,
have considerable experience not only in instrumenting but also in
configuring manageable hubs.  Although much of this draft is directed toward
instrumentation for fault and performance management,  considerable effort
was also expended to model the real repeater products that exist in the
marketplace.

A repeater is frequently implemented as one or more cards in a  modular
hub having  multiple backplane connections and with a single agent
managing the hub.  These hubs may contain multiple repeaters and have the
```

ability to dynamically create and delete groups of ports or individual ports. While not all products have all these features, we did reach a consensus on features in the repeater MIB that correctly and usefully model either a high-end or low-end repeater without unduly burdening the simpler repeaters.

Two years ago only a minority of the task force supported attributes that were primarily for configuration, but as we realized (from discussion and implementation) that it was both practical and desirable to provide such attributes, an overwhelming and persistent consensus developed in their favor.

One example that has recently been controversial in the IETF is the use of hubID (now repeaterID) to distinguish one of many repeaters within a hub enclosure. We have found that this provides a simple, inexpensive, standard, interoperable, and useful way of allowing a single agent to address multiple repeaters, and thus urge that it be retained.

We, as members of the IEEE 802.3 Repeater Management task force, therefore hope that the RM MIB Working Group will consider preserving not only the IEEE attributes directed towards fault and performance instrumentation, but also those provided for configuration management.

Attendees

Miriam Amos Nihart	miriam@decwet.zso.dec.com
Karl Auerbach	karl@eng.sun.com
Steve Bostock	steveb@novell.com
David Bridgham	dab@asylum.sf.ca.us
Jeffrey Case	case@cs.utk.edu
James Codespote	jpcodes@tycho.ncsc.mil
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Michael Erlinger	mike@lexcel.com
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Satish Joshi	sjoshi@synoptics.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Yoav Kluger	ykluger@fibhaifa.com
Cheryl Krupczak	cheryl@cc.gatech.edu
Ron Lau	rlau@synoptics.com

Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
David Minnich	dwm@fibercom.com
Lynn Monsanto	monsanto@sun.com
Anil Rijsinghani	anil@levers.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Iris Tal	437-3580@mcimail.com
June-Kang Yang	natadm!yang@uunet.uu.net
John Ziegler	ziegler@artel.com

3.3.6 Internet Accounting (acct)

Charter

Chair(s):

Cyndi Mills, cmills@bbn.com
Gregory Ruth, gruth@bbn.com

Mailing Lists:

General Discussion: accounting-wg@wugate.wustl.edu
To Subscribe: accounting-wg-request@wugate.wustl.edu
Archive:

Description of Working Group:

The Internet Accounting Working Group has the goal of producing standards for the generation of accounting data within the Internet that can be used to support a wide range of management and cost allocation policies. The introduction of a common set of tools and interpretations should ease the implementation of organizational policies for Internet components and make them more equitable in a multi-vendor environment.

In the following accounting model, this Working Group is primarily concerned with defining standards for the Meter function and recommending protocols for the Collector function. Individual accounting applications (billing applications) and organizational policies will not be addressed, although examples should be provided.

Meter <-> Collector <-> Application <-> Policy

First, examine a wide range of existing and hypothetical policies to understand what set of information is required to satisfy usage reporting requirements. Next, evaluate existing mechanisms to generate this information and define the specifications of each accounting parameter to be generated. Determine the requirements for local storage and how parameters may be aggregated. Recommend a data collection protocol and internal formats for processing by accounting applications.

This will result in an Internet Draft suitable for experimental verification and implementation.

In parallel with the definition of the draft standard, develop a suite of test scenarios to verify the model. Identify candidates for prototyping and implementation.

Goals and Milestones:

Done	Policy models examined.
Done	Internet Accounting Background Working Draft written.
Done	Collection Protocols Working Papers written.
Done	Internet Accounting Background final draft submitted as an informational document.
Done	Collection protocol working papers reviewed.
Done	Collection protocol recommendation.
Mar 1992	Architecture submission as Internet Draft.
Jul 1992	Architecture submission as RFC.
Done	Architecture working papers written.

Request For Comments:

RFC 1272 "Internet Accounting: Background"

CURRENT MEETING REPORT

Reported by Gregory Ruth/BBN

Minutes of the Internet Accounting Working Group (ACCT)

Internet Accounting Background

The Wednesday session reviewed the Internet Accounting Background document which had recently moved to the status of RFC (1272). The major changes to this document since the July IETF were in the areas of security requirements and counting strategy.

The security concerns for internet accounting were discussed and fundamental requirements were found to be data integrity and data confidentiality. It was recommended that, to the extent possible, SNMP security services should be used to satisfy these requirements.

The counting strategy discussion revolves around how packets (datagram fragments) should be counted: on entry to a network or upon successful delivery. Since there are good arguments for both methods (depending on the intended use of the accounting information), the capability for both should be included in an internet accounting system.

Working Group participants offered comments, criticisms and suggestions that will be incorporated into either a new version of the RFC or follow-on documents. Two new items were suggested: (1) it should be mentioned that, in addition to the uses already listed, internet accounting may also be used to monitor the correct operation of the network (i.e., it may reveal problems/anomalies); (2) among the values that an internet accounting system could report for a flow might be a binary value indicating whether a flow was active or not in the measured time period.

Internet Accounting Architecture

On Thursday the Working Group discussed the latest draft of the Internet Accounting Architecture. Although this document has existed for several months now and has undergone three or four extensive revisions, it still needs work, both in form and content. An intensive session was spent going over the document section by section and paragraph by paragraph to refine both form and content. In particular the Working Group worked on:

- Tightening up the statement of scope that this document will address.

- More carefully and clearly defining the Internet Accounting model (and its difference from the OSI accounting model) and the interactions of its components.

Numerous detailed (but important) changes were suggested and will be incorporated into the next version of the Architecture document. Among them:

- An explanation that we intend to develop a draft MIB and accounting control function definition, but not a complete protocol specification for accounting.

- A clear statment about which packet processing layer accounting is done at, namely the IP layer.
- The addition of a security section to the architecture document (and the role to be played by SNMP security services).
- Definition of "subscriber", "flow start time" and other loosely used terms.

The Working Group intends to conduct a dialog over the changes and a review of this document over the Internet in the next couple of months with a view to advancing it to the status of Internet Draft by the next IETF conference.

General

The Working Group has decided to combine the metering services document (formerly intended to be separate) with the architecture document and to announce our intention to produce a draft MIB document (separately) before the Working Group's effort is done.

It was agreed that it is time once again to check what progress, if any, the OSI effort on accounting is making.

Attendees

Larry Blunk	ljb@merit.edu
David Bolen	db31@nis.ans.net
Ronald Jacoby	rj@sgi.com
Kenneth Laube	laube@bbn.com
Mike Marcinkevicz	mdm@csu.net
Cyndi Mills	cmills@bbn.com
Rebecca Nitzan	nitzan@nersc.gov
Allan Rubens	acr@merit.edu
Gregory Ruth	gruth@bbn.com
Sharad Sanghi	sharad@ans.net
Timon Sloane	peernet!timon@uunet.uu.net
Sally Tarquinio	sally@gateway.mitre.org
Frank Welch	

3.3.7 OSI Internet Management (oim)

Charter

Chair(s):

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

Mailing Lists:

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

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.

Request For Comments:

RFC 1095 "Common Management Information Services and Protocol over TCP/IP CMOT"

RFC 1189 "The Common Management Information Services and Protocols for the Internet"

RFC 1214 "OSI Internet Management: Management Information Base"

3.3.8 Remote LAN Monitoring (rmonmib)

Charter

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

Mailing Lists:
General Discussion: rmonmib@lexcel.com
To Subscribe: rmonmib-request@lexcel.com
Archive:

Description of Working Group:

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

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

Goals and Milestones:

Done Mailing list discussion of Charter and collection of concerns.

Done Discussion and final approval of Charter; discussion and agreement on models and terminology. Make writing assignments.

Done Discussion of the first draft document. Begin work on additional drafts if needed.

Mar 1991 Review latest draft of the first document and if OK give to IESG for publication as an RFC.

Internet Drafts:

"SNMP Trap Definitions For Remote Network Monitoring", 08/22/1991, Steven Waldbusser <draft-ietf-rmon-trap-00.txt>

Request For Comments:

RFC 1271 "Remote Network Monitoring Management Information Base"

CURRENT MEETING REPORT

Reported by Michael Erlinger/Lexcel

Minutes of the Remote LAN Monitoring Working Group (RMONMIB)

The Group congratulated itself on the acceptance of the RMON MIB as a Proposed Standard and its having been published as RFC 1271.

Inter-Operation Testing

The Group discussed the possible venues for testing of various RMON MIB implementations. There seemed to be at least four possibilities:

1. Internet - RMON MIB implementations could be made available via the Internet. Those wishing to make available a particular implementation could do so by announcing via the RMON mailing list the location of the RMON device. Those wishing to test that device could access it via the Internet. The discussion centered on the possible Internet load created by such devices. It was concluded that this load should be minimal as this is only a test environment, not a management environment.

2. IETF - It might be possible to create a RMON test environment at the next IETF. The Chair will look into the possibilities of using CERFnet or USD facilities for creation of such a test environment which would be open to all those wishing to test RMON tools.

3. RMON Meeting - Although token ring had not been discussed, it was suggested that if there are any token ring meetings outside of the IETF meeting, then an RMON testing environment could be staged at the same time. The Chair indicated that this would be considered in the scheduling of any such meetings.

Discovery

There had been a BOF the prior evening associated with device discovery. At the BOF there seemed to be a consensus that the RMON Working Group should investigate device discovery as a possible RMON MIB extension. Much discussion ensued as to the definition of discovery, current MIBs associated with discovery, and priority within RMON. It was decided that the Chair should get together with Fred Baker and come to a better understanding of what is being requested. In particular, detailed requirements need to be created.

Token Ring

It was decided that creating RMON token ring extensions should be the top priority for the Group. The current mailing list would continue to serve the RMON Group (no separate token ring mailing list would be created). It was decided that before January 1, 1992, the Chair would publish a proposed Charter and a proposed schedule which would include a meeting prior to the March IETF.

Other

Other RMON issues were discussed. In particular row creation. It was suggested that the row creation reference within the RMON specification be clarified by adding additional examples, (e.g., what happens when a row contains a read only value?).

Attendees

Robert Austein	sra@asylum.sf.ca.us
Steve Bostock	steveb@novell.com
Jeffrey Buffum	buffum@vos.stratus.com
Lida Carrier	lida@apple.com
James Codespote	jpcodes@tycho.ncsc.mil
James Davin	jrd@ptt.lcs.mit.edu
Michael Erlinger	mike@lexcel.com
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Darrell Furlong	
Shawn Gallagher	gallagher@quiver.enet.dec.com
Robin Iddon	robini@spider.co.uk
Ronald Jacoby	rj@sgi.com
Mark Kepke	mak@cnd.hp.com
Ron Lau	rlau@synoptics.com
Jim McQuaid	
David Perkins	dperkins@synoptics.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Timon Sloane	peernet!timon@uunet.uu.net
Bruce Taber	taber@interlan.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Mark Therieau	markt@python.eng.microcom.com
Maurice Turcotte	dnmrt@interlan.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jeremy Wilson	
June-Kang Yang	natadm!yang@uunet.uu.net

3.3.9 Simple Network Management Protocol (snmp)

Charter

Chair(s):
Marshall Rose, `mrose@dbc.mtview.ca.us`

Mailing Lists:
General Discussion: `snmp-wg@nisc.nyser.net`
To Subscribe: `snmp-wg-request@nisc.nyser.net`
Archive:

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.

Internet Drafts:

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

"Use of the Community String for SNMP Proxys", 10/05/1990, Richard Fox <draft-ietf-snmp-proxys-01.txt>

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

"Definitions of Managed Objects for the SIP Interface Type", 11/07/1990, Tracy Cox, Kaj Tesink <draft-ietf-snmp-smdssipmib-06.txt>

Request For Comments:

RFC 1155 "Structure and Identification of Management Information for TCP/IP-based Internets"

RFC 1156 "Management Information Base for Network Management of TCP/IP-based internets"

RFC 1157 "A Simple Network Management Protocol (SNMP)"

RFC 1158 "Management Information Base for Network Management of TCP/IP-based internets: MIB-II"

RFC 1161 "SNMP over OSI"

RFC 1162 "Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542) Management Information Base"

RFC 1212 "Concise MIB Definitions"

RFC 1213 "Management Information Base for Network Management of TCP/IP-based internets: MIB-II"

RFC 1215 "A Convention for Defining Traps for use with the SNMP"

RFC 1229 "Extensions to the Generic-Interface MIB"

RFC 1230 "IEEE 802.4 Token Bus MIB"

RFC 1231 "IEEE 802.5 Token Ring MIB"

RFC 1232 "Definitions of Managed Objects for the DS1 Interface Type"

RFC 1233 "Definitions of Managed Objects for the DS3 Interface Type"

RFC 1238 "CLNS MIB - for use with Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542)"

RFC 1283 "SNMP over OSI"

RFC 1284 "Definitions of Managed Objects for the Ethernet-like Interface Types"

CURRENT MEETING REPORT

Reported by Marshall Rose/DBC

Minutes of the Simple Network Management Protocol Working Group (SNMP)

1. **Resolution of the Ether-like MIB Process**

   The "Pax Davin" solution was reviewed. During the three-week comment period, no objections were raised on the mailing list. As a result, the ether-like MIB, as put forth by the IESG, will be published as a Proposed Standard. Further, a new Working Group, the EtherMIB Working Group has been chartered to carry out the remaining terms of the solution. This Working Group met jointly with the SNMP Working Group.

2. **Termination of the SNMP Working Group**

   The history of the SNMP Working Group was reviewed. As the Working Group has completed its Charter, it has now officially disbanded. Thank you one and all.

Attendees

Steve Alexander	stevea@i88.isc.com
Miriam Amos Nihart	miriam@decwet.zso.dec.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Philip Budne	phil@shiva.com
Jeffrey Buffum	buffum@vos.stratus.com
James Codespote	jpcodes@tycho.ncsc.mil
Tracy Cox	tacox@sabre.bellcore.com
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
William Jackson	jackson@manta.nosc.mil
Ole Jacobsen	ole@csli.stanford.edu
Ronald Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Scott Kaplan	scott@ftp.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com

Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
Ron Lau	rlau@synoptics.com
Evan McGinnis	bem@3com.com
Lynn Monsanto	monsanto@sun.com
David Perkins	dperkins@synoptics.com
Anil Rijsinghani	anil@levers.enet.dec.com
M. Richard Rose	
Marshall Rose	mrose@dbc.mtview.ca.us
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@interlan.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Frank Welch	
Jeremy Wilson	
June-Kang Yang	natadm!yang@uunet.uu.net
Wengyik Yeong	yeongw@psi.com
John Ziegler	ziegler@artel.com

3.3.10 X.25 Management Information Base (x25mib)

Charter

Chair(s):
Dean Throop, throop@dg-rtp.dg.com

Mailing Lists:
General Discussion: x25mib@dg-rtp.dg.com
To Subscribe: x25mib-request@dg-rtp.dg.com
Archive: dg-rtp.dg.com:x25mib/Current.Mail

Description of Working Group:

This Working Group will produce a set of three documents that describe the Management Information Base for X.25. The first document will specify the objects for the X.25 Link Layer. The second document will specify the objects for the X.25 Packet Layer. The third document will specify the objects for managing IP over X.25. The Working Group need not consider the Physical Layer because the "Definition of Managed Objects for RS-232-like Hardware Devices" already defines sufficient objects for the Physical Layer of a traditional X.25 stack. Any changes needed at the Physical Layer will be addressed as part of that activity.

The X.25 object definitions will be based on ISO documents 7776 and 8208 however nothing should preclude their use on other similar or interoperable protocols (i.e., implementations based on CCITT specifications).

The objects in the Link and Packet Layer documents, along with the RS-232-like document, should work together to define the objects necessary to manage a traditional X.25 stack. These objects will be independent of any client using the X.25 service. Both of these documents assume the interface table as defined in MIB-II contains entries for the Link and Packet Layer interfaces. Thus these documents will define tables of media specific objects which will have a one to one mapping with interfaces of ifType ddn-x25, rfc877-x25, or lapb. The objects for the IP to X.25 convergence functions will be defined analogously with the ipNetToMedia objects in MIB II.

The Working Group will endeavor to make each layer independent from other layers. The Link Layer will be independent of any Packet Layer protocol above it and should be capable of managing an ISO 7776 (or similar) Link Layer provider serving any client. Likewise the X.25 Packet Layer objects should be independent of the Link Layer below it and should be capable of managing an ISO 8208 (or similar) Packet Layer serving any client.

The Working Group will also produce a third document specifying the objects for managing IP traffic over X.25. These objects will reside in their own table but will be associated with the X.25 interfaces used by IP. These objects will not address policy decisions or other implementation specific operations associated with X.25 connection management decisions except as explicitly described in existing standards. These objects will manage the packet flow between IP and the X.25 Packet Layer specifically including observation of packet routing and diagnosis of error conditions. Progress on the Link and Packet Layer documents will not depend on progress of the IP over X.25 document. The IP over X.25 document will proceed on a time available basis after work on the Link and Packet Layer documents and as such the Link and Packet Layers may be completed before the IP over X.25 work.

All documents produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions (such as Concise MIB format). To the extent feasible, the object definitions will be consistent with other network management definitions. In particular ISO/IEC CD 10733 will be considered when defining the objects for the X.25 Packet Layer.

Goals and Milestones:

Done	Distribute first draft of documents and discuss via E-mail.
Done	Working Group meeting as part of IETF to review documents.
Sep 1991	Distribute updated documents for more E-mail discussion.
Nov 1991	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Jan 1992	Documents available with specified changes incorporated.

Internet Drafts:

"SNMP MIB extension for HDLC", 10/07/1991, Dean Throop, Fred Baker <draft-ietf-x25mib-hdlcmib-00.txt>

"SNMP MIB extension for IP over X.25", 10/07/1991, Dean Throop <draft-ietf-x25mib-ipox25mib-00.txt>

"SNMP MIB extension for the X.25 Packet Layer", 10/07/1991, Dean Throop <draft-ietf-x25mib-x25packet-00.txt>

CURRENT MEETING REPORT

Reported by Dean Throop/Data General

Minutes of the X.25 Management Information Base Working Group (X25MIB)

The X25mib Working Group met at the IETF meet in Santa Fe on Monday, November 18th, 1991. All draft documents were discussed and all were referred back to the editor for further changes.

Inter-MIB Structural Issues

Fred Baker raised the issue of a single X.25 packet layer running over multiple LAPB sessions. The current X.25 MIB will not support such a structure. Since inverse multiplexing occurs in other situations, it was agreed a general solution would be better than putting direct support for such structures into the X.25 MIB. Fred Baker volunteered to draft a MIB that would allow one level_X_MIB to be redirected to identify multiple level_X-1_MIBs.

To facilitate this, the SYNTAX of the x25InfoDataLinkId will be changed from INTEGER to OBJECT IDENTIFIER. That object identifier will identify an instance of the index for the first table of the MIB for the layer under that X.25. For X.25 running over LAPB, it will be an instance of lapbParamIndex. For X.25 running over multiple link layer entities, it will be an instance in the table of the MIB that Fred will draft (see above). For X.25 running over interfaces that don't have specific MIBs, it could also be the ifIndex for an interface.

A similar change will be made in how the HDLC MIB identifies the port below it. The hdlcParamPortIndex will be deleted and the SYNTAX of hdlcParamPortId will be changed to OBJECT IDENTIFIER. The object identifier will identify an instance of the index for the first table of the MIB for the port under LAPB. In general this will be an instance of rs232PortIndex.

Dave Perkins said he had a new tool which identified several syntax problems with the current drafts. It was agree the MIBs should be changed to correct these problems and the issue was referred to the editor to complete.

HDLC MIB

It was agreed to change the name of HDLC back to LAPB because the MIB wasn't broad enough to cover all variants of HDLC. It didn't cover basic HDLC framing, nor SDLC, nor LAPD. It is indeed a LAPB MIB and should be so named.

The LAPB MIB will be expanded to include support for ISO 8885 XID negotiations. Some of the attendees with European experience indicated that XID negotiations are important for that community.

Conformance

The issue of conformance was discussed. It was agreed that the MIBs will contain tables that are mandatory and optional. A vendor must implement all mandatory tables to claim conformance. The optional tables will be present to allow management of implementations that implement more than a minimal X.25/LAPB stack.

X.25 MIB

With the conformance issue in mind, the X.25 MIB will be examined to restructure the tables to make some tables optional. The objects required by a minimal X.25 implementation should be in required tables and all other objects should be in optional tables.

The X.25 MIB will also be examined to determine if some objects can be eliminated. Herve Goguely from LIR Corporation volunteered to review the current MIB in light of his European experience and to develop a list of objects to consider deleting.

The Group discussed recording error conditions from the last closed connection. It was agreed a table should be added to record the reason for the last abnormal close. The table should allow recording the last N conditions however vendors will only be required to keep 1 condition; vendors may choose to keep more if resources permit. The RMON MIB will be examined for a possible paradigm for structuring the table.

IP over X.25

Andrew Malis informed the Working Group that the IPLPDN Working Group has started writing a new RFC to replace RFC 877. He said there were several aspects of that draft that were inconsistent with the IP over X.25 MIB. The IP over X.25 MIB will be examined to align it with the revised RFC on IP over X.25 coming from the IPLPDN Working Group. Andrew Malis and Fred Baker will serve as liaisons between the IPLPDN and X25MIB Working Groups.

Attendees

Fred Baker	fbaker@emerald.acc.com
James Davin	jrd@ptt.lcs.mit.edu
Herve Goguely	lircorp1!roger@uunet.uu.net
Cheryl Krupczak	cheryl@cc.gatech.edu
Andrew Malis	malis@bbn.com
Evan McGinnis	bem@3com.com
David Perkins	dperkins@synoptics.com
Philippe Roger	lircorp1!roger@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com

3.4 OSI Integration Area

Director(s):

- Ross Callon: callon@bigfut.enet.dec.com

Area Summary reported by Ross Callon/DEC

Network OSI Operations

The Network OSI Operations Working Group, chaired by Sue Hares, met at the IETF meeting.

Bill Biagi gave an overview of the Corporation for Open Systems (COS). He described COS as an "International vehicle for accelerating the introduction of multi-vendor products conforming to OSI and ISDN". COS is involved in a number of OSI-related projects, including performance and conformance testing, and operation of OSINET.

There is thought about having OSINET connected to the Internet CLNP pilot.

Sue Hares gave a tutorial on IDRP (Inter-Domain Routing Protocol). IDRP is basically BGP with a number of enhancements, altered to support CLNP (with the possibility to also support IP). Advantages over BGP include: Confederations; Internal reliability (runs over CLNP or IP directly, rather than over TCP); Uses real authentication of routing packets based on MD4); Has a check for memory corruption (also based on MD4); Allows for route servers.

A number of issues related to CLNP deployment in the Internet were discussed. Problems relating to robustness of the current CLNP software primarily come down to maturity of software – it just takes time and effort to bash out the problems with software before we can get very reliable service (we often forget how long it took to get IP to be as reliable as it is). The needs for filtering mechanisms and management tools for CLNP were also discussed, along with the CLNP hookups to Interop. Sue announced that the NSFNET T3 network plans to use Integrated IS-IS based on the Wisconsin public domain implementation, and that IDRP will be used for OSI intra-domain routing.

X.400 Operations

The X.400 Operations Working Group also met several times during the IETF meeting, co-chaired by Rob Hagens and Alf Hansen.

The X.400 Operations Working Group is working on several issues related to X.400 deployment, including X.400 routing (routing between message transfer agents), Naming, and X.400–X.500 interoperation (this last topic is being pursued in cooperation with the OSI-Directory Services Working Group). The current status of the X.400 pilot was also discussed.

The X.400 Operations Working Group is also working on a document defining the minimum requirements for offering X.400 service in the Internet, as part of the ongoing internet X.400 pilot project. This paper is international in scope, and is being worked on jointly with the RARE X.400 Working Group.

Two other X.400-related documents have been approved by the Working Group. One is on compatibility and interoperation between 1988 and 1984 versions of X.400. The other defines gatewaying between X.400 1988/1984 messages and RFC 822 messages.

OSI Directory Services

The OSI-DS Group met at INTEROP on October 8, 1991. Full details are available in the published Minutes. The meeting was well attended.

The following Internet Drafts of the Group were recommended for publication as RFCs by the IESG, and all, bar one, have since been progressed.

- "An Interim Approach to Use of Network Addresses" (RFC 1277)
- "A String Encoding of Presentation Address" (RFC 1278)
- "Domains and X.500" (RFC 1279)
- "Using the OSI Directory to Achieve User Friendly Naming" (<draft-ietf-osids-friendlynaming-03.txt, ps> still under discussion)
- "Replication Requirement to provide an Internet Directory using X.500" (RFC 1275)
- "The COSINE and Internet X.500 Schema" (RFC 1274)
- "Replication and Distributed Operations Extensions to Provide an Internet Directory using X.500" (RFC 1276)

The document "Naming Guidelines for Directory Pilots" was discussed briefly, and pending minor edits is ready to be submitted as an RFC. The document "DSA Naming" was discussed. It is intended to attempt to progress this document prior to the next meeting, as it is important for the next stage of expansion. Pilot experiments on "Handling QOS (Quality of service) in the Directory" are ongoing, and recommendatios on this Internet Draft will be deferred until we have some practical experience.

The document "An Access Control Approach for Searching and Listing" was presented. It was agreed that this should be submitted privately by the authors as an informational RFC. The area was of interest, and this function should be considered later for Internet standardisation. If done, this should probably be based on the 92 access control.

It was agreed that following successful experiments (Russ Wright, Tim Howes, et. al.) that pictures in the directory should migrate for G3Fax from JPEG. Definitions would be added to the schema to allow for this.

A draft document "A Strategic Plan for Deploying an Internet Directory Service" was discussed. The next version of this document will be an Internet Draft. This Working Group will take on active review of the document. There were many comments, but broad concensus on the direction proposed.

A lengthy discussion on postal addresses was avoided by scheduling this item at 18:00.

Office Document Architecture

The Office Document Architecture (ODA) Working Group met in Santa Fe, chaired by Peter Kirstein.

The ODA Group is coordinating an ODA pilot project. The Group has been working on obtaining and documenting the availability and interoperability of ODA software from several sources (currently from BBN, Bull, and DEC, future software expected from ICL and other sources). The currently available ODA software supports most aspects of structured text, as well as bit-mapped graphics. They are currently in the process of distributing software to users. They also have an ODA testing capability at University College London.

Currently ODA can be exchanged using either X.400 or UU encoded SMTP/822 mail. There are plans for future support of ODA using the SMTP extensions.

3.4.1 Network OSI Operations (noop)

Charter

Chair(s):
Susan Hares, skh@merit.edu

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

Description of Working Group:

The Working Group is chartered to work on issues related to the deployment of CLNP in the Internet. Initial activities include both deployment planning and education of regional and other connected networks.

Initial planning efforts include the development of routing and management plans.

Goals and Milestones:

Aug 1991	Create tutorials for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.
Aug 1991	Collect OSI Routing and Addressing plans into a Repository. Make the plans available at Merit.edu:/pub/iso/noop/plan.
Ongoing	Provide a forum to discuss these OSI Routing plans by email or in Group discussions.
Nov 1991	Collect a list of OSI Network Utilities available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.
Nov 1991	Collect list of OSI Network Layer NOC tools and publish a list.
Nov 1991	Collect Methods of OSI Network Layer Debugging and write a document describing these methods.

CURRENT MEETING REPORT

Reported by Dan Blum/Rapport and April Merrill/NCSC

Minutes of the Network OSI Operations Working Group (NOOP)

The Network OSI Operations Group met on Tuesday night, and Wednesday morning and notes were taken by Dan Blum and April Merrill respectively. A third NOOP session was held on Wednesday evening and was listed as the NOOP OSI Network Tools BOF.

The Network OSI Operations Group convened at 1930 on November 19, 1991. The Group heard presentations from Richard Colella (NIST) and William Biagi (Corporation for Open Systems). After the presentations, Sue Hares (Chair) gave an IDRP tutorial, then described the OSI activities at INTEROP '91, after which some general discussions were held.

NIST

Richard Colella of NIST discussed NIST's OSI Routing Laboratory. The current focus of the laboratory is IS-IS routing software. The lab provides interoperability testing in an informal environment between participating vendors. There is as yet no formal conformance testing or OSINET-style publication of results. Participating vendors included DEC, Proteon, and 3COM.

The lab was open beginning in August 1991 to assist with INTEROP preparations, and is tentatively scheduled to re-open in January of 1992.

COS

William Biagi of COS provided background on his organization and discussed its plans for fostering OSI internetworking. Although COS was originally founded to promote OSI and ISDN, some members are beginning to look at the requirements for coexistence with TCP/IP. Most members are running proprietary networks that support neither OSI nor TCP/IP.

COS operates the OSINET interoperability testing organization as a non-profit corporation. OSINET is linked to ACCUNET and TYMNET via X.25. The following plans are afoot:

- Expand OSINET to incorporate an Internet connection with some of the regional networks.
- Establish X.400 mail relays, possibly supporting TP0, TP4, and RFC 1006.
- Establish CLNS systems in each COS member organizations by the end of 1992.
- Act as a repository for CCITT blue books and other documents.

- Possibly explore the X windows over OSI.
- Explore other Internet-like services.

IDRP Tutorial

Sue Hares discussed the IDRP design goals, protocol, and other features. She indicated that Mitre has an IDRP prototype implementation running today. Details of the presentations are available on the slides presented.

OSI At INTEROP '91

Sue discussed the OSI network set up for the November 1991 INTEROP demonstration.

The show featured worldwide connectivity (North America, Australia, Europe). Proteon, Wellfleet, Network Systems, cisco, 3com, DEC participated. IS to IS was used in the booth. It was actually "stress tested" by physical outages, to which it reacted well.

A great many practical lessons were learned. While network administrators and various regionals and various corporations were willing to install OSI routers, they did not enable CLNS on their production routers. While the multiprotocol routers claim to support both CLNP and IP, sometimes there are bugs when both are simultaneously active. A big stumbling block to installing CLNP on production routers in production networks is the lack of CLNP router security "filtering" capabilities.

Overall, however, the experience at INTEROP seems to indicate that critical mass in the product arena is at hand. What is urgently needed at this time is a user/application base to provide additional field testing of CLNP and associated protocols.

Wednesday Morning

The objective of the session was discussed during the earlier meeting. The notes do not reflect the actual topics. The following topics were discussed:

Introductions and OSI at your site

People attending the meeting gave their name and described their interest in OSI and what is happening at their site for OSI.

OSI INTEROP '91

3COM, DEC, Proteon, and Wellfleet tested IS-IS interoperability at INTEROP '91. 3COM, Proteon, and Wellfleet provided IS-IS support for the INTEROP '91 OSI Demonstration booth.

The IDRP prototype was developed for INTEROP '91 by Dave Katz of Merit. Merit expects work on prototype to continue. Mitre is also working on a prototype for IDRP.

Twenty-five network service providers and twenty-five OSI vendors participated to provide INTEROP '91 connectivity for four OSI applications: FTAM, VT, X.400, and X.500.

Changing the Charter

Sue Hares wants to change the Charter. The Group felt there was significant work to still be accomplished. Sue will re-write the Charter and send it to the Working Group. This Charter should include:

1. Survey form for OSI service.
2. RFC on OSI tools (late addition by Sue Hares Working Group Chair).
3. Tools work re-defined.
4. FYI RFC on OSI in the Internet - Frequently asked questions.
5. Pilot project activities.
6. National Test bed activities.
7. Routing Plans still written and reviewed.

Routing Plans

The Group had a lot of concerns about how to make CLNS software work. The routing software is not being well tested for CLNS and IP together. Problems show up immediately or after software has been running for a week or two weeks.

The Routing Plans are good tools. However, there are three types of routing plans:

1. Initial test routing plan.
2. Medium term routing plan.
3. Long-term routing plan.

Most people working for INTEROP '91 had initial test routing plans. Most regional networks are going back and adapting medium-term routing plans. This medium routing plan needs to try to look at the long range routing ideas, but it needs to try to work something out for now.

Sue Hares mentioned that any type of routing plan is okay to send to the list. She was concerned that the superb routing plans already sent to the NOOP people have stopped people from sending initial attempts at routing plans to the NOOP list. Any thoughts or initial versions of routing plans are welcomed on this list.

Additional regional routing plan discussion was tabled until after the NSFNET T1 and T3 OSI routing plan was described by Sue Hares.

NSFNET T1 OSI Routing Plan

Architecture of T1 NSS: 9 RTs on Dual Token Ring History of OSI in NSFNET T1 Network: By INTEROP '89 - prototype demonstration By August '90 - full deployment on T1 Network INTEROP '91 - used as part of OSI Internet demonstration

How OSI routing works:

MAP NSAP prefix -> IP address
then
MAP IP address -> AS
then
MAP AS -> next hop

When node reaches edge of NSFNET (external interface of remote NSS) then:

NSAP Prefix is mapped to Network Entity Title (NET) of router packets are to be sent to a regional network.

Both the NSAP prefix -> IP address and NSAP Prefix-> NET are static mappings. Each nss processor (rcp and psp) has the NSAP prefix ->IP address mapping. Only the E-PSP for which is attached to the NET has the mapping between the NSAP Prefix -> NET.

NSAP Addresses: NSFNET uses GOSIP format for NSAP address assignment.

NSFNET NSAP has GOSIP IDP (47 0005), AAI of FFFF00. The Notes have the full NSAP address.

OSI Routing Strategy: Transit all OSI traffic that conforms to acceptable use strategy.

Proposed T3 OSI Routing Architecture

Time for implementation: as soon as possible to accomplish the following things:

1. T3 network stability and transition from T1 network.
2. Change some of the T3 hardware to hardware that allows higher throughput.
3. Change to software base to switch software to gated.

Current plan is that the NSFNET T1 network will remain for OSI until the NSFNET T3 has OSI working.

Proposed Software Support:

OSI support:

- CLNP support
- ES-IS support
- IDRP support in gated

Dual IS-IS in gated

(further details in the notes from the talk)

DECNET Phase 4 -> Phase 5 Problems

DEC is shipping Phase 5 with VAX/vmx and Ultrix. Tony Hain recommended bringing up Ultrix as Phase 4. Don't try Phase 5 yet on Ultrix if you already have a Phase 4. The DECNET Program for address assignment has lots of power, but may assign addresses you do not want if the User does not understand the addressing questions. DEC is working on additional refinements.

ESNET, DEC and cisco are working on plans to solve transition problems in transitioning between large DECNET Phase 4 areas (like HEPNET) to Phase 5 areas.

The ESNET routing plan will be out in January and has some details on transition between Phase 4 and Phase 5.

Next step in CLNS routing Regional

OSI Infrastructure set-up for INTEROP '91 was for the OSI demonstration. The European WG-CLNS-4 has been running both IP and CLNP since 1990.

John Curran shared about NEARNET's network. The NEARNET client sites are close together and have a high probability of getting calls from DECNET Phase 5 sites. NEARNET would like advice. Not enough people in regionals have worked in this area. Sites can change from one concentrating router to another hub router in the NEARNET backbone. This switching is possible due to the small distance between sites. However, it complicates the OSI routing.

John wonders if there is a way to assign them a long-term NSAP prefix. One can assign NSAP prefixes for them to get somethinging running for DECNET Phase 5, but NEARNET is in the process of drawing up a plan for the NEARNET network for OSI. NEARNET expects to have a routing plan by December.

Trying NSAP allocation in the NEARNET network is not clear since their environment is changing. John would like to see a few OSI Networks operational before he decides how NEARNET will handle OSI.

John defined some problems (with discussion from the Group) we face in OSI::

- Customers have networks, but the OSI applications are not being used.
- Who will educate the people within the networks on OSI?
- We need to know which regionals are doing CLNS even if it is a partial CLNS service.
- NOOP should do a survey of regional networks.
- Where do we get payback on the investment in learning and technology for OSI service?

John Curran also defined four stages in getting CLNS from testing to production:

1. Stage 1 - Initial experimentation

   Spare routers and test machines are used to try out the CLNS network layer and OSI applications.

2. Stage 2 - CLNS trial

   Need a national test bed to try out CLNS code on different routers. These routers need to be tested under applications loads. The Testbed could be glued together with IP networks using encapsulation. **Note:** the Group decided to query the NOOP Working Group to see who is interested in getting a Testbed.

3. Stage 3 - CLNS and IP in production network

   CLNS needs to run a test service in production routers. However, the CLNS service is considered a "trial" service and may encounter some down time. **Editor's note:** This test service in prodution routers was taken by NSFNET in the T1 Network, and has proven very effective. The IP service is given priority in problem solving, but the CLNS is exercised)

4. Stage 4 - CLNS and IP production in network

   Both CLNS and IP are production services.

Users are reluctant to migrate to OSI due to the need to cut over applications or work with new OSI applications. It is important to get those OSI applications running in the Internet to run over CLNP. The X.500 and X.400 Working Groups in IETF should be encouraged to get their applications running over TP4 and CLNP as well as TCP/IP. Also, we hope to see applications like X-windows transition to OSI.

National Test Bed: The NOOP Group wants to organize a national test bed for CLNP and applications over CLNP.

Survey for OSI Service: The NOOP Group wants to have a survey of who will provide CLNP service. Sue Hares, Linda Winkler, and John Curran will put together a list of questions.

Security Concerns: The NOOP noticed that none of the routers can filter packets based on OSI addresses or OSI application information. Companies which use IP filter to provide some security for their company networks will not let OSI traffic in from the Internet due to the lack of security filters. Sometimes OSI packets can flow to the router at a company, but no further due to the lack of OSI filtering.

OSI Tools

A NOOP session will be held to discuss network tools. We need to start making recommendations on what OSI network layer tools need to be in routers.

NSAP Addresses

Applications addresses will be in:

- X.500
- A flat file by Sue
- Visual representation of OSI -> IP mapping for quick look-up in OSI debugging.
- Domain Name Service

Hitchiker's Guide to OSI in the Internet

Sue Hares is writing an FYI on OSI in the Internet. Sue Hares will circulate this to the NOOP Working Group.

Summary of Action Items:

Action item 1: Re-write Charter to update to current work: Sue Hares

Action item 2: Query NOOP and other lists to see who wants to start a working sub-group on a testbed for CLNP testing, and Phase 4 to Phase 5 transitions: Sue Hares

Action item 3: Put together a survey on OSI services: Sue Hares, Linda Winkler, and John Curran

Action item 4: Write up the security concerns: Walt Lazear

Action Item 5: Collect addresses and publish a list of file names for the flat files and their anonymous FTP location: Sue Hares

Action Item 6: Write and circulate for comment an FYI or Hitchhiker's Guide to OSI in the Internet: Sue Hares

These Minutes cover the Wednesday evening NOOP OSI Network Tools session.

The objective of the session was to begin work on identifying what tools are needed and available for assisting in the deployment and management of OSI protocols in the Internet. Five tools were discussed. Sue Hares and Cathy Wittbrodt (ESNET) will produce an Internet Draft expanding on the information below. The intent is to end up with an Internet Standard for ad-hoc OSI tools. The Internet Draft will specify the required and recommended tools. However, the listing of tool implementations will be in the NOC Tools Catalogue.

OSI Ping

Ping is critical and is already supported on most intermediate systems (ISs). All known implementations use the method adopted by ISO (i.e., the new PDU type) except for Sun, which implemented the selector-based ping. The ISO method is the one to be mandated by the Internet Draft in all ISs.

Most end systems (ESs) do not implement osi ping, but the Internet Draft will mandate it for all ESs. Ping needs to use some name-to-address mapping, such as an '/etc/isohosts' table. The Internet Draft should consider including features from IP ping other than the basic mechanism, such as 'fill'.

OSI Traceroute

Traceroute is also important to debugging network problems. Traceroute is not widely available on popular platforms. We need at least some simple functions first, then features like source route later. Traceroute should also use the '/etc/isohosts' file for name-to-address mapping.

Ping Monitor

This is a useful tool, but not as important as the other tools. It will be suggested, but not mandated.

Routing Table Dumper

Both ESs and ISs should have a capability to locally dump the routing tables (the moral equivalent of 'netstat -rn'). We should specify in the Internet Draft what information is useful to see. Note that the information should be consistent with the MIBs.

Transport Ping

In addition to verifying that the network entity is alive via ping, there was some feeling that it would be useful to have a transport ping as well. Currently, FTAM is used to verify layer 4+ connectivity, but this relies on getting all selectors right. Due to lack of familiarity with OSI and its terminology, selectors and other higher-layer configuration information is not always understood and correctly configured. A transport ping to a well-known transport selector might be useful. This tool needs more thought and will not be included in the Internet Draft.

Platforms of Interest

The following platforms are (non-exhaustive) lists of the ES and IS environments of interest.

ES: Banyan, BSD 4.x, CDC, DECNET Phase V (Ultrix & VMS), HP IBM RS6000, Retix, Sun Wollongong

IS: 3Com, cisco, DEC, routers, NSC, Proteon, Wellfleet

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
William Babson	bill%penril@uunet.UU.NET
William Barns	barns@gateway.mitre.org
William Biagi	bbiagi@cos.com
Daniel Blum	4108980@mcimail.com
Ross Callon	callon@bigfut.enet.dec.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Colella	colella@osi.ncsl.nist.gov
Rob Coltun	rcoltun@ni.umd.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
John Curran	jcurran@bbn.com
Tim Dixon	dixon@nikhef.nl
Tom Easterday	tom@cic.net
Robert Enger	enger@ans.net
Dino Farinacci	dino@cisco.com
Vince Fuller	vaf@stanford.edu
Darrell Furlong	
Jisoo Geiter	geiter@gateway.mitre.org
Jim Ghadbane	jimgh@newbridge.com
Alf Hansen	Alf.Hansen@delab.sintef.no
Susan Hares	skh@merit.edu
Christine Hemrick	hemrick@cisco.com
Kevin Jordan	kej@udev.cdc.com
Darren Kinley	kinley@crim.ca
Dale Land	land@lanl.gov
Walter Lazear	lazear@gateway.mitre.org
Peter Liebscher	plieb@sura.net
April Merrill	abmerri@tycho.ncsc.mil
Glenn Michel	gym@lanl.gov
Mark Needleman	mhn@stubbs.ucop.edu
David O'Leary	oleary@sura.net
Yakov Rekhter	yakov@watson.ibm.com
Ron Roberts	roberts@jessica.stanford.edu
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Ursula Sinkewicz	sinkewic@decvax.dec.com
Sally Tarquinio	sally@gateway.mitre.org
Kannan Varadhan	kannan@oar.net
Carol Ward	cward@spot.colorado.edu
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Jessica Yu	jyy@merit.edu

3.4.2 OSI Directory Services (osids)

Charter

Chair(s):
Steve Hardcastle-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`
Archive:

Description of Working Group:

The OSI-DS Group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this Group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

Done — Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.

Done — Study the relationship of the OSI Directory to the Domain Name Service.

Ongoing — Maintain a Schema for the OSI Directory on the Internet.

Ongoing — Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

Internet Drafts:

"Building an Internet Directory using X.500", 11/19/1990, S. Kille <draft-ietf-osix500-directories-01.txt, or .ps>

"Using the OSI Directory to Achieve User Friendly Naming", 11/26/1990, S. Kille <draft-ietf-osids-friendlynaming-03.txt, or .ps>

"Handling QOS (Quality of service) in the Directory", 03/20/1991, S.E. Kille <draft-ietf-osids-qos-01.txt, or .ps>

"Naming Guidelines for Directory Pilots", 03/21/1991, P. Barker, S.E. Hardcastle-Kille <draft-ietf-osids-dirpilots-04.txt, .ps>

"DSA Naming", 03/21/1991, S.E. Hardcastle-Kille <draft-ietf-osids-dsanaming-02.txt, or .ps>

"Schema for Information Resource Description in X.500", 06/14/1991, Chris Weider <draft-ietf-osids-resdescripx500-00.txt>

"Schema for NIC Profile Information in X.500", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-osids-nicprofilex500-00.txt>

"Interim Directory Tree Structure for Network Infrastructure Information", 06/14/1991, Chris Weider, Mark Knopper, Ruth Lang <draft-ietf-osids-treestructure-00.txt>

"Directory Requirements for COSINE and Internet Pilots (OSI-DS 18)", 07/09/1991, S.E. Hardcastle-Kille <draft-ietf-osids-requirements-00.txt, .ps>

"Generic Security Service Application Program Interface Overview and C bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-00.txt>

"An Access Control Approach for Searching and Listing", 09/23/1991, S.E. Hardcastle-Kille, T. Howes <draft-ietf-osids-accesscntrl-00.txt, .ps>

"Representing Public Archives in the Directory", 12/04/1991, Wengyik Yeong <draft-ietf-osids-archdirectory-00.txt>

"A String Representation of Distinguished Names", 01/30/1992, S. E. Hardcastle-Kille <draft-ietf-osids-distnames-00.txt, .ps>

Request For Comments:

RFC 1275 "Replication and Distributed Operations Extensions to Provide an Internet Directory using X.500"

RFC 1276 "Replication and Distributed Operations Extensions to Provide an Internet Directory"

RFC 1277 "Encoding Network Addresses to Support Operation Over Non-OSI Lower Layers"

RFC 1278 "A String Encoding of Presentation Address"

RFC 1279 "X.500 and Domains"

INTERIM MEETING REPORT

Reported by Steve Hardcastle-Kille/UCL

Minutes of the OSI Directory Services Working Group (OSIDS)
October 8, 1991

Previous Minutes

The Minutes from the July meeting in Atlanta were accepted without change.

Matters Arising

Updates were given on some of the action items from the previous meeting.

- George Brett was to find out how to get the CNI documents and send this information to the osi-ds list. Steve Hardcastle-Kille will prompt George.

- Mark Knopper & Chris Weider were writing a paper on storing NSAP information, and were going to update various other documents. Mark has moved on to other duties, and doesn't have much time available. Chris is working on them when he can. They're happening, but slowly, due to the work necessary for the FOX project.

- The status of the JPEG and QOS experiments will be given later in the meeting.

- Steve Hardcastle-Kille & Paul Barker have a draft of the Naming Guidelines document ready for discussion during this meeting.

- Steve Hardcastle-Kille has produced a first draft of a strategy/overview document.

Liaisons

1. RARE Working Group 3 - (Steve Hardcastle-Kille)

   The size of the Working Group 3 tutorials has grown. The last one was around 70 people, whereas previous ones were about 20-30. Earlier ones were strictly technical but the latest one wasn't.

   There are now high-level liaisons between RARE and the IAB and CCIRN.

   Work in the Research Community is US-Centric and is being RFC-driven. A European equivalent of RFC is being set up. (CFC?)

   Further technical items will be presented at relevant points during the remainder of the meeting.

2. OIW - (Russ Wright & Richard Colella & Ken Rossen)

   The OIW retracted their earlier "postal attribute" decision and spent a fair bit of time on discussion of it.

   Work is progressing on implementation agreements on Access Control, Replication and Distributed Operations. It will be aligned with the output of the Berlin meeting. The Distributed Operations work is based on 88, not on 92 extensions.

   There are no conformant subclasses for 92 Access Control, so it will be all or nothing for implementations. The OIW work is concentrating on error handling.

   92 replication work is on a shadow protocol. There will be some leeway in conformance. OIW is working on that, as well as error situations.

   There are some OIW documents available online via FTP & FTAM. Richard Colella will send an index of the documents to the osi-ds list.

3. ISO/CCITT - (Ken Rossen)

   There is a hope to include X.500 in GOSIP 3.

   92 Extensions Work

   - There are 13 documents: 12 PDAM (proposed document amendments) and 1 new part on replication.
   - The replication work is on the second draft and looks stable.
   - Access control is on the third draft, but there is still a lot to discuss. The U.S. has a need for minimal access control (restrictive access control).
   - The Berlin output should be fairly stable.

   CCITT PICS PRO FORMA (?)

   - This will be 1988 based. It is (or will contain) a checklist of conformance issues, both required and optional. Youbong is the editor; we should see if she can make it available.

   The protocol version was not advanced in 1992.

   Distributed entry work has been dropped for now. It will not be in 1992.

   There was some detail given on a defect in 1988 X.500

   - Multicasting / parallel chaining with > 1 DSA proceeding

- Problem with NSSRs (Non Specific Subordinate References)

4. NADF - (Einar Stefferud)

   Stef gave a brief overview of the NADF. It is comprised of public directory service providers (or future providers). The members may not actually offer directory services. One large issue they're struggling with is dealing with different information for the same entity, and having the different pieces managed by different service providers.

   A pilot will be getting underway soon, maybe in January. Participation will be limited to NADF members.

   NADF documents are openly available. There are currently 218 documents.

   ANSI is now registering alphanumeric names for private organizations. There are four tables under the C=US arc:

   - FIPS 5 - regions (states/provinces)
   - Organizations and private institutions created by Congress. (this is a point of NADF tension as NADF will honor these, but ANSI may not register them, or even keep a list of them)
   - ANSI names
   - ANSI national standards

   There has been some confusion regarding registration versus listing in the directory. It turns out to be a somewhat emotional issue. Some details are included in a later section.

5. FOX - (Ruth Lang)

   They are making technical progress. Merit is working on a draft document on how to store NSAPs in the directory. SRI is working on a User Agent to access the WHOIS information.

   Not a lot of progress in the "futures thought" area since Atlanta.

   Steve Hotz has gone back to school, but may continue doing the US DSA reports. (Since the meeting, Tom Tignor (tpt2@isi.edu) has assumed responsibility for the Directory Services Activities Report).

6. PSI WPP - (Wengyik Yeong)

   There are currently 74 organizations, with 165K entries. The composition is: 34 Universities (105K), 19 Government/Non-Profit (25K) and 21 Commercial (40K). More applications are needed, but drawing people/applications requires better reliability. Reliability is a problem. It's more of a manager problem than a software problem. Security and access control are also needed to draw more participants.

   They will be converting to the new US naming scheme (RFC 1255). This involves a DIT conversion.

7. Paradise - (Steve Hardcastle-Kille)

   Steve had some more Paradise handouts, although not many. There is a new glossy brochure coming from the helpdesk, probably in November. (how to get copies of these?) It's focus is global, rather than European.

   Steve outlined some of the Paradise services:

   - DSA service. This provides national and top-level DSAs, and replicas of national pilots. It also handles relaying between various networks (TCP/IP, IXI, PSS)

   - DUA service. DE is now offered. (see *draft* help card) The help card will be online ... and comments are welcome, please! DE will be available with the latest QUIPU patch.

   - Support for small to medium organizations, that would rather not run their own DSA. A simple interface is provided. These are typically for small numbers of entries.

8. AARNet - (Mark Prior)

   The AARN network has been up for 18 months now. The directory will have been up for twelve months by the end of December. There are currently four DSAs being run by the AARNet project (out of 21 total in Australia), one for each of four educational sites. They're working on naming guidelines. Problems encountered are availability (due to Internet link problems?), and size. They noted:

   (a) We were concerned about the use of non standard attributes, and especially the creation of new syntaxes due to the non-extensibility of some commercial directories.

(b) During the life of the project the AU DIT has grown from 10k to 30k. This growth was not just from the Project members sites so we are hopeful that this level of growth will continue through '92.

(c) We are currently shadowing the country level information for all countries (except FR and PT) on a local DSA and are using preferDSA to improve lookup speed. Previously we had tried to just have the information locally but without being an official slave DSA but this meant that all DSAs needed this information and this caused serious image size problems.

(d) We complained about Belgium's inaccessibility, and problems in contacting Giant Tortoise across the Internet.

(e) Standards Australia has issued a set of naming guidelines (SAA MP59-1991) and we will be following them, as much as possible. We will also start using an official OID, and will be applying for a NSAP and PRMD for AARNet.

(f) We held a demonstration of the Directory at the QUESTNet Winter Workshop, during July. Where we used color photo's (encoded in GIF and using a private image attribute). This was very successful, and popular, although not practical in long haul links when using GIF, as the image size was of the order of 50-100k. We intend to hold a similar demonstration at the Australian Networkshop during December, but this time use JPEG. We intend to start migrating our photo attributes to JPEG ASAP.

(g) We have strong views on the need for a preferredName attribute, and will deploy one, as we see it as necessary for telephony, if OSI-DS does not define it.

(h) We still see regular updating of the Directory as a serious problem and work needs to be done by someone to provide more tools in this area.

Document Progression to RFC

The IESG has recommended to the IAB that the bulk of the OSI-DS papers should move to RFC status (some information and some standards track).

Strategy Document

This is a new document. It is not on the RFC track right now. It's a controversial document, and has need for broader input than just OSI-DS. The first draft is just out, with a lot of it being Steve's view. It will be worked on by various members of IESG, IAB and OSI-DS.

Some discussion highlights:

- Support for a standard API is needed. Also support for XDS/XOM interface for directory services. (this will get more applications using it)

- One question heard is "will it be small and fast and compare with the DNS?"
- We should learn from the NADF experience in name assignment, although this will be something done on a country by country basis. We should be aligning to future national X.500 services.
- This is a strategy document, which means non-specific, which means it's hard to add specifics. The lack of timeframes may affect perceived appropriateness of some items (for example, when/how to intercept replication)
- There was concern over some of the extensions mentioned. One issue was whether this was trying to nail down QUIPU (the answer is "no") We need to deal with the issue of standards (e.g., core X.500) versus functions (e.g., lightweight protocols). Core X.500 should probably be emphasized. Extensions MUST be documented, and not just by implementation. They should then be fed back into the standards process.
- Two areas to document really: how to deploy a directory service, and where we're going with it (operational aspects). Should this be one document or two?

In broad terms, the Working Group supports this document and will edit/review it. The emphasis should be to do this as a directory service, not as an OSI service.

Access Control for Searching and Listing

Steve and Tim summarized the paper. There was some discussion of implementation problems, as well as whether we want to make this a standard, or just to implement it and make it an informational RFC. The leaning was for the latter.

There was some thought that this had some overlap with the 92 work, so we should at least consider retrofitting to 92. We should also send this work to the standards stream as a work item.

The solution outline makes the problem difficult, but not impossible. Is this a sufficient solution?

DIXIE, DAS, and Lightweight Protocols - RFCs 1202 and 1249

These are two informational RFCs. DIXIE and DAS are different. DIXIE provides more of a DAP-type function, while DAS provides more user level (user interface) function.

Should we fix on a single implementation? To confuse things, many manufacturers have done similar protocols. Standardizing seems to be a good thing, but neither of the two specifications is good enough.

There was some discussion and clarification between lightweight APIs and lightweight protocols. The two issues are whether it's easy to use for an application, and whether it's easy to use with respect to the OSI stack.

There was a preference for ASN.1 over ASCII in the protocol. It was felt a list of requirements was needed before we do any subsetting of the protocol.

No volunteers willing to take on the writing chore. Weng and Tim will investigate further.

Presentation of New US Naming Scheme - Wengyik Yeonag

The new naming scheme is specified in NADF 175 / RFC 1255. One of the key points is that listing (where you expect to be found) is separate from standing (a function of where you're registered).

There was much discussion on whether the standing/listing separation is the right thing to do. No consensus was reached on whether it was "Free Market Rules" or "A Mess".

Naming Guidelines

Paul Barker's new text was accepted with one small change.

Steve reported that from Zurich and post-Zurich discussions, it was decided to drop all of O=Internet, O=Cosine, L=Europe and L=North America from the root. The problem with having these objects directly under the root is that there is no formal registration process and there will be problems with PTT connections.

However, the issue of whether O=Internet should be dropped or not came back up and was resolved that it should stay where it is for now (at the root).

The JPEG Experiment: Pictures in the Directory - Russ Wright

This was the result of an old action item. Russ talked about some of the experiments done, which were successful.

The decisions made:

- Use JPEG for photo attribute.
- Keep FAX around for a transition period only, and then deprecate.
- Use a new attribute type, rather than overload existing one. Russ will liaise with Schema group to do this.
- Timescale on transition: the next version of QUIPU will have both FAX and JPEG. The version after that will not have FAX. The transition period will be twelve months.

The QOS Experiment

Time Howes, Paul Barker and Geir Pedersen have been working on this. There has been some progress in code, but not much in deployment.

Mark Smith will continue putting this into the Mac DUA. Colin will send a note (after some experience) to the list on how to install the necessary pieces in the DSA, and then use it. Weng & Colin will clarify and rectify any inconsistencies between the QOS specification and code, where code mean QUIPU.

DSA Operations for Paradise: Managing the root of the DIT
Colin Robbins

Besides the services that Steve mentioned earlier, they also do a fair bit of probing. It turns out not to be as simple as once thought. Simply connecting is not enough. If you do authentication, then you get a loopback, which is better. There are also network islands out there, so you need to do multiple probes on a DSA. Failures can be due to a number of things: network, host, dsa, protocol (hopefully not!), or authentication.

They've started "passive probing" to cut down on the load. In this case, they only probe the DSAs they haven't connected to.

Any suggestions on how to improve things are *welcome*.

They also do some counting of the DIT. This isn't really needed for operations. There was some discussion on whether counting should be done by DSA, or by subtree and propagated up the tree. This is a non-standard mechanism, since there is nothing in X.500.

DSA Naming. Presentation and discussion.

Steve summarized and provided some clarification for the current Internet Draft on this. There were some scalability concerns. It's easy to add another tree, though there is some time to go before that's needed. Adding new trees may be political rather than technological.

There was a strong suggestion that more pictures and more examples be included to help clarify what's being described.

It was decided that some discussion on the list was required. People would go and "Think Hard" about this.

AOB

Andrew MacPherson's message regarding a new personal attribute. Paul and Steve will respond to the list as promised. There has been some discussion on surname and Scandinavian names on the list, but there were no volunteers to try to summarize it.

- Postal addresses

  The 6x30 format is fairly well entrenched (for one, in heavy metal in mailing houses). Steve made a proposal that we add a new Oversized-Postal-Address, which is 6x60, and have Postal-Address as a subtype of it. This was perceived by some not to be a solution. We ran out of time without a resolution.

Next Meeting

The next meeting will be at the IETF in San Diego (March 16-20, 1992).

Summary of Action Items

- Steve Harcastle-Kille: Prompt George Brett to send information on accessing CNI documents to the osi-ds list.
- Richard Colella: Send index of available OIW documents to the osi-ds list.
- Wengyik Yeong & Tim Howes: Continue investigating DIXIE and DAS.
- Russ Wright: Liaise with Schema group to instantiate new attribute type for jpeg photo.
- Tim Howes: continue fitting the QOS stuff into the Mac DUA.
- Colin Robbins: send a note to the osi-ds list on how to install the necessary QOS pieces into the DSA. (and then how to use it)
- Wengyik Yeong & Colin Robbins: clarify & rectify inconsistencies between the QOS draft and the QUIPU code.
- Everyone: discuss strategy and DSA naming documents via email.

Attendees

Alyson Abramowitz	ala@hpindda.hp.com
David Brent	brent@CDNnet.ca
Stuart Cain	scain@hpindeg.cup.hp.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Colella	colella@osi.ncsl.nist.gov
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Tim Howes	Tim.Howes@umich.edu.
Burton Kaliski	burt@rsa.com
Paul Koski	koski@hpindeg.cup.hp.com
Ruth Lang	rlang@nisc.sri.com
Sylvain Langlois	sylvain@cli53an.edf.fr
Andy Linton	andy.linton@comp.vuw.ac.nz
John Mann	
Daniel Molinelli	moline@trw.com

Julian Onions	J.Onions@cs.nott.ac.uk
Colin Robbins	c.robbins@xtel.co.uk
Kenneth Rossen	kenr@i88.isc.com
Einar Stefferud	stef@ics.uci.edu
Charles Wolverton	ctw@aero.org
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov
Wengyik Yeong	yeongw@psi.com

3.4.3 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
Archive: janeb.cs.wisc.edu:/pub/archives/ietf-osi

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:

Done	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.

Request For Comments:

RFC 1139	"Echo function for ISO 8473"

3.4.4 Office Document Architecture (oda)

Charter

Chair(s):

Peter Kirstein, kirstein@cs.ucl.ac.uk

Mailing Lists:

General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk
Archive:

Description of Working Group:

The ODA Working Group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g., Spreadsheets. The Working Group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liaison with the SMTP and X.400 Working Groups.

This Working Group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Done	Inaugural meeting.
Done	Produce a paper stating what ODA standards or profiles still need completing.
Jul 1991	Produce paper on how both SMTP and X.400 message systems should be supported.
Done	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Dec 1991	Provide first feedback on the ODA Pilot.
Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.

CURRENT MEETING REPORT

Reported by Peter Kirstein/UCL

Minutes of the Office Document Architecture Working Group (ODA)

Current Status of Implementations:

As a background for the discussions on Pilots, the current status of implementations was reviewed. A document on the current status had been circulated prior to the meeting; is is called "The ODA Document Cnvertors" [ODC]. This document reviewed the status of three of the available packages:

Provider	Package
BBN/UCL	SLATE/ODA
Bull	WORD for WINDOWS/ODA
DEC	DECWRITE/ODA

Each of the three are now available for immediate use, and the suppliers have agreed to make three copies of each available to individual organizations participating in the IETF-ODA Pilots. The status of each is discussed below:

BBN/UCL-SLATE/ODA

There has been a Release of v1.2 of the BBN SLATE/UCL ODA software; it converts between SLATE v1.2 and ODA/ODIF Q112. The software is made available currently on SPARCstations, but is believed to be easily portable to IBM RISC 6000 machines and DEC Ultrix workstations. There is documentation for the system on the normal ietf-osi-oda infoserver. At present the system operates with the UCL PP message (v5.x or later versions), and thus can operate over SMTP (with UUencode) or X.400; later versions will work with the extensions to SMTP proposed in the 822EXT Working Group. It is possible to interoperate with any other SMTP mail systems which do UUencoding.

An agreement has been reached with BBN, that they will provide for the IETF Pilot 250 copies of SLATE v1.2, and will maintain it with later releases. It is restricted to "academic and research institutes only"; others must purchase SLATE from BBN. The software will be updated as later releases of SLATE become available. The whole documentation will be provided by UCL - who will include the BBN SLATE documentation. The BBN portion of the software will be provided to US participants by a "Shrink-wrapped Licence"; non-US sites will have to sign a BBN license supplied by UCL. In both cases, UCL will keep a register of copies supplied, and must furnish that to BBN. UCL will exact a small handling charge for the distribution. Details of the license agreements are given in [ODC].

Bull-Word FOR WINDOWS/ODA

This software is also available to the Working Group; it will run on a DOS PC, and must be integrated by the using site with a mail system. The Bull software is designed for conversion between RTF and ODA Q112, but they use it only with WORD for WINDOWS.

The software requires a PC/AT with PC-DOS or MS-DOS v 3.10 or above with at least 1 MB of EMS, hard disc and floppy. It requires also MS-WINDOWS v3.0, WORD for Windows v1 or other editor supporting RTF, and font scaling for a CRT such as ATM Adobe. The programs include Q112 <-> RTF convertors, a formatter, filter, and a browser. It also includes filters and test documents. There is appropriate documentation from Bull.

The license agreement specifies the use of up to three copies of the software on DOS systems; the usage of the software is restricted to R&D purposes. The licensees should provide a report on the usage. The intention is to provide the software to Universities and Public Research laboratories for evaluation, research and demonstration. The period of the agreement is initially until June 1992.

Bull will distribute the software and documentation, within a month of users returning a signed license agreement to Bull. Details of the license agreement are given in [ODC].

DEC

This package is regarded as a Gateway product between their CDA products and OSI. The VMS release was made in April, the ULTRIX release is on Extended Field Trial (EFT). They run on all current DEC machines. Again details of the license agreement are given in [ODC].

[ODC] describes also the limitations in current interworking between the three implementations.

We expect that there will also be available a version of WORDPERFECT/ODA from UPC/ICL. This software has not yet been tested fully for compatibility, and its license arrangements are still under discussion. We would expect to provide further information on this software at the next IETF meeting.

Interest was also keen in MAC software, and the Chair agreed to contact Apple since it was believed they had software in some relevant state. It was agreed that in view of the imminence of so much of the software, it was important to update this list regularly. The Chair would provide an updated status for the next meeting at the next IETF.

The participants in the meeting expressed an interest in having a reasonably up-to-date directory of who is using ODA for the Pilot. The Chair agreed to put up a list of mailboxes on an X.500 Directory system to which all those interested had access. He will also keep a list of active users on the Infoserver.

The Pilot

The document describing the implementations available was sent to the Working Group mailing list only just before the meeting. A number of people have requested one or other of the implementations available, so that there should be feedback from early use by the time of the next IETF meeting. Currently we expect to use UUencoded SMTP or X.400 for document transfer. However, in view of the excellent progress being made in the 822EXT Working Group, and the alpha release availability of a package from Marshall Rose, we expect also to be testing shortly with that package.

Next Meeting

It was proposed to hold the next meeting in San Diego during the week of March 16-20, 1992.

Attendees

Harald Alvestrand	herald.alvestrand@delab.sintef.no
Cyrus Chow	cchow@ames.arc.nasa.gov
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Ursula Sinkewicz	sinkewic@decvax.dec.com
Andrew Veitch	aveitch@bbn.com

3.4.5 X.400 Operations (x400ops)

Charter

Chair(s):

Alf Hansen, Alf.Hansen@delab.sintef.no

Mailing Lists:

General Discussion: ietf-osi-x400ops@pilot.cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@pilot.cs.wisc.edu
Archive:

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this Group is to insure interoperability between Internet X.400 management domains and the existing Internet mail service. The specific task of this Group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

Done	Initial meeting, produce internal outline.
Done	Working draft, circulate to interested people.
Jul 1991	Internet Draft available.
Dec 1991	Document ready for publication.

CURRENT MEETING REPORT

Reported by Kevin Jordan/CDC

Minutes of the X.400 Operations Working Group (X400OPS)

Welcome

The meeting was chaired by Alf Hansen, and Kevin Jordan volunteered as secretary.

There were no additional comments against the Atlanta meeting Minutes.

Action list from Atlanta meeting

- Rob Hagens and Alf Hansen were to revise draft RFC and distribute to the Working Group. **Done**

  NOTE: At the Atlanta meeting, we discussed the need for a separate document which would describe the strategy for X.400 Operations in the international X.400 internet. In Santa Fe, we decided that this document is not needed.

- Kevin Jordan was to update white paper on use of X.500 for support of X.400 routing and address mapping and distribute to the Working Group. **Done**

- Claudio Allocchio and Urs Eppenberger were to write a white paper on use of DNS for support of X.400 routing and address mapping. **Not Done.**

  They wrote software instead! The software will be made available to the RARE/COSINE and XNREN communities.

- Steve Hardcastle-Kille was to update 88->84 downgrading draft RFC and work with EWOS to make support of DD.COMMON well defined and mandatory. **Draft RFC Updated**

- Peter Yee was to do some research into North American groups such as EMA and NADF and make recommendations for liaison with these groups. **Yee was unable to attend the Santa Fe meeting. Peter plans to email his findings to the Working Group.**

IETF X.400 Operations Working Group Business

It was decided that the following changes should be made to the Charter:

- The Charter should be updated to include references to other documents in progress, e.g., the Routing and Mapping documents.

- The Charter should reflect that our work on X.400 operations and deployment will not be complete by December 1992.

- The Charter will probably be upated occasionally as X.400 operational requirements evolve and as real experience in X.400 operations becomes broader.

Relations to other groups. Significant changes were made to the draft RFC as a result of comments made against it at the RARE WG1 meeting which took place shortly before the Santa Fe meeting. While most of these changes were technically justified, and the authors were given authorization to make such changes at the Atlanta meeting, it was strongly recommended that this sort of change not be undertaken in the future without the review and consensus of the IETF Working Group. The RFC is supposed to be the product of the IETF Working Group. The IETF Working Group respects and welcomes contributions from RARE WG1, but North American members of IETF are not eligible to be members of RARE WG1, so they are unable to express their views through votes at RARE WG1 meetings. Therefore, significant changes to the draft should not be made without review and approval of the IETF Working Group membership.

X.400 Service Milestones

Each member of the Working Group presented highlights and milestones of X.400 service provided at his/her home site.

XNREN Project. More and more sites are joining the XNREN Project. However, X.400 traffic continues to be relatively light. Very little progress has been made on establishing connections to public ADMD service providers. The University of Wisconsin has established an experimental and publicly available X.400-based fax service. The fax service imposes some constraints and limitations. Contact Rob Hagens and/or Allan Cargille for details.

Norway. The Norwegian R&D X.400 network currently serves over 5000 active users. The principal Norwegian WEP carries between 20,000 and 40,000 X.400 messages per month.

COS. The Corporation for Open Systems has installed PP and SunLink/MHS internally. COS is planning to connect its X.400 service to the Internet and wants to use OSI CLNS in addition to RFC1006.

Navy. The U.S. Navy is aggressively pushing X.400 internally. It is deploying various types of X.400 gateways. Transport/network services provided include X.25 and CLNS.

Merit. Merit drove the OSI infrastructure demonstration at INTEROP '91, and managed to use CLNS to interconnect virtually every regional network of the U.S. Internet successfully. Sites in Europe (especially Finland) were also interconnected using CLNS. X.400 mail was successfully exchanged between a variety of sites over Internet using CLNS. Merit also provides a gateway between NSFNet and SprintMail.

ESNet. ESNet continues to implement and deploy X.400 internally. ESNet plans to make X.400 mail a production-oriented service by January 1, 1992.

CDNNet. X.400 traffic levels continue to grow. The primary CDNNet MTA currently exchanges between 10,000 and 15,000 X.400 messages per day. CDNNet is subscribed as a PRMD to ADMD Telecom Canada. CDNNet is seeking approval to become an ADMD itself. CDNNet maintains the EAN X.400 mail software and has recently developed an X Window System based X.400 user agent.

Slovenia. The X.400 R&D network in Slovenia currently serves over 2000 active users.

GARR. X.400 traffic continues to increase. GARR is connected to the public X.400 networks in Italy. GARR provides a centralized gateway service to a variety of other email networks including HEPNet, SPAN, EARN, and Internet. GARR supports multiple protocol stacks including X.25, RFC1006, DDCMP, and CLNS.

NORDUNet. NORDUNet has initiated a project to improve the reliability of the email services in the Nordic countries. Alf has been appointed as the official NORDUNet Mail Inspector.

Review of "Requirements for X.400 Management Domains (MDs) Operating in the Global R&D X.400 Service"

Revisions to the draft RFC will include the following:

- Title change to "Operational Requirements for X.400 Management Domains", and
- References to "Global R&D X.400 Service" will be changed to "International X.400 Service".

The References

Urs will distribute a new revision of his Routing Coordination paper. The new revision will reflect comments made at the recent RARE WG1 meeting.

Harald Alvestrand will polish his "Routing Policy" draft and distribute it to the Working Group. It was agreed that this paper should become one of the RFC's in the X.400 set. It will be referenced by the base RFC.

Use of an X.500 Infrastructure for Routing Purposes

Keving Jordan's X.500 white paper was generally well accepted. However, the following recommendations were made against it:

- As an optimization to the route determination algorithm, take advantage of the fact that a failed directory read operation will return a distinguished name prefix in the case that part of a distinguished name is matched. This can be used to locate the longest match of an O/R name in one read, and a second read can then be used to obtain desired attributes.

- Update the document to allow for PRMD's explicitly under ADMD's and propose that the X.400 tree be rooted under a new object occuring under country (rather than rooting the X.400 tree directly under country).

Status and necessary actions for implementation of experiments with the draft RFC for use of the DNS system for address mapping purposes.

Claudio Allochio has implemented a scheme for using existing PTR resource records to store address mapping information. He has also implemented a scheme for using MX resource records to store X.400 routing information.

Tools have been implemented for extracting PTR and MX records and producing RARE tables from them.

The Italian PARADISE Project is also implementing Kevin Jordan's recommendations for using X.500 to support X.400 routing and address mapping.

Summary of conclusions and actions

P. Yee	Peter will distribute his recommendations for liaisons with other groups.
R. Hagens, A. Hansen	The editors will review section 3.1, rewrite it, and distribute it to the Working Group for review and comment.
	The RFC authors will revise the document in accordance with the comments and conclusions generated at this meeting. A new draft will be distributed prior to the next IETF meeting, no later than January 15.
J. Geiter	Jishoo will write a recommendation for the construction of X.400 names based upon relevant RFC's and Implementor's Agreements.
A. Hansen	Alf will formally propose to RARE WG1 that mapping coordination procedures be published as RFC's.
All	The issue of ADMD=" " versus ADMD=0 will be discussed via email after the text about this issue from the recent RARE WG1 meeting is distributed.

K. Jordan	Kevin will rewrite his paper on use of X.500 for support of X.400 as a pair of draft RFC's: one related to use of X.500 for X.400 routing purposes, and one related to use of X.500 for address mapping purposes. NOTE: This action should be reconsidered in light of Steve Hardcastle-Kille's comprehensive paper on the same subject. I propose that we adopt Steve's paper as the basis for further work in this area.
U. Eppenberger	Urs will update his paper on static routing and mapping procedures and present it as a draft RFC.

Other Business

Borka Jerman-Blazic and Harald Alvestrand each made presentations on national character set issues and suggested alternatives for solving this problem with respect to X.400. The Working Group made no conclusions but agreed that this issue needs further discussion at future meetings.

Future Meetings

The next general IETF meeting is scheduled for the week of March 16th in San Diego, California. The X.400 Operations Working Group will meet on March 17 and March 18.

Attendees

Claudio Allocchio	claudio.allocchio@elettra-ts.infn.it
Harald Alvestrand	herald.alvestrand@delab.sintef.no
William Biaggi	bbiagi@cos.com
Ken Carlberg	carlberg@cseic.saic.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Colella	colella@osi.ncsl.nist.gov
Curtis Cox	ccox@wnyose.nctsw.navy.mil
John Demco	demco@cs.ubc.ca
Tim Dixon	dixon@nikhef.nl
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@es.net
Robert Hagens	hagens@cs.wisc.edu
Alf Hansen	Alf.Hansen@delab.sintef.no
Susan Hares	skh@merit.edu
Christian Huitema	christian.huitema@sophia.inria.fr
Borka Jerman-Blazic	jerman-blazic@ijs.ac.mail.yu
Kevin Jordan	kej@udev.cdc.com

Scott Kaplan	scott@ftp.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Jack Liu	liu@koala.enet.dec.com
Linda Winkler	lwinkler@anl.gov
Robert Woodburn	woody@cseic.saic.com
Russ Wright	wright@lbl.gov

3.5 Operational Requirements Area

Director(s):

- Susan Estrada: estradas@cerf.net
- Phill Gross: pgross@nis.ans.net
- Bernhard Stockman: boss@sunet.se

Area Summary reported by Susan Estrada/CERFnet and Bernhard Stockman/NORDUnet

During this IETF five working groups met. There were three BOF's on operations related subjects. The Operational Requirements Area Directorate (ORAD) met together with FARNET which met in Santa Fe at the beginning of the IETF week.

User Connectivity Problems

The User Connectivity Working Group met twice this week, in the true tradition of operators being overcommitted to these things, and actually came up with some really good outputs. They decided how to do a NOC phone book, standardized network status reports and standardized total ticket hand-off, which is the mechanized procedure. There should be some implementations happening in the next six months, which will actually make our lives a lot simpler.

Those interested in getting on the mailing list, send a request to ucp-request@nic.near.net

Network Joint Management

Network Joint Management (NJM) met once this week. Following the FARNET theme of "Hardening the Mid-level Networks", the Group discussed fifty simple things one can do to help the Internet be hard. The operators were encouraged to subscribe to nwg@merit.edu, which is going to be the open discussion list for what's going on in the networking community.

Network Status Reports

Around thirty people attended this session. Network status reports were given from:

- ESnet (Tony Hain)
- NSI (Milo Medin)
- MILNET (Katherine Huber)
- EBONE-92 (Bernhard Stockman)

Phill Gross has been organizing the network status report sessions for some time. However, at this meeting, Phill turned the organization of this Group over to Gene Hastings. The choice of Gene as the new Chair was an indication of the similar subjects covered by the

Network Status Report (NETSTAT) sessions and the NJM Working Group, also chaired by Gene. It is hoped that both NJM and NETSTAT will benefit from this new close coordination.

Router Requirements Checklist

The idea behind a router requirements checklist is to take that router requirements document and turn it into something that may be used as guidance for purchasing router equipment. The Group decided that this was a useful thing to do. A strawman checklist will be constructed soon. To subscribe to the mailing list send a request to rcl@cerf.net. It's not clear that this work will be done within an IETF working group. The idea is rather to bash this out, and just get it issued as an informational RFC, without having to form a working group.

Quality of Service Measurements

This BOF only concerned quality of service measurements for wide area networks. Basically the idea here is that as regionals, or as networks, there is no need to find measurement criteria available. The base line is to find the right questions to ask and that is a good way to start. A working group will be formed and a mailing list set up for discussing this subject.

Benchmarking and Methodology

The Benchmarking folks met this week. They word-smithed the benchmarking document. They're going to have one more video meeting in January, and a draft document will be available by the next IETF.

The Operations Requirments Area Directorate

The Operations Requirements Area Directorate session was chaired by Susan Estrada, Phill Gross and Bernhard Stockman. Around fifty people attended. The meeting was a joint session between ORAD and FARNET people.

Presentation of the Intercontinental Engineering and Planning Group (IEPG). Geoff Huston, co-Chair of IEPG, gave an overview of the current IEPG work. The IEPG met in Santa Fe the week before IETF. Major topics of interest for the IEPG Group were:

- Interactions between network regions.
- Protocol infrastructure.
- Multi-lingual applications.
- Network minimal service levels.
- Global traffic flows.
- Information services.

There is the need to define operation tools to the vendors. For example there is a need to make the SNMP displays used today a little more meaningful and a lot more helpful to use in the long run. A working group will be initiated, probably at the next IETF, that will define recommendations for the operational folks to give to vendors, to help them design better interfaces.

Operational Statistics

Operational Statistics met during two session with around thirty participants chaired by Bernhard Stockman.

The main topic was a simplified version of earlier documents describing the gathering, storage and presentation of statistical data. The major time was spent on discussing the storage format and polling periods. Prior to this there had been a discussion on 5-15 minutes polling periods. It was concluded that one single polling period could not be recommended. The polling period has to be dependent on the type of polling being performed so the meeting defined a set of polling periods for different situations. The intention is to have the simplified version ready for Internet Draft during December 1991.

BGP Deployment and Applications

The BGP Deployment and Applications BOF had approximately thirty participants, and was chaired by Jessica Yu. The reason for this BOF to investigate the need and interest of forming an IETF working group around this concept. Topics that were treated:

- A review of today BGP implementation and usage.
- Presentation by cisco on current implementations and future plans.
- Discussion around the NSFnet T3 and T1 BGP implementations.
- A review of midlevel networks currently using BGP.
- The need for an IETF working group to facilitate an inter-operability test and to act as a forum for knowledge transfer.

CURRENT MEETING REPORT

Minutes of the Operational Requirements Area Directorate (ORAD)

Report not submitted. Refer to Area Report for a brief summary.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Robert Aiken	raiken@nsf.gov
Thomas Bajzek	twb+@andrew.cmu.edu
Charles Bazaar	bazaar@emulex.com
Jordan Becker	becker@nis.ans.net
Robert Blokzijl	K13@nikhef.nl
David Bolen	db31@nis.ans.net
Alison Brown	alison@oar.net
John Curran	jcurran@bbn.com
Kent England	kwe@bbn.com
Hans Eriksson	hans@sics.se
Michael Erlinger	mike@lexcel.com
Susan Estrada	estradas@sdsc.edu
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Farrell Gerbode	farrell@rice.edu
John Gong	jgong@us.oracle.com
Phillip Gross	pgross@nis.ans.net
Tony Hain	hain@es.net
Eugene Hastings	hastings@psc.edu
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
Eric Hood	ehood@nwnet.net
Fred Howlett	fhowlett@attmail.com
Geoff Huston	g.huston@aarnet.edu.au
Dale Johnson	dsj@merit.edu
James Jokl	jaj@virginia.edu
Dan Jordt	danj@nwnet.net
Daniel Long	long@nic.near.net
E. Paul Love	loveep@sdsc.edu
Donald Morris	morris@ucar.edu
David O'Leary	oleary@sura.net
Ari Ollikainen	ari@es.net
Joe Ragland	jrr@concert.net
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Sharad Sanghi	sharad@ans.net

Robert Shirey	shirey@mitre.org
Ursula Sinkewicz	sinkewic@decvax.dec.com
Frank Solensky	solensky@clearpoint.com
Bernhard Stockman	boss@sunet.se
George Strawn	gstrawn@nsf.gov
Roxanne Streeter	streeter@nsipo.nasa.gov
Claudio Topolcic	topolcic@nri.reston.va.us
Carol Ward	cward@spot.colorado.edu
Linda Winkler	lwinkler@anl.gov
Wengyik Yeong	yeongw@psi.com
William Yundt	gd.why@forsythe.stanford.edu

CURRENT MEETING REPORT

Reported by Pushpendra Mohta/CERFNet

Minutes of the Router Requirements Check List BOF (RREQLIST)

The need for preparing such a checklist by an IETF working group was acknowledged. This will help prospective buyers of routers.

In addition, it was decided that separate checklists were in order to identify the functionality sought in the router, [Exterior, Interior (Dumb, Smart, Big, Small) Leaf-node etc.]. The checklists will extract RFC Compliance features from the replacement for RFC 1009 (Requirement for IP Routers, Philip Almquist, Editor). A mailing list (rrcl@cerf.net, rrcl-request@cerf.net) was created for this purpose.

The BOF (to be turned into a working group) will meet at the San Diego IETF to finalize the checklists. Meanwhile, discussions will commence on the mailing list.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Kurt Dobbins	dobbins@ctron.com
Susan Estrada	estradas@sdsc.edu
Eugene Hastings	hastings@psc.edu
Mike Marcinkevicz	mdm@csu.net
Donald Morris	morris@ucar.edu
Miguel Sasson	sasson@xylogics.com
Frank Solensky	solensky@clearpoint.com
Osamu Takada	takada@sdl.hitachi.co.jp

CURRENT MEETING REPORT

Minutes of the Developing Operational Measurement Criteria BOF (OPMEAS)

Report not submitted. Refer to Area Report for a brief summary.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
John Curran	jcurran@bbn.com
Susan Estrada	estradas@sdsc.edu
Darrell Furlong	
Jack Hahn	hahn@sura.net
Ellen Hoffman	esh@merit.edu
Dale Johnson	dsj@merit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Timothy Salo	tjs@msc.edu
Frank Solensky	solensky@clearpoint.com
Iris Tal	437-3580@mcimail.com

3.5.1 Benchmarking Methodology (bmwg)

Charter

Chair(s):
Scott Bradner, sob@harvard.edu

Mailing Lists:
General Discussion: bmwg@harvisr.harvard.edu
To Subscribe: bmwg-request@harvisr.harvard.edu
Archive:

Description of Working Group:

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

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

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

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the Group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

Done Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Done The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.

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

Request For Comments:

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices"

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

Minutes of the Benchmarking Working Group (BMWG)

The Working Group met on Wednesday, November 20th. The current draft for the testing methodology was reviewed line by line and a number of changes were suggested. There are tentative plans to hold another video conference sometime in late January or early February to do a final pre-IETF run-through of a full document. The results of the video conference will be submitted as an Internet Draft and will be reviewed at the March IETF meeting.

Attendees

Steve Alexander	stevea@i88.isc.com
Atul Bansal	bansal@wile.nac.dec.com
James Beers	beers@nr-tech.cit.cornell.edu
Scott Bradner	sob@harvard.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Jim DeMarco	jdemarco@ftp.com
Barbara Denny	denny@sri.com
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Darrell Furlong	
Olafur Gudmundsson	ogud@cs.umd.edu
Jack Hahn	hahn@sura.net
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Mark Lewis	mlewis@telebit.com
Jim McQuaid	
Bill Melohn	melohn@auspex.com
David Minnich	dwm@fibercom.com
Frank Solensky	solensky@clearpoint.com
Iris Tal	437-3580@mcimail.com
Dean Throop	throop@dg-rtp.dg.com
Jonathan Wenocur	jhw@shiva.com
Jeremy Wilson	
June-Kang Yang	natadm!yang@uunet.uu.net
Nancy Yeager	nyeager@ncsa.uiuc.edu
John Ziegler	ziegler@artel.com

3.5.2 Network Joint Management (njm)

Charter

Chair(s):
Gene Hastings, hastings@psc.edu

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

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.

Goals and Milestones:

None specified

CURRENT MEETING REPORT

Reported by Kent England/BBN

Minutes of the Network Joint Management Working Group (NJM)

Gene Hastings passed out an Agenda and we started talking about simple things we can do now [to make the world better and safer].

Canonical Trouble Mailbox

What if we all defined a mailbox named "trouble@your.net" for receipt of network trouble reports?

What if we all defined a finger-name "noc@your.net" so users could receive a small bit of information about our respective NOCs?

There was some general discussion of DNS records needed for this. Many nets don't have a machine called "your.net", but use MX records. Some don't use MX records for "your.net", but instead use a machine named "noc.your.net" or similar. Finger requires an address record, but an alias record could provide some indirection for "your.net"

Vikas Aggrawal, JvNCnet, noted that he had posted a note on the namedroppers list discussing this. [Vikas, repost to NJM?]

So what names should we use? Joe Ragland, CONCERT, doesn't like "noc" and Carol Ward from WestNet doesn't like "trouble". Gene suggested the name "net-trouble@noc.your.net" for the canonical mailbox, but this is too long for finger, so use "noc@noc.your.net" for finger.

Dan Long, NEARnet, will add a field to his NOC PhoneBook entry for nets to include the preferred mailbox name.

Then Gene raised the issue of DNS inverse address-to-name lookups (PTR records). Gene suggested that all router interfaces should inverse lookup to a descriptive name, and, in addition, the host zero address should invert to something descriptive [of what?]. Vince Fuller noted that DNS already maps host zero to the gateway name. This is left as an unresolved issue.

John Curran, NEARnet, noted upcoming Responsible Person DNS records that we will find useful. Gene Hastings reminded us all to fill out information for Dan Long's NOC PhoneBook entries. Dan said he would send a note to NJM with changes to the entry information.

Gene raised the issue of "nsr" mailing list usage, and suggested a new Merit list for discussion of internet woes called "internet-ops@merit.edu". Ittai Hershman, ANS, said there had been some discussion in Merit/ANS about moving discussions from nsr to nwg or Gene's suggested new list. The concern is about too large a group of readers engaging in too much traffic of discussion and diluting the quality of the nsr list for operational people. John Curran noted that individuals will use whatever list they can find to report trouble and we need to educate users as these misdirections are corrected. Ittai noted that this had been discussed this morning in UCP in the context of machine parseable messages, but these lists don't scale well. Dale Johnson, Merit, noted that this issue parallels the issue of usage of the IETF list and that we need to create "nsr-discuss" as has been suggested for IETF.

Joe Ragland said that one reason we need this list is to relieve Merit of unrelated traffic. FARNET members need an independent channel for communication. Cathy Wittbrodt, ESnet, asked if users aren't going to use this list and Ittai said UCP has addressed this concern. Dan Long noted that service providers need a discussion list for themselves.

We agreed to use the existing nwg mailing list as an "nsr-discuss" list and to use NJM as a meta problem discussion list. [So nsr remains the channel for Merit to send out announcements, nwg is for discussion of operational problems, and NJM is for meta discussions about ops. -kwe]

Tricks of the Trade

Vince Fuller, BARRNET, noted that he is tired of reports to BARRNET from users that say that BARRNET is broken, when in actuality these users are simply unable to traceroute across the MILNET. How can users be made more aware of the limitations of traceroute?

Jordan Becker asked if every AS has a reliable host for pinging and tracing? Dan Long noted he will include such an entry in his NOC PhoneBook.

What about test servers to test telnet, mail, etc? Dan Long noted the success of the NEARnet mail bouncer [bouncer@nic.near.net] as a very useful tool for site contacts to use to test mailers. This automatic bouncer has reduced the workload on NEARnet operations and analyst staff tremendously and is seen as a very valuable service [almost free].

How do we associate network numbers to the AS announcing them? Jessica Yu of Merit noted that they have this file at Merit [net.now?] Ittai noted that new tools for getting and updating this information are under development [ANS?] and we should be hearing more about this in the near future.

Cathy, ESnet, noted that she builds router access control lists from this file and Merit should announce changes to this file format in advance to avoid Cathy having problems like the day she lost NEARnet when it went over to the T3 backbone and the file format changed.

Gene noted that Van Jacobson has a new path characteristics analysis tool that analyzes paths per hop.

John Curran noted that the NNSC is doing another issue of the Internet Managers PhoneBook on paper and electronically.

Dial-ups and Serial Port Servers

Dave OLeary, SURANet, noted that he had sent a note to regional-techs asking for information on dial-up service and hadn't gotten much response. This started a discussion of dial-up servers, or serial port servers. Gene Hastings noted that PSC has a NetBlazer dial-up with SLIP for schools to use. Lines are shared.

Dave OLeary asked about the Livingston product and Brian Lloyd and Vince Fuller seemed to know most about this new product.

There was some discussion about the difficulty of using a NetBlazer as a router. Seems NetBlazers don't do dynamic routing very well. And then there is the difficulty of tracking hosts amongst serial ports.

Milo Medin, NSI, noted that dial-up servers could use OSPF and advertise host routes in order to solve the host tracking problem.

The question was asked about who sells serial port servers and the list looks like:

- cisco TRouter (NEARnet uses)
- Xylogics Annex
- NetBlazer
- Xyplex has something
- NAT cheap router (see Vince Fuller for more info)

Other Business

As we seemed to have used up all our time, the other Agenda items were deferred.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Jordan Becker	becker@nis.ans.net
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Kent England	kwe@bbn.com
Vince Fuller	vaf@stanford.edu
Jack Hahn	hahn@sura.net
Eugene Hastings	hastings@psc.edu
Ittai Hershman	ittai@nis.ans.net
Dale Johnson	dsj@merit.edu
James Jokl	jaj@virginia.edu
Dan Jordt	danj@nwnet.net
Mark Knopper	mak@merit.edu
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@nic.near.net

Milo Medin	medin@nsipo.nasa.gov
Donald Morris	morris@ucar.edu
David O'Leary	oleary@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Joe Ragland	jrr@concert.net
Ron Roberts	roberts@jessica.stanford.edu
Harri Salminen	hks@funet.fi
Tom Sandoski	tom@concert.net
Bernhard Stockman	boss@sunet.se
Roxanne Streeter	streeter@nsipo.nasa.gov
Claudio Topolcic	topolcic@nri.reston.va.us
Carol Ward	cward@spot.colorado.edu
Cathy Wittbrodt	cjw@nersc.gov
Jessica Yu	jyy@merit.edu

3.5.3 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, boss@sunet.se
Phillip Gross, pgross@nis.ans.net

Mailing Lists:

General Discussion: oswg-l@wugate.wustl.edu
To Subscribe: oswg-l-request@wugate.wustl.edu
Archive:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The Working Group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Done	Agreement on a model.
Done	Survey for most useful and popular metrics.
Done	Survey for most useful and popular presentation formats.
Dec 1990	Identify similar efforts being performed by other groups.
Done	Define a common minimal set of metrics.
Mar 1991	Propose a MIB for metrics not already there.
Done	Define a common storage format to facilitate data sharing.
Done	Define common presentation formats to make data comparable.
Mar 1991	Develop outline, and make writing assignments for paper (Opstat1) documenting March 91 milestones.
May 1991	Complete paper Opstat1.

May 1991	Possible mid-term meeting to review Opstat1.
May 1991	Submit Opstat1 as Internet Draft.
Jul 1991	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Jul 1991	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Jul 1991	Propose old tools to be retrofitted.
Jul 1991	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools.
Sep 1991	Complete paper Opstat2.
Sep 1991	Possible mid-term meeting to review Opstat2.
Sep 1991	Submit Opstat2 as Internet Draft.
Dec 1991	Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Claudio Topolcic/CNRI and Bernhard Stockman/NORDUnet

Minutes of the Operational Statistics Working Group (OPSTAT)

Monday's Session

The purposes of this meeting were:

- Review the current status of the OPSTATS activities.
  - Bernhard's papers
  - Other related efforts, specifically, Susan Estrada's BOF
- Decide what can be progressed now and progress it.
  - Model
  - Set of metrics (simple SNMP only)
  - Display formats
  - Simple collection, storage, and exchange
- Define what is still left to do.
  - MIB for new SNMP variables
  - Exchange protocol
  - More sophisticated storage formats
  - Develop publicly available collection tools
  - Display formats for weekly and instantaneous reports
- Specific actions to be taken in this meeting were:
  - Decide polling period
  - Agree on what to progress
  - Edit Bernhard's papers, review on Thursday, submit as Internet Draft

The model was presented for people who were new to the Group. A fundamental part of this model is the agreement on a common minimal set of metrics that will be collected. It was noted that some of these may be difficult to obtain.

It had been proposed that there would be three report formats that would be produced; a monthly report, a weekly report, and an instantaneous display. A format for the monthly report had been agreed to. It was described as a "McDonalds" report because it would contain only total aggregates. It was felt that this report would support management activities, whereas the weekly report would support engineering planning, and the instantaneous display would support problem resolution. However, it was realized that the real distinction was not the time-frame but the degree of aggregation of the data. The data in the management reports would be more aggregated than that in the engineering reports, regardless of the time it covered.

Bernhard's documents described the data that would be collected from each router, both for each of the router's interface, and for the router itself. These are all MIB variables. It was at first assumed that the per interface variables were specific to IP, but it was pointed out that the loading data needs to be total, not IP specific, or the link loading could not be determined. It was also pointed out that the MIB interface variables are multi-protocol anyway, so there is no problem. However, it was also pointed out that if the router variables are IP only, then they do not give a measure of the router's loading.

It was noted that the loading information that is important is not related to any interface, but to the links. Links are occasionally re-homed when interfaces fail. Currently, the data is processed by hand to compensate for such re-homing. The documents do not make this distinction and need to be clarified.

Dropping the "storage requirements" section of Bernhard's document was considered, but it was decided to keep it in, since dropping it would give the mis-impression that the Group hadn't thought about the problem.

It had been proposed that the client-server model not be covered in the current documents. The reason, in part, was that the original purpose of the Working Group was to get the various network operators to produce consistent reports that could be compared, not to exchange information, and that exchanging information is not required very often.

The data storage format was discussed. The format impacts what will be stored and what can be done with it. To reduce storage requirements, several people proposed that raw data could be kept for some period of time, and then aggregated somewhat and kept for some other period of time, and then further aggregated. The proposals differed in the time periods, and the form of aggregation. However, it was pointed out that although engineering requirements tend to be common, so common non-aggregated data will be useful, management requirements tend to differ, so common aggregated data is not useful. In the end, it was realized that how much data is retained, and how long, are local decisions that cannot be standardized.

The data format should support the process that the data will undergo. The process was identified as:

- Collect status data about routers and interfaces.
- Collect "resource" data, for example, about the mapping of links to interfaces.
- Process the data to merge 1 and 2, decreasing the quantity of data but without loss of information.
- Produce reports from the above reduced data.

It was understood that the processing in Step three would not lead to sufficient reduction in quantity to address long-term data storage problems. However, it was felt that this processing should not be combined with the report generation.

Bernhard proposed a raw data format, which was discussed. He will incorporate suggestions into his document.

It was suggested that the monthly reports be based on a matrix that identified all the variables that would be collected and processing functions that could be applied to them. This would not only clearly delimit the scope of the report generation process, but would also allow new variables to be added easily. However, this approach would not support functions that are based on multiple variables, and although the matrix could be relatively full, any network operator might select only a few possibilities, and worse, the different operators might select different sets.

It was felt that the Working Group should recommend a specific polling period. Two were on the table; 5 minutes and 15 minutes. Concern was expressed that 5 minutes or less might result in excessive overhead or be impossible to implement with a poller that polls one router at a time. For variables describing link loading, such as bytes transmitted, the polling period is a function of the line speed. A one minute polling period will miss the interesting peaks of a T1 line, but will show the individual packets on a 1200 baud line. For variables not describing link loading, such as packets dropped, the polling interval can generally be very long, until the value changes, at which time the polling period should be shortened to help identify the problem. So it may be that a 15 minute polling period is sufficient for anything other than link utilization. This discussion was deferred until the next meeting on Thursday.

Geoff Huston suggested a different approach. He proposed that the link utilization parameter that is most closely correlated to the clients' dissatisfaction is the mean standard deviation of inter-packet arrival times of evenly spaced (when transmitted) TCP packets. He suggested that this parameter explodes as soon as congestion appears.

Thursday's Session

During the second OPSTAT session the storage format and the polling periods were discussed in more detail.

The Storage Format

The placeholder for the header section is suggested to be within the log-file. However, there might be use for both separate and in-band headers. The need for multiple header sections within one log-file was expressed. When closing and reopening the same log-file there is the need for close and start time specifications. When changing log-source there is the need for specifying a new device. Three delimiter pairs were suggested:

```
BEGIN_TIME   - END_TIME
BEGIN_DEVICE - END_DEVICE
BEGIN_DATA   - END_DATA
```

There are currently two storage formats. The version presented by Bernhard Stockman and and earlier version produced by Chris Myers. Chris Myers volunteered to produce a second version of his storage format strawman.

The generic log data format is:

```
timestamp, tag, delta_sample_interval, data1, data2, data3, ..., dataN
```

where the tag defines the logged variables.

The Polling Period.

The reason for the polling is to achieve statistics to serve as a base for trend and capacity planning. From the operational data it shall be possible to derive engineering and management data.

It will not be sufficient with a polling period of 15 minutes to detect variations in peak-behavior. It was suggested that a maximum period of one minute would be needed. Using such a tight polling period will create a need for aggregating stored data. Aggregation here means that over a period with logged entries, a new aggregated entry is created by taking the first and last of the previously logged entries over some aggregation period and computing a new entry.

A method of displaying both average and peak-behaviors in the same bar-diagram is to compute both the average value over some period and the peak value during the same period. The average and peak values are then displayed in the same bar.

A problem here is how to aggregate peak values. There is the possibility of creating a new peak value being the peak of all the peaks, the average of all the peaks, etc.

Another reason for aggregation is the differentiation of needed polling periods depending on the reason for and source of the polling.

What is foreseen is that over a relatively short period, polled data will be logged at the tightest polling period (one minute) regularly these data will be pre-processed into the actual files being stored. The pre-processing may include steps such as the computation of percent samples above a certain limit, average of all samples during the aggregation period, cumulative histograms. This pre-processing will then not only serve as storage compacting but will also provide some initial statistical processing.

Recommendation on polling period:

```
Basic polling period    1 minute (60 seconds).
```

Recommendation on aggregation periods:

```
Over a

    24 hour period        aggregate to 15 minutes,
    1 month period        aggregate to 1 hour,
    1 year period         aggregate to 1 day
```

Aggregation is the computation of new average and maximum values for the aggregation period based on the previous aggregation period data.

Recommendation for saving periods of logged and aggregated data:

```
    15 minute aggregation period      saved 1 week.
    1 hour aggregation period         saved 1 month.
    1 day aggregation period          saved 1 year.
```

Finally it was decided that, as the current document will not contain the protocol specification of the client-server model, it will be sufficient to put the coming RFC into the informational track.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Miriam Amos Nihart	miriam@decwet.zso.dec.com
Jordan Becker	becker@nis.ans.net
Robert Blokzijl	K13@nikhef.nl
Steve Bostock	steveb@novell.com
Randy Butler	rbutler@ncsa.uiuc.edu
John Gong	jgong@us.oracle.com
Phillip Gross	pgross@nis.ans.net
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Kathleen Huber	khuber@bbn.com
Geoff Huston	g.huston@aarnet.edu.au
Walter Lazear	lazear@gateway.mitre.org
April Marine	april@nisc.sri.com
Robert Morgan	morgan@jessica.stanford.edu
Dennis Morris	morrisd@imo-uvax.dca.mil
Chris Myers	chris@wugate.wustl.edu
Rebecca Nitzan	nitzan@es.net
Marsha Perrott	mlp+@andrew.cmu.edu
Ron Roberts	roberts@jessica.stanford.edu
Timothy Salo	tjs@msc.edu

Bernhard Stockman	boss@sunet.se
Joanie Thompson	joanie@nsipo.nasa.gov
Claudio Topolcic	topolcic@nri.reston.va.us
Andrew Veitch	aveitch@bbn.com
Wengyik Yeong	yeongw@psi.com
Osmund de Souza	osmund.desouza@att.com

3.5.4 User Connectivity (ucp)

Charter

Chair(s):
Dan Long, long@nic.near.net

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

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:

Done — 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.

Internet Drafts:

"FYI on an Internet Trouble Ticket Tracking System for addressing Internet User Connectivity Problems", 02/11/1991, M. Mathis, D. Long <draft-ietf-ucp-connectivity-01.txt>

Request For Comments:

RFC 1297 "NOC Internal Integrated Trouble Ticket System Functional Specification Wishlist ("NOC TT REQUIREMENTS")"

CURRENT MEETING REPORT

Reported by Gene Hastings/PSC and Dan Long/BBN

Minutes of the User Connectivity Problems Working Group (UCP)

Summary

A presentation on the UCP work-in-progress was made at the FARNET meeting on Monday. Useful discussion ensued and continued in the Operational Requirements Area Directorate (ORAD) meeting Tuesday. The consensus was that the UCP work is important and should be actively pursued so that initial implementations are in place in the next 6-12 months.

Encouraged by good suggestions and the support of FARNET and ORAD, the UCP Group met twice and made progress on these three projects:

1. A NOC PhoneBook

   - Reviewed collection efforts to date.
   - NEARnet database forms-based entry being used.
   - Twenty-Two NOCs have registered so far.
   - Suggested new fields and formats (now incorporated in database).
   - Will advertize to wider audience soon.
   - Plan to distribute PhoneBook with a finger-based search tool.
   - All searches will also return caution that information is for NOCs only.

2. Standardized Network Status Reports

   - Developed a syntax for standard email-based outage reports.
   - NOCs will generate these reports which will contain information about current, past, or planned outages.
   - These reports will be sent to a mailing list which anyone can subscribe to.
   - People can develop their own tools for parsing and providing interactive access to this information.
   - Ideally, end-users and NOCs could use such tools to get more information about connectivity problems.

3. Standardized Trouble Ticket Handoffs

   - Revised the UCP Trouble Ticket Tracking draft to allow Network Service Centers to limit who they are required to accept calls from (from last meeting).
   - Network Service Centers will hand off tickets to other Network Service Centers but will stay in the loop with the reporting user.

- Developed a syntax for standard email-based trouble ticket hand-offs between Network Service Centers.
- Several groups are interested in participating in trials of this system.

For more information, join the list: ucp-request@nic.near.net

DETAILED NOTES of the November 19th meeting by Gene Hastings

The Distributed Agenda (roughly)

- Status
- NSC Phonebook
  - NISI
  - Current
  - Future
- Reducing Need for Tickets
  - Notification Schemes
  - Database of Network Status
  - New Working Group?
- UCP Ticket Sharing
  - What Information to Exchange
  - Method/Format for Exchange
  - Implementations

Dan Long (NEARnet) gave an overview of FARNET's interest in UCP topics.

NSC Phonebook - as of the meeting, there were 18 entries in the pilot NSC Phonebook. Note was made of the parallel NISI effort to collect similar listings for NICs. Vikas Aggarwal (JvNC) recommended that contact information be included in DNS TEXT records.

The present NSC Phonebook Database is in Informix. Dan Long volunteered to continue to maintain it and to deploy a finger <keyword> query agent for it. Dale Johnson (Merit) offered a second Informix host for it if someone else would maintain the actual database.

Vikas volunteered to produce a DNS entry template.

Group consensus was an acknowledgement that these are interim efforts.

Reducing the need for tickets - that is, reducing the need for calls from users which require the opening of tickets. This might be secured through:

- Notification Schemes
- Database of Network Status
- End-User diagnostic tools

These are all in keeping with the idea that if users are better informed about the state of the world, and have easier means of learning for themselves the nature of difficulties, they will have less need to call and talk to a person. An example of this at a department or campus level might be a bulletin board which lists scheduled outages and includes explanations of what services will be affected, along with pointers for further inquiries and help files explaining the nature of some classes of failure.

A strawman proposal was made to distribute email with a standard format, initially based on the NEARnet trouble tickets. Discussion followed as to which problem this proposal was intended to solve. Uses for standard format mail include: ease of information extraction when read; ease of parsing for inclusion in a database or for triggering alarms; assurance of completeness of information in report; and the possibility of making many reports machine-generated. Desired fields were felt to include:

- ASN#
- Net#
- Net Name
- Host Address/NSAP
- Host Name
- Affected protocol or service
- Start/End Date & Time
- Responsible Person or NOC
- Ticket Cross-reference
- Last Update
- Reporting NOC
- Perspective/Scope
- EXPLANATION
- FURTHER EXPLANATION

There is still confusion and some disagreement concerning what things are or aren't tickets. There was, and will continue to be, discussion on use, control and interpretation. For example, whether these messages should be intelligible to, and distributed to, end users.

Michael Patton (MIT) observed that poorly formed information distributed to the public would generate more calls, not less.

Ittai Hershman (ANS) reported that nsr <network-status-reports@merit.edu> is now carried in a PSI newsgroup, so the mechanism for end-users to see those messages is in place.

Further work on mail format was deferred until the meeting of the 20th.

Discussion returned to the NSC Phonebook. New fields to add to listings:

- Administrator to escalate to
- Domain name of NOC
- Bigger net # field (allow listing of multiple net numbers)
- Cross references to other nets, centers
- Bigger phone number fields (multiple numbers)
- FAX #
- Discussion of upper vs. mixed case for organization and net names. [there is a practical limitation of the pilot database, in that it will not fold case for searches.]

Questions were raised as to what limitations should be placed on the distribution of this information, if published. Following objections to having internal operations numbers available to arbitrary end-users, Ittai Hershman proposed limiting the distribution of the information to NOCs & NSCs, with harshly-worded boilerplate against indiscriminate release. A quick hack to limit availability is to include an access string in the finger query, acting as a pseudo password; Instead of "finger psinet-nsc@nic.near.net", something like "finger psinet-nsc-abqothl@nic.near.net".

DETAILED NOTES of the November 20th meeting by Dan Long

The second meeting focussed on mechanisms for hand-offs of tickets between NSC's. We agreed that a similar format to that described above should be used to allow hand-offs to be generated and parsed either manually or automatically.

The Group brainstormed a list of fields that would be of interest:

- Description of problem
- Description of solution
- Location (or Source/Destination) of problem: AS#, Net#, Host Address, Service Description, etc.
- Problem Start/End Date/Time
- Ticket Open/Close Date/Time
- Ticket Number (made unique by prepending a unique NSC identifier)
- NSC List (list of NSCs that have handled this problem)
- Notifications (who should be kept informed about this problem?)
- Contact Info (who should be worked with to resolve this problem?)
- Notes: number, date/time, author, text

The Group agreed that this list of items will likely need to evolve but that we should be conservative in the addition of fields so that the syntax remains simple and that the burden on human operators is minimized.

There was a fair amount of discussion about notifications and whether end-users should be notified about steps taken by NSCs other than the originating NSC. Organizations have different policies about how much detail to reveal. The consensus was that the originating NSC can use the Notifications field to include the user (or not) as they see fit and that other NSCs working on the problem should honor the notifications field to report progress.

In the original paper by Matt Mathis, the idea was for the entire ticket to be handed off to the appropriate NSC and for the new NSC to deal with the user. We agreed on a change whereby the originating NSC maintains the contact with the user and keeps its own ticket open on the problem. It may, as the document describes, hand the problem off to another NSC but that NSC must then report back when it is done to the originating NSC who, in turn, will obtain closure with the user. The hand-off will be handled much as the original document describes.

The general format of the mail message will be:

```
To: trouble-ticket-handoff@destination  (the specific address for any given
 NSC is in the NSC PhoneBook)
Subject: ticket-number {hand off, update, close} note-number
(note-number is 0 on 1st hand-off,
        1 on 1st note,
        ...
        N+1 on close)
And in the body of the message:
Fieldname: (contents)
   ...
...
Note: 1 Date Time Author
(note text)
...
...
```

Several people volunteered to begin using these formats for status updates and ticket hand-offs. Dan Long will publish detailed writeup of formats so people can get started.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Thomas Bajzek	twb+@andrew.cmu.edu
Robert Blokzijl	K13@nikhef.nl
Jeff Erwin	
Susan Estrada	estradas@sdsc.edu
Farrell Gerbode	farrell@rice.edu
John Gong	jgong@us.oracle.com
Eugene Hastings	hastings@psc.edu

Ittai Hershman	ittai@nis.ans.net
Dale Johnson	dsj@merit.edu
James Jokl	jaj@virginia.edu
Dan Jordt	danj@nwnet.net
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@nic.near.net
Donald Morris	morris@ucar.edu
David O'Leary	oleary@sura.net
Michael Patton	map@lcs.mit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Joe Ragland	jrr@concert.net
Ron Roberts	roberts@jessica.stanford.edu
Tom Sandoski	tom@concert.net
Bernhard Stockman	boss@sunet.se
Carol Ward	cward@spot.colorado.edu
Cathy Wittbrodt	cjw@nersc.gov
Paul Zawanda	zawanda@ncsa.uiuc.edu

3.6 Routing Area

Director(s):

- Bob Hinden: hinden@bbn.com

Area Summary reported by Bob Hinden/BBN

There continued to be a lot of routing related activities at the Santa Fe IETF. Two talks were given on inter-domain routing. One which represents the work of the IDPR Working Group titled "Inter-Domain Policy Routing" was presented by Martha Steenstrup. The other talk described an approach which combines features of IETF Inter-Domain Policy Routing Protocol (IDPR) and the ISO Inter-Domain Routing Protocol (IDRP). This talk entitled "A Unified Approach to Inter-Domain Routing" was presented by Deborah Estrin.

Seven Routing working groups met and there were three Routing related BOF's. Their reports are as follows:

Border Gateway Protocol

The Group discussed subnet mask support in BGP. Issues of advertising subnet and supernet masks between autonomous systems were discussed. The discussion evolved into the need to develop an overall internet routing architecture. This lead into the formation of a group which evolved into the ROAD Group.

The BGP Group discussed the Internet Draft of describing OSPF and BGP interactions, with the OSPF Working Group.

IP Over Large Public Data Networks

The IPLPDN Group agreed on bridging formats and protocol identification on IEEE 802.6 LAN's. The Group will forward its recommendations to the IEEE 802.6 committee. This was done in response to a request from 802.6.

Protocol identification and encapsulation for circuit ISDN (B channels) were discussed with members of the Point-to-Point Protocol Extensions Working Group. The Group agreed to use the XID procedure to determine the encapsulation format. A code point will be obtained to indicate PPP in ISDN. The Group will document this in an Internet Draft.

An Internet Draft updating IP over X.25 was discussed. The comments will be incorporated in a new version of the document.

There was a long discussion on how routing should work over large public data networks. Proposals for "Directed ARP" and "Short Cut" routing were discussed. One of the major issues is whether there will be different IP network addresses on each public network or if there will be one (i.e., class A) network number for each network. In the former case should traffic be required to go through a router.

The Group decided to recommend to the IESG that the "IP over SMDS" RFC be advanced to Draft Standard.

IP Routing for Wireless/Mobile Hosts

This BOF met and was well attended. Dan Duchamp gave a presentation on the work done at Columbia University to support roaming wireless IP Hosts. Ken Carlberg described his work at SAIC on connectionless ISO CLNP support for mobile end systems. Steve Deering reported on the current work of the IEEE 802.11 Wireless LAN committee.

The attendees agreed there was enough interest to form a working group. Steve Deering will Chair the Group and write a Charter.

IS-IS for IP Internets

The Group met and updated the RFC for Dual IS-IS to be compatible with the current international standard for IS-IS. An Internet Draft reflecting these changes will be produced soon. The Group also discussed the IS-IS MIB. An Internet Draft is now available.

Work was started on BGP-ISIS interactions. The approach taken is similar to what has been done with OSPF and BGP.

Inter Domain Policy Routing

The Group reviewed the current Internet Drafts describing the IDPR protocol. They also discussed how IDPR would be used in the Internet and how it interacts with existing routing protocols. The Group also reviewed the status of the IDPR "gated" implementation.

The Group met jointly with the Router Requirements Working Group. The goal was to make sure that IDPR was not inconsistent with the router requirements specification. The result of the discussion was that there were not any inconsistencies.

Multicast Extensions to OSPF

The Group completed reviewing the first draft of MOSPF specification (60+ pages). The main concerns raised were in the area of inter-autonomous system interactions and co-existence of multiple multicast routing protocols in the same router.

An implementation of MOSPF is underway at Proteon and a "gated" implementation is planned by Cornell.

New Internet Routing and Addressing Architecture

The BOF discussed Noel Chiappa's proposal for a new internet routing and addressing architecture. The first half of the meeting was spent reviewing the proposal and the last half was spent getting a better understanding of the overall dimensions of the architecture. There was general interest in forming a working group.

Open Shortest Path First IGP

The Group discussed the "Virtual Link" problem found during multi-vendor OSPF interoperability testing in Foxborough, MA. A backward compatible fix has been developed. The OSPF specification is being updated to reflect the fix.

The Group reviewed a set of changes to the OSPF MIB and the new OSPF Trap MIB. The later will be published as an Internet Draft.

Work continued on the proposal for a new OSPF option "Not So Stubby Area". This option will provide for improved support for RIP clouds attached to OSPF domains and help with the transition of domains from RIP to OSPF. This option is now fully fleshed out.

The Group also discussed the possibility of defining a reduced subset of OSPF for autonomous systems, such as stub domains, not needing the full functionality.

RIP Version II

A BOF was held to see if there was sufficient interest to form a working group to develop a new version of RIP. Features to be added to the protocol included:

- Subnet Masks
- Authentication
- Autonomous System Numbers
- MTU & Link Speed indications

The attendees agreed that a working group should be formed.

CURRENT MEETING REPORT

Reported by Dino Farinacci/cisco

Minutes of the Border Gateway Protocol Deployment and Application BOF (BGPDEP)

Agenda

- Coordinate BGP deployment
- Coordinate BGP policy implementations

1. Implementation testing for interoperability and verification before production deployment is desirable.

2. RFC 1267 & 1268 refers to the latest BGP protocol and usage.

3. Tony Li from cisco volunteered to coordinate joint testing. He is reachable at cs@cisco.com. Jessica Yu volunteered Merit as the place where the testing could occur.

4. Who's using BGP?

   - PSC - gated on T1 NSS and T3 ENSS.
   - SDSC - gated on T1 NSS.
   - CA*Net - gated with Dennis Ferguson's extensions.
   - Cornell/ICM - cisco.
   - Merit - cisco 8.2(6) on T1 and T3 NSS.
   - VPI - cisco on T3
   - IBM Information Network - cisco on T3
   - ANS connection to T3 backbone

5. Who's testing it?

   - Sesquinet
   - Argonne
   - ConcertNet
   - ESNet
   - Nearnet
   - EASINet

6. Who has implemented it?

   - cisco BGP2 in software release 8.3.

- NSFnet which supports BGP1 and BGP2.
- Gated which supports BGP1 only.
- Implementation from Canada.
- BBN in the T20 router - BGP3.

7. BGP Gated Status - by Scott Brim.

   - BGP2 & BGP3.
   - Flexible definition of neighbors.
   - Internal, External, and Test neighbor definition. Test dumps everything to peer but believes nothing.
   - Internal BGP works. Can run with OSPF and synchronizes correctly. Currently there is no flexibility to adapt to other IGPs.
   - Shared subnet between external neighbors is enforced.
   - If between two ASs, can have different AS numbers, and AS paths will be correct.
   - Improved performance.
   - RFC 1164 specifies a policy language that has been implemented in gated.

8. BGP MIB

   - Will be updated to reflect the latest protocol version 1Q92

9. cisco Implementation Status - by Tony Li.

   - BGP2 available in 8.3(1) and 8.2(7) soon.
   - 9.0 available 1Q92 which has BGP2 & BGP3.
     - AS-path access lists are supported. Uses regular expressions.
   - BGP to do list - for future, no commitments.
     - Multiple BGP processes.
     - Use AS number filter for redistribution into OSPF.
     - Performance enhancements.
     - Convergence time improvements to IP forwarding table for BGP routes.

10. NSFnet's implementation status - by Yakov Rekhter.

    - Current version is BGP2. Plans to use gated's implementation.

- Plans to carry BGP information in IS-IS as an alternative to internal BGP.

11. What sites tried BGP but backed it out and reasons why.

    - NEARnet
      - Redistribution problems.
    - ESNet
      - Access list problems.
    - MILnet at FIX-E
      - Ran into memory shortage problem due to usage of CSC2 board. In the process of upgrading it to CSC3 board.
    - BARRNet
      - Multiple vendors interoperating with half running BGP had consistency problems.
    - Sesquinet
      - Redistribution problems; will try new software from cisco.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
William Babson	bill%penril@uunet.UU.NET
Jordan Becker	becker@nis.ans.net
Robert Blokzijl	K13@nikhef.nl
David Bolen	db31@nis.ans.net
Scott Bradner	sob@harvard.edu
Scott Brim	swb@nr-tech.cit.cornell.edu
Ronald Broersma	ron@nosc.mil
Rob Coltun	rcoltun@ni.umd.edu
Dino Farinacci	dino@cisco.com
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Kathleen Huber	khuber@bbn.com
Dan Jordt	danj@nwnet.net
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Mark Knopper	mak@merit.edu
Tony Li	tli@cisco.com
Daniel Long	long@nic.near.net
Andrew Malis	malis@bbn.com

April Merrill	abmerri@tycho.ncsc.mil
Donald Morris	morris@ucar.edu
David O'Leary	oleary@sura.net
Yakov Rekhter	yakov@watson.ibm.com
Harri Salminen	hks@funet.fi
Tom Sandoski	tom@concert.net
Sharad Sanghi	sharad@ans.net
Stephen Shew	sdshew@bnr.ca
Brad Steinka	brad@python.eng.microcom.com
Bernhard Stockman	boss@sunet.se
Sally Tarquinio	sally@gateway.mitre.org
Kannan Varadhan	kannan@oar.net
Carol Ward	cward@spot.colorado.edu
Jonathan Wenocur	jhw@shiva.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Robert Woodburn	woody@cseic.saic.com
Jessica Yu	jyy@merit.edu

CURRENT MEETING REPORT

Minutes of the New Internet Routing and Addressing BOF (NIMROD)

Report not submitted. Refer to Area Report for brief summary.

Attendees

Karl Auerbach	karl@eng.sun.com
Robert Austein	sra@asylum.sf.ca.us
William Babson	bill%penril@uunet.UU.NET
Fred Baker	fbaker@emerald.acc.com
Atul Bansal	bansal@wile.nac.dec.com
William Barns	barns@gateway.mitre.org
Daniel Blum	4108980@mcimail.com
Scott Bradner	sob@harvard.edu
David Bridgham	dab@asylum.sf.ca.us
Scott Brim	swb@nr-tech.cit.cornell.edu
Ronald Broersma	ron@nosc.mil
Ken Carlberg	carlberg@cseic.saic.com
J. Noel Chiappa	jnc@lcs.mit.edu
Richard Cogger	rhx@cornellc.cit.cornell.edu
Rob Coltun	rcoltun@ni.umd.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Steve Deering	deering@xerox.com
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Kurt Dobbins	dobbins@ctron.com
Kent England	kwe@bbn.com
Dino Farinacci	dino@cisco.com
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Vince Fuller	vaf@stanford.edu
Chris Gunner	gunner@osicwg.enet.dec.com
Jack Hahn	hahn@sura.net
Eugene Hastings	hastings@psc.edu
Christine Hemrick	hemrick@cisco.com
Robert Hinden	hinden@bbn.com
Randy Hoebelheinrich	jrc@lanl.gov
Bob Jeckell	robert_jeckell@nso.3com.com
Dan Jordt	danj@nwnet.net
Scott Kaplan	scott@ftp.com
Michael Karels	karels@cs.berkeley.edu
Frank Kastenholz	kasten@europa.clearpoint.com
Tom Kessler	kessler@sun.com
Stev Knowles	stev@ftp.com

Anthony Lauck	lauck@tl.enet.dec.com
Walter Lazear	lazear@gateway.mitre.org
John Lekashman	lekash@nas.nasa.gov
Mark Lewis	mlewis@telebit.com
Tony Li	tli@cisco.com
Brian Lloyd	brian@ray.lloyd.com
E. Paul Love	loveep@sdsc.edu
Carl Malamud	carl@malamud.com
Gary Malkin	gmalkin@ftp.com
Glenn McGregor	ghm@merit.edu
Bill Melohn	melohn@auspex.com
Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.dca.mil
John Moy	jmoy@proteon.com
Mark Needleman	mhn@stubbs.ucop.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Andy Nicholson	droid@cray.com
David O'Leary	oleary@sura.net
Zbigniew Opalka	zopalka@roswell.spartacus.com
J. Bradford Parker	brad@cayman.com
Michael Patton	map@lcs.mit.edu
Yakov Rekhter	yakov@watson.ibm.com
Michael Roberts	roberts@educom.edu
Philippe Roger	lircorp1!roger@uunet.uu.net
Harri Salminen	hks@funet.fi
Sharad Sanghi	sharad@ans.net
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Stephen Shew	sdshew@bnr.ca
William Simpson	Bill_Simpson@um.cc.umich.edu
Ursula Sinkewicz	sinkewic@decvax.dec.com
Frank Solensky	solensky@clearpoint.com
Martha Steenstrup	msteenst@bbn.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Osamu Takada	takada@sdl.hitachi.co.jp
Iris Tal	437-3580@mcimail.com
Sally Tarquinio	sally@gateway.mitre.org
William Townsend	townsend@xylogics.com
Mike Truskowski	truskowski@cisco.com
Kannan Varadhan	kannan@oar.net
Lee Wade	wade@nsipo.arc.nasa.gov
Carol Ward	cward@spot.colorado.edu
Jonathan Wenocur	jhw@shiva.com
Scott Williamson	scottw@nic.ddn.mil
Walter Wimer	walter.wimer@andrew.cmu.edu
C. Philip Wood	cpw@lanl.gov

Robert Woodburn woody@cseic.saic.com
Jessica Yu jyy@merit.edu

3.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):
Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:
General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu
Archive:

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done	Complete development of Version 2 of the Border Gateway Protocol (BGP).
Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
Done	Develop a MIB for BGP.
Done	Work with the Security Area to enhance the provision for security in BGP.
Done	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

Internet Drafts:

"Border Gateway Protocol NEXT-HOP-SNPA Attribute", 04/15/1991, Paul Tsuchiya <draft-ietf-bgp-nexthop-01.txt>

"Default Route Advertisement In The Border Gateway Protocol", 08/09/1991, Dimitry Haskin <draft-ietf-bgp-defaultroute-00.txt>

"Multicast Communications Using BGP", 08/26/1991, Scott Brim <draft-ietf-bgp-multicast-01.txt>

"BGP OSPF Interaction", 10/25/1991, Kannan Varadhan <draft-ietf-bgp-ospfinteract-02.txt>

"A Unified Approach to Inter-Domain Routing", 12/06/1991, D. Estrin, Y. Rekhter, S. Hotz <draft-ietf-bgp-unirouting-00.txt>

Request For Comments:

RFC 1105	"Border Gateway Protocol BGP"
RFC 1163	"A Border Gateway Protocol (BGP)"
RFC 1164	"Application of the Border Gateway Protocol in the Internet"
RFC 1265	"BGP Protocol Analysis"
RFC 1266	"Experience with the BGP Protocol"
RFC 1267	"A Border Gateway Protocol 3 (BGP-3)"
RFC 1268	"Application of the Border Gateway Protocol in the Internet"
RFC 1269	"Definitions of Managed Objects for the Border Gateway Protocol (Version 3)"

CURRENT MEETING REPORT

Reported by Yakov Rekhter/IBM

Minutes of the Border Gateway Protocol Working Group (BGP)

The BGP Working Group had three meetings during the Santa Fe IETF. The first meeting started with a presentation by Phill Gross (ANS) on the need for supporting subnet masks in BGP. Phill's presentation was followed by an extensive discussion on the subject of IP addressing and routing. Discussion rapidly steered from the initially proposed subject (subnet masks in BGP) towards general issues of IP addressing and routing. It was recognized that the problems and issues that surfaced during the discussion were not BGP problems, and thus could not be solved by introducing subnet masks support in BGP. It had been also recognized that before introducing subnet masks support in BGP we need a much better understanding of the problem(s) we are trying to solve. At the end of the first meeting it was suggested that a new Working Group should be formed. This Working Group should look at the general issues of addressing and routing in IP internets.

The second meeting was dedicated to the discussion of the OSPF-BGP Interaction document. The author of the document received numerous comments. The comments were incorporated into the draft, and a new version of the draft is available (draft-ietf-bgp-ospfinteract-01.txt). The second meeting involved active participation from members of the OSPF and Router Requirements Working Groups. We expect to turn the draft into a Proposed Standard by the next IETF.

The third meeting was given to Jessica Yu (NSFnet/Merit) who chaired the BGP Deployment and Application BOF. Minutes of this BOF will be available separately.

Attendees

Steve Alexander	stevea@i88.isc.com
Philip Almquist	almquist@jessica.stanford.edu
Nagaraj Arunkumar	nak@3com.com
William Babson	bill%penril@uunet.UU.NET
David Bolen	db31@nis.ans.net
Scott Brim	swb@nr-tech.cit.cornell.edu
Ronald Broersma	ron@nosc.mil
Gregory Bruell	gob@shiva.com
Theodore Brunner	tob@thumper.bellcore.com
Ross Callon	callon@bigfut.enet.dec.com
Ken Carlberg	carlberg@cseic.saic.com
Rob Coltun	rcoltun@ni.umd.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
John Damiano	
Tim Dixon	dixon@nikhef.nl
Dino Farinacci	dino@cisco.com

Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Phillip Gross	pgross@nis.ans.net
Christine Hemrick	hemrick@cisco.com
Robert Hinden	hinden@bbn.com
Kathleen Huber	khuber@bbn.com
Ronald Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Mark Knopper	mak@merit.edu
Ron Lau	rlau@synoptics.com
Tony Li	tli@cisco.com
April Merrill	abmerri@tycho.ncsc.mil
Dennis Morris	morrisd@imo-uvax.dca.mil
John Moy	jmoy@proteon.com
Radia Perlman	perlman@radia.enet.dec.com
Thomas Pusateri	pusateri@cs.duke.edu
Yakov Rekhter	yakov@watson.ibm.com
Manoel Rodrigues	manoel.rodrigues@att.com
Sharad Sanghi	sharad@ans.net
Frank Solensky	solensky@clearpoint.com
Ravi Srinivasan	ravi@eng.vitalink.com
Martha Steenstrup	msteenst@bbn.com
Sally Tarquinio	sally@gateway.mitre.org
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Kannan Varadhan	kannan@oar.net
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	walter.wimer@andrew.cmu.edu
June-Kang Yang	natadm!yang@uunet.uu.net
Yueli Yang	yueli@bnr.ca
Jessica Yu	jyy@merit.edu
Osmund de Souza	osmund.desouza@att.com

3.6.2 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

George Clapp, clapp@ameris.center.il.ameritech.com

Mailing Lists:

General Discussion: iplpdn@nri.reston.va.us
To Subscribe: iplpdn-request@nri.reston.va.us
Archive: /ietf.mail.archives/iplpdn.mail.archive

Description of Working Group:

The IP over Large Public Data Networks Working Group (IPLPDN) will specify the operation of the TCP/IP protocol suite over public data networks (PDNs) such as SMDS, ISDN, X.25 PDNs, and Frame Relay. The Working Group will develop and define algorithms for the resolution of IP addresses and for the routing of IP datagrams over large, potentially global, public data networks.

The IP over SMDS Working Group has defined the operation of the Internet protocols when SMDS is used to support relatively small virtual private networks, or Logical IP Subnets (LISs). Issues arising from public and global connectivity were delegated to the IPLPDN Working Group.

The IPLPDN Working Group will also continue the work of the Private Data Network Routing Working Group (pdnrout) on X.25 PDNs. This work will be extended to include call management and the use of the ISDN B channels for the transport of IP datagrams.

Address resolution and routing over Frame Relay will also be discussed.

Goals and Milestones:

Done	Establish priorities and dates of completion for documents.
TBD	Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.
TBD	Routing of IP datagrams across very large internets implemented SMDS and on other PDNs.
TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.

Internet Drafts:

"Discovery and Routing over the SMDS Service", 06/17/1991, Paul Tsuchiya <draft-tsuchiya-routingsmds-01.txt>

"Management Information Base for Frame Relay DTEs", 06/17/1991, Caralyn Brown, Fred Baker, Charles Carvalho <draft-ietf-iplpdn-frmib-05.txt>

Request For Comments:

RFC 1293 "Inverse Address Resolution Protocol"

RFC 1294 "Multiprotocol Interconnect over Frame Relay"

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

Minutes of the IP over Large Public Data Networks Working Group (IPLPDN)

The IPLPDN Working Group covered the following topics:

- The IEEE 802.6i draft standard specifying protocol identification and encapsulation formats for remote bridging over 60 bit address DQDB subnetworks was reviewed and approved for release to a confirmation ballot by the IEEE 802.6 Working Group.

- Protocol identification and encapsulation for circuit ISDN (B channels) was discussed with members of the PPPEXT Group. The Group agreed to use the XID procedure to determine the encapsulation format. Andy Malis volunteered to contact Joel Halpern and Lyman Chapin to obtain a code point indicating PPPEXT. William Jolitz volunteered to write an Internet Draft on the procedure.

- An Internet Draft updating IP over X.25 was reviewed. Comments were received by Andy Malis, who will incorporate the comments in a new version of the draft for review at the next IETF meeting.

- It is anticipated that multiple IP subnetworks will be implemented on public data networks such as SMDS, Frame Relay, and X.25. Currently, two IP stations on different IP subnetworks implemented on the same PDN can only communicate via a router. The Group discussed two mechanisms which would allow such stations to communicate directly. These mechanisms would also resolve IP addresses to the corresponding PDN addresses (e.g., E.164 and X.121).

- The Group agreed to ask the IAB to advance RFC 1209 to Draft Standard. RFC 1209 describes IP over SMDS.

- Parameter negotiation and procedures for automatic configuration across Frame Relay permanent virtual circuits was discussed. The Group agreed to define a procedure to negotiate data link parameters, but thought that work on automatic configuration should be done within the Network Management Area.

Issues for the next meeting:

- Continued work on issues concerning multiple IP subnetworks on the same PDN.

- Specify a mechanism to identify the encapsulate over circuit ISDN (acquisition of a code point to identify PPP).

- Review and approval of the draft on IP over X.25.

- Specify parameter negotiation over FR PVCs.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Mary Artibee	artibee@sgi.com
William Babson	bill%penril@uunet.UU.NET
Fred Baker	fbaker@emerald.acc.com
William Barns	barns@gateway.mitre.org
Art Berggreen	art@acc.com
David Bolen	db31@nis.ans.net
George Clapp	clapp@ameris.center.il.ameritech.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Kenneth Crepea	crepea@cisco.com
Kent England	kwe@bbn.com
Hans Eriksson	hans@sics.se
Craig Fox	foxcj@network.com
John Garrett	jwg@edsel.att.com
Jim Ghadbane	jimgh@newbridge.com
Herve Goguely	lircorp1!roger@uunet.uu.net
Chris Gunner	gunner@osicwg.enet.dec.com
John Hagan	hagan@dccs.upenn.edu
Kenneth Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Robert Hinden	hinden@bbn.com
William Jolitz	william@okeeffe.cs.berkeley.edu
Dan Jordt	danj@nwnet.net
David Kaufman	dek@proteon.com
Tom Kessler	kessler@sun.com
Cheryl Krupczak	cheryl@cc.gatech.edu
Joseph Lawrence	jcl@sabre.bellcore.com
Eliot Lear	lear@sgi.com
Tony Li	tli@cisco.com
Andrew Malis	malis@bbn.com
Gary Malkin	gmalkin@ftp.com
Mike Marcinkevicz	mdm@csu.net
Keith McCloghrie	kzm@hls.com
David McDysan	0002806303@mcimail.com
Dean Morris	morris@marvin.dec.com
Chandy Nilakantan	csn@3com.com
Rebecca Nitzan	nitzan@es.net
David O'Leary	oleary@sura.net
Ari Ollikainen	ari@es.net

Zbigniew Opalka	zopalka@roswell.spartacus.com
J. Bradford Parker	brad@cayman.com
Radia Perlman	perlman@radia.enet.dec.com
David Piscitello	dave@sabre.bellcore.com
Ron Roberts	roberts@jessica.stanford.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Philippe Roger	lircorp1!roger@uunet.uu.net
Allan Rubens	acr@merit.edu
Sharad Sanghi	sharad@ans.net
Miguel Sasson	sasson@xylogics.com
Mark Schaefer	schaefer@davidsys.com
Stephen Shew	sdshew@bnr.ca
Mike Spengler	mks@msc.edu
Ravi Srinivasan	ravi@eng.vitalink.com
Michael St. Johns	stjohns@umd5.umd.edu
Brad Steinka	brad@python.eng.microcom.com
Sally Tarquinio	sally@gateway.mitre.org
Dean Throop	throop@dg-rtp.dg.com
Mike Truskowski	truskowski@cisco.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Andrew Veitch	aveitch@bbn.com
Yuan Wang	natadm!ycw@uunet.uu.net
Scott Wasson	sgw@sgw.xyplex.com
Robert Woodburn	woody@cseic.saic.com
Richard Woundy	rwoundy@ibm.com
L. Michele Wright	uncng!michele@uunet.uu.net
Jessica Yu	jyy@merit.edu
Paul Zawanda	zawanda@ncsa.uiuc.edu

Directed ARP

John W. Garrett
AT&T Bell Laboratories
jwg@edsel.att.com

John Dotts Hagan
University of Pennsylvania
Hagan@UPENN.EDU

Jeff Wong
AT&T Bell Laboratories
jaw@io.att.com

Overview

Directed ARP is a procedure to route address resolution information among different IP (sub-)networks on the same link level network.

Within an IP (sub-)network, the mechanism used to resolve IP address to link level addresses is the responsiblility of the administrative authority that manages the IP (sub-)network. Directed ARP provides a way for a host in one IP (sub-)network to resolve an IP address in another IP (sub-)network.

Directed ARP may be used by routing to test whether a potential next-hop is local.

The Problem

Currently, IP addresses are assigned to indicate administrative boundries <u>and</u> to indicate the scope of link level connectivity. What do you do when the link level connectivity spans multiple administrative domains, as is likely in a Public Data Network?

* If we put everyone on the same IP (sub-)network, we have to all agree how to run it.
* How do we adapt to changing link level connectivity in the evolving Public Data Network? Ongoing reassigment of IP addresses?

<u>SOLUTION:</u> Use Directed ARP to determine the scope of lthe ink level connectivity.

Directed ARP Procedure

An ARPrequestaddr is associated with each route table entry, and is the link level address of the source of that route table entry.

```
ReceiveARP (ARPrequest, targetIPaddr, fromInterface)
{
  if targetIPaddr = my_IPaddr
    RespondToARP(ARPrequest)
  else
    {
    if (route=GetRouteTo(targetIPaddr)) != NULL
      {
      if NextHopInt(route) = fromInterface
         AND IsLocal(route)
        {
        if (ARPrequestaddr=
           AssociatedARPrequestaddr(route)) != NULL
          {
          ForwardARP(ARPrequest, ARPrequestaddr)
          }
        else
          {
          if (targetLLaddr=
             ARPtablelookup(targetIPaddr)) != NULL
            {
            ForwardARP(ARPrequest,targetLLaddr)
            }
          else
            Discard(ARPrequest)
          }
        }
      else
        Discard(ARPrequest)
      }
    else
      Discard(ARPrequest)
  }
}
```

Routing

Routing discovers new and better next-hops to IP (sub-)networks, but must make sure an advertised next-hop is local before entering it in the routing table:

* ARP for proposed next-hop to test if local

To use routing to discover link level connectivity, a router must determine a next-hop is local to a routing peer, and convey that information:

* A next-hop may be advertised as local to a routing peer if it is reached through the same interface.

Routers can tell hosts of new link level connectivity by sending an ICMP Host Redirect:

* A next-hop may be advertised as local to a host if the host is local and reached through the same interface.

Example: SMDS

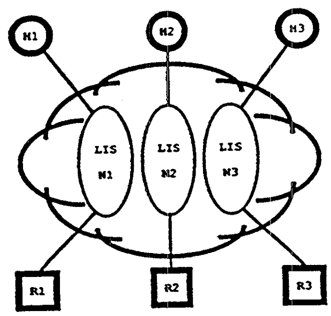

	DESTINATION	NEXT-HOP	INTERFACE	ARP REQUEST ADDRESS
H1:	N1	local	smds	G1
	default	R1	smds	G1
H2:	N2	local	smds	G2
	default	R2	smds	G2
H3:	N3	local	smds	G3
	default	R3	smds	G3
R1:	N1	local	smds	G1
	N2	local	smds	smds (R2)
	N3	local	smds	smds (R3)
R2:	N1	local	smds	smds (R1)
	N2	local	smds	G2
	N3	local	smds	smds (R3)
R3:	N1	local	smds	smds (R1)
	N2	local	smds	smds (R2)
	N3	local	smds	G3

Example: X.25

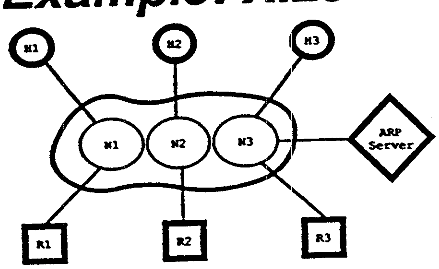

Address resolution on N1 is achieved by administering tables
Address resolution on N2 is algorithmic
Address resolution on N3 is achived using Directed ARP to an ARP Server

	DESTINATION	NEXT-HOP	INTERFACE	ARP REQUEST ADDRESS
H1:	N1	local	x.25	NULL
	default	R1	x.25	NULL
H2:	N2	local	x.25	NULL
	default	R2	x.25	NULL
H3:	N3	local	x.25	ARP Server
	default	R3	x.25	ARP Server
R1:	N1	local	x.25	NULL
	N2	local	x.25	x.25(R2)
	N3	local	x.25	x.25(R3)
R2:	N1	local	x.25	x.25(R1)
	N2	local	x.25	NULL
	N3	local	x.25	x.25(R3)
R3:	N1	local	x.25	x.25(R1)
	N2	local	x.25	x.25(R2)
	N3	local	x.25	ARP Server

Observations

Directed ARP is a tool:

* Extends address resolution across administrative bounderies
* A test to determine if a next-hop is local

Utilizing link level connectivity across administrative bounderies is a routing problem:

* Need a way to test if a next-hop is local even if it is across administrative boundries
* Directed ARP performs that test

Directed ARP is easy to deploy:

* Directed ARP works with targets that only implement ARP

3.6.3 ISIS for IP Internets (isis)

Charter

Chair(s):
Ross Callon, `callon@bigfut.enet.dec.com`

Mailing Lists:
General Discussion: `isis@merit.edu`
To Subscribe: `isis-request@merit.edu`
Archive:

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.

Internet Drafts:

"Integrated IS-IS Management Information Base", 11/05/1991, Chris Gunner <draft-ietf-isis-mib-00.txt>

Request For Comments:

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

CURRENT MEETING REPORT

Reported by Ross Callon/DEC

Minutes of the ISIS for IP Internets Working Group (ISIS)

Agenda

1. Update to ISIS MIB (Chris Gunner)
2. Update to RFC 1195 (Ross Callon)
3. ISIS-BGP Interactions (Yakov Rekhter)

Update to ISIS MIB Draft

Chris Gunner presented an updated draft of ISIS MIB. The major changes from the previous draft were:

- The number of tables were reduced from the previous version - now there are roughly half as many tables as before.

- No restriction on Set PDU's contents in the MIB specification. An agent, however, can impose one on the Set PDU's contents.

There was a suggestion to link the IP Destination Object and the IP Forwarding Table.

Additional detail-level reviewers of the ISIS MIB would be appreciated. It is expected that this will occur as the MIB is implemented. The IS-IS MIB is currently an Internet Draft.

Update to RFC 1195

Ross Callon presented an updated version of RFC 1195. The changes to RFC 1195 are listed in Section Six of the new draft (which has been distributed to the Working Group).

Following changes and topics were discussed:

- Reference to final International Standard of ISO 10589. This is the biggest change to the draft. This allows several sections of RFC 1195 to be removed as they are redundant with corrections and improvements that have been made to ISO 10589. For example, Annexes on encoding of sequence number packets and on authentication are now redundant with ISO 10589.

- The specification now allows announcement of the IP Router ID over unnumbered links. This is needed for Strict Source Routing, network management, and for locally originated IP packets over unnumbered links. The specification will be updated to specify that for routers which have only unnumbered links, the router ID must be announced in the LSP's as a Host Route. The specification should probably also include a brief description of what a router ID is.

- What should be done when a router is an L1 and L2 router, doing RIP, but L2 is not IP capable? The specification now describes this in some detail, but some editorial clarification is needed (see the "mixed operation" section of the update to RFC 1195).

- NSAP address for IP-only routers was discussed. There are several ways in which these can be obtained. This is currently being pursued in several other places for uses which include, but which go well beyond use in IS-IS. Therefore this should be removed from the update to RFC 1195.

- There was a discussion of how to transition two instances of IS-IS which are operating in "Ships in the Night" mode to a single instance of Integrated IS-IS. Ross proposed two possible transition methods, one of which was well received and the other of which was quickly rejected. Implementations will not be required to be able to run two instances of IS-IS in this manner. However, if an implementation does implement the capability of running two versions of IS-IS in SIN mode, then the implementation must also implement the controls needed to be able to transition from SIN mode to integrated mode and vice versa.

- An optimization of when to leak routes from L1 to L2 was discussed and approved. This would optionally allow selective leaking of routes from level 1 to level 2 LSPs, in a manner which does not effect routes (except for an improvement in routes in one obscure case) but which would reduce the amount of information in level 2 LSPs, at the cost of slightly more work for the routers doing the route leaking. This feature would work well even when implemented by only some routers, and therefore can be optionally implemented and deployed.

- There was a discussion of redundant manually configured summary routes. It was agreed that this issue was not particularly important, but that the specification should be complete and unambiguous. The decision was that when redundant summary addresses are manually configured, both are announced.

- Dino Farinacci suggested that we can use the LSP protocols supported field to avoid creating a black hole when all routers within an area are not the same type (all OSI, all IP, or all Dual). Again this was a feature which will work well even when implemented by only some routers in an area (routers which do not implement this will interwork with those that do). This proposal was accepted.

Ross agreed to update the specification based on this discussion, and to have this issued as an Internet Draft when available.

BGP - ISIS Interaction

Yakov Rekhter presented the issues of interaction between BGP and IS-IS. After the discussion, Sharad Sanghi of ANS and Atul Bansal volunteered to write the BGP-IS-IS draft.

Leakage of routes between BGP and IS-IS was discussed, and it was agreed that this should be the same as in the OSPF-BGP case.

The relationship between BGP router IDs and IS-IS router IDs was discussed.

Piggybacking of BGP information in IS-IS packets was discussed. In those cases where all or most level 2 routers are border routers running BGP, this makes sense (IS-IS solves the n-square BGP link problem, and provides reliable multicast mechanism). However, in those cases where very few level 2 routers are border routers, the n-square link problem is not significant, and piggybacking requires non-border routers to store BGP information. It was therefore agreed that whether to piggyback BGP information on IS-IS packets or to run internal BGP will depend upon the network environment, and therefore both possibilities should be allowed. If a network has very few BGP speakers then I-BGP is a good solution. If a network has lots of BGP speakers and very few non-BGP speaking L2 routers then Piggybacking is most efficient.

Auto-configuration of I-BGP neighbors was also discussed. Auto I-BGP configuration optimization was suggested as an efficient mechanism for discovering I-BGP neighbors. This feature eliminates the nightmare - manual configuration of I-BGP neighbors. This auto-configuration can be piggybacked on IS-IS.

Tagging is currently defined by RFC 1195. This should continue to be available.

We also discussed how to pass BGP information between two I-BGP neighbors when one is doing OSPF and the other is doing ISIS? This required close cooperation with the folks working on BGP-OSPF interactions.

Attendees

Atul Bansal	bansal@wile.nac.dec.com
Robert Blokzijl	K13@nikhef.nl
Daniel Blum	4108980@mcimail.com
David Bolen	db31@nis.ans.net
Ross Callon	callon@bigfut.enet.dec.com
Ken Carlberg	carlberg@cseic.saic.com
Dean Cheng	dean@sunz.retix.com
Chi Chu	chi@sparta.com
Kurt Dobbins	dobbins@ctron.com
Dino Farinacci	dino@cisco.com
Chris Gunner	gunner@osicwg.enet.dec.com
Christine Hemrick	hemrick@cisco.com
Tony Li	tli@cisco.com
April Merrill	abmerri@tycho.ncsc.mil
Dave Monachello	dave@pluto.dss.com
Dennis Morris	morrisd@imo-uvax.dca.mil
Yakov Rekhter	yakov@watson.ibm.com

Manoel Rodrigues	manoel.rodrigues@att.com
Sharad Sanghi	sharad@ans.net
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Ravi Srinivasan	ravi@eng.vitalink.com
Andrew Veitch	aveitch@bbn.com
Cathy Wittbrodt	cjw@nersc.gov
Osmund de Souza	osmund.desouza@att.com

3.6.4 Inter-Domain Policy Routing (idpr)

<u>Charter</u>

Chair(s):

Martha Steenstrup, msteenst@bbn.com

Mailing Lists:

General Discussion: idpr-wg@bbn.com
To Subscribe: idpr-wg-request@bbn.com
Archive:

Description of Working Group:

The Inter Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Ongoing	Gain experience with the prototype in "real networks".
TBD	Develop gated version.
TBD	Add a small set of additional features and submit protocol into IETF standards process.

Internet Drafts:

"An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-idpr-architecture-03.txt>

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

"Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)", 07/22/1991, R.A. Woodburn <draft-ietf-idpr-mib-00.txt, .ps>

"Inter-Domain Policy Routing Configuration and Usage", 07/25/1991, H. Brown, M. Steenstrup <draft-ietf-idpr-configuration-00.txt>

Request For Comments:

RFC 1126 "Goals and functional requirements for inter-autonomous system routing"

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

Minutes of the Inter-Domain Policy Routing Working Group (IDPR)

The IDPR Working Group met for three sessions during the Santa Fe IETF meeting. Currently, the top priorities are to complete implementation of the gated version of IDPR and to submit the IDPR protocols as a Proposed Standard.

Our first session was designed to elicit active participation in experimenting with the forthcoming gated version of IDPR. We began with a short review of the IDPR approach to policy routing followed by a status report of work in progress. However, we spent the majority of the session describing how IDPR will fit into the current Internet and how to configure one's networks to take advantage of IDPR.

Our second session was intended for those interested in actual IDPR implementations. Woody Woodburn, who is leading the development effort for the gated version of IDPR gave an overview of the software architecture as well as a report on the status of the implementation thus far.

During our third session, we met jointly with the Router Requirements Working Group. The main purpose of the meeting was to determine if IDPR flagrantly violated any of the current router requirements. Our concerns centered around interactions of IDPR with IP, in particular with source routing, TTL, and trace route. We also discussed IDPR's expectations of the intra-domain routing procedures, in particular that intra-domain routes remain within a domain. To our relief, the meeting turned up no major incompatibilities of IDPR with router requirements. However, we encourage attendees to think carefully about the issues and bring forth any problems they discover.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Nagaraj Arunkumar	nak@3com.com
William Babson	bill%penril@uunet.UU.NET
Daniel Blum	4108980@mcimail.com@mcimail.com
David Bolen	db31@nis.ans.net
David Bridgham	dab@asylum.sf.ca.us
Scott Brim	swb@nr-tech.cit.cornell.edu
Ken Carlberg	carlberg@cseic.saic.com
Dean Cheng	dean@sunz.retix.com
Chi Chu	chi@sparta.com
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Barbara Denny	denny@sri.com
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Craig Fox	foxcj@network.com

Jim Ghadbane	jimgh@newbridge.com
Kathleen Huber	khuber@bbn.com
Bob Jeckell	robert_jeckell@nso.3com.com
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Scott Kaplan	scott@ftp.com
Michael Karels	karels@cs.berkeley.edu
David Kaufman	dek@proteon.com
John Lekashman	lekash@nas.nasa.gov
Mark Lewis	mlewis@telebit.com
Brian Lloyd	brian@ray.lloyd.com
Donald Merritt	don@brl.mil
Paul Mockapetris	pvm@darpa.mil
Dave Monachello	dave@pluto.dss.com
Dennis Morris	morrisd@imo-uvax.dca.mil
Radia Perlman	perlman@radia.enet.dec.com
Manoel Rodrigues	manoel.rodrigues@att.com
Sharad Sanghi	sharad@ans.net
Miguel Sasson	sasson@xylogics.com
John Seligson	johns@ultra.com
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Stephen Shew	sdshew@bnr.ca
Robert Shirey	shirey@mitre.org
Frank Solensky	solensky@clearpoint.com
Ravi Srinivasan	ravi@eng.vitalink.com
Michael St. Johns	stjohns@umd5.umd.edu
Martha Steenstrup	msteenst@bbn.com
Osamu Takada	takada@sdl.hitachi.co.jp
Iris Tal	437-3580@mcimail.com
Scott Wasson	sgw@sgw.xyplex.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Robert Woodburn	woody@cseic.saic.com
June-Kang Yang	natadm!yang@uunet.uu.net
Jessica Yu	jyy@merit.edu
Paul Zawanda	zawanda@ncsa.uiuc.edu

3.6.5 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):
Steve Deering, deering@xerox.com

Mailing Lists:
General Discussion: mospf@devvax.tn.cornell.edu
To Subscribe: mospf-request@devvax.tn.cornell.edu
Archive:

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.

Done 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.

Done 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.

Internet Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-01.ps>

CURRENT MEETING REPORT

Reported by Scott Brim/Cornell University

Minutes of the Multicast Extensions to OSPF Working Group (MOSPF)

1. **Agenda**
   - Roster, Introductions, Notetaker
   - Reports on Related Activities
     - X3S3.3?
     - BGP?
     - Router Requirements?
   - Review of Latest MOSPF Draft
     - Scott's Comments
     - Forwarding Algorithm
     - Extent of Reverse Paths
     - Inter-AS Interactions
     - "Host" Behavior of MOSPF Routers
   - Token Ring Address Mapping
   - Multicast Scope Proposal
   - Implementation Status
   - Future Work
     - MIB
     - Standards Track
     - Field Tests/Interoperability Tests

2. **Reports on Related Activities**

   X3S3.3: Recently started work on "advanced services," including multicast. Steve Deering addressed them on multicasting model and MOSPF. Dave Marlowe has draft extension to CLNS for join & leave group, proposal for NSAP assignment; other stuff. Nobody knows what happened at the last meeting in Boston.

   BGP: Internet Draft issued on alternative approaches. Only one person signed up to implement (Scott Brim) and he's not going to do it until after he finishes MOSPF. Not this year.

   Router Requirements: Multicast will not be in the forwarding MIB because ??? it uses source address and nothing else in there does right now. They're going to wait. Someone is going to have to write a multicast routing MIB in addition to the different MC protocol MIBs. John Moy contributed a section on multicasting router requirements which will have to be revised, and soon.

3. Review of Latest MOSPF Draft

3.1. Discussion of Scott Brim's Previously Emailed Comments

How do you tell what the previous hop of a packet was? You can't without looking at the previous hop's link-level address. The issue is that on a border network you need to determine whether a packet you receive was being multicast *through* your autonomous system or just *getting* to your autonomous system? (See figure in draft.) That's why we came up with datalink unicasting. It looks like this area of interactions between OSPF and RPF protocols isn't finished. Datalink unicasting is a start, but doesn't cover everything. We're going to have to study this one more. Do we have to encapsulate when crossing an AS boundary? Right now the BGP model is straight RPF, and RPF has no idea what an AS boundary is.

Perhaps if there's a host sitting on the border LAN, then you only accept unicasts *unless* the packet originates on that LAN. Datalink unicasting is for transits, not for locally-originated traffic. What if someone is *sending* to a group that host belongs to?

In BGP, only one AS announces the shared net. Should we combine the flags that say you shouldn't listen to multicasts with the one that says to do datalink unicasts?

One definite conclusion is that you shouldn't base *forwarding* on whether something came from another AS. In *building* the FIB that's important, but should not be used after that.

Caching negative results is already in the document.

What if a vertex is not labelled? Yes, document needs a statement saying go to the next section.

Yes, there should be justifications for why we did *not* do things in some way.

3.2. Forwarding Algorithm

Moy: Spell out preprocessing. When called (directly from IP forwarding), first check:

- If 224.0.0.1 to 224.0.0.255, never forwarded, only sent to internal applications.
- If IGMP message, send to IGMP process, don't forward.
- Then follow rest of section.

Internally generated multicast packets must be handled differently – in John's design at least. This is *not* true in Steve's design, and a significant amount of time was spent comparing them.

Steve: host specification (RFC 1112) says group membership is associated with an interface. Forwarding sends to a set of outgoing interfaces. As *part* of forwarding to an interface, in the per interface code, if this host is a member of the destination group *on that interface*, this host receives a copy, not by interface loopback but in memory. An application which joins on multiple interfaces receives multiple copies. Also, if this host *sends* a packet, if this is a forwarder, the packet is looped back in the interface handler for the interface the packet is being sent on, and given to the forwarder which forwards it as necessary as if it *came* from that interface. A multicast packet, when it hits the forwarder, is always associated with an interface. The forwarding function is thus relatively pure.

John says if you're only doing MOSPF, membership can be associated with the router, not with a particular interface, so local delivery is hung off the packet forwarder. If originated locally, a packet goes directly to the forwarding process, which knows which interface you want to forward it out of, and decides whether to deliver it locally. If an interface goes down, with the Deering scheme, then the application has to rejoin on another interface or it doesn't receive any traffic. Steve's model is necessary for a multihomed host; John's is possible on an MOSPF router because of its complete knowledge of the topology. However, the programmer's interface shouldn't change depending on whether MOSPF is running or not, so maybe you should still do it with interfaces.

The time to join on more than one interface is, for example, when you are doing an expanding ring search, and you want to get a hit on any interface. Also, Steve's model gives you the possibility of making sure you only receive a packet for a particular group on *one* *particular* interface. John's model has the *router* being a member, on *any* interface, so the router as a whole gets a copy of a packet. Steve was forced into his approach to make multihomed hosts work. If we allow both models, then yes, the environment does change for applications – applications can't receive multiple copies with John's approach. An artifact of Steve's approach is that the packet goes out on the intended interface with the intended TTL, and goes out on other interfaces (if it needs forwarding immediately) with the TTL one less. Steve's gut reaction is that applications won't care if they don't get multiple copies, but he doesn't know for sure. John *can* emulate all of Steve's behavior, delivering duplicate packets – but would it be better if he didn't.

3.3. Extent of Reverse Paths

Within the area where the source is, you use forward costs. Everywhere else you use all reverse costs. If you don't use *all* links in the *reverse* direction, you get pockets of non-delivery of datagrams. The problem occurs when you have asymmetric reachability or costs on links within a receiving area. Steve thinks this is a problem due to the way John stores his information and due to his decision that a multicast routing table entry is simply a pointer to a unicast entry and a group address. Steve thinks the information for using forward costs is there, but not used. This discussion was not really concluded.

3.4. Inter-AS Interactions

Covered already in the section on "Scott's comments".

3.5. "Host" Behavior of MOSPF Routers

Covered already in the section on "Forwarding algorithm".

4. Token Ring Address Mapping

A functional address for carrying IP multicasts on token ring has not yet been obtained. Steve could write a one-page RFC on how to use it if he only had the address. Coltun will follow up on it.

5. Multicast Scope Proposal

Steve's proposal reviewed. (1) local wire, already allocated as 224.0.0.1,255; (2) site-wide – start allocating from the bottom up at 224.0.1.0; (3) global, allocated from 249.255.255.255 downward. Thus we can decide about the middle later. This would require the number czar to ask multicast group requestors just what they are going to be used for and make an intelligent allocation based on what they say – this might not be acceptable.

6. Implementation Status

Not covered.

7. Future Work

7.1. MIB

No volunteers came forward.

7.2. Standards Track

Not covered.

7.3. Field Tests / Interoperability Tests

Not covered, except to say that we should try to line up some test beds.

Attendees

Nagaraj Arunkumar	nak@3com.com
William Babson	bill%penril@uunet.UU.NET
Atul Bansal	bansal@wile.nac.dec.com
James Beers	beers@nr-tech.cit.cornell.edu
Scott Brim	swb@nr-tech.cit.cornell.edu
Yee-Hsiang Chang	yhc@concert.net
Dean Cheng	dean@sunz.retix.com
Rob Coltun	rcoltun@ni.umd.edu
Steve Deering	deering@xerox.com
Kurt Dobbins	dobbins@ctron.com
Dino Farinacci	dino@cisco.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Jim Ghadbane	jimgh@newbridge.com
Christine Hemrick	hemrick@cisco.com
Ronald Jacoby	rj@sgi.com
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Stev Knowles	stev@ftp.com
Mark Lewis	mlewis@telebit.com
April Merrill	abmerri@tycho.ncsc.mil
Greg Minshall	minshall@wc.novell.com
Dave Monachello	dave@pluto.dss.com
Dean Morris	morris@marvin.dec.com
John Moy	jmoy@proteon.com
Andy Nicholson	droid@cray.com
Stephen Shew	sdshew@bnr.ca
Ravi Srinivasan	ravi@eng.vitalink.com
Michael St. Johns	stjohns@umd5.umd.edu
William Townsend	townsend@xylogics.com
Scott Wasson	sgw@sgw.xyplex.com
L. Michele Wright	uncng!michele@uunet.uu.net

MULTICAST EXTENSIONS TO OSPF

Agenda

- Roster, Intros, Notetaker
- Reports on Related Activities
  - X353.3?
  - BGP?
  - Router Requirements?
- Review of Latest MOSPF Draft
  - Scott's Comments
  - Forwarding Algorithm
  - Extent of Reverse Paths
  - Inter-AS interactions
  - "HOST" behavior of MOSPF Routers
- Token Ring Address Mapping
- Multicast Scope Proposal
- Implementation Status
- Future Work
  - MIB
  - Standard Track

Mods to Section 8, "Detailed...forward"

- Called directly from IP forwarding process, first check:
  - 224.0.0.1-224.0.25.5 → appl. only
  - IGMP → IGMP process
  - Then rest of section
- Must also handle internally generated multicast datagrams differently
- Need a way to incorporate other multicast routing protocols into forwarding process
  - All use same forwarding cache?

Proposal for Scoped Multicast Addresses

- Define 3 levels for now: Subnet, Site, Global

  Subnet: 224.0.0.0 - 224.0.0.255 (same as now)
  Site: Allocate from 224.0.1.0, Upwards
  Global: Allocate from 239.255.255.255, Downwards

- Pick Half-way point (232.0.0.0) as dividing line, for now.

  => Currently assigned groups become site-wide only (not too disruptive, will be worse if delayed)

  => May Change boundary or add levels in the future - only boundary routers need know

- Complicates job of IANA.

3.6.6 Open Shortest Path First IGP (ospf)

Charter

Chair(s):
Mike Petry, petry@ni.umd.edu
John Moy, jmoy@proteon.com

Mailing Lists:
General Discussion: ospfigp@trantor.umd.edu
To Subscribe: ospfigp-request@trantor.umd.edu
Archive:

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

Done	Design the routing protocol, and write its specification.
Done	Develop multiple implementations, and test against each other.
Done	Obtain performance data for the protocol.
Done	Make changes to the specification (if necessary) and publish the protocol as a Draft Standard RFC.
TBD	Gather operational experience with the OSPF protocol and submit the document as a Standard.

Internet Drafts:

"OSPF Version 2 Traps", 07/23/1991, Rob Coltun <draft-ietf-ospf-trapmib-00.txt>

Request For Comments:

RFC 1131	"OSPF specification"
RFC 1245	"OSPF Protocol Analysis"
RFC 1246	"Experience with the OSPF Protocol"

RFC 1247	"OSPF Version 2"
RFC 1248	"OSPF Version 2 Management Information Base"
RFC 1252	"OSPF Version 2 Management Information Base"
RFC 1253	"OSPF Version 2 Management Information Base"

CURRENT MEETING REPORT

Reported by John Moy/Proteon

Minutes of the Open Shortest Path First IGP Working Group (OSPF)

The OSPF Working Group met Tuesday November 19, 1991 at the Santa Fe IETF. In addition, at this IETF work was performed (and decisions were made) in other working groups affecting OSPF. This related work is summarized below in the liaison section.

1. Liason with Other Working Groups

- In the Open IESG meeting, it was announced that the IAB had approved the OSPF Applicability Statement, which recommends the use of OSPF as the Common Inter-domain Gateway Protocol (IGP). It is expected that the Applicability statement will be published as an RFC.

- The wording of the router requirements document now reads: "if a router implements dynamic routing, it must implement OSPF" as an aside, it also must implement RIP. Router requirements has also made Type of Service (TOS) in OSPF optional (this was part of a more general discussion of whether further subsets of OSPF are possible and/or useful, which was continued at the OSPF Working Group meeting; see Section 2 (e) below). The Router Requirements Working Group has also asked that the behavior of OSPF in the face of database overflows be written down. Finally, an IP Forwarding Table MIB has been defined allowing network management stations to dump equal-cost routes, and routes that depend on TOS (both of which are possible with OSPF).

- The BGP Working Group has been working on a document specifying the interaction between BGP and OSPF. A first draft of this document, written by Kannan Varadhan of OARnet, had been published as an Internet Draft before the Santa Fe IETF. At the IETF the sections describing route exchange, the setting of BGP IDs and OSPF Router IDs, and the setting of the BGP NEXT_HOP attribute and the OSPF forwarding address were pretty much agreed upon. The setting of the tag field in type 5 AS external LSAs was more controversial, and several different proposals were floated. An updated Internet Draft should appear shortly.

2. Working Group Minutes

The following items were discussed in the Working Group session. All items on the Agenda were covered, except for a planned discussion of OSPF's non-broadcast network support (which is a hot topic currently because of all the activity in the IPLPDN Group).

(a) A Problematic Virtual Link Configuration

A handout was provided describing a configuration of virtual links that was found to create routing loops. This configuration was discovered during the last round of OSPF testing, immediately prior to INTEROP '91. Basically, the problem arises because, in a virtual link's transit area, the area border routers may have a different view of routing than the area's internal routers. The current OSPF specification tries to deal with this by having the endpoints of a virtual link run an extra computation: the "resolution of virtual next hops" described in Section 13.3 of the specification.

However, this is not enough to avoid loops in all configurations, as the handout showed. The handout also presented a fix to the specification, whereby any router bordering transit areas would a) keep track of all transit areas that are traversed en route to any particular destination and b) for such a destination, run the "resolution of virtual next hops" using summary links belonging to each of the traversed areas.

It was generally felt that the handout's fix was too complicated. An alternative fix, involving less bookkeeping while potentially running the "resolution of virtual next hops" process on more destinations, was proposed. This simpler fix is being investigated.

The handout, augmented with a discussion of the simpler fix, will be published as an Internet Draft. Eventually, a new (but backward-compatible) version of the OSPF specification will have to be published. Besides having a fix for the virtual link problem, it was proposed to at that time add the following: a) make the origination of summary-LSAs into stub areas optional and b) add text describing how to avoid originating summary-LSAs into an area when you know that they will never be used (i.e., when the first hop for the destination belongs to the area itself; this is sort of equivalent to split horizon in a Bellman-Ford algorithm).

(b) Proposed Changes to the OSPF MIB

The following changes to the OSPF MIB were proposed. It is the intent that all these changes be backward-compatible with the present MIB:

- Change the range of the ospfIfRtrDeadInterval, ospfIfPollInterval and ospfVirtIfRtrDeadInterval variables from 0-0xffffffff to 0-0x7fffffff. This is being done to make life easier for MIB compilers, realizing that it doesn't really make any sense to set the variables higher than 0x7fffffff anyway.

- Remove the TOSType definition from the OSPF MIB, and instead refer to a TOS definition in the new IP Forwarding Table MIB.

- Add a separate table for type 5 AS externals, removing them from the current ospfLsdbTable. At the moment it is not clear just where in the ospfLsdbTable the type 5 AS externals should go.

- Add type 6 (group-membership-LSAs) and type 7 (the new NSSA externals) LS types to the ospfLsdbTable. This will allow us to monitor the OSPF extensions (somewhat) from the base OSPF MIB.

- Add a boolean to the Area Table allowing you to turn on or off the origination of summary-LSAs into stub areas.

- Somehow figure out how to represent OSPF type 1 and type 2 metrics, and also the four level OSPF routing hierarchy (intra-area, inter-area, type 1 external and type 2 external) in the new IP Forwarding Table MIB. This may be done entirely with comments.

There was an additional proposal on the table to clean up/rationalize the ASCII names of some of the OSPF MIB variables. It was decided to ask the Network Management Directorate whether this would be too large a change to make at this time.

(c) The OSPF Trap MIB

Rob Coltun reported on the state of the OSPF Trap MIB. There are currently twelve traps: Interface state change (regular and virtual), Neighbor state change (regular and virtual), Configuration error (over real and virtual links), Receive bad packet (over regular and virtual links), Packet retransmission (over regular and virtual links), Originate LSA and MaxAge LSA. Each trap can be enabled and disabled separately. Trap origination is rate-limited, and traps are inhibited for the first 2*DeadInterval seconds after a router comes up.

It was decided to add two more traps. The first indicates that the link state database has overflowed. The second indicates that the link state database is close to overflowing, because available resources have dropped below some configurable threshold (units of the threshold being number of LSAs).

After making these additions, the document will be published as an Internet Draft.

(d) Current Proposal for OSPF Not so stubby area (NSSA) Areas

Rob Coltun presented the current proposal for OSPF NSSA areas. His view-graphs will appear in the IETF Proceedings. A brief summary of his presentation follows:

- An NSSA area is a new kind of area which does not process type 5 external LSAs (reducing memory resource requirements) but which can originate external routing information and pass it on to the rest of the OSPF system. For example, an NSSA area can be used where you wanted to use an OSPF stub area, but couldn't because hanging off of the area was a RIP cloud.

- External routes are originated into an NSSA area via a new link state advertisement: type 7 LSAs. The format of type 7 LSAs are identical to type 5 LSAs. However, type 7s are specific to a single NSSA area only. There will be a propagate bit in the type 7 LSA's Options field which indicates whether the type 7 LSA should be translated into a type 5 LSA at the NSSA border. Those type 7 LSAs which are to be translated MUST specify a forwarding address (this makes translation into type 5 LSAs simple, and also enables a simple already specified tie-breaking mechanism ensuring that only one border router does the translation).

- Area border routers attached to NSSAs originate a type 7 LSA specifying the default route (with the propagate bit off) into the NSSA. This compensates for the fact that type 5 LSAs are not flooded into NSSAs. Also, to maintain the OSPF routing hierarchy area border routers attached to NSSAs must summarize the internal (intra-area and inter-area) OSPF routes into the NSSA (for OSPF stub areas this summarization is optional).

Several other possible NSSA features were discussed, namely: a) allowing type 7 information to be collapsed (instead of directly translated) at NSSA boundaries and b) allowing selective reverse translation at NSSA boundaries (i.e., type 5 LSAs into type 7 LSAs for propagation into the NSSA). It was decided to leave both features outside the scope of the NSSA option.

(e) Defining a Minimal Subset of OSPF

We spent some time discussing whether it was useful to subset OSPF beyond simply making TOS optional. It was generally agreed that this would probably not be a commercially viable product, since the router would be limited to only certain places in the topology. However, it did appear that it might be viable for those products that naturally reside at the edge of the IP routing domain, for example, the Shiva FastPath box.

3. Action Items

- Revise the OSPF specification with a fix for the virtual link problem [John Moy]

- Revise the OSPF MIB [Fred Baker]
- Publish the OSPF Trap MIB as an Internet Draft [Rob Coltun]
- Document the NSSA option and publish as an Internet Draft [Rob Coltun and Vince Fuller]
- Outline the possibilities for a minimal OSPF implementation [John Moy with help from Shiva and Cayman Systems].
- Document the OSPF response to database overflow [John Moy]

Attendees

Steve Alexander	stevea@i88.isc.com
Fred Baker	fbaker@emerald.acc.com
James Barnes	barnes@xylogics.com
James Beers	beers@nr-tech.cit.cornell.edu
David Bolen	db31@nis.ans.net
Gregory Bruell	gob@shiva.com
Dean Cheng	dean@sunz.retix.com
Kenneth Crepea	crepea@cisco.com
John Damiano	
Kurt Dobbins	dobbins@ctron.com
Christine Hemrick	hemrick@cisco.com
Ronald Jacoby	rj@sgi.com
April Merrill	abmerri@tycho.ncsc.mil
Dean Morris	morris@marvin.dec.com
John Moy	jmoy@proteon.com
Thomas Pusateri	pusateri@cs.duke.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Sharad Sanghi	sharad@ans.net
Stephen Shew	sdshew@bnr.ca
Richard Smith	smiddy@pluto.dss.com
Frank Solensky	solensky@clearpoint.com
Ravi Srinivasan	ravi@eng.vitalink.com
Yuan Wang	natadm!ycw@uunet.uu.net
Scott Wasson	sgw@sgw.xyplex.com
Osmund de Souza	osmund.desouza@att.com

NSSA

External Routes Are Imported As Type 7 LSA

- Identical To Type 5 LSAs
- Only Flooded Within NSSA
- At ABR Type 7 LSA Is Translated To Type 5.

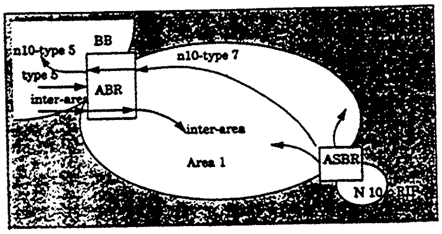

"Don't Propagate" Bit In Type 7 LSA Hdr

- Loop Advoidance

Forwarding Address

- Network Between Autonomous System Is OSPF Net Use Next Hop
- Else Use Any OSPF Interface

Must Import Summary (Inter-Area) Routes Into NSSA

- Summarys Must Be Preferred To Externals
- Optional Summary Import Into Stub Areas (New Option)

3 Cases: Single ABR, Multiple ABR, ABR Is ASBR

- Single ABR Easy Case - ABR Imports Default As Type 7

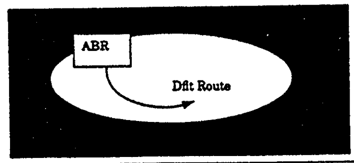

Multiple ABRs

- ABR Choose Who Is To Export As Type 5

1. Identical Route?
2. Lowest Router ID Exports

- Type 7 Dflt Route Has Don't Propagate Bit Set
- Cannot Be Transit Area

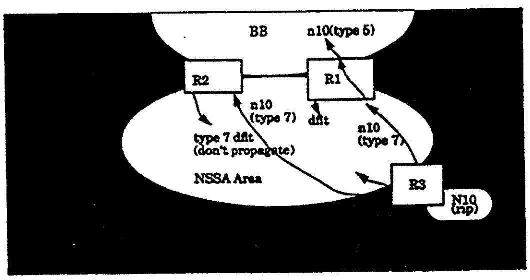

ABR Is ASBR

- Routes Are Converted To Type 5 (for BB) And Type 7 (for NSSA)
- Don't Propagate Bit Set In Hdr

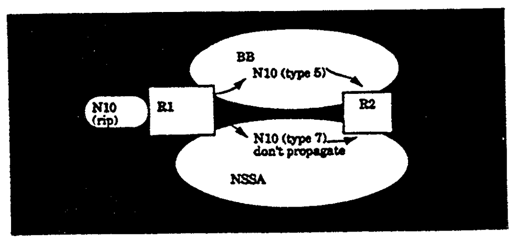

Filters At ASBR?

- Type 7 Collapsed
  - RIP Subnets Summarized
- Type 7 Selectively Propagated
- Type 5 Propagated Into NSSA As Type 7
  - Don't Propagate Bit Set

Support Objects

- ospfSetTrap
  - Bit map to enable/disable traps
- ospfConfigErrorType
  - badVersion
  - areaMismatch
  - unknownNBMANbr -- initiator of trap is DR eligible
  - unnkownVirtualNbr
  - authTypeMismatch
  - authFailure
  - netMaskMismatch
  - helloIntervalMismatch
  - deadIntervalMismatch
  - optionsMismatch
- ospfPacketType

OSPF TRAPS

ospfIfStateChange

- ospfRouterId, ospfIfIpAddress, ospfAddressLessIf, ospfIfState

ospfVirtIfStateChange

- ospfRouterId, ospfVirtAreaId, ospfVirtIfNeighbor, ospfVirtIfState

ospfNbrStateChange

- ospfRouterId, ospfNbrIpAddress, ospfNbrAddressLessIndex, ospfNbrRtrId, ospfNbrState

ospfVirtNbrStateChange

- ospfRouterId, ospfVirtNbrArea, ospfVirtNbrRtrId, ospfVirtNbrState

ospfConfigError

- ospfRouterId, ospfIfIpAddress, ospfAddressLessIf, IpAddress (src) ospfConfigErrorType, ospfPacketType

ospfVirtConfigError

ospfRxBadPacket

- ospfRouterId, ospfIfIpAddress, ospfAddressLessIf, IpAddress (src) ospfPacketType

ospfVirtIfRxBadPacket

ospfTxRetransmit

- ospfRouterId, ospfIfIpAddress, ospfNbrRtrId, ospfPacketType

ospfVirtIfTxRetransmit

ospfOriginateLSA

- ospfRouterId, ospfLsdbAreaId, ospfLsdbType, ospfLsdbLSID, ospfLSDBRouterID

ospfMaxAgeLSA

Multiple OSPFs

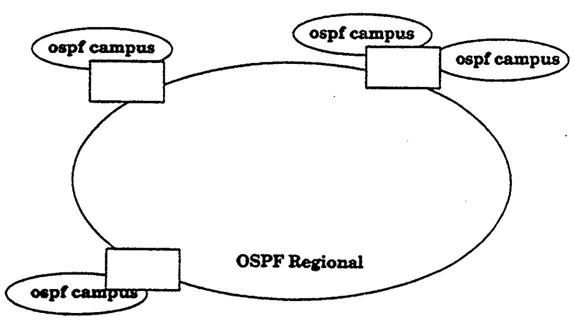

Multiple OSPFs

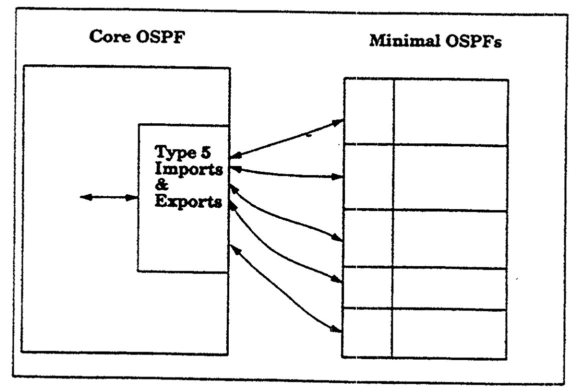

Minimal OSPF

Single Area Only

Not Dr Eligible

Only Request Intra-Area Routes And Externals To Be Propagated Into Regional

Only Add External LSAs And Intra-Area Routes To LSDB

- Only Add Links Associated With Directly Connected Network
- Worst Path To Everything Else

3.6.7 RIP Version II (ripv2)

Charter

Chair(s):
Gary Malkin, gmalkin@ftp.com

Mailing Lists:
General Discussion: ietf-rip@ftp.com
To Subscribe: ietf-rip-request@ftp.com
Archive: gmalkin/rip/archive@vax.ftp.com

Description of Working Group:

The RIPV2 Working Group is chartered to expand the RIP protocol, as defined in RFC 1058. The expansion will include the addition of subnet masks to the routing entries. The expansion may also include authentication, AS numbers, next hop address, MTU, or linkspeed. Since all routing protocols are required to have a MIB, one will be defined. The primary issue is the maintainance of backwards compatibility, which must be preserved.

The purpose of improving RIP is to make a simple, widely available protocol more useful. It is not intended that RIP-II be used in places where OSPF would be far better suited.

Goals and Milestones:

Mar 1991	Review of RIP-II Internet Draft to ensure the additions are useful and backwards compatible. Also ensure that the additions cannot cause routing problems.
Jul 1991	Final review of RIP-II Internet Draft and submission into the standards track. First review of RIP-II MIB.
TBD	Review of implementations. Final review of MIB.
TBD	Given successful implementation experience, advancement of RIP-II to Draft Standard. Submission of MIB into the standards track.
TBD	Final meeting to achieve closure on any pending issues.

Internet Drafts:

"RIP Version 2 Addition of Subnet Masks", 08/14/1991, Gary Malkin <draft-ietf-malkin-rip-01.txt>

CURRENT MEETING REPORT

Reported by Gary Malkin/FTP Software

Minutes of the RIP Version II Working Group (RIPV2)

This meeting was devoted to the review of the Internet Draft. These are the proposed changes. Note that backward compatibility with RIP-I is maintained in all cases.

The 2-byte Must Be Zero (MBZ) field following the family will contain an AS number. Details are TBD (when someone is found who understands IGP/EGP interactions).

Authentication will be supported. If the first RIP entry in the packet has a family of -1, then the two bytes following the family will indicate the type of authentication. Only plain-text password will be defined for now. It will operate as OSPF does and will use the remaining 16 bytes (or fraction thereof).

The remaining MBZ field will be split into 2 2-byte fields: MTU and linkspeed. These will NOT be used for routing; just for information by hosts which snoop RIP packets. The linkspeed will be encoded, in some TBD fashion, since 16 bits is insufficient for LAN speeds.

There are additional issues which must be discussed. They include: subsuming of routes, and route leaking.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Mary Artibee	artibee@sgi.com
Art Berggreen	art@acc.com
Ronald Broersma	ron@nosc.mil
Philip Budne	phil@shiva.com
Dean Cheng	dean@sunz.retix.com
J. Noel Chiappa	jnc@lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Michael Karels	karels@berkeley.edu
Holly Knight	holly@apple.com
Tony Li	tli@cisco.com
Gary Malkin	gmalkin@ftp.com
Frank Solensky	solensky@clearpoint.com
William Townsend	townsend@xylogics.com
Scott Wasson	sgw@sgw.xyplex.com
Jonathan Wenocur	jhw@shiva.com
Jessica Yu	jyy@merit.edu

3.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker

The Security Area within the IETF is responsible for development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet.

Much of the work of the Security Area is performed in coordination with working groups in other areas. The Security Area Advisory Group (SAAG) is a group of security experts which provides both consulting help to other areas and direct management of working groups within the Security Area.

Security Area Advisory Group

The main bulk of work for this Group consists of a set of formal work items. These work items correspond to four types of activities:

1. Working groups within the IETF Security Area. These are marked as "Security."

2. Working groups in allied organizations that function as part of the IETF Security Area. These are marked either "PSRG" for the Privacy and Security Research Group, or "TSIG" for working groups within the Trusted Systems Interoperability Group.

3. Security relevant developments within working groups in areas other than security. These are marked according to the relevant area, viz., Applications, Internet Services, Network Management, OSI Integration, Operational Requirements, Routing, Transport and Services, or User Services.

4. Internal work items. These are topics which do not merit the creation of a formal working group but which do need some level of attention. These are assigned to a SAAG member and followed for one or more SAAG meetings. These are marked as "SAAG".

The Security Area Advisory Group met during the first and last working group period of the Santa Fe IETF. The first meeting is used to coordinate the activities for the week and the second meeting is used to report on the activities that have occurred.

During the week, of the twenty-three open work items on Monday, five work items were closed and four new work items were opened. Eight work items received no attention. The key activities for the week to report are all working groups in the Security Area: SNMP Security, Common Authentication Technology, and Privacy Enhanced Mail.

SNMP Security

There are three documents which have been reissued. There are four implementations, three of which have been demonstrated to interoperate with each other. The final actions are cleaning up the documents, reviewing them, and submitting them to the IESG for consideration as Proposed Standards.

One of the important technical issues to be discussed was the choice to be made between message digests: MD4 and MD5. It is clear that MD5 is the right choice for standards actions or something that you want to invent for some stability. MD5 does run slower by some amount than MD4, but the overall equation makes a fairly modest impact. There will probably be a lot of performance measurements showing up, but it is pretty clear that performance is not the critical issue.

This decision effects a number of other working groups, each of which has decided to adopt whatever choice is made by SNMP Security. These include the 822 Extensions and PPP Extensions Working Groups.

Common Authentication Technology

The basic idea is you have a set of applications that want access to one or more authentication mechanisms, for example Kerberos or the Distributed Authentication Security Service (DASS). There is a common program interface, a general security services application program interface, that has been defined such that these applications can be written to be neutral with respect to which mechanism is actually employed. The binding with a mechanism takes place at some later time, currently compile time.

The feature of the mechanisms currently proposed is they depend upon a global identification scheme, i.e., you have a name that exists outside the context of the machine you are trying to connect from or connect to. The name identifies a set of credentials that are forwarded on your behalf. This raises the question of what happens when there is an application of the technology on a machine on which your credentials do not exist, for example a terminal server at an IETF meeting. Does it makes sense for one of the underlying mechanisms to be the use of passwords?

This opened up the discussion in multiple directions, but the critical question is what is the ambition level of the CAT technology, with respect to a much larger set of issues in security, authentication, identification, and in particular with respect to authorization. For now, CAT will continue down the path it is on. There will be subsequent activity to serve these larger functions.

There is a set of documents in preparation. Two Internet Drafts exist that describe the interface, a basic functional description, and specific C structures. An Internet Draft exists for each of Kerberos and DASS.

Privacy Enhanced Mail

There was a great deal of controversy during the Atlanta meeting, but a number of meetings and interactions have occurred since then, resulting in substantial progress and resolution of major issues. It is worth noting there was a booth and demonstrations at INTEROP, where multiple interoperating implementations were heavily attended during the whole period. It was a big success for everybody.

The key decision that came out of the meetings following Atlanta was to open the range of policies. There had been a single policy coming into existence emphasizing very high assurance that the binding of the name and public key in the certificate actually represented a real person and had not been forged. The controversy focused on whether or not the high assurance was appropriate in all cases. The resolution was to push the notion of assurance down one level in the certificate validation hierarchy and create what are now called Policy Certification Authorities, which are bound together by a policy neutral administrative function called the Internet Certificate Authority. The current candidate policies are a continuation of the high assurance policy, a mid-range policy, some support for residential users, and support for persona users.

There will be two bodies of software available. There is an implementation Trusted Information Systems has been developing for some time, that will include modifications to MH and a general purpose filter, which will be distributed in source code form with an object version of the cryptography supplied by RSA Data Security Incorporated (RSADSI). RSADSI will separately make available a limited size tool kit, in source form, on a you-build-it basis. Neither of these is intended for commercial applications.

CURRENT MEETING REPORT

Reported by Larry J. Blunk/Merit

Minutes of the Terminal Server Accounting and Authentication BOF (TERMACCT)

The meeting began with a presentation of the Authentication, Authorization, and Accounting services currently provided in Merit terminal servers, and how these features are lacking in commercial terminal server equipment.

Authentication was discussed and there seemed to be a consensus that Kerberos would be the way to go. There was some question about whether terminal servers with limited resources would be able to implement Kerberos (such as, how much ROM would it take?).

Authorization was mentioned as being a difficult issue. Kerberos V5 has hooks for authorization, but currently provides no definitions. OSF DCE apparently provides some authorization capabilities using Kerberos V5, but it is not clear how suitable it would be for terminal servers.

Accounting and billing issues were discussed among which was the need to define accounting and billing variables. There also may need to be interaction between the authorization and accounting systems (to deny authorization for someone who has exceeded a usage quota, for example). It was mentioned that the cost, in resources and real dollars, of accounting needs to be weighed against the actual value of the service.

There was much interest in the notion of a "connection manager" which could provide a common or customizable user interface. Such a manager would be run on a host machine and would likely interact with the authentication, authorization, and accounting services.

The consensus of the BOF participants seemed to be that Merit should come up with a requirements document for further discussion. It could then be determined whether a new working group should be formed.

Attendees

Steve Alexander	stevea@i88.isc.com
James Barnes	barnes@xylogics.com
Charles Bazaar	bazaar@emulex.com
Larry Blunk	ljb@merit.edu
Philip Budne	phil@shiva.com
Stephen Crocker	crocker@tis.com
Erik Fair	fair@apple.com
Barbara Fraser	byf@cert.sei.cmu.edu
Jack Hahn	hahn@sura.net
Neil Haller	nmh@bellcore.com

Ellen Hoffman	esh@merit.edu
Borka Jerman-Blazic	jerman-blazic@ijs.ac.mail.yu
Charles Kaufman	kaufman@dsmail.enet.dec.com
Mark Knopper	mak@merit.edu
Kenneth Laube	laube@bbn.com
Mark Lewis	mlewis@telebit.com
John Linn	linn@zendia.enet.dec.com
Brian Lloyd	brian@ray.lloyd.com
Glenn McGregor	ghm@merit.edu
Dave Monachello	dave@pluto.dss.com
J. Bradford Parker	brad@cayman.com
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
William Simpson	Bill_Simpson@um.cc.umich.edu
Frank Slaughter	fgs@shiva.com
Eric Smith	eric@telebit.com
Richard Smith	smiddy@pluto.dss.com
Michael St. Johns	stjohns@umd5.umd.edu
Theodore Tso	tytso@mit.edu
Wengyik Yeong	yeongw@psi.com

Authentication, Authorization, and Accounting issues for terminal/network servers

BOF

MERIT

Introduction

• Why do we use the term terminal/network server?

- The term "terminal server" has become somewhat antiquated; may not reflect true capabilities
- We refer to a device which provides traditional connection-oriented, protocol translation services, for example,

  Async <-> Telnet X.25 <-> Telnet Async <-> X.25 Async <-> LAT LAT <-> Telnet

- Such a device may also offer connectionless services such as PPP and SLIP.

• Why did we call this BOF?

- Merit is evaluated vendor equipment to replace its current proprietary terminal servers, and we would like to preserve the functionality of our current terminal servers.
- There seems to be a lack of standards in this area.
- The standards that are currently in use are severely lacking in security (i.e. TACACS)

MERIT

Authentication Issues

• Would like to be able to authenticate in a distributed fashion.
  User should be able to specify a trusted authentication server.

• A secure authentication mechanism is desired.
  Encryption should be employed to prevent passive attacks.
  Mechanism should also be resistant to active attacks.

• Authentication may be implied for certain trusted interfaces.

• CAT needs to be consulted on these issues.

• Issues for connection-oriented sessions
  • Would be desirable to allow non-authenticated sessions based upon which host a user requests (i.e. systems not requiring passwords).

• Issues for connectionless services
  • Automatic authentication is desirable (e.g. PAP or CHAP for PPP).
  • May also want to provide a basic, non-authenticated service (i.e an anonymous service).

MERIT

Authorization Issues

An access control list could be used to define authorization for an entity. This list would specify allowable hosts and/or nets.

• Issues for connection-oriented sessions

Authorization may be done by the remote system instead of the local t/n server (or in addition to).

An access control list could additionally specify allowable telnet ports.

• Issues for connectionless services

An access control list could define allowable tcp and udp ports.

MERIT

Accounting Issues

• Accounting is envisioned as being performed by a separate server.

• Such a server would collect statistical, billing, and auditing information.

• Again, security should be a priority.

• Redundancy (multiple servers) would be important for reliability.

• Accounting WG needs to be consulted.

MERIT

Other Issues

• Traditionally, terminal/network servers have been "resource poor".
  Can vendors be expected to implement complex mechanisms?

• The ability to present a customizable user interface may be important.
  How should this be implemented? Some sort of "connect manager"?

• Which working groups should be consulted?
  CAT, Accounting, and Telnet have been identified. Any others?
  Should a new working group be formed?

MERIT

3.7.1 Commercial Internet Protocol Security Option (cipso)

Charter

Chair(s):

Ron Sharp, rls@neptune.att.com

Mailing Lists:

General Discussion: cipso@wdl1.wdl.loral.com
To Subscribe: cipso-request@wdl1.wdl.loral.com
Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Commercial Internet Protocol Security Option (CIPSO) Working Group is chartered to define an IP security option that can be used to pass security information within and between security domains. This new security option will be modular in design to provide developers with a single software environment which can support multiple security domains.

The CIPSO protocol will support a large number of security domains. New security domains will be registered with the Internet Assigned Numbers Authority (IANA) and will be available with minimal difficulty to all parties.

There is currently in progress another IP security option referred to as IPSO (RFC 1108). IPSO is designed to support the security labels used by the U.S. Dept of Defense. CIPSO will be designed to provide labeling for the commercial, U.S. civilian and non-U.S. communities.

The Trusted Systems Interoperability Group (TSIG) has developed a document which defines a structure for the proposed CIPSO option. The Working Group will use this document as a foundation for developing an IETF CIPSO specification.

Goals and Milestones:

Done — Review and approve the Charter for the IETF CIPSO Working Group. Review revised TSIG CIPSO Specification.

Done — Review outstanding comments/issues from mailing list. Continue work on specification and prepare it for submission as an Internet Draft by the end of May.

Jul 1991 — Review outstanding comments/issues from mailing list. The specification will be submitted to the IESG for consideration as a Proposed Standard.

Mar 1992 — Submit specification to the IESG for consideration as a Draft Standard. There must be at least two interoperable implementations by this time.

Ongoing Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard.

Internet Drafts:

"Commercial IP Security Option", 12/03/1991, Trusted Sys Interop. Group (TSIG) <draft-ietf-cipso-ipsecurity-00.txt>

3.7.2 Common Authentication Technology (cat)

Charter

Chair(s):

John Linn, linn@zendia.enet.dec.com

Mailing Lists:

General Discussion: cat-ietf@mit.edu
To Subscribe: cat-ietf-request@mit.edu
Archive: /cat-ietf/archive@bitsy.mit.edu

Description of Working Group:

The goal of the Common Authentication Technology Working Group is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security-oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accommodating evolution of underlying technologies.

In support of these goals, the Working Group will pursue several interrelated tasks. We will work towards agreement on a common service interface allowing callers to invoke security services, and towards agreement on a common authentication token format, incorporating means to identify the mechanism type in conjunction with which authentication data elements should be interpreted. The CAT Working Group will also work towards agreements on suitable underlying mechanisms to implement security functions; two candidate architectures (Kerberos V5, based on secret-key technology and contributed by MIT, and X.509-based public-key Distributed Authentication Services being prepared for contribution by DEC) are under current consideration. The CAT Working Group will consult with other IETF working groups responsible for candidate caller protocols, pursuing and supporting design refinements as appropriate.

Goals and Milestones:

Done	Preliminary BOF session at IETF meeting, discussions with Telnet and Network Printing Working Groups.
Done	Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet Draft process.

Done	First IETF meeting as full Working Group: review Charter distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.
Oct 1991	Update mechanism-independent Internet Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.
Nov 1991	Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.
Dec 1991	Submit service interface specification to RFC standards track.
Ongoing	Progress Internet Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms.

Internet Drafts:

"Generic Security Service Application Program Interface", 06/12/1991, John Linn <draft-ietf-cat-genericsec-00.txt, .ps>

"The Kerberos Network Authentication Service", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-00.txt, .ps>

"Distributed Authentication Security Service", 11/04/1991, Charles Kaufman <draft-ietf-cat-dass-00.txt, .ps>

CURRENT MEETING REPORT

Reported by John Linn/DEC

Minutes of the Common Authentication Technology Working Group (CAT)

The meeting began with a review of the planned Agenda. The first session was devoted to mechanism-oriented discussion, including presentation and discussion of public-key Distributed Authentication Security Services (DASS) architecture and consideration of weaker-level authentication schemes which might be considered in support of CAT. The second session was primarily devoted to interface questions and issues pending from the Atlanta meeting.

To this point, CAT has emphasized authentication mechanisms which provide authentication in terms of global names but which also require deployment of significant supporting infrastructure. Interest has been expressed in enabling entry to CAT through simpler alternative mechanisms (e.g., passwords, hand-held authenticators, Yellow Pages (YP)), which generally authenticate in terms of local (per-host) names rather than a global structure. This prospect was controversial for two basic reasons: (1) in terms of the level of portability that would actually be supportable for subsequent migration to stronger mechanisms, and (2) because of concern that support within CAT could result in institutionalizing the current weak state of authentication within the Internet. Evaluation and debate on these questions will continue.

DASS Architecture

Charlie Kaufman gave a presentation on the DASS architecture, which was recently submitted to Internet-Drafts and accompanied by a letter from Digital Equipment Corporation to the IAB ceding change control to the IETF process. The general scope was described as strong mutual interactive authentication, with functionality analogous to Kerberos (V4) but extended for elimination of the on-line Key Distribution Center (KDC), limitation of dictionary attacks against passwords, delegation support, hierarchic realm support, and support for various types of principals (user, node, combination). A login agent protocol using two hash algorithms was incorporated to provide password guessing protection. DASS fits under the GSS-API, providing all CAT services as well as additional functions.

DASS credentials cannot, if intercepted, be used to permanently impersonate the principal they represent. Temporary impersonation (for credentials' lifetime, normally corresponding to the duration of a login session) is possible in the case of an overrun workstation. It was also observed that execution of rlogin with the delegation option set results in transfer of credentials to the rlogin target, and concern was expressed that this poses danger in the case of a temporarily unattended workstation.

Several aspects were contrasted against Kerberos. DASS tokens are built by using a certificate chain and the target's public key, but repeated use of public key operations is not needed to build successive authenticators on the same context. Address data is placed into the authenticator, not the predecessor ticket, permitting a deferred, application-specific binding. Timestamps and Kerberos-like authenticator caches are employed to determine authenticator acceptance.

The motivation for DASS's login agent was questioned. This agent was described as a means to provide password guessing protection; it was noted that other key and password protection schemes can also be used, offering different tradeoffs. The absence of Certificate Revocation Lists (CRLs) from the architecture was also questioned; it was noted that the intent was to trust the certificate store as a primary and rapid revocation mechanism, leading to a discussion of the recognized (though not currently implemented) need for authentication of the certificate store. It was also noted that hybrid models accommodating CRL as well as store-based revocation were also possible.

The relation between DASS and Privacy-Enhanced Mail (PEM) was discussed. At the moment, DASS diverges from the most recent PEM selection of signature algorithm representation within X.509 certificates; DASS will likely align with PEM. Different hierarchic traversal rules are employed (including DASS's use of uplink as well as downlink certificates), but DASS and PEM should be able to use a common infrastructure. Sharing of keys and certificate stores should also be possible, given resolution of credential management issues.

The DASS usage of uplink as well as downlink certificates has trust implications, and builds on a premise that closer points in the trust hierarchy will generally be viewed by users and administrators as more trusted than more remote points. Pairwise cross-certification makes it possible to manifest pairwise relationships between different Certification Authorities (CAs), even if remote from each other in the namespace. Compromise of a high-level CA can compromise a large number of authentication paths, but does not impact local or cross-certified authentications lower in the tree.

DASS futures include: DASS/PEM alignment, replacement of the Certificate Distribution Center (CDC) with a standard directory, serverless "PEM-like" modes of operation in which certificates are transferred between peers, and supplemental options to the login agent mechanism, allowing different security vs. convenience tradeoffs (it was noted that standardization in this area, while useful, is less critical than standardization of tokens. A question arose as to whether DASS and PEM should share long-term private keys, given DASS's goal of minimizing such keys' exposure and PEM's requirement (unless, e.g., a password is demanded for each processed PEM message) to keep such keys available and accessible for use. Questions also exist about the logistics of infrastructure sharing with PEM.

Discussion was given to revocation, and how storage and use of CRLs could reduce the need to trust the certificate store. It was asserted that store- based revocation is better suited to rapid revocation (e.g., of a terminated employee) than is the (generally schedule-based) CRL model. While unscheduled CRLs can be generated at any time, it is hard to assure their propagation to all necessary points. Multi-tiered revocation, including CRLs for highly trusted mid- to long-term revocation and store-based short-term revocation, may be an appropriate hybrid.

Discussion was given to partial (limited) delegation. It is desirable to constrain the set of delegated rights, but difficult to predetermine a useful set of restrictions to be supported or to identify what rights particular servers will require in order to carry out user requests. Group affiliations are one possibility (as employed, e.g., in the OSF DCE). It was noted that delegation crosses the boundary from authentication towards authorization. Kerberos V4 requires password re-entry to delegate; in V5, login-time flags permit various alternatives, but there is yet little operational experience with what flag options will be most used. Vint Cerf cited a digital library service example, motivating the need for delegation by the fact that a requester cannot generally determine where actions must be taken in order to satisfy their requests; for this example, a controllable charging right is desired.

Lower-Function Mechanisms

There are a large range of authentication schemes with lower function than the powerful cryptographic schemes so far emphasized within CAT. A key controversial question arose: should such schemes, even at the level of unprotected passwords, be construed or explicitly supported within the CAT model? Arguments in favor include easy caller adoption with potential migration path to later use of stronger mechanisms. Arguments opposed include technical issues which could constrain later migration, and the prospect that institutionalization of weak mechanisms could in fact deter deployment of stronger security mechanisms within the Internet (conflicting with the goal of facilitating deployment of stronger authentication within the Internet).

In discussion, most Working Group attendees opposed recommendation of unprotected passwords as a CAT mechanism. It was observed, for example, that "CAT should provide security services matching caller expectations", and that extension down to the level of unprotected passwords was not perceived as qualifying. There was also an assertion that CAT integration within protocol implementations was unlikely to be performed if no security benefits would directly result. Extension to intermediate mechanisms providing enhancement over passwords, but requiring little infrastructure for deployment, was received more positively.

Many members of the lower-function mechanism class raise technical concerns for CAT integration. They do not normally authenticate in terms of global names, but rather in terms of names local to the verifier system. While it is fairly straightforward to distinguish mechanisms to callers in terms of the security services they provide, there is no comparable means to rank mechanisms providing a particular service in terms of the quality with which that service is provided. It was observed that different classes of mechanisms might be admissible into mutually-trusting threat environments such as those for which RFC-931 was designed. It may be appropriate to recognize the distinction and ordering between two suggested equivalence classes: "non-disclosing" (cryptographically strong) and "disclosing" mechanisms, even though metrics for ordering of strengths within these classes are lacking.

Accommodation of hand-held authenticators and like technologies within CAT would require the ability for such a CAT mechanism to call out for user input at context establishment time. The input required varies on a basis which is target-specific, in contrast to Kerberos or DASS credentials which are typically established in conjunction with user login in a target-independent fashion. Simple passwords could also be user-entered at context establishment time on a target-specific basis, or an encrypted password file (containing multiple target-specific entries) could be unlocked at credential establishment time.

It was noted that Kerberos is the only presently-proposed mechanism which does not require the use of patented public-key technology. NNTP (not developed on a product basis) was cited as an interested client effectively barred from access to such technology. [Note: Plans announced at the IETF Privacy Enhanced Mail Working Group by RSA Data Security to provide a freely-available public-key implementation may modify this situation, should this implementation's interfaces and characteristics prove suitable as a basis for CAT usage]. It was noted that users lacking source code for their operating systems are impeded from authentication system integration requiring, e.g., modification to /bin/login.

A desire was voiced for a "Strategic Plan for CAT Deployment" document to be developed, documenting the pieces and steps required for this process. It was noted that a perception exists that integration of CAT is being construed within the IETF as a prerequisite for advancement of an application protocol on the standards track, and that other working groups may not be fully cognizant of CAT scope, directions, and schedule. It was also noted that a claim of "CAT conformance" is not in itself meaningful, but that "CAT with specific mechanism(s)" is well-formed.

Discussion of Issues List

We discussed identified issues flagged on the CAT mailing list, and considered the interface specification suitable for advancement as a basis for follow-on work.

(D1) Suggestion that CAT mechanisms should incorporate additional token exchanges into context establishment sequences so as to avoid returning COMPLETE status before it is known that the CAT peer has successfully accepted the context. It was accepted as a desirable recommendation to mechanism designers that context establishment should be self-contained and modular, providing full bi-directional peer-entity authentication (and assurance of cryptographic token acceptance) without need to invoke CAT per-message protection primitives in order to validate context setup.

(D2) Desire to make identification of set of intermediaries involved in context establishment available to CAT caller. Such a CAT extension would be technically feasible, but its value for mechanism-independent interpretation was questioned. Since its primary advocate was not available for discussion, the topic was tabled for the present.

(D3) Suggested optional overlay of calls to integrate CAT authentication with data stream calls, analogous to Kerberos' send_auth interface. No new status was reported on this work item.

(D4) Discussion of alternative coding schemes (character sets, etc.) for CAT tokens. This suggestion had been intended as a means to support CAT-based integration of password mechanisms in a manner which would be interoperable with non-CAT peers implementing like schemes. It was recognized in discussion that CAT's scope cannot extend in general to interoperation with peers not supporting CAT and its token exchange paradigm.

(D5) Specifics of shared-mechanism determination approaches, including combinations of negotiation, directory entries, configuration data, and user/caller input. It was proposed that negotiation schemes be considered in follow-on work on an identified "negotiated" mechanism, which would itself exchange tokens in order to identify a shared mechanism and then perform authentication under that shared mechanism.

(D6) CAT naming portability issues and approaches, in advance of IETF-level agreement as cited in (H1). Discussion explored aspects of this problematic area and of the GSS-API facilities incorporated for portability support absent agreement on a common global naming format.

Attendees

Jim Barnes	barnes@Xylogics.COM
Charles Bazaar	bazaar@emulex.com
Larry Blunk	ljb@merit.edu
Thomas Boorman	tmb@lanl.gov
David Borman	dab@cray.com
Ken Carlberg	carlberg@cseic.saic.com
Lida Carrier	lida@apple.com
Vinton Cerf	vcerf@nri.reston.va.us
Jim Clifford	jrc@lanl.gov
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Stephen Crocker	crocker@tis.com
Jim DeMarco	jdemarco@ftp.com
Steve Dusse	spock@rsa.com
Barbara Fraser	byf@cert.sei.cmu.edu
L. Dain Gary	ldg@cert.sei.cmu.edu
Jisoo Geiter	geiter@gateway.mitre.org
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Neil Haller	nmh@bellcore.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Stephen Kent	kent@bbn.com
Peter Kirstein	kirstein@cs.ucl.ac.uk
Deidre Kostick	dck2@sabre.bellcore.com
John Linn	linn@zendia.enet.dec.com
Ellen McDermott	emcd@osf.org
Glenn McGregor	ghm@merit.edu

Bill Melohn	melohn@auspex.com
Andy Nicholson	droid@cray.com
Brad Passwaters	bjp@sura.net
Robert Purvy	bpurvy@us.oracle.com
Jeffrey Schiller	jis@mit.edu
William Simpson	Bill_Simpson@um.cc.umich.edu
Richard Smith	smiddy@pluto.dss.com
Sven Tafvelin	tafvelin@ce.chalmers.se
Theodore Tso	tytso@mit.edu
Sally Wilkins	sfw@lanl.gov
Preston Wilson	preston@i88.isc.com
C. Philip Wood	cpw@lanl.gov

3.7.3 Privacy-Enhanced Electronic Mail (pem)

Charter

Chair(s):
Stephen Kent, `kent@bbn.com`

Mailing Lists:
General Discussion: `pem-dev@tis.com`
To Subscribe: `pem-dev-request@tis.com`
Archive: `pem-dev-request@tis.com`

Description of Working Group:

PEM is the outgrowth of work by the Privacy and Security Research Group (PSRG) of the IRTF. At the heart of PEM is a set of procedures for transforming RFC 822 messages in such a fashion as to provide integrity, data origin authenticity, and optionally, confidentiality. PEM may be employed with either symmetric or asymmetric cryptographic key distribution mechanisms. Because the asymmetric (public-key) mechanisms are better suited to the large scale, heterogeneously administered environment characteristic of the Internet, to date only those mechanisms have been standardized. The standard form adopted by PEM is largely a profile of the CCITT X.509 (Directory Authentication Framework) recommendation.

PEM is defined by a series of documents. The first in the series defines the message processing procedures. The second defines the public-key certification system adopted for use with PEM. The third provides definitions and identifiers for various algorithms used by PEM. The fourth defines message formats and conventions for user registration, Certificate Revocation List (CRL) distribution, etc. (The first three of these were previously issued as RFCs 1113, 1114 and 1115. All documents have been revised and are being issed first as Internet Drafts.)

Goals and Milestones:

Done	Submit first, third, and fourth documents as Internet Drafts.
Done	Submit second document as Internet Draft.
Done	First IETF Working Group meeting to review Internet Drafts.
Sep 1991	Submit revised Internet Drafts based on comments received during Working Group meeting, from pem-dev mailing list, etc.
Nov 1991	Submit Internet Drafts to IESG for consideration as Proposed Standards.

Ongoing Revise Proposed Standards and submit to IESG for consideration as Draft Standard, and repeat for consideration as Internet Standard.

Internet Drafts:

"Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", 03/26/1991, John Linn <draft-ietf-pem-msgproc-01.txt>

"The MD5 Message-Digest Algorithm", 07/08/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md5-01.txt>

"The MD2 Message-Digest Algorithm", 07/10/1991, B. Kaliski <draft-rsadsi-kaliski-md2-00.txt>

"The MD4 Message-Digest Algorithm", 07/10/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md4-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV: Notary, Co-Issuer, CRL-Storing and CRL-Retrieving Services", 07/10/1991, B. Kaliski <draft-ietf-pem-notary-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", 07/17/1991, Steve Kent <draft-ietf-pem-keymgmt-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", 08/22/1991, David Balenson <draft-ietf-pem-algorithms-00.txt>

CURRENT MEETING REPORT

Minutes of the Privacy-Enhanced Mail Working Group (PEM)

Report not submitted. Refer to Area Report for brief summary.

Attendees

Mary Artibee	artibee@sgi.com
Charles Bazaar	bazaar@emulex.com
Vinton Cerf	vcerf@nri.reston.va.us
James Conklin	conklin@bitnic.educom.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Stephen Crocker	crocker@tis.com
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Steve Dusse	spock@rsa.com
Barbara Fraser	byf@cert.sei.cmu.edu
Ned Freed	ned@innosoft.com
James Galvin	galvin@tis.com
L. Dain Gary	ldg@cert.sei.cmu.edu
Neil Haller	nmh@bellcore.com
William Jackson	jackson@manta.nosc.mil
Mitchell Kapor	mkapor@eff.org
Charles Kaufman	kaufman@dsmail.enet.dec.com
Stephen Kent	kent@bbn.com
Peter Kirstein	kirstein@cs.ucl.ac.uk
Vincent Lau	vincent.lau@eng.sun.com
Kenneth Laube	laube@bbn.com
Anthony Lauck	lauck@tl.enet.dec.com
John Linn	linn@zendia.enet.dec.com
Ellen McDermott	emcd@osf.org
Daniel Newman	dan@innosoft.com
Brad Passwaters	bjp@sura.net
Robert Purvy	bpurvy@us.oracle.com
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
Theodore Tso	tytso@mit.edu
Preston Wilson	preston@i88.isc.com
Wengyik Yeong	yeongw@psi.com

3.7.4 SNMP Security (snmpsec)

Charter

Chair(s):

James Galvin, galvin@tis.com
Keith McCloghrie, kzm@hls.com

Mailing Lists:

General Discussion: snmp-sec-dev@tis.com
To Subscribe: snmp-sec-dev-request@tis.com
Archive: snmp-sec-dev-request@tis.com

Description of Working Group:

The SNMP Security Working Group is chartered to determine the set of security services needed by the SNMP. The specification of those services, the supporting mechanisms, and the adjunct infrastructure will become an enhancement to the SNMP and eventually an Internet standard.

The specification must not alter the fundamental SNMP network management philosophy and must not entail changes to existing SNMP standards or framework.

Goals and Milestones:

Done	Publish Internet Draft specifications.
Jul 1991	Submit specification to IESG for consideration as a Proposed Standard.
Dec 1991	Submit specification to IESG for consideration as a Draft Standard.
Ongoing	Submit specification to IESG for consideration as a Standard.

Internet Drafts:

"SNMP Administrative Model", 04/09/1991, James Davin, James Galvin, Keith McCloghrie <draft-ietf-snmpsec-admin-02.txt, .ps>

"Definitions of Managed Objects for Administration of SNMP Parties", 04/09/1991, Keith McCloghrie, James R. Davin, James M. Galvin <draft-ietf-snmpsec-mib-02.txt>

"SNMP Security Protocols", 04/09/1991, James M. Galvin, Keith McCloghrie, James R. Davin <draft-ietf-snmpsec-protocols-02.txt, .ps>

CURRENT MEETING REPORT

Reported by James Galvin/TIS and Keith McCloghrie/Hughes

Minutes of the SNMP Security Working Group (SNMPSEC)

Agenda

- Document Finalization
- Interoperability Reports
- Other Comments
- Steps to Publication

The Working Group wanted to see revised documents and implementation experience before it would consider recommending the documents for publication.

Two of the three documents had been revised and distributed prior to the meeting: SNMP Security Protocols and Definitions of Managed Objects for Administration of SNMP Parties. There were no substantive changes to be made to the SNMP Administrative Model document so it was not revised for this meeting.

Document Finalization

Two editorial changes had been suggested on the mailing list for the revised SNMP Security Protocols document. These changes were noted for the Working Group.

The editorial changes required of the SNMP Administrative Model document were noted for the Working Group.

Interoperability Reports

There are four known implementations of the suite of documents; the only feature not implemented in any of them was support for proxy. Three of them have interoperated with each other, using noAuth/noPriv, using MD4, and using DES. The Working Group requested that the implementations be upgraded to include support for proxy. [Editors' note: two of the implementations were so upgraded within a few days of the meeting.]

A number of minor changes were suggested as feedback from the implementation experience, the most significant being: changing the units of the party clock to be in seconds, and adding a new MIB object to the party table to specify the largest SNMP message size that a party would accept. These changes were presented to the Working Group and all were approved. A suggestion that additional MIB objects were required to support proxy to non-SNMP-party based proxied agents was agreed to, but also that these additional objects were considered to be the subject of separate follow-on document(s).

In addition, some performance data was presented comparing the use of MD4 and MD5 as authentication digest algorithms. The data indicated that using MD5 took 15% longer than no authentication, whereas using MD4 took 5% longer than no authentication. However, it was noted that the MD4 implementation was an "optimized" implementation, while the MD5 implementation was the one directly out of the Internet Draft. This suggests that the reported difference should be a worst case scenario.

Next, it was reported to the meeting that the authors of MD4 have decided that the MD4 algorithm is suitable for use in all applications except those which are long-lived. In particular, a protocol standard is considered long-lived. Consequently, the Working Group decided to adopt MD5 instead of MD4.

Other Comments

A number of other wording changes to the documents were suggested and adopted.

Steps to Publication

The Working Group agreed that its work was ready for publication. The following steps were specified.

- The documents would be revised according to the comments discussed at the meeting by Friday, November 22.
- The documents will be submitted as Internet Drafts by Monday, November 25.
- The three weeks immediately following their availability as Internet Drafts will be set aside for final review of the documents by the Working Group.
- At the end of three weeks, the documents will be revised (if necessary) according to any discussions on the mailing list, and submitted to the IESG with a recommendation they be published as a Proposed Standard.

Attendees

Steve Alexander	stevea@i88.isc.com
James Barnes	barnes@xylogics.com
Larry Blunk	ljb@merit.edu
Steve Bostock	steveb@novell.com
David Bridgham	dab@asylum.sf.ca.us
Theodore Brunner	tob@thumper.bellcore.com
Philip Budne	phil@shiva.com
Jeffrey Buffum	buffum@vos.stratus.com
Jeffrey Case	case@cs.utk.edu
Richard Cherry	rcherry@wc.novell.com
James Codespote	jpcodes@tycho.ncsc.mil
Stephen Crocker	crocker@tis.com
Dave Cullerot	cullerot@ctron.com

James Davin	jrd@ptt.lcs.mit.edu
Michael Erlinger	mike@lexcel.com
Jeff Erwin	
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
James Galvin	galvin@tis.com
William Jackson	jackson@manta.nosc.mil
Ole Jacobsen	ole@csli.stanford.edu
Ronald Jacoby	rj@sgi.com
Satish Joshi	sjoshi@synoptics.com
Frank Kastenholz	kasten@europa.clearpoint.com
David Kaufman	dek@proteon.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
Deidre Kostick	dck2@sabre.bellcore.com
Ron Lau	rlau@synoptics.com
Kenneth Laube	laube@bbn.com
Walter Lazear	lazear@gateway.mitre.org
John Linn	linn@zendia.enet.dec.com
Keith McCloghrie	kzm@hls.com
Ellen McDermott	emcd@osf.org
Evan McGinnis	bem@3com.com
David Minnich	dwm@fibercom.com
Lynn Monsanto	monsanto@sun.com
David Perkins	dperkins@synoptics.com
David Piscitello	dave@sabre.bellcore.com
Robert Purvy	bpurvy@us.oracle.com
Anil Rijsinghani	anil@levers.enet.dec.com
Marshall Rose	mrose@dbc.mtview.ca.us
Gregory Ruth	gruth@bbn.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
William Simpson	Bill_Simpson@um.cc.umich.edu
Timon Sloane	peernet!timon@uunet.uu.net
Bruce Taber	taber@interlan.com
Iris Tal	437-3580@mcimail.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jeremy Wilson	
Preston Wilson	preston@i88.isc.com
John Ziegler	ziegler@artel.com

- Document Finalization
- Interoperability Reports
- Other Comments
- Steps to Publication

Security Protocols

- Change
  party Auth Prot → party Auth Protocol
  party Priv Prot → party Priv Protocol
- Security Considerations
  - if you lose network while changing keys
- MD5
- Time ticks → "seconds"
- Add party MaxMsgSize (484..'FF')
  UDP: 2^{16-1}-hdr
- Setting secrets s/b consistent with expected length
  ie. 10 octets: conformance + 3.4

Administrative Model

- Remove 'mutually disjoint'
- Change
  party Auth Prot → party Auth Protocol
  party Priv Prot → party Priv Protocol
- Timeticks → "seconds"
- Clarify ASN.1 is abstraction
- Add party MaxMsgSize (484..'FF')
  "UDP:2^{16-1}-hdr"
- Page 21 - agent send response to MS MIB

Performance Results
4000 GetRequests

	Short (1 varbind)	long (27 varbinds)
NoAuth/noPriv	21	68
md4Auth/noPriv	29	79
md5Auth/noPriv	32	87
md4Auth/desPriv	68	276
md5Auth/despriv	73.6	-

cpu time in Manager

	Short	Long		
noAuth	7.2	38.3	0.0978	
md6Auth	10.1	42.9	0.103	5%
md5Auth	10.9	46.7	0.1126	15%

Steps to Publication

1. Finalize documents
   - Friday
2. Post Internet Draft
   - Submitted Friday
3. Final WG Review via email
   - (3 weeks)
4. Submit Recommendation to IESG

3.8 Transport and Services Area

Director(s):

- David Borman: dab@cray.com

In the Transport and Services Area, there were six meetings: one BOF and five working group meetings. One of the Working Groups (Audio/Video Transport) was a new Group formed during the IETF meeting.

Service Location Protocol

The scope of the Working Group was narrowed. The Group is now limiting itself to locating services within a single administrative domain. To locate services outside of that area, directory services that are already in place should be used.

Domain Name System

Currently, the thrust of this Working Group is the development of a MIB. A proposed MIB document was distributed at the meeting, but it will need a bit of work before any agreement can be reached.

Distributed File system Group

The DFS Working Group has not been meeting very regularly. It was originally chartered to address the questions of what impact distributed file systems would have in the Internet, and are they going to cause things to break? So far nothing's really broken, so the Group has been taking a defensive position, meeting just when there are interesting things to talk about. If something should arise that would require the Working Group's attention, it will be ready to spring into action at that time. So far the vendors have been good at putting in the necessary changes to keep their distributed file systems from causing problems in the Internet, and not requiring intervention by the Working Group.

Two presentations were made at this meeting: One on Cached NFS from Sun, and one on the Andrew File System Version 3 Specification that has recently been released by Transarc.

New Technology TCP

This Group met as a BOF. The new TCP law, was a presentation by Sean O'Malley, from Arizona, talking about how to do protocol evolution. There was much spirited and lively discussion. The basic outcome was that the Arizona people were encouraged to continue doing their research, but to stop calling it TCP, and instead get a new protocol number for their experiments.

TCP Large Windows

The TCP Large Windows Working Group had a very fruitful meeting. The entire document was reviewed. There were two main changes to the document, both with regard to the Selective Acknowledgement (SACK) option. Originally it has two 16-bit fields, these were expanded to 32 bits. This allows it to operate with the window scale option, but not have a dependence on it. The second change was to make the SACK a hard ACK of data; once the data receiver has sent a SACK, it must retain that data. Once the data sender receives a SACK, it can discard that data from its retransmit queue. The rest of the document was accepted as is. After making the changes, the document will be passed to the IESG for approval to be published as a Proposed Standard.

Audio/Video Transport

This is a new Working Group that was formed during the Teleconferencing BOF. It's Charter is to specify protocols for real time transmission of audio and video over UDP. They want to promote interoperation between the current experimentation that is going on in this area, so that people will be doing things that interoperate. The steps the Working Group will be taking are:

- Solicit contributions from those who have developed packet audio and video.
- Distill appropriate protocol features.
- Test connectivity among Working Group members using DARTnet software.
- Hold the next meeting via packet audio teleconference in mid-January.

CURRENT MEETING REPORT

Minutes of the New Technology TCP BOF (NTTCP)

Report not submitted. Refer to Area Report for brief summary.

Attendees

Steve Alexander	stevea@i88.isc.com
Mary Artibee	artibee@sgi.com
David Borman	dab@cray.com
Yee-Hsiang Chang	yhc@concert.net
Jim DeMarco	jdemarco@ftp.com
Hans Eriksson	hans@sics.se
Karen Frisa	karen.frisa@andrew.cmu.edu
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Peter Honeyman	honey@citi.umich.edu
Christian Huitema	christian.huitema@sophia.inria.fr
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Michael Karels	karels@cs.berkeley.edu
Ron Mackey	rem@dsiinc.com
Allison Mankin	mankin@gateway.mitre.org
Paul Milazzo	milazzo@bbn.com
Greg Minshall	minshall@wc.novell.com
Mark Needleman	mhn@stubbs.ucop.edu
Sean O'Malley	sean@cs.utexas.edu
Tim Seaver	tas@mcnc.org
Frank Solensky	solensky@clearpoint.com
Mike Spengler	mks@msc.edu
Osamu Takada	takada@sdl.hitachi.co.jp
Kathleen Wilde	wilde@decvax.dec.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Nancy Yeager	nyeager@ncsa.uiuc.edu

3.8.1 Audio/Video Transport (avt)

Charter

Chair(s):
Stephen Casner, casner@isi.edu

Mailing Lists:
General Discussion: rem-conf@es.net
To Subscribe: rem-conf-request@es.net
Archive: rem-conf/rem-conf-archive:nic.es.net

Description of Working Group:

The Audio/Video Transport Working Group was formed to specify protocols for real-time transmission of audio and video over UDP and IP multicast. The result may be independent protocols specific to each medium, or a common, lightweight, real-time transport protocol may be extracted.

UDP transmission of audio/video is only sufficient for small-scale experiments over fast portions of the Internet, but the transport protocols produced by this Working Group should be useful on a larger scale in the future when network-level resource management mechanisms are deployed to provide low-delay service and to guard against unfair consumption of bandwidth by audio/video traffic.

Similarly, initial experiments can work without any connection establishment procedure so long as a priori agreements on port numbers and coding types have been made. To go beyond that, we will need to address simple control protocols as well. Since IP multicast traffic may be received by anyone, the control protocols must handle authentication and key exchange so that the audio/video data can be encrypted. More sophisticated connection management is the subject of current research, and should be the topic of a follow-on working group.

Goals and Milestones:

Nov 1991 Define the scope of the Working Group, and who might contribute. Our first step will be to solicit contributions of potential protocols from projects that have already developed packet audio and video. From these contributions we will distill the appropriate protocol features.

Jan 1992 Conduct a teleconference Working Group meeting using a combination of packet audio and telephone. The topic will be a discussion of issues to be resolved in the process of synthesizing a new protocol. Make writing assignments for first-draft documents.

Mar 1992 Review first draft documents, determine necessary revisions. Follow-up discussion will occur on mailing list. Plan implementations.

May 1992 Teleconference meeting using implementations of draft protocols. Discuss draft revisions based on implementations, submit as Internet Drafts.

Jul 1992 Review updated draft, and assess whether these protocols should enter the standards track or be published only as experimental protocols. Make final revisions to drafts and give to IESG for publication as RFCs of appropriate type.

CURRENT MEETING REPORT

Reported by Stephen Casner/ISI

Minutes of the Audio/Video Transport BOF (AVT)

Scope of this Working Group

In the Teleconferencing BOF two days earlier, several areas of interest were identified. One area was deemed ready for a Working Group: the Audio/Video Transport Working Group was formed to specify protocols for real-time transmission of audio and video over UDP and IP multicast, and this kickoff meeting was scheduled.

The purpose of this Working Group is to foster interoperation among several packet audio/video experiments already underway using UDP. Many hosts on the Internet have the capability now to do 64Kb/s PCM audio, and many could inexpensively add frame-grabbers and simple software compression for low-frame-rate video. We want to enable desktop teleconferencing among these hosts. Therefore, the focus of the Working Group is short-term: our goal is to have the protocols defined and experimental implementations running by the July 1992 IETF meeting.

UDP transmission of audio/video is only sufficient for small-scale experiments over fast portions of the Internet. Scaling up will require network resource management and more sophisticated connection management; these are subjects of current research that we will defer to later working groups.

Who might contribute?

Our first step will be to solicit contributions of potential protocols from those projects that are already experimenting with packet audio and video. From these contributions we will distill the appropriate protocol features. Many projects are working on "multimedia", but that is a wide topic. We are particularly interested in those who are working on packet transmission of real-time media. Some 25-30 projects were suggested as candidates by meeting participants.

Brief descriptions of some existing protocols

Of the projects represented by the participants, three have audio and/or video protocols already in use.

- Paul Milazzo from BBN described the protocol used in the Desktop Video Conference (DVC) program. DVC uses the low-cost VideoPix frame-grabber card for SPARCstations plus software compression to generate video at about 5 frames per second. The DVC protocol communicates sequences of video sub-image blocks over either TCP or UDP.

- Hans Eriksson from SICS described the PicturephoneTalk program. It uses TCP connections to carry uncompressed images from the VideoPix card and 64Kb/s PCM audio from the built-in audio device. The protocols include timestamps to allow synchronization of the streams at the receiver.

- Steve Casner from ISI described the Network Voice Protocol (NVP) and Packet Video Protocol (PVP) that have been in use for several years in the multimedia conferencing system operating on the TWBnet, and more recently on DARTnet. The data packet formats that are part of these protocols are fairly simple. They have been used over both ST (Stream Protocol) and UDP.

Short descriptions will be prepared for each of these three protocols to serve as models in the solicitation for contributions from other projects.

Issues to be addressed

A number of protocol issues have already been identified from differences among the protocols known to the participants. In this first session, we did not get to any serious discussion of these issues, but some are included here to prompt thinking:

- Are there functions needed for both audio and video that should be extracted into a lightweight real-time transport protocol layer? One candidate might be timestamp information for synchronization.

- Should timestamps be based on the media sample clock or on real time?

- Should sequence numbers count packets or smaller units such as samples?

- What should be the expected lifetime of the protocols produced by this Working Group? Are they only to be used in the short-term with UDP, or perhaps later directly over a real-time internet protocol?

- NVP and PVP include checksums over the header because they run over ST which provides no checksum. Should such a feature be included in our new protocols in anticipation of their use over protocols other than UDP?

- The CCITT recommendation G.764 Packetized Voice Protocol includes a field that records the cumulative variable queueing delays experienced by a packet in traversing the network. This may be useful for deadline-scheduling of packet forwarding, but is it practical to expect Internet routers to update this field?

- The Xerox PARC Phoenixphone protocol includes an “energy” value for each data packet; should this be included?

- Should we provide a mechanism for communicating control functions within data packets?

It is expected that additional issues will be identified when additional protocols are contributed. Further discussion will occur in subsequent teleconference meetings and by e-mail.

Next Meeting

We've set an ambitious goal to hold the next meeting of the Working Group by packet audio teleconference in mid-January. The initial plan is to use the UDP packet audio software for SPARCstations now used on DARTnet, and to employ DARTnet as an IP Multicast backbone. From DARTnet end nodes we can use IP Multicast Tunneling to other Working Group member sites. We need to begin immediately to test whether network performance on these potential tunnel paths is sufficient.

It is unlikely that all Working Group members will be able to participate by packet audio. It will be possible to set up a telephone conference call for some sites and to patch that into the packet audio conference. We may also face the problem of too much timezone spread among the interested parties (Australia to Europe) to be able to include everyone.

The topic for the teleconference meeting will be a discussion of the protocol issues listed above plus others that arise as we learn more about additional existing protocols that should be considered in the design process. The solicitation for descriptions of these potential contributing protocols will be sent to the Working Group mailing list (rem-conf@es.net) and individually to contacts for each of the 25-30 projects already identified, in case they are not on the list. Anyone wishing to make a contribution is invited to send a short note to the Chair (casner@isi.edu) requesting the solicitation.

Attendees

Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@concert.net
Barbara Denny	denny@sri.com
Hans Eriksson	hans@sics.se
Russ Hobby	rdhobby@ucdavis.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Holly Knight	holly@apple.com
Paul Milazzo	milazzo@bbn.com
Ari Ollikainen	ari@es.net
Allan Rubens	acr@merit.edu
Claudio Topolcic	topolcic@nri.reston.va.us
Andrew Veitch	aveitch@bbn.com

3.8.2 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
Archive:

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 distributed file systems 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.

CURRENT MEETING REPORT

Reported by Peter Honeyman/UMich

Minutes of the Distributed File System Working Group (DFS)

The DFS Working Group met for the third time on November 19, 1991 at the Santa Fe IETF.

Agenda

- NFS Developments
- AFS-3 Documents
- AFS-3 Congestion control
- Announcements

NFS Developments

Tom Kessler (kessler@sun.com) described work at Sun to add local disk caching to NFS.

The Cache File System (CFS) is a generic mechanism that caches files and directories from other VFS systems. The principal cache repository is UFS, i.e., the Berkeley FFS.

A principal design goal to boost NFS server performance by reducing load, but CFS helps reduce network load as well if the cache hit rate is high. CFS is also useful for improving CD-ROM performance.

Like AFS-3, CFS caches chunks of files. Unlike AFS-3, there is a one-to-one correspondence between cached files and files on the server. Missing chunks are represented by "holes" in the cached file.

Consistency checking has not been implemented; CFS is a client-only modification, so the consistency checking can be no stronger than that in the VFS system being cached. The consistency check mechanism is modular and offers hooks for a CFS developer to provide alternate enforcement mechanisms.

"Blot-out" mode lets you overlay files with local copies. The unit of blot-out is a complete file. The local overlay is not purged from the cache by ordinary LRU replacement policy. Other files can be marked to make them "sticky" in the cache.

CFS supports numerous write modes:

- Write-through. Synchronous with server.

- Blot-out. Write to cache only, make local copy sticky. Useful for writing CD-ROM.

- Write-around. Modify actual file only. Useful if cache is scarce resource. [I may not have this right.]

Write-through is the normal mode. CFS helps READ, READDIR, READLINK, LOOKUP performance, does not help GETATTR, directory modifications, WRITE, SETATTR.

As for the bottom line, Tom Kessler, who uses CFS on his home computer, was asked "How does it feel?" According to Tom, "It feels pretty good."

Chris Silveri's foils, which Tom Kessler used in his CFS presentation, can be obtained in PostScript form via anonymous FTP from citi.umich.edu.

- See /afs/umich.edu/user/h/o/honey/IETF/cfs-vg.ps.

Other NFS Developments

There has been some progress on the part of vendors in tuning the NFS parameters (tsize, RTO, RTT measurements) in systems they ship to better conserve network resources. A number of people reported that they find NFS/UDP over the Internet satisfactory. [At least one person was surprised to hear this.]

NFS/TCP is commercially available, and is under development by many vendors. Connection maintenance is not entirely a solved problem.

Sun/RPC over UDP has problems with accurate RTT because the network latency is smeared by the upper-layer (i.e., NFS) service times. (See "Transport Issues in the Network File System" by Bill Nowicki, Computer Communication Review 19(2), pp. 16-20 (April, 1989) for related work.)

Watch Connectathon for further activity in the NFS/TCP arena.

AFS-3 Documents

There was some discussion of the four-or-so inches of AFS-3 documents made available by Transarc. It is not clear what advantage there is in putting an RFC imprimatur on them. Nor is Transarc enthusiastic about reformatting the documents to conform to RFC 1111.

AFS-3 Congestion Control

Peter Honeyman (honey@citi.umich.edu) described his recent work on congestion control for Rx. (Joint work with Dave Bachmann and Larry Huston.) The goal has been to make AFS usable over slow links, down to about 10 Kbits/sec. Much has been accomplished so far, work continues.

Announcements

dfs-wg@citi.umich.edu is a mailing list for ongoing discussions of the Working Group. Administrative matters, such as requests to be added or dropped from the list, should be addressed to dfs-wg-request@citi.umich.edu, not to the list as a whole.

There is a Workshop on File Systems to be held in Ann Arbor on May 21-22, 1992. Contact fsworkshop@citi.umich.edu for further information.

Attendees

Mary Artibee	artibee@sgi.com
David Borman	dab@cray.com
Philip Budne	phil@shiva.com
Randy Butler	rbutler@ncsa.uiuc.edu
Lida Carrier	lida@apple.com
Yee-Hsiang Chang	yhc@concert.net
Richard Cherry	rcherry@wc.novell.com
Jim DeMarco	jdemarco@ftp.com
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Peter Honeyman	honey@citi.umich.edu
Tom Kessler	kessler@sun.com
Holly Knight	holly@apple.com
Vincent Lau	vincent.lau@eng.sun.com
Tony Mason	mason@transarc.com
Bill Melohn	melohn@auspex.com
Paul Milazzo	milazzo@bbn.com
Greg Minshall	minshall@wc.novell.com
Robert Morgan	morgan@jessica.stanford.edu
J. Bradford Parker	brad@cayman.com
Eric Smith	eric@telebit.com
Mike Spengler	mks@msc.edu
Sven Tafvelin	tafvelin@ce.chalmers.se
Kathleen Wilde	wilde@decvax.dec.com
Preston Wilson	preston@i88.isc.com
Nancy Yeager	nyeager@ncsa.uiuc.edu

3.8.3 Domain Name System (dns)

Charter

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

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

Description of Working Group:

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

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

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

Goals and Milestones:

TBD	Adding DNS variables to the MIB.
TBD	Hints, tips, and operations guide for DNS software.
TBD	Implementation catalog for DNS software.
TBD	Discussion of adding load balancing capability to the DNS.
TBD	Discussion of adding a Responsible Person Record.
TBD	Discussion of adding network naming capability to the DNS.

CURRENT MEETING REPORT

Reported by Michael Reilly/DEC

Minutes of the Domain Name System Working Group (DNS)

The primary purpose of this meeting was to discuss the DNS MIB and issues related to its development. The following topics were put up for discussion:

- Grouping of the Objects
- Selection of Object to include
- Read-write status of the Objects
- Additional Objects to include

An early draft was handed out to focus discussion. It contained all of the groups currently proposed as well as many of the objects for each of the groups.

One of the first items was a discussion of what imperative operations would be allowed via SNMP. After some discussion we agreed to include an object which would tell the server to "do a zone update now".

It was also agreed that we would provide a mechanism which would allow for cache entries to be removed - this may be accomplished by setting the TTL to zero for the specified entry. There was a strong desire to have the cache table have an entry for the source of the cache information to aid in debugging.

The most difficult subject was whether to allow SNMP to add or modify records in permanent storage (i.e., not just in cache). Since the exact sequence of events to take place on a server when a record is modified was never specified, it is difficult to write MIB elements until this is understood. It was agreed that Bob Austein would attempt to write up the "rules" which define the sequence of steps to be taken when an action such as the creation of an A record is specified. Jon Saperia will work with Bob to write a set of corresponding MIB objects which reflect these "rules".

There was also a discussion of how various counter information would be tabulated in the MIB. Several of the requests were to be able to see request not only by record type, but also by:

- Requests via UDP
- Requests via TCP
- Number of local requests
- Number of remote requests

In addition to these "cuts" at the data, several suggestions were made about the specific items to count. There are now lists of counters suggested by both Win Treese and Philip

Almquist. They will be reconciled and put into the appropriate counters groups in the next version of the MIB which will be posted to the mailing list.

Bob Austein also brought up the general issue of how to deal with the changes that may be coming up in IP addressing and subnet mask sizes. He will work on some ideas and present that information to the Working Group for discussion.

The group expressed a desire for the mailing list for the Working Group to be the name-droppers mailing list.

Attendees

Steve Alexander	stevea@i88.isc.com
Philip Almquist	almquist@jessica.stanford.edu
Mary Artibee	artibee@sgi.com
Robert Austein	sra@asylum.sf.ca.us
Ronald Broersma	ron@nosc.mil
Lida Carrier	lida@apple.com
Robert Enger	enger@seka.scc.com
Scott Kaplan	scott@ftp.com
Darren Kinley	kinley@crim.ca
Walter Lazear	lazear@gateway.mitre.org
Gary Malkin	gmalkin@ftp.com
Milo Medin	medin@nsipo.nasa.gov
Paul Mockapetris	pvm@darpa.mil
J. Bradford Parker	brad@cayman.com
Brad Passwaters	bjp@sura.net
Michael Patton	map@lcs.mit.edu
Robert Purvy	bpurvy@us.oracle.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Tim Seaver	tas@mcnc.org
Frank Slaughter	fgs@shiva.com
Michael St. Johns	stjohns@umd5.umd.edu
Martha Steenstrup	msteenst@bbn.com
John Veizades	veizades@apple.com
Scott Williamson	scottw@nic.ddn.mil
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov
Robert Woodburn	woody@cseic.saic.com
Nancy Yeager	nyeager@ncsa.uiuc.edu

3.8.4 Service Location Protocol (svrloc)

Charter

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

Mailing Lists:
General Discussion: srv-location@apple.com
To Subscribe: srv-location-request@apple.com
Archive: pub/srv-location/svr-loc-archive

Description of Working Group:

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

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

Goals and Milestones:

Done Open discussion and determine if a Working Group should be formed.

Done Continue discussion trying to refine the problem statement and possible resolutions.

Jul 1991 Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the Service Location Protocol Working Group (SVRLOC)

Significant progress was made in narrowing the scope of this Working Group and the protocol that it is trying to design.

A decision was reached to limit the scope of this Group to an administrative domain that is defined as an area which is equivalent to the area served by a binding broker (a central site to manage the service topology of the area). Outside of that area the scaling and naming issues are left to the directory services in place.

Scott Kaplan from FTP Software and John Veizades from Apple will be working on an architectural document for presentation to the Working Group list by the next meeting. Scott Bradner from Harvard and Greg Bruell from Shiva offered to comment on preliminary versions of the document.

Scott and John met after the meeting and made significant progress on the architecture.

Attendees

Steve Bostock	steveb@novell.com
Scott Bradner	sob@harvard.edu
David Bridgham	dab@asylum.sf.ca.us
Gregory Bruell	gob@shiva.com
Jeffrey Buffum	buffum@vos.stratus.com
Randy Butler	rbutler@ncsa.uiuc.edu
Lida Carrier	lida@apple.com
Richard Cherry	rcherry@wc.novell.com
Jim DeMarco	jdemarco@ftp.com
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Peter Honeyman	honey@citi.umich.edu
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@ftp.com
Holly Knight	holly@apple.com
Joshua Littlefield	josh@cayman.com
Greg Minshall	minshall@wc.novell.com
Thomas Pusateri	pusateri@cs.duke.edu
Michael Ritter	mwritter@applelink.apple.com
Tim Seaver	tas@mcnc.org
Frank Slaughter	fgs@shiva.com
Timon Sloane	peernet!timon@uunet.uu.net
John Veizades	veizades@apple.com

Lee Wade	wade@nsipo.arc.nasa.gov
Preston Wilson	preston@i88.isc.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov
Wengyik Yeong	yeongw@psi.com

3.8.5 TCP Large Windows (tcplw)

Charter

Chair(s):
David Borman, dab@cray.com

Mailing Lists:
General Discussion: tcplw@cray.com
To Subscribe: tcplw-request@cray.com
Archive:

Description of Working Group:

The TCP Large Windows Working Group is chartered to produce a specification for the use of TCP on high delay, high bandwidth paths. To this end, this Working Group recommended RFC 1072 "TCP extensions for long-delay paths" and RFC 1185 "TCP Extension for High-Speed Paths" be published jointly as a Proposed Standard. Deficiencies in the technical details of the documents were identified by the End-to-End Research Group of the IRTF. Rather than progress the standard with known deficiencies, the IESG tasked the End-to-End Research Group to fix and merge these two documents into a single protocol specification document. This review was done on the eze-interest@isi.edu mailing list.

The TCP Large Windows Working Group is being resurrected for a one time meeting, to review and if appropriate, approve this new document.

Goals and Milestones:

Nov 1991 Review the TCP Extended Window Size proposal from the IRSG End to End Research Group and if acceptable, recommend it for standards status.

Internet Drafts:

"TCP Extensions for High Performance", 11/12/1991, V. Jacobson, R. Braden, D. Borman <draft-ietf-tcplw-tcpext-01.txt>

CURRENT MEETING REPORT

Reported by David Borman/Cray Research

Minutes of the TCP Large Windows Working Group (TCPLW)

The TCP Large Windows Working Group met for the first time in over a year and a half. Since the last time the group met, RFCs 1072 and 1185 had been proposed as Internet Standards, but problems with the specifications were discovered, and a new version, combining the two documents, was drafted and modified to address the problems. This document was circulated for comments and discussion, and the Working Group was scheduled to review the current status of the document, and hopefully come up with something that could be recommended for the standards track.

Bob Braden gave an overview of the proposed options. (The overheads that were presented are attached to this report.) One of the assumptions made when the options were being designed was that vendors would not want to change the TCP header.

Discussion was then broken up into several segments, to allow discussion of each proposed option, and the merits of each proposed option versus other ways of solving the same problems.

It was agreed that the Window Scale (WS) options was well-defined, and discussion of a window scale option versus a 32 bit window option in each packet was deferred until later in the meeting.

The timestamp option was then discussed, just as a timestamp. The Protection Against Wrapped Sequence (PAWS) discussion was deferred to later. It was agreed that the definition was sufficient.

The next item for discussion was the Selective Acknowledgement (SACK) option. The first item of discussion was whether the offset and length values should be 16 or 32 bits. There was also discussion as to whether the SACK option needed to have more than one SACK field. The decisions were that the fields would be 32 bit fields; it simplifies the specification and processing of the option, and eliminates any dependencies on the window-scale option. It was also felt that nothing was gained by limiting the number of SACKs in a single SACK option.

The other item of discussion on SACK was whether SACK is advisory or a real acknowledgement. It was agreed that it was a real ACK; hence once a piece of data has been SACKed, the receiver of the data has committed to accepting the data, and the sender is free to discard its copy of the data that it was saving for possible re-transmission.

Having approved the three options as being adequately defined, with the agreed upon changes, the discussion turned back to window scale option vs. a 32 bit window in each packet, and the PAWS mechanism. With the caveat that the WS option would have no effect on other options (done by expanding the SACK option to 32 bits), and that the initial value for the shift is recommended to be based on the size of the receivers buffer, it was agreed to go with the window scale option.

The final discussion was on PAWS. First, the question was, did the PAWS mechanism provided adequate protection, and was it clearly defined? The answer was "yes". The final discussion then rested on PAWS vs. expanding the sequence space to 64 bits. Since PAWS works, and 64 bit sequence space eats up option space, and there were some concerns about the extra processing needed to deal with a 64 bit sequence space, it was decided to accept the PAWS method instead of expanding the sequence space.

Bob Braden will incorporate the changes into the document, and get it published as an Internet Draft. Editoral comments on the document are to be sent to Bob.

Attendees

David Borman	dab@cray.com
Robert Braden	braden@isi.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Richard Cherry	rcherry@wc.novell.com
Jim DeMarco	jdemarco@ftp.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Kenneth Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Jean-Michael Jouanigot	jimi@cernvax.cern.ch
Darren Kinley	kinley@crim.ca
Ron Mackey	rem@dsiinc.com
Donald Merritt	don@brl.mil
Glenn Michel	gym@lanl.gov
Sean O'Malley	sean@cs.utexas.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Miguel Sasson	sasson@xylogics.com
Tim Seaver	tas@mcnc.org
John Seligson	johns@ultra.com
Mike Spengler	mks@msc.edu
Kathleen Wilde	wilde@decvax.dec.com
Johnathan Wilson	wilson@ddnuvax.af.mil
Nancy Yeager	nyeager@ncsa.uiuc.edu

TCP EXTENSIONS FOR HIGH-PERFORMANCE

Van Jacobson

Bob Braden

Dave Borman

TCP Extensions using new TCP options:

- RFC-1072: For large bandwidth*delay product paths
  - Extended Windows
  - Selective Acknowledgments
  - TCP Timestamps
- RFC-1185: For high-speed paths (> O(100 Mbps))

RFC-1072: TCP EXTENSIONS FOR LONG-DELAY PATHS

A. EXTENDED WINDOW

TCP Window Scale Option

3	L=3	S

Effective Window = $2^S \times$ Transmitted Window

$0 \leq S \leq 14$

Both sides must send Window Scale in SYN segments.

RFC-1072: TCP EXTENSIONS FOR LONG-DELAY PATHS

B. SELECTIVE ACKNOWLEDGMENTS

TCP SACK-Permitted Option

4	L=2

In SYN: "I can receive SACKs"

TCP SACK (Selective ACK) Option

5	L	S_1	N_1	...	S_K	N_K
					origin	size

Block of Q'd data

RFC-1072: TCP EXTENSIONS FOR LONG-DELAY PATHS

C. MEASURING RTT - *RTTM*

TCP Timestamps Option: TSopt

8	L=10	Timestamp Value *TSval*	TS Echo Reply *TSecr*

Both sides must send TSopts in SYNs, to allow TSopts in data segments.
and ack

RTTM

TCP A   TCP B

1. ---> < SYN, TSval=1 > --->
2. <--- < SYN, ACK(SYN), TSval=17, TSecr=1 > <---
3. ---> < ACK(SYN), TSval=2, TSecr=17 > --->
4. ---> < data1, TSval=5, *TSecr=17* > ---> (ignore TSecr)
5. <--- < ACK(data1), TSval=29, TSecr=5 > <---
6. ---> < data2, TSval=8, *TSecr=29*> ---> (ignore TSecr)
7. <--- < ACK(data1), TSval=33, TSecr=8 > <---

RTTM SPECIAL CASES

	TCP A		TCP B
1.		—> < dataA, TSval=1 > —>	(ACK delayed)
2.		—> < dataB, TSval=2 > —>	
3.	(A)	<— < ACK(dataB), TSecr=1 > <—	
3.		—> < dataC, TSval=3> —> X *(lost)*	
4.		—> < dataD, TSval=4 > —>	(Queued)
5.	(B)	<— < ACK(dataB), TSecr=1> <—	
5.	(Retransmit)	—> < dataC, TSval=8 > —>	
6.	(C)	<— < ACK(dataD), TSecr=8 > <—	

RFC-1185: TCP CORRECTNESS AT HIGH SPEEDS

Must discard old duplicate segments, despite sequence number wrap-around.

Bandwidth	*Time to wrap 32-bit sequence #*
1.5 Mbps	> 3 hrs
10 Mbps	~ 30 mins
45 Mbps	~ 6.5 mins
100 Mbps	~ 3 mins (Getting risky...)

Old data segments: can cause undetected errors.

Old ACK segments: can cause connection lock-up/failure.

PAWS — *Protest Against Wrapped Sequence #s*

VJ: Use the TCP timestamp option [RFC-1072] as a logical extension of the sequence number field.

Discard (non-SYN) segments whose timestamps are earlier than the timestamp of the last segment received at the left window edge.

- Timestamps initialized by TSval in SYN segments.
- Still use 3-way handshake to validate SYN segments.
- Send timestamps on both ACK and data segments.
- Can use same timestamp that RFC-1072 used for RTT.

TIMESTAMP CLOCK

- RTTM [RFC-1072]: *Timestamp clock is:*
  - Proportional to realtime, with frequency for RTT measurement.
  - Not monotonic across crashes or new connections.
- PAWS [RFC-1185]: *Timestamp clock is:*
  - Not necessarily proportional to realtime, but MONOTONIC within a connection.
  - Frequency: ticks 1/sec – 1/ms.
  - Not monotonic across crashes or new connections. [assuming TIME-WAIT and 'Quiet Time' of 2*MSL].

TIMESTAMP CLOCK

- Making monotonic clock:
  - Hardware clock
  - Count clock interrupts
  - Add variable offset to system (software) clock that is subject to being reset.

TIMESTAMP CLOCK

- Old-age timestamp values

  If connection is IDLE for 25 days, timestamp value saved in TCPCB wraps around, and all new timestamps will be rejected (!)

- Invalidate PAWS timestamp if connection is idle for more than 25 days.

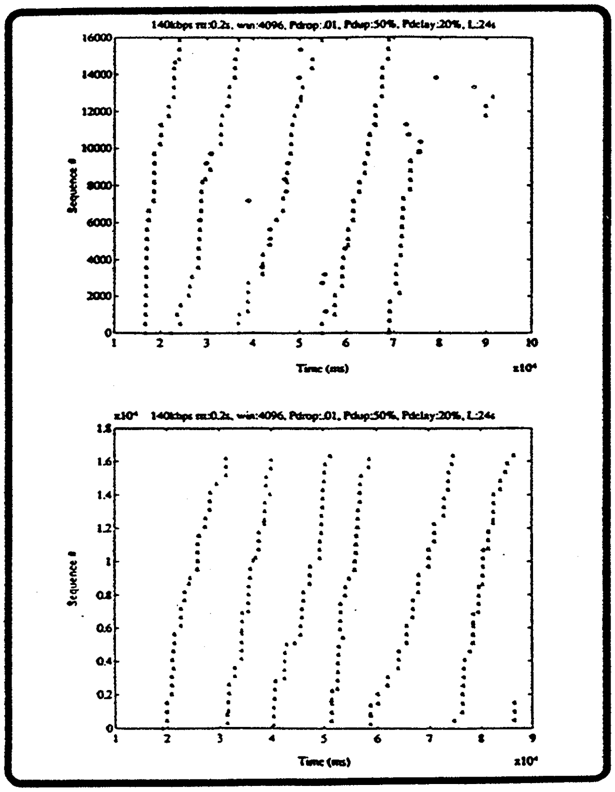
140kbps rtt:0.2s, win:4096, Pdrop:.01, Pdup:50%, Pdelay:20%, L:24s
Sequence #
Time (ms)
x10^4
x10^4
140kbps rtt:0.2s, win:4096, Pdrop:.01, Pdup:50%, Pdelay:20%, L:24s
Sequence #
Time (ms)
x10^4

3.8.6 Trusted Network File Systems (tnfs)

Charter

Chair(s):

Fred Glover, fglover@decvax.dec.com

Mailing Lists:

General Discussion: tnfs@wdl1.wdl.loral.com
To Subscribe: tnfs-request@wdl1.wdl.loral.com
Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Trusted Network File System (TNFS) Working Group is chartered to define protocol extensions to the Network File System (NFS) Version 2 protocol which support network file access in a multilevel secure (MLS) Internet environment. MLS functionality includes mandatory access control (MAC), discretionary access control (DAC), authentication, audit- ing, documentation, and other items as identified in the Trusted Computer System Evaluation Criteria (TCSEC) and Com- partmented Mode Workstation (CMW) documents.

The primary objective of this Working Group is to specify extensions to the NFS V2 protocol which support network file access between MLS systems. It is intended that these extensions should introduce only a minimal impact on the existing NFS V2 environment, and that unmodified NFS V2 clients and servers will continue to be fully supported.

Transferring information between MLS systems requires exchanging additional security information along with the file data. The general approach to be used in extending the NFS V2 protocol is to transport additional user context in the form of an extended NFS UNIX style credential between a Trusted NFS (TNFS) client and server, and to map that con- text into the appropriate server security policies which address file access. In addition, file security attributes are to be returned with each TNFS procedure call. Other- wise, the NFS V2 protocol remains essentially unchanged.

The Trusted System Interoperability Group (TSIG) has already developed a specification which defines a set of MLS exten- sions for NFS V2, and has also planned for the future integration of Kerberos as the authentication mechanism. The TNFS Working Group should be able to use the TSIG Trusted NFS document as a foundation, and to complete the IETF TNFS specification within the next 3-6 months.

Goals and Milestones:

Done Review and approve the TNFS Working Group Charter, review revised TSIG TNFS Specification, and publish a proposed standard following the July meeting.

Jul 1991	Review revised TSIG TNFS Specification.
Nov 1991	Publish a Proposed Standard following the July meeting.
Oct 1991	Review outstanding comments/issues from mailing list.
Oct 1991	Make any final revisions to TNFS document based on comments, issues, and interoperability testing.
Mar 1992	Request IESG to make the revised document a Draft Standard.
Mar 1991	Verify the interoperability of TNFS implementations at the 1992 NFS Connectathon.

Internet Drafts:

"A Specification of Trusted NFS (TNFS) Protocol Extensions", 07/23/1991, Fred Glover <draft-ietf-tnfs-spec-00.txt, .ps>

INTERIM MEETING REPORT

Reported by Fred Glover/DEC

Minutes of the Trusted Network File Systems Working Group (TNFS) (July 1991)

Overview

The TNFS Working Group met for the first time as an IETF working group during the July 1991 IETF meeting in Atlanta. A recent reorganization of the IETF Applications Area, directed by Russ Hobby, resulted in the formation of a new Transport and Services area, which is directed by Dave Borman. The Trusted NFS IETF Group is administratively positioned within the the Transport and Services Area, and thus falls under the direction of Dave Borman. Steve Crocker, Security Area Director, will provide the TNFS Group with security assistance, and Richard Basch, from the Security Area, is currently participating in our Working Group to help us get started.

Our TNFS Charter and current draft of the TNFS specification were posted in the Internet-Drafts Directories. The base name of the specification is:

`<draft-ietf-tnfs-spec-00.(ps, txt)>`

Although the TSIG trusted NFS Group has been meeting for some time, this was the first official meeting of the IETF Trusted NFS Group. As an IETF working group, we will plan to follow the IETF guidelines and procedures for IETF working group operation. Our primary objectives are identified in our Charter. We decided to continue to meet on the TSIG meeting schedule, which meets more frequently than IETF, and also allocates two full days for the Working Group effort. This will permit the Group to continue to make good progress towards meeting its objectives.

We will plan to meet concurrently with the IETF on at least an annual basis. During those joint meetings, we will plan to schedule time to present a status overview of the TNFS Working Group effort, with an opportunity to collect feedback and comments from attendees. Due to the large number of IETF sessions, and the limited amount of time allocated to each session, it seems reasonable to focus the joint meeting on a general status and Q & A, and to focus the Working Group meetings on actual work items. Note that the Working Group is an IETF working group, and is open to all who have interest in participating in the Working Group activities.

Our TNFS meeting Minutes will be distributed through the IETF, and our documents will be updated and registered in the IETF directories.

Meeting Summary

At this first meeting of the IETF TNFS Working Group, we allocated time to review the current status of our TSIG efforts, and to identify our IETF objectives. During the meeting, we:

- Reviewed the TNFS Charter.
- Reviewed the primary objectives of the TNFS Working Group.
- Presented a slide overview of the TNFS specification.
- Inspected (page by page review) the TNFS specification, and
- Inspected the token mapping specification.

Charter Review

As a result of the review of the TNFS Charter document, we:

- Agreed to pursue the standards track for both the TNFS specification and the Token Manager specification (but, see open issues below).

- Agreed to plan for interoperability testing of TNFS implementations in March, 1992.

Token Manager Review

Recommendations from the review of the Token Manager document included:

- Include additional examples for clarity.

- Add a rationale section and include rationale for the omission of generation numbers.

- Add the RPC program number now assigned for this service (390087 decimal).

- Present the service as required for TNFS, but applicable to other clients as well (trusted sessions for example).

- Place the updated document into the Internet-Drafts Directory to obtain additional comments.

TNFS Specification Review

Recommendations from the review of the TNFS document included:

- Add a new section to the document identifying any expectations and/or requirements including:

  - The use of authentication and message integrity for commercial environments.

  - Restricting client side mount operations to the server's export point only; a request to mount a subdirectory which resides below the export point in the exported directory would be denied; rationale: symbolic links could be used on the client to provide a similar effect (to that of mounting subdirectories); without this restriction, access permission from clients is not checked against the higher level components in the server's exported tree.

- Other related RPC services in the trusted environment, such as the lock manager and mount daemon, must be able to support the AUTH_MLS credential flavor, but may also make use of other policies; the token manager must also support the AUTH_MLS credential format.

Open Issues/Action Items

A number of issues which came up during the review of the documents resulted in action items to be addressed prior to the next meeting:

- Update the Token Manager specification; place into the Internet-Drafts Directory. [Fran Fadden]

- Update the TNFS specification; place update into the Internet-Drafts Directory. [Fred Glover]

- Communicate with Area Director regarding policies for standards versus experimental track; obtain recommendation regarding requirement for authentication, message integrity prior to standards track. [Fred Glover]

- Document tradeoffs in the sender/receiver based mapping models. [Charlie Watt]

- Develop proposal to add a flags field to the attribute structures in support of additional optional policies, such as multi-level directories and the two person rule. [Charlie Watt]

Next Meeting

The TNFS Group will plan to meet as both a TSIG and an IETF Working Group at the Hewlett-Packard facilities in Cupertino, California, (October 15-17, 1991). At that meeting, we will plan to:

- Review the TNFS Implementation Hints - Carl Smith author.
- Review the TNFS Administration - Ali Gohshan author.
- Review the updated RPC based Token Proposal - Fran Fadden, Dave Summers authors.

  - NOTE: a Token Manager presentation could be made at the plenary session as well; there appeared to be interest in this from the trusted sessions group at the July meeting.

- Review the updated TNFS RFC Document.
- Review the sender/receiver based token models.
- Review the attribute structure flags field.

Attendees

Richard Basch	basch@mit.edu
Hardy Doelfd	
Fran Fadden	fran@decvax.dec.com
Fred Glover	fglover@decvax.dec.com
Ali Gohshan	
Barry Miracle	miracle@sctc.com
Andy Nicholson	droid@cray.com
Mark Saake	saake@llnl.gov
Carl Smith	cs@eng.sun.com
Mark Stein	marks@eng.sun.com
Charles Watt	
Larry Wikelius	lwikeliu@convex.com

INTERIM MEETING REPORT

Reported by Fred Glover/DEC

Minutes of the Trusted Network File Systems Working Group (TNFS) (October 1991)

General Summary

The TNFS Working Group met in October 1991 as a joint IETF/TSIG Working Group. Our goal is to develop an RFC standard for Trusted NFS, and to develop additional documentation to support implementation and management of the TNFS environment. We have working documents for TNFS, Token Management, Implementation Guidelines, and System Management which we are reviewing for completeness at present. Our plan is to demonstrate basic TNFS interoperability at the March '92 Connectathon.

Meeting Summary

During the October meeting, we:

- Reviewed the TNFS Charter.
- Inspected (page by page review) the updated Trusted NFS (TNFS), token mapping (TKM), implementation guidelines, and administration documents, and
- Reviewed a proposal for the introduction of multi-level directories (MLDs).

Charter Review

At the July meeting, there were a couple of questions raised concerning the specific RFC "track" to be used for the TNFS specification, and the need to include support for authentication and message integrity into the TNFS draft. These questions were later discussed with Dave Borman, our Area Director, to obtain Dave's input on our options. Based on Dave's comments and our Group discussion during the October meeting, we agreed:

- To move the TNFS specification forward on the standards track,
- To address the issues of authentication and message integrity in a future effort, and
- To tighten the relationship of the supporting documents with the TNFS specification: Token Manager, Implementation, and Management documents.

The Group Charter and TNFS specification will be updated to explicitly identify these decisions.

Token Manager Review

Recommendations from the review of the Token Manager document included:

- Include support for the use of the AUTH_MLS credential within the Token Manager.
- Additional discussion/review is required to understand:
  - Requirements for the support for a reverse mapping mechanism to support the mapping of attributes to tokens as well as the expansion of tokens into attributes.
  - Requirements for diskless initialization.
  - Requirements for distinguished value tokens, and
  - Requirements for tokenized UID/GID (and ability to map UID/GIDs).
- Place the updated document into the Internet-Drafts Directory to obtain additional comments.

TNFS Specification Review

Recommendations from the review of the TNFS document included:

- Update the references
- Improve the auditing policy statements
- Replace the national caveot field with a vendor specified field
- Use the portmapper to identify the port to be used for TNFS
- Additional discussion/review is required for the following:
  - Use of distinguished value tokens for purpose of "policy not in effect"
  - Requirements for multiple privilege sets (POSIX?)
  - Requirements for multi-level directories

TNFS Implementation Guidelines Specification Review

Recommendations from this review included:

- Update to reflect decision on use of vendor specific fields
- Add references section and include TNFS and TKM documents
- Plan to include multi-level directory section

TNFS Administration Specification Review

Recommendations from this review included:

- Add host specific export control

- Add additional context and/or guidelines for the use of export control variables

Next Meeting

The TNFS Group will plan to meet as both a TSIG and an IETF Working Group in January (28th-30th) at the Sheraton Execuport in Fort Lauderdale, Florida. At that meeting, we will plan to:

- Continue the discussion of the proposed document updates identified above
- Continue the discussions on multi-level directories
- Review the SecureWare Token Mapping mechanism (Joy Leima presentation)
- Review Interoperability Test Plans:
  - Readiness to participate in Connectathon '92
  - Updated NFS test suite
  - Review use of unmapped attributes from earlier TSIG meeting
  - Updated "nfs.h" file

Attendees

Fran Fadden	`fran@decvax.dec.com`
Fred Glover	`fglover@decvax.dec.com`
Ali Gohshan	
Joy Leima	`joy@sware.com`
Narayan Makaram	
Mark Saake	`saake@llnl.gov`
Carl Smith	`cs@eng.sun.com`

3.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Eight working groups met at the IETF in Santa Fe:

Directory Information Services Infrastructure

DISI is a Working Group that provides a forum to define user requirements in X.500. It is an offshoot of the OSI Directory Services Group and is a combined effort of the User Services Area and the OSI Integration Area of the IETF.

- Paper 1, "Executive Summary" (Weider, Reynolds, Heker). Defines issues DISI should be working on. This current draft has been revised twice since the Atlanta IETF. It is ready for Internet Draft submission, and on to FYI RFC publication.

- Paper 2, "Survey" (Lang, Wright). This document will undergo one last modification before publication. Additional DUAs were added that were inadvertently left out, as well as additional submissions. This document will be reissued as an Internet Draft, then submitted to the RFC Editor for FYI RFC publication.

Internet Anonymous FTP Archives

This is a new Working Group which met for the first time in Santa Fe. It is chartered to define a set of recommended standard procedures for the access and administration of anonymous FTP archive sites on the Internet.

IAFA attendees agreed on the Charter. Discussion then focused on two documents this Group intends to produce:

- "Anonymous FTP Site Administrator's Guide"
- "Anonymous FTP User's Guide"

The contents of these two documents were discussed, as well as a discussion on new technology issues. Newer technology issues were tabled for further discussion at a later date.

John Curran (BBN), Ellen Hoffman (Merit), and April Marine (SRI) volunteered to work on the "Anonymous FTP User's Guide" document.

Internet School Networking

This is a new Working Group which met for the first time in Santa Fe. It is chartered to facilitate the connection of the United States' K-12 (Kindergarten-12th Grade) schools, public and private, to the Internet, and school networking in general.

ISN's session gathered educators and Internet folks together. This meeting primarily focused on going over and refining the Charter, the goals and projected milestones.

Network Information Services Infrastructure

Dana Sitzler has resigned as co-Chair of NISI. April Marine has accepted the co-Chair position. Patricia Smith will remain as co-Chair.

The final review of this Group's Internet Draft, "Building a Network Information Services Infrastructure" was discussed. Inclusion of a security "verification" section in the document has been placed and agreed upon.

Continued discussion on where this Group should go from here - there was justification of additional action items/tasks that fall in NISI's realm.

NOC-Tool Catalogue Revisions

Gary Malkin has resigned as co-Chair to work with Tracy LaQuey Parker on the User-Glossary document. Darren Kinley has accepted the co-Chair position. Robert Enger will remain as co-Chair.

The "Son of NOCTools" Working Group is updating and revising their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

This Group has "one last call" out for submissions, and is continuing to accept additional "vendor gathering" for one more month. The document will be submitted as an Internet Draft, then on to the RFC Editor for FYI RFC publication.

User Documentation

The Userdoc Working Group will be preparing a revised bibliography of on-line and hard copy documents, reference materials, and training tools addressing general networking information and how to use the Internet. The target audience includes those individuals who provide services to end users and end users themselves.

(See the USWG minutes below for further information on this Group's current progress.)

User Glossary

Karen Roubicek has resigned as co-Chair of UserGloss. Gary Malkin has accepted the co-Chair position. Tracy LaQuey Parker will remain as co-Chair.

The User Glossary Working Group met and decided on the document format and updated goals and milestones. A draft document will be ready for review at the next IETF in

San Diego. A review and final draft will be presented at the IETF in Boston. The final document will be published shortly after the Boston IETF.

User Services Working Group

The User Services Working Group provides a regular forum for people interested in all user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet.

Agenda items included:

- Report on the RARE WG3 meetings held in Zurich, Switzerland. Reported by Joyce K. Reynolds.

- SIGUCCS draft - Presented by Martyne Hallgren. Written by ACM Siguccs Networking Taskforce. Document title - "Connecting to the Internet - what connecting institutions should anticipate",

- Revision of Userdoc Working Group - Presented by Lenore Jackson and Ellen Hoffman, Userdoc co-Chairs. Discussion focused on the revision of its Charter, objectives, future goals, and establishing procedures on updating the bibliography.

- QUAIL - presented by Gary Malkin. Gary Malkin and April Marine held a brief discussion with the USWG on the updating of the "Questions and Answers for New Internet Users".

3.9.1 Directory Information Services Infrastructure (disi)

Charter

Chair(s):
Chris Weider, clw@merit.edu

Mailing Lists:
General Discussion: disi@merit.edu
To Subscribe: disi-request@merit.edu
Archive: pub/disi-archive@merit.edu

Description of Working Group:

The Directory Information Services (pilot) Infrastructure Working Group (DISI) is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

The DISI Working Group is an offshoot of the OSI Directory Services Group, and, accordingly, is a combined effort of the OSI Integration Area and User Services Area of the IETF. The current OSIDS Working Group was chartered to smooth out technical differences in information storage schema and difficulties in the interoperability and coherence of various X.500 implementations. The DISI Group is concerned solely with expanding the Directory Services infrastructure. As DISI will be providing infrastructure with an eye towards truly operational status, DISI will need to form liaisons with COSINE, Paradyse, and perhaps the RARE WG3.

As a final document, the DISI Working Group shall write a Charter for a new working group concerned with user services, integration, maintenance, and operations of Directory Services, the Internet Directory User Services Group.

Goals and Milestones:

Done First IETF Meeting: review and approve the Charter making any changes necessary. Examine needs and resources for the documentation to be produced, using as a first draft a document produced by Chris Weider, Merit, which will be brought to the IETF. Assign writing assignments. Further work will be done electronically.

Jul 1991 Second IETF Meeting: review and approve documentation; review and approve Charter for the IDUS Group.

Aug 1991 Electronically review final draft of documentation, and, if acceptable, submit to IESG for publication.

Dec 1991 Third IETF Meeting: Declare success and reform DISI Group as IDUS group.

Internet Drafts:

"Interim Schema for Network Infrastructure Information in X.500 New name: Encoding Network Addresses to support operation ov", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-disi-netinfrax500-00.txt>

"An Executive Introduction to Directory Services Using the X.500 Protocol", 12/18/1991, Chris Weider, Joyce Reynolds, Sergio Heker <draft-ietf-disi-execdir-01.txt>

Request For Comments:

RFC 1292 "A Catalog of Available X.500 Implementations"

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Directory Information Services Infrastructure Working Group (DISI)

The Agenda was as follows:

1. Discussion and possible modification of Atlanta Minutes.
2. Ruth and Russ's document, mention work done and final form as the paper should be out as an Internet Draft and started up the track.
3. Chris, Joyce, and Sergio's paper. This has been revised twice since Atlanta, and is on its last round of reviews.
4. Discussion of and possible assignment of the "How to Join a Pilot" paper.
5. Discussion of and possible assignment of the "How to Set Up a DSA" paper.
6. Future Work???

How it went:

1. The Atlanta Minutes were discussed and Chris Weider gave an outline of the changes proposed to the Minutes. Ruth Lang noted that her comments on the Atlanta Minutes had not been reflected in the Minutes. She took an action item to forward these comments to Chris Weider [Ed. Note: she forwarded and Chris incorporated.]
2. Ruth Lang and Russ Wright's document was discussed next. They mentioned that they had made the changes which had been suggested at Atlanta, and it was noted by Joyce Reynolds that Jon Postel was eager to see this paper advanced to RFC status as soon as possible. Several modifications/additions were also discussed. Chris Weider wanted to know why the authors did not flesh out entries which had no information provided by the authors but which DISI would be able to supply the relevant information; Ruth mentioned that DISI should not be responsible for the implementation descriptions, as there were some concerns about DISI's liability. However, DISI should be responsible for the sections written by Ruth and Russ. Accordingly, we should not provide input to the implementation descriptions without explicit authorization from the responsible organization or implementation description author. Christian Huitema suggested the inclusion of benchmarks for implementations. After some positive discussion, it was decided that this was something that we should do, but time would be needed to develop benchmarking tests, distribute the tests to implementors, and gather results for inclusion. Therefore, this version of the catalog document will be issued without benchmarks, but the tests should be developed and recommmended for inclusion in the next version of the document. Richard Colella suggested adding

a 'date of last update' field to each entry, to let the reader determine how recent the information is. This suggestion was accepted by DISI.

3. Chris, Joyce, and Sergio's document was discussed next. Many people had not yet had a chance to read it, but those who did had several comments. The final upshot was that the people who read it thought that it should focus less on technical details and more on "What can it do for me??". Chris pointed out that the focus on technical details was mandated to some extent by the necessity of explaining what makes X.500 so powerful, but agreed that the focus should and would be shifted.

4. Chris started discussion of the "How to Join a Pilot" and "How to Set Up a DSA" papers. Several people were of the opinion that there should also be an "X.500 Case Studies" paper written to show people how to set up and use X.500, and illustrate some of the stumbling points in the deployment. It was also suggested that the focus of DISI's "How to Join a Pilot" paper should really be a "Pilot Catalog", listing points of contact to join each pilot, and some additional information. After discussion of all three papers, (in addition to the paper which had been mentioned on benchmarking), Ruth Lang mentioned that she thought that we should not assign any more papers at this time, and that she would take an action item to get the European versions of these papers and post summaries to the list. The Group seemed to agree with this sentiment.

5. The Group also discussed the creation of an "X.500 Bibliography" document; i.e., a document that contains pointers to Internet-relevant technical papers, books, and Internet Drafts, or RFCs. The writing of this document was not assigned.

Attendees

William Biagi	bbiagi@cos.com
Robert Blokzijl	K13@nikhef.nl
Steve Bostock	steveb@novell.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Colella	colella@osi.ncsl.nist.gov
Robert Cooney	cooney@wnyose.nctsw.navy.mil
John Demco	demco@cs.ubc.ca
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@seka.scc.com
Michael Erlinger	mike@lexcel.com
Jisoo Geiter	geiter@gateway.mitre.org
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Susan Hares	skh@merit.edu
Russ Hobby	rdhobby@ucdavis.edu
Alton Hoover	hoover@nis.ans.net
Christian Huitema	christian.huitema@sophia.inria.fr
Lenore Jackson	jackson@dftnic.gsfc.nasa.gov

Darren Kinley	kinley@crim.ca
Edward Krol	e-krol@uiuc.edu
Ruth Lang	rlang@nisc.sri.com
Walter Lazear	lazear@gateway.mitre.org
Peter Liebscher	plieb@sura.net
April Marine	april@nisc.sri.com
Glenn Michel	gym@lanl.gov
Robert Purvy	bpurvy@us.oracle.com
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov
Sze-Ying Wuu	syww@thumper.bellcore.com
Wengyik Yeong	yeongw@psi.com

3.9.2 Internet Anonymous FTP Archives (iafa)

Charter

Chair(s):

Peter Deutsch, peterd@cc.mcgill.ca
Alan Emtage, bajan@cc.mcgill.ca

Mailing Lists:

General Discussion: iafa@cc.mcgill.ca
To Subscribe: iafa-request@cc.mcgill.ca
Archive: pub/iafa-archive@archive.cc.mcgill.ca

Description of Working Group:

The Internet Anonymous FTP Archives Working Group is chartered to define a set of recommended standard procedures for the access and administration of anonymous ftp archive sites on the Internet. Such a set of procedures will provide a framework for:

(a) allowing the inexperienced Internet user the ability to more easily navigate the hundreds of publically accessible archive sites; and,

(b) allowing users and network-based tools to retrieve specific site information such as access policies, contact information, possible areas of information specialization, archived package descriptions, etc., in a standardized manner.

Particular emphasis will be placed on the possible impact of these procedures on the FTP site administrators.

Attention will be paid to the impact of newer archive indexing and access tools on the operation of such archive sites. A set of suggestions will be offered to allow archive site administrators to better integrate their offerings with such tools as they are developed.

The security of the anonymous FTP site configuration will also be considered to be an integral part of this document. It is expected that remote management of the archives will be adequately handled by existing network management procedures.

Goals and Milestones:

Nov 1991 First IETF Meeting: review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended procedures and impact on site administrators. Assign writing assignments for the first draft of the documents.

Mar 1991 Review first draft and determine necessary revisions. Follow up discussion will occur on mailing list.

Jun 1991 Make document an Internet Draft. Continue revisions based on comments at IETF and on the mailing list.

Nov 1992 Fourth IETF meeting. Review final drafts and if OK, give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Alan Emtage/McGill

Minutes of the Internet Anonymous FTP Archives Working Group (IAFA)

The IAFA meeting Agenda was approved. The Charter was discussed and a sentence of the following type was suggested and approved for addition to the Charter:

"This Working Group will also look into the FTP Protocol functionality (though not the protocol itself) as it concerns anonymous FTP to see if changes can be recommended and forwarded to the appropriate Working Group or Area Director".

The "User's Guide To Anonymous FTP" document was discussed. It was decided that many such documents already exist and that writing a new one from scratch would be redundant. John Curran, Ellen Hoffman and April Marine volunteered to draw together the current documents and provide a draft User's Guide to be distributed on the IAFA mailing list before the next IETF meeting.

The "Anonymous FTP Site Administrators' Guide" was discussed as a useful document for the recommended operating procedures for an anonymous ftp site. Alan Emtage volunteered to write a draft to be distributed on the IAFA mailing list before the next IETF meeting.

Attendees

Charles Bazaar	bazaar@emulex.com
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@seka.scc.com
Hans Eriksson	hans@sics.se
Barbara Fraser	byf@cert.sei.cmu.edu
Neil Haller	nmh@bellcore.com
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
J. Paul Holbrook	holbrook@cic.net
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Darren Kinley	kinley@crim.ca
Edward Krol	e-krol@uiuc.edu
Eliot Lear	lear@sgi.com
Peter Liebscher	plieb@sura.net
April Marine	april@nisc.sri.com
Glenn McGregor	ghm@merit.edu

Chris Myers	chris@wugate.wustl.edu
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Harri Salminen	hks@funet.fi
William Simpson	Bill_Simpson@um.cc.umich.edu
Dana Sitzler	dds@merit.edu
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Wengyik Yeong	yeongw@psi.com

3.9.3 Internet School Networking (isn)

Charter

Chair(s):

John Clement, clement@educom.edu
Arthur St. George, stgeorge@bootes.unm.edu
Connie Stout, CStout@tea.texas.gov

Mailing Lists:

General Discussion: cosndisc@bitnic.educom.edu
To Subscribe: listserv@bitnic.educom.edu (Sub: cosndisc)
Archive: listserv@bitnic.educom.edu

Description of Working Group:

The Internet School Networking Working Group is chartered to facilitate the connection of the United States' K-12 (Kindergarten-12th Grade) schools, public and private, to the Internet, and school networking in general.

It is critically important that national networking for K-12 education proceed along established lines of protocol, using existing network structures. The Working Group's first priority will be to establish guidelines for specialized user interfaces. K-12 networking will also require other support services, such as directories, online and hotline help, specialized training programs and collaborative projects with instructional and curriculum groups, disciplinary groups and postsecondary institutions.

While the initial focus is school networking in the U.S., the Working Group will coordinate its efforts with similar activities in other countries and regions of the world.

Goals and Milestones:

Nov 1991 Meet for the first time at IETF and establish approval of Charter. Examine the status of projects in process when Working Group was created. Begin work on list of deliverables.

Jan 1991 Release X.500 "K-12 People Directory" version in collaboration with Merit. Develop plans and milestones for K-12 Resources Directory.

Mar 1991 First draft of information packet document for computing directors to assist them in connecting K-12 schools. First draft of user interface guideline statement.

May 1991 Release X.500 K-12 Resource Directory version in collaboration with Merit. Present final draft guideline statement.

CURRENT MEETING REPORT

Reported by Arthur St. George/U New Mexico

Minutes of the Internet School Networking Working Group (ISN)

The Internet School Networking Working Group is chartered to facilitate the use of the Internet by United States' K-12 (Kindergarten-12th Grade) schools, public and private, and school networking in general.

It is critically important to establish guidelines for the integration of K-12 educational networking into the Internet. The Working Group's first priority will be to establish connectivity models, including costs, for distribution to all parties concerned with connection of K-12 schools to the Internet. This Working Group will also provide guidelines to technical solutions for other support services such as directories, online and hotline help, specialized training programs and collaborative projects with instructional and curriculum groups, disciplinary groups and post-secondary institutions.

In pursuit of these goals, the Working Group will work closely with other organizations such as the Consortium for School Networking, and the Coalition for Networked Information. Finally, the Working Group will act as technical consultant to K-12 organizations and groups seeking connection and access to the Internet.

While we agreed that the first task for the Working Group was to produce a document which outlined multiple connectivity models for K-12 connection to the Internet, we also briefly discussed one other project which should be raised, that is the user interface document. The idea here is to discuss and come to agreement on what would constitute standards for K-12 interface(s) to the Internet. These run the gamut from e-mail to bulletin boards. Competing with this project for number 2 priority status is an RFC which provides the 25 most commonly asked questions about K-12 connection to and use of the Internet.

Attendees

Joe Blackmon	blackmon@ncsa.uiuc.edu
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Jack Hahn	hahn@sura.net
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Jeff Hayward	j.hayward@utexas.edu
Russ Hobby	rdhobby@ucdavis.edu
J. Paul Holbrook	holbrook@cic.net
Ole Jacobsen	ole@csli.stanford.edu
Kenneth Klingenstein	kjk@spot.colorado.edu
David Korns	David_Korns@auo.mts.dec.com
Edward Krol	e-krol@uiuc.edu
Ruth Lang	rlang@nisc.sri.com

Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Frederick Mueller	fmueller@telebit.com
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Miguel Sasson	sasson@xylogics.com
William Schrader	wls@psi.com
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Arthur St. George	stgeorge@bootes.unm.edu
Joanie Thompson	joanie@nsipo.nasa.gov
Andrew Veitch	aveitch@bbn.com
David Wasley	dlw@berkeley.edu
Chris Weider	clw@merit.edu
Jonathan Wenocur	jhw@shiva.com

3.9.4 Internet User Glossary (userglos)

Charter

Chair(s):

Tracy LaQuey Parker, tracy@utexas.edu
Gary Malkin, gmalkin@ftp.com

Mailing Lists:

General Discussion: usergloss@ftp.com
To Subscribe: usergloss-request@ftp.com
Archive:

Description of Working Group:

The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the Charter as necessary. Discussion of Userglos Working Group Chair nominations submitted by USWGers.
TBD	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
TBD	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.

3.9.5 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@ans.net
Darren Kinley, kinley@crim.ca

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu
Archive:

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.
- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.
- Identify additional new or improved tools that have become apparent since the last compilation of the reference document.
- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Done — Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.

Aug 1991 — Draft of catalog will be prepared, draft to be reviewed and modified. Initiate IETF Internet Draft review process by submission of a "Son of NOCTools" catalog draft to IESG Secretary.

Dec 1991 — Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI RFC for publication.

CURRENT MEETING REPORT

Reported by Darren Kinley/CRIM

Minutes of the NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

Agenda

The NOCTOOL2 Working Group met at the 22nd IETF in Santa Fe. The meeting served primarily to report on the revision of the RFC 1147/FYI 2, and again served as a forum for discussion of the "living documents" problem. The Agenda was as follows,

- Working Group Co-Chair shuffling
- Current status
- NOCtools' resolution of the "living documents" problem
- Planned wrap up of revision process

Discussions

Co-Chair Shuffling

The replacement of Gary Malkin by Darren Kinley during the month of October was announced. Thanks to Gary for his efforts.

Current Status

Bob Enger, Darren Kinley, and Gary Malkin each fumbled putting things behind schedule. After much whining, pleading, and shaming their lives were spared. Schedules have been revised.

Revised Schedules

Internet Draft to be prepared for the month of February, its availability (including TOC) widely announced, late submissions to be included as required, and finally document to be submitted to the RFC editor. The delay from announcement as Internet Draft to submission to the RFC editor will depend on the number of late entries received, the feedback from the community, and will be left to the discretion of the RFC editors.

"Living Documents"

Technical and philosophical aspects of this problem were discussed extensively. Everyone agreed that "high-tech" solutions were still not available nor would they be in the very near future. The Group decided to schedule a BOF on this topic at the upcoming IETF in San Diego.

Two sets of mechanisms and procedures to aid in the maintenance and distribution of this document were put forward, improved, and adopted by the Group.

1. Anonymous FTP retrieval with automated email submission.

   The official/edited document will be made available via anonymous FTP as well as an automated electronic distribution tool (listserv) as a whole or in pieces. Entries can be submitted via electronic mail and made immediately available in an unofficial/unedited portion of the distribution. An editor will occasionally verify these last entries and include them in the official document. Possible homes might include Merit or Washington University.

2. Usenet news.

   A news hierarchy similar to this was discussed.

   alt.noctools.announce (moderated) A small number of entries will be posted daily so that during a month the entire document will be posted once.

   alt.noctools.wanted (unmoderated) Questions of the nature "I'm looking for..." will be posted to this Group.

   alt.noctools.new (unmoderated) New tools can be announced here. A Gateway between this Group and the automated email submission address should be put in place.

   alt.noctools.bugs (unmoderated) Any patchs to tools can be posted here. Hopefully, they will also be sent to the tool creator.

   alt.noctools.d (unmoderated) General discussion.

   Resources needed to support some of these functions could be located very near to the FTP retrieval and automated submissions service.

Other Items

- The majority of cases are handled in the current plan, but what about the small number of sites without even UUCP access?

- BOF on "living documents" problem.

Action Items

Robert Enger, Darren Kinley: Contact remaining people with entries, continue to solicit and accept new entries, prepare draft document.

Joyce Reynolds: Consider funding for maintenance of "living documents" problem vis-a-vis the Internet Society.

Darren Kinley, Joyce Reynolds: Compile list where NOCtools catalog availability announcement will get the widest possible distribution.

Robert Enger, Darren Kinley, Chris Myers, Mike Patton: Properly define Usenet newsgroups to be created and write charters for these groups as required.

Chris Myers: Make a home for anonymous FTP and automated electronic submissions and distribution tool.

Attendees

Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@seka.scc.com
Darren Kinley	kinley@crim.ca
Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Chris Myers	chris@wugate.wustl.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Brad Passwaters	bjp@sura.net
Michael Patton	map@lcs.mit.edu
Joyce Reynolds	jkrey@isi.edu
Brad Steinka	brad@python.eng.microcom.com

3.9.6 Network Information Services Infrastructure (nisi)

Charter

Chair(s):
April Marine, april@nisc.sri.com
Pat Smith, psmith@merit.edu

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

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. The work will begin with existing NIC functions and services and should build upon work already being done within the Internet community. It should address areas such as common information formats, methods of access, user interface, and issues relating to security and privacy of Internet databases.

Goals and Milestones:

Done	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.
Done	Complete draft for phase 2 suggesting cooperative agreements for NICs.
Done	Revised draft document ready for Working Group review. Document defines NIC functions and suggests some standardizations for NIC services, as well as offers new mechanisms for exchanging information between NICs.
Done	Document submitted as Internet Draft for comment from a wider internet audience.
Done	Working Group discussed current Internet draft and suggested minor revisions. Decision made to continue Working Group activity beyond this document.
Nov 1991	First document released as informational RFC. Outline and discuss new NISI tasks at IETF meeting.
Jul 1992	Write a document explaining the security issues of privacy and accuracy in Internet databases. Publish as an informational RFC.

Internet Drafts:

"Building a Network Information Services Infrastructure", 07/15/1991, D. Sitzler, P. Smith, A. Marine <draft-ietf-nisi-infrastructure-00.txt>

CURRENT MEETING REPORT

Reported by Patricia Smith/Merit

Minutes of the Network Information Services Infrastructure Working Group (NISI)

NISI Internet Draft Document

All comments are in on this draft. April has made the final tweaks and will format it and submit it as an RFC. Yea!

NIC Profiles X.500 Directory

At the Atlanta meeting in July 1991 the information being sought for the Network Information Center (NIC) Profiles had been discussed at length. At the Santa Fe meeting a suggestion was made that we not get hung up with that type of discussion again but rather make an effort to put information on more NICs up in X.500 and allow folks on the NISI and DISI lists to play around with it. Pat indicated she would be pro-active in this area and begin to knock on NIC doors to solicit their participation.

In addition, the effort to create a friendly user interface environment in order to facilitate entry and updating of information will continue.

Overlap with User Connectivity Problems Working Group Chaired by Dan Long

The UCP Working Group is looking at gathering NOC information via template for inclusion in some sort of directory. There is general agreement that NISI should have some liaison activity with UCP in order to avoid duplication of effort. Pat sat in on the UCP meeting on Wednesday and told them of the NISI interest. UCP members agree that we should work together and Pat agreed to stay in touch with Dan Long as the NIC directory activity develops.

Request for a Database Security Document

After the Atlanta IETF, the IESG tasked the NISI Working Group with writing an informational document that makes some recommendations regarding the need for accuracy and privacy of data in databases maintained by network information centers. After some discussion, it was decided that April Marine, J. Paul Holbrook, and John Curran will work up a draft of this RFC which is expected to be only a few pages long. In addition, the section on database accuracy that was added to the current NISI Internet Draft document, which addresses much the same concerns, will be retained in that document.

Internet Society (ISOC) Request from Vint Cerf

Vint Cerf suggested that the NISI Working Group collect information regarding various Internet products and services. Evidently this request was prompted by the fact that the ISOC has received questions about the Internet and has felt the need of such a collection. The Working Group was reluctant to take this on in the form suggested because the task is one that each NIC does currently for its own constituency and the task seemed redundant. In addition, the NIC.DDN.MIL is putting together a directory of directories, lending more weight to the fact that NISI should not do so as well. The Working Group agreed that one good strategy for ISOC to use with such callers would be to refer them directly to an existing Internet information center (the NNSC was specifically mentioned) for the answers to such general questions as "What is the Internet?" "What's on the Internet?" and "How do I join the Internet?"

However, while this suggestion solved part of the problem, it led to a related discussion regarding the problem of easy discovery of information available about and via the Internet. Currently, there is no means for either a NIC or a user to easily determine what information is available and where. Neither is there an easy means for alerting users to information newly available. This discussion was a natural lead in to the following Agenda topic.

Should NISI be Dissolved?

There was discussion of whether or not NISI should be terminated at the next IETF and emerge as a new working group. We decided to take to the mailing list a discussion of whether or not NISI has accomplished what it set out to do. The consensus at first glance appears to be that if at the next meeting in March the RFC has been published and other projects are either completed or well underway, then the San Diego meeting would be the last one for the NISI Group. The related issues of information discovery and delivery could be handled in a new working group in the User Services Area which would tentatively be called Network Information Delivery (NID).

NID is seen as filling a very important and timely need since, at present, everyone is trying to figure out the best and most efficient means of locating and delivering information (X.500, WAIS, etc.). It is felt that we can provide critical direction in this area as far as understanding and application of the various types of directory services currently available.

Over the next couple of months, then, the mailing list needs to discuss the questions of whether the NISI group has fulfilled its Charter or what else it has on its plate, what the Charter of this new working group would be, and various matters related to the work we've recommended in the NISI document. Plus we'll need feedback and comments on the database security draft. April and Pat will take it upon themselves to try to get the discussions going on the list.

Attendees

Miriam Amos Nihart	miriam@decwet.zso.dec.com
Robert Blokzijl	K13@nikhef.nl
James Codespote	jpcodes@tycho.ncsc.mil
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
J. Paul Holbrook	holbrook@cic.net
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Alton Hoover	hoover@nis.ans.net
Geoff Huston	g.huston@aarnet.edu.au
Darren Kinley	kinley@crim.ca
Carol Lambert	cjlx@cornella.cit.cornell.edu
Ruth Lang	rlang@nisc.sri.com
Peter Liebscher	plieb@sura.net
April Marine	april@nisc.sri.com
Ellen McDermott	emcd@osf.org
David Minnich	dwm@fibercom.com
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Harri Salminen	hks@funet.fi
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Nancy Yeager	nyeager@ncsa.uiuc.edu

3.9.7 User Services (uswg)

Charter

Chair(s):

Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:

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

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.

Request For Comments:

RFC 1150 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. notes"

RFC 1177 "FYI on Questions and Answers - Answers to Commonly Asked "New Internet User" Questions"

RFC 1206 "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions"

RFC 1207 "Answers to Commonly asked "Experienced Internet User" Questions"

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

Minutes of the User Services Working Group (USWG)

Announcements

- New Working Group - Internet School Networking (ISN)
- New Working Group - Internet Anonymous FTP Archives (IAFA)
- New Working Group - User-Doc Revised

Discussions/Reports:

1. Report on the RARE WG3 meetings held in Zurich, Switzerland. Reported by Joyce K. Reynolds.

   The RARE WG3 USIS Working Group and the IETF User Services Working Group are working in parallel. Jill Foster, Chair of USIS, concurs. The RARE WG3 Working Group provides first level services to end users. The IETF User Services Working Group has traditionally provided second level services (i.e., providing documentation to people who provide services to end users). RARE WG3 will continue to provide first level services, while IETF User Services will provide "information packets" and other user services documentation. RARE will delegate their members to actively work on IETF working groups.

2. Revision of User-Doc Working Group - Presented by Lenore Jackson & Ellen Hoffman, User-Doc Co-Chairs.

   Discussion focused on the revision of its Charter, establishing procedures on updating the bibliography, additional objectives, and future goals.

3. QUAIL - Presented by Gary Malkin.

   Brief discussion and comments of an updated Quail document for "new Internet users". The updating of this document is necessary due to the transition of DDN NIC services from SRI to GSI. The current structure/formatting of the document was also discussed.

4. SIGUCCS Draft - Martyne Hallgren. Written by ACM SIGUCCS Networking Taskforce.

   Martyne's taskforce intends to submit this draft into the "Internet-Drafts" process, with the end result being an FYI/RFC. She brought this document into the USWG forum for comments and feedback.

   Document title - "Connecting to the Internet - What Connecting Institutions Should Anticipate".

   Abstract: This memo outlines the major issues an institution should consider in the decision and implementation of a campus connection to the Internet. The list of issues is not exhaustive but rather this document should alert decision makers to major concerns to be addressed in the critical phases of an institution's full participation in the Internet community.

Martyne received positive comments on this draft. This draft is akin to the "INCH" task on the USWG list of "things to do".

Attendees

Robert Blokzijl	K13@nikhef.nl
Vinton Cerf	vcerf@nri.reston.va.us
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@seka.scc.com
Hans Eriksson	hans@sics.se
John Gong	jgong@us.oracle.com
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Ellen Hoffman	esh@merit.edu
J. Paul Holbrook	holbrook@cic.net
Alton Hoover	hoover@nis.ans.net
Geoff Huston	g.huston@aarnet.edu.au
Lenore Jackson	jackson@dftnic.gsfc.nasa.gov
Ole Jacobsen	ole@csli.stanford.edu
Scott Kaplan	scott@ftp.com
Darren Kinley	kinley@crim.ca
Carol Lambert	cjlx@cornella.cit.cornell.edu
Ruth Lang	rlang@nisc.sri.com
Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Tom Sandoski	tom@concert.net
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Wengyik Yeong	yeongw@psi.com

Chapter 4

Network Status Briefings

4.1 DDN MILNET Report

Presented by Kathleen Huber/BBN

DDN MILNET

Kathleen Huber

November, 1991

BBN Communications
A Division of Bolt, Beranek and Newman, Inc.

AGENDA

- Internet Growth
- Mailbridges

INTERNET GROWTH

INTERNET GROWTH SUMMARY

- 3161 Networks Advertised By BMIL~~DEC~~ DCEC on 11/8/91
- 10,633 Networks Registered
- 37,397 Network Numbers Assigned

NUMBER OF NETWORKS
LINEAR
DECEMBER 1983 - NOVEMBER 1991

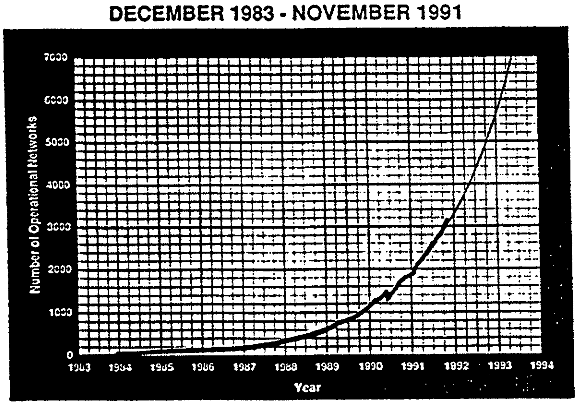

NUMBER OF NETWORKS
LOGARITHMIC
DECEMBER 1983 - NOVEMBER 1991

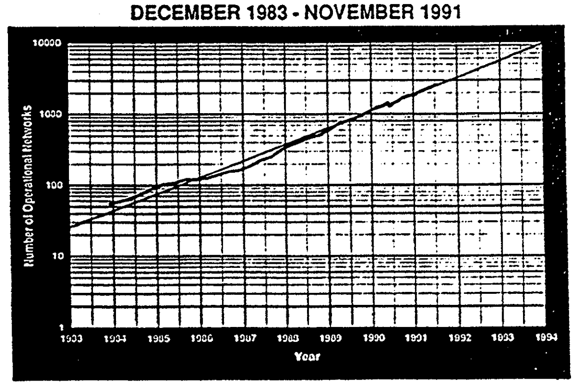

INTERNET PROBLEMS

- FIX-East cisco was CPU-bound while interacting with NSFNET NSS, possibly due to a NSS fragmentation problem in EGP updates.
- T1 line problems between MITRE-cisco and FIX-East cisco.
- NIC moved from SRI on the West Coast to GSI on the East Coast. It did not obtain a T1 line to NSFNET until November 7.

MAILBRIDGES

EGP NEIGHBOR COMPARISON

	DIRECT NEIGHBORS				
	Feb 91'	Apr	Jul	Aug	Oct
BMILAMES	61	52	54	57	48
BMILBBN	62	78	89	88	104
BMILDCEC	88	95	155	159	150
BMILISI	60	58	63	77	127
BMILOAK (BMILLBL)	78	114	177	181	168
BMILMTR	56	57	57	64	51
BMILRAN	50	70	76	79	40

TRAFFIC SUMMARY COMPARISON

	Avg. Pkts/Day Forwarded		Avg. Pkts Dropped	
	March - July	Aug - Oct	March - July	Aug - Oct
BMILAMES	4,417,258	3,757,744	0.06%	2.21 %
BMILBBN	302,469	221,243	0.0%	.30 %
BMILDCEC	432,644	316,666	0.0%	.50 %
BMILISI	238,492	319,078	0.0%	.45 %
BMILOAK (BMILLBL)	300,352	215,333	0.05%	.25 %
BMILMTR	3,661,255	3,181,024	0.03%	4.77 %
BMILRAN	232,519	205,527	0.02%	.12 %

MAILBRIDGE PROBLEMS

- "Limited Route Distribution" fails when the default gateway fails to obtain routes.

  (This will be fixed in Patch 13)

- Periodic resource problems on the Mailbridges are causing SPREAD updates to be improperly assembled

SHORT-TERM SOLUTIONS

- Temporarily by-passing the NSFNET framentation problem by filtering out AS 701; thereby reducing the EGP update size. (Connectivity to AS 701 is maintained through an additional EGP session with TWB gateway at FIX-East)
- Attempting to run BGP between the FIX-East cisco on the SURANET, though unable to do so at this time due to memory limitations in that cisco.
- Preparing a Mailbridge patch to fix the "Limited Route Distribution" bug.

LONG - TERM SOLUTIONS

- Implementing strategy for long-term growth effects
  - Advertisement of network 0 to limit size of EGP updates
  - Integrating T/20s with the Mailbridges
- Fix-Points
  - Upgrade Fix-East ciscos hardware so that BGP can be configured
  - Conversion to DISN-NT -- Ft. Detrick
  - Similar conversion planned for BMILAMES on a later date

4.2 Ebone Report

Presented by Bernhard Stockman/NORDUnet

Networking and Electronic Mail in Eastern European Countries
Claudio Allocchio, INFN

This small note summarizes the results of a survey performed among participants to the First International School on Network Management and Analysis, held in Trieste (December 4-16, 1990). The situation can differ a lot from country to country: some already have networking structures available, some others still lack the basics, but in all countries there is a strong push to establish connections and to open contacts with other nations. The major common problems are the bad financial situation and licensing. Also know-how is a lacking point in many situations. A final table tries to compare the different situations. This note is only a first and non-exhausting survey, but I hope it can help. Let's now look more in details the various situations.

Hungary:

There is currently available a national experimental X.25 network called IIF "Information Infrastructure Network" of the Hungarian Academy of Sciences. This network is a semi-private infrastructure for a closed community (the R&D community) but it is operated by Hungarian PTT. On this network volume charging is applied. There are currently one EARN node and one EUnet node, both interfaced with IIF network. There are plans to increase a lot of networking capabilities within the country in the near future; during 1991 the public X.25 service, interconnected with other PSDNs will be established. The intention is to run OSI services on the X.25 network, but also a number of tcp/ip point to point connections are foreseen. A tcp/ip connection to HEPnet via CERN is also in preparation. E-mail in Hungary is available via the ELLA system, with an x.400-like user agent developed locally, running on top of IIF X.25. There is a gateway connecting ELLA with EUnet and EARN. The plans are to implement an X.400 MHS as soon as the public x.25 service will be available.

Poland:

Since a few years there is an experimental x.25 network connecting 8 sites via leased lines. This is supposed to be the kernel of the Country Academic Computer Network (KASK) which is supposed to be completed by mid 1991, covering 16 sites, and having gateways to EARN and DFN. There is already an EARN connection in Warsaw and Wroclaw, and two DECnet connections to CERN from Cracow and Warsaw. E-mail: EARN and HEPnet nodes are reachable with their native protocols, but there is no current gateway to the national x.25 infrastructure.

Romania:

There are currently only a few and poor quality connections in the country at 4800 bps joining single institutions, and no international links are available. However the PTT started a program to make available networking capabilities (x.25) in the country. It will take some years. For international connectivity satellite links are probably the correct solution.

USSR:

The USSR situation is very fragmented and there is not what can be called a 'nation wide' networking strategy. Public X.25 service (IASNET) is available and is connected to the major PSDNs in the world, using both ground and satellite links. The research community had free access to this facility, but starting from 1991 each Institute will have to pay the service and this is a major problem. The largest research institutes have some small local private X.25 networks, and DECnet technology is also used locally. Leased lines are obtainable, but due to the country's geographical situation satellite links are more reliable. Any institute is trying to provide its user with connectivity with the foreign countries, but often local connectivity is much less implemented, showing real cooperation problems exist. The PAD access to remote sites providing services is the most diffused situation. Apart from public X.25 the cooperative "DEMOS", settled in Moscow, is the recognized Internet entry point for domain '.su' and has a link to Helsinki (Finland). It provides access using UUCP to about 30 sites. Many other institutes join the services via dial-up connections (remote login) to DEMOS and using Kermit to transfer data to their remote account. The DEMOS services are expanding to reach more sites and to improve quality. There is also a planned EARN link from Moscow to Poland and plans to establish DECnet connections on top of public x.25 links to the HEPnet community. The most used connectivity method however is still remote login to some foreign institute to access its services. Electronic mail is available from DEMOS sites, via EUnet, but most of other institutes access to it indirectly via remote login to collaborating remote partners.

Czecholslovakia:

The networking in the country is currently based on 9.6 Kbps lines used for the national EARN backbone and on 2.4 Kbps dial up lines used within the EUnet community. International connectivity is assured by an EARN link joining Prague to Linz (Austria) and an EUnet link joining Bratislava to Vienna. In 1992 the public X.25 service will be available with international connectivity. The intention is to build a national backbone enabling multi- protocol transport (at least SNA, TCP/IP and X.25) and to have a 64 Kbps link to Internet. Electronic mail is available directly, using both EARN ans EUnet facilities.

```
                 Summary of Networking is Eastern European Countries

                 | Hungary              | Poland              | Romania
--------------------------------------------------------------------------------
 Public X.25     | Expected 1991        | NO (experimental    | NO
```

available	IIF (semi-private network for academic community, 200 DTEs)	network 8 nodes)	
Pad Access available	YES	NO	NO
Leased lines available	YES (but takes long time)	YES	YES (very poor quality)
Available Speed	up to 9600 bps	up to 9600 bps	up to 4800 bps
Satellite/ Ground lines	ground	ground	ground
DECnet net exists	some Local implementations	some local one link to CERN	NO
TCP/IP net exists	some local implementations	1 local implementation	NO
EARN net exists	one link to TU-Wien	one int'l link, 5 nodes	NO
EUnet net exists	one link	NO	NO
Other net exists	IIF nation wide X.25	exper. X.25 nation wide	NO
Planned DECnet	YES nation wide	YES some sites	--
Planned TCP/IP	YES nation wide	YES nation wide	--
Planned EARN	YES nation wide	YES (SNA) nation wide	YES
Planned EUnet	YES nation wide	--	--
LAN technology	NOVELL, DECnet	NOVELL	NOVELL

E-mail available	YES directly	YES directly	NO
E-mail protocols	ELLA, UUCP RSCS	VMSmail, RSCS	N/A

	URSS	Czechoslovakia
Public X.25 available	YES, IASNET connected to most of PSDN in Europe and USA	NO (expected in 1992)
Pad Access available	YES (x.21) but not from all sites	NO
Leased lines available	YES, but sometimes with poor quality	YES
Available Speed	up to 9600 bps	up to 9600 bps
Satellite/ Ground lines	ground & satellite	ground
DECnet net exists	some Local implementations	NO
TCP/IP net exists	NO	NO
EARN net exists	NO	one int'l link, 5 nodes
EUnet net exists	YES, link to Finland	YES
Other net exists	--	--
Planned DECnet	YES nation wide	NO

Planned TCP/IP	YES nation wide	YES nation wide
Planned EARN	YES nation wide	YES nation wide
Planned EUnet	YES nation wide	--
LAN technology	Ethernet, DECnet	Ethernet
E-mail available	YES, directly from some sites	YES directly
E-mail protocols	UUCP	UUCP, RSCS

IP-Connections to Eastern Europe - Peter Rastl

I got a copy of your recent correspondence with Stefan Fassbender concerning EASInet and Eastern Europe. As I think, you should know about our Austrian activities in this respect, let me first introduce myself: I am the Director of the Vienna University Computer Center. Our computer center is one of the EASI sites and uses for this reason an EASInet line (currently still 64 kbit) to CERN. We provide Internet connectivity to all Austrian universities and coordinate TCP/IP networking in Austria. (Scientific networking in Austria in general has still some need for better coordination - e.g. Austria's national EARN node is for historical reasons located elsewhere, at the University of Linz, and has its own leased lines).

You will know about IBM's "Academic Initiative for Central and Eastern Europe", which has brought an IBM 3090 to Prague, Budapest and Warsaw, with other locations to follow. IBM has also announced to finance network connections from theses sites to Vienna University, where they should be connected to the international networks. I have signed a contract with IBM to support these universities in Eastern Europe with the establishment of their computing infrastructure, which means that our staff has helped with the installation of the 3090 systems in those countries and given numerous courses on topics of relevance, including of course the various aspects of networking.

In addition to IBM's offer to pay for network connections to these countries our Austrian Ministry of Science and Research responded positively to my proposal to establish 64 kbit TCP/IP connections to Czechoslovakia, Hungary and Poland and is willing to pay the

Austrian part of the PTT tariffs of these lines and to support each of these countries with an appropriate CISCO router.

I have sent this offer to the Technical University Prague (Ohera, Gruntorad), the Technical University Bratislava (Horvath), the University of Economy Budapest (Daruhazy) and the University of Warsaw (Wegrzynowski). All of them are thoroughly interested in our initiative, the feasibility of the requested 64 kbit links is probable, but still not clear, as of today. Although we cannot expect that the universities in Eastern Europe become EASI sites (with a 3090-600 VF) - currently they are not even allowed by COCOM to get any access to vector processing - it seems quite likely that these countries can get reasonable access to Internet and that this could be realized with the aid of IBM's academic initiative. In this case I am quite sure that IBM will support these countries with their EASInet infrastructure.

I hope that you regard our activities for Eastern Europe useful and I would like to give you more information on these activities if necessary. If I can offer further cooperation, please let me know.

WAN in Poland - Rafal Pietrak

Introduction

Although this report is supposed to reflect the status of Wide Area Networking in Poland, I am involved only in IP activities, so other protocols my not be properly treated here.

X.25

Following large academic effort during the 80's, to develop packet switching network in Poland there is a significant X.25 infrastructure among academic sites in southwest region, in particular in Wroclaw and Katowice area. The project was called National Academic Computer Network (KASK in Polish) and it provided mail and remote login services among major mainframe computers there. The PAD support was missing.

There is a proposal for a Germany government sponsored project to connect Wroclaw to the DFN X.25. It would include outfitting of three sites in southwest area of Poland with X.25 packet exchanges (frame relays) together with PADs.

Polish PTT is now running project called POLPAK that aims at providing widely available public X.25. PTT doesn't say whether it will be available at all before the end of next year.

EARN

EARN is present in Poland for one year now. The national node – PLEARN – was established 17th July 1990. Since then it grown steadily within the country. A strategy of stretching EARN network to all interested cities instead of widening the usage of the established nodes, created a network of long distance leased lines dedicated to a computer use of Polish academic community. An Academic and Research Computer Network group (NASK in Polish) was established to manage it.

October 1991, IBM East European Countries Initiative provided us a 3090 mainframe to replace a slower one for the PLEARN node. The 3090 came without the TCP/IP software, contrary to what was in contract. We are now working on providing a wider access to the machine without the networking software on it.

Together with the 3090, IBM claims: it'll provide us an international leased line to Viena University. Unlike the 3090 computer, the line will not be accepted unless it'll carry Internet Protocol (, too).

Internet

There is a number of IP LANs in Poland resulted from recent relaxation of COCOM restrictions concerning Poland and flood of computers with banded TCP/IP software – like SUN's – that followed. These LANs however until recently were not interconnected and usually even didn't have proper net-numbers from hostmaster@nic.ddn.mil, neither.

This changes now; We start getting net-numbers for our LANs. We are cooperating with other European Internet networkers within RIPE framework. Our networks are properly recognized/routed within the Europe. There are some exceptions though with sites that are connected to the Europe via the USA. Polish networks are not yet recognized there.

The following network-numbers are now assigned to some Polish organizations:

```
192.102.225.0 - ICPOLIP   -- InterCity Polish-IP leased lines
148.81.0.0    - WAWPOLIP  -- Warsaw Area Polish-IP
149.156.0.0   - KRAKPOLIP -- Krakow Area Polish-IP
150.254.0.0   - ?         -- Poznan Area
```

None of thees networks are registered in registrar@nic.ddn.mil, yet. NSFNET management promised us connected status, but this is not yet done. Now we have only mail access to the USA, danpost.uni-c.dk was defined to be a MX for most of our hosts.

Warszawa – Krakow line was booted in June 1991, Warszawa – Kobenhavn line started in September 1991. All our lines use PD software on IBM-PC clones for IP routing. For some key nodes routers were ordered, and now we are waiting for Us Department of Commerce export licenses to come.

Our IP backbone uses NASK 9600 leased lines, and now it consists of:

```
Kobenhavn -- Warszawa,  -- just turning into a 64k Stokholm--Warszawa line
```

```
Krakow -- Warszawa
Poznan -- Warszawa
Katowice -- Warszawa,    -- single node; they still don't have
                            registered net-number.
Torun -- Warszawa
```

There is a pending project of first MAN in Poland. In near future (3-5 months) two locations at Warsaw – namely Computer Center of Warsaw University and Physics Dep. of Warsaw University – are going to be connected by means of an optical cable. The distance is ca. 4km long and we expect to use Fiber Optics extension to the Ethernet there, in the beginning.

DECNET

It was a HEPNET initiative to connect Poland to CERN. A 9600baud line was established last year (1990) and was running DECnet since then. Last month, September 1991, it was equipped with a CISCO router so both DECnet and IP protocols are there now.

During the following weeks this HEPNET DECnet will be extended to Physics Department of Warsaw University in Warsaw by means of 9600baud leased line. This location still leak a router, so the line will be utilized as CERN-Krakow line previously was.

NASK – National Academic Computer Network (in Polish)

Established as a technical/management/support body for initial Polish EARN long distance leased lines, the group aims at providing a WAN backbone network that could be utilized by more protocols then just the EARN RSCS. Initially financed directly by ministry of education, it's slowly turning into something like foundation for Polish academic backbone network.

Now, the backbone network is constructed by means of dedicated lines, leased from PTT. These lines are capable of operating at 9600 baud. The line sharing among protocols is done by Time Division Multiplexers (TDM) hooked into a fairly simple (and cheap) synchronous modems. Actual TDM used have four ports, can compress data and make statistical use of the bandwidth resulting in effectively doubling the available line capacity.

It is not clear yet how these TDM ports will be utilized. Most probable scheme includes private X.25, RSCS, Internet and DECnet. Currently only RSCS is widespread;

The NASK currently consists of the following lines:

```
9600baud together with TDM --------------------------------------------
          line               |  protocols other then RSCS
-----------------------------------------------------------------------
       Warszawa -- Kobenhavn     X.25, Internet
```

```
     Warszawa -- Torun
     Warszawa -- Bialystok
     Warszawa -- Lublin
     Warszawa -- Krakow         Internet
     Warszawa -- Katowice       X.25, Internet
     Warszawa -- Wroclaw        X.25
     Warszawa -- Ludz
     Warszawa -- Poznan         X.25, Internet

2400baud no TDM ----------------------------------------------------------
       line                 |  only RSCS there
--------------------------------------------------------------------------
     Poznan   -- Szczecin
     Poznan   -- Wroclaw
     Wroclaw  -- Katowice
     Katowice -- Krakow
     Torun    -- Bydgoszcz
```

NASK is going to move Warszawa–Kobenhavn line onto a satellite Warszawa– –Stockholm connection which offers 64kbaud and a better entry point to the Internet for the very same price as the 9600 baud line. The operation is going to be done next week.

Finances

All these activities are fully financed by the Polish academic community. Our international lines are paid by us, too. As Poland is entering European Networking Community we are willing to learn how these expenses can be shared by both connected parties. We would like to provide commercial IP in Poland also. It's going to be based on NASK long haul trunks but it's still unclear how to share their budget.

Persons

```
Tomasz Hofmokl <fdl50@plearn.bitnet> -- EARN Polish Director
Daniel J.Bem <bem@plwrtu11.bitnet>   -- KASK, X.25 - Wroclaw area
Krzysztof Heller <uiheller@plkrcy11.bitnet> -- Krakow area IP, will apply
                                        for '.pl' domain authority
Rafal Pietrak <rafal@fuw.edu.pl>     -- Warszaw area IP
```

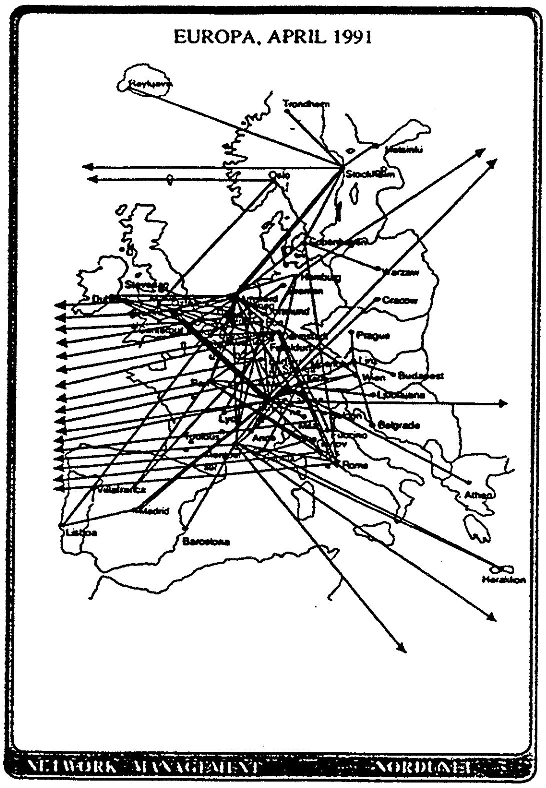

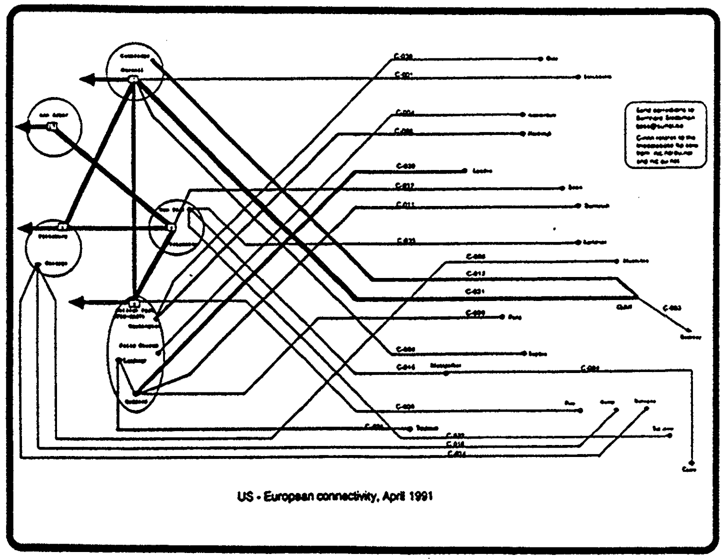

CSECKOSLOVAKIA

9.6 Kbps EARN lines inside the country

2.4 Kbps EUnet dial-up lines

EARN link Prague - Linz

EUnet link Bratislava - Vienna

International public X.25 available 1992

Plans for a multi-protocol backbone with at least 64 Kbps connectivity to Internet.

HUNGARY

National experimental priovate X.25 network

One EARN node

One EUnet node

Publix X.25 during 1991.

Planned TCPIP connection to CERN

X.400 mail service

POLAND

Experimental X.25 network connecting 8 sites

Planned computer network (KASK) connecting 16 sites.

International connections to CERN and Copenhagen.

Plans to connect to DFN.

One EARN node (PLEARN)

Planned IBM EASI line to Vienna

POLAND NETWORKING

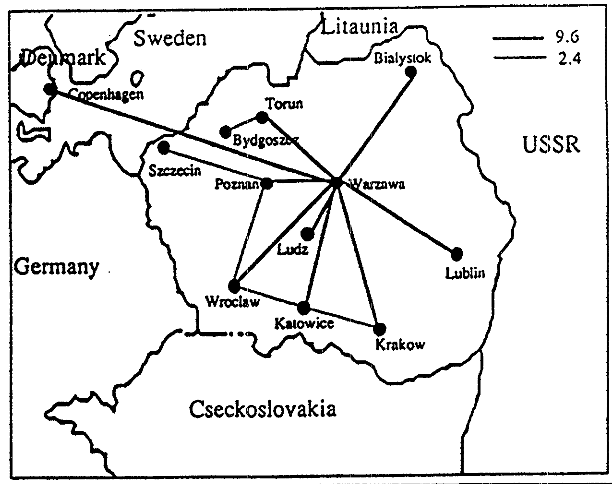

ROMANIA

A few 4800 connections inside the country

No international links

PTT program for X.25 recently started

Sattelite links a possibility

USSR

No nation-wide strategy

Public X.25 (IASNET) available

Sattelite links more reliable

Link Moscow (DEMOS) - Helsinki

DEMOS provide dial-up access

Planned link USSR - Poland

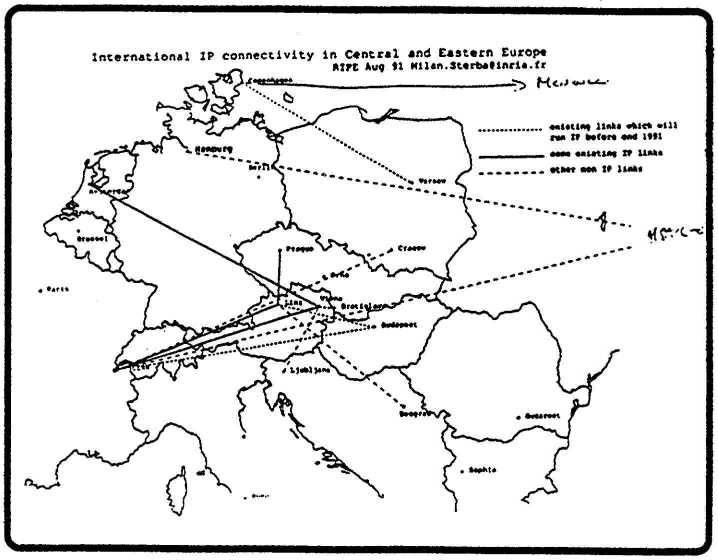

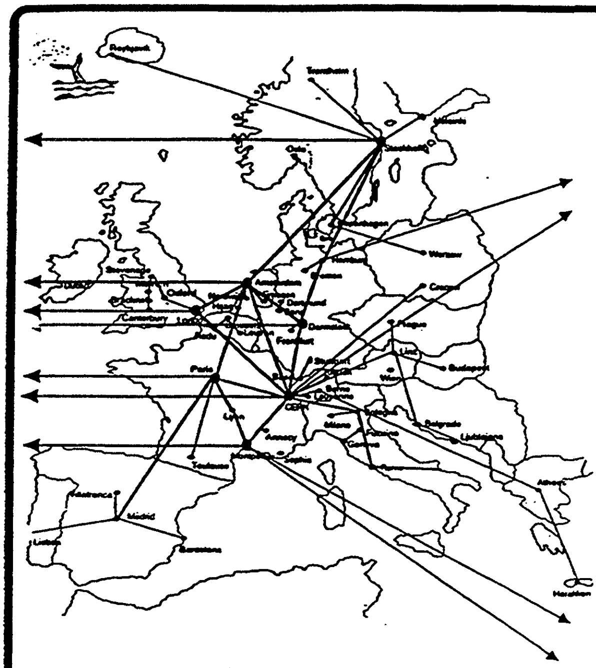

Optimized European International Connectivity.

NOTE. National lines are just put in as examples and have nothing to do with actual national infrastructures. The map is only intended as an example of an optimized pan-European backbone.

Bernhard Stockman, NORDUnet/SUNET.

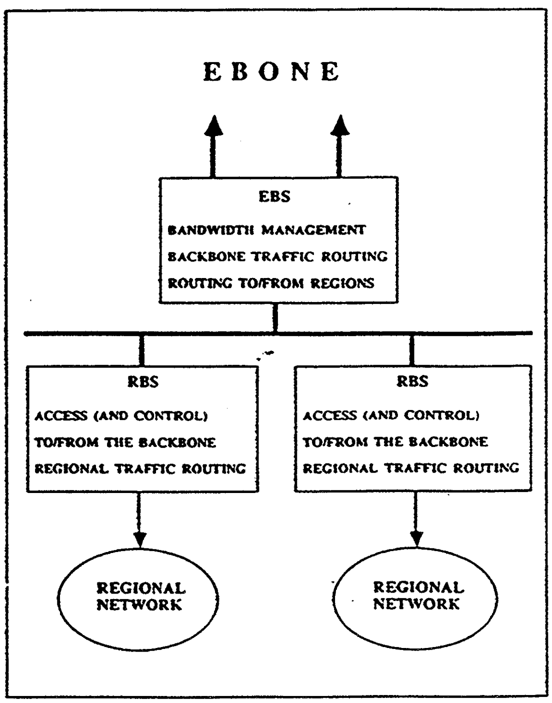

FIGURE 1

4.3 ESnet Report

Presented by Tony Hain/LLNL

ESnet Status Report

IETF - Santa Fe

November 1991

Anthony L. Hain

Associate Network Manager

ESnet / NERSC

PAST ACTIVITIES

Participation in Inter-Op OSI Demonstration

Frame Relay Trial between LLNL/PPPL/GA

Connected DOE Ops Offices SAN/CHI

STATS

35 Routers Managed

96 Directly Connected Networks
887 Regional Connected Networks
873 Networks via Other Backbones

1.5G Packets Received

88% IP / 12 % DECnet

PLANNED ACTIVITIES

Peer with JVNC

IP on existing DECnet line to INFN

New lines to KEK & JAERI-NAKA

T3 line - LLNL to LBL

RFP for Cell based service T3 -> OC12

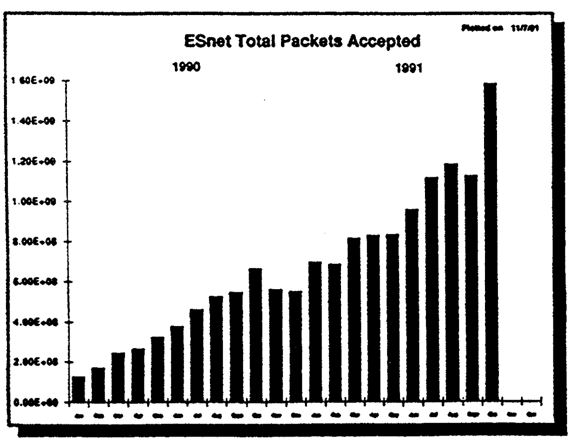

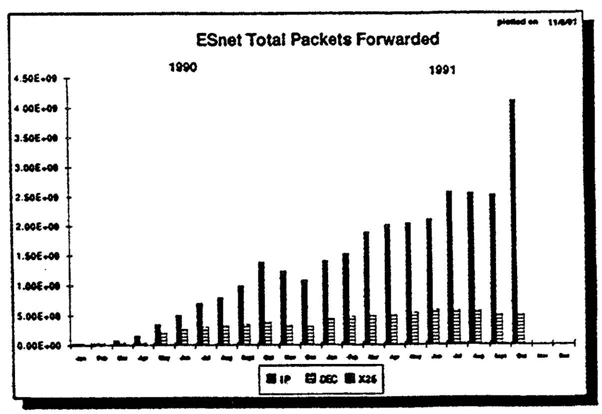

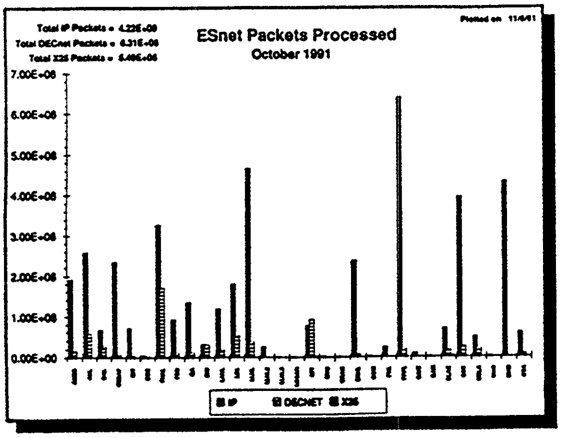
ESnet Packets Processed
October 1991
7.00E+06
6.00E+06
5.00E+06
4.00E+06
3.00E+06
2.00E+06
1.00E+06
0.00E+00
IP
DECNET
X25

4.4 NSI Report

Presented by Milo Medin/NASA

NSI Status Update

Internet Engineering Task Force Meeting
November 20, 1991
Santa Fe, New Mexico

Milo Medin
NASA Science Internet Office

page 1

NSI Status Report

NSI Multiprotocol Network

- 65 Proteon routers
- 41 Support DECNET IV
- 62 Support IP (for user data routing)
- 6 Support ISO 8473 with DIS 10589 IS-IS routing
- No Appletalk, XNS, IPX or other support

page 2

NSI Status Report

NSI-DECnet (formally known as SPAN) modernization

- MultiNet DECNet in IP encapsulation (ncap24, ncap5, ncap6, ncap56)
- Replacement of all DEC equipment with multiprotocol routers
- Deletion and consolidation of 9.6 Kbps links
- Upgrade of remaining 9.6Kbps links to 56Kpbs

page 3

NSI Status Report

Network Operations Support

- 24x7 Network operations coverage
- 1-800-424-9920 Trouble reporting hotline
- Proteon Overview used for Multiprotocol network management
- Homegrown Sun DECNet IV monitoring support
- Transition to DEC MCC/MSU product for integrated management
- Out of band access to router console ports

page 4

NSI Status Report

Recent Events

- Multi-T1 Backbone upgrades in progress
- "Hardening" backbone connectivity
- Installation of T1 Internet access to GSI NIC from FIX-E
- Extension to NASA 70 Meter Deep Space Net sites
- UARS mission launch and virtual net support
- Joint NSF-NASA link to US Antarctic Base - McMurdo Sound (12/20)

page 5

NSI Status Report

Current International Access

- 56 Kbps VSAT to CTIO in La Serena, Chile
- 128 Kbps terrestrial to UK (ULCC, RAL, Oxford Univ)
- 56 Kbps link to ESOC
- 56 Kbps link to CNES
- 256 Kbps to Australia (PACCOM sat.)
- 64 Kbps to New Zealand (PACCOM sat.)
- 64 Kbps to Hong Kong (PACCOM sat.)
- 768 Kbps to Japan, Korea via Hawaii (PACCOM terrestrial)

page 6

NSI Status Report

Future International connectivity

- 56 Kbps to SAR in Greenland
- 512 Kbps access to UK (re-engineered FatPipe)
- 256 Kbps to ESA/ESOC
- 256 Kbps to ESA/ESTEC

page 7

NSI Status Report

NSI OSI Activities

- GSA allocation of NASA NSI AAI (005900)
- Allocation of routing domains to NASA Centers
- Deployment of DIS 10589 IS-IS between ARC, JPL, GSFC
- Support of NASA X.500 White Pages service
- Future deployment of 2 DEC 5000 DNS/DTS servers
- Routing interoperation with NSFNet, and ESNet

page 8

NSI Status Report

FIX -West support

- Generator backup (w/500 gal. fuel tank) online
- December deployment of FIX FDDI concentrator
- Facility upgrade during post-Christmas holiday
- Improved statistics collection

page 9

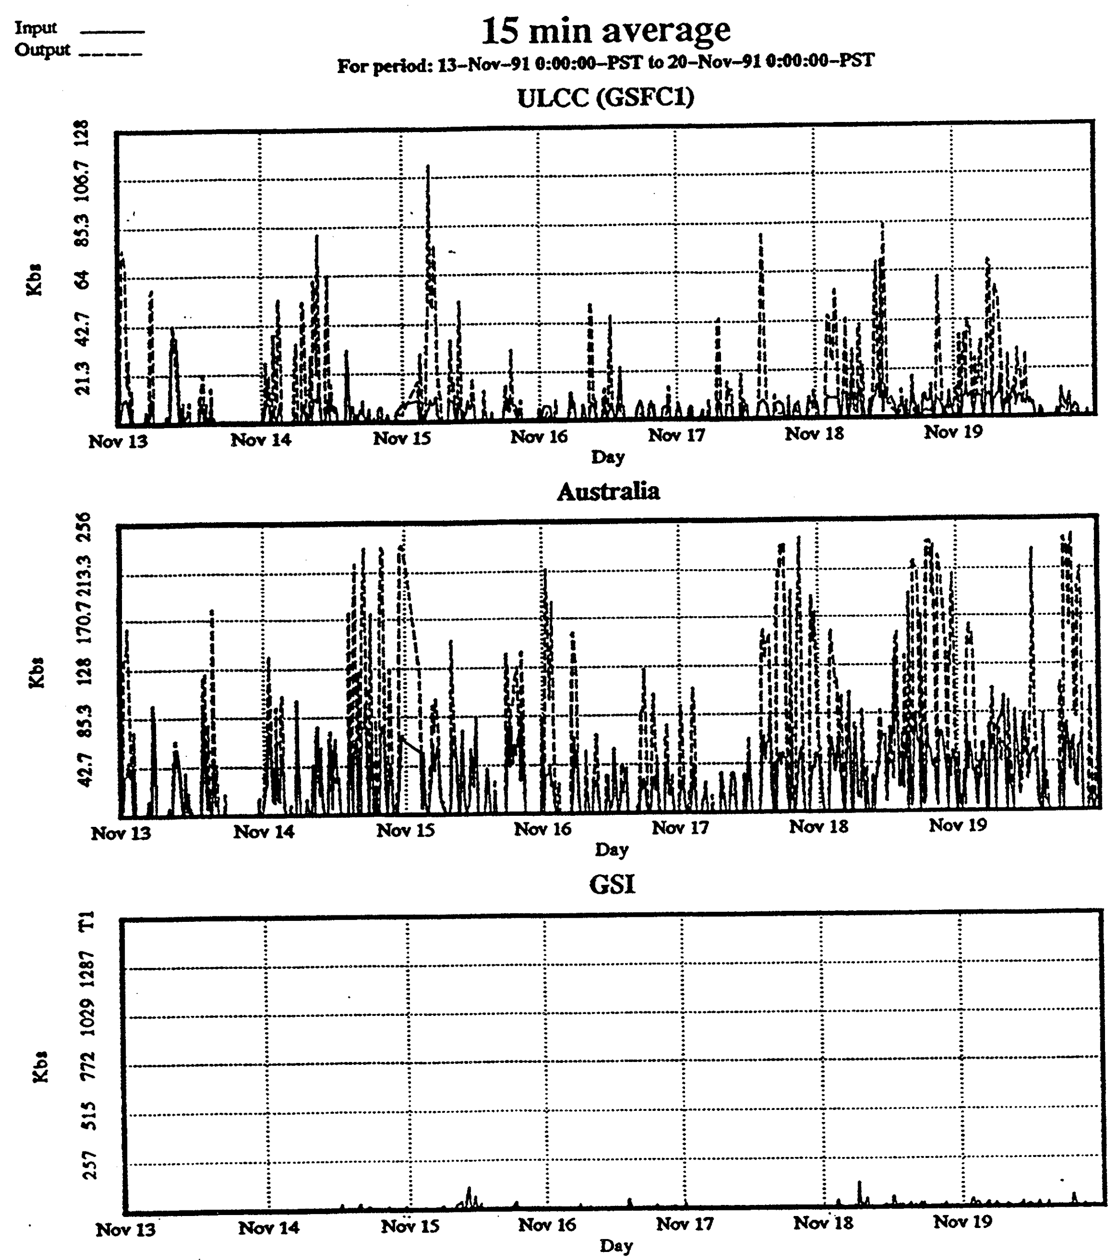
Input
Output
15 min average
For period: 13-Nov-91 0:00:00-PST to 20-Nov-91 0:00:00-PST
ULCC (GSFC1)
Kbs
21.3
42.7
64
85.3
106.7
128
Nov 13
Nov 14
Nov 15
Nov 16
Nov 17
Nov 18
Nov 19
Day
Australia
Kbs
42.7
85.3
128
170.7
213.3
256
Nov 13
Nov 14
Nov 15
Nov 16
Nov 17
Nov 18
Nov 19
Day
GSI
Kbs
257
515
772
1029
1287
T1
Nov 13
Nov 14
Nov 15
Nov 16
Nov 17
Nov 18
Nov 19
Day

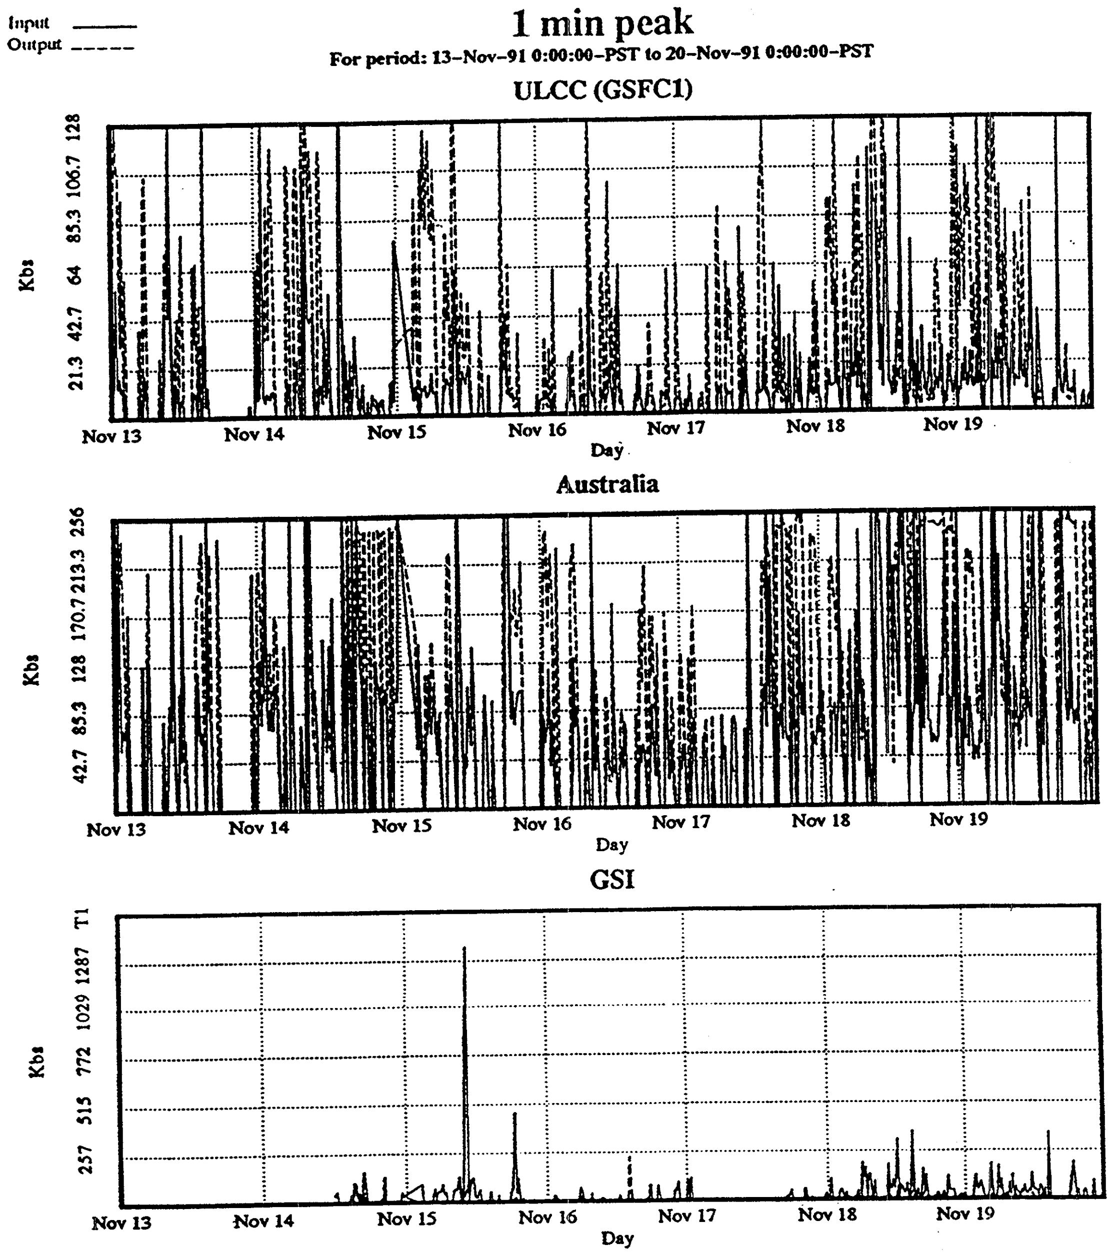

Input
Output
1 min peak
For period: 13-Nov-91 0:00:00-PST to 20-Nov-91 0:00:00-PST
ULCC (GSFC1)
Kbs
21.3 42.7 64 85.3 106.7 128
Nov 13 Nov 14 Nov 15 Nov 16 Nov 17 Nov 18 Nov 19
Day
Australia
Kbs
42.7 85.3 128 170.7 213.3 256
Nov 13 Nov 14 Nov 15 Nov 16 Nov 17 Nov 18 Nov 19
Day
GSI
Kbs
257 515 772 1029 1287 T1
Nov 13 Nov 14 Nov 15 Nov 16 Nov 17 Nov 18 Nov 19
Day

Chapter 5

IETF Protocol Presentations

5.1 Inter-Domain Policy Routing

Presented by Martha Steenstrup/BBN

Inter-Domain Policy Routing

M. Steenstrup
BBN Communications

and members of the
IDPR Working Group

Policy Routing

- Motivation
- Our Approach
- How to Reap the Benefits
- Implementation Status

The Internet

- Administrative domain: collection of contiguous hosts, networks, and gateways under a single administration
- Large number of domains
- Arbitrary interconnectivity among domains
- Diverse service offerings and restrictions among domains
- Heterogeneous routing and addressing schemes among domains

Goal of Policy Routing

To provide routes that:

- Satisfy the service requirements stipulated by the sources
- Respect the service constraints imposed by the domains transited

Policies

- Service:
  access restrictions
  quality
  monetary cost
- Source Policies:
  service requirements
  private
- Transit Policies:
  service constraints
  public

IDPR Approach

- Link-state distribution of routing information:
  domain transit policies and interconnectivity
- Source-specified routes:
  path setup at the domain level
- Encapsulation:
  within a domain, use local routing and addressing schemes
- Security:
  integrity checks, authentication, and timestamps

Routing Entities

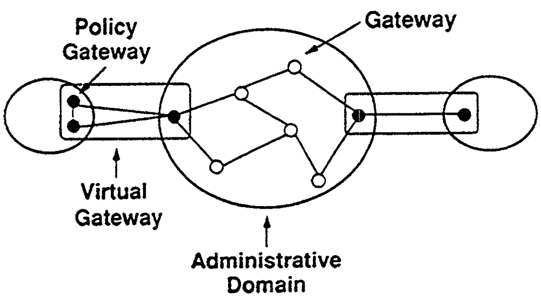

Size Accommodation

- Routing at the level of administrative domains
- Virtual gateways as domain interconnections
- Special purpose servers separate from policy gateways for configuration and route generation
- Distributed databases with partial information
- Super domains consisting of several contiguous domains with similar transit policies

More Size Accommodation

- Distribution of link-state information only after a change occurs
- Retention of usable link-state information only
- Route precomputation and caching
- Multiplexing many host flows onto a single path
- Network management tools for policy tracking

Using IDPR

- Deciding when to use IDPR
- Interactions with other inter-domain routing procedures
- Proxy domains
- Interactions with intra-domain routing procedures
- Configuration and management

When to Use IDPR

- Host applications with strict service requirements
- Cost containment for users
- Transit domains with varied services
- Transit domains with access restrictions

Inter-Domain Routing Interactions

- Multiple inter-domain routing procedures can coexist
- IDPR is most effective in contiguous IDPR domains
- Can construct partial IDPR routes, relying on other inter-domain routing procedures to complete the route
- Host traffic flows are selectively configured for IDPR routing at the source domain

Proxy Domains

- Stub (non-transit) domains need not support IDPR
- Proxy domains generate policy routes on behalf of hosts in non-IDPR domains
- Stub domains must provide proxy domains with source policy configuration
- Stub domains must be able to reach proxy domains via some inter-domain routes

Intra-Domain Routing Interactions

- Provides IDPR entity reachability and quality of service information within a domain
- To cross a transit domain, IDPR traffic should always use intra-domain routing information
- To reach first policy gateway for external destinations, IDPR traffic should use alternate inter-domain routing information or configured routing information

Configuration

- Adjacent domain connectivity
- Entity identifier/address maps
- Source policies
- Transit policies
- Protocol parameters
- Domain identifiers in DNS
- Internet coordination of policy

Status

- Architecture, protocol specifications, configuration and usage guide, and MIB available as Internet Drafts
- Working prototype for Suns completed in early 1991 (USC, SAIC, and BBN)
- Experiments performed with USC laboratory network, Sparta, Mitre, and DCA networks, and DARTNET and the TWBnet
- Gated version in progress (R. Woodburn, C. Chu, and H. Bowns)

Benefits of IDPR

Control over:

- Paths traversed by your hosts' traffic
- The use of your networks' resources

Hence:

- Service consistent with your applications' requirements
- Traffic consistent with your resource constraints

More Benefits of IDPR

Friendly:

- Can be added to existing Internet
- Interoperates with existing protocols
- No host changes required
- Only policy gateways require IDPR software

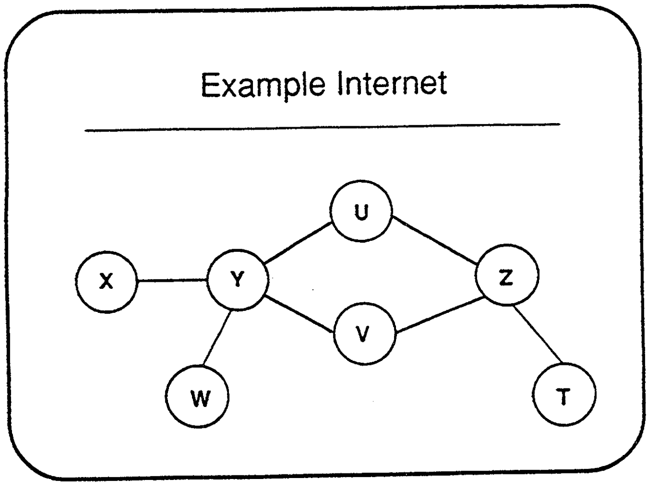
Example Internet
U
X
Y
Z
V
W
T

5.2 Router Discovery

Presented by Steve Deering/Xerox

Router Discovery

Steve Deering

deering@parc.xerox.com

Router Discovery — Current

- manual configuration (incl. BOOTP)
  - — administrative burden
  - — can't track dynamic availability
- eavesdropping on routing traffic (RIP)
  - — doesn't work with new routing routing protocols
- proxy ARP
  - — often, slowest router wins
  - — can get 3 pkts for every one sent
- "logical address", e.g. <net,subnet,1>

Candidate Protocols

- stub RIP — non-RIP routers could still send default RIP route
- cisco GDP — periodic UDP broadcasts by routers
- Deering ICMP extensions — periodic ICMP multicasts by routers
- Prindeville protocol #1 — initially, send IP unicasts as LAN multicasts
- Prindeville protocol #2 —multicast ICMP queries for dest. & TOS
- ES-IS subset, with IP addresses
- proxy ARP
- <net, subnet,1>
- BOOTP, BootParam, NIT

ICMP Router Advertisement

Type = 9	Code = 0	Checksum
Num Addrs	Ent Size = 2	Lifetime
Router Address		
Preference Level		
Router Address		
Preference Level		
⋮		

IP destination = 224.0.0.1 or
255.255.255.255

default transmit interval = 7—10 minutes

default lifetime = 30 minutes

default preference level = 0

ICMP Router Solicitation

Type = 10	Code = 0	Checksum
Reserved = 0		

IP destination = 224.0.0.2 or
255.255.255.255

max # transmissions = 3, at startup only

retransmission interval = 3 seconds

(routers that fail to respond or that appear later are discovered from their advertisements)

Note:

- discovers **default** routers only; depend on Redirects for particular dest & TOS
- can configure non-advertisement of some addresses
- **not** intended for black hole detection
  - — hosts should already have other mechanisms, for non-broadcast nets and nets without rtr discovery
  - — means advertising rate can be very low => negligible overhead
  - — can configure higher rate, if desired

Security Considerations

Proposed protocol allows any neighbor to masquerade as a default router

- — can eavesdrop on off-subnet traffic
- — can deny forwarding service
- — can modify forwarded traffic, by insertion/deletion/modification

Note that these threats already exist on subnets that use ARP

Can configure hosts to ignore router advertisements, if desired

Protocol is extensible, would allow addition of authentication fields in future

Current Status

- published as Proposed Standard in September, 1991 — RFC 1256
- Sun implementation for BSD-based systems (user-level, host or router) to be available by anonymous FTP
- Host implementations planned by Apple (Mac) and FTP Software (PC).
- Anyone else?

Chapter 6

Technical Presentations

6.1 FARNET Overview

Presented by Laura Breeden/FARNET

FARNET
Federation of American Research Networks
100 Fifth Avenue
Waltham, MA 02154

1-800-72-FARNET
or ftp *nic.cerf.net*, cd *farnet*

Why a national association of networks?

- To exchange information and ideas
- To advocate for networking at the national level
- To provide a single point of contact for information about midlevel networks
- To support and develop the members of FARNET

Mission and strategy

FARNET's *mission* is to advance research and education through the use of data networking.

FARNET has adopted four primary *strategic focuses:*

1. Improvement of customer and network information services for Internet users
2. Advocacy
3. Programmatic ("user group") activities
4. Negotiation of group discounts for members

Membership profile

National service providers	3
Regional (multi-state) networks	7
Supercomputer center networks	3
State and provincial networks	15
University campuses	1
Local service providers	1

Serving more than 1,000 organizations, in every state and several foreign countries

History

- Founded in 1986 by leaders of NSF midlevel networks
- Quarterly meetings since 1987
- Began collecting dues in 1988
- Incorporated as non-profit in 1990
- Hired first staff member in 1991

So, what does FARNET *do??*

- FARNET meeting on K-12 networking activities (May 1991)
- FARNET workshop and recommendations to NSF on backbone structure post-1992
- FARNET meeting on "hardening the regionals" (Nov. 1991)
- Liaison activities with IETF, SIGUCCS, CCIRN, IEPG, FNC, NTTF, etc.
- FARNET *Gazette* and position papers

How FARNET works

Members elect Board of Directors

Board sets policy

Executive Director responsible for day-to-day operations

How it works (continued)

Seven committees provide guidance, organize projects:

Technical	Gene Hastings
Program	Paul Love
User Services	Martyne Hallgren
K-12	Jim Luckett
External Affairs	Richard Mandelbaum
Memb. & Bylaws	Jim Williams
Nominating	Glenn Ricart

YOU can join a committee or be on the mailing list.

FARNET and the IETF

√ Lots of overlap already!

√ Network operators need to work *smarter*, not *harder*!

√ Protocols, standards, tools, and procedures should help us do this!

And finally...

√ Good coordination between FARNET and IETF is essential!

FARNET networks are in production environment, with real-world economic and operational constraints

Plenty of notice is needed for major changes, such as OSPF and BGP

New standards and procedures have to be workable and affordable

For more information

Ftp to nic.cerf.net

Cd *farnet*

Cd *farnet_info* or *farnet_docs*

Members, papers, committees, agendas, gazettes...

MEMBERSHIP LIST

Advanced Networks and Services (ANS)
BARRnet (Bay Area Regional Research Network)
CERFnet (California Educational and Research Foundation Network)
CICnet (Committee on Institutional Cooperation Network)
Colorado SuperNet
CONCERT (Communication for North Carolina Education, Research and Technology)
Cornell University
CREN (Corporation for Research and Educational Networking)
CSUnet (California State Universities Network)
JvNCNet
MichNet/MERIT
Midnet
MRnet (Minnesota Research Network)
National Center for Atmospheric Research (NCAR)*
NEARnet (New England Academic and Research Network)
NETillinois
NevadaNet
NorthWestNet
NYSERNet (New York State Educational and Research Network)
OARnet (Ohio Academic and Research Network)
Onet (Ontario Network)*
PREPnet (Pennsylvania Research and Economic Partnership Network)
PSCnet (Pittsburgh Supercomputer Center Network)
PSInet (Performance Systems International Network)
Sesquinet
SDSCNet (San Diego Supercomputer Center Network)
SURAnet (Southeast Universities Research Association Network)
THEnet (Texas Higher Education Network)
VERnet (Virginia Education and Research Network)
Westnet

* Associate Member

6.2 Overview of ATM

Presented by George Clapp/Ameritech

What is Broadband ISDN?

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

George H. Clapp

AMERITECH SERVICES

Broadband ISDN Interface: *Asynchronous Transfer Mode (ATM)*

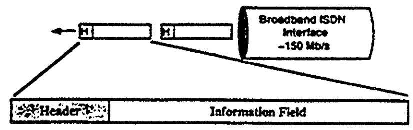

Characteristics:

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

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

George H. Clapp

AMERITECH SERVICES

Why ATM?

Flexibility for the End User:

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

Flexibility for the Network Operator:

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

George H. Clapp

AMERITECH SERVICES

B-ISDN User-Network-Interface (UNI) ATM Cell Format

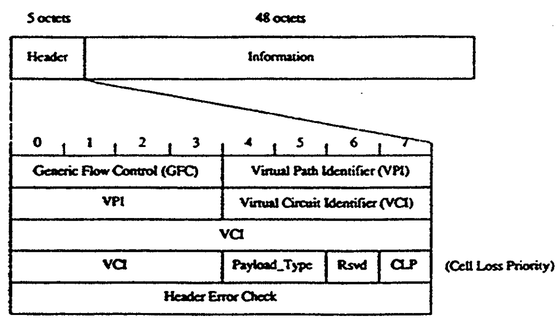

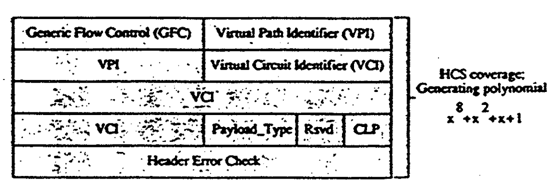

George H. Clapp

AMERITECH SERVICES

T1S1.5 PAYLOAD TYPE

RESERVED BIT ADDED TO FIELD.

000		USER NORMAL
001		USER SPECIFIC
010		USER NORMAL, CONGESTION EXPERIENCED
011		USER SPECIAL, "
100	}	RESERVED
101		
110	}	NETWORK OA+M
111		

B-ISDN Network-Network-Interface (NNI) ATM Cell Format

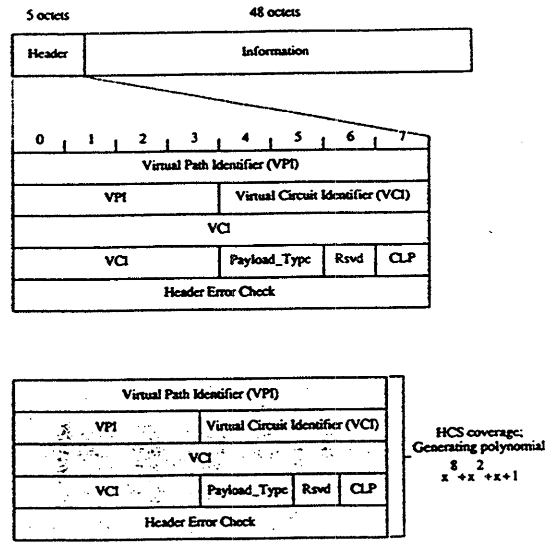

George H. Clapp

AMERITECH SERVICES

Virtual Path Concept

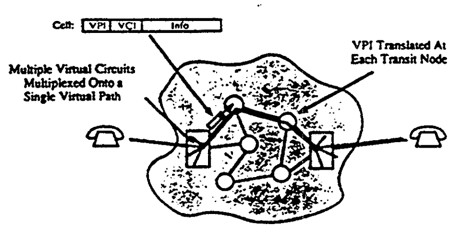

Advantages:

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

George H. Clapp

AMERITECH SERVICES

Header Check Sequence

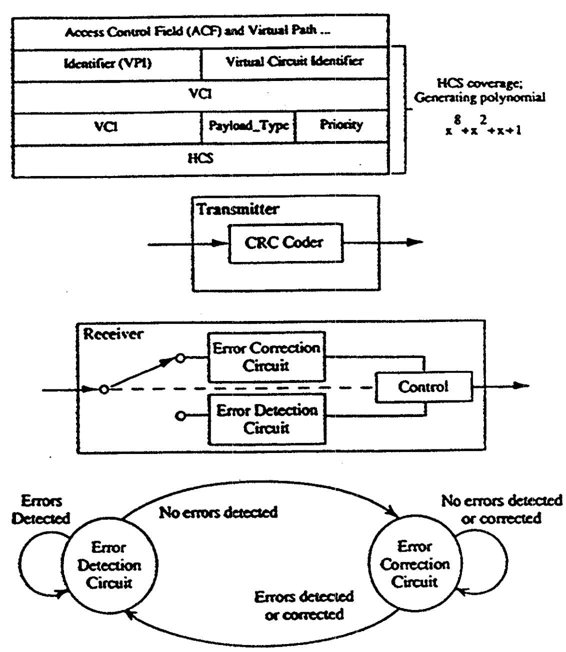

Transitions occur upon reading a segment.

George H. Clapp

AMERITECH SERVICES

ATM Protocol Architecture

	Variable Bit Rate Services				
ATM Adaptation Layer	Signaling	CO	CL	Other VBR Services	CBR Services
	Convergence Sublayer (CS)			CS	CS
	Segmentation and Reassembly (SAR)			SAR	SAR
	ATM				
	Physical Layer (STS-3c)				

ATM ADAPTATION LAYERS
SERVICE CLASSIFICATION

	CLASS A	CLASS B	CLASS C	CLASS D
TIMING RELATION BETWEEN SOURCE AND DESTINATION	REQUIRED		NOT REQUIRED	
BIT RATE	CON-STANT	VARIABLE		
CONNECTION MODE	CONNECTION ORIENTED			CONNEC-TION-LESS

CLASS A : CIRCUIT EMULATION, CBR VIDEO
CLASS B : VARIABLE BIT RATE VIDEO/AUDIO
CLASS C : CO DATA
CLASS D : CL DATA

ATM ADAPTATION LAYERS (AALs)

SUBLAYERS

- CONVERGENCE SUBLAYER (CS)
- SEGMENTATION AND REASSEMBLY SUBLAYER (SAR)

AAL 1 : CBR (CIRCUIT EMULATION) CLASS A; UNDER STUDY.

AAL 2 : VBR, TIMING INFO, ERROR CONTROL; CLASS B; UNDER STUDY.

AAL 3 : VBR, CONNECTION-ORIENTED, CLASS C. SIGNIFICANT PROGRESS

AAL 4 : VBR, CONNECTIONLESS, CLASS D; MAINTAIN COMPATIBILITY WITH IEEE 802.6 STD.

802.6 and BISDN Protocol Layers

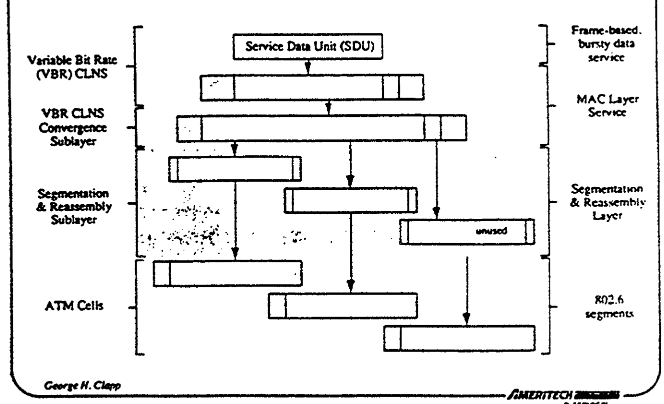

AAL 3+4 COMMON PART CS (CPCS)
SERVICE SPECIFIC CS (SSCS)

Adaptation Layer
VBR CLNS Convergence Protocol Sublayer, aka, Common PDU Header & Trailer
(CPCS-PDU FORMAT)

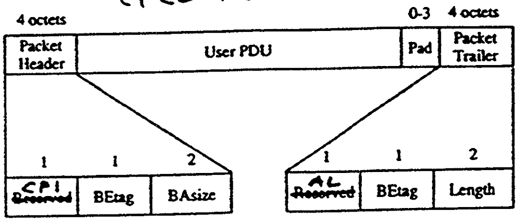

CPI COMMON PART INDICATOR, USED TO INTERPRET FOLLOWING FIELDS, 00 = CL.

BEtag Same value is placed in the header and trailer fields; used to associate header and trailer of the same PDU for error control. Values 0-255 are cycled through.

BAsize Used by receiver for buffer management; either...
Length in octets of ~~802.6 IMPDU~~ PAYLOAD (header and information, inclusive), or...
Greater than or equal to the true PDU length.

Length Length in octets of the ~~user PDU~~ CPCS-PDU PAYLOAD (less the Pad).

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

AL ALIGNMENT (TO 4 OCTETS)

George H. Clapp

AMERITECH SERVICES

IEEE 802.6 Adaptation Layer
Segmentation and Reassembly (SAR) Sublayer

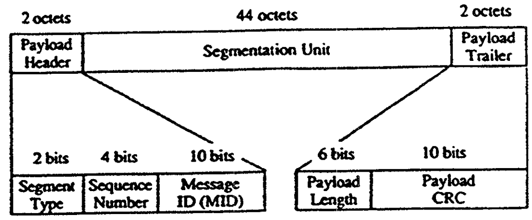

Segment Type Used for packet delineation;
Encoding
Beginning of Message (BOM): 10
Continuation of Message (COM): 00
End of Message (EOM): 01
Single Segment Message (SSM): 11

Sequence Number Values 0-15 are cycled through to number consecutive segments of a packet; value is incremented relative to the previous value for a particular VCI/MID value.

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

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

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

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

George H. Clapp

AMERITECH SERVICES

CCITT BISDN Adaptation Layer
Segmentation and Reassembly (SAR) Sublayer (the same)

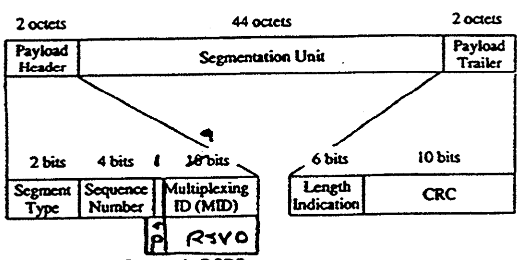

Segment Type Same as in DQDB.

Sequence Number Values 0-15 are cycled through to sequentially number consecutive segments of a packet; used for detection of out-of-sequence segments.

P PRIORITY (FOR AAL 5)

Multiplexing ID Shortened from 14 to 10 bits. Used as in DQDB to associate CLNS segments of a single packet. (FOR AAL 4)

RSVO FOR AAL 3 (IF NO MID MUXING DONE).

Length Indication Same as in DQDB.

CRC Same as in DQDB.
Generating polynomial: $G(x) = x^{10} + x^{9} + x^{5} + x^{4} + x + 1$

George H. Clapp

AMERITECH SERVICES

CRC per Cell
Sources of Error

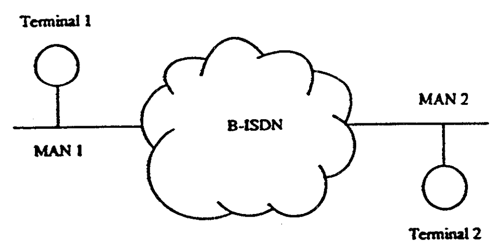

3 sources of error from terminal 1 to terminal 2 in a fiber optic network:†

- Random Errors: 10^{-12}.
- Burst errors (protection switching): 0.24 events per day on 1000 mile system; 20-40 ms duration.
- Buffer overflow: engineering parameter.

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

George H. Clapp

AMERITECH SERVICES

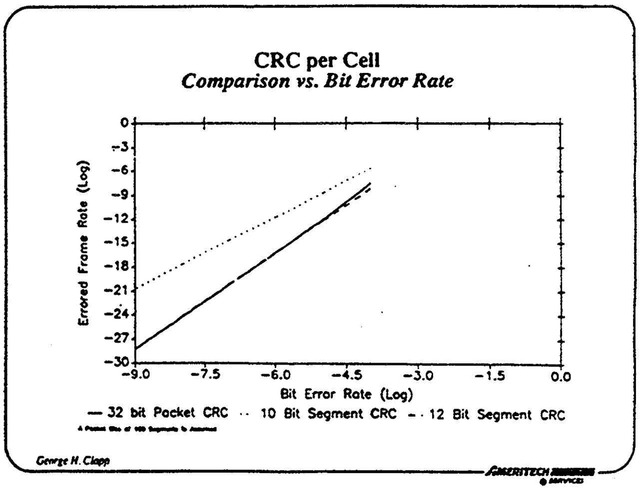
CRC per Cell
Comparison vs. Bit Error Rate
Errored Frame Rate (Log)
Bit Error Rate (Log)
— 32 bit Packet CRC ·· 10 Bit Segment CRC –· 12 Bit Segment CRC
George H. Clapp
AMERITECH SERVICES

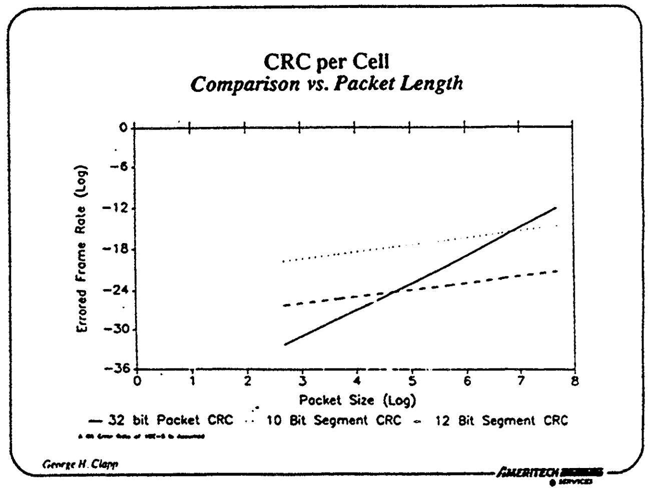
CRC per Cell
Comparison vs. Packet Length
Errored Frame Rate (Log)
Packet Size (Log)
— 32 bit Packet CRC ·· 10 Bit Segment CRC – 12 Bit Segment CRC
George H. Clapp
AMERITECH SERVICES

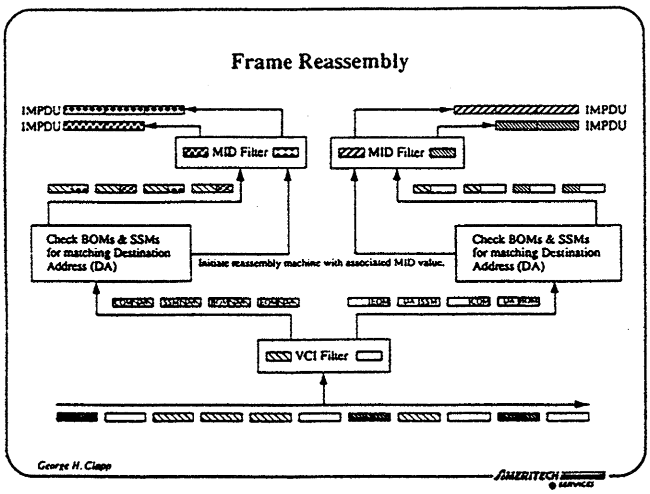
Frame Reassembly
IMPDU
IMPDU
MID Filter
MID Filter
IMPDU
IMPDU
Check BOMs & SSMs for matching Destination Address (DA)
Initiate reassembly machine with associated MID value.
Check BOMs & SSMs for matching Destination Address (DA)
VCI Filter
George H. Clapp
AMERITECH SERVICES

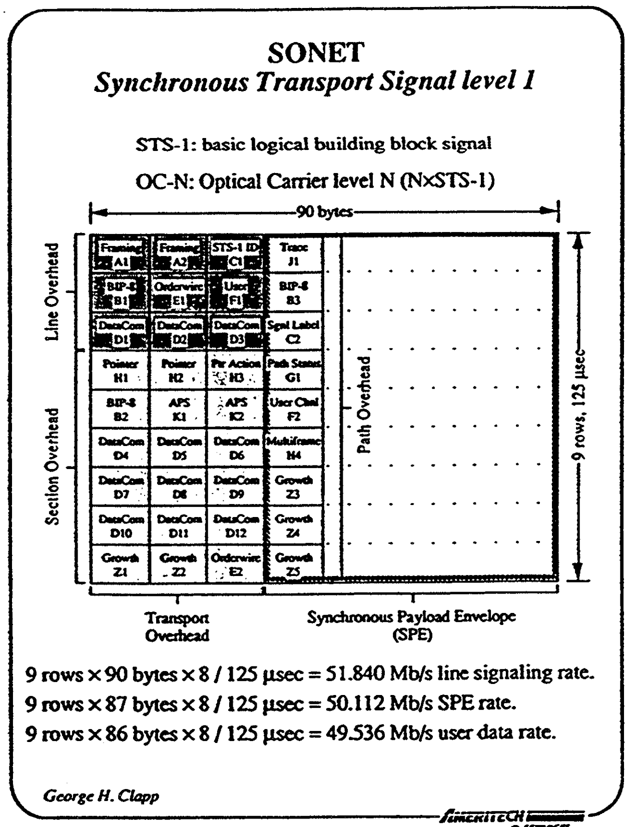
SONET
Synchronous Transport Signal level 1
STS-1: basic logical building block signal
OC-N: Optical Carrier level N (N×STS-1)
90 bytes
Line Overhead
Section Overhead
Path Overhead
9 rows, 125 μsec
Framing A1
Framing A2
STS-1 ID C1
Trace J1
BIP-8 B1
Orderwire E1
User F1
BIP-8 B3
DataCom D1
DataCom D2
DataCom D3
Sgnl Label C2
Pointer H1
Pointer H2
Ptr Action H3
Path Status G1
BIP-8 B2
APS K1
APS K2
User Chnl F2
DataCom D4
DataCom D5
DataCom D6
Multiframe H4
DataCom D7
DataCom D8
DataCom D9
Growth Z3
DataCom D10
DataCom D11
DataCom D12
Growth Z4
Growth Z1
Growth Z2
Orderwire E2
Growth Z5
Transport Overhead
Synchronous Payload Envelope (SPE)
9 rows × 90 bytes × 8 / 125 μsec = 51.840 Mb/s line signaling rate.
9 rows × 87 bytes × 8 / 125 μsec = 50.112 Mb/s SPE rate.
9 rows × 86 bytes × 8 / 125 μsec = 49.536 Mb/s user data rate.
George H. Clapp
AMERITECH SERVICES

SONET
Synchronous Hierarchical Rates

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

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

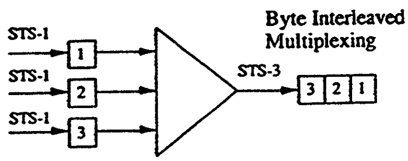

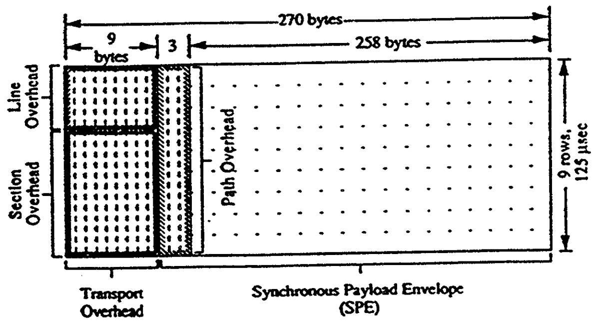

George H. Clapp

AMERITECH SERVICES

Standard CCITT SONET Rates
Synchronous Transport Module Levels

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

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

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

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

STM-1 is the CCITT basic building block.

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

†Candidate rates which are not standardized.

George H. Clapp

AMERITECH SERVICES

CCITT STM-1 Format

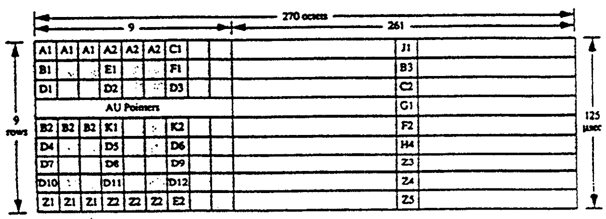

A1, A2	Framing	B3	BIP-8
AU Pointers	Administrative Unit Pointers	C2	STM Identifier
B1	BIP-8 (Bit Interleaved Parity)	F2	User-Defined Channel
B2	BIP-24	G1	Path Status
C1	STM Identifier	H4	Multiframe Indicator
D1-D12	Data Communications Channel	J1	Path Trace
E1, E2	Order Wire	Z3-Z5	Growth (reserved as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

George H. Clapp

AMERITECH SERVICES

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

Three mechanisms:

- Pointer.
  - "Multiframe Indicator" octet (H_4) of the Path Overhead points to the beginning of the first complete cell following the H_4 octet.
- Cell counting.
- Header Check Sequence (HCS) calculation.
  - Following cell boundary identification via the H_4 octet, the HCS is checked to verify cell delineation.

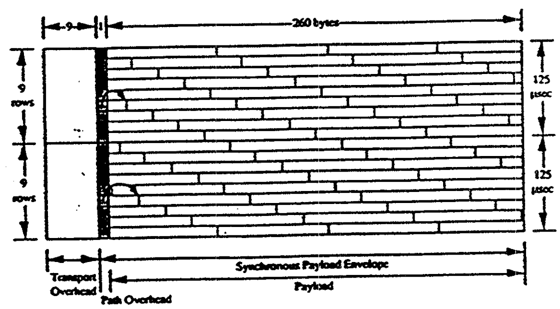

$$\frac{260 \text{ bytes} \times 8 \times 9}{125\ \mu\text{sec}} = 149.760 \text{ Mb/s} \qquad \frac{(260 \times 9) \text{ bytes / frame}}{53 \text{ bytes / cell}} = 44\tfrac{8}{53} \text{ cells / frame}$$

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

George H. Clapp

AMERITECH SERVICES

6.3 ATM in LAN's

Presented by Tom Lyon/Sun Microsystems

ATM and LANs

Tom Lyon

Sun Microsystems, Inc.

pugs@sun.com

LAN Trends

- Ethernet 10Base–T concentrators turn logical bus into physical star.
- FDDI concentrators turn logical ring into physical star.
- Bandwidth is shared, but wires are not!
- If star topologies are prevalent, why have a MAC layer? Use switching instead.

LAN Problems

- Speeds beyond FDDI will be needed, and soon!
- No hope of real–time guaranteed delivery on current LANs (e.g., video teleconference)
- Too hard to build bus–based routers/bridges of adequate bandwidth. Use switching instead.

Why ATM?

- Need for high aggregate bandwidth – it makes switching easy.
- Need for multi–service/multi–media capabilities.
- Need for *scaleable* solution – same architecture at ever higher speeds.
- Same needs for LAN as for B–ISDN, only sooner!

Why ATM?

- Leverage billions of $ of ATM/B–ISDN R&D, ensure convergence of LAN & WAN
- 40 Gb/s switches already demonstrated

The Vision

- 155 Mbps to the desktop, Nx155 to servers, star wiring only
- CPU motherboard connection @ $100 parts
- Wiring closet switch/concentrator @ $1000/port
- *Short* run from desktop to wiring closet – be FDDI compatible for now, use UTP later
- Later generations at higher speed, same cost

Challenges

- IP is fundamentally connectionless, ATM is fundamentally connection–oriented
- Per–cell processing overhead in 2.8 uSec – think big bytes, not small packets!
- No MAC layer means simultaneous transmitters can overrun receiver
- Move cost assumptions from WW–III survivable C.O. switch to commodity wire–closet concentrator

Usage Assumptions

- Traffic will be predominantly TCP/IP – but model holds for OSI CLNS, XNS, IPX, Appletalk, etc.
- ATM is just one of many data link level networks
- Both end systems and routers must work well
- Evolution of apps to request QOS/reservations – pure datagram service no longer sufficient
- All true for both B–ISDN and ATM LANs

B–ISDN vs. ATM LANs

- Don't assume B–ISDN solution until the problem is agreed upon.
- Physical Layers
- Signalling
- Congestion
- Management, Multicast, OA&M, ...

AAL3/4 vs. ATM LANs

- CRC–10 not good enough for UTP.
- Another layer of multiplexing (MID) adds much complexity to end systems.
- 48 byte payloads more efficient than 44 because fewer bus bursts needed.

SEAL/AAL5 Proposal

- Simple: 48 byte payloads, last cell padded, length and CRC–32 added.
- Requires end of packet indication in ATM header.
- CRC handles missing, corrupted, re–ordered cells.

6.4 A Unified Approach to Inter-Domain Routing

Presented by Deborah Estrin/USC

A Unified Approach to Inter-Domain Routing

Deborah Estrin

Computer Science Department
University of Southern California
estrin@usc.edu

Yakov Rekhter
IBM
yakov@ibm.com

Opinions expressed are mine ... not necessarily the same as those of my co-author...

Motivation

Internet growth necessitates **Policy Routing:**

Transit restrictions imposed by source, destination and transits.

Multiple types of service (TOS).

Multiple carriers, charging schemes.

Internet growth also necessitates algorithms with good complexity characteristics (avoid n^2 growth).

Inter-Domain Routing

IDR must advertise policies and compute routes based on them.

There is no ONE best route to a destination.

Transit policies govern who/what may traverse resources and which path to destination is preferred.

Source route selection policies determine which legal paths are selected/used/preferred.

Two parallel efforts: hop-by-hop routing (e.g. BGP/IDRP) and source-demand routing (e.g., IDPR).

Hop-by-Hop: Node Routing

BGP/IDRP successor to EGP.

Basic Mechanisms:
- Hop-by-hop routing.
- Path Vector (PV) avoids loops.
- Path attributes, distrib. lists.

Scalability:
- Reachability info--hierarchical address assignment.
- Transit policies and topology--confederations.

NOT good for combinatorial explosion of multiple routing criteria.

Replicated routing tables for each possible combination.

Source-Demand Routing

IDPR designed to support policies of diverse ADs and TOSs:
On-demand installation of routes initiated by source.

Basic Mechanisms:
- Source routing avoids loops (w/setup)
- Link state style route computation with policy terms.

Flexibility from source-installation of legal, preferred, specialized, loop-free routes.

Avoids replicated routing tables; route installed only when needed.

But, generic routes (i.e., widely used) not as efficient with source demand routing--no aggregation.

Proposed Unified Architecture

Hop-by-hop, node routing (NR) for (pre) computing and installing generic routes.

Source Demand Routing (SDR) for more specialized routes installed on demand.

Architecture adapts to changing traffic/demand over time:
- e.g., as a TOS becomes widely used it can be precomputed instead of installed on demand.

Analogous to "off-the-shelf" vs. special order products.

Unified Architecture Disclaimer

#1
Will **not** support a special route for every source-destination pair in the Internet, **simultaneously**.

#2
We have **not** worked out all the details yet...

#3
This presentation describes more of a **hybrid** than a **unification**; expect to move more in the latter direction.

Evaluation Criteria for Design of NR and SDR components

Complexity:
- Storage Overhead (RIB, FIB)
- Route Computation Complexity
- Bandwidth overhead

Aggregation:
- Hide parts of topology,
- Combine reachable dests,
- Express constraints on sets.
- **Flexibility--multiple views:**
  - Overlapping confederations,
  - Aggregated or disaggregated,
  - Avoid centrally-coordinated, hierarchy of ADs.

Evaluation Criteria (continued)

Routing Policies Expressible on all routes

Individual Transit and Route Selection control based on : Endpoints, TOS, Path, User_Class, $

Information hiding

Not: policies **automatically** contingent on external state or behavior

TOS support

Does **not** imply choosing route based on instantaneous resource-availability (load)

Commonality of NR and SDR components.

Node Routing: PV vs. LS

Complexity:

Storage comparable, ignoring aggregation.

Computational complexity
Greater for LS with dissimilar route selection rules for each AD.

Bandwidth comparable and not as critical.

Aggregation

HBH routing with LS requires consistant network maps at all routing nodes--can not support multiple views, i.e., "inconsistant" aggregation.

Path Vector (PV) uses FP to eliminate looping, not consistant map.

SDR: PV vs. LS

Source routing used to avoid loops--do not need consistant network map.

LS provides source with more flexibility--available routes are not determined by arbitrary route selection policies of intermediate nodes.

...Functionality for SDR and aggregation for NR outweigh desireability of common mechanism.

Summary: Node Routing (NR)

Distributed computation of routes.

Full AD-path and policy attributes carried along with each distance vector. (AD or Confeds in path)

Routes installed and computed reflect most recent info re. operational status of routing facilities.

Route changes triggered by changed status (operational or policy).

Optimized for common case (steady) traffic.

NR nodes select most generic route available--multiple routes if widely used and predictable.

Summary: SDR

Need for special route is unpredictable--determined by source.

Inappropriate to burden intermediate nodes with selecting and storing (long-term) special routes.

SDR computes route at source and may use setup to install route at time of demand in intermediate nodes; else SR in pkt.

SDR uses configured information to compute routes; up-to-date operational status may or may not be available.

If computed route is not available, source notified and may receive update.

Hierarchy

NR groups domains to form hierarchy--good for scaling.

SDR treats domains as individuals--a single domain in a confederation may have privilges not shared with others, and may express policy different from others, in the domain.

Feature: SDR supports non-aggregated, non-hierarchical policies.

Cost: more challenging scaling problem.

Limit the number of **simultaneously** installed SDR routes to avoid n^2 problem.

Protocols

NR is approximately BGP/IDRP (nodes select generic routes).

SDR is approximately IDPR but...

#1 Alternative methods for **distribution/acquisition** of connectivity/policy info.

Reverse path update for **dynamic** information--exploit locality.

Selective (on demand) retrieval of **configured** connectivity info.

SDR Protocols (cont.)

#2 **Setup**
May map SDR to (completely) equivalent NR route.

Explicit SR (AD-level) in pkt option.

Setup flags: RPU, notification, default to NR on failure, ...

#3 **Network layer**
May be able to use IP+ (or CLNP) instead of special network layer header with unique path ID.

#4 **Intra-Domain routing**
May avoid encapsulation when IGP supports route injection.

#5 **Reachability Information**
May be included in updates, as alternative to using DNS lookup

Unification: Preliminary Thoughts

Opportunities to unify protocols to avoid doubling of implementation, operation, overhead, and management efforts.

#1 RIB: Directed graph for NR routes with generic policy bits, augmented by SDR links with non-generic policy bits.

#2 NR Path Vectors may include address(es) of repositories for SDR-updates for each AD.

#3 Replace IDPR global path ID with IP+ header fields; unifies packet header parsing and forwarding.

#4 Reachability information (Network# to AD# map) shared.

Status

Not intended for deployment in very near future.

Goal to **guide** current operational and research efforts.

Can be introduced incrementally.

Most critical new work needed:

> Experimentation and small-scale deployment:
> SDR Route installation and pkt forwarding.
>
> Research: connectivity information distribution.
>
> Deployment: BGP/IDRP

Proposed Plan

Continue IDPR research, prototype efforts: experience with SDR style routing **needed!!**

Continue IDRP standardization efforts.

BUT, start steering each effort towards unified architecture.

Orthogonal, but critical issue must be resolved (by IAB):

> Independent (S.I.N.) IP and OSI routing stacks, i.e., proceed with both BGP and IDRP, and SDR for each.
>
> Integrated IP and OSI, i.e., proceed with merged BGP/IDRP, and SDR.

6.5 Foiling the Denial of Service Threat

Presented by Radia Perlman/DEC

Certain security concerns such as preventing unauthorized disclosure or modification of information are well understood, and there are reasonable techniques, performed at the upper layer protocols, for addressing these threats. The denial of service threat is one in which an intruder can disable the network. All current Network Layer protocols can be disabled by a single router. The router might have been compromised by a saboteur, or it may simply be misconfigured, implemented improperly, or be experiencing hardware problems. When we think of a router failing, we usually think of a fail-stop failure, in which a router instantly makes a transition from operating perfectly to halting. A "Byzantine failure", named after the famous computer science problem known as the "Byzantine Generals problem" is one in which a node does not simply halt, but rather behaves erratically. We can assume a node with a Byzantine failure might send badly formed control messages, send well formed control messages with erroneous information, corrupt messages from other nodes, or simply flood the network with so much information that all resources become depleted. We can assume, either due to Murphy's law or due to the fact that malice might be involved, that the faulty router can intelligently choose the most disruptive course of action. However, it cannot do anything supernatural. For instance, if we have a secure cryptographic system, a faulty router cannot circumvent the cryptographic security.

We would like to design a network that will continue to deliver data between two nonfaulty nodes, provided that at least one path of nonfaulty routers and links connects the nodes. Furthermore, we would like our design to be practical. It should not be significantly more expensive or inconvenient than a conventional network. Our design must be reasonable to configure.

In this presentation I will present a design for a network that meets these goals. The first part of the design involves a robust scheme for reliable wide area broadcast, in which a source sends a packet that is received by all the routers. We use this for distribution of Link State Packets and for distribution of public keys for all the routers. To make the scheme work, buffers are reserved at each router for a packet from each other source. The source cryptographically signs the packet so that no other node can generate a packet that might occupy the buffer reserved for that source. The only configuration necessary is that each node will need to know the public keys of the (one or two) "trusted" nodes that broadcast the public keys of all the routers, and the trusted nodes need to be configured with the public keys of all the routers.

The robust broadcast scheme could be used for delivery of data traffic, but it has two disadvantages. One is that the broadcast scheme makes inefficient use of bandwidth when the data is intended for a single destination. The other is that routers must verify cryptographic signatures every time they receive a packet, which would considerably slow down forwarding rates. Thus instead of using the robust scheme for data traffic, we instead use a link state routing scheme. However, instead of having routers make hop by hop decisions for data forwarding, the source router selects a path and sets it up with a special cryptographically protected packet. Once a path is set up, assuming the source router was lucky enough to select a correctly functioning path, data can flow along that path without any cryptographic protection. If the source was unlucky, the path will not work. The upper layers will complain and the source can choose a different path.

Foiling the Denial of Service Threat

Radia Perlman

Digital Equipment Corporation

November 1991

Today

- In all current protocols, a single misbehaving router (or bridge) can disable the whole "network"
- Misbehavior can be caused by
  - -- software bugs
  - -- misconfiguration
  - -- ambiguous specification
  - -- incorrect protocol
  - -- flaky hardware
  - -- hackers

Misbehavior

- Things a router might do
  - -- Send garbage
  - -- Generate misleading control info
  - -- Delay other routers' control info
  - -- Corrupt control info
  - -- Send too many messages
- Example -- ARPANET disaster
  - -- One router was sick. Just before it died it generated a few LSPs from itself with random sequence numbers
  - -- LSPs became a virus
  - -- Miraculously, it was fixable:
    - . protocol designers were the implementers and field service
    - . all rtrs identical hardware and software

What I will present

- Network will deliver data between A and B provided that at least one nonfaulty path connects A and B
- The design is practical
  - -- Configuration is not much worse than conventional network
  - -- Memory, bandwidth requirements are not much worse than conventional networks
  - -- CPU is not much worse than conventional networks
  - -- CPU involved in forwarding a data packet is NO worse than in a conventional network

Outline of Solution

- First design a mechanism for broadcasting information. Source's packet gets delivered to everyone.
- That scheme requires knowledge of public keys of all the routers
- Use the robust broadcast mechanism for two things, both of which require broadcast
  - -- Distribution of public keys of all the routers
  - -- Distribution of LSPs
- With the LSP database, have the source router compute a route
- Issue a route setup packet
- Have data packets follow the path chosen by the source router

Robust Broadcast

- Flooding guarantees delivery provided that sufficient resources exist
  - -- Bandwidth, Memory, CPU
- Routers can be engineered with sufficient CPU -- worst case processing requirement is known -- link bandwidth is finite
- Memory is preallocated, with a buffer for each source
- Bandwidth is allocated fairly by running through memory round robin
- To ensure packet occupying FOO's buffer is really from FOO, use public key signature scheme
- Assume all routers know the source's public key
- To ensure it's the latest packet from FOO, use sequence numbers

Database

	nbr 1	nbr 2	nbr 3	nbr 4	nbr 5	nbr 6
Src X seq # a signat.	T		T	T	T	A
Src Y seq # b signat.		T		T		T
Src Z seq # c signat.						
Src W seq # d signat.		T				A

T = transmit pkt   A = send Ack

Events/Actions

- Receive packet from source S, via neighbor N, with sequence number k. Current packet stored from S has # j.
  - -- is signature valid?
  - -- if k > j, overwrite packet in memory and set T flag for all neighbors but N. Set A flag for N
  - -- if k = j, set A flag for N
  - -- if k < j, set T flag for N
- Receive Ack from neighbor N, about source S, sequence number k
  - -- if k = j, clear flag for N
  - -- else, set T flag for N
- Link to neighbor N is free -- continue round robin scan. If find flag, send packet or ACK (pkt if flag=T, ack if flag=A)

Knowing Public Keys

- We could require configuration of all keys into all routers -- that would be horrible
- Instead, have a "Trusted public key distributor".
  - -- Everyone has to be configured with its public key
  - -- It has to be configured with all routers' public keys
  - -- It broadcasts a packet containing the current list of all nodes and public keys
- What about bad trusted PKD? Use multiple PKDs
  - -- majority rules (requires at least 3)
  - -- Set aside a buffer for every possible public key. Then if only one PKD is OK, some resources will be reserved for good guy

What about data packets?

- LSPs and public key lists need to go to all routers
- It is inefficient to deliver a packet from S to D via broadcast
- Conventional hop by hop route decisions require synchronized databases at all routers. May be possible, but I don't know how
- Instead, use source routing, as computed by source router
  - -- databases need not agree
  - -- can use complex routing algorithm, like avoiding routers on past failed paths

Route Setup

- When S chooses a route to D, it puts the chosen route into a route setup packet
- Route setup packet is cryptographically signed by S
- The route setup packet travels along the specified path
- Routers along the path remember for the pair S/D the direction from which a packet should arrive, and the direction to forward a packet. They also reserve a buffer for data packets from S to D
- When a data packet arrives from neighbor N, with source S and destination D
  - -- check if N is where S/D traffic should arrive
  - -- if so, forward the packet in the set up direction

Costs of this scheme

- Configuration
  - -- Each router needs
    - . its own private key
    - . PKD's public key
  - -- PKD needs public keys of all routers
- Need storage for public key list
- LSPs need to have signature
  - -- 500 bits
  - -- CPU to verify signature
- Data exchange requires route setup first
- But, data packets no larger than before, and no cryptography required to forward data packets

Problems

- Computed route may not work
  - -- can try a different route
  - -- can try a node disjoint route
  - -- can try to do failure analysis
  - -- can use broadcast as fallback, so in usual case things are as efficient as now, and in the presence of malice, routes unlucky enough to trip over bad guys will resort to broadcast
- Sick source can run out of sequence numbers
  - -- No matter how large the space, this is possible
  - -- Fix is to change that source's public key!

Summary

- Even with a significant percentage of malicious routers, data can flow between S and D provided at least one path connects S and D
- The network does not provide authentication -- garbage might be delivered also -- that's a problem for the higher layers
- The design can be extended to hierarchical networks
- The design is practical
  - -- configuration
  - -- memory
  - -- bandwidth
  - -- forwarding speed

6.6 Development of ATM and the ATM Standards

Presented by George Clapp/Ameritech

SMDS, Metropolitan Area Networks, & Broadband ISDN

George H. Clapp

Ameritech Services, Inc.
Science and Technology
The Meadows Corporate Center
2820 West Golf Road
Rolling Meadows, IL 60008-4014
phone: 708-806-8318
fax: 708-806-8292
email: meritec!clapp@uunet.uu.net

AMERITECH SERVICES

What is SMDS?[SM]

- High-speed, connectionless, public, packet switching service which will extend LAN-like performance beyond the subscriber's premises.
- Defined in a set of Technical Advisories (e.g., TA-TSY-000772) released by Bell Communications Research (Bellcore).
- Transmission rates
  - DS3 (44.736 Mb/s line signaling rate with 44.209 Mb/s payload).
  - DS1 (1.544 Mb/s line signaling rate with 1.536 Mb/s payload).
  - Anticipated that it will operate at the CCITT Synchronous Digital Hierarchy (SDH) Synchronous Transport Module (STM-1) speed (155.520 Mb/s line signaling rate with 149.760 Mb/s payload).
- Issue 3 was released in October of 1989; updated in December of 1990.

SMDS is a service mark of Bellcore

George H. Clapp

AMERITECH SERVICES

What is an IEEE 802.6 MAN?

- Standardization of the ***Distributed Queue Dual Bus (DQDB)*** Medium Access Control (MAC) algorithm, which resolves contention to a shared medium, broadcast network.
- Extension of a Local Area Network in speed, distance, and number of users.
- Integrated transport of high speed data, voice, and compressed video.
- ≥50 km in diameter.
- Primary service is high speed connectionless data transport and switching.
- Initial transmission line signaling rate will DS3 (45 Mb/s) with extension to SONET (Synchronous Optical NETwork) rates (155.52 Mb/s).
- Focus was on the public network.
- Standardization work began in April 1981.
- Standard was approved in December 1990.

George H. Clapp

AMERITECH SERVICES

What is Broadband ISDN?

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

George H. Clapp

AMERITECH SERVICES

MAN Architecture
Introduction of a Central Switch

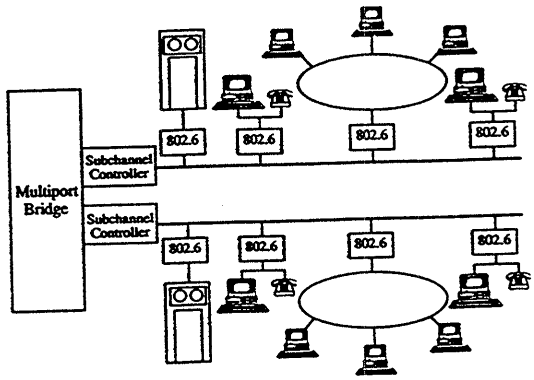

George H. Clapp

AMERITECH SERVICES

Target Broadband Data Architecture

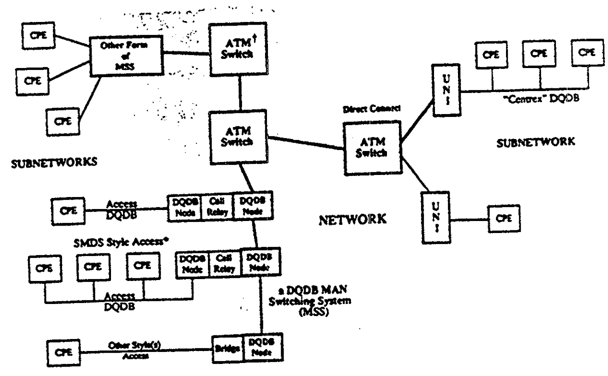

* Per Current Bellcore TA-772
† ATM Switches can initially be DQDB Bridges

George H. Clapp

AMERITECH SERVICES

SMDS Network Architecture

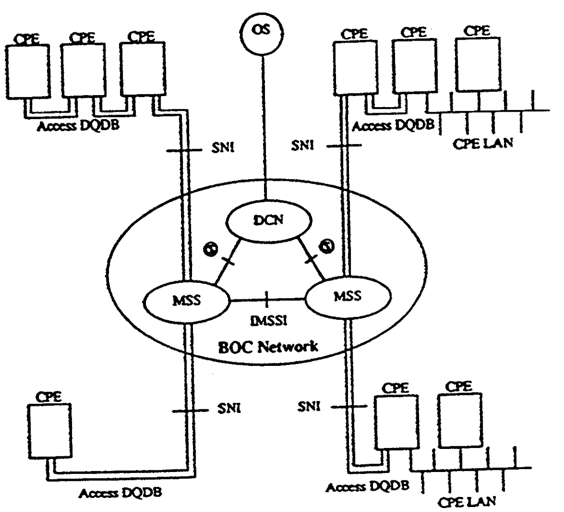

SNI: Subscriber Network Interface
MSS: MAN Switching System
IMSSI: Inter-MAN Switching System Interface
Ⓢ: Generic Interface for Operations
OS: Operations System

DCN: Data Communications Network
CPE: Customer Premises Equipment
CPE LAN: CPE Local Area Network
DQDB: Distributed Queue Dual Bus

George H. Clapp

AMERITECH SERVICES

Broadband ISDN Interface:
Asynchronous Transfer Mode (ATM)

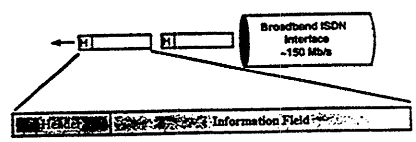

Characteristics:

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

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

George H. Clapp

AMERITECH SERVICES

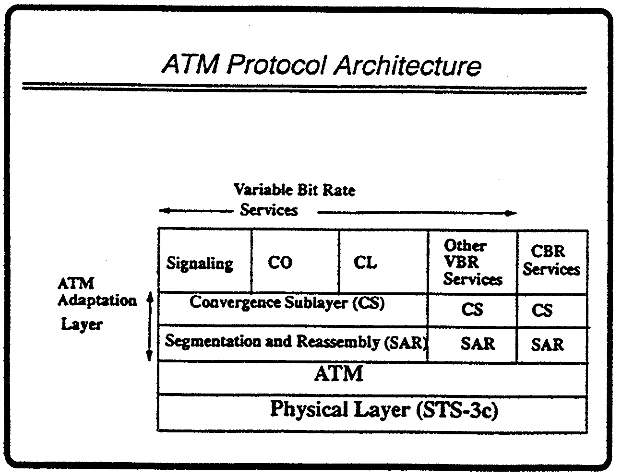
ATM Protocol Architecture
Variable Bit Rate Services
ATM Adaptation Layer
Signaling
CO
CL
Other VBR Services
CBR Services
Convergence Sublayer (CS)
CS
CS
Segmentation and Reassembly (SAR)
SAR
SAR
ATM
Physical Layer (STS-3c)

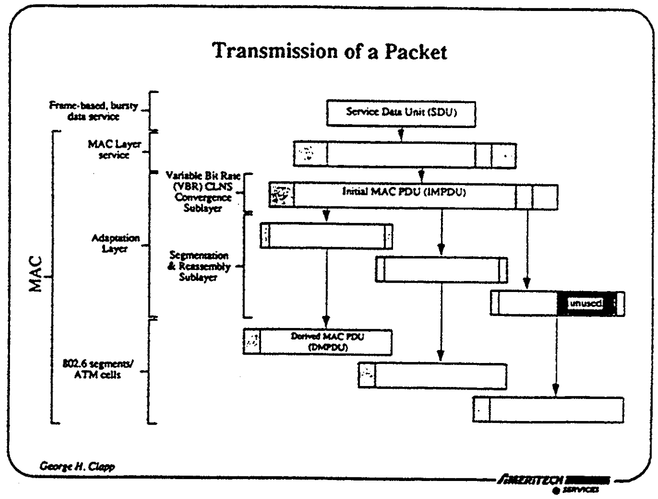
Transmission of a Packet
Frame-based, bursty data service
MAC Layer service
MAC
Variable Bit Rate (VBR) CLNS Convergence Sublayer
Adaptation Layer
Segmentation & Reassembly Sublayer
802.6 segments/ ATM cells
Service Data Unit (SDU)
Initial MAC PDU (IMPDU)
unused
Derived MAC PDU (DMPDU)
George H. Clapp
AMERITECH SERVICES

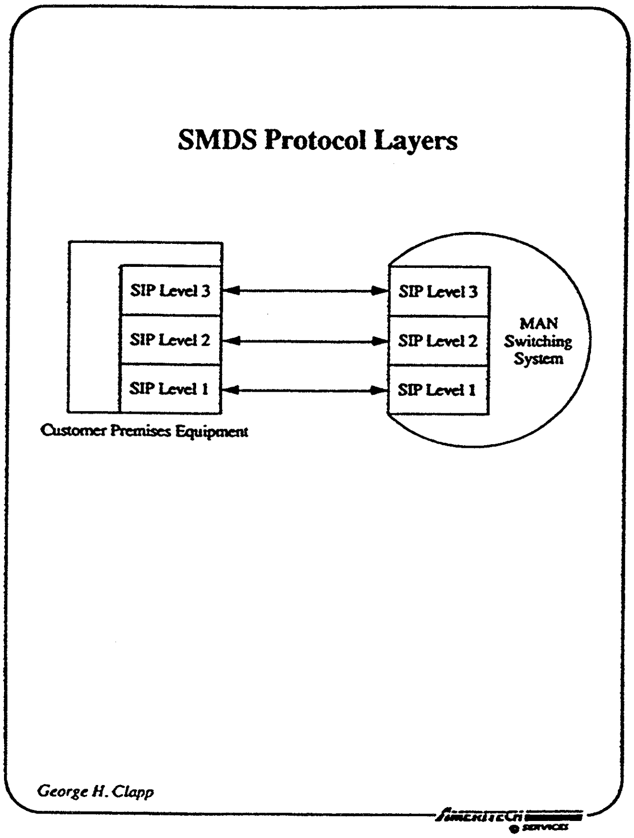
SMDS Protocol Layers
SIP Level 3
SIP Level 2
SIP Level 1
SIP Level 3
SIP Level 2
SIP Level 1
MAN Switching System
Customer Premises Equipment
George H. Clapp
AMERITECH SERVICES

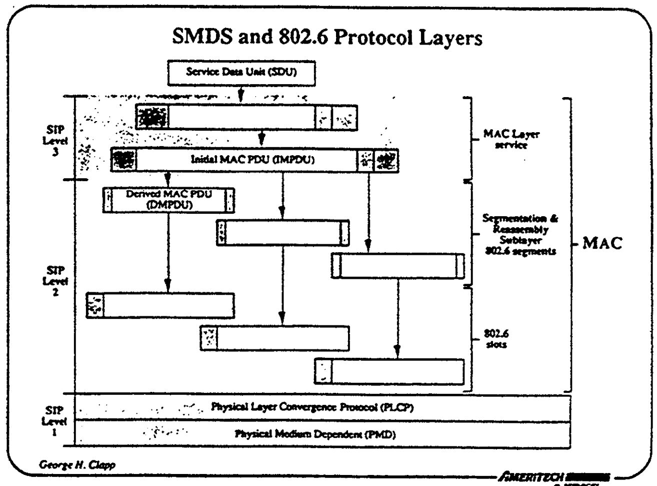
SMDS and 802.6 Protocol Layers
Service Data Unit (SDU)
SIP Level 3
Initial MAC PDU (IMPDU)
Derived MAC PDU (DMPDU)
SIP Level 2
SIP Level 1
MAC Layer service
Segmentation & Reassembly Sublayer 802.6 segments
MAC
802.6 slots
Physical Layer Convergence Protocol (PLCP)
Physical Medium Dependent (PMD)
George H. Clapp
AMERITECH SERVICES

802.6 and BISDN Protocol Layers

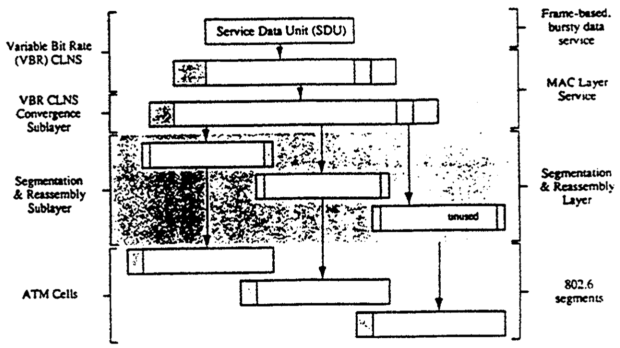

George H. Clapp

AMERITECH SERVICES

B-ISDN User-Network-Interface (UNI) ATM Cell Format

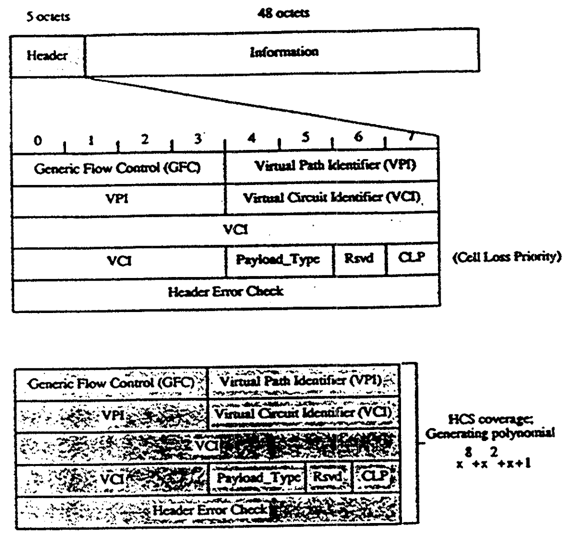

George H. Clapp

AMERITECH SERVICES

T1S1.5 Payload Type

RESERVED BIT ADDED TO FIELD.

000	USER NORMAL
001	USER SPECIFIC
010	USER NORMAL, CONGESTION EXPERIENCED
011	USER SPECIAL, "
100	RESERVED
101	RESERVED
110	NETWORK OA+M
111	NETWORK OA+M

EFFECTS OF DELAY

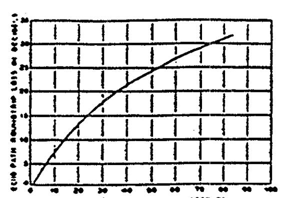

Fig. 6 — Average listener judgment of minimum echo attenuation necessary for just-tolerable talker echo conditions on modern telephone sets.

NOTES:

1. 68% of customers tolerate ± 5dB of curve.
2. Test used a typical loop, delay/loss measured from CO to point of reflection and back.

SOURCE: "The Effects of Time Delay and Echoes on Telephone Conversation." J. W. Emling and D. Mitchell, BSTJ, Nov. 1963, p. 2869 - 2892.

"Echo and Its Effects on the Telephone User." Bell Laboratories Record, G. M. Phillips Aug. 1954, p. 281-284.

"Echo Suppressors for Long Telephone Circuits." Bell Laboratories Monograph #165, June 1925.

PRIMARY SOURCE OF ECHO

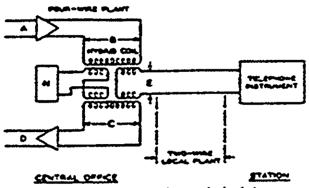

Fig. 4—Occurrence of echo in two-wire local plant.

1. Normal hybrids provide 11db echo return loss, ± 3db for all loops by a compromise impedance, N.
2. Selection of N based on loop type (loaded/non-loaded) significantly improves echo return loss (15db?). This is common practice.

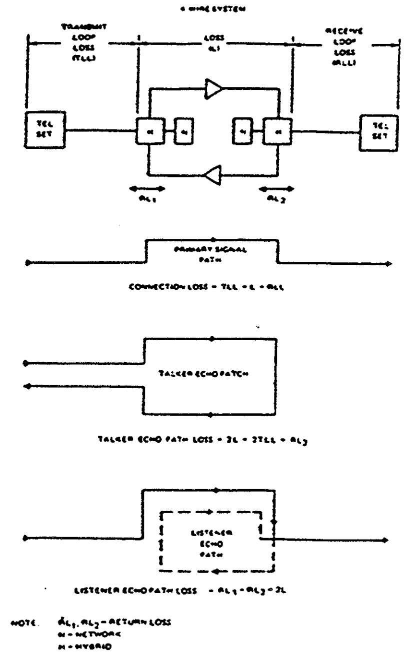

Figure 7.1-2. Talker and Listener Echo Paths

Voice-Related Objectives

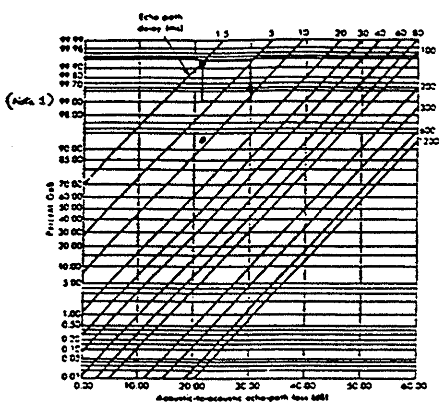

Source: Transmission Systems for Communications, AT&T Bell Laboratories, 1982, p. 151

Notes: Go B = "Good or better"

1. See "Performance Effect of Clear Channel Alternatives," T. C. Spang, T1Q1 Submittal, Oct. 1985.

ATM / ECHO OPTIONS

Block Size (Octets)	Echo Delay (Msec)	Echo Problem	Echo Control
≥ 152	≥ 38	Talker echo	Cancel (note 1)
32 - 152	8 - 38	Listener	Attenuate or cancel (note 2)
1 - 32	≤ 8	None*	None (note 3)

* Best case possible. Assumed only delays are packetization delays

NOTES:

1. Standard North-American practice for inter-exchange.
2. See "Voice Transmission over mixed packet and circuit switched networks", R. P. Singh and S. Singhal, Infocom Proceedings Mar. 1988., p. ZA.3.1.
3. True for intra-office only. Inter-office may require echo cancelers.

802.6 Segment Format

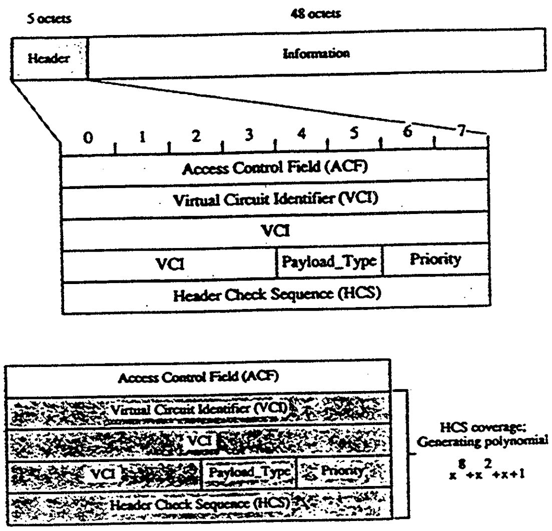

George H. Clapp

B-ISDN Network-Network-Interface (NNI) ATM Cell Format

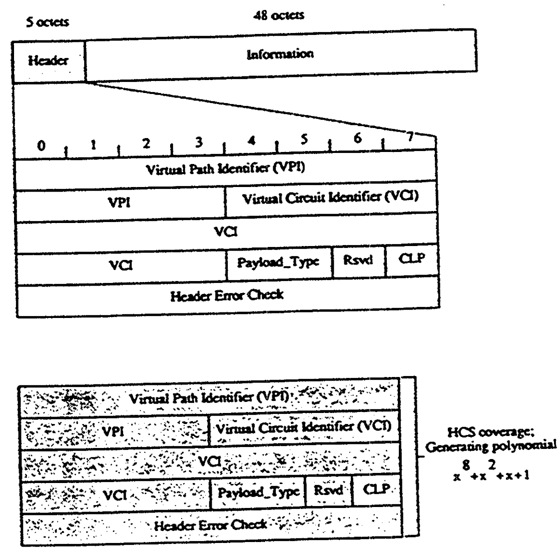

George H. Clapp

Virtual Path Concept

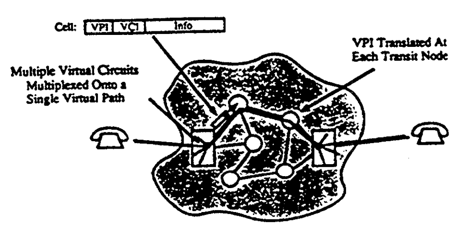

Advantages:

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

George H. Clapp

Header Check Sequence

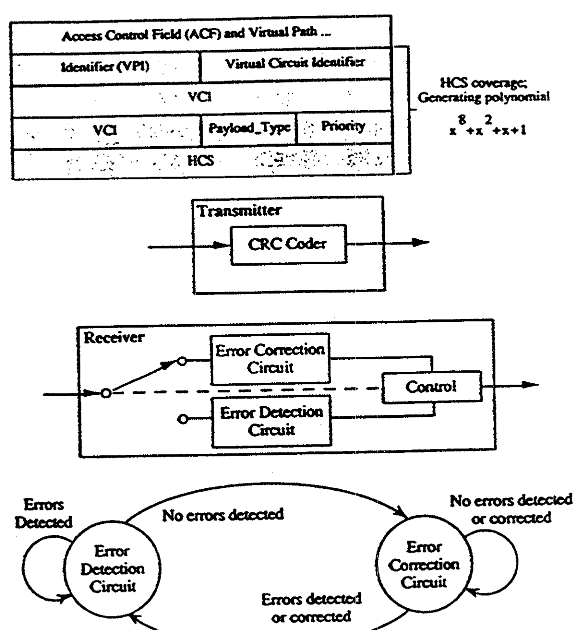

Transitions occur upon reading a segment.

George H. Clapp

SMDS SIP L_2 PDU & IEEE 802.6 DQDB DMPDU

802.6 DMPDU

1 bit	1 bit	1 bit	2 bit	3 bits	20 bits	2 bits	2 bits	8 bits	2 bits	4 bits	10 bits	44 octets	6 bits	10 bits
Busy	SL type	PSR	rsvd	REQs	VCI	Payload Type	Priority	HCS	Segment Type	Seq No	MID	Payload	Payload Length	Payload CRC

From CPE towards the network*

Busy	†	†	†	REQs	11111111111111111111	00	00	00100010	Segment Type	Seq No	MID	Payload	Payload Length	Payload CRC

From the network towards CPE

Busy	0	0	00	000	11111111111111111111	00	00	00100010	Segment Type	Seq No	MID	Payload	Payload Length	Payload CRC

*The network may operate in one of four modes (0 through 3):
- In modes 0 through 2, the network operates at REQ priorities 0 through 2, respectively.
- In mode 3, the network *disregards* the REQ field and sends an L_2 PDU whenever one is ready for transmission.

†These fields are not inspected by the network.

George H. Clapp

AMERITECH SERVICES

IEEE 802.6 Adaptation Layer
Segmentation and Reassembly (SAR) Sublayer

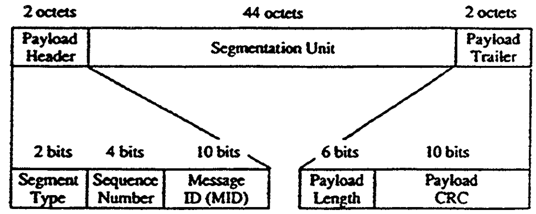

Segment Type — Used for packet delineation;
Encoding
- Beginning of Message (BOM): 10
- Continuation of Message (COM): 00
- End of Message (EOM): 01
- Single Segment Message (SSM): 11

Sequence Number — Values 0-15 are cycled through to number consecutive segments of a packet; value is incremented relative to the previous value for a particular VCI/MID value.

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

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

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

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

George H. Clapp

AMERITECH SERVICES

CCITT BISDN Adaptation Layer
Segmentation and Reassembly (SAR) Sublayer (the same)

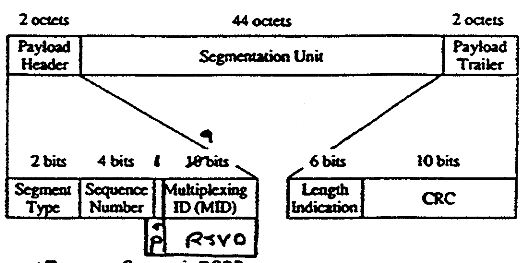

Segment Type — Same as in DQDB.

Sequence Number — Values 0-15 are cycled through to sequentially number consecutive segments of a packet; used for detection of out-of-sequence segments.

P — PRIORITY (FOR AAL 3)

Multiplexing ID — Shortened from 14 to 10 bits. Used as in DQDB to associate CLNS segments of a single packet. (FOR AAL 4)

RSVD — FOR AAL 3 (IF NO MID MUXING DONE).

Length Indication — Same as in DQDB.

CRC — Same as in DQDB.
Generating polynomial: $G(x) = x^{10} + x^{9} + x^{5} + x^{4} + x + 1$

George H. Clapp

AMERITECH SERVICES

AAL 3+4 COMMON PART CS (CPCS)
SERVICE SPECIFIC CS (SSCS)

Adaptation Layer
VBR CLNS Convergence Protocol Sublayer, aka, Common PDU Header & Trailer
CPCS-PDU FORMAT

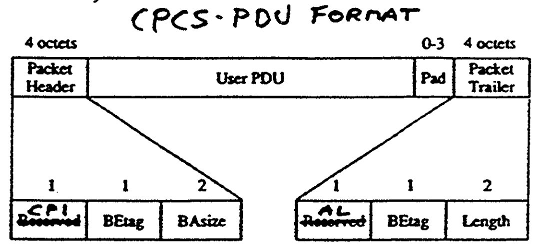

CPI — COMMON PART INDICATOR, USED TO INTERPRET FOLLOWING FIELDS, 00 = CL.

BEtag — Same value is placed in the header and trailer fields; used to associate header and trailer of the same PDU for error control. Values 0-255 are cycled through.

BAsize — Used by receiver for buffer management; either...
Length in octets of ~~802.6 IMPDU~~ PAYLOAD (header and information, inclusive), or...
Greater than or equal to the true PDU length.

Length — Length in octets of the ~~user PDU~~ CPCS-PDU PAYLOAD (less the Pad).

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

AL — ALIGNMENT (TO 4 OCTETS)

George H. Clapp

AMERITECH SERVICES

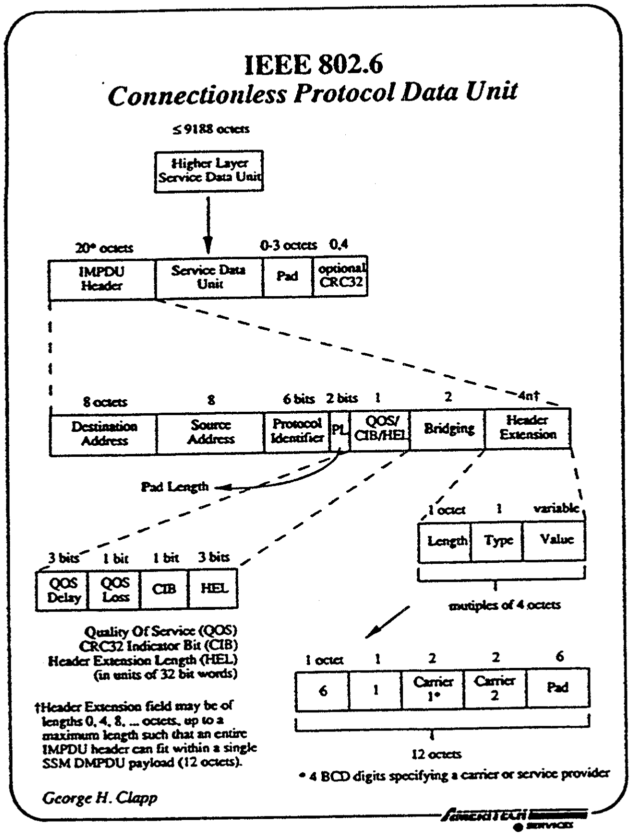
IEEE 802.6
Connectionless Protocol Data Unit
≤9188 octets
Higher Layer Service Data Unit
20* octets
0-3 octets
0,4
IMPDU Header
Service Data Unit
Pad
optional CRC32
8 octets
8
6 bits
2 bits
1
2
4n†
Destination Address
Source Address
Protocol Identifier
PL
QOS/ CIB/HEL
Bridging
Header Extension
Pad Length
1 octet
1
variable
Length
Type
Value
multiples of 4 octets
3 bits
1 bit
1 bit
3 bits
QOS Delay
QOS Loss
CIB
HEL
Quality Of Service (QOS)
CRC32 Indicator Bit (CIB)
Header Extension Length (HEL) (in units of 32 bit words)
1 octet
1
2
2
6
6
1
Carrier 1*
Carrier 2
Pad
12 octets
†Header Extension field may be of lengths 0, 4, 8, ... octets, up to a maximum length such that an entire IMPDU header can fit within a single SSM DMPDU payload (12 octets).
* 4 BCD digits specifying a carrier or service provider
George H. Clapp
AMERITECH SERVICES

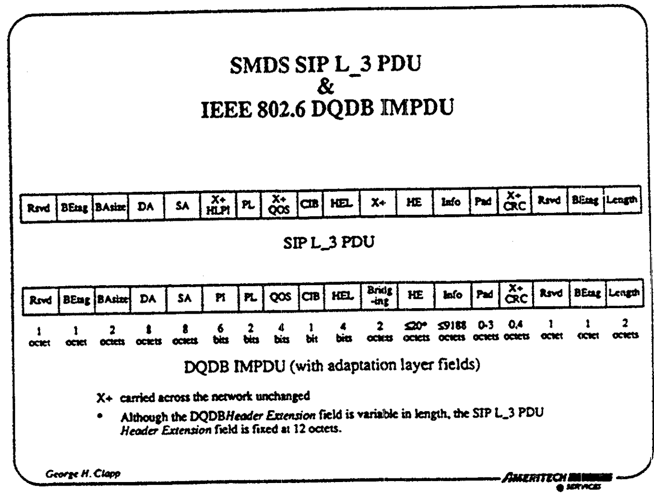
SMDS SIP L_3 PDU
&
IEEE 802.6 DQDB IMPDU
Rsvd
BEtag
BAsize
DA
SA
X+ HLPI
PL
X+ QOS
CIB
HEL
X+
HE
Info
Pad
X+ CRC
Rsvd
BEtag
Length
SIP L_3 PDU
Rsvd
BEtag
BAsize
DA
SA
PI
PL
QOS
CIB
HEL
Bridg -ing
HE
Info
Pad
X+ CRC
Rsvd
BEtag
Length
1 octet
1 octet
2 octets
8 octets
8 octets
6 bits
2 bits
4 bits
1 bit
4 bits
2 octets
≤20* octets
≤9188 octets
0-3 octets
0,4 octets
1 octet
1 octet
2 octets
DQDB IMPDU (with adaptation layer fields)
X+ carried across the network unchanged
* Although the DQDB Header Extension field is variable in length, the SIP L_3 PDU Header Extension field is fixed at 12 octets.
George H. Clapp
AMERITECH SERVICES

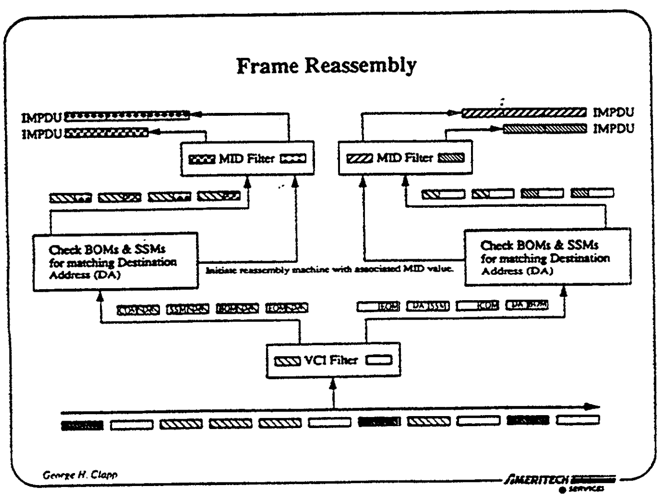
Frame Reassembly
IMPDU
IMPDU
IMPDU
IMPDU
MID Filter
MID Filter
Check BOMs & SSMs for matching Destination Address (DA)
Check BOMs & SSMs for matching Destination Address (DA)
Initiate reassembly machine with associated MID value.
VCI Filter
George H. Clapp
AMERITECH SERVICES

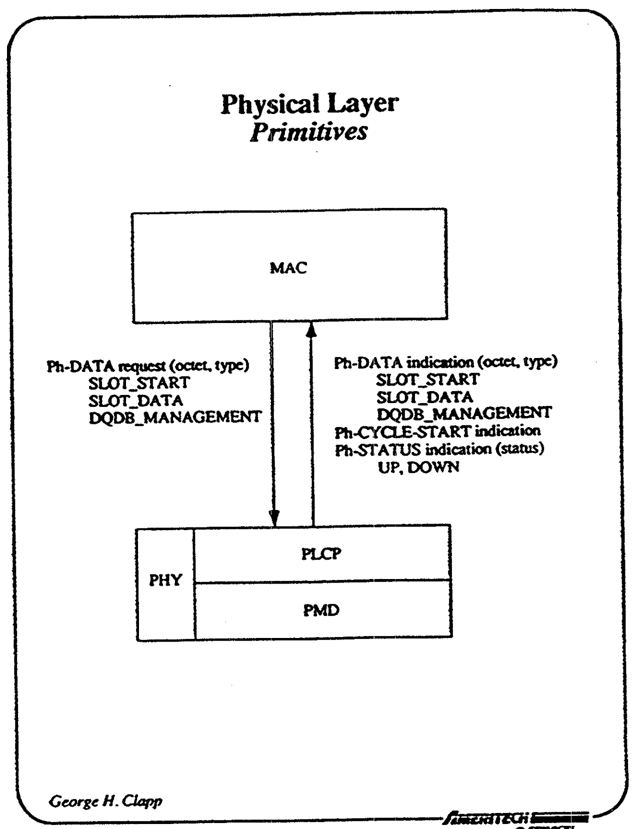
Physical Layer
Primitives
MAC
Ph-DATA request (octet, type)
SLOT_START
SLOT_DATA
DQDB_MANAGEMENT
Ph-DATA indication (octet, type)
SLOT_START
SLOT_DATA
DQDB_MANAGEMENT
Ph-CYCLE-START indication
Ph-STATUS indication (status)
UP, DOWN
PHY
PLCP
PMD
George H. Clapp
AMERITECH SERVICES

Physical Layer Convergence Procedure (PLCP)

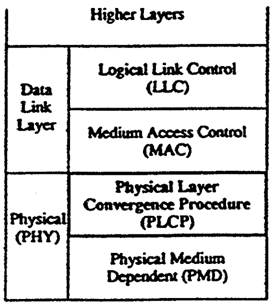

Functions of PLCP:

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

George H. Clapp

Digital Hierarchy Bit Rates
(CCITT G.702)

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

†European Conference of Post and Telecommunication Administrations

George H. Clapp

North American Digital Signal (DS) Transmission Hierarchy

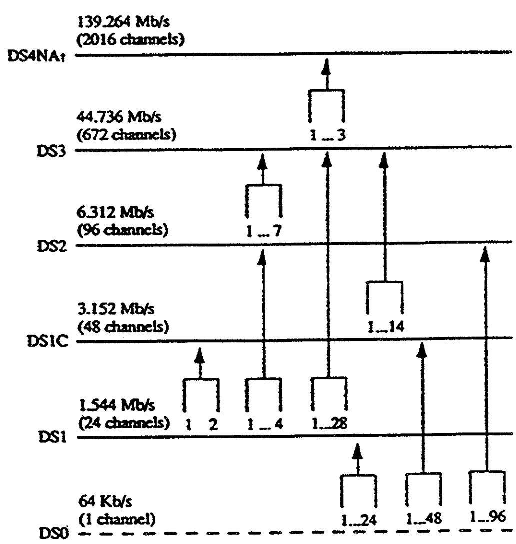

George H. Clapp † DS4NA: Digital Signal 4 North American

DS1 Transmission
Extended Superframe

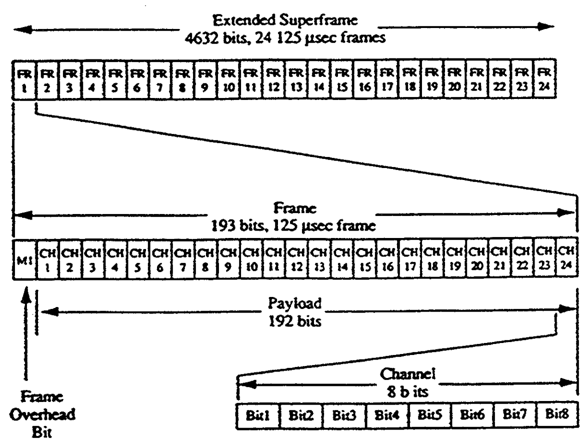

1.536 Mb/s payload.

First bit of each frame is used for transmission overhead.

George H. Clapp

Proposed DS1 Mapping
SMDS

A1	A2	P9	Z4	802.6 segment 1	
A1	A2	P8	Z3	802.6 segment 2	
A1	A2	P7	Z2	802.6 segment 3	
A1	A2	P6	Z1	802.6 segment 4	
A1	A2	P5	F1	802.6 segment 5	3 msec
A1	A2	P4	B1	802.6 segment 6	
A1	A2	P3	G1	802.6 segment 7	
A1	A2	P2	M2	802.6 segment 8	
A1	A2	P1	M1	802.6 segment 9	
A1	A2	P0	C1	802.6 segment 10	6 octets

A1 Framing Byte (F6 hex)
A2 Framing Byte (28 hex)
P0-P9 Path Overhead Identifier Octets
B1 Bit Interleaved Parity 8 (BIP-8)
G1 PLCP Path Status
F1 PLCP Path User Channel
Z1-Z2 Growth
M1-M2 DQDB layer management information
T1-T2 Timing Source Counter
C1 Cycle/Stuff Counter

10 slots × 48 octets/slot × 8 bits/octet / 3 msec = 1.280 Mb/s total payload.
10 slots × 44 octets/slot × 8 bits/octet / 3 msec = 1.173 Mb/s payload less SAR sublayer.

George H. Clapp

AMERITECH SERVICES

DS3 Transmission

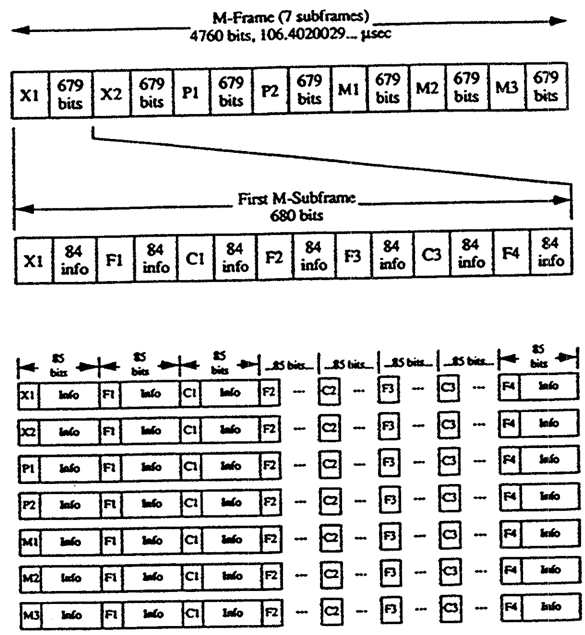

George H. Clapp

AMERITECH SERVICES

DS3 Mapping
IEEE 802.6 Working Group

Impose 125 µsec frame upon DS3 payload:
84/85 × 44.736 Mb/s = 44.2097 Mb/s payload
44.2097 Mb/s × 125 µsec = 5526.211765... bits
= 690.776... octets
690.776 / 53 octets/cell = 13 cells/cycle + 1.776 octet (too little)
12 slots × 48 × 8 / 125 µsec = 36.864 Mb/s total payload
12 slots × 44 × 8 / 125 µsec = 33.792 Mb/s user payload (less SAR sublayer)

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

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

George H. Clapp

AMERITECH SERVICES

SONET
Synchronous Transport Signal level 1

STS-1: basic logical building-block signal

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

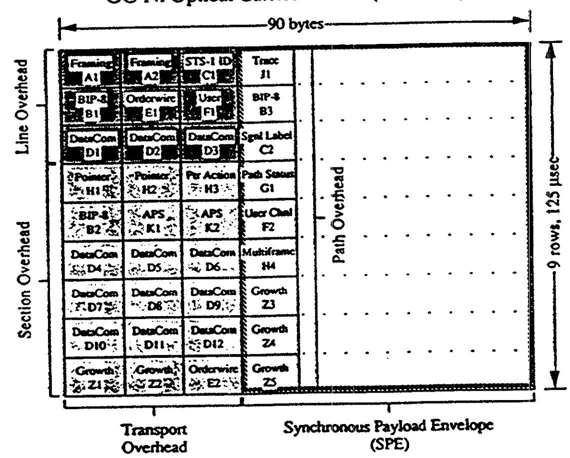

9 rows × 90 bytes × 8 / 125 µsec = 51.840 Mb/s line signaling rate.
9 rows × 87 bytes × 8 / 125 µsec = 50.112 Mb/s SPE rate.
9 rows × 86 bytes × 8 / 125 µsec = 49.536 Mb/s user data rate.

George H. Clapp

AMERITECH SERVICES

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

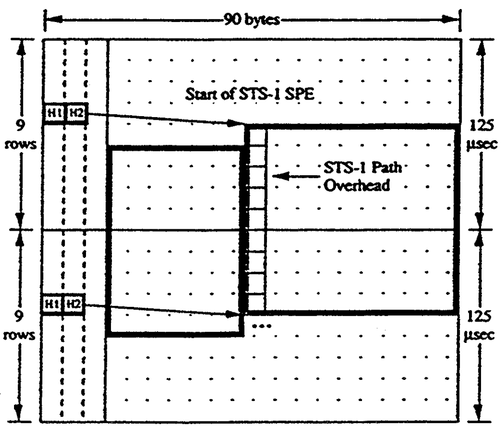

H1 & H2 pointers in Line Overhead used to adjust for "slips" and "stuffs" across plesiochronous boundaries.

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

George H. Clapp

AMERITECH SERVICES

SONET
Synchronous Hierarchical Rates

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

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

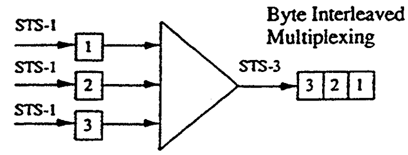

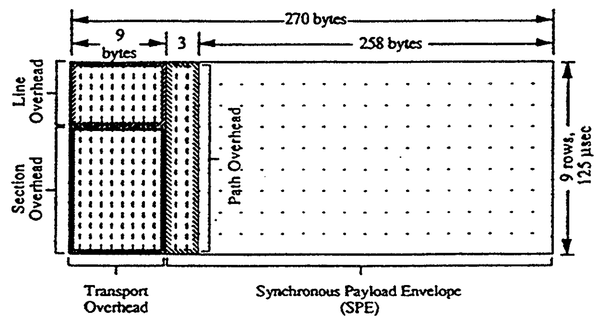

George H. Clapp

AMERITECH SERVICES

Standard CCITT SONET Rates
Synchronous Transport Module Levels

Supports "super-rate" services which require multiples of the STS-1 payload capacity, e.g., Broadband ISDN H4 channel.

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

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

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

STM-1 is the CCITT basic building block.

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

†Candidate rates which are not standardized.

George H. Clapp

AMERITECH SERVICES

CCITT STM-1 Format

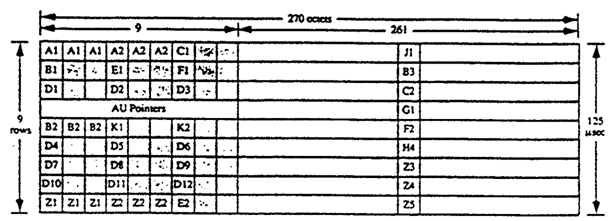

A1, A2	Framing	B3	BIP-8
AU Pointers	Administrative Unit Pointers	C2	STM Identifier
B1	BIP-8 (Bit Interleaved Parity)	F2	User-Defined Channel
B2	BIP-24	G1	Path Status
C1	STM Identifier	H4	Multiframe Indicator
D1-D12	Data Communications Channel	J1	Path Trace
E1, E2	Order Wire	Z3-Z5	Growth (reserved as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

George H. Clapp

AMERITECH SERVICES

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

Three mechanisms:

- Pointer.
  - "Multiframe Indicator" octet (H4) of the Path Overhead points to the beginning of the first complete cell following the H4 octet.
- Cell counting.
- Header Check Sequence (HCS) calculation.
  - Following cell boundary identification via the H4 octet, the HCS is checked to verify cell delineation.

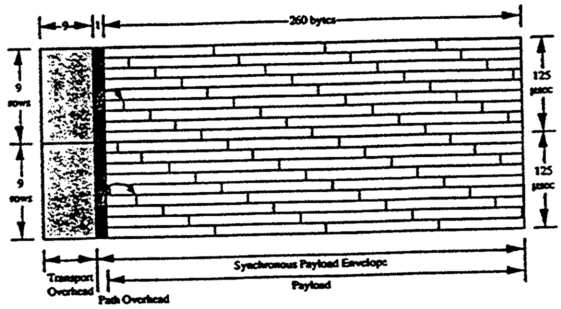

$$\frac{260 \text{ bytes} \times 8 \times 9}{125\ \mu\text{sec}} = 149.760 \text{ Mb/s} \qquad \frac{(260 \times 9) \text{ bytes / frame}}{53 \text{ bytes / cell}} = 44\frac{8}{53} \text{ cells / frame}$$

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

George H. Clapp

AMERITECH SERVICES

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

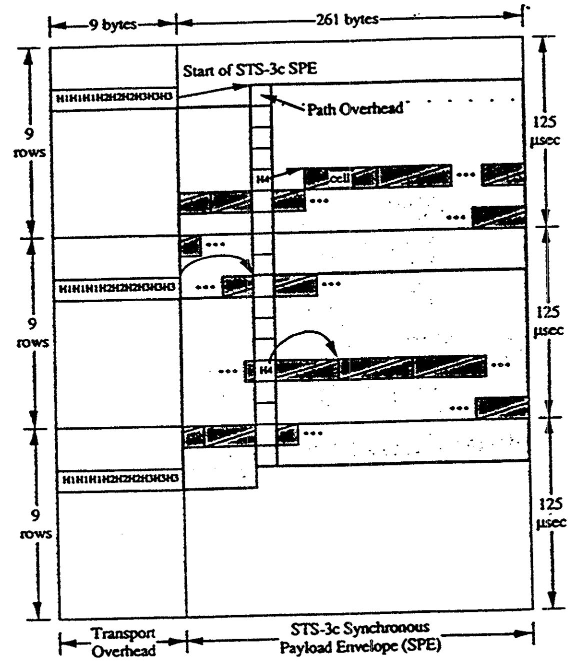

George H. Clapp

AMERITECH SERVICES

SONET
ATM Cell Delineation

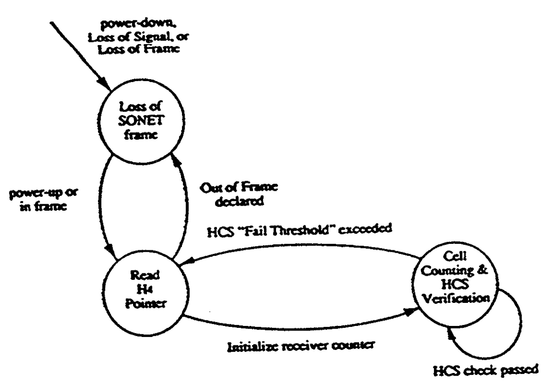

George H. Clapp

AMERITECH SERVICES

6.7 Advanced Network Activities at Los Alamos

Presented by John Morrison/LANL

Advanced Networking Activities at Los Alamos

John Morrison

jfm@lanl.gov
+1-505-667-3310

Outline

- HIPPI
- Crossbar Interface
- Multiple Crossbar Network Testbed
- CASA

High Performance Parallel Interface (HIPPI)

- 800 Mbit/s (100 MByte/s) speed range
- Point-to-point
- Simplex (use 2 for full-duplex)
- Intended for memory-to-memory transfers
- Supports both circuit and packet switching
- ANSI Standard September 1991

HIPPI Features

- 100 & 200 Mbyte/s options (800 and 1600 Mbit/s)
- 25 meters using copper cable
- Simplex channel
- Flow control for full speed operation over long distances
- Byte parity in 32-bit or 64-bit words
- Checksum every 256 words
- Very simple signalling sequences
- Easily implemented in off-the-shelf parts
- No bidirectional signal lines
- Supports addressing for networking

HIPPI Protocols

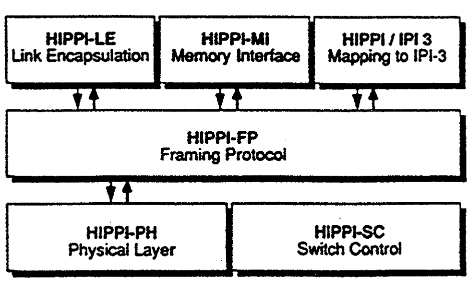

HIPPI Standards Status

Document	Revision	Status
HIPPI-PH	8.2	ANSI Standard X3.183
HIPPI-FP	4.2	Out for X3 Public Review
HIPPI-LE	3.2	X3 Public Review in Nov 91
HIPPI-MI	2.2	Under development
HIPPI-SC	2.2	Forward to X3T9 in Oct 91
HIPPI-IPI-3		Under development

Documents available via anonymous ftp from nsco.network.com

HIPPI Network

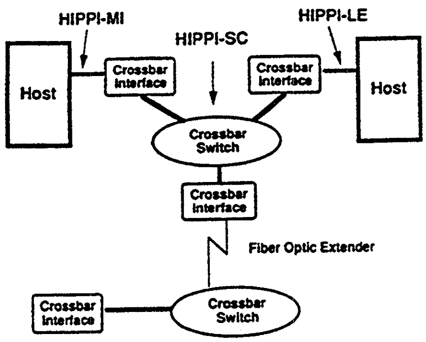

CBI Functions

- Network Management
  - Configuration management
  - Performance measurement
- Protocol Support
  - Hardware support for TCP/IP checksum calculation
  - TCP/IP Engine
- Buffering
  - Decouples host from network data rates
  - Provides multiplexed packet stream
- Switch Control/Routing
  - Bind HIPPI-LE addresses to HIPPI-SC addresses
  - Extends network to multiple switches transparently

Crossbar Interface (CBI)

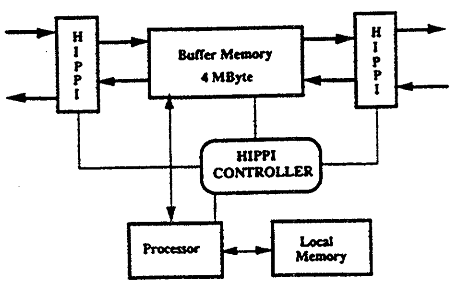

CM-2 HIPPI

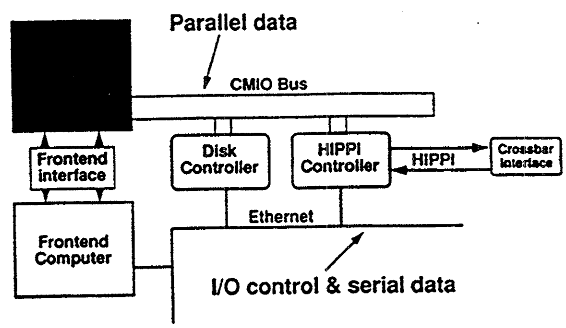

Multiple Crossbar Network

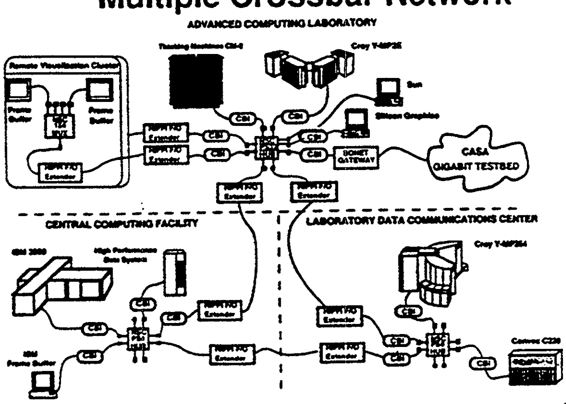

Visualization

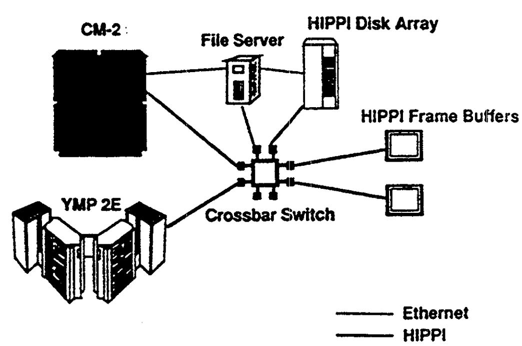

HIPPI-SONET Interface BLOCK DIAGRAM

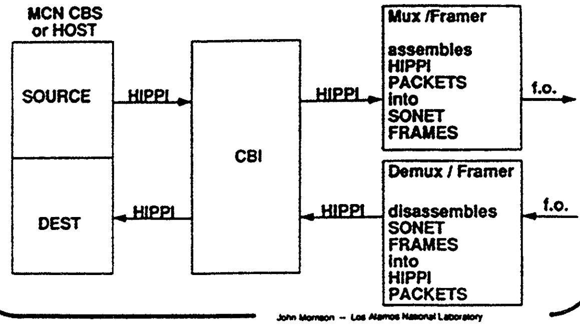

HIPPI-SONET Framer

HIPPI-SONET interface to the SONET LTE or A/D MUX will consist of eight (8) OC-3 tributaries.

Advantages to us of using OC-3 as our local interface:

- **Use Bellcore OC-3 integrated circuits**
- **Only moderate speed electrical signalling required**
- **Reduced rate or degraded service possible**

The carrier interface to the LTE or A/D MUX will be OC-48.

Chemical Reaction Dynamics

- Model F + H reaction
- Relevant to power lasers
- Large Matrix Operations

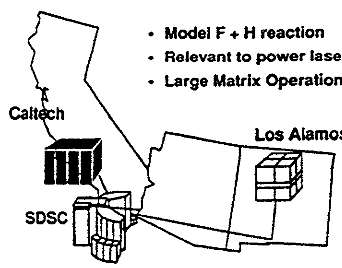

Calcrust

- Combine Landsat, seismic & topological databases
- Identify fault zones, surface erosion & predict earthquake magnitude

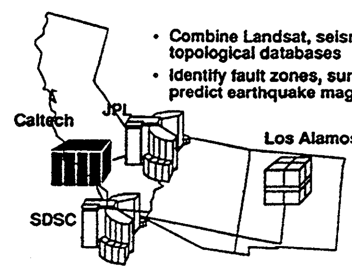

Global Climate Modeling

- Couple ocean and atmospheric models
- Used for climate studies
- Investigate nonlinear speedup

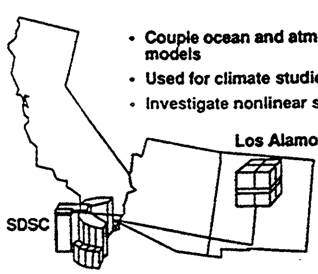

6.8 Router Requirements

Presented by Philip Almquist

I. Introduction

Outline of This Talk

- Introductory material
  - History
  - Goals
- Overview of the Router Requirements document
- Technical Issues
- Issues with "Requirements for ..." documents

History

- 8/79: IEN-109 (How to Build a Gateway)
- 5/86: RFC-985 (Requirements for Internet Gateways - draft)
- 6/87: RFC-1009 (Requirements for Internet Gateways)
- 10/89: RFC's 1122 and 1123 (Requirements for Internet Hosts)
- 2/90: Router Requirements WG formed

Goals

- Create a replacement for RFC-1009 which is more complete, more precise, and up-to-date
- Be consistent (both in style and content) with the Host Requirements RFC's
- Correct and clarify protocol standards
- Reconsider questionable requirements in RFC-1009
- Ensure that compliant routers will be interoperable
- Ensure that compliant routers provide features needed in an increasingly complex Internet
- Identify, and in some cases resolve, open questions
- Finish up on schedule!

II. Document Overview

Requirements

- "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", "MAY"
- "MUST IMPLEMENT", "SHOULD IMPLEMENT"
- Compliance should be objectively verifiable
- Requirements must be on routers, not their users
- Routers may be configurable to violate requirements, but must not do so by default

Contents

- Introduction
- Internet Architecture
- Link Layer
- Internet Layer -- Protocols
- Internet Layer -- Forwarding
- Transport Layer
- Application Layer -- Routing Protocols
- Application Layer -- Network Management Protocols
- Application Layer -- Miscellaneous Protocols
- Operations and Maintenance

III. Technical Issues

TOS

- *Type of Service in The Internet Protocol Suite* defines the details of how the IP TOS mechanism works:
  - TOS is an integer (not independently-settable bits)
  - Some values currently have defined semantics:
    - 1000 -- minimize delay
    - 0100 -- maximize throughput
    - 0010 -- maximize reliability
    - 0001 -- minimize monetary cost
    - 0000 -- normal service
  - Hosts and routers don't care about semantics
- The definition is (mostly) consistent with Host Requirements, OSPF, and Dual IS-IS

IP Precedence

- Prescribes the precedence value of some types of locally-originated datagrams
- Defines how Precedence should affect:
  - output queueing
  - discard policy (congestion)
  - use of Link Layer precedence facilities
- Defines an optional precedence cutoff mechanism

Forwarding Algorithm

- Describes in detail the steps involved in forwarding packets
- Conformance is defined as achieving the same result as the steps described

Route Choice and Route Leaking

- For a router in multiple routing domains, defines rules for:
  - *route choice*: deciding which routing domain to use to reach a particular destination
  - *route leaking*: passing routes from one routing domain to the another

Route Choice and Route Leaking II

- Rules are important because:
  - network managers may need precise control of these functions
  - consistency between routers is sometimes needed to prevent routing problems
- Rules are difficult to establish because:
  - different routing protocols use different criteria for picking the "best" route to a destination
- For the gory details:
  - *Ruminations on the Next Hop*
  - *Ruminations on Route Leaking*
  - *Some Thoughts on Multi-Domain Routing*

Network Management

- Requires SNMP (of course)
- Out-of-the box config includes read-only community "public"
- *IP Forwarding Table MIB* defines the Forwarding Table MIB

Operations and Maintenance

- A variety of requirements reflect the view that:
  - A router ought to be able to reboot and configure itself even if everything else is hosed
  - A network manager ought to be able to monitor and manage a router even if everything else is hosed

Traffic Filtering

- Defines:
  - Packet filtering (by source or destination address or port)
  - "Martian Filtering"
  - "Source Address Validation"

Unnumbered Serial Lines

- Describes two models in current use:
  1. "half-routers"
  2. node addresses ("router-id")
- We chose node addresses:
  - more flexible approach, but
  - less widely implemented
- Good topic for a separate RFC...

Routing Protocols

- None are required
- Implementing any routing protocol => MUST IMPLEMENT OSPF
- Implementing any routing protocol => MUST IMPLEMENT RIP
- Implementing EGP => MUST IMPLEMENT BGP

For Future Study

- Variable width subnets
  - Required by Router Requirements
  - Problems of ambiguous cases are the subject of a future draft by Robert Elz
  - Problems of administration ignored
- Multiple (sub)nets sharing a wire
  - Requirements level TBD
  - The ways in which this breaks the IP Architecture are the subject of a future draft by Lars Poulson
- Congestion control
  - SHOULD NOT send Source Quench
  - Study continues in IRTF

For Future Study II

- Router engineering
  - Buffer management
  - Resource allocation
  - Forwarding table lookup (Patricia, Cecilia, ...)
- SNMP Security
- IP Architecture changes for IPLPDN routing

Other Technical Issues

- Security
- Destination Unreachable codes
- Route Filtering
- TTL
- All-subnets broadcast ({ <net>,-1,-1 })
- Fragmentation algorithms

IV. Requirements Documents

Conformance Issues

- Is achieving conformance too difficult?
- Is the current scheme ("unconditional conformance", "conditional conformance", or "non-conformance" adequate?
- Is there a need for or a way to get independent assessments of compliance?
- Are checklists more useful than they are harmful?

Purpose of Requirements Documents

- Ensure protocol correctness?
- Ensure the box has the features needed in:
  - at least a few places in the Internet?
  - most parts of the Internet?
  - all imaginable topologies and administrative policies?
- Provide implementation hints?
- Inform consumers?
- Be useful in RFP's?

Requirements Document Architecture

- Current:

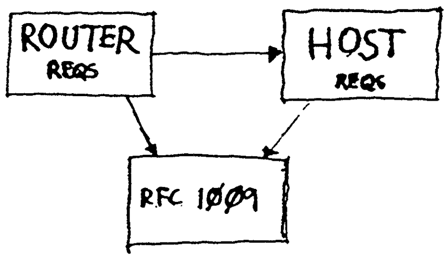

Requirements Document Architecture

- Future(?):

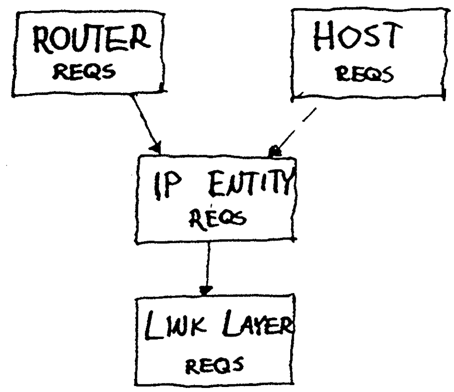

6.9 FARNET and the IETF

Presented by Susan Estrada/CERFnet

"FARNET and the IETF"

Susan Estrada
Vice President, FARNET
Co-Area Director, ORAD
Executive Director, CERFnet
And a Hell of a Good Person

FARNET
Federation of American Research Networks
100 Fifth Avenue
Waltham, MA 02154

1-800-72-FARNET
or ftp *nic.cerf.net*, cd *farnet*

Why a national association of networks?

To exchange information and ideas

To advocate for networking at the national level

To provide a single point of contact for information about midlevel networks

To support and develop the members of FARNET

Mission and strategy

FARNET's *mission* is to advance research and education through the use of data networking.

FARNET has adopted four primary *strategic focuses:*

1 Improvement of customer and network information services for Internet users
2 Advocacy
3 Programmatic ("user group") activities
4 Negotiation of group discounts for members

Membership profile

National service providers	3
Regional (multi-state) networks	7
Supercomputer center networks	3
State and provincial networks	15
University campuses	1
Local service providers	1

Serving more than 1,000 organizations, in every state and several foreign countries

History

Founded in 1986 by leaders of NSF midlevel networks

Quarterly meetings since 1987

Began collecting dues in 1988

Incorporated as non-profit in 1990

Hired first staff member in 1991

So, what does FARNET *do??*

FARNET meeting on K-12 networking activities (May 1991)

FARNET workshop and recommendations to NSF on backbone structure post-1992

FARNET meeting on "hardening the regionals" (Nov. 1991)

Liaison activities with IETF, SIGUCCS, CCIRN, IEPG, FNC, NTTF, etc.

FARNET *Gazette* and position papers

How FARNET works

Members elect Board of Directors

Board sets policy

Executive Director responsible for day-to-day operations

How it works (continued)

Seven committees provide guidance, organize projects:

Technical	Gene Hastings
Program	Paul Love
User Services	Martyne Hallgren
K-12	Jim Luckett
External Affairs	Richard Mandelbaum
Memb. & Bylaws	Jim Williams
Nominating	Glenn Ricart

YOU can join a committee or be on the mailing list.

FARNET and the IETF

√ Lots of overlap already!

√ Network operators need to work *smarter*, not *harder*!

√ Protocols, standards, tools, and procedures should help us do this!

And finally...

√ Good coordination between FARNET and IETF is essential!

FARNET networks are in production environment, with real-world economic and operational constraints

Plenty of notice is needed for major changes, such as OSPF and BGP

New standards and procedures have to be workable and affordable

FARNET recommendations to NSF re the NSFNET backbone after 1992

1. The multi-tier model for providing network services is valid and should be preserved.

2. Strong NSF support for top-level backbone services must continue.

3. As the agency responsible for NREN facilities coordination and deployment under the HPCC plan, NSF must assume a strong management role *vis a vis* the NREN.

4. The operation of the backbone network should be recompeted in the GFY92 timeframe with multiple awardees.

5. NSF should ensure that new technology is deployed in the backbone very carefully, to protect the quality of service to the end-user.

6. Midlevel networks should be able to exercise choice among vendors of top-level backbone services.

7. The backbone awardees should not be able to take advantage of their position to inhibit competition or to compete unfairly.

8. Provider accountability for performance should be ensured through the NSF award process.

9. NSF should take a leadership role in developing mechanisms to permit commercial traffic on the IINREN.

10 NSF should explore the feasibility of connecting midlevel networks using a FIX or CIX model as an alternative to a traditional backbone.

11. NSF should support the development of software tools for end-user applications and network management and operations.

12. NSF should issue a new solicitation aimed at midlevel and campus providers, with award criteria based on policy goals such as improving the ease of use of the network and leveraging private and non-Federal public funds.

6.10 SMDS Showcase and Applicable RFC's

Presented by Tate Jennings/Bellcore

SMDS is a connectionless, public, packet-switched data service whose operation and features are similar to those found in high-speed data networks. SMDS provides datagram packet transfer, exhibiting high throughput and low delay, and provides the transparent transport and delivery of up to 9188 octets of user information in a single transmission. Both individually and group-addressed (multicast) packets can be transferred. In addition to these LAN-like features, a set of addressing-related service features (source address validation, source and destination address screening) are provided to enable a subscriber or set of subscribers to create a logical private network, or closed user group, over SMDS.

INTEROP '91 in San Jose was an excellent example of technology teaming. Twenty-seven companies (including all the RBOCs) participated in the SMDS Showcase Demonstration. Twenty nodes in two buildings were connected via eighteen DS1 and two DS3 access links. Two switches trunked by a DS3 digital microwave radio link were employed. Additionally, an SMDS link comprised of DS1 access through seven switches also connected the demonstration with Telecom '91 in Geneva and delivered 64 octet packets round-trip in 230 milliseconds. The entire network was monitored by a prototype SNMP agent.

The implementation of Draft RFC 1209, The Transmission of IP Datagrams over the SMDS Service, formed the basis of the demonstration network. The network was configured as a single logical IP subnetwork (LIS), supporting both IP transmission and ARP. The SMDS group address was used to provide IP multicast capability, and the SNAP encapsulation method was followed. Six separate vendor implementations of the draft standard interoperated successfully over the three day period.

The SMDS Showcase also demonstrated the SMDS Customer Network Management service, showing interoperability between one SNMP agent developed by Bellcore and AT&T and five Network Management Systems (NMSs). The NMSs were from Cabletron, DEC, HP, SNMP Research, and Sun. The SNMP agent supported the applicable portions of MIB-II (RFC1213), DS1 (RFC1232), SMDS Subscription MIB (Bellcore.txt at venera.isi.edu), and the SIP MIB (draft-ietf-snmp-smdssipmib-01.txt).

Interop '91 SMDS Showcase Demonstration

Tate B. Jennings
Bellcore
tbj@sabre.bellcore.com

SMDS
Interop Announcements
Demonstration/Standards
Network
SMDS CNM
SMDS CNM Demo at Interop '91
Conclusion

SMDS

- Public switched packet service
- High speed, low latency
- LAN-like features
- Metropolitan areas

Interop Announcements

- Breadth of support
- Service support (30 cities)
- Tariffs - (pre-service = $500/mo.)

Demonstration/Standards

- SMDS access (IEEE 802.6)
- IP over SMDS (RFC 1209)
- SNMP/SMDS (RFC 1213, 1232, Bellcore.txt and snmpsmdsip-0.1.txt)

Network

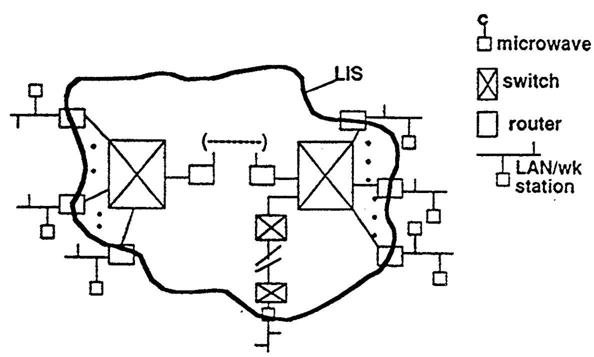

SMDS CNM

Supporting SMDS CNM in the BCC Network

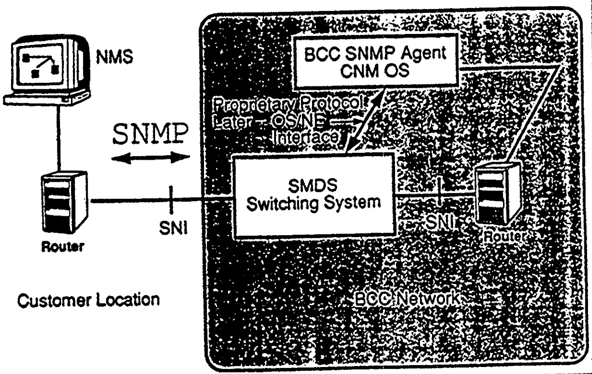

SMDS CNM DEMO AT INTEROP '91

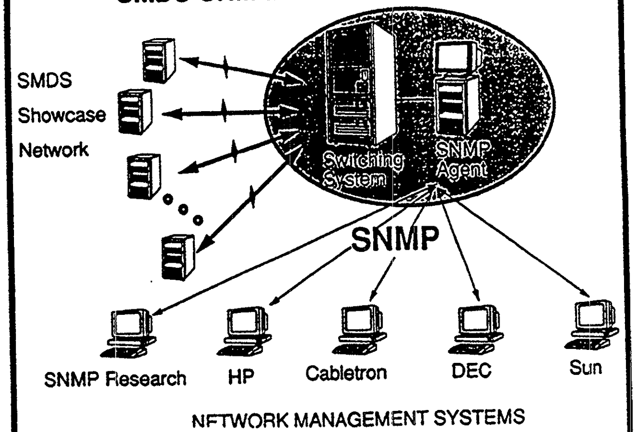

Conclusion

- Timely
- Real standards
- Real products
- Real service

6.11 NSFnet Update

Presented by Mark Knopper/Merit and Jordan Becker/ANS

ANSNET/NSFNET Backbone: Activities and Status

Mark Knopper
Jordan Becker
November 18-21, 1991
IETF/FARNET

- Traffic statistics
- Achieving Stability:
  - T3 backbone improvements
  - T3 backbone remaining problems
  - T1 backbone congestion
  - "Stability period" and Cutover
- Merit/ANS NOC/IE restructuring
- Policy routing database development
- Phase III network planning

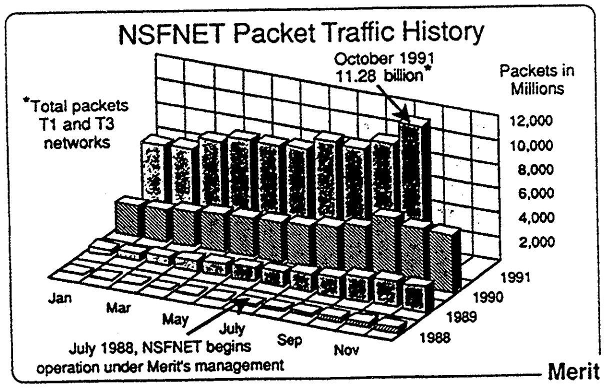

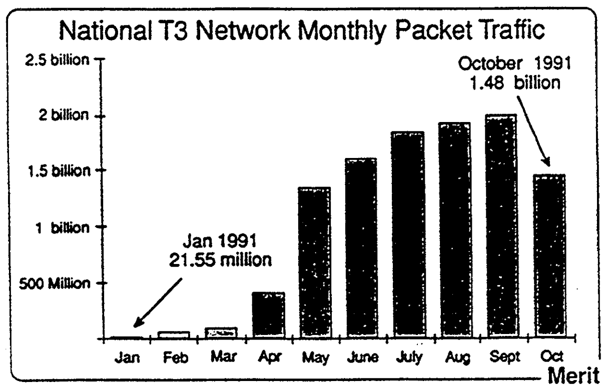

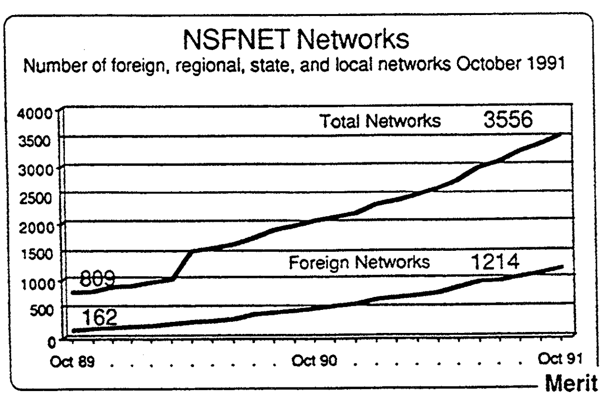

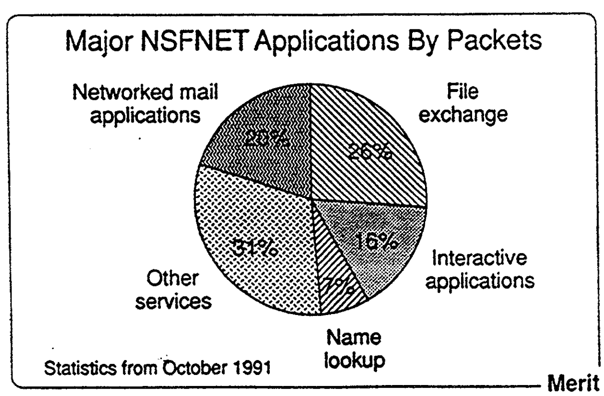

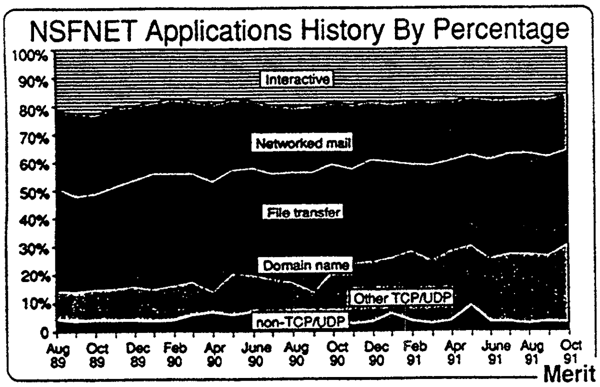
NSFNET Applications History By Percentage
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0
Interactive
Networked mail
File transfer
Domain name
Other TCP/UDP
non-TCP/UDP
Aug 89
Oct 89
Dec 89
Feb 90
Apr 90
June 90
Aug 90
Oct 90
Dec 90
Feb 91
Apr 91
June 91
Aug 91
Oct 91
Merit

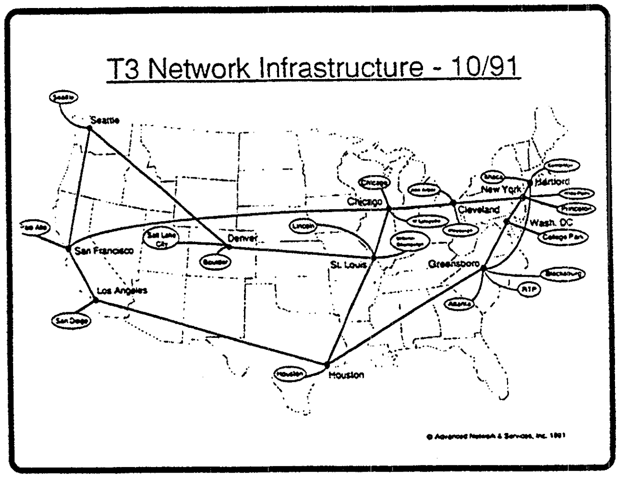
T3 Network Infrastructure - 10/91
Seattle
San Francisco
Los Angeles
Denver
Chicago
Cleveland
New York
Hartford
Wash. DC
St. Louis
Greensboro
Houston

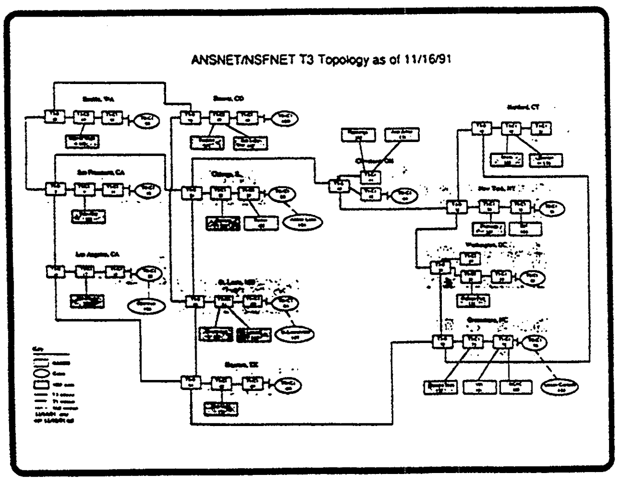
ANSNET/NSFNET T3 Topology as of 11/16/91

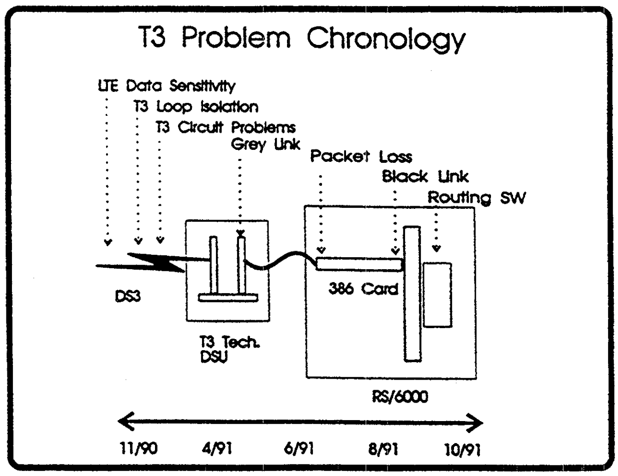
T3 Problem Chronology
LTE Data Sensitivity
T3 Loop Isolation
T3 Circuit Problems
Grey Link
Packet Loss
Black Link
Routing SW
DS3
T3 Tech. DSU
386 Card
RS/6000
11/90
4/91
6/91
8/91
10/91

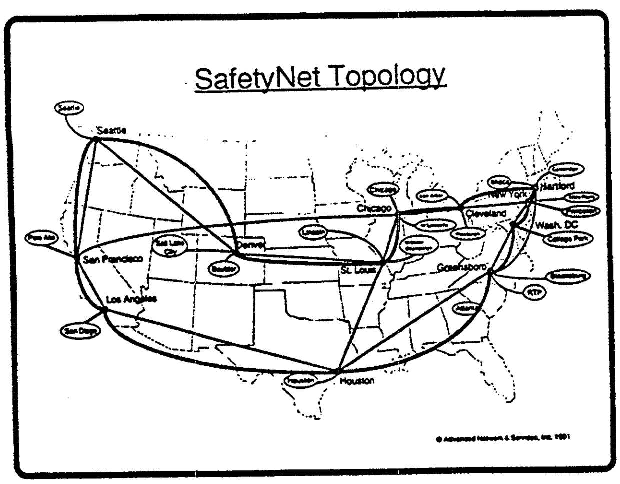
SafetyNet Topology
Seattle
San Francisco
Los Angeles
Denver
Chicago
Cleveland
New York
Hartford
Wash. DC
St. Louis
Greensboro
Houston

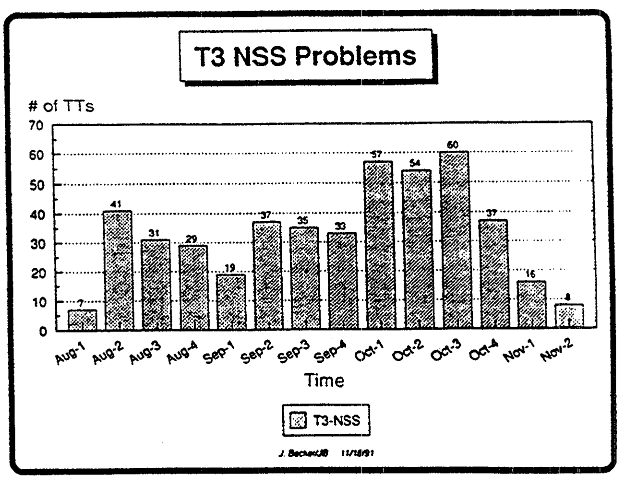
T3 NSS Problems
# of TTs
70
60
50
40
30
20
10
0
7
41
31
29
19
37
35
33
57
54
60
37
16
8
Aug-1
Aug-2
Aug-3
Aug-4
Sep-1
Sep-2
Sep-3
Sep-4
Oct-1
Oct-2
Oct-3
Oct-4
Nov-1
Nov-2
Time
T3-NSS

T3 Network Problem Status

- o T3 Problems Addressed
  - Grey Link Problem, Packet Loss Problems Fixed
  - Black Link Problem Work Around Applied
  - Multiple T1/T3 Interconnect Gateways Enabled
  - Core Network "Safety Net" Installation in Progress
  - Eliminated Network Unreachable Messages Internal to T3 Net.
- o Internal BGP Message Corruption Problem Unresolved
- o T960 T1/Ethernet Interface Problems Fixed
  - Route Caching, Ethernet Freeze, Sticky Ethernet
  - Route Download Hang
- o Configuration Control and Administration
  - Improved Control for System Software, Drivers, Microcode
  - Increase ANS & IBM Staffing at in Ann Arbor Supporting Merit

T1 Network Problem Status

- Congestion Related Packet Loss
  - Dual E-PSP Configuration Deployed
- EGP Loss Due to 5 Fragment Message Ethernet Truncation
  - Fixed in Build 270
- Intermittent Kernel Crashes
  - Fixed in Build 270
- Large IS-IS Messages Dropped Due to Token Ring Buffer Size
  - Fixed in Build 277

Merit/ANS NOC/IE restructuring

- Goal: solve inefficient problem resolution due to NOC staffing issues and IE on-call "burnout".
- New structure

NOC: 1. First level technical support
2. Shift management

- NOC function to become more formalized. Emphasis on reporting, escalation, tracking, and procedure execution.

New group: National Network Attack Force

- Managed by IE manager, but group works mainly in NOC.
- 6 people, 7/24 coverage
- Group to work on problem resolution process, tools, training, testing.

IE:
- Time freed for true Engineering activities.
- Third level escalation.

6.12 NIC Services

Presented by Scott Williamson/Network Solutions

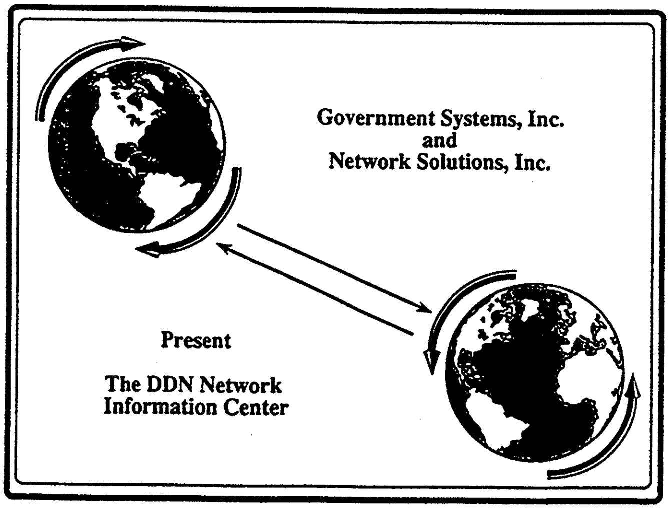

INTRODUCTION TO THE "NEW" NIC

WHOIS:

Government Systems, Inc. and Network Solutions, Inc.

Address:
Network Information Center
4200 Park Meadow Drive
Chantilly, VA 22020

Phone:
1 (800) 365-DNIC
1 (703) 802-4535

FAX:
1 (703) 802-8376

IP Address: 192.112.36.5

Hostname: NIC.DDN.MIL

Hardware: SunOS 470 running UNIX

NIC SERVICES

- General User Assistance: Help Desk Services
- MILNET TAC and non-TAC User Registration
- MILNET Security Matters: Reporting, Coordinating
- Host, Domain, and IP Network Registration
- Online Document Source Via FTP
  - RFCs, FYIs, TACNEWS
  - Protocols Information
  - General Information for Network Users (Netinfo)
- SERVICE: The NIC's Automatic Mail Server
- WHOIS Database: A Network "White Pages"
- NIC QUERY: A General Reference Lookup Service

REGISTRATION SERVICES

★ In addition to registering users in the WHOIS database,the NIC registers the following network entities:

- Domains
- IP Network Addresses
- Autonomous System Numbers (ASNs)
- Hosts
- TACs/Mini-TACs
- Gateways and PSNs

★ Registration software must be able to handle many variations of the standard registration data. Registration templates are submitted by applicants from all over the world.

★ Under normal operating system conditions and if all the required information is present on the template when it is received, processing a registration takes 8 working days from the date of receipt.

★ Processing and issuing TAC Cards and maintaining the "Hotlist" are also the responsibility of the REGISTRAR and HOSTMASTER staff members.

The Congestion Problem

- At the outset of the new NIC's operations, all Internet traffic was chanelled directly through the MILNET into a single 56 Kb line to the NIC.
- Because of the traffic bottleneck, performance was degraded for both Internet and MILNET users, as both user communities had to compete for the same limited connection resources.
- Because of the congestion problem, other seemingly unrelated problems arose, e.g., Root Domain Server behaved atypically due to instability of connections.
- The WHOIS software, already burdened by INGRES server inadequacy, appeared to be even more slow and unreliable because of the increased congestion.

The Congestion Solution

- In cooperation with NSF and NASA, a new T1 connection has now been established between FIX East and the NIC.
- This new connection is available to Internet users and gives them a direct pipeline to the NIC, thus enhancing both connectivity and performance.

THE CONGESTION SOLUTION

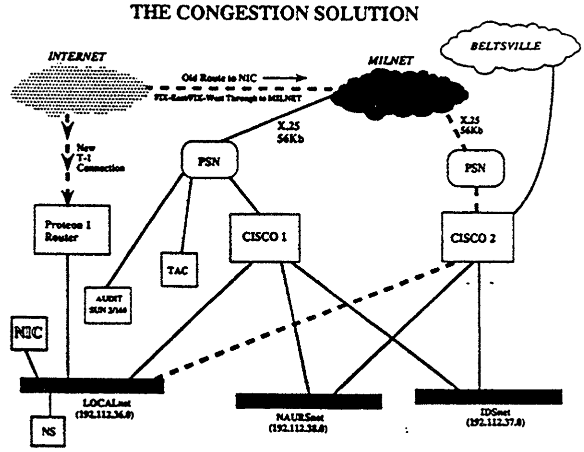

The WHOIS Problem

- The WHOIS software for the new NIC was originally implemented in ESQL/C under SUN's version of INGRES.
- INGRES implementations on other vendors' operating systems exceeded the NIC's projected contract requirements, but the SUN implementation proved to be unsatisfactory.
- The SUN INGRES server was inadequate for handling the NIC's heavy load and could not achieve an acceptable response time for database search requests.

The WHOIS Solution

- A two-phased approach was chosen to implement the WHOIS software enhancement effort.
- Phase I involved rewriting the LOOKUP portion of the software and its interface with INGRES. This Phase is now in beta testing and the software is in daily use on the Internet.
- Phase II involves the total re-engineering of the WHOIS software to make it totally independent of the INGRES server.
- Phase II is currently in testing. Deployment is expected in late November or early December.

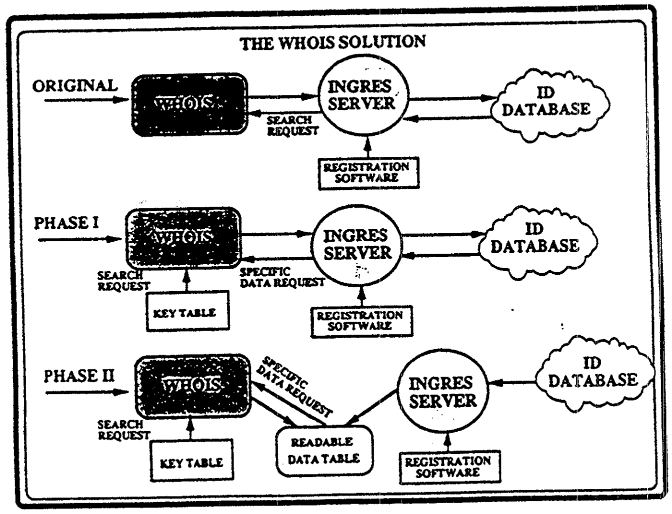

Currently Implemented Interface Enhancements

Here's what we've done so far:

- Moved all NIC Software to the UNIX Environment
- Developed a UNIX-based Command Interpreter that emulates the TOPS-20 Interpreter
- Updated Online Services Software to Prevent Abrupt Logouts from the NIC

Appendix A

Attendees

Vikas Aggarwal
(W)609-258-2403
(WF)609-258-2424
aggarwal@jvnc.net

Robert Aiken
(W)202-357-9717
(WF)202-357-0320
raiken@nsf.gov

Steve Alexander
(W)708-505-9100
(WF)708-505-9133
stevea@i88.isc.com

Claudio Allocchio
(W)+39 40 3758523
(WF)+39 40 226338
claudio.allocchio@elettra-ts.infn.it

Philip Almquist
(W)415-752-2427
almquist@jessica.stanford.edu

Harald Alvestrand
(W)+47 7 59 70 94
(WF)+47 7 53 25 86
harald.alvestrand@delab.sintef.no

Miriam Amos Nihart
(W)206-865-8549
(WF)206-865-8890
miriam@decwet.zso.dec.com

Mary Artibee
(W)415-335-1766
(WF)415-969-2314
artibee@sgi.com

Nagaraj Arunkumar
(W)408-264-5104
nak@3com.com

Karl Auerbach
(W)415-336-2090
(WF)415-969-9131
karl@eng.sun.com

Robert Austein
(W)617-942-0915
sra@asylum.sf.ca.us

William Babson
(W)301-921-8600
(WF)301-921-0161
bill%penril@uunet.UU.NET

Fred Baker
(W)805-963-9431
(WF)805-962-8499
fbaker@emerald.acc.com

Atul Bansal
(W)508-486-5355
(WF)508-486-7417
bansal@wile.nac.dec.com

James Barnes
(W)617-272-8140
(WF)617-273-5392
barnes@xylogics.com

William Barns
(W)703-883-6832
(WF)703-883-7142
barns@gateway.mitre.org

Charles Bazaar
(W)714-668-5416
(WF)714-668-6818
bazaar@emulex.com

Jordan Becker
(W)914-789-5315
(WF)914-789-5310
becker@nis.ans.net

James Beers
(W)607-255-2627
(WF)607-255-5771
beers@nr-tech.cit.cornell.edu

Art Berggreen
(W)805-963-9431
(WF)805-962-8499
art@acc.com

Param Bhamrah
(W)703-620-8990
(WF)703-620-0913
pbhamrah@nri.reston.va.us

William Biagi
(W)703-883-2806
(WF)703-848-8933
bbiagi@cos.com

Robert Blokzijl
(W)+31 20 5920413
(WF)+31 20 5925155
K13@nikhef.nl

Daniel Blum
(W)301-270-6405
(WF)301-270-2187

Larry Blunk
(W)313-936-3335
(WF)313-747-3745
ljb@merit.edu

David Bolen
(W)914-789-5327
(WF)914-789-5310
db3l@nis.ans.net

Thomas Boorman
(W)505-617-8503
tmb@lanl.gov

Nathaniel Borenstein
(W)201-829-4270
(WF)201-829-7019
nsb@thumper.bellcore.com

David Borman
(W)612-683-5571
(WF)612-683-5599
dab@cray.com

Steve Bostock
(W)408-473-8203
(WF)408-435-1706
steveb@novell.com

Robert Braden
(W)310-454-8852
(WF)310-823-6714
braden@isi.edu

Scott Bradner
(W)617-495-3864
(WF)617-495-0914
sob@harvard.edu

Craig Brenner
(W)408-974-5861
Brenner2@applelink.apple.com

David Bridgham
(W)617-942-0915
(WF)617-942-7846
dab@epilogue.com

Scott Brim
(W)607-255-5510
(WF)607-255-5771
swb@cornell.edu

Ronald Broersma
(W)619-553-2293
(WF)619-553-6811
ron@nosc.mil

Jack Brown
(W)602-533-2834
(WF)602-533-5676
jbrown@huachuca-emh8.army.mil

Gregory Bruell
(W)617-252-6300
(WF)617-252-6852
gob@shiva.com

Theodore Brunner
(W)201-829-4678
(WF)201-267-9298
tob@thumper.bellcore.com

Philip Budne
(W)617-252-6300
(WF)617-252-6852
phil@shiva.com

Jeffrey Buffum
(W)508-490-6822
(WF)508-624-7488
buffum@vos.stratus.com

John Burnett
(W)415-780-7768
(WF)415-364-0123
jlb@adaptive.com

Randy Butler
(W)217-244-8285
(WF)217-244-1987
rbutler@ncsa.uiuc.edu

Ross Callon
(W)508-486-5009
(WF)508-486-5279
callon@bigfut.enet.dec.com

Ken Carlberg
(W)703-448-0210
carlberg@cseic.saic.com

Lida Carrier
(W)408-974-8325
(WF)408-974-0892
lida@apple.com

Jeffrey Case
(W)615-573-1434
(WF)615-573-9197
case@cs.utk.edu

Stephen Casner
(W)310-822-1511
(WF)213-823-6714
casner@isi.edu

Peter Caswell
(W)508-898-3300
(WF)508-366-1356
pfc@pacer.soft.com

Vinton Cerf
(W)703-620-8990
(WF)703-620-0913
vcerf@nri.reston.va.

Yee-Hsiang Chang
(W)919-248-1169
(WF)919-248-1405
yhc@concert.net

A. Lyman Chapin
(W)617-873-3133
(WF)617-873-4086
lyman@bbn.com

Dean Cheng
(W)213-828-8811
dean@sunz.retix.com

Richard Cherry
(W)510-975-4520
rcherry@wc.novell.com

J. Noel Chiappa
(H)804-898-8183
(HF)804-898-7663
jnc@lcs.mit.edu

Cyrus Chow
(W)415-604-6843
(WF)415-604-4003
cchow@ames.arc.nasa.gov

Chi Chu
(W)703-448-0210
(WF)703-739-3323
chi@sparta.com

George Clapp
(W)708-248-5462
(WF)708-248-3775
clapp@ameris.center.il.ameritech.com

Jim Clifford
(W)505-667-2596
jrc@lanl.gov

James Codespote
(W)301-684-7302
(WF)301-850-7166
jpcodes@tycho.ncsc.mil

Richard Cogger
(W)607-255-5510
(WF)607-255-5771
rhx@cornell.cit.bitnet

Richard Colella
(W)301-975-3627
(WF)301-590-0932
colella@osi.ncsl.nist.gov

Rob Coltun
(W)301-240-2168
rcoltun@ni.umd.edu

James Conklin
(W)202-872-4200
(WF)202-872-4318
conklin@bitnic.educom.edu

Robert Cooney
(W)202-433-2448
(WF)202-433-5613
cooney@wnyose.nctsw.navy.mil

Curtis Cox
(W)202-433-4026
(WF)202-433-3188
ccox@wnyose.nctsw.navy.mil

Tracy Cox
(W)908-758-2107
(WF)908-758-4177
tacox@sabre.bellcore.com

Stephen Coya
(W)703-620-8990
scoya@nri.reston.va.us

Kenneth Crepea
(W)407-273-8433
(WF)407-273-0576
crepea@cisco.com

Mark Crispin
(W)206-543-5762
(WF)206-543-3909
mrc@cac.washington.edu

Stephen Crocker
(W)301-854-6889
(WF)301-854-5363
crocker@tis.com

Dave Cullerot
(W)603-332-9400
(WF)603-335-3467
cullerot@ctron.com

John Curran
(W)617-873-3400
jcurran@bbn.com

John Damiano
(W)919-848-9129

Megan Davies
(W)703-620-8990
mdavies@nri.reston.va.us

James Davin
(W)617-253-6020
(WF)617-258-8682
jrd@ptt.lcs.mit.edu

Osmund de Souza
(W)908-949-1393
(WF)908-949-1726
osmund.desouza@att.com

Steve Deering
(W)415-812-4839
(WF)415-812-4471
deering@xerox.com

Marty Del Vecchio
(W)617-252-6300
(WF)617-252-6852
marty@shiva.com

Jim DeMarco
(W)617-246-0900
(WF)617-245-7943
jdemarco@ftp.com

John Demco
(W)604-822-6724
(WF)604-822-5485
demco@cs.ubc.ca

Barbara Denny
(W)415-859-2812
(WF)415-859-4812
denny@sri.com

Peter Deutsch
(W)514-398-3709
(WF)514-398-6876
peterd@cc.mcgill.ca

Peter DiCamillo
(W)401-863-7582
(WF)401-863-7329
cmsmaint@brownvm.brown.edu

Tim Dixon
(W)+31 20 592 5078
(WF)+31 20 592 5043
dixon@nikhef.nl

Kurt Dobbins
(W)603-332-9400
(WF)603-335-5917
dobbins@ctron.com

Daniel Duchamp
(W)212-854-8115
(WF)212-666-0140
duchamp@liberty.cs.columbia.edu

Steve Dusse
(W)415-595-8782
(WF)415-595-1873
spock@rsa.com

Tom Easterday
(W)614-292-4843
(WF)614-292-7081
tom@cic.net

Alan Emtage
(W)514-398-8117
(WF)514-398-6876
bajan@cc.mcgill.ca

Robert Enger
(W)914-789-5355
(WF)914-789-5310
enger@ans.net

Kent England
(W)617-873-3617
(WF)617-873-3776
kwe@bbn.com

Hans Eriksson
(W)+46 8 752 1528
(WF)+46 8 751 7230
hans@sics.se

Michael Erlinger
(W)714-870-2117
(WF)714-680-0279
mike@lexcel.com

Jeff Erwin
(W)503-645-5400

Susan Estrada
(W)619-534-5067
(WF)619-534-5167
estradas@sdsc.edu

Erik Fair
(W)408-996-1010
(WF)408-974-3103
fair@apple.com

Bill Fardy
(W)603-332-9400
(WF)603-335-5917
fardy@ctron.com

Dino Farinacci
(W)415-688-4696
(WF)415-688-8282
dino@cisco.com

Jeff Fitzgerald
(W)703-342-6700
(WF)703-342-5961
jjf@fibercom.com

Peter Ford
(W)505-665-0558
(WF)505-665-4939
peter@lanl.gov

Craig Fox
(W)612-424-4888
foxcj@network.com

Barbara Fraser
(W)412-268-7090
(WF)412-268-6989
byf@cert.sei.cmu.edu

Ned Freed
(W)714-624-7907
(WF)714-621-5319
ned@innosoft.com

Karen Frisa
(W)412-268-8576
(WF)412-268-4987
karen.frisa@andrew.cmu.edu

Vince Fuller
(W)415-723-6860
(WF)415-723-0010
vaf@stanford.edu

Darrell Furlong
(W)508-460-4646
(WF)508-481-9772

Shawn Gallagher
gallagher@quiver.enet.dec.com

James Galvin
(W)301-854-6889
(WF)301-854-5363
galvin@tis.com

John Garrett
(W)908-580-4719
(WF)908-580-4721
jwg@edsel.att.com

L. Dain Gary
(W)412-268-7090
(WF)412-268-6989
ldg@cert.sei.cmu.edu

John Gawf
(W)303-444-9532
(WF)303-444-9595
gawf@compatible.com

Jisoo Geiter
(W)703-883-6571
(WF)703-883-7142
geiter@gateway.mitre.org

Tony Genovese
(W)415-423-2471
(WF)415-422-0435
genovese@es.net

Farrell Gerbode
(W)713-527-4988
(WF)713-527-6099
farrell@rice.edu

Jim Ghadbane
(W)613-591-3600
(WF)613-591-3680
jimgh@newbridge.com

Joseph Godsil
(W)217-244-7715
(WF)217-244-1987
jgodsil@ncsa.uiuc.edu

Herve Goguely
(W)408-436-7600
(WF)408-436-6060
lircorp1!roger@uunet.uu.net

John Gong
(W)415-641-8308
(WF)415-506-7226
jgong@us.oracle.com

John Greer
(W)505-665-4496
(WF)505-665-4361
jyg@lanl.gov

Phillip Gross
(W)914-789-5341
(WF)914-789-5310
pgross@nri.reston.va.us

Olafur Gudmundsson
(W)301-405-2767
(WF)301-405-6707
ogud@cs.umd.edu

Chris Gunner
(W)+44 734 203516
(WF)+44 734 876384
gunner@osicwg.enet.dec.com

John Hagan
(W)215-898-9192
(WF)215-898-9348
hagan@dccs.upenn.edu

Robert Hagens
(W)608-262-5776
(WF)608-262-9777
hagens@cs.wisc.edu

Jack Hahn
(W)301-982-4600
(WF)301-982-4605
hahn@sura.net

Tony Hain
(W)510-422-4200
(WF)510-422-0435
hain@es.net

Neil Haller
(W)201-829-4478
(WF)201-829-2504
nmh@bellcore.com

Martyne Hallgren
(W)607-255-5510
(WF)607-255-5771
martyne@nr-tech.cit.cornell.edu

Alf Hansen
(W)+47 7 592982
(WF)+47 7 532586
Alf.Hansen@delab.sintef.no

Susan Hares
(W)313-936-2095
(WF)313-747-3745
skh@merit.edu

Eugene Hastings
(W)412-268-4960
(WF)412-268-5832
hastings@psc.edu

Kenneth Hayward
(W)613-763-4042
(WF)613-763-8581
crm57d@bnr.ca

Frank Heath
(W)714-833-6864
(WF)714-833-6863
heath@cmc.com

Sarah Hebert
(W)505-667-3310
seh@noc-gw.lanl.gov

Christine Hemrick
(W)415-688-8164
(WF)415-688-4555
hemrick@cisco.com

Ittai Hershman
(W)914-789-5337
(WF)914-789-5310
ittai@nis.ans.net

Robert Hinden
(W)617-873-3757
(WF)617-491-1496
hinden@bbn.com

Russ Hobby
(W)916-752-0236
(WF)916-752-9154
rdhobby@ucdavis.edu

Randy Hoebelheinrich
(W)505-667-5421
(WF)505-665-4361
jrc@lanl.gov

Ellen Hoffman
(W)313-936-3335
(WF)313-747-3745
esh@merit.edu

J. Paul Holbrook
(W)313-998-7680
(WF)313-998-6105
holbrook@cic.net

Greg Hollingsworth
(W)301-953-6065
(WF)301-953-1093
gregh@mailer.jhuapl.edu

Peter Honeyman
(W)313-763-4413
(WF)313-763-4434
honey@citi.umich.edu

Eric Hood
(W)206-562-3000
(WF)206-562-4822
ehood@nwnet.net

Alton Hoover
(W)914-789-5341
(WF)914-789-5310
hoover@nis.ans.net

Kathleen Huber
(W)617-873-2520
(WF)617-873-4086
khuber@bbn.com

Christian Huitema
(W)+33-93657715
(WF)+33-93657766
christian.huitema@sophia.inria.fr

Geoff Huston
(W)+61-6-249-3385
(WF)+61-6-247-3425
g.huston@aarnet.edu.au

Robin Iddon
(W)+31 554 9424
robini@spider.co.uk

Lenore Jackson
(W)301-286-7251
(WF)301-286-5251
jackson@dftnic.gsfc.nasa.gov

William Jackson
(W)619-553-2291
(WF)619-533-6811
jackson@manta.nosc.mil

Ole Jacobsen
(W)800-INTEROP
(WF)415-949-1779
ole@csli.stanford.edu

Van Jacobson
(W)510-486-6411
van@ee.lbl.gov

Ronald Jacoby
(W)415-335-1371
(WF)415-969-2314
rj@sgi.com

Bob Jeckell
(W)408-764-6049
robert_jeckell@nso.3com.com

Borka Jerman-Blazic
(W)+38 61 159 199
(WF)+38 61 161 029
jerman-blazic@ijs.ac.mail.yu

Dale Johnson
(W)313-936-3335
(WF)313-747-3745
dsj@merit.edu

James Jokl
(W)804-924-0616
(WF)804-982-4715
jaj@virginia.edu

William Jolitz
(W)510-420-1406
(WF)510-420-0174
william@okeeffe.cs.berkeley.edu

Kevin Jordan
(W)612-482-6835
(WF)612-482-4455
kej@udev.cdc.com

Dan Jordt
(W)206-562-3000
(WF)206-562-4822
danj@nwnet.net

Satish Joshi
(W)408-764-1578
(WF)408-988-5525
sjoshi@synoptics.com

Jean-Michael Jouanigot
(W)+41 22 767 44 17
(WF)+41 22 767 71 55
jimi@cernvax.cern.ch

Scott Kaplan
(W)415-543-9001
(WF)415-543-9002
scott@ftp.com

Mitchell Kapor
(W)617-864-1550
(WF)617-864-0866
mkapor@eff.org

Michael Karels
(W)415-642-4948
(WF)415-642-5775
karels@cs.berkeley.edu

Frank Kastenholz
(W)508-435-2000
(WF)508-435-7504
kasten@europa.clearpoint.com

Neil Katin
(W)415-336-3718
(WF)415-961-0992
katin@eng.sun.com

David Kaufman
(W)508-898-2800
(WF)508-366-8901
dek@proteon.com

Charles Kaufman
(W)508-486-7329
(WF)508-486-5279
kaufman@dsmail.enet.dec.com

Manu Kaycee
(W)603-332-9400
(WF)603-335-5917
kaycee@ctron.com

Stephen Kent
(W)617-873-3988
(WF)617-873-4086
kent@bbn.com

Mark Kepke
(W)303-229-4201
(WF)303-229-2692
mak@cnd.hp.com

Tom Kessler
(W)415-336-3145
(WF)415-336-9131
kessler@sun.com

Darren Kinley
(W)514-340-5768
(WF)514-340-5777
kinley@crim.ca

Peter Kirstein
(W)+44 71 380 7286
(WF)+44 71 387 1397
kirstein@cs.ucl.ac.uk

John Klensin
(W)617-253-1355
(WF)617-491-6266
klensin@infoods.mit.edu

Yoav Kluger
(W)+97 24 31 3313
(WF)+97 24 51 6266
ykluger@fibhaifa.com

Holly Knight
(W)408-974-5110
(WF)408-974-0892
holly@apple.com

Mark Knopper
(W)313-763-6061
(WF)313-747-3745
mak@merit.edu

Jim Knowles
(W)415-964-9900
jknowles@trident.arc.nasa.gov

Stev Knowles
(W)617-246-0900
(WF)617-245-7943
stev@ftp.com

David Korns
(W)512-343-4657
David_Korns@auo.mts.dec.com

Deidre Kostick
(W)908-758-2642
(WF)908-758-4177
dck2@sabre.bellcore.com

Edward Krol
(W)217-333-7886
(WF)217-244-7089
e-krol@uiuc.edu

Cheryl Krupczak
cheryl@cc.gatech.edu

Louise Laier
(W)408-974-6034
(WF)408-257-7982
laierl@applelink.apple.com

Carol Lambert
(W)607-255-8966
(WF)607-254-5222
cjlx@cornella.cit.cornell.edu

Dale Land
(W)505-667-9923
land@lanl.gov

Ruth Lang
(W)415-859-5608
(WF)415-859-6028
rlang@nisc.sri.com

Ron Lau
(W)408-764-1444
(WF)408-988-5525
rlau@synoptics.com

Vincent Lau
(W)415-336-1069
vincent.lau@eng.sun.com

Kenneth Laube
(W)617-873-3176
(WF)617-873-4086
laube@bbn.com

Anthony Lauck
(W)508-486-7644
(WF)508-486-5279
lauck@tl.enet.dec.com

Joseph Lawrence
(W)908-758-4146
jcl@sabre.bellcore.com

Walter Lazear
(W)703-883-6515
(WF)703-883-7142
lazear@gateway.mitre.org

Eliot Lear
(W)415-335-1056
lear@sgi.com

Gordon Lee
gordon@ftp.com

Debra Legare
(W)703-620-8990
(WF)703-620-0913
dlegare@nri.reston.va.us

Barry Leiner
(W)415-960-7540
(WF)415-960-7500
bleiner@ads.com

John Lekashman
(W)415-604-4359
(WF)415-604-4377
lekash@nas.nasa.gov

Mark Lewis
(W)408-745-3232
(WF)408-745-3810
mlewis@telebit.com

Tony Li
(W)415-688-8186
(WF)415-688-8282
tli@cisco.com

Peter Liebscher
(W)301-982-4600
(WF)301-982-4605
plieb@sura.net

John Linn
(W)508-486-6588
(WF)508-486-6014
linn@zendia.enet.dec.com

Joshua Littlefield
(W)617-494-1999
(WF)617-494-5167
josh@cayman.com

Jack Liu
(W)603-881-2590
(WF)603-881-0120
liu@koala.enet.dec.com

Brian Lloyd
(W)916-676-1147
brian@ray.lloyd.com

Daniel Long
(W)617-873-2766
(WF)617-873-5620
long@nic.near.net

E. Paul Love
(W)619-534-5043
(WF)619-534-5152
loveep@sdsc.edu

Ron Mackey
(W)708-665-4693
(WF)708-665-4706
rem@dsiinc.com

Carl Malamud
carl@malamud.com

Andrew Malis
(W)617-873-3419
(WF)617-873-4086
malis@bbn.com

Gary Malkin
(W)617-246-0900
(WF)617-245-7943
gmalkin@ftp.com

Allison Mankin
(W)703-883-7907
(WF)703-883-7142
mankin@gateway.mitre.org

Mike Marcinkevicz
(W)714-773-3009
(WF)714-773-3990
mdm@csu.net

April Marine
(W)415-859-5318
(WF)415-859-6028
april@nisc.sri.com

Tony Mason
(W)412-338-4400
(WF)412-338-4404
mason@transarc.com

Keith McCloghrie
(W)415-966-7934
(WF)415-960-3738
kzm@hls.com

Ellen McDermott
(W)617-621-7266
emcd@osf.org

David McDysan
(W)214-918-5178
(WF)214-918-6019
0002806303@mcimail.com

Evan McGinnis
(W)408-764-6064
(WF)408-764-6135
bem@3com.com

Glenn McGregor
(W)313-936-3000
(WF)313-747-3745
ghm@merit.edu

Donna McMaster
(W)408-764-1206
(WF)408-764-1299
mcmaster@synoptics.com

Jim McQuaid
(W)919-941-5730
(WF)919-941-5751

Milo Medin
(W)415-604-6440
(WF)415-604-0063
medin@nsipo.nasa.gov

Bill Melohn
(W)408-492-0900
(WF)408-492-0909
melohn@auspex.com

Greg Merrell
(W)408-496-3107
merrell@greg.enet.dec.com

April Merrill
(W)301-684-7302
(WF)301-850-7166
abmerri@tycho.ncsc.mil

Donald Merritt
(W)410-278-6808
(WF)301-278-5075
don@brl.mil

Glenn Michel
(W)505-667-1845
(WF)505-665-4361
gym@lanl.gov

Paul Milazzo
(W)617-873-4325
(WF)617-873-5620
milazzo@bbn.com

David Minnich
(W)703-342-6700
(WF)703-342-5961
dwm@fibercom.com

Greg Minshall
(W)510-975-4507
(WF)510-975-1238
minshall@wc.novell.com

Paul Mockapetris
(W)703-696-2262
(WF)703-696-2202
pvm@darpa.mil

Dave Monachello
(W)201-438-2400
(WF)201-438-7767
dave@pluto.dss.com

Lynn Monsanto
(W)415-336-2946
(WF)415-336-3156
monsanto@sun.com

Keith Moore
(W)615-546-4566
(WF)615-974-4404
moore@cs.utk.edu

Robert Morgan
(W)415-723-9711
(WF)415-723-0010
morgan@jessica.stanford.edu

Dennis Morris
(W)703-487-3329
(WF)703-487-3190
morrisd@imo-uvax.dca.mil

Donald Morris
(W)303-497-1282
(WF)303-497-1137
morris@ucar.edu

Dean Morris
(W)+44 0734 868711
(WF)+44 0734 867969
morris@marvin.dec.com

John Morrison
(W)505-667-3310
(WF)505-665-6361
jfm@cutthroat.lanl.gov

John Moy
(W)508-898-2800
(WF)508-366-8901
jmoy@proteon.com

Frederick Mueller
(W)408-745-3028
(WF)408-734-3333
fmueller@telebit.com

Chris Myers
(W)314-935-7390
(WF)314-935-7302
chris@wugate.wustl.edu

Mark Needleman
(W)415-987-0530
(WF)415-839-3573
mhn@stubbs.ucop.edu

Michael Newell
(W)202-453-8752
(WF)202-755-2637
mnewell@nhqvax.hg.nasa.gov

Daniel Newman
(W)714-624-7907
(WF)714-621-5319
dan@innosoft.com

Andy Nicholson
(W)612-683-5473
(WF)612-683-5599
droid@cray.com

Chandy Nilakantan
(W)408-764-5213
(WF)408-764-5003
csn@3com.com

Rebecca Nitzan
(W)510-423-4055
(WF)510-422-0435
nitzan@es.net

David O'Leary
(W)301-982-3214
(WF)301-982-4605
oleary@sura.net

Sean O'Malley
(W)602-621-6613
(WF)602-621-4246
sean@cs.utexas.edu

Ari Ollikainen
(W)510-423-5962
(WF)510-422-0435
ari@es.net

Zbigniew Opalka
(W)508-937-1600
(WF)508-937-0445
zopalka@roswell.spartacus.com

J. Bradford Parker
(W)617-494-1999
(WF)617-494-5167
brad@cayman.com

Brad Passwaters
(W)301-982-4600
(WF)301-982-4605
bjp@sura.net

Michael Patton
(W)617-253-6061
(WF)617-258-8682
map@lcs.mit.edu

David Perkins
(W)408-764-1516
(WF)408-988-5525
dperkins@synoptics.com

Radia Perlman
(W)508-486-7648
(WF)508-486-5279
perlman@radia.enet.dec.com

Marsha Perrott
(W)412-268-7871
(WF)412-268-7875
mlp+@andrew.cmu.edu

David Piscitello
(W)908-758-2286
(WF)908-758-4177
dave@sabre.bellcore.com

Jon Postel
(W)213-822-1511
(WF)213-823-6714
postel@isi.edu

Venkat Prasad
(W)408-764-5226
(WF)408-764-5003
vsp@3com.com

Robert Purvy
(W)415-506-2972
bpurvy@us.oracle.com

Thomas Pusateri
pusateri@cs.duke.edu

Joe Ragland
(W)919-248-1404
(WF)919-248-1405
jrr@concert.net

Christopher Ranch
(W)408-473-8667
cranch@novell.com

Yakov Rekhter
(W)914-945-3896
yakov@watson.ibm.com

Joyce Reynolds
(W)310-822-1511
(WF)310-823-6714
jkrey@isi.edu

Anil Rijsinghani
(W)508-952-3520
anil@levers.enet.dec.com

Michael Ritter
(W)408-862-8088
(WF)408-257-7982
mwritter@applelink.apple.com

Ron Roberts
(W)415-723-3352
(WF)415-723-0010
roberts@jessica.stanford.edu

Michael Roberts
(W)202-872-4200
(WF)202-872-4318
roberts@educom.edu

Manoel Rodrigues
(W)908-949-4655
(WF)908-949-1726
manoel.rodrigues@att.com

Philippe Roger
(W)408-436-7600
(WF)408-436-6060
lircorp1!roger@uunet.uu.net

Marshall Rose
(W)415-968-1052
(WF)415-968-2510
mrose@dbc.mtview.ca.us

Karen Roubicek
(W)617-329-3350
(WF)617-326-5484
roubicek@faxon.com

Allan Rubens
(W)313-936-3335
(WF)313-747-3745

Gregory Ruth
(W)617-873-3150
(WF)617-873-8202
gruth@bbn.com

Harri Salminen
(W)+358-0-457-2288
(WF)+358-0-457-2302
hks@funet.fi

Timothy Salo
(W)612-626-0347
(WF)612-624-6550
tjs@msc.edu

Tom Sandoski
(W)919-248-1104
(WF)919-248-1405
tom@concert.net

Eric Sandoval
(W)505-857-7224
sandoval@hotair.enet.dec.com

Sharad Sanghi
(W)914-789-5300
(WF)914-789-5310
sharad@ans.net

Jonathan Saperia
(W)508-496-8333
saperia@tcpjon.enet.dec.com

Miguel Sasson
(W)617-272-8140
sasson@xylogics.com

Mark Schaefer
(W)408-720-6804
(WF)408-720-1337
schaefer@davidsys.com

Jeffrey Schiller
(W)617-253-8400
(WF)617-258-8736
jis@mit.edu

Tim Seaver
(W)919-248-1973
(WF)919-248-1455
tas@concert.net

John Seligson
(W)408-922-0100
(WF)408-433-9287
johns@ultra.com

Harvey Shapiro
(W)202-433-4048
(WF)202-433-5613
shapiro@wnyose.nctsw.navy.mil

Daisy Shen
(W)914-945-3632
(WF)914-945-4439
daisy@watson.ibm.com

Stephen Shew
(W)613-763-2462
(WF)613-763-2626
sdshew@bnr.ca

Robert Shirey
(W)703-883-7210
(WF)703-883-1397
shirey@mitre.org

Keld Simonsen
(W)+45 33 130023
(WF)+45 33 911828
keld@dkuug.dk

William Simpson
(W)313-543-3269
Bill_Simpson@um.cc.umich.edu

Ursula Sinkewicz
(W)508-486-5604
(WF)508-486-7417

Dana Sitzler
(W)313-936-3335
(WF)313-747-3745
dds@merit.edu

Frank Slaughter
(W)617-252-6300
(WF)617-252-6852
fgs@shiva.com

Timon Sloane
(W)408-727-4111
(WF)408-727-4410
peernet!timon@uunet.uu.net

Patricia Smith
(W)313-936-3335
(WF)313-747-3745
psmith@merit.edu

Richard Smith
(W)201-438-2400
(WF)201-438-7767
smiddy@pluto.dss.com

Eric Smith
(W)408-745-3129
(WF)408-745-3810
eric@telebit.com

Frank Solensky
(W)508-435-2000
(WF)508-435-7386
solensky@clearpoint.com

Evan Solley
(W)503-292-1374
(WF)503-226-6543
solley@applelink.apple.com

Mike Spengler
(W)612-626-1757
(WF)612-624-6550
mks@msc.edu

Ravi Srinivasan
(W)510-745-6061
ravi@eng.vitalink.com

Arthur St. George
(W)505-277-8046
(WF)505-277-8101
stgeorge@bootes.unm.edu

Michael St. Johns
(W)301-688-6741
stjohns@umd5.umd.edu

Martha Steenstrup
(W)617-873-3192
(WF)617-491-1496
msteenst@bbn.com

Brad Steinka
(W)617-551-1209
(WF)617-551-1015
brad@python.eng.microcom.com

Bob Stewart
(W)508-264-9900
(WF)508-264-9930
rlstewart@eng.xyplex.com

David S.A. Stine
(W)415-688-4681
(WF)415-688-8282
dstine@cisco.com

Bernhard Stockman
(W)+46-8-790-6519
boss@sunet.se

George Strawn
(W)202-357-9717
(WF)202-357-0320
gstrawn@nsf.gov

Roxanne Streeter
(W)415-604-4845
(WF)415-604-0063
streeter@nsipo.nasa.gov

Bruce Taber
(W)508-263-9929
(WF)508-263-8655
taber@interlan.com

Sven Tafvelin
(W)+46 31 721706
(WF)+46 31 721561
tafvelin@ce.chalmers.se

Osamu Takada
(W)+81 44 966 9111
(WF)+81 44 966 6853
takada@sdl.hitachi.co.jp

Iris Tal
(W)+97 23 5447361
(WF)+97 23 5447368
437-3580@mcimail.com

Sally Tarquinio
(W)703-883-7987
(WF)703-883-7142
sally@gateway.mitre.org

Kaj Tesink
(W)908-758-5254
(WF)908-530-6862
kaj@nvuxr.cc.bellcore.com

Mark Therieau
(W)617-551-1195
(WF)617-551-1015
markt@python.eng.microcom.com

Joanie Thompson
(W)415-604-4550
(WF)415-604-0063
joanie@nsipo.nasa.gov

Dean Throop
(W)919-549-8421
(WF)919-248-5942
throop@dg-rtp.dg.com

Claudio Topolcic
(W)703-620-8990
(WF)703-620-0913
topolcic@nri.reston.va.us

William Townsend
(W)617-272-8140
(WF)617-273-5392
townsend@xylogics.com

Mike Truskowski
(W)314-275-9174
(WF)314-542-3176
truskowski@cisco.com

Theodore Tso
(W)617-253-8091
(WF)617-258-8736
tytso@mit.edu

Paul Tsuchiya
(W)201-829-4484
(WF)201-829-2504
tsuchiya@thumper.bellcore.com

Maurice Turcotte
(W)404-781-3728
(WF)404-804-5190
dnmrt@interlan.com

Kannan Varadhan
(W)614-292-4137
(WF)614-292-7168
kannan@oar.net

Gregory Vaudreuil
(W)703-620-8990
(WF)703-620-0913
gvaudre@nri.reston.va.us

Andrew Veitch
(W)617-873-2588
aveitch@bbn.com

John Veizades
(W)408-974-2672
(WF)408-974-9083
veizades@apple.com

Sudhanshu Verma
(W)415-857-8196
(WF)415-852-3762
verma@hpspdla.hp.com

Lee Wade
(W)415-604-4789
(WF)415-604-0063
wade@nsipo.arc.nasa.gov

Steven Waldbusser
(W)412-268-6628
(WF)412-268-4987
waldbusser@andrew.cmu.edu

Yuan Wang
(W)408-370-4366
(WF)408-370-4222
natadm!ycw@uunet.uu.net

Carol Ward
(W)303-492-5860
(WF)303-492-4198
cward@spot.colorado.edu

David Wasley
(W)510-642-3478
(WF)510-642-5385
dlw@berkeley.edu

Scott Wasson
(W)508-264-9900
(WF)508-264-9930
sgw@sgw.xyplex.com

Chris Weider
(W)313-936-3335
(WF)312-747-3745
clw@merit.edu

Frank Welch
(W)602-869-4059
(WF)602-869-3509

Jonathan Wenocur
(W)617-252-6300
(WF)617-252-6852
jhw@shiva.com

Kathleen Wilde
(W)603-881-0387
(WF)508-486-5279
wilde@decvax.dec.com

Sally Wilkins
(W)505-667-1391
sfw@lanl.gov

Scott Williamson
(W)703-802-8405
(WF)703-802-8376
scottw@nic.ddn.mil

Jeremy Wilson
(W)503-645-5400

Preston Wilson
(W)708-505-9100
(WF)708-505-9133
preston@i88.isc.com

Walter Wimer
(W)412-286-6252
(WF)412-268-4987
walter.wimer@andrew.cmu.edu

Linda Winkler
(W)708-972-7236
(WF)708-972-5983
lwinkler@anl.gov

Cathy Wittbrodt
(W)415-466-4016
(WF)415-422-0435
cjw@nersc.gov

C. Philip Wood
(W)505-667-2598
(WF)505-665-4361
cpw@lanl.gov

Robert Woodburn
(W)703-734-9000
(WF)703-749-5997
woody@cseic.saic.com

Richard Woundy
(W)203-783-7996
(WF)203-783-5046
rwoundy@ibm.com

Russ Wright
(W)510-486-6965
(WF)510-486-6363
wright@lbl.gov

L. Michele Wright
(W)213-443-5743
(WF)213-443-5582
uncng!michele@uunet.uu.net

Sze-Ying Wuu
(W)201-829-5201
(WF)201-984-5837
syww@thumper.bellcore.com

June-Kang Yang
(W)408-370-4363
(WF)408-370-4222
natadm!yang@uunet.uu.net

Nancy Yeager
(W)217-244-8024
(WF)217-244-1987
nyeager@ncsa.uiuc.edu

Wengyik Yeong
(W)703-620-6651
(WF)703-620-4586
yeongw@psi.com

Jessica Yu
(W)313-936-2655
(WF)313-747-3745
jyy@merit.edu

William Yundt
(W)415-723-3104
(WF)415-723-0010
gd.why@forsythe.stanford.edu

Paul Zawanda
(W)217-244-4728
(WF)217-244-7396
zawanda@ncsa.uiuc.edu

John Ziegler
(W)508-562-2100
(WF)508-562-6942
ziegler@artel.com

Corporation
for
National
Research
Initiatives

*Research
and
Development
for
the
National
Information
Infrastructure*

1895 Preston White Drive
Suite 100
Reston, VA 22091
USA

Tel: (703) 620-8990
Fax: (703) 620-0913