Proceedings of the Twenty-Fourth Internet Engineering Task Force

Massachusetts Institute of Technology
NEARnet
Cambridge
July 13-17, 1992

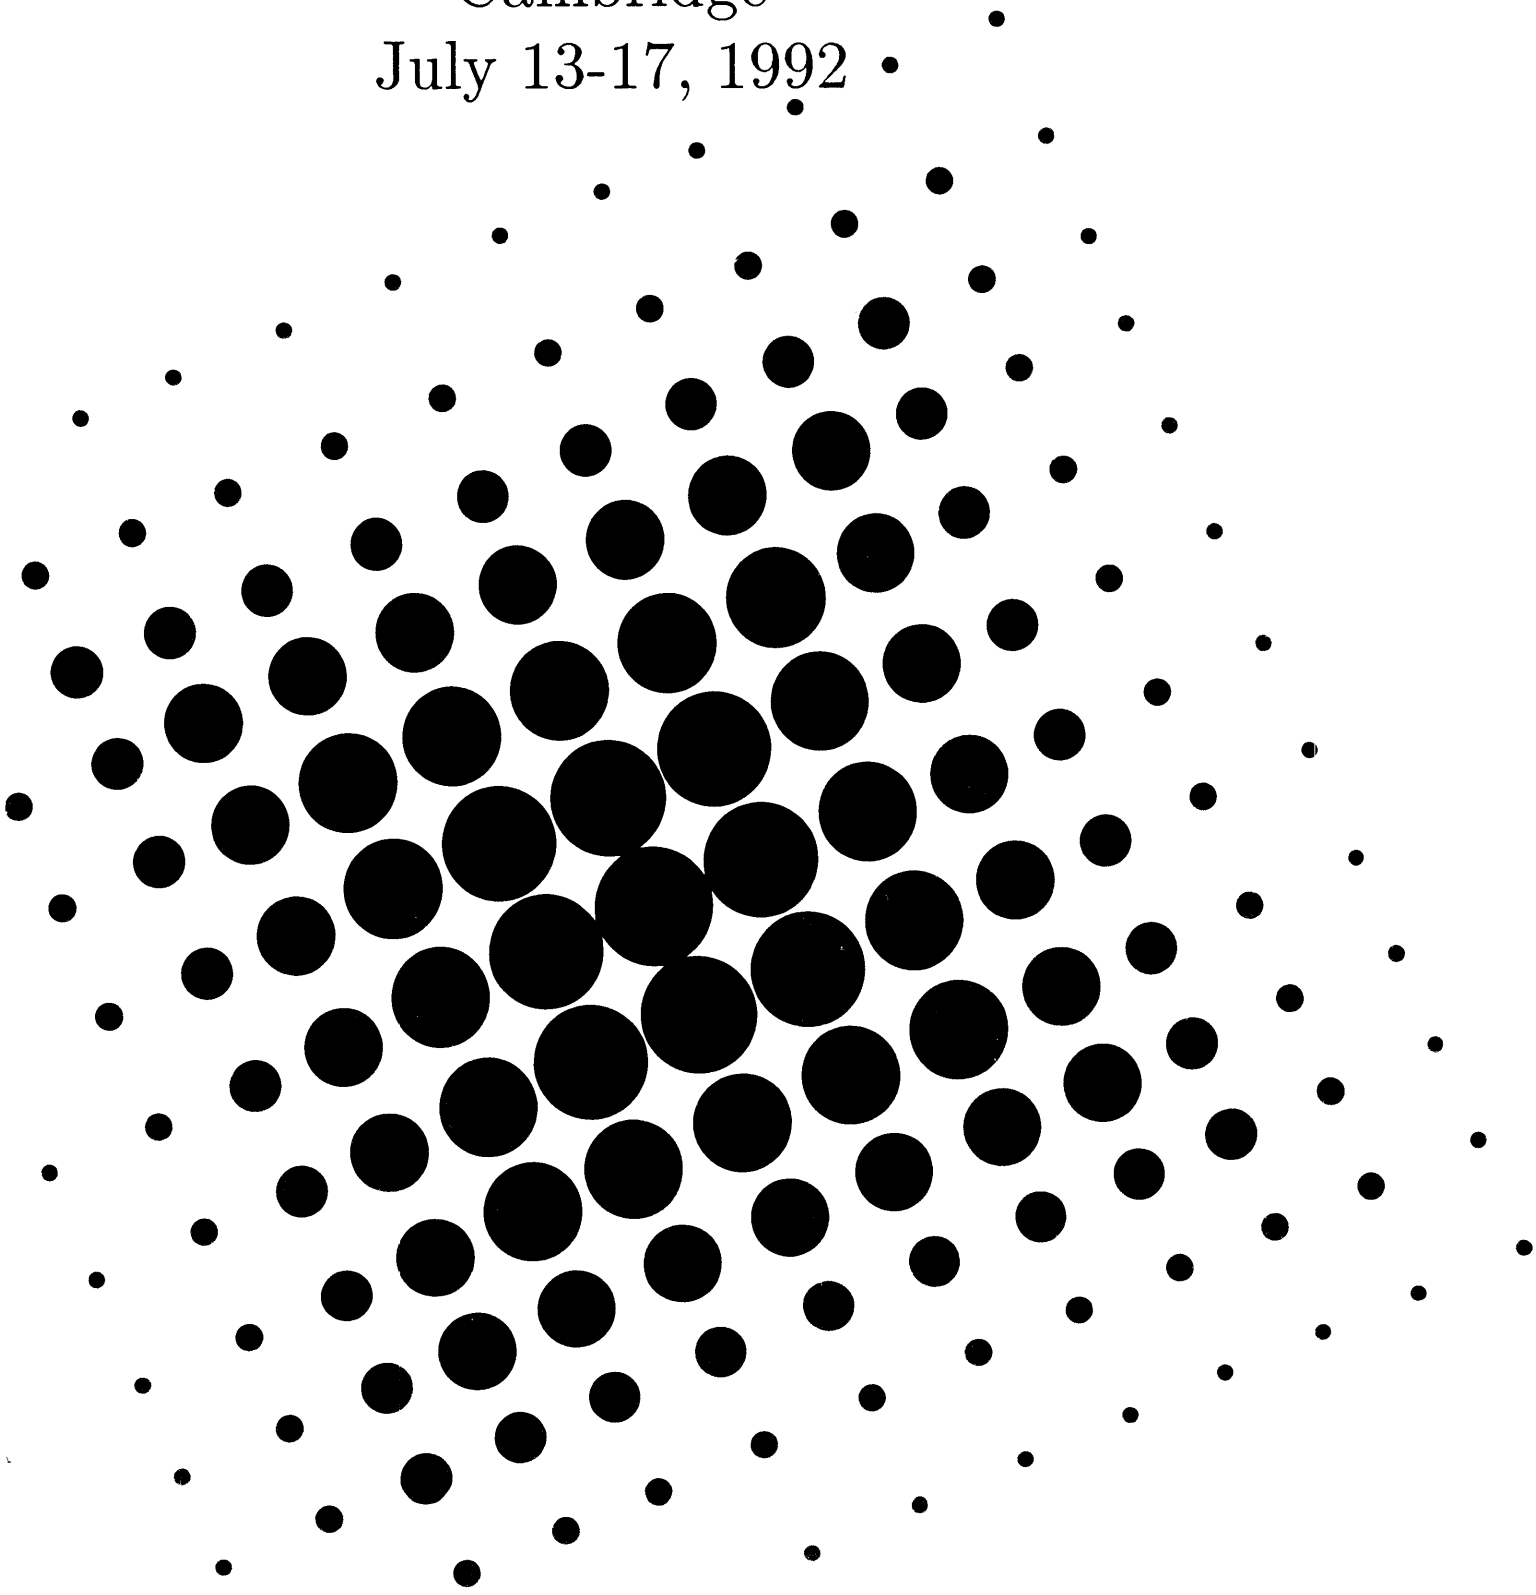

PROCEEDINGS OF THE TWENTY-FOURTH INTERNET ENGINEERING TASK FORCE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
NEARnet
CAMBRIDGE, MASSACHUSETTS
July 13-17, 1992

Compiled and Edited by
Megan 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 cooperative agreement with the National Science Foundation (NCR-8820945).

ACKNOWLEDGEMENTS

The Twenty-Fourth Internet Engineering Task Force was held at the Hyatt Regency in Cambridge, Massachusetts during the week of July 13th and set a new attendance record of 687. Though smaller in size than other hotels we've been in, the Hyatt managed to accommodate more than Eighty Working Group and Birds of a Feather sessions in its meeting space. In addition, the Hyatt staff were most helpful in assisting Jeff Schiller (MIT) and John Curran (NEARnet) with the setup of the Terminal Room.

Jeff and John, and a host of others, provided us with remarkable access to the Internet via a LaserLink hookup. With more than thirty terminals and four "u-bring-em" connections there were seldom any lines. In addition to Jeff and John, the following individuals deserve our thanks and appreciation: Richard Basch, James Bruce, Jerry Burke, Tom Coppeto, Joanne Costello, Occie Elder, Ron Hoffman, Jay McSweeney, Terri MacRae, Paul Milazzo, Chris Murphy, Ted Ts'o, Bruce Ottomano, and Frank Smith.

As the success of any IETF meeting is dependent on the combined efforts of interested individuals, we wish also to thank the following for their generous contributions:

Terminal Room Equipment

Apple Computer	Macintosh Computers
Cabletron Systems	Terminal Room LAN
NEARnet	Internet Link
New England Digital Data	LCI *Lace* 10Mb (Ethernet) Laser

Technical Presenters

Bob Braden, Dave Clark, Dave Crocker, Fred Glover, Bob Hinden, Mitch Kapor, Paul Tsuchiya and Tony Villasenor.

Our final note of appreciation is extended to Debra Legare and Cynthia Clark. In view of the increasing size of the IETF meetings, their ability to efficiently manage the registration process, while at the same time preserving its personal tone, is quite remarkable.

Megan Davies/CNRI

Contents

Director's Message

The IETF met in Cambridge from July 13-17, 1992, and was co-hosted by MIT and NEARnet. Our thanks and appreciation go out to Chuck Davin, Jeff Schiller, Ted T'so and John Curran, and to all the others that helped with the largest IETF meeting to date.

We are very pleased to have been able to invite the members of the Trusted Systems Interoperability Group (TSIG) to meet jointly with the IETF in Cambridge, especially since a number of IETF working groups are co-sponsored by TSIG. As the IETF continues to grow with the Internet itself, we expect to see more co-sponsored working groups with TSIG and other organizations. Efforts began in San Diego to co-sponsor working groups with RARE and to make it a standard practice; no longer handled as a special, one-time, case.

As expected, the main issue of discussion and comment was the growth of the Internet, and the challenges that need to be addressed, continuing the efforts initiated in 1991 with the establishment of the ROAD Group and the attention this topic received at both the Santa Fe and San Diego IETF meetings. One of the highlights of our week long meeting was a presentation by Dr. David Clark, a former member of the IAB and long time leader in the area of Internet Architecture, who spoke on the Internet of today, visions of tomorrow's Internet, and the tradition of the IETF.

The Cambridge meeting also saw presentations of two new proposals for handling the growth of the Internet: PIP (the "P" Internet Protocol) by Paul Tsuchiya, and IPAE (IP Address Encapsulation) by Bob Hinden and Dave Crocker. These two proposals join NIMROD, TUBA, and others for consideration by the IETF. As recommended by the IESG prior to the Cambridge meeting, these alternatives will be examined and considered by the IETF during the Washington D.C. meeting in November of 1992.

Growth of the IETF

As the Internet grows, so apparently does the IETF itself. The Cambridge meeting was another of record breaking proportions. Yes, it was another one of those! There were over 680 registered attendees at the 24th IETF meeting, an increase of just over 28% from the San Diego meeting! Additionally, there were over 80 Working Group, BOF, and directorate meetings held, an increase of almost 34% from San Diego levels! Interesting, isn't it, that the percentage increase in the number of groups meeting is higher than the percentage increase in the number of attendees.

And, as the Internet suffers from growth, so does the IETF. Along with the growth of the IETF is the increase in the number of suggestions for handling the growth.

Indeed, the first half of the Open Plenary almost turned into the IETF Growth BOF which had to be cancelled due to scheduling restrictions.

A number of topics and suggestions were discussed during the plenary, including the idea of higher hurdles for BOFs, an idea to restrict submitting Internet Drafts 2-4 weeks prior to an IETF meeting, and the ever popular topics of organizational hierarchy and the procedural processes.

The one suggestion which met with clear, almost unanimous agreement from the IETF (an almost unheard of condition) was the desire to create a moderated or controlled IETF mailing list to be used for announcements only, separate from the un-moderated IETF mailing list which should continue to exist in its un-moderated form. The consensus was overwhelmingly in favor of such an action that I am taking it upon myself to ensure this new separate list is set up before the November IETF.

Packet Video

Following up on the success in transmitting packet audio across the Internet to remote sites at the San Diego IETF meeting, Steve Casner and Steve Deering, with a number of other folks who aren't named Steve, came to the Cambridge IETF meeting with another experiment: transmitting packet video across the Internet! This time, there were over 90 audio sites and 45 video sites located in 10 different countries. While the San Diego setup was outgoing only, in Cambridge they were able to establish 2-way audio communications with the remote sites, which allowed for actual participation from the remote attendees during the meeting. In fact, there was even time allocated to permit questions and discussions from the remote participants. Our thanks and appreciation go to both Steves (and those behind the scenes) for extending the reach of the IETF meetings.

Future Meetings

The next plenary meeting of the IETF will be held in Washington, D.C. from November 16-20, 1992, and is being hosted by Sprint International. Following that, we will meet in Columbus, Ohio from March 29-April 2, 1993. This meeting is co-hosted by OARnet and The Ohio State University.

The July 1993 meeting will also be another first in that it is being held in Europe, hosted by RARE and SURFnet. The IETF will be meeting in Amsterdam from July 12-16, 1993.

Stephen J. Coya
Executive Director, IETF

IETF Progress Report

Between the IETF meetings in San Diego and Boston, there were seven new Working Groups created:

1. Mobile IP Working Group (mobileip)
2. Token Ring Remote Monitoring (trmon)
3. SNMP over a Multi-protocol Internet (mpsnmp)
4. Host Resources MIB (hostmib)
5. MIME-MHS Interworking (mimemhs)
6. TCP Client Identity Protocol (ident)
7. OSI IDRP for IP over IP (ipidrp)

and five working groups that were concluded:

1. OSI Internet Management (oim)
2. OSI General (osigen)
3. DECnet Phase IV MIB (decnetiv)
4. Special Host Requirements (shr)
5. Remote LAN Monitoring (rmonmib)

Additionally, there were 43 RFC's published since the San Diego IETF meeting in March, 1992:

RFC	Status	Title
RFC1305	PS	Network Time Protocol (v3)
RFC1313	I	Today's Programming for KRFC AM 1313 Internet Talk Radio
RFC1314	PS	A File Format for the Exchange of Images in the Internet
RFC1315	PS	Management Information Base for Frame Relay DTEs
RFC1316	PS	Definitions of Managed Objects for Character Stream Devices
RFC1317	PS	Definitions of Managed Objects for RS-232-like Hardware Devices
RFC1318	PS	Definitions of Managed Objects for Parallel-printer-like Hardware Devices
RFC1319	I	The MD2 Message-Digest Algorithm
RFC1320	I	The MD4 Message-Digest Algorithm
RFC1321	I	The MD5 Message-Digest Algorithm
RFC1322	I	A Unified Approach to Inter-Domain Routing
RFC1323	PS	TCP Extensions for High Performance
RFC1324	I	A Discussion on Computer Network Conferencing
RFC1325	I	FYI on Answers to Commonly asked "New Internet User" Questions
RFC1326	I	Mutual Encapsulation Considered Dangerous
RFC1327	PS	Mapping between X.400(1988) / ISO 10021 and RFC 822
RFC1328	PS	X.400 1988 to 1984 downgrading

RFC1329	I	Thoughts on Address Resolution for Dual MAC FDDI Networks
RFC1330	I	Recommendations for the Phase I Deployment of OSI Directory Services (X.500) and OSI Message Handling Services (X.400) within the ESnet Community
RFC1331	PS	The Point-to-Point Protocol (PPP) for the Transmission of Multi-protocol Datagrams over Point-to-Point Links
RFC1332	PS	The PPP Internet Protocol Control Protocol (IPCP)
RFC1333	PS	PPP Link Quality Monitoring
RFC1335	I	A Two-Tier Address Structure for the Internet: A Solution to the Problem of Address Space Exhaustion
RFC1336	I	Who's Who in the Internet Biographies of IAB, IESG and IRSG Members
RFC1337	I	TIME-WAIT Assassination Hazards in TCP
RFC1338	I	Supernetting: an Address Assignment and Aggregation Strategy
RFC1339	E	Remote Mail Checking Protocol
RFC1340	I	ASSIGNED NUMBERS
RFC1341	PS	MIME (Multipurpose Internet Mail Extensions): Mechanisms for Specifying and Describing the Format of Internet Message Bodies
RFC1342	PS	Representation of Non-ASCII Text in Internet Message Headers
RFC1343	I	A User Agent Configuration Mechanism For Multimedia Mail Format Information
RFC1344	I	Implications of MIME for Internet Mail Gateways
RFC1345	I	Character Mnemonics & Character Sets
RFC1346	I	Resource Allocation, Control, and Accounting for the Use of Network Resources
RFC1347	I	TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing
RFC1348	E	DNS NSAP RRs
RFC1349	PS	Type of Service in the Internet Protocol Suite
RFC1350	S	THE TFTP PROTOCOL (REVISION 2)
RFC1351	PS	SNMP Administrative Model
RFC1352	PS	SNMP Security Protocols
RFC1353	PS	Definitions of Managed Objects for Administration of SNMP Parties
RFC1354	PS	IP Forwarding Table MIB
RFC1357	I	A Format for E-mailing Bibliographic Records

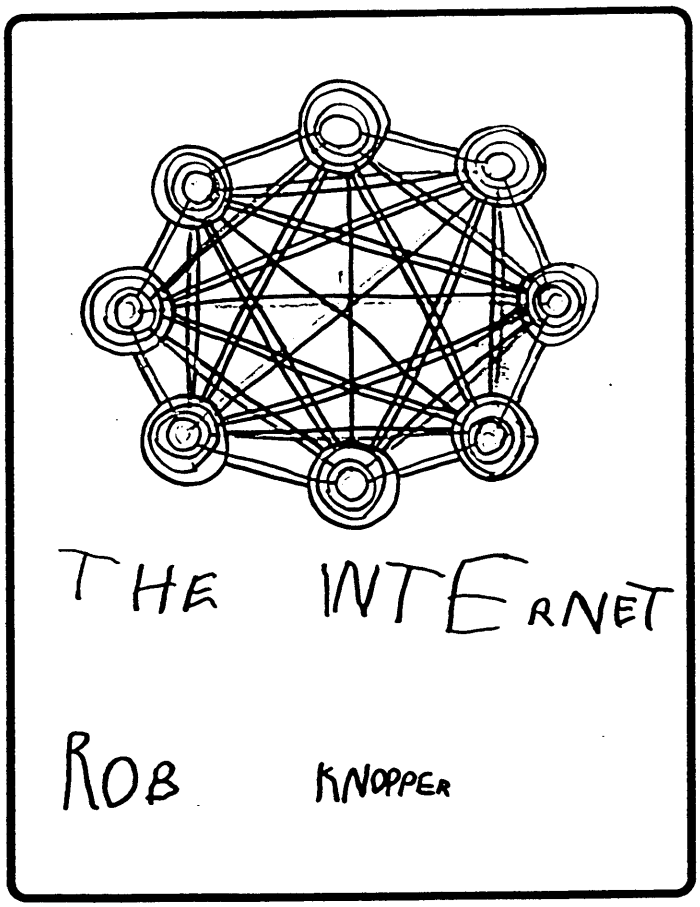

"This is the Internet...

...This is the Internet on drugs."

(Quote by Phill Gross.)

Agenda of the Twenty-Fourth IETF

(July 13-17, 1992)

MONDAY, July 13, 1992

7:30-8:30 am	IETF Registration and Continental Breakfast	
8:30-10:00 am	Introductions and Presentations	
	• "NREN Update" (Tony Villasenor/NASA)	
	• "IETF Internet Audio/Videocast" (Steve Casner/ISI and Steve Deering/Xerox PARC)	
10:00-12:00 noon	Morning Sessions	
	APP	Internet SMTP Extensions WG (smtpext) (John Klensin/MIT)
	INT	IP over Appletalk WG (appleip) (John Veizades/Apple)
	INT	IP over ATM WG (atm) (Bob Hinden/Sun)
	OPS	Network Status Reports (netstat) (Gene Hastings/PSC)
	OSI	OSI Directory Services WG (osids) (Steve Hardcastle-Kille/ISODE)
	RTG	OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)
	RTG	Open Shortest Path First IGP WG (ospf) (John Moy/Proteon)
	SEC	Security Area Advisory Group (saag) (Stephen Crocker/TIS)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	New Internet Routing and Addressing Architecture BOF (nimrod) (Noel Chiappa)
	BOF	Remote Conferencing BOF (remconf)(Jack Drescher/MCNC and Ari Ollikainen/LLNL)
	BOF	Universal Document Identifiers BOF (udi) (Tim Berners-Lee/CERN)
	INT	IP over Appletalk WG (appleip) (John Veizades/Apple)
	OSI	OSI Directory Services WG (osids) (Steve Hardcastle-Kille/ISODE)

1:30-3:30pm Monday, July 13, 1992 - Afternoon Sessions I (cont'd.)

RTG	Border Gateway Protocol WG (bgp) * (Yakov Rekhter/IBM)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
RTG	Open Shortest Path First IGP WG (ospf) * (John Moy/Proteon)
SEC	Network Access Server Requirements WG (nasreq) (Allan Rubens/Merit)
TSV	Domain Name System WG (dns) (Mike Reilly/DEC)
USV	Internet School Networking WG (isn) (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

BOF	Email Requirements BOF (mailreq) (Russ Hobby/UCDavis)
BOF	SNMP Security Implementors BOF (snmpseci) (Keith McCloghrie/Hughes and Jim Galvin/TIS)
OPS	Benchmarking Methodology WG (bmwg) (Scott Bradner/Harvard)
OSI	MHS-DS WG (mhsds) (Kevin Jordan/CDC and Harald Alvestrand/SINTEF DELAB)
OSI	Network OSI Operations WG (noop) (Sue Hares/Merit)
RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
RTG	IP Routing for Wireless/Mobile Hosts WG (mobileip) (Steve Deering/Xerox)
SEC	TSIG/IETF Coordination Meeting
TSV	Service Location Protocol WG (svrloc) (John Veizades/Apple)

* BGP and OSPF will meet jointly

TUESDAY, July 14, 1992

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	IETF Technical Presentations	
	• "The Futures of the Internet" (Mitch Kapor/EFF)	
9:30-12:00 noon	TSIG Plenary Session	
9:30-12:00 noon	Morning Sessions	
	BOF	Router Requirements Checklist BOF (rreqlist) (Pushpendra Mohta/CERFnet)
	APP	Network Database WG (netdata) (Daisy Shen/IBM)
	APP	Telnet WG (telnet) (Steve Alexander/INTERACTIVE Systems)
	INT	IP over ATM WG (atm) (Bob Hinden/Sun)
	INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant)
	MGT	FDDI MIB WG (fddimib) (Jeff Case/UTenn)
	MGT	Internet Accounting WG (acct) (Cyndi Mills/BBN and Gregory Ruth/BBN)
	OSI	MHS-DS WG (mhsds) (Kevin Jordan/CDC and Harald Alvestrand/SINTEF DELAB)
	RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)
	USV	User Documents WG (userdoc2) (Ellen Hoffman/UMich and Lenore Jackson/NASA)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	New Internet Routing and Addressing Architecture BOF (nimrod) (Noel Chiappa)
	MGT	Host Resources MIB WG (hostmib) (Steve Waldbusser/CMU)
	OPS	Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)

1:30-3:30pm Tuesday, July 14, 1992 - Afternoon Sessions I (cont'd.)

BOF	Shared Whois Project (swip) (Sheri Repucci/Merit)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
SEC	Common Authentication Technology WG (cat) (John Linn/DEC)
SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
TSG	Trusted Administration WG (tadmin) (Jeff Edelheit/MITRE)
TSG	Trusted Sessions WG (tsess) (Julie LeMoine/MITRE)
TSG	Trusted X WG (txwg) (Mark Smith/AT&T)
TSV	Trusted Network File Systems WG (tnfs) (Fred Glover/DEC)
USV	Internet User Glossary WG (userglos) (Tracy LaQuey Parker/UTexas and Gary Malkin/Xylogics)

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

4:00-6:00 pm Afternoon Sessions II

BOF	IP Addressing Plan (ipaddr) BOF (Bernhard Stockman/SUNET)
BOF	IP Security BOF (ipsec) (Steve Crocker/TIS)
OPS	Operational Area Directorate (orad) (Phill Gross/ANS, Susan Estrada/CERFnet, Bernhard Stockman/SUNET)
OSI	SNMP over a Multi-protocol Internet WG (mpsnmp) (Ted Brunner/Bellcore)
RTG	OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
TSG	Trusted Administration WG (tadmin) (Jeff Edelheit/MITRE)
TSG	Trusted Sessions WG (tsess)(Julie LeMoine/MITRE)
TSG	Trusted X WG (txwg) (Mark Smith/AT&T)

1:30-3:30pm Tuesday, July 14, 1992 - Afternoon Sessions II (cont'd.)

TSV Trusted Network File Systems WG (tnfs)
(Fred Glover/DEC)

USV Internet User Glossary WG (userglos)
(Tracy LaQuey Parker/UTexas and Gary Malkin/Xylogics)

7:00-10:00 pm Tuesday, July 14, 1992 - Evening Sessions

BOF Internet Society Q&A (isoc) (Vint Cerf/CNRI)

BOF Networked Information Retrieval BOF (nir)
(Jill Foster/UNewcastle-Upon-Tyne and George Brett/MCNC)

BOF OSF Distributed Computing Environment BOF (dce)
(Doug Hartman/OSF)

BOF Remote Conferencing BOF (remconf) (Jack Drescher/MCNC and Ari Ollikainen/LLNL)

BOF Uninterruptable Power Supply BOF (upsmib)
(Jeff Case/UTenn and Bob Stewart/Xyplex)

BOF Xwindows over OSI and Skinny Stack OSI (BOF)
(Dave Piscitello/Bellcore)

INT Point-to-Point Protocol Extensions WG (pppext)
(Brian Lloyd/Consultant)

OPS BGP Deployment and Application WG (bgpdepl)
(Jessica Yu/Merit)

RTG Multicast Extensions to OSPF WG (mospf)
(Steve Deering/Xerox PARC)

SEC TCP Client Identity Protocol WG (ident)
(Mike St. Johns/DOD)

WEDNESDAY, July 15, 1992

8:30-9:00 am — Continental Breakfast

9:00-9:30 am — Technical Presentations

- "Pip: The 'P' Internet Protocol" (Paul Tsuchiya/Bellcore)

9:00-12:00 noon — TSIG Plenary

9:30-12:00 noon — Morning Sessions

INT	IP over ATM WG (atm) (Bob Hinden/Sun)
INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant)
MGT	Chassis MIB WG (chassis) (Jeff Case/UTenn and Bob Stewart/Xyplex)
MGT	DS1/DS3 MIB WG (trunkmib) (Fred Baker/ACC and Tracy Cox/Bellcore)
OSI	Office Document Architecture WG (oda) (Peter Kirstein/UCL)
RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
TSG	Trusted Administration WG (tadmin) (Jeff Edelheit/MITRE)
TSG	Trusted Sessions WG (tsess) (Julie LeMoine/MITRE)
TSG	Trusted X WG (txwg) (Mark Smith/AT&T)
TSV	Trusted Network File Systems WG (tnfs) (Fred Glover/DEC)
USV	User Services WG (uswg) (Joyce Reynolds/ISI)

Breaks — Coffee available throughout morning.

1:30-3:30 pm — Afternoon Sessions I

BOF	Remote Conferencing BOF (remconf) (Jack Drescher/MCNC and Ari Ollikainen/LLNL)
MGT	Token Ring Remote Monitoring WG (trmon) (Mike Erlinger/Lexcel)

1:30-3:30pm Wednesday, July 15, 1992 - Afternoon Sessions I (cont'd.)

OPS	BGP Deployment and Application WG (bgpdepl) (Jessica Yu/Merit)
OSI	X.400 Operations WG (x400ops) (Alf Hansen/SINTEF DELAB)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
RTG	RIP Version II WG (ripv2) (Gary Malkin/Xylogics)
SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
TSG	Trusted Administration WG (tadmin) (Jeff Edelheit/MITRE)
TSG	Trusted Sessions WG (tsess) (Julie LeMoine/MITRE)
TSG	Trusted X WG (txwg) (Mark Smith/AT&T)
TSV	Trusted Network File Systems WG (tnfs) (Fred Glover/DEC)
USV	Internet Anonymous FTP Archives WG (iafa) (Peter Deutsch/McGill and Alan Emtage/McGill)

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

4:00-6:00 pm Afternoon Sessions II

MGT	IEEE 802.3 Hub MIB WG (hubmib) (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
OPS	Operational Area Directorate (orad) (Phill Gross/ANS, Susan Estrada/CERFnet, Bernhard Stockman/SUNET)
OSI	X.400 Operations WG (x400ops) (Alf Hansen/SINTEF DELAB)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
TSG	Trusted Administration WG (tadmin) (Jeff Edelheit/MITRE)
TSG	Trusted Sessions WG (tsess) (Julie LeMoine/MITRE)
TSG	Trusted X WG (txwg) (Mark Smith/AT&T)
TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)

4:00-6:00pm Wednesday, July 15, 1992 - Afternoon Sessions II (cont'd.)

TSV Trusted Network File Systems WG (tnfs)
(Fred Glover/DEC)

USV Network Information Services Infrastructure WG (nisi)
(April Marine/SRI and Pat Smith/Merit)

7:00-10:00pm Wednesday, July 15, 1992 - Evening Session

BOF Authorization and Access Control BOF (aac)
(Clifford Neuman/ISI)

BOF A New Internet Protocol BOF (pip)
(Paul Tsuchiya/Bellcore)

BOF Directory Resources Engineering Group BOF (dregs)
(Joan Gargano/UCDavis and Joyce K. Reynolds/ISI)

BOF Simple Management Protocol (SMP) Framework BOF
(smpframe) (Marshall Rose/DBC)

BOF Perspectives on the Next Generation of the NSFnet
(nsfnet) (Laura Breeden/FARNET)

THURSDAY, July 16, 1992

8:30-9:00 am — Continental Breakfast

9:00-9:30 am — Technical Presentations

- "The DARPA Research Testbed Network (DARTnet): A Progress Report" (Bob Braden/ISI)

9:30-12:00 noon — TSIG Plenary

9:30-12:00 noon — Morning Sessions

BOF	IP Address Encapsulation BOF (ipae) (Bob Hinden/Sun and Dave Crocker/TBO)
INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant)
MGT	Bridge MIB WG (bridge) (Fred Baker/ACC)
OPS	User Connectivity Problems WG (ucp) (Dan Long/BBN)
RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
SEC	Privacy-Enhanced Electronic Mail WG (pem) (Steve Kent/BBN)
TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)

Breaks — Coffee available throughout the morning.

1:30-3:30 pm — Afternoon Sessions I

APP	Network Database WG (netdata) (Daisy Shen/IBM)
BOF	TCP/UDP over CLNP-addressed Networks BOF (tuba) (Ross Callon/DEC and Peter Ford/LANL)
MGT	Internet Accounting WG (acct) (Cyndi Mills/BBN and Gregory Ruth/BBN)
MGT	X.25 Management Information Base WG (x25mib) (Dean Throop/Data General)
OPS	Network Joint Management WG (njm) (Gene Hastings/PSC)

1:30-3:30pm Thursday, July 16, 1992 - Afternoon Sessions I (cont'd.)

OSI MIME-MHS Interworking WG (mimemhs)
(Steve Thompson/SoftSwitch)

RTG IP Routing for Wireless/Mobile Hosts WG (mobileip)
(Steve Deering/Xerox)

SEC Security Area Advisory Group (saag)
(Stephen Crocker/TIS)

USV Directory Information Services Infrastructure WG (disi)
(Chris Weider/Merit)

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

4:00-6:00 pm Technical Presentations

- "Trusted NFS: Protocol Extensions for MultiLevel Security" (Fred Glover/DEC)
- "IP Encapsulation" (Bob Hinden/Sun and Dave Crocker/TBO)
- "A Cloudy Crystal Ball - Visions of the Future" (Dave Clark/MIT)

7:00pm-9:00pm Open Plenary and IESG

FRIDAY, July 17, 1992

8:30-9:00 am	Continental Breakfast	
9:00-9:30am	Technical Presentations	
9:30-12:00pm	Summary Reports	
	APP	Applications Area (Russ Hobby/UC Davis)
	INT	Internet Area (Noel Chiappa and Philip Almquist/Consultant)
	MGT	Network Management Area (Chuck Davin/MIT)
	OPS	Operations Area (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/SUNET)
	OSI	OSI Integration Area (Erik Huizer/SURFnet and David Piscitello/Bellcore)
	RTG	Routing Area (Bob Hinden/Sun)
	SEC	Security Area (Steve Crocker/TIS)
	TSV	Transport and Services Area (David Borman/Cray Research)
	USV	User Services Area (Joyce K. Reynolds/ISI)
12:00 pm	Concluding Remarks (Phill Gross/ANS)	
12:30 pm	Adjourn	

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
TSG	Trusted Sytems Interoperability Group (TSIG)
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	Philip Almquist/Consultant
Network Management:	James Davin/ Bellcore
OSI Integration:	Dave Piscitello/Bellcore
	Erik Huizer/SURFnet
Operational Requirements:	Phill Gross/ANS
	Bernhard Stockman/SUNET
Routing:	Robert Hinden/Sun
Security:	Steve Crocker/TIS
Transport and Services	David 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 include 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

Fall 1992

Washington, DC
U.S. Sprint
Host: Robert Collet
November 16-20, 1992

Spring 1993

Columbus, OH
OARnet and The Ohio State University
Host: Kannan Varadhan
March 29-April 2, 1993

Summer 1993

Amsterdam, Netherlands
SURFnet and RARE
Host: Erik Huizer
July 12-16, 1993

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 must be submitted to the IESG or the RFC editor, or they 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.

Internet Drafts are generally in the format of an RFC, although it is expected that the documents may be "rough" drafts. This format is specified fully in RFC 1111. In brief, an Internet Draft shall be submitted in ASCII text, limited to 72 characters per line and 58 lines per page followed by a formfeed character. Overstriking to achieve underlining is not acceptable.

Postscript is acceptable, but only when submitted with a matching ASCII version (even if figures must be deleted). Postscript should be formatted for use on 8.5x11 inch paper. If A4 paper is used, an image area less than 10 inches high should be used to avoid printing extra pages when printed on 8.5x11 paper.

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 the upper left hand corner of the first page. The document should NOT refer to itself as an RFC or a Draft RFC.

The Internet Draft should neither state nor imply that it is a Proposed Standard. To do so conflicts with the role of the IAB, the RFC Editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the Internet Draft. All Internet Draft should include a section containing the following verbatim statement:

> This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.
>
> Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

> To learn the current status of any Internet Draft, please check the 1id-abstracts.txt listing contained in the Internet-Drafts Shadow Directories on nic.ddn.mil, nnsc.nsf.net, nic.nordu.net, ftp.nisc.sri.com, or munnari.oz.au.

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.

A document expiration date must appear on the first and last page of the Internet Draft. The expiration date is always six months following the submission of the document as an Internet Draft. Authors can calculate the six month period by adding five days to the date when the final version is completed. This should be more than enough to cover the time needed to send the document or notification of the document's availability to internet-drafts@nri.reston.va.us.

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

Area and Working Group Reports

2.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UC Davis

Area Overview

The general goal of the area is to define the protocols to create an interoperable multimedia distributed computing environment for the Internet. The Internet is now a global communications resource and people want more applications than the standard Telnet, FTP and SMTP provide.

At the Boston meeting there were demonstrations of how video and audio could be carried between workstations over the Internet today. The Teleconferencing Architecture BOF discussed what is necessary to create a multi-workstation multimedia conference environment. With the greater international interest in the IETF there is a desire to be able to extend the "face-to-face" meetings over the network around the world.

Another interest that generated several BOFs was information storage and retrieval over the Internet. There is a growing amount of information available over the network, but how does one locate and retrieve it?

Email Requirements BOF (MAILREQ)

The defined protocols (SMTP, RFC 822, MIME, ...) do not begin to address the complexity of the email environment of the Internet and beyond. We need a set of documents that can give "the big picture" and provide guidelines for implementors, operators and users. The BOF discussed problems and needs for the email world and some possible documents to address these issues.

OSF Distributed Computing Environment BOF (DCE)

Doug Hartman from OSF presented and overview of DCE and answered questions. DCE addresses several of the "holes" that we currently have in the Internet Protocols. Future work in the IETF to fill these holes may be able to use the work of OSF as a strong starting place.

Remote Conferencing BOF (REMCONF)

The BOF continued work to define an overall architecture to cover all aspects of remote conferencing. This included things such as session management and groupware. Separate Working Groups may be generated to work on specific protocols specified by the archi-

tecture. There was discussion on the current work on session management. This may be spawned off as a Working Group.

Internet Mail Extensions Working Group (SMTPEXT)

The Working Group did the final review of the document that extends SMTP to allow the transport of 8-bit characters and provides some additional capabilities to improve efficiency, especially when very large files are being transmitted. Working Group has submitted the document to be a Proposed Standard.

Network Database Working Group (NETDATA)

Two approaches to provide SQL over TCP/IP networks were reviewed. One approach has been developed by the Working Group Chair, Daisy Shen over the past few IETF meetings. The other approach has been developed by the SQL Access Group and was presented to the IETF for the first time at the Boston meeting. The Working Group agreed that a single approach needs to be developed for standardization and would evaluate both approaches. Security work is still needed on both approaches.

TELNET Working Group (TELNET)

The TELNET Working Group made further progress on authentication and decided that the document should be put forth as an Experimental Protocol. Dave Borman presented an extension to remote flow control that the Group reviewed and will be submitted to be a Proposed Standard. The Working Group reviewed the Environment Option and will put it forward to become a Proposed Standard.

CURRENT MEETING REPORT

Reported by Russ Hobby/UC Davis

Minutes of the Email Requirements BOF (MAILREQ)

The BOF was started with several questions about working on email requirements.

WHY? Because the protocols do not specify everything. No one document puts all the information in one place.

WHAT? Scope to look at: Internet? Email World? Gateways? Audience: Implementors? Operators? Users? Documents: How many and what?

HOW? Next step beyond the BOF?

WHO? There needs to be a Working Group Chair. There needs to be writers for the documents.

The Group came up with a gripe/wish list to focus on what work should be done and what documents would be written. The list was:

1. Email flow control
2. Clarify 822 ambiguities
3. Better/standard error messages and tracking
4. Guidelines for handling violations
5. Management of list exploders
6. Duplicate suppression
7. Server command sets
8. Distributed message storage
9. Email API
10. Verify that recipient exists
11. Common Terminology
12. Email MIB
13. Automatic verification of delivery/reading
14. Multipath elimination
15. Email routing protocol
16. Vacation mail requirements
17. Retry and reporting requirements
18. Know when a message comes from a list
19. Postmaster behavior requirements
20. Dynamic and more detailed MX
21. Human readable errors
22. Email/directory interface.

There was discussion of several possible documents that could be produced. There was

general agreement that there really needs to be a document or group of documents that define an overall email architecture. This would help us describe where we are with email today and provide direction for work. It was also suggested that we think of email messages as "Big Packets" and use some of the technology that we now use for our packet networks in the email world.

Suggested documents were:

1. Email Architecture (three possible documents)
   - The Big Email Picture (General email components and terminology)
   - Internet Email Architecture (SMTP/822/MIME/...)
   - Email Gateways
2. "Packet" Email System
3. Email MIB
4. Listserve Command Syntax
5. List Behavior and Management
6. Error Messages and Handling
7. Implementor's Email Requirements
8. Email/Directory Interface
9. How to find an Email Address
10. Quality of Email Service

There was agreement that these issue urgently need addressing and that a Working Group should be formed. The following mail list will be created for further discussion:

ietf-mailreq@ucdavis.edu for Group discussion

ietf-mailreq-request@ucdavis.edu for list adds and drops

Attendees

Mark Bokhan	bokhan@abitok.enet.dec.com
Luc Boulianne	lucb@cs.mcgill.ca
James Conklin	jbc@bitnic.educom.edu
David Crocker	dcrocker@mordor.stanford.edu
Roger Fajman	raf@cu.nih.gov
Jill Foster	jill.foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
Tony Genovese	genovese@nersc.gov
Terry Gray	gray@cac.washington.edu
Alf Hansen	Alf.Hansen@delab.sintef.no
Russ Hobby	rdhobby@ucdavis.edu
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Todd Kaehler	kaehler@zk3.dec.com

Neil Katin	katin@eng.sun.com
David Katinsky	dmk@rutgers.edu
John Klensin	klensin@infoods.mit.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Keith Moore	moore@cs.utk.edu
John Murray	murray@premenos.sf.ca.us
Michael Patton	map@lcs.mit.edu
Michael Powell	mdpowel@pacbell.com
Anthony Rutkowski	amr@sprint.com
Tim Seaver	tas@concert.net
Mark Smith	mcs@umich.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Wagner	jwagner@princeton.edu
Fred Ziegler	ziegler@aspentec.com

CURRENT MEETING REPORT

Reported by George Brett/MCNC

Minutes of the Networked Information Retrieval BOF (NIR)

The Group agreed that the mailing list, nir@cc.mcgill.ca, which was set up by Peter Deutsch would continue to be used by the NIR Working Group and all names on the roster would be added to that list. The Group also discussed the elements of the Charter necessary to develop this BOF into for a Working Group. The draft of the Charter has been posted to the NIR mailing list. The goals and objectives are still up for discussion.

Other items of discussion from the floor included the frequency of publications, future definition of networked information retrieval standards, and the scope of the documents produced by the NIR Working Group.

Updates from Brett & Foster

Jill Foster and George Brett presented information about projects that they are working on. Jill described the RARE Information Services and User Support Working Group. George spoke about the Coalition for Networked Information Directories and Networked Resources Services Working Group and the Clearinghouse for Networked Information Discovery and Retrieval. For more information on these projects, Jill and/or George can be reached via email (see addresses below).

Templates

A majority of the session was spent talking about the two templates that will be used to gather information about networked information retrieval applications and about organizations which are working on projects related to NIR. The draft copies of the templates have been posted to the NIR mailing list and are available at the archive kept there (pub/mailing-lists/nir at archives.cc.mcgill.ca). Comments and suggestions about the templates are welcome, but we do plan to begin implementing them before the November 1992 IETF meeting.

Attendees

George Abe	abe@infonet.com
Ed Albrigo	ealbrigo@cos.com
George Brett	ghb@jazz.concert.net
Mitchell Charity	mcharity@lcs.mit.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Alan Clegg	abc@concert.net
Richard Colella	colella@osi.ncsl.nist.gov
Peter Deutsch	peterd@cc.mcgill.ca

Alan Emtage	bajan@cc.mcgill.ca
Roger Fajman	raf@cu.nih.gov
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Alf Hansen	Alf.Hansen@delab.sintef.no
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
J. Paul Holbrook	holbrook@cic.net
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
David Katinsky	dmk@rutgers.edu
John Klensin	klensin@infoods.mit.edu
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Walter Lazear	lazear@gateway.mitre.org
Thomas Lenggenhager	lenggenhager@switch.ch
Hock-Koon Lim	lim@po.cwru.edu
April Marine	april@nisc.sri.com
Charlotte Mooers	mooers@nnsc.nsf.net
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Rakesh Patel	rapatel@hardees.rutgers.edu
Michael Patton	map@lcs.mit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Mel Pleasant	pleasant@cs.rutgers.edu
Sheri Repucci	smr@merit.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Karen Roubicek	roubicek@faxon.com
Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Vincent Sgro	sgro@cs.rutgers.edu
Mark Smith	mcs@umich.edu
Patricia Smith	psmith@merit.edu
Simon Spero	ses@cmns.think.com
Chris Weider	clw@merit.edu
Yung-Chao Yu	yy@qsun.att.com

CURRENT MEETING REPORT

Reported by Doug Hartman/OSF

Minutes of the OSF Distributed Computing Environment BOF (DCE)

The Open Software Foundation (OSF) hosted a BOF to discuss the Distributed Computing Environment (DCE) work by OSF. It was attended by about forty people (see list). Doug Hartman, DCE engineering director, gave a quick tour of the OSF organization, what OSF does, and what DCE is. (See following slides.) The most important points are:

- OSF is a research and development cooperative that develops things like Motif, OSF/1, DCE and DME in conjunction with our 300+ members.

- DCE is a set of libraries and programs to be used with existing networks and transports to create production-quality distributed systems and applications.

- DCE provides services for RPC, multithreading, network security, global naming of objects, and a distributed file system. These services are integrated to provide a consistent foundation.

- The DCE specifications will be available from bookstores within a few months. DCE code is licensable now. (Implementations before specifications.) Many companies are planning DCE deployments.

- OSF invites IETF attendees to learn more about OSF, DCE, and our other activities. Contact hartman@osf.org for more information.

The audience asked a number of good questions throughout the presentation to clarify points or to understand how this information might affect them. Some highlights were:

- Question: What about OSI? This doesn't look like OSI, and many OSF members are active in OSI.

  - Answer: DCE does not replace OSI protocols for use by government organizations and telecomm providers. DCE does address some of the same technology areas, but is more focused on production distributed applications which need high performance, security and interoperability today. OSF is working with standards organizations to give them the benefit of our "existing practice"; we would be happy to see DCE protocols included in OSI upper layers. (DCE can use OSI lower layers already.)

- Question: How are decisions made at OSF? Don't the OSF sponsors get involved?

  - Answer: OSF members give us input as to what they would like to see us do. We need to cover our costs from license revenue, so we are motivated to meet

the needs of our licensees. While the OSF board (including sponsors) initially made the recommendation to pursue interoperability technology in 1988, OSF employees have made the decisions about what that technology would be since then. All OSF members are regularly given the chance to provide their recommendations to us.

- Question: How will we proceed with future OSF <-> IETF discussions?
  - Answer: BOFs sometimes lead to working groups. This does not seem like the right thing for DCE, since the "DCE Working Group" is already in place at OSF, where it is called the DCE SIG; it meets about every two months. The most likely course of action at this point is to have OSF work with the Applications Area of the IETF. We would be able to present details about DCE solutions in the areas of RPC, naming, security, and so forth.

Attendees

Doug Barlow	barlow@decwet.dec.com
James Beers	beers@nr-tech.cit.cornell.edu
Natan Borshansky	natanb@turisc.timeplex.com
Randy Butler	rbutler@ncsa.uiuc.edu
Edward Cande	cande@zk3.dec.com
John Chang	jrc@uswest.com
Isaac Choi	ichoi@synoptics.com
Scott Cossette	sdc@concord.com
Fran Fadden	fran@decvax.dec.com
Bill Fardy	fardy@ctron.com
Kathryn Fuller	kathryn@codex.com
Fred Glover	fglover@zk3.dec.com
Terry Gray	gray@cac.washington.edu
Steve Hardcastle-Kille	s.kille@isode.com
Douglas Hartman	hartman@osf.org
Arthur Lin	yalin@srv.pacbell.com
Ellen McDermott	emcd@osf.org
Ed Menze	menze@cs.arizona.edu
Keith Moore	moore@cs.utk.edu
John Murray	murray@premenos.sf.ca.us
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
Kary Robertson	kr@concord.com
Dan Romascanu	dan@lannet.com
Henry Sinnreich	hsinnreich@mcimail.com
Carl Smith	cs@eng.sun.com

Dean Throop	throop@dg-rtp.dg.com
Sylvia Tom	tom_sylvia@tandem.com
John Vollbrecht	jrv@merit.edu
Justin Walker	justin@apple.com
Peter Williams	p.williams@uk.ac.ucl.cs
Joseph Zur	zur@fibhaifa.com

OSF Distributed Computing Environment

IETF BOF

July 14, 1992

Doug Hartman

hartman@osf.org

Introduction to OSF

NCRA organization for open software R&D

300+ organizational members provide input

membership is open; there's an annual fee

250+ OSF employees make decisions, do work

(of course, we gets lots of help!)

We work with standards organizations

OSF technology offerings

Building blocks for open systems
- Motif
- OSF/1
- DCE
- DME

Implementations in ANSI C for UNIX-like OS's

Specifications follow implementation (IETF model)

Specifications (AES) in bookstores (on-line?) include protocols and programming interfaces

Implementation licenses cover our costs

The DCE offering

Libraries, servers, tools for distributed computing

Core services: RPC, naming, security, for creating distributed applications

Portable to a variety of OS and hardware platforms

Builds on existing network transports

Distributed UNIX file system included

Environment: bring your own e-mail, spreadsheet

DCE acceptance

DCE implementations (in whole or in part) in progress for OSF/1, SVR4, AIX, HP-UX, VMS, MVS, Windows, NT (partial list)

DCE developer kits shipping from DEC, Transarc, Gradient; many others by end of 1992

Active participation by system vendors:IBM, DEC, HP, SNI, NCR, Bull, Cray, Stratus, Hitachi, ICL
And end users: Boeing, EDS, Bellcore, Citicorp, Mitre (partial list)

Explorations at Novell, Apple, Tandem, Sunsoft

We're exploring use in embedded applications

DCE Design Principles

Use one base communication vehicle: RPC

Standardize configurations with common servers

Combine namespaces to name everything

Build in authentication and access control

No built-in scaling limits

Address production use: threads, replication

Use existing networks, OS, infrastructure

Client/Server Architecture

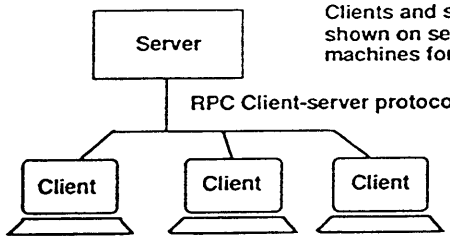

Programming and user interfaces define clients

RPC interfaces define servers (except X.500, which uses message protocols)

RPC Client/Server Interfaces

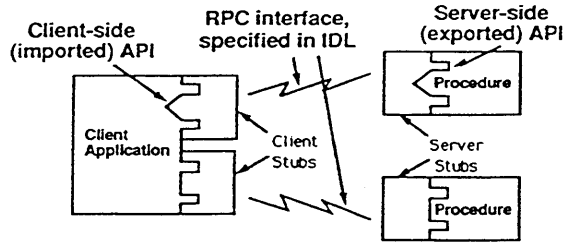

- Interface definition stays constant
- Stub behavior can vary
- Allows creation of 'network library' by compiling IDL into ordinary API

Use of Transport Protocols

dfs	sec	time	cds gda	X.500 gds	
rpc interface				OSI presentation	
CL rpc		CO rpc		OSI session	DCE
UDP or other datagram transport		TCP, TP0-4, or other connection-based transport			OS

Transport interface (e.g. sockets or XTI) requires multi-event wait (e.g. select)

Connection-based transports assume a reliable, sequenced byte stream; message frames are OK

DCE technical issue #1

Why use RPC? Why not messages?

RPC is fast, intuitive, heteregenous, extensible

RPC uses transport-level messages

RPC can provide high-level queued delivery

Use threads to address concurrency

Multithreaded Design

DCE servers handle concurrent requests from clients

Multithreaded clients: concurrent requests use threads built into DCE RPC

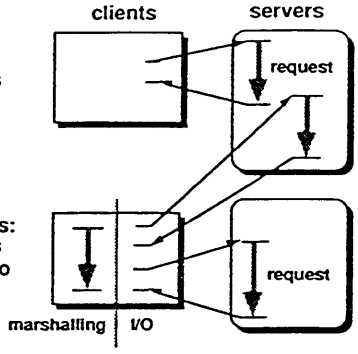

Cells

Cell: an independent set of clients & servers, managed as a group

DCE's idea of domain

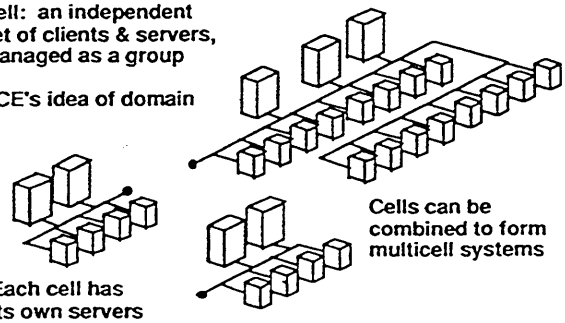

Cells can be combined to form multicell systems

Each cell has its own servers (shown as machines)

Creating Multicell Systems

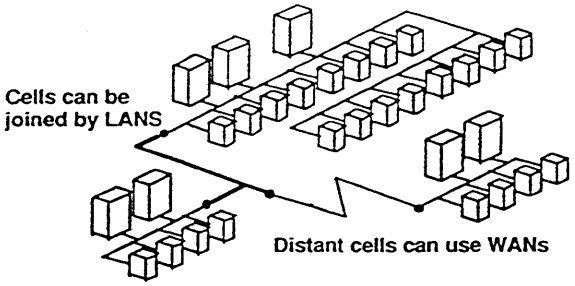

Cells are autonomous, but cooperate (for example, sharing cell security keys)

Designed for an internet

DCE uses underlying addressing, routing

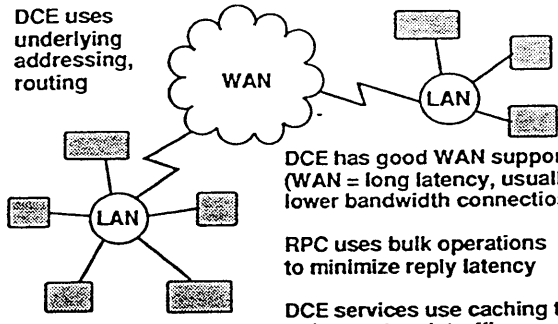

DCE has good WAN support (WAN = long latency, usually lower bandwidth connection)

RPC uses bulk operations to minimize reply latency

DCE services use caching to reduce network traffic

Integrated Global Namespace

Each object has a unique hierarchical name

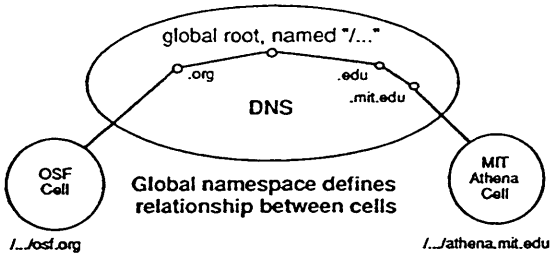

Global namespace defines relationship between cells

Cells can appear in either DNS or X.500

Cell Namespace

Each cell has same namespace organization

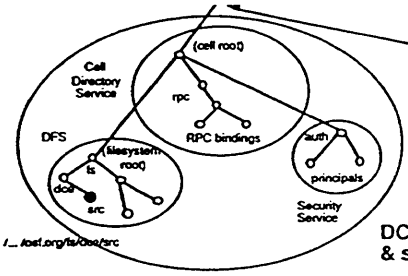

Global Directory Agent manages global directory junction

DCE supports DFS & security junctions

Location transparency is managed within cells

DCE technical issue #2

Should we use X.500 to name everything?

No: use junctions (federated naming)

Name requirements/syntaxes will always vary

Provide several services for unique needs

Allow extensibility

Security Architecture

Clients log in, authenticate, get Privilege Attribute Certificate:

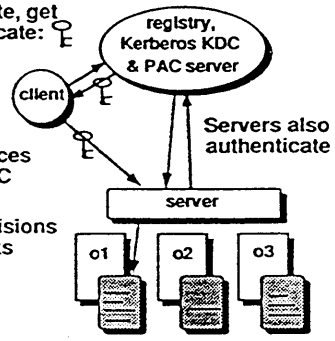

Servers also authenticate

Clients request services via authenticated RPC

Servers make access decisions using Access Control Lists attached to objects (files, names in the namespace)

Security and Multicell Access

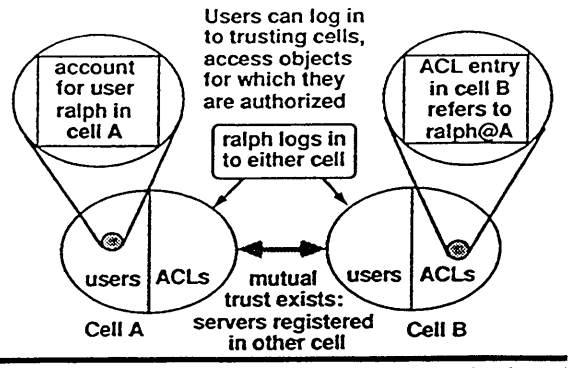

DCE technical issue #3

What about public key?

Doesn't replace need to use fast (DES) keys

Common server design works well for secret key

We hope to integrate public key in the future

Public key patent/license issues still being worked

File System Architecture

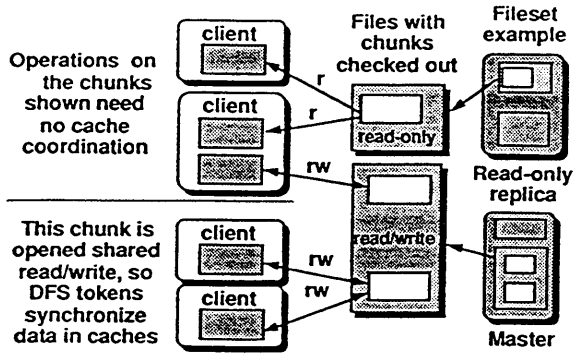

Relationship to NFS

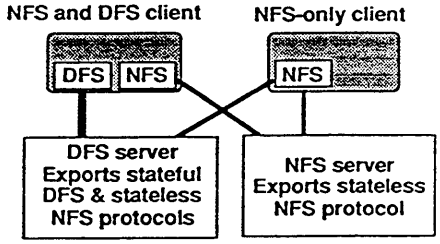

DFS protocol provides global file namespace, security, caching, replication, administration

NFS protocol provides backwards compatibility

DCE technical issue #4

Why would I want to use this file system?

Consistent global file naming

Integrated security

Improved caching, concurrency, replication

Better performance for large networks

Operating System Kernel Interfaces

The DCE DFS is designed for a UNIX-based OS kernel:

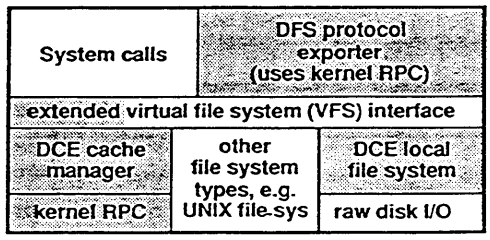

Shaded parts are provided with DCE, assume presence of other parts in the OS

Operating system requirements

These functions are needed for full DCE:

basic multitask support, 1 MB+ address space, memory protect	file and process operations per POSIX & XPG3, including time and IPC	native threads or support for DCE threads (nonblocking I/O and time ticks)

Partial DCE support (e.g. client-only) may be possible on systems without these functions.

"Where's Waldo?"

Kerberos: Security KDC mechanism

DNS: one way to locate cells

NTP: time service can accept NTP time

X.500: Client, server, OSI layer 5-7

AFS: precedessor to DCE DFS

NCS: RPC is NCS 2.0

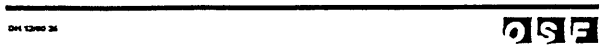

DCE status

DCE 1.0 technology released January 1992

Update release with bugs fixes, late features coming out in three weeks

Vendor ports released to pilot installations late this year

Work continues to enhance technology to address future areas

DCE application areas

"DME apps": print managers, license managers

OLTP extensions (e.g. Transarc)

Object oriented extensions (e.g. HP)

Distributed DBMS (e.g. Oracle, Informix)

Distributed development environment tools, debuggers, tracing tools, CASE tools

"Groupware": document mgmt, project mgmt

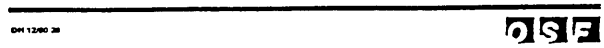

Future Directions

Internationalization

Performance tuning

Tracing and auditing

Easier administration

Easier programming

Better bridges to non-DCE environments

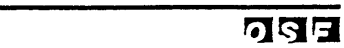

CURRENT MEETING REPORT

Reported by Jack Drescher/MCNC and Ari Ollikainen/LLNL

Minutes of the Remote Conferencing BOF (REMCONF)

Two BOF sessions were held on Remote Conferencing as a follow-on to the Teleconferencing Architecture BOFs held at the San Diego IETF in March. A preparatory meeting was added to the schedule between the first and second full BOFs. Sixty one different people attended either one or both BOF sessions.

The stated purpose of the BOF was to present/solicit feedback on the first draft of the Remote Conferencing Architecture Outline/Structure, which had been sent to list prior to the IETF. Copies were also handed out prior to the first session. Other purposes were to solicit other major research efforts which might be relevant to the Remote Conferencing Architecture and to develop ties with these efforts as appropriate.

Full Session of July 13, 1992

A brief statement of continuity from previous BOF sessions was reviewed along with the proposed Working Group Milestones. It was noted that the Group had not yet been able to achieve IESG approval as a Working Group. The main issue is "broadness of Charter".

Yee-Hsiang Chang of MCNC presented the highlights of the first draft of the Remote Conferencing Architecture Outline. This architecture has a strong dependence on, and proposes extensions to, the Connection Control Protocol work being done at ISI. Among the points made during the ensuing discussion were:

- Detailed configuration definitions are a good subject for this Group to work on.
- Conference Directory Services should be based on X.500.
- We need to develop a working relationship with the Security Working Group.
- Capability for the "loosely controlled conference" (similar to the IETF broadcasts done from Boston) needs to be added to today's CCP.
- We need to define what conferencing function APIs will be needed.
- A restatement that this architecture proposes distributed conference servers.
- There was a request to "define the problem to be solved" or "what function is desired"?
- There was a request to have a tutorial on CCP in the next session.

- It was suggested that we investigate the SMPTE header being developed by video and motion picture engineers.

There was additional discussion on the proposed Working Group Charter and the "sharpening" of the proposed Charter required to gain Working Group status. A a special "prep for Wednesday" session, involving a smaller group, was scheduled on Tuesday evening, July 14.

Special Session of Tuesday, July 14

A group of twelve people came to agreement on a statement of desired function and an approach to rewording of the proposed Working Group objectives, both to be presented at the 7/15 full session. A heavy emphasis on connection management and interface detail as a deliverable were the main discussion points on the revised Working Group objectives. Eve Schooler of ISI agreed to produce charts and give the CCP tutorial on 7/15, a repeat from the San Diego BOF. There was also consensus that the Group should discuss the subject of Conference Session Dynamics as a high priority item.

Full session of July 15, 1992

The statement of desired function was presented and seemed to be well accepted as a first pass. A marked up REMCONF Working Group Goals chart was shown, with the changed emphasis highlighted. A reminder point was made that the "Focus on" chart notation means "detailed deliverable".

> Statement of Desired Function
>
> People can interactively conference, at their workstations, across the Internet.
>
> As part of the conference process, the following media can be accommodated in workstation windows:
>
> - Motion Video - Audio - Still Images - Data, Text... via shared workspace
>
> Users can easily schedule, initiate, manage and terminate the conferences. Users can easily find out about scheduled public conferences.
>
> Tightly controlled conferences provide "n to n" conferencing from 2 to approximately 15 people. Loosely controlled conferences provide 1 to n capability where n can be something less than infinity.
>
> The Remote Conferencing Architecture will encourage multi-vendor interoperability.
>
> The Remote Conferencing Architecture will define, for tightly controlled conferences, security mechanisms to control:
>
> - Who can find out about a conference. - Who can join that conference.
>
> NOTE: conference Data will be secured by whatever security mechanisms the IESG/IAB approves.

The Remote Conferencing Architecture will anticipate and allow for the exploitation of future IP network enhancements and Internet environments in order to define quality of service levels.

The Connection Control Point tutorial was given. A point was made that other approaches to connection management, if available, should be evaluated along with CCP.

A very brief discussion of the Touring Machine project at Bellcore, (presented at the San Diego BOF) was held and a request was made to include the description and API(new charts) in the minutes of the BOF.

Future of the Group

The organizational niche for the Remote Conferencing BOF has not been decided. Independently, we intend, at this time, to continue the discussions on the additions to CCP and to come to the November IETF with the first draft of the overall proposed Remote Conferencing Architecture, as called for in the proposed Working Group milestones. We would appreciate input from the Internet community on the referenced outline that's available via FTP and future expansions to it that will be made available via announcement to rem-conf@es.net.

Attendees

Sam Ayers	s.ayers@sybus.com
James Berry	beri@sandia.llnl.gov
David Bolen	db31@nis.ans.net
Robert Braden	braden@isi.edu
Scott Brim	swb@cornell.edu
Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@concert.net
Andrew Cherenson	arc@sgi.com
Robert Clements	clements@bbn.com
Tom Coppeto	tom@mit.edu
Chas DiFatta	chas@sei.cmu.edu
Jack Drescher	drescher@concert.net
Hans Eriksson	hans@sics.se
Roger Fajman	raf@cu.nih.gov
Paul Fox	pgf@cayman.com
Osten Franberg	euaokf@eua.ericsson.se
Ron Frederick	frederick@parc.xerox.com
Jerry Friesen	jafries@sandia.llnl.gov
Robert Gilligan	Bob.Gilligan@eng.sun.com
Simon Hackett	simon@internode.com.au
Robert Hagens	hagens@ans.net

Russ Hobby	rdhobby@ucdavis.edu
J. Paul Holbrook	holbrook@cic.net
Mark Hollinger	myth@ssg.lkg.dec.com
Christian Huitema	christian.huitema@sophia.inria.fr
Scott Kaplan	scott@ftp.com
David Katinsky	dmk@rutgers.edu
Nick Kawaguchi	mamster@lanai.cs.ucla.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Padma Krishnaswamy	kri@sabre.bellcore.com
Joshua Littlefield	josh@cayman.com
E. Paul Love	loveep@sdsc.edu
Carl Malamud	carl@malamud.com
Allison Mankin	mankin@gateway.mitre.org
Kim Mayton	mayton@wg.com
Donald Merritt	don@brl.mil
Paul Milazzo	milazzo@bbn.com
Robert Mines	rfm@sandia.llnl.gov
Donald Morris	morris@ucar.edu
Ari Ollikainen	ari@es.net
Larry Palmer	lp@decvax.dec.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Allan Rubens	acr@merit.edu
Eve Schooler	schooler@isi.edu
Henry Sinnreich	hsinnreich@mcimail.com
Louis Steinberg	louiss@vnet.ibm.com
Scott Stein	scotts@apple.com
Sally Tarquinio	sallyt@gateway.mitre.org
Mark Uhrmacher	maui@tabasco.lcs.mit.edu
Andrew Veitch	aveitch@bbn.com
John Wagner	jwagner@princeton.edu
Sandro Wallach	sandro@elf.com
Abel Weinrib	abel@bellcore.com
Jeff Wong	jaw@io.att.com
John Wroclawski	jtw@lcs.mit.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

Remote Conferencing Architecture

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

email: yhc@concert.net

Proposed IETF Work Group Goals

REMOTE CONFERENCING

- "End-To-End" Remote Conferencing Architecture
- Session Planning, Initiation, Maintenance, Change etc.
- Relate Audio, Video, Groupware.....
- "Tight" And "Loose" Control Conference Sessions
- Standards, Interoperability, Open Architecture
- "Incorporate" Work Of Other Groups
- H. Multicast, A/VT, Resource Reservation/Allocation
- -Commision New Work As Needed

Architecture Milestones: Outline 7/92 Complete 3/93

Remote Conferencing Architecture (Overview)

Central Issues:

How to send messages over Internet on time to all the members?
This is tough, which leads to the development of

- **Coding**
- **Multicast**
- **Resource management (end systems and networks)**
- **Fast transport protocol (Audio/Video Transport)**

How do we start a conference? What information I must know? (addresses, coding scheme used, ...)

Remote Conferencing Architecture (Overview) (cont.)

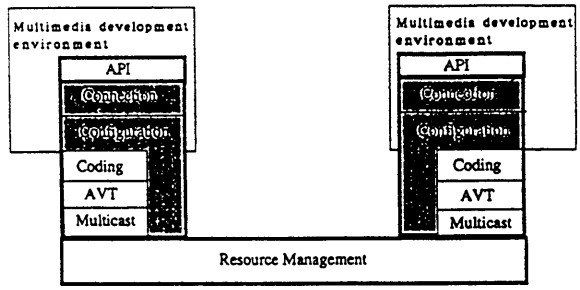

Research has been doing on coding, multicast, and resource management. The goal in this WG is to provide the connection and configuration control and interface/requirements to coding/multicast/resource management.

Interface/Requirements to Multicast, Coding, and Resource Management

Multicast:

- **Global multicast address administration**

Coding:

- **Standard coding schemes and codec configuration parameters**
- **Feedback information to codec**

Resource management

- **Admission control**
- **Resource parameters**

Remote Conferencing Styles and Their Connection Requirements

Loose Control

Example: IETF to the Internet world.

- **No detailed information is available about the members in the multicast group.**
- **Joining and leaving the group is easy.**
- **Connection setup is minimal.**
- **Configuration information must be carried out by all the packets or supported by the network architecture.**
- **Minimal security is available.**
- **It is suitable for a large group audience.**

Remote Conferencing Styles and Their Connection Requirements (cont.)

Tight Control (Eve Schooler's Connection Control Protocol)

Example: collaboration

- Members of the group are identified.
- All members must be informed when someone joins or leaves the group.
- Connection set-up can carry the configuration information and security information.
- Minimal configuration information will be carried by each packet.
- It is suitable for a small group.

Remote Conferencing Styles and Their Connection Requirements (cont.)

The Goal for Connection and Configuration Control:

Combine both loose and tight control.

Key: Flexible configuration control.

Configuration Parameters

With networks --

- Available network bandwidth
- Delay requirement
- Jitter requirement
- Multicast addresses

\- - -

With end systems --

- Coding/compression scheme (JPEG, MPEG, H.261, PCM) and associated parameters (bit rate).
- Security level (encryption).
- Media equipment (stereo/mono/cd, audio mixing)

\- - -

Suggested End System's Modules

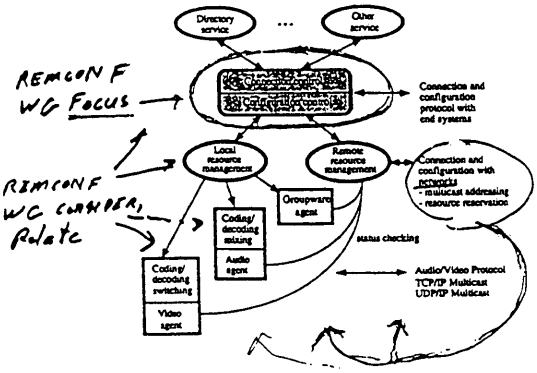

Suggested Protocols

Connection/Configuration Control Protocol:

- Extend CCP to include the configuration control and loose control style.

Transport Protocol: Audio/Video Transport WG

TCP/IP multicast for groupware.

Our Design Philosophy

- Include all the existing efforts.
- Provide the flexibility in implementation.
- Build the architecture, and use the architecture to drive the associated protocol development.
- Try to combine tight and loose control in connection control.

One Proposed Solution (Network Architecture)

Remote Conference BOF

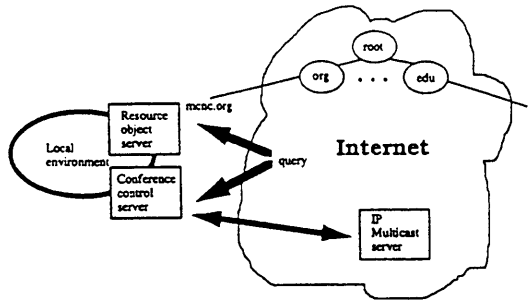

- Conference servers provides configuration information for query.
- It is easy to do the address lookup if we maintain the conference control server address as the Internet name server today.

One Proposed Solution (cont.)

Remote Conference BOF

For tight control, the conference initiator can set up the filter for the access right of the participants. Various connection style can be supported. For example,we can use the invitation style (CCP), which the conference initiator will send message to invite each participant to join. Participants can negotiate the configuration parameters with, and obtain security information from the server. Participants will input their address on the server, which will allow information exchange.

For loose control, the participants only need to get one piece information from the conference control server of the initiator -- multicast address. Then they can join the group directly. Other configuration information such as video/audio coding type, and groupware type can be obtained also from the server.

An Architecture for Multimedia Connection Management

Eve M. Schooler
Multimedia Conferencing Project
USC/ Information Sciences Institute
Marina del Rey, CA 90292

I. A Connection Management Architecture

- A framework for multiple participant, multiple media sessions
- Connection Manager (CM)
  - Central component that orchestrates connections
  - Layer below User Interfaces, but above *Media Agents*
  - Avoids duplication of effort: participation and authentication
  - Coordinates presentation of shared information
  - Facilitiates inter-media and inter-site synchronization
    - *Conduit through which control info flows (locally and remotely)*

Coordinated Management of Separate Services

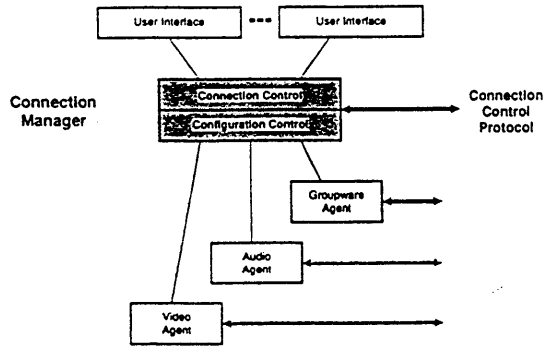

II. The Connection Control Protocol (CCP)

- An application layer protocol used by connection managers
- Includes provisions for:
  - Flexible group transaction services
  - Robustness mechanisms for WAN operation
  - Negotiation for heterogeneous site configurations
  - Conference pre-arrangement
  - Remote control capabilities
  - An interface across which timing info may be passed

A Distributed, Peer-to-peer Model

- Peer connection managers reside throughout the Internet
  - Each acts as both client and server
- Conference orchestration entails:
  - The *initiator* is designated leader for duration of setup
  - Communication with peer CMs
  - Communication with local agents
- Four-phase connection establishment procedure
  1. Negotiate a common set of *capabilities*
  2. Request others' participation
  3. Initiate underlying voice, video and groupware data flows
  4. Propagate info among peers, then revert to having no special status

Other CCP Attributes

- Others may be invited, join, or leave at any time after setup
  - Disconnection of either party during a 2-party call disconnects both
- Support for an extensible set of UI and/or media agent operations
- Detection and correction of state mismatches
  - Exchange of state info with every message
  - Trigger active state queries
  - Employ a resynchronization algorithm
- Resolution of *connection collisions*

III. Configuration Management

- As the number of WAN teleconferencing sites scales up, so does the likelihood for heterogeneity among them

Several mechanisms proposed to combat heterogeneity:

- Configuration language
- Distributed locator service
- Resource synthesizer

Dynamic Configuration Management of Shared Resources

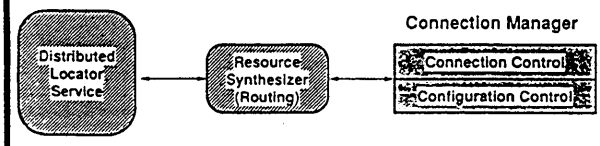

The Site Wide Operating Environment

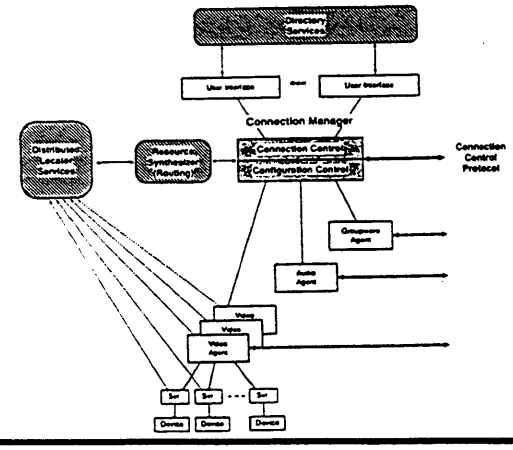

What's Missing?

- CCP is incomplete
  - For Scalability one might forego:
    - Distribution of state information
    - Full communication among conferees
    - Roving leader (initiator)
    - Reliable multicast
  - Loose-style control
  - Merge conference capabilities
  - Clean way to provide media asymmetries
- Interfaces with UI and media agents
- A detailed configuration language

In My Opinion...

This effort is not mature enough for standardization

- Other architecture schemes MUST be examined

For More Information

- Multimedia Conferencing (MMC) documents available on-line
- Anonymous FTP from venera.isi.edu in the /pub directory
- CCP-related papers:
  - *mmc-ccp-arch.ps* (v1.0 ♦ v1.1)
  - *mmc-ccp-spec.ps* (v1.1 ♦ v1.2)
- Other MMC papers:
  - *mmc-**

2.1.1 Automated Internet Mailing List Services (list)

Charter

Chair(s):

David Lippke, lippke@utdallas.edu

Mailing Lists:

General Discussion: ietf-list-wg@utdallas.edu
To Subscribe: ietf-list-wg-request@utdallas.edu
Archive: pub/ietf-list-wg@ftp.utdallas.edu

Description of Working Group:

This Working Group will concern itself with "list servers", i.e., advanced mail exploders/reflectors which provide services such as automated subscription, archive maintenance, and coordination with similar systems on the network.

The Group will initially focus its activities towards establishing a baseline user interface. Although most current systems support a command set patterned after Eric Thomas' BITNET LISTSERV, there is wide variance in the options supported and in the general patterns of interaction. This results in a great deal of user confusion. The Working Group's interface definition will address this by establishing a set of commands, options, interactions, and procedures which will (hopefully) be supported by all list servers as a subset of their full repertoire.

As a part of the user interface work, the Group will also define an authentication service for users' list server transactions. Toward this end, and to address the privacy issue, the Group will consult with the Security Area Advisory Group (SAAG).

The second phase of the Group's work will be to provide for the interconnection and coordination of list servers. Experience with the BITNET LISTSERV has shown that it is important for users be able to view the collection of list servers on the network as an integrated whole. Ideally, users should only have to deal with their local mailing list service—which knows where all public lists are, what they are, and is able to act on the user's behalf with respect to them. Interconnecting list servers allows this "integrated user view" to be created and also lets issues such as traffic minimization, timely distribution, and load sharing to be more easily addressed. Consequently, the Working Group will define the conceptual models, communication methods, and extensions to prior work which are necessary to bring this interconnection and coordination about.

It is anticipated that further work on issues of authentication and privacy will continue in parallel with the "integration" effort — perhaps manifesting itself as a separate RFC which extends the user interface definition produced during the first phase.

Goals and Milestones:

Done	Review the Group's Charter and begin work on the user interface definition.
Nov 1991	Resolve outstanding issues with the user interface definition and prepare document for IESG submission. Begin work to address the interconnection/coordination issue.
Jan 1992	Submit user interface definition document to IESG as a Proposed Standard.
Mar 1992	Focus the interconnection/coordination work. Finalize and document settled issues.
TBD	Submit interconnection/coordination definition document to the IESG for publication as a Proposed Standard.

2.1.2 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 interoperability. This discussion will be held by email.
Done	Discuss the elimination of the 7 bit restrictions in SMTP, and the implications of removing this restriction in terms of interoperation.
Done	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.
Done	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 Messages", 07/10/1991, John Klensin <draft-ietf-smtpext-8bittransport-06.txt>

CURRENT MEETING REPORT

Reported by John Klensin/MIT

Minutes of the Internet Mail Extensions Working Group (SMTPEXT)

A copy of the working draft was published as an Internet Draft (draft-ietf-smtpext-8bittransport-05.txt) at the end of June, following an earlier version published about a month earlier. These two versions were of the character of "tying up the loose ends", since most significant issues had been resolved by the close of the San Diego meeting or in list discussion shortly thereafter. There was no discussion on the list between the time that draft was announced and the time of the Working Group meeting that would have implied protocol changes; the limited discussion that did occur focused on explanatory and specificity improvements to the document text.

The Working Group meeting itself consequently was brief and quite focused, resolving the few remaining outstanding issues (about which there had been little disagreement and substantially no discussion), and then agreeing to recommend that IESG recommend a revised document as a Proposed Standard.

The Working Group session also generated some informal discussions that led to further specific language in the document and some clarified features. A revised document version was prepared after the meeting and made available to Working Group participants both at the IETF and on the list, and comments on it resulted in some additional minor changes. Specific Issues addressed and resolved included:

- Format and keywords for additional trace field information. After discussion of interactions with MIME body parts, the conclusion was to leave the level of detail at that specified prior to the San Diego meeting, relying on additional MIME headers to document per-body-part transformations. In summary, the trace information inserted by the transport in the message headers will document that a MIME transformation occurred and the specific changes made to individual body parts should be documented with those body parts. The Working Group strongly recommends that syntax, semantics, and requirements for the per-body-part audit documentation be added to MIME in the process of its going to draft standard.

- Some additional tracing keywords were added to permit documenting the cases in which a transport agent or gateway performed a conversion to make an invalid message or address form valid. Tracing these activities may make it possible to identify and fix some of the historically-most-difficult problems with electronic mail.

- Agreement was reached on additional clarification of the relationship of EHLO to commands and keywords not specified as part of either this enhanced protocol or RFC821. In summary, "old" (RFC821-only) implementations are not expected to support EHLO at all, nor are they retroactively bound by any of the specific provisions

of the enhanced protocol (although they are strongly encouraged to start registering keywords). Implementations that support the enhanced protocol and, hence, EHLO, must return keywords for all of the non-experimental commands that they provide, and all of those keywords must be registered. All experimental commands must start in "X"; no keywords will be registered or otherwise specified that start in "X".

- The requirement the EHLO return a LIMIT line (permitted message size information) was reaffirmed and explicitly documented.

A new Internet Draft was submitted during the IETF meeting and has been published as draft-ietf-smtpext-8bittransport-06.txt. The Working Group recommends that the content of this draft be published as an RFC with "Proposed Standard" status.

This concludes the present phase of the Working Group's work. Closing out the document at this point defers action on several outline proposals, discussed in the San Diego Minutes but never acted upon or proposed in any detail, for future efforts as the need arises.

Attendees

Robert Austein	sra@epilogue.com
Mark Baushke	mdb@cisco.com
Alan Clegg	abc@concert.net
James Conklin	jbc@bitnic.educom.edu
Ned Freed	ned@innosoft.com
Tony Genovese	genovese@nersc.gov
Paul Hill	pbh@mit.edu
Todd Kaehler	kaehler@zk3.dec.com
Neil Katin	katin@eng.sun.com
John Klensin	klensin@infoods.mit.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Marjo Mercado	marjo@cup.hp.com
Keith Moore	moore@cs.utk.edu
Hank Nussbacher	hank@vm.tau.ac.il
Michael Patton	map@lcs.mit.edu
John Payne	jop@wang.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Richard Schmalgemeier	rgs@merit.edu
Jane Smith	jds@jazz.concert.net
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Wagner	jwagner@princeton.edu

2.1.3 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.
Done	Review and finalize the draft document.
Done	Submit the document as a Proposed Standard.

Internet Drafts:

"Japanese Character Encoding for Internet Messages", 08/25/1992, Jun Murai, Mark Crispin, Erik van der Poel <draft-ietf-822ext-iso2022jp-01.txt>

Request For Comments:

RFC 1341	"MIME (Multipurpose Internet Mail Extensions): Mechanisms for Specifying and Describing the Format of Internet Message Bodies"
RFC 1342	"Representation of Non-ASCII Text in Internet Message Headers"

2.1.4 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 into 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
Done	First draft to be completed.
Done	Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting.
Done	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-03.txt>

"Network Database Implementation Information Internet Draft", 12/16/1991, Daisy Shen <draft-ietf-netdata-implement-02.txt>

CURRENT MEETING REPORT

Reported by Daisy Shen/IBM

Minutes of the Network Database Working Group (NETDATA)

The Netdata Working Group, Chaired by Daisy Shen met in Boston. The Working Group met on Tuesday morning and Thursday afternoon. Both sessions were well attended. The purpose of these sessions was to revise the current Internet Drafts and discuss other vendors' efforts.

During the first session, Daisy Shen presented the first half of Netdata Internet Drafts in details, and led the discussion on various issues that are related to the netdata drafts. Netdata members made some changes on the drafts. Daisy Shen will revise the drafts according to the suggestions. It will be reflected in the next version. Scott Newman presented a high level view of the effort of SQL Access Group's work.

During the second session, Daisy Shen finished presenting the drafts and did a demo which was implemented according to the drafts. Both John Wagner and Hank Nussbacher volunteered to go over the current drafts and make them look more like Internet drafts. Members have set a criteria for the protocol. They are

1. Standard SQL
2. Local & remote transparency
3. Security
4. Authentication
5. Data Compression
6. Specification Completeness
7. Integrity
8. Recovery

Due to the limited time during the meeting sessions, all members agreed to have discussion over the netdata mailing list. The topics are:.

- Unit of work.
- Security for data.
- Mapping.
- Pros and cons of Netdata's drafts and the SQL Access Group's work.
- What's missing.
- Standardize error messages.
- Adapt the pros and put them together.
- Any issue related to Network Database.
- Join meetings between IETF and the SQL Access Group.

Attendees

John Batzer	
Joel Berson	berson@brake.enet.dec.com
Atul Garg	agarg@synoptics.com
Russ Hobby	rdhobby@ucdavis.edu
Scott Newman	newman@broke.enet.dec.com
Hank Nussbacher	hank@vm.tau.ac.il
Richard Schmalgemeier	rgs@merit.edu
Vincent Sgro	sgro@cs.rutgers.edu
Daisy Shen	daisy@watson.ibm.com
Henry Sinnreich	hsinnreich@mcimail.com
Chuck Townsend	townsend@ctron.com
John Vollbrecht	jrv@merit.edu
John Wagner	jwagner@princeton.edu

2.1.5 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
Done	Continue work on definition of issues and protocols. Work to be conducted on mailing list.
Done	First draft of RFC to be completed. To be discussed at IETF meeting and revised as necessary.

Done 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.

Request For Comments:

RFC 1314 "A File Format for the Exchange of Images in the Internet"

2.1.6 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 initial list of changes to be considered are the following:

(1) User level and site designated authentication methods; (2) Binary transfer capability; (3) Minimization of line turnaround; and (4) 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.
Done	Submit Internet Draft for review and comment.
Done	Possibly meet at USENIX for further comment.
Done	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>

2.1.7 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.

2.1.8 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.

(1) Re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol.

(2) 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

(3) 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-04.txt>

"Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-04.txt>

"Telnet Environment Option", 03/03/1992, D. Borman <draft-ietf-telnet-environment-03.txt>

"Telnet Authentication: Kerberos Version 4", 03/03/1992, D. Borman <draft-ietf-telnet-authker-v4-01.txt>

"Telnet Authentication: Kerberos Version 5", 03/03/1992, D. Borman <draft-ietf-telnet-authker-v5-00.txt>

"Telnet Authentication : SPX", 07/09/1992, Kannan Alagappan <draft-ietf-telnet-authspx-00.txt>

Request For Comments:

RFC 1116 "Telnet Linemode option"

RFC 1184 "Telnet Linemode Option"

CURRENT MEETING REPORT

Reported by Steve Alexander/INTERACTIVE Systems

Minutes of the TELNET Working Group (TELNET)

The Telnet Working Group met on July 14th in Cambridge. The primary topic of discussion was (once again) encryption. A compromise position was reached in which the following will occur:

> We will move forward with folding encryption into the authentication option as discussed in San Diego. This document will be put forth as a Proposed Standard. The older encryption document will be put forth as an Experimental RFC when the newer one is available. The older one is not felt to be secure in the face of active attacks, and therefore should not be "blessed."

In other business, the Group will recommend that Environment Option be put forth as a Proposed Standard, and that Authentication Option, Kerberos IV Authentication Suboption, and SPX Authentication Suboption be put out as Experimental RFCs. The Kerberos V Authentication Suboption is still being revised by Ted Ts'o, and will also be put out as an Experimental RFC when done.

Attendees

Steve Alexander	stevea@i88.isc.com
Robert Austein	sra@epilogue.com
Mark Baushke	mdb@cisco.com
David Borman	dab@cray.com
Geetha Brown	geetha@decvax.dec.com
Michael DeAddio	deaddio@thumper.bellcore.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Roger Fajman	raf@cu.nih.gov
Robert Gilligan	Bob.Gilligan@eng.sun.com
Neil Haller	nmh@thumper.bellcore.com
John Linn	linn@erlang.enet.dec.com
Kent Malave	kent@chang.austin.ibm.com
Marjo Mercado	marjo@cup.hp.com
Clifford Neuman	bcn@isi.edu
Jeffrey Schiller	jis@mit.edu
Jeremy Siegel	jzs@nsd.3com.com
Rajesh Srivastava	
Theodore Ts'o	tytso@mit.edu

2.2 Internet Area

Director(s):

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

Area Summary reported by Greg Vaudreuil/CNRI

Below are summaries of the Internet Area Working Groups which have made notable progress since the last IETF meeting in San Diego.

Dynamic Host Configuration Working Group (DHC)

The Dynamic Host Configuration Working Group did not meet this week. They are essentially finished with their work and are expected to submit a set of documents to the IESG in the next few weeks.

IP over Asynchronous Transfer Mode Working Group (ATM)

Things are really interesting as the IP over ATM Working Group debates the various proposals doing multi-protocol networking over ATM. It is not clear what the better proposals are and the Working Group is hard at work trying to make progress.

Point to Point Protocol Extensions Working Group (PPPEXT)

The PPP Extensions Working Group made a new IETF record for the most productive session. Because of good attendance, several long-standing log jams were broken and several documentes are expected. They discussed and made progress on data compression, checksum negotiations, and authentication.

2.2.1 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 ro
Done	We will summarize those mechanisms already in place for managing the information identified by Objective 1.
Done	We will suggest new mechanisms to manage the information identified by Objective 1.
Done	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 c
TBD	Write a bootp extensions document.

Internet Drafts:

"Clarifications and Extensions for the Bootstrap Protocol", 05/03/1991, Walt Wimer <draft-ietf-dhc-bootp-01.txt>

"Dynamic Host Configuration Protocol", 07/09/1991, R. Droms <draft-ietf-dhc-protocol-04.txt, .ps>

"DHCP Options", 06/30/1992, R. Droms <draft-ietf-dhc-options-01.txt>

"Interoperation Between DHCP and BOOTP", 06/30/1992, R. Droms <draft-ietf-dhc-between-bootp-01.txt>

2.2.2 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 Internet Packets Over AppleTalk Networks", 03/08/1991, John Veizades <draft-ietf-appleip-MacIP-01.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)

A proposal was made to limit the discussion to only old items, new business should be brought up either in future BOFs or when a specific Working Group is formed to discuss the issue.

MacIP

- Status and comments on the Evans-Ranch Document.
- Could the Thursby DecNET model be used for MacIP-2?
- Specific protocol description without any implementation details.
- Implementation hints should be in the appendix of the document.
- Goal for MacIP 2 to make customers happy.
- Is there a desire or need to do anything more than document the current working model?
- Time line for completion of this effort Proposed Standard by November IETF to the Working Group Chair.

IP over Localtalk

IP over the localtalk Chris Ranch (Novell) wants to experiment with it.

General AppleTalk MIB

- Discussion of DDP forwarding table and an understanding of the next hop gateway.
- Add Diffs to RFC 1243.

AFP

- Little to say.
- Posted to pacer ftp in the next few weeks with differences from the previous specification.

MacIP MIB

Presented and will be published to the list for comments.

SNMP

Status of SNMP over DDP

- Being discussed in the Multiprotocol SNMP Group.
- Going to Proposed Standard.

AURP

- Status of implementations.
- Document to be published as an informational RFC.
- Added Device hiding, ZIP storm, Backup paths, etc.
- Implementation notes to be added to document.
- Roll out at WWDC, InterOp DC, MacTivity, InterOp SF.
- Shiva, Novell, Cayman and Apple are interoperating.
- Contact Craig Brenner or Alan O at Apple to interoperate at InterOp SF.

PPP

- Really close to being a 100 percent done to be proposed as a Standard.
- Interoperability ?? Cayman has some of this done.

Greg Burell and Routers

Codification of the bringing up of a router.

AppleTalk Routing Issues

What's next? Apple wants to come up with an advance routing protocol (local). Should have scalability, load balancing, type of service routing, no hop count limit, multicast routing. Link state vs. distance vector? Based on a Standard with a transition strategy. Apple will work independently, with Standards groups and with individual organizations. No commitment to use what comes out of the Standards Group.

Working Group Status- Greg Minshall

- NO forum for the discussion of AppleTalk issues in a public forum.
- How does Apple fit into this structure. How does the IETF fit into this.
- We have a Group of Working Groups that could take over this work. Have to present this to the IESG.

Charters and presentation to the IESG by November IETF

- NO decision making apparatus.
- Need to meet with Phill Gross and Vint Cerf.

Noel and AppleTalk in the IETF

- History of the AppleTalk proposed at the Santa Fe IETF and was discussed with Apple.
- Where is AppleTalk going in the long run... Do we always maintain a completely separate protocol stack?
- Proposed to Vint - do we move to a common layer three or even four?
- There must be a statement of where AppleTalk is going in the long run to get buy in from the IESG.
- Protocols that relate closely to the IETF sphere of protocols then the protocol can fall under the IETF-AppleTalk and IS to IS falls far a field of this.

Attendees

Atul Bansal	bansal@wile.enet.dec.com
James Beers	beers@nr-tech.cit.cornell.edu
Rich Brown	Richard.E.Brown@Dartmouth.edu
Gregory Bruell	gob@wellfleet.com
Philip Budne	phil@shiva.com
Cyrus Chow	cchow@orion.arc.nasa.gov
Tim Cleveland	tac@lkg.dec.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Bob Jeckell	rrj@3com.com
Miriam Kadansky	mckadansky@eng.xyplex.com
Hock-Koon Lim	lim@po.cwru.edu
Joshua Littlefield	josh@cayman.com
Greg Merrell	merrell@caldec.enet.dec.com
Beth Miaoulis	beth@cayman.com
Greg Minshall	minshall@wc.novell.com
Alan Oppenheimer	oppenheime1@applelink.apple.com
J. Bradford Parker	brad@cayman.com
Christopher Ranch	cranch@novell.com
Tom Reuther	reuther@nac.dec.com
Mike Ritter	mwritter@applelink.apple.com

Sam Roberts	sroberts@farallon.com
Guenter Roeck	roeck@conware.de
Frank Slaughter	fgs@shiva.com
Timon Sloane	peernet!timon@uunet.uu.net
Evan Solley	evan@farallon.com
Brad Steinka	brad@microcom.com
David Stine	dstine@cisco.com
Wayne Tackabury	wft@pacersoft.com
John Veizades	veizades@apple.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jesse Walker	walker@eider.lkg.dec.com
Jonathan Wenocur	jhw@shiva.com

2.2.3 IP over Asynchronous Transfer Mode (atm)

Charter

Chair(s):

Robert Hinden, hinden@eng.sun.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.

Done Circulate drafts of IP over ATM Specifications.

Mar 1992 Review approaches to running IP over ATM.

Internet Drafts:

"Multiprotocol Interconnect over ATM Adaptation Layer 5", 06/12/1992, Juha Heinanen <draft-ietf-atm-multipro-02.txt>

CURRENT MEETING REPORT

Reported by Bob Hinden/Sun

Minutes of the IP over Asynchronous Transfer Mode Working Group (ATM)

Agenda

- Review Agenda.
- Discussion of "Multiprotocol Interconnect over ATM" Internet Draft.
- Status of other ATM Groups.
- MAC Layer Proposal by John Burnett.
- Presentation on "Network Layer Architecture for ATM Networks" by Fong Liaw, Sun.

The first half of the meeting was spent discussing the "Multiprotocol Interconnect over ATM" Internet Draft written by Juha Heinanen. The document describes three approaches to encapsulating datagrams in ATM. These can be divided into two classes: Vitual Circuit (VT) Based Encapsulation, and Multiplexing Encapsulation.

There is general agreement on the VC Based encapsulation method. In this approach, one protocol per VC, the VC identifier is used to identify the protocol being carried.

The two approaches proposed for the multiplexing are roughly equivalent. They differ in the manner used to identify the protocols being encapsulated. One uses NLPID (from Frame Relay) and the other uses LLC's (from 802.x LAN). The Group was not able to agree on which was preferred. There was some agreement that it would be better to select only one and that it was not good for uses of ATM to have to choose between the two, but no consensus was reached. The Group agreed to change the name of the document to "Multiprotocol Interconnect over ATM using AAL5" because it better describes the method chosen to do the encapsulation. A separate document may be written that describes doing encapsulation using other ATM adaptation layer protocols.

Dan Grossman gave a presentation on the recent CCITT meeting. He reported that the differences between AAL3 and AAL4 have been resolved. The have been merged to from AAL3/4. AAL5 specification has been completed and all major technical issues have been resolved.

Doug Hunt gave a presentation on the status of the ATM Forum. The forum completed its first major Document "UNI Specification". They have begun working on signaling for switched VC's and traffic management.

Fong Liaw gave a presentation titled "Network layer Architecture for ATM Networks". It addressed issues which have not been resolved in private ATM network architectures. The topics included addressing, routing, multicast, signaling and connection management, and connection-less services.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
David Arneson	arneson@ctron.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Sam Ayers	s.ayers@sybus.com
Ashar Aziz	ashar.aziz@eng.sun.com
Douglas Bagnall	deb@proteon.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
Tony Bates	tony@ean-relay.ac.uk
Ken Benstead	kbenstead@coral.com
C. Douglas Brown	cdbrown@sandia.gov
Caralyn Brown	cbrown@wellfleet.com
Alan Bryenton	bryenton@bnr.ca
John Burnett	jlb@adaptive.com
Robert Candela	candela@sceng.ub.com
James Carlson	carlson@xylogics.com
John Chang	jrc@uswest.com
Dilip Chatwani	dilip@synoptics.com
Cheng Chen	ctc@huey.bellcore.com
Dean Cheng	dean@sun2.retix.com
Isaac Choi	ichoi@synoptics.com
George Clapp	clapp@ameris.center.il.ameritech.com
Henry Clark	henryc@oar.net
James Cobban	jcobban@bnr.ca
Rob Coltun	rcoltun@ni.umd.edu
Jim Comen	comenj@interlan.interlan.com
Michael Conn	0043874510@mcimail.com
Eric Cooper	ecc@fore.com
Tracy Cox	tacox@sabre.bellcore.com
Michael DeAddio	deaddio@thumper.bellcore.com
Michael DiGioia	digioia@interlan.interlan.com
Hans Dittler	dittler@conware.de
Jack Drescher	drescher@concert.net
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@mrbill.canet.ca
James Forster	forster@cisco.com
Craig Fox	foxcj@network.com
Osten Franberg	euaokf@eua.ericsson.se
Cliff Frost	cliff@cmsa.berkeley.edu
Kathryn Fuller	kathryn@codex.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Atul Garg	agarg@synoptics.com
Elise Gerich	epg@merit.edu

David Getchell	getchell@cisco.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
William Haggerty	haggerty@ctron.com
William Hawe	hawe@erlang.enet.dec.com
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Ittai Hershman	ittai@nis.ans.net
Robert Hinden	hinden@eng.sun.com
Doug Hunt	dhunt@bbn.com
David Husak	dave@synnet.com
J. Lance Jackson	jackson@noc.harvard.edu
George Kajos	kajos@coral.com
Nick Kawaguchi	mamster@lanai.cs.ucla.edu
Kenneth Key	key@cs.utk.edu
Deidre Kostick	dck2@sabre.bellcore.com
Mark Lewis	mlewis@telebit.com
Fong-Ching Liaw	fong.liaw@eng.sun.com
Arthur Lin	yalin@srv.pacbell.com
Ron Mackey	rem@dsiinc.com
Andrew Malis	malis@bbn.com
Allison Mankin	mankin@gateway.mitre.org
Bill Manning	bmanning@rice.edu
Keith McCloghrie	kzm@hls.com
Paul Moran	Paul_Moran@3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Jack Ng	jkn@network.com
Luc Nguyen	luc@bellsouth.com
Erik Nordmark	nordmark@eng.sun.com
Ari Ollikainen	ari@es.net
Roger Osmond	bytex!rfo@uunet.uu.net
Philippe Park	ppark@bbn.com
Ed Perreault	ed@python.eng.microcom.com
Roy Perry	rperry@uswat.uswest.com
Eric Peterson	elpeterson@eng.xyplex.com
Venkat Prasad	vsp@3com.com
Russell Pretty	pretty@bnr.ca
Scott Quinn	smq@aries.bellcore.com
K. K. Ramakrishnan	rama@erlang.enet.dec.com
Richard Ramos	ramos@mtunm.att.com
David Reiser	dsr@codex.prds.cdx.mot.com
Salvatore Ricci	ricci@mtunm.att.com
Anil Rijsinghani	anil@levers.enet.dec.com
Manoel Rodrigues	manoel_rodrigues@att.com
Christine Ross	ross@warp.interlan.com
Allan Rubens	acr@merit.edu

Anthony Rutkowski	amr@sprint.com
Timothy Salo	tjs@msc.edu
Jonathan Sandberg	Jss@mitl.com
Miguel Sasson	sasson@xylogics.com
John Sax	jsax@bbn.com
Hellen Sears	sears@interlan.interlan.com
Henry Sinnreich	hsinnreich@mcimail.com
Keith Sklower	sklower@cs.berkeley.edu
Lansing Sloan	ljsloan@llnl.gov
Andrew Smith	asmith@synoptics.com
Frank Solensky	solensky@clearpoint.com
Rajesh Srivastava	rajesh@flash.tandem.com
Brad Steinka	brad@microcom.com
Scott Stein	scotts@apple.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Terrance Sullivan	terrys@newbridge.com
George Swallow	gswallow@bbn.com
Bruce Taber	taber@interlan.com
Marten Terpstra	terpstra@ripe.net
Linda Tom	toml@interlan.interlan.com
Sylvia Tom	tom_sylvia@tandem.com
Stephen Tsun	snt@nsd.3com.com
Ahmet Tuncay	atuncay@synoptics.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Ross Veach	rrv@uiuc.edu
Justin Walker	justin@apple.com
James Watt	jamesw@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Liming Wei	liming.wei@eng.sun.com
Jehu Westmark	westmark@ds1.scri.fsu.edu
Chris Wheeler	cwheeler@cac.washington.edu
Rick Wilder	wilder@ans.net
Cathy Wittbrodt	cjw@nersc.gov
Jeff Wong	jaw@io.att.com
Robert Woodburn	woody@cseic.saic.com
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com
Joseph Zur	zur@fibhaifa.com

IP over ATM

Working Group

Robert Hinden

July 13, 1992

AGENDA FOR WEEK

- Monday
  - Review Agenda
  - Discussion of *Multiprotocol Interconnect over ATM* Internet Draft, Juha Heinanen
- Tuesday
  - Continuation of Monday Discussion
  - Status of other ATM Groups
- Wednesday
  - *Network Layer Architecture for ATM Networks* Presentation, F. Liaw, Sun
  - Discussion

ENCAPSULATION PROPOSAL SUMMARY

- VC Based Encapsulation
  - One Protocol per VC
- Multiplexing Approaches
  - NLPID Approach
  - LLC Approach
- No Agreement on Multiplexing Approaches
  - Desire to be compatible with Divergent Networking Technologies
  - Support Various forms of Bridging (a.k.a. Interworking)

COMPARISON OF MULTIPLEXING APPROACHES

IP

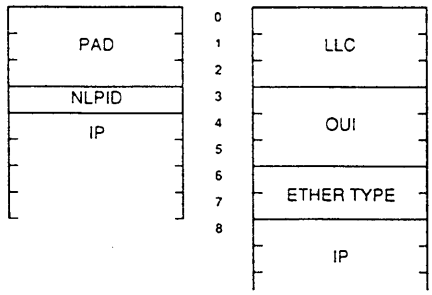

COMPARISON OF MULTIPLEXING APPROACHES

CLNP

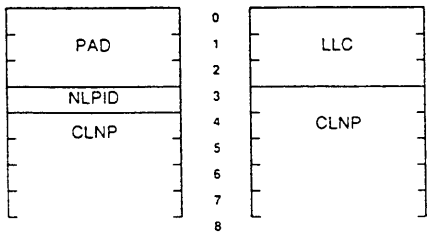

NEW PROPOSAL

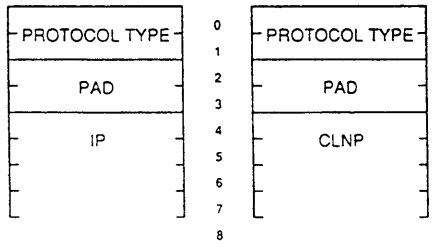

24th Meeting
Internet Engineering Taskforce

IP over ATM Working Group

Standards Status

Dan Grossman
Motorola Codex
20 Cabot Blvd, M/S C2-15
Mansfield, MA 02048

dan@merlin.dev.cdx.mot.com

1992 CCITT Recommendations

I.113 Vocabulary of terms for broadband aspects of ISDN
I.121 Broadband Aspects of ISDN
I.150 B-ISDN Asynchronous Transfer Mode Functional Characteristics
I.211 B-ISDN Service Aspects
I.311 B-ISDN General Network Aspects
I.321 B-ISDN Protocol reference model and its applications
I.327 B-ISDN Functional architecture
I.361 B-ISDN ATM Layer Specification
I.362 B-ISDN ATM Adaption Layer (AAL) Functional Description
I.363 B-ISDN ATM Adaption Layer (AAL) Specification
I.364 Support of Broadband Connectionless Data Service on B-ISDN
I.371 Traffic control and Congestion Control in B-ISDN
I.414 B-ISDN User-network Interface
I.432 B-ISDN UNI - Physical Layer Specification
I.610 B-ISDN Operation and Maintenance Principles and Functions

AAL Related Issues (I.362, I.363. I.364)

- All substantive differences between AAL3 and AAL4 protocols resolved; merged to form AAL3/4
- AAL5 specification drafted and attached to meeting report; new §6/I/363 added, with content 'currently being studies' (signal to industry)
- all but one substantive technical issue in AAL5 resolved; remaining point tenuous
- selection of an AAL for signalling (3/4 or 5) deferred
- text for Service Specific Convergence Sublayer drafted

AAL Related Issues (I.362, I.363. I.364) (cont.)

- I.364 (ex-I.cls) approved; aligned with 802.6 and SMDS, but several items FFS (carrier selection, QoS selection, encoding of QoS field, subaddress field, use of header extension, use of CRC-32 field by network)
- AAL 1 approved

ATM Layer issues (I.150, I.361)

- Closure on bidirectional VPI/VCI assignment
- Point-to-point signalling virtual channel assigned
- Standoff continues on GFC (DQDB based vs Orwell Ring based)
- Discussion of performance of ATM in errored environment

Traffic Management (I.371)

- I.371 approved for plenary Assembly
- Definition of peak cell rate - inverse of minimum intercell time ± CDV tolerance
- source traffic descriptors for high and low priority traffic apply to CLP = 0 and ((CLP = 0) + (CLP = 1) traffic, respectivel
- agreements on treatment of virtual path connections
- definition of traffic shaping broadened

Network Layer Architecture for ATM Networks

Sun Microsystems
July 15, 1992

Overview

- Multiple Protocol Network Addressing
- Datagram Overlay Network
- Standard Dynamic Topology Discovery Protocol

Addressing

Problem

- Need addressing scheme to route packets or connections through ATM networks

Network Layer Address is Appropriate

- No need for link layer address

New Addressing is unnecessary and inconvenient

- More administrative overhead
- Already exits good network addressing schemes
- Need address translation
- Direct ATM reachability

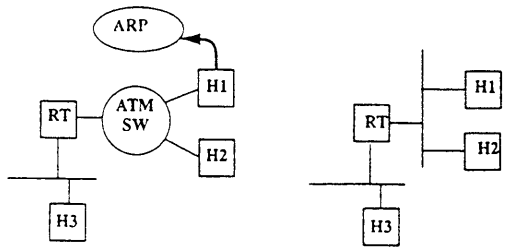

Example of ATM Address Translation (Resolution)

Addressing (cont.)

Multiple protocol ATM Switches/Routers

- Functionality
  - Similar to today's **multi-protocol routers**
  - Accept multiple addressing schemes
  - Choose a route through the network
  - Determine network topology
  - Establish ATM VC connections
- Advantages
  - No address translation
  - No need to determine direct ATM reachability
  - Transparent to non-ATM networks
  - Supporting advanced network functionality
  - Code reuse

Public Network Compatibility

Problem

- Exclusive E.164 address in public network

Implementation

- In the Local, application's IP in called address, and null sub-address
- In the public, "exit" point E.164 in the called address, and application's IP in the sub-address
- Subaddress uses NSAP format

Public Network Compatibility (cont.)

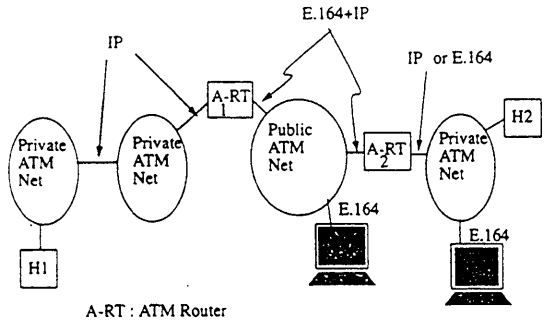

A-RT : ATM Router

Datagram Overlay Network

- We need both datagram and connections !
  - Provides more flexibility
  - Satisfy most application's needs
- Why datagram services ?
  - Server location
  - Keep-alive packets.
  - Single request-response applications
- Why connections ?
  - Intermediate node pkt reassemble increases delay
  - Real-time applications need low delay
  - Large volume traffic such as FTP

Datagram Overlay Network (cont.)

- Datagram needs to find "next hop" to forward packets
- Connection needs to find "next hop" to route connections
- Share the same routing table

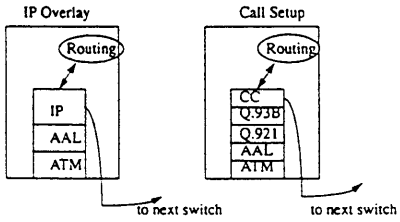

Datagram Overlay Network (cont.)

Examples

- Pure datagram
  - BOOTP, ICMP, ...
- Pure connection
  - Video, Audio,...

Datagram Overlay Network (cont.)

Fast Setup

- Send Data before the setup confirmation
  - Requires switches to buffer the cells
  - Lost/Corrupted signaling cells cause black hole

An alternative, while you wait, use datagram

- Send to datagram overlay network before the setup confirmation
- Send data to the connection after the setup confirmation

Heuristic setup (traffic monitoring)

- Send to datagram overlay network
- Set up connection when load warrants the connection

Switch to Switch Interface

- Why ?
  - Multi-vendor switching fabric
  - Plug in and play
- Switch-to-Switch Interface
  - Standard Connection management protocol
  - Standard dynamic topology discovery protocol

Switch to switch interface (cont.)

Why dynamic topology discovery protocol ?

- Need to route around failure in **real-time**
- End system mobility

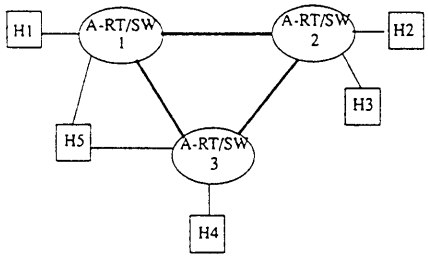

A-RT/SW : ATM Router/Switch

Why standard protocol ?

- The only way to achieve efficient routing in multiple multivendor ATM LAN environment

Which protocol ?

- OSPF, IS-IS, and so on ...

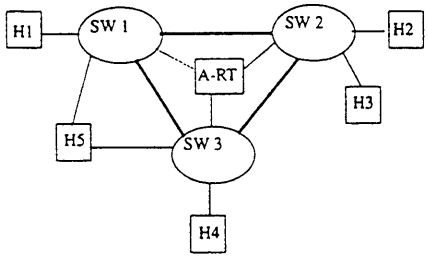

A-RT : ATM Router

Summary

- ATM network as an optimized packet switched network
- Switch as an optimized multi-protocol router
- VC as a cached route
- Datagram overlay network as an implementation convinience of connection establishment
- Standard dynamic topology discovery protocol as a vehicle to efficient (inter-) networking of ATM LANs.

2.2.4 IP over FDDI (fddi)

Charter

Chair(s):

Dave Katz, dkatz@cisco.com

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.

Internet Drafts:

"Transmission of IP and ARP over FDDI Networks", 09/14/1992, D. Katz <draft-ietf-fddi-ipoverfddi-00.txt>

Request For Comments:

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

2.2.5 Multi-Media Bridging (mmb)

Charter

Chair(s):

Jeffrey Fitzgerald, jjf@fibercom.com

Mailing Lists:

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

2.2.6 Point-to-Point Protocol Extensions (pppext)

Charter

Chair(s):

Brian Lloyd, brian@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:

"The PPP DECnet Phase IV Control Protocol (DNCP)", 06/04/1991, Steven Senum <draft-ietf-pppext-decnet-03.txt>

"The PPP AppleTalk Control Protocol (ATCP))", 07/08/1991, Brad Parker <draft-ietf-pppext-appletalk-03.txt>

"The PPP OSI Network Layer Control Protocol (OSINLCP)", 07/25/1991, D. Katz <draft-ietf-pppext-osinlcp-02.txt>

"PPP Authentication Protocols", 07/25/1991, B. Lloyd, W.A. Simpson <draft-ietf-pppext-authentication-06.txt>

"IPX PPP Internetwork Packet Exchange Control Protocol [IPXCP]", 06/10/1992, Michael Allen <draft-ietf-pppext-ipxcp-01.txt>

"The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol", 06/22/1992, Frank Kastenholz <draft-ietf-pppext-ipcpmib-01.txt>

"The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol", 06/22/1992, Frank Kastenholz <draft-ietf-pppext-lcpmib-01.txt>

"The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol", 06/22/1992, Frank Kastenholz <draft-ietf-pppext-secmib-01.txt>

"The Definitions of Managed Objects for the Bridge Network Control Protocol of the Point-to-Point Protocol", 06/22/1992, Frank Kastenholz

<draft-ietf-pppext-bridgemib-01.txt>

Request For Comments:

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

RFC 1331 "The Point-to-Point Protocol (PPP) for the Transmission of Multi-protocol Datagrams over Point-to-Point Links"

RFC 1332 "The PPP Internet Protocol Control Protocol (IPCP)"

RFC 1333 "PPP Link Quality Monitoring"

CURRENT MEETING REPORT

Reported by Constance Fleenor Lloyd/Lloyd and Associates

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

Brian Lloyd opened the meeting stating that there is a small problem with with proceeding with RFC 1334 (Authentication) that he hoped to clear up this week in a meeting he had scheduled with Steve Kent of the IAB.

General

Collectively, RFCs 1331, 1332, 1333, and 1334 will completely replace RFCs 1171, and 1172.

Document Status

Appletalk, and OSI documents are currently at last call. Copies of the Appletalk document can be obtained from Brad Parker.

A show of hands was requested for how many are building routers. About twelve hands were raised, and a second show of hands indicated about nine were planning on CLNP. It was stated that these people need to read the document <draft-ietf-pppext-osinlcp-02.txt>.

TELEBIT PPP-a-thon

Mark Lewis reviewed the results of the PPP-a-thon, in which ten vendors made their various PPP options interoperate. Vendors mapped protocols to media, and implementations.

Work done at the PPP-a-thon resulted in an update to RFC 1220; clarifying negotiation of all mac types. The wording in RFC 1220 is not clear on results of mac negotiations. The proposal is to yank mac type, and LAN type from 1220, and add option to state mac address to the RFC. Fred Baker, the author of RFC 1220, is going to make the changes to the RFCs.

Another PPP-a-thon is planned for the week prior to Fall INTEROP '92. This will also serve as a hot stage for the INTEROP PPP demo/solutions showcase which is planned if Brian Lloyd can get confirmation from at least 15 vendors for participation. It is hoped that some new implementations will be available for this week.

INTEROP PPP Solutions Showcase Demo

If enough vendors sign for participation there will be a PPP demo at INTEROP. Charge for participation in the INTEROP PPP demo will be $4,000. Brian Lloyd feels that a centralized demonstration is best, rather than the distributed demo of 1990. The demo booth should be as self-contained as possible. All possible aspects of PPP would be part of

the demo; dial-up, and leased-line asynchronous; synchronous 56K to T3; switched services including dial-in ISDN, and switched 56; etc.

Enough participants present at the meeting confirmed interest in the INTEROP demo to provide the required 15 vendors.

LAPB, and Compression

Fred Baker led discussion on the subject of LAPB, and compression;

There was lively discussion on using more than one compression algorithm concurrently; and if so should one NCP identifier be used, or subprotocol identifiers.

Note: The issues were later addressed in BOFs, resulting in consensus on a complete compression configuration protocol. Fred Baker is writing a draft document that describes the compression mechanisms and negotiations.

IPX

Chris Ranch, from Novell, led discussion on IPX over PPP. Novell is publishing an information RFC on IPXWAN; available via ftp novell.com. IPXWAN will be essential for interoperability with Netware.

They will work on IPXCP options for a paper that will obsolete IPXWAN in the future.

Twelve people raised their hands when Chris asked who is doing IPXWAN with PPP.

A small group consisting of Mark Lewis, Bill Simpson, Brian Lloyd, Glenn McGregor, and Chris Ranch met with Marty DelVechio at Shiva in a BOF to discuss Bill Simpson's document that attempts to converge the Shiva and Novell NCP documents. Bill Simpson will modify his convergence document based on the results of the discussions at the BOF and repost it.

DECNET

Craig Fox led discussion on the DECNET draft. He made some changes to the explanatory passages in the draft, removed a section, and moved a section of text. The text was posted as a new draft the next day.

116/32 Bit FCS

Tony Lauck said that DEC, in compliance with standards policy, will charge a "reasonable" fee to use for DEC's 48 bit FCS. Numbers on the order of $1,000 or more were bandied about by Mr. Lauck. This was not well received by the Group. The general consensus is to seek a different mechanism to negotiate 16/32 bit FCS. Both Craig Fox and Karl Fox suggested mechanisms based on sending two initial configuration requests using both 16

and 32 bit FCS. A new document that does not make use of any of the DEC mechanisms is likely.

MIBs

Frank Kastenholz described how the single large MIB document had been divided into separate MIB documents to match with the existing LCP, NCP, LQM, and authentication documents. He asked if there were any objections to submitting MIBs for Proposed Standard. There were no significant objections.

Attendees

Steve Alexander	stevea@i88.isc.com
Sam Ayers	s.ayers@sybus.com
Fred Baker	fbaker@acc.com
Jim Barnes	barnes@xylogics.com
Luc Boulianne	lucb@cs.mcgill.ca
Steve Buchko	stevebu@newbridge.com
David Carr	dcarr@gandalf.ca
Greg Celmainis	gregc@newbridge.com
Dean Cheng	dean@sun2.retix.com
Chris Chiotasso	chris@artel.com
Henry Clark	henryc@oar.net
James Cobban	jcobban@bnr.ca
Andy Davis	andy@spider.co.uk
Marty Del Vecchio	marty@shiva.com
Pierre Dupont	dupont@mdd.comm.mot.com
Robert Finn	rfinn@wellfleet.com
Constance Fleenor Lloyd	cfleenor@lloyd.com
Craig Fox	foxcj@network.com
Scott Ginsburg	ginsburg@wellfleet.com
Martin Gray	mg@spider.co.uk
Chris Gunner	gunner@dsmail.enet.dec.com
James Harville	jbh@unh.edu
Dwight Jamieson	djamies@bnr.ca
Bob Jeckell	rrj@3com.com
Frank Kastenholz	kasten@ftp.com
David Kaufman	
Alan Kullberg	akullber@bbn.com
David Kumpf	kumpf@hprnd.rose.hp.com
Anthony Lauck	lauck@tl.enet.dec.com
Gordon Lee	gordon@ftp.com
Mark Lewis	mlewis@telebit.com
Brian Lloyd	brian@lloyd.com

Keith Mader	keith@microcom.com
Glenn McGregor	ghm@merit.edu
Greg Minshall	minshall@wc.novell.com
Mary Christine O'Connor	oconnor@interlan.com
Larry Palmer	larry%acts@uunet.uu.net
J. Bradford Parker	brad@cayman.com
Jay Pasco-Anderson	jay@merlin.dev.cdx.mot.com
Lars Poulsen	lars@cmc.com
Venkat Prasad	vsp@3com.com
Christopher Ranch	cranch@novell.com
John Shriver	jas@proton.com
William Simpson	Bill_Simpson@um.cc.umich.edu
Mark Smith	mcs@umich.edu
Chris Sullivan	mm@gandalf.ca
Richard Sweatt	rsweatt@synoptics.com
Kannan Varadhan	kannan@oar.net
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Veizades	veizades@apple.com
Jesse Walker	walker@eider.lkg.dec.com
Scott Wasson	sgwasson@eng.xyplex.com
James Watt	jamesw@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Honda Wu	natadm!honda@uunet.uu.net

2.2.7 Router Requirements (rreq)

Charter

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.
Done	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>

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

Request For Comments:

RFC 1349 "Type of Service in the Internet Protocol Suite"

RFC 1354 "IP Forwarding Table MIB"

2.3 Network Management Area

Director(s):

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

Area Summary reported by James Davin/MIT

During the Twenty-Fourth plenary meeting of the IETF, nine working groups and two Birds-of-a-Feather (BOF) sessions met in the Network Management Area.

Simple Management Protocol (SMP) Framework BOF (SMPFRAME)

The SMPFRAME BOF addressed a recent proposal for evolutionary enhancements to the SNMP network management framework. During this session, developers of the SMP (Simple Management Protocol) proposal gave a comprehensive technical presentation of their ideas, and a period of detailed discussion ensued. Several common themes emerged from that discussion:

- There was general agreement that a single transition from existing SNMP technology to the next stage of SNMP evolution is highly desirable; multi-stage or protracted transitions were generally felt to be undesirable.

- There was general agreement that minimizing the number of distinct management technologies deployed in the Internet is highly desirable.

- Most members of the community felt that the SMP proposals addressed many of the perceived problems in the current SNMP framework, although many members suggested that certain adjustments to the SMP work could increase its value and acceptance to the community.

- There was general agreement that an aggressive schedule for standardizing the next generation of SNMP technology is appropriate.

- There was general agreement that security aspects of new SNMP technologies should be considered separately from purely network management aspects, although consideration of both aspects must be coordinated carefully in terms of schedule.

Based on the community discussion during the BOF session, SMP proponents agreed to contribute their work to the process of SNMP evolution that was set in motion in March.

A tentative plan has been formulated to further that process in response to community sentiment expressed during the Cambridge IETF meeting:

- As described in the plan for SNMP evolution, detailed technical specifications which

are contributed to the SNMP evolution process will be published as Internet Drafts.

- The first phase of the evolution process (the call for contributions) will be concluded soon. Community members with detailed, written technical proposals for SNMP evolution are encouraged to contribute those proposals as Internet Drafts in the very near future or at least to inform the IESG of their intent to do so. Announcement of the closing date will be made to the IETF mailing list, and the relevant working group(s) will be not be obligated to consider contributions after that date.

- Consistent with the process set in motion in March, an open IETF working group will be chartered to consider the written contributions to the SNMP evolution process. Consistent with the community sentiment for timely progress, the first meeting of that working group will be sometime in September. As usual, the organization and charter for this working group will be announced on the IETF mailing list.

- This working group will conclude its business not later than the plenary meeting of the IETF in the spring of 1993 and its schedule will be closely coordinated with any related activity within the existing working group on SNMP security.

Uninterruptable Power Supply BOF (UPSMIB)

Among the Birds-of-a-Feather sessions that met during the week was a session on SNMP instrumentation for uninterruptible power supplies (UPSs). Approximately 8–10 UPS vendors and a number of other interested parties met to discuss the substance and format of possible work in this area. The Group reviewed a strawman MIB document (available as an Internet Draft) and concluded that a working group effort to pursue a UPS MIB is desirable.

Bridge MIB Working Group (BRIDGE)

The Bridge MIB Working Group met after a period of inactivity to consider alignment of the existing Bridge MIB work with recent IEEE work on source routing bridges. The Group discussed a revision to the existing MIB that would support identification of the protocol or protocols in use by a bridge device. The Group also decided to address the recent IEEE work on source routing by beginning work on a new MIB devoted to those functions. The new MIB would include portions of the existing Bridge MIB relevant to source routing together with any new objects that may be required for alignment with the IEEE work. The existing MIB document is expected to be considered soon for Draft Standard status, whereas the newer MIB document is expected to enter the standards track when the working group effort is complete.

Chassis MIB Working Group (CHASSIS)

The Chassis MIB Working Group met to continue discussion of a MIB that instruments collections of traditional network devices that may be comprised by a single physical package.

The Group continued its discussion of an existing Internet Draft document and addressed in particular the mapping of MIB views to managed devices within a chassis. The Group also reviewed the desirability and priority of its several work items.

DS1/DS3 MIB Working Group (TRUNKMIB)

The TRUNKMIB Working Group met to continue its discussion of revisions to the existing DS1 and DS3 MIBs as they are considered for Draft Standard status. The goal of these revisions is to reflect implementation experience with the existing versions and to align with ANSI work in this area. Revised Internet Draft documents reflecting the discussion during this meeting are expected soon.

FDDI MIB Working Group (FDDIMIB)

The FDDI MIB Working Group met to discuss alignment of the existing FDDI MIB (RFC 1285) with version 7.2 of the SMT work recently produced by ANSI. The Group began discussion of what changes were desirable to accomplish that alignment.

Host Resources MIB Working Group (HOSTMIB)

The Host Resources MIB Working Group met for the first time in Cambridge. An initial draft of a host resources MIB was discussed, and there was a consensus to adopt that draft for use as a baseline document. Discussion of the developing host resources MIB will continue at an interim Working Group meeting sometime in September. The time and place of that meeting will be announced on the Working Group mailing list.

IEEE 802.3 Hub MIB Working Group (HUBMIB)

The Hub MIB Working Group also met in Cambridge. This Group affirmed a minor change to the 802.3 Repeater MIB document (involving a change to an enumerated type) before its presentation for consideration as a Proposed Standard. The Group began its discussion of a MIB for 802.3 MAUs and decided that this new MIB will cover both DTE and repeater devices.

Internet Accounting Working Group (ACCT)

The Internet Accounting Working Group met at the Cambridge meeting to conclude its business. At the meeting, it was reported that an initial implementation of the MIB drafted by the Working Group is underway at the University of Auckland. A second implementor is being sought. The Group is seeking publication of its MIB as an experimental RFC, and verified that the MIB could be subsequently considered for the standards track if community interest and need warranted. The Group contemplated opening a discussion with members of the Router Requirements Working Group to assess interest in this work among router vendors. The Group also recommended that work be undertaken within the IETF Network

Management Area to provide common identifiers for various link-layer media and network protocols, as this would facilitate any future accounting work.

Token Ring Remote Monitoring MIB Working Group (TRMON)

The Token Ring Remote Monitoring MIB Working Group met to continue discussion of extensions to the mechanisms of RFC 1271 to support remote monitoring of IEEE 802.5 token ring LANs. Discussion will continue, and closure on these token ring extensions will be sought via email. Moreover, implementation experience with the existing RMON MIB (RFC 1271) will be discussed via email with the goal of reaching consensus on what changes may be appropriate as the RMON MIB is considered for elevation to Draft Standard status.

X.25 Management Information Base Working Group (X25MIB)

The X.25 MIB Working Group met to work in earnest on the last of the three MIBs in its Charter. The Group completed deliberation on a MIB to instrument multi-protocol over X.25 convergence functions. The Group had previously completed work on MIBs for instrumenting LAPB and the X.25 packet layers. These latter MIBs were reviewed at the Cambridge meeting and will soon be presented to the IESG for a recommendation.

CURRENT MEETING REPORT

Reported by Marshall Rose/DBC

Minutes of the Simple Management Protocol (SMP) Framework BOF (SMPFRAME)

The BOF began with a 1-1/2 hour technical presentation by the four SMP authors. This was followed by 2 hours of discussion. The outcome:

1. The SMP specification will be submitted to Internet-Drafts.

2. There was strong consensus that a working group should be formed to consider SMP as the basis for SNMP version 2. The Area Director was asked to develop a timeline for such an activity.

3. Work on SNMP Security will be independent from, yet coordinated with, the proposed SNMP version 2 Working Group.

If you missed getting a copy of the presentation notes, the PostScript version is available via anonymous FTP:

	#1	#2	#3
host	ics.uci.edu	case.cs.utk.edu	lancaster.andrew.cmu.edu
area	mrose/isode-smp/	pub/smp/	pub/smp/
file	smp-bof.ps	smp-bof.ps	smp-bof.ps
or	smp-bof.ps.Z	smp-bof.ps.Z	smp-bof.ps.Z

You can retrieve either the .ps file (275K) or the compressed .ps file (75K). In the latter case, be sure to do a binary transfer.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
Steve Alexander	stevea@i88.isc.com
Michael Anello	mike@xlnt.com
David Arneson	arneson@ctron.com
Karl Auerbach	karl@empirical.com
Karen Auerbach	karen@epilogue.com
Robert Austein	sra@epilogue.com
Sam Ayers	s.ayers@sybus.com

Fred Baker	fbaker@acc.com
Jim Barnes	barnes@xylogics.com
William Barns	barns@gateway.mitre.org
Andy Bierman	bierman@davidsys.com
Jonathan Biggar	jon@netlabs.com
Uri Blumenthal	uri@watson.ibm.com
Natan Borshansky	natanb@turisc.timeplex.com
Steve Bostock	steveb@novell.com
David Bridgham	dab@epilogue.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
John Chang	jrc@uswest.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Richard Colella	colella@osi.ncsl.nist.gov
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Tom Coppeto	tom@mit.edu
Tracy Cox	tacox@sabre.bellcore.com
Stephen Crocker	crocker@tis.com
Dave Cullerot	cullerot@ctron.com
Cathy Cunningham	cmc@microcom.com
James Davin	jrd@ptt.lcs.mit.edu
Michael Davison	davison@cs.utk.edu
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Peter DiCamillo	Peter_DiCamillo@brown.edu
Chas DiFatta	chas@sei.cmu.edu
Michael DiGioia	digioia@interlan.interlan.com
David Engel	david@ods.com
Fred Engel	engel@concord.com
Michael Erlinger	mike@lexcel.com
Roger Fajman	raf@cu.nih.gov
Bill Fardy	fardy@ctron.com
Stefan Fassbender	stf@easi.net
Margaret Forsythe	mrf@ftp.com
Craig Fox	foxcj@network.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Kathryn Fuller	kathryn@codex.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Shawn Gallagher	gallagher@quiver.enet.dec.com
Atul Garg	agarg@synoptics.com
Doron Gartner	doron@chipcom.com
David Getchell	getchell@cisco.com
Rob Graham	robert_graham@protools.com

Pria Graves	priag@nsd.3com.com
Simon Hackett	simon@internode.com.au
William Haggerty	haggerty@ctron.com
James Halpin	halpin@turkey.sw.stratus.com
Arturo Heinsen	heinsen_a@msm.cdx.mot.com
Ittai Hershman	ittai@nis.ans.net
Jeff Hughes	jeff@col.hp.com
Erik Huizer	huizer@surfnet.nl
Robin Iddon	axon@cix.compulink.co.uk
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@ftp.com
Frank Kastenholz	kasten@ftp.com
David Kaufman	
Nick Kawaguchi	mamster@lanai.cs.ucla.edu
James Keller	j.keller@sprint.com
Mark Kepke	mak@cnd.hp.com
Kenneth Key	key@cs.utk.edu
Zbigniew Kielczewski	zbig@eicon.qc.ca
Deidre Kostick	dck2@sabre.bellcore.com
Cheryl Krupczak	cheryl@cc.gatech.edu
Walter Lazear	lazear@gateway.mitre.org
Thomas Lenggenhager	lenggenhager@switch.ch
Nina Lewis	nina@cam.unisys.com
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Arun Mahajan	axm@proteon.com
Carl Malamud	carl@malamud.com
Kent Malave	kent@chang.austin.ibm.com
Andrew Malis	malis@bbn.com
J. Scott Marcus	smarcus@bbn.com
Kim Mayton	mayton@wg.com
Cindy Mazza	
Keith McCloghrie	kzm@hls.com
Vern McGeorge	vern@hpda.cup.hp.com
Thomas McGinty	mcginty_t*corp_m@msm.cdx.mot.com
John McKenna	mckenna@ralvm12.vnet.ibm.com
Donna McMaster	mcmaster@synoptics.com
Shehzad Merchant	merchant@erg.sri.com
Beth Miaoulis	beth@cayman.com
Edwin Mier	mier@tigger.jvnc.net
Douglas Miller	dmm@telebit.com
Kenneth Miller	
David Minnich	dwm@fibercom.com
Greg Minshall	minshall@wc.novell.com
Rohit Mital	rm@protools.com

Lynn Monsanto	monsanto@sun.com
Paul Moran	Paul_Moran@3com.com
John Murray	murray@premenos.sf.ca.us
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Tom Nisbet	nisbet@tt.com
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
Roger Osmond	bytex!rfo@uunet.uu.net
Bob Page	bob.page@eng.sun.com
Andrew Patka	apatka@wellfleet.com
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
David Piscitello	dave@sabre.bellcore.com
Bill Pope	zpope@osf.org
Mary Quinn	quinn@maui.InterLan.Com
Richard Ramos	ramos@mtunm.att.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
David Reiser	dsr@codex.prds.cdx.mot.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Salvatore Ricci	ricci@mtunm.att.com
Anil Rijsinghani	anil@levers.enet.dec.com
Kary Robertson	kr@concord.com
Sam Roberts	sroberts@farallon.com
Guenter Roeck	roeck@conware.de
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jon Saperia	saperia@tcpjon.ogo.dec.com
Michael Sapich	sapich@conware.de
Michael Scanlon	scanlon@interlan.com
Jim Scott	scott@kali.enet.dec.com
Allwyn Sequeira	allwyn@nmu.ub.com
Koichiro Seto	seto@hitachi-cable.co.jp
Robert Shirey	shirey@mitre.org
Timon Sloane	peernet!timon@uunet.uu.net
Michael Slocombe	slocombe@bbn.com
Rajesh Srivastava	rajesh@flash.tandem.com
Michael St. Johns	stjohns@umd5.umd.edu
Einar Stefferud	stef@nma.com
Louis Steinberg	louiss@vnet.ibm.com
Bob Stewart	rlstewart@eng.xyplex.com
Adam Stolinski	
Chris Sullivan	mm@gandalf.ca
Wayne Tackabury	wft@pacersoft.com
Kaj Tesink	kaj@cc.bellcore.com
Rodney Thayer	rodney@world.std.com

Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Chuck Townsend	townsend@ctron.com
Stephen Tsun	snt@nsd.3com.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Desikan Venkatrangan	venkat@geoduck.matrix.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jesse Walker	walker@eider.lkg.dec.com
Ray Wasson	
James Watt	jamesw@newbridge.com
Gerard White	
Walter Wimer	walter.wimer@andrew.cmu.edu
Steven Wong	wong@took.enet.dec.com
Kiho Yum	kxy@nsd.3com.com
Uwe Zimmerman	metrix!uwe@uunet.uu.net
Joseph Zur	zur@fibhaifa.com

Simple Management Protocol (SMP) Framework
Birds-of-a-Feather

Jeffrey D. Case, SNMP Research, Inc.

Keith McCloghrie, Hughes LAN Systems, Inc.

Marshall T. Rose, Dover Beach Consulting, Inc.

Steven L. Waldbusser Carnegie Mellon University

June 15, 1992

FORMAT

- Three hours!
- Two parts

  Technical Presentation

  Questions and Discussion

  (each is approximately $1\frac{1}{2}$ hours in duration)
- To ensure enough time for both presentation and discussion:

  Please limit questions during the first part to matters of clarification

HISTORICAL SETTING

- In March, the IESG issued a call for proposals on evolving the Internet-Standard Network Management Framework
- Key observation:

  Existing framework provides stable and effective network management for the Internet, which is used pervasively and continuously

INTRODUCING THE SMP

- In response, the SMP specification was published

  Simple Management Protocol (SMP) Framework
- Authored by the "usual suspects"
- Eight documents:

  Introduction to SMP
  Structure of Management Information for SMP
  Textual Conventions for SMP
  Protocol Operations for SMP
  Transport Mappings for SMP
  Management Information Base for SMP
  Manager to Manager MIB for SMP
  SNMP/SMP Coexistence

TOPICS

Topic I: Structure Of Management Information and SMP MIB

Topic II: Transport Mappings and Protocol Operations

Topic III: Manager to Manager Interactions and Implementation Experience

Topic IV: Administrative Framework and Coexistence

TOPIC I

Structure of Management Information and SMP MIB

Topics

- New (and unchanged) Data Types
- ASN.1 macros
- Textual Conventions
- SMP MIB
- MODULE-COMPLIANCE for SMP agents

NEW DATA TYPES

- BIT STRING — enumerated only
- Counter64 — 64-bit, use only when wraps in under 1 hour
- NsapAddress — for OSI addresses
- Integer32 — to cap signed integers in $[-2^{31}, 2^{31} - 1]$
- Counter32, Gauge32 — identical to Counter/Gauge

UNCHANGED DATA TYPES

- OCTET STRING, OBJECT IDENTIFIER
- INTEGER — enumerated or with specified range
- TimeTicks, IpAddress
- Opaque — kept for backward compatibility

ASN.1 MACROS

- OBJECT-TYPE extended
- OBJECT-GROUP
- MODULE-COMPLIANCE
- AGENT-CAPABILITIES
- TRAP-DEFINITION (replaces TRAP-TYPE)

OBJECT-TYPE MACRO

- SYNTAX
- UNITS — e.g., `UNITS "seconds"` (new)
- ACCESS — MAX-ACCESS
  - Specifies maximum which makes "protocol sense"
  - Use "not-accessible" for auxiliary objects
  - Use "read-create" for creatable columnar objects
  - No more "write-only"
- STATUS
  - No more "optional"
- DESCRIPTION (now mandatory)
- REFERENCE
- INDEX or AUGMENTS (now mandatory for conceptual rows)
  - IMPLIED keyword in INDEX clause
- NUM-ENTRIES — e.g., `NUM-ENTRIES { ifNumber }` (new)
- DEFVAL

OBJECT-GROUP MACRO

- No longer in ASN.1 commentary
- Contained objects are specified explicitly

MODULE-COMPLIANCE MACRO

- The (minimum) requirements for compliance
- Specified in terms of MIB modules and GROUPs
- Objects specified only if minimum SYNTAX or ACCESS is different from MIB

AGENT-CAPABILITIES MACRO

- Which MIB modules/objects/values are actually implemented
- Evolution of RFC-1303's MODULE-CONFORMANCE
- Omission of CREATION-REQUIRES has different meaning

TRAP-DEFINITION MACRO

- Assigns an OBJECT IDENTIFIER to a trap
- No ENTERPRISE clause

TEXTUAL CONVENTIONS

- Defined using TEXTUAL-CONVENTION macro
- Value is anything which can go in a SYNTAX clause
- The usual DESCRIPTION, REFERENCE clauses
- DISPLAY-HINTS clause, e.g., `DISPLAY-HINTS "1x:"`

DEFINED TEXTUAL CONVENTIONS

- DisplayString, PhysAddress, MacAddress
- TruthValue, AutonomousType, InstancePointer
- TimeStamp, TimeInterval
- TestAndIncr
  - Provides for atomic, or sequenced, operations
  - When set, supplied value must match current value
  - After successful set, value increments

DEFINED TEXTUAL CONVENTIONS (cont.)

- RowStatus
  - The basis for creation/deletion of conceptual rows
  - Similar to RMON's EntryStatus
  - Writing underCreation(1) fails if already exists
  - Writing underModification(3) allows changes
  - Writing underDestruction(2) deletes conceptual row

SMP MIB

- Four groups:
  - smpInOut — extension of snmp group
  - smpOR — extension of sysObjectID/sysDescr
  - smpTrap — counters of #-traps sent to each party
  - smpSet — smpSetSerialNo, for sequencing Sets

MODULE-COMPLIANCE for SMP agents

- system, snmp
- partyTable, partySecretsTable, aclTable, viewTable
- smpInOut, smpOR, smpTrap, smpSet

TOPIC II

Transport Mappings and Protocol Operations

Topics

- Transport Mappings
- Introduction to Protocol Operations
- Protocol Entities
- Exceptions and Error Codes
- Bulk Retrieval Mechanism
- New Set Features
- Removal of TCP/IP-centric Features
- Traps and Event Notifications

Transport Mappings

Several are defined:

- smpUDPDomain: SMP over UDP ("Normal")
- smpOSIclnsDomain: SMP over connection-less NS
- smpOSIconsDomain: SMP over connection-oriented NS
- smpIPXDomain: SMP over IPX
- smpDDPDomain: SMP over Appletalk's datagram protocol DDP
- restartDomain: local configuration storage
- entityDomain: a device

USE UDP! and connection-less works best

Proxy used to translate between / among

Transport Mappings (cont.)

Other tidbits in Transport Mappings document:

- Minimum maximum message sizes
- Well known ports / transport selectors / etc
- BER serialization rules – including the possibility of use of other serialization rules with other transport domains although no others are initially defined

Introduction to Protocol Operations

First, what hasn't changed?

Answer: A lot!

- This is the same old S[N]MP you know and love
- Some enhancements
- Some problems corrected
- More alike than different
- While the changes are minor, the results are dramatic

Some Protocol Related Goals

- Improve performance
- Improve set operations, especially row creation and deletion
- Disambiguate error responses
- Silence the complaints about unsupported variables on the formerly wimpy but now not so wimpy get operator
- Reduce code size
- Remove TCP/IP-centric aspects
- Tighten Language
  - replace old ambiguities with new ones
  - replace old frequent errors with new ones
  - replace old frequently asked questions with new ones

Protocol Entities

- No real changes here
- Slightly new nomenclature in recognition of hierarchical management schemes which are expected to expand under SMP
- Entities can be:
  - SMP entity acting in an agent role
  - SMP entity acting in a manager role
  - SMP entity acting in a dual manager/agent role supporting Manager to Manager communications either hierarchically or in a mesh

Communications Between Protocol Entities

Communications can be:

- Request/response communications between an entity acting in a manger role and an entity acting in an agent role
- Unconfirmed event notification (a.k.a. a trap) from an entity acting in an agent role to a an entity acting in a manger role
- Request/response communications between two entities, both acting in manager roles (Manager to Manager communications)

Exceptions and Error Codes

Three classes:

- Exceptions
- Error responses
- Timeouts

Exceptions

Three kinds of exceptions, per variable binding, on responses to Get, GetNext, and GetBulk operations

- noSuchObject exception
- noSuchInstance exception
- endOfMibView exception

So stop complaining (-:

Error Responses

Need to disambiguate several overloaded error conditions so can stop flying blind, especially on sets

Instead of "NO!" you get "NO! because..."

New error types include:

noAccess	wrongType
wrongLength	wrongEncoding
wrongValue	noCreation
inconsistentValue	resourceUnavailable
commitFailed	undoFailed
authorizationError	notWritable

Timeouts

- Timeouts are another error condition
- A goal was to reduce the number of occasions in which a manager sends a query and gets back NOTHING
- That is, we need to disambiguate the "timeout" error condition
- Only so much can be done here
- Especially important now because will have new sources of the "timeout" error condition as a result of new security "features"
- New rules regarding authentication framework
- Empty variable bindings list on tooBig replies
- Also have new SMP counters to instrument silent drops

Bulk Retrieval Mechanism

Goals

- To retrieve large quantities of information
- To retrieve sparse tables
- Minimize round trips
- Efficiency
- Full packets
- Keep it simple with no new PDU format or connection-oriented transport

Bulk Retrieval Mechanism (cont.)

Network management personnel often need to retrieve single instances of some variables along with repeated instances of other variables.

Example: periodically retrieve interface statistics like:

ifInOctets, ifOutOctets, ifInErrors, IfOutErrors, ifInNUCastPkts, ifInUCastPkts, ifOutNUCastPkts, and ifOutUCastPkts

for each interface plus sysUpTime in order to compute packets per second, bytes per second and percent errors for each interface

Other uses include the retrieval of entire tables or sections of tables.

Example: retrieve the entire ipNetToMediaTable

Parameters Controlling The GetBulk Operation

Bulk retrieval operation is controlled by the combination of the:

- non-repeaters: number of variables for which single instances are requested
- max-repetitions: maximum number of instances requested for other variables
- variable bindings list in the request
- manager's party maximum message size
- agent's maximum message size

The values for non-repeaters and max-repetitions are communicated in the request in the fields normally associated with error-status and error-index

GetBulk Example

- GetBulkRequest [non-repeaters = 1, max-repetitions = 2]
  ((sysUpTime = NULL),
  (ipNetToMediaPhysAddress = NULL),
  (ipNetToMediaType = NULL))
- Response [error-status = 0, error-index = 0]
  ((sysUpTime.0 = "123456"),
  (ipNetToMediaPhysAddress.1.9.2.3.4 = "000010543210"),
  (ipNetToMediaType.1.9.2.3.4 = "dynamic"),
  (ipNetToMediaPhysAddress.1.10.0.0.51 = "000010012345"),
  (ipNetToMediaType.1.10.0.0.51 = "static"))

GetBulk Example (cont.)

- GetBulkRequest [non-repeaters = 1, max-repetitions = 2]
  ((sysUpTime = NULL),
  (ipNetToMediaPhysAddress.1.10.0.0.51 = NULL),
  (ipNetToMediaType.1.10.0.0.51 = NULL))
- Response [error-status = 0, error-index = 0]
  ((sysUpTime.0 = "123466"),
  (ipNetToMediaPhysAddress.2.10.0.0.15 = "000010987654"),
  (ipNetToMediaType.2.10.0.0.15 = "dynamic"),
  (ipNetToMediaNetAddress.1.9.2.3.4 = "9.2.3.4"),
  (ipRoutingDiscards.0 = "2"))

GetBulk Observations

- If non-repeaters = number of varbinds ↦ same as GetNext
- If non-repeaters = 0 and max-repetitions = 1 ↦ same as GetNext
- You never get a tooBig from GetBulk
- I intend to use GetBulk for a general replacement for GetNext which I used for nearly everything
- Typically requires changes to main routine only, not to each method routine
- Can change method routines if desired for better performance
- The GetNext operation is **powerful**
- The GetBulk operation is **awesome**

Sets

Sets will be much more important now that have security

- Problems cleaned up
- Richer error codes
- Row creation and deletion

Row Creation Using RowStatus

- Create by writing underCreation(1)
- Read the row with GetRequest
- Returned values can be modified
- For noSuchInstance, value must be written
- For noSuchObject, value must not be written
- Finally, writing active(4) brings row into use

Removal of TCP/IP-centric Features

- Transport Mappings document already mentioned
- Trap PDU had field with a NetworkAddress type
- NetworkAddress allows you to specify any choice of address type as long as it is an IP address (a choice from a list of size one is really no choice at all)

Non-TCP/IP-centric Trap

- Replace the Trap PDU
- Replacement looks amazingly similar to an unsolicited Response PDU
- Information that was in the headers is now in specified positions in the variable bindings list
- This is how it should have always been
- Since we now have only one PDU format instead of two, it allows less code size to generate and parse PDUs
- Traps are now named by an OBJECT IDENTIFIER

Trap Configuration

Manager stations can remotely configure trap destinations by manipulating the party information in the aclTable and viewTable (from SNMP Security)

Traps are sent to all parties for which:

- The aclSubject = the SMP protocol entity
- The (aclPrivileges & trapmask) = true
- The name of the trap is present in the viewTable of the aclSubject
- All of the variable bindings in the trap are present in the aclSubject's view

InformRequest PDU

- Used for Manager to Manager communications
- Looks amazingly like the other PDUs
- Configured at the request to a manager by a remote manager
- Sent from one manager to another
- Can result in:
  - acknowledgement response back to the original sender
  - error response back to the original sender
  - timeout and retransmission using parameters configured via the Manager to Manager MIB

TOPIC III

Manager to Manager Interactions and Implementation Experience

Manager to Manager Functions

- "Inform you of information in my view" — Inform PDU
- Event Notification (Acknowledged)
- Data transfer between managers

Dual Role

- Entity in manager role
  - Collects data
- Entity in agent role
  - Allows remote configuration of management functions
- Manager to Manager MIB defines this configuration

Alarm Generation

- Polls local or remote parties
- Detects threshold crossings or variable not accessible
- Flow control
  - Hysteresis
  - Maximum of 1 event per second (per threshold)
  - Acknowledgment of Informs

Configuration of Notifications

Manager to Manager MIB configures:

- Source and destination parties for Inform
- Destination party specifies destination address
- Retransmission parameters
  - Timeout
  - Retries

Access control

Alarm Entry queries data with party/secrets of dual-role manager, so access control is necessary.

- Access control based on MIB views
- Destination party in index of alarmEntry
  - Access excluded by default
  - Particular destination parties can be included in MIB view
- A single view entry allows Manager A to cause Agent C to be queried

Potential Uses

- Distribution of polling function
- Distributed diagnosis
- Minimize impact of security on agents (fewer parties and keys)

Implementation Experience

Size + Speed

- Size
  - Elimination of old trap PDU format makes code smaller
  - Authentication is simpler and smaller
- Speed
  - Bulk provides orders of magnitude improvement
  - Don't underestimate importance of "Add varbinds 'till full"

SNMP Security implementation problems solved

- Reordering protection was damaging to operational needs
  - Dropped messages reordered by network (even Get Requests)
  - SMP omits this algorithm
  - SMP MIB provides protection from reordering of SETs
- Party proliferation in SNMP Security
  - Projection for campus net was 10^5 parties
  - Clear and Present Danger to "low-impact" management
- Clock resync is simpler
  - Clock sync used to require an SNMP SET
  - In SMP, clock sync is part of normal operations

Row creation algorithm makes configuration easier

- Expanded error codes mean less guesswork
- Allows unimplemented/inaccessible objects in row
- MS can use, modify, or ignore default values from agent
  - Allows Agent to suggest values appropriate to situation

Applications are easier to write

- Packets return info even if varbinds have errors
- Tighter wording in protocol specification allows less code in applications
- Bulk replaces special case code in MS

TOPIC IV

Administrative Framework and Coexistence

Topics

- Administrative Framework
- Management Information
- Protocol Operations

ADMINISTRATIVE FRAMEWORK

- Based on SNMP Security, but with some exceptions
- `snpUDPdomain` — new transport domain
- Compliant behavior requires DES for party creation only
- Instance-level granularity of access is optional
- Source/Destination parties must use same authentication protocol

ADMINISTRATIVE FRAMEWORK (cont.)

- `snpMD5AuthProtocol` — new digest-auth protocol
- No ordered delivery mechanism
  - Not needed for retrieval
  - Not adequate for multiple managers
  - Instead, provide replay protection outside of "lifetime"
- Simplified clock synchronization
  - More "automatic" synchronization
  - SET operation not needed anymore (case #1)

MANAGEMENT INFORMATION

- Changes are fairly straight-forward
- IMPORT from SMP-SMI
  instead of RFC1155-SMI and RFC-1212
- Use new OBJECT-GROUP macro to define object groups
  instead of ASN.1 commentary
- "Standard" MIB modules should include a compliance definition
  instead of ASN.1 commentary
- You don't need to change object definitions to use SMP!

OBJECT TYPES

- SYNTAX:
  INTEGER (w/o range/enumerations) ⊢ Integer32
  Counter ⊢ Counter32
  Gauge ⊢ Gauge32
- ACCESS ⊢ MAX-ACCESS
  value should make "PROTOCOL SENSE"
  use "read-create" for creatable variables
  use "not-accessible" for auxiliary variables
- DESCRIPTION clause mandatory
- For conceptual rows:
  INDEX (or AUGMENTS) clause mandatory

AND MAY WE SUGGEST THAT YOU...

- Use RowStatus object for tables which support creation
- Cap OCTET STRINGs
  when at all possible
- Use the TEXTUAL-CONVENTIONS macro
  instead of ASN.1 commentary
- Add the UNITS clause
  when appropriate
- Use the AUGMENTS clause instead of the INDEX clause
  when appropriate

TRAP DEFINITIONS

- IMPORT from SMP-SMI
  instead of RFC-1215
- TRAP-TYPE ⊢ TRAP-DEFINITION
- Remove ENTERPRISE clause
- VARIABLES ⊢ OBJECTS
- TRAP-DEFINITION value is an OBJECT IDENTIFIER

PROTOCOL OPERATIONS

- Two approaches:
  Proxy agent
  Bi-lingual manager
- Mapping rules are the same

SMP ⊢ SNMP

- GET, GET-NEXT, and SET are unaltered
- GET-BULK ⊢ GET-NEXT
  Zero the non-repeater/max-repetitions fields

SNMP ↦ SMP

- GET-RESPONSE is unaltered
  - Preserve noSuchName, badValue, readOnly errors
  - On a tooBig error, remove the variable-bindings field
- Trap-PDU ↦ SMP-Trap-PDU
  - First varbind: sysUpTime.0
    - Use the timestamp field
  - Second varbind: smpTrapOID.0,
    - Use generic trap OID, or
    - Use enterprise.0.specific-trap
  - Last varbind: smpEnterpriseOID.0
    - Use the enterprise field

HOW TO GET A COPY (SMP SPECIFICATION)

- via FTP:

host:	lancaster.andrew.cmu.edu	ics.uci.edu	case.cs.utk.edu
area:	pub/smp/	mrose/isode-smp/	pub/smp/
files:	smp-*.txt	smp-*.txt	smp-*.txt

- via AFS:

  `/afs/andrew.cmu.edu/netdev/smp/`

- via E-MAIL:

  `smp-spec@dbc.mtview.ca.us`

HOW TO GET A COPY (SMP SOFTWARE)

- Two commercial packages:
  - Hughes LAN Systems and SNMP Research
- Two openly available packages:

what:	CMU SNMP 2.0.1	4BSD/ISODE SMP
host:	lancaster.andrew.cmu.edu	ics.uci.edu
area:	smp-dist	mrose/isode-smp/
file:	README	isode-smp.tar.Z

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Uninterruptible Power Supply BOF (UPSMIB)

Agenda

- To identify the scope of the problem of monitoring and controlling uninterruptible power supplies;
- To discuss an Internet Draft containing an initial proposal for such a MIB, and
- To assess the interest and commitment towards ongoing work, including the possibility of creating a Working Group to prepare and advance proposals for standardization in this area. If there is sufficient interest and commitment, the Working Group Charter and timetable will be discussed.

For this meeting, Jeff Case presided and Bob Stewart recorded. The meeting was well-attended, about 30 people, with representatives of about 10 UPS vendors, many becoming involved in the Internet and the IETF for the first time. After considerable discussion and review of a proposal, the meeting decided to request startup of a Working Group, with most of the work being done via a mailing list.

Goal

Efforts to make uninterruptible power supplies to be monitorable and controllable via the Internet Standard Management Framework have already begun. In the past, when MIB standardization has trailed product development, as it did for terminal servers, intelligent repeaters, and MAC bridges, users have been faced with the difficulties associated with the unnecessary proliferation of similar, but different, enterprise-specific (vendor) MIBs. As a result, it is desirable to begin standardization efforts as soon as possible.

A draft document has been prepared as an introduction to the problem. It states:

> This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing uninterruptible power supply (UPS) systems.

The document which is in the Internet-Drafts Directory, as previously announced, is a new version of the memo which contains incorporates the suggestions received by the authors since the initial document was published, plus one new group, the upsTruck group.

Jeff presented the organizations and procedures that oversee Internet Standard Development.

- Mailing list discussion a bit of a problem due to lack of mail access by UPS vendors.
- Stressed structure above IAB, current and changing.
- All-volunteer lower structure.

Jeff stated goals.

- Decide whether to proceed.
- Identify sufficiently interesting set of common attributes with no optional objects. The answer to a concern over response to unimplemented objects was that groups are the unit of conformance for interoperability.

Jeff presented highlights of the strawman proposal, which caused several points of discussion.

- Volts and Minutes are too coarse, tenths of volts and seconds better.
- Concern over debating individual objects deferred for later detailed assessment of proposal.
- Traps are to be in a separate document to ease passage of MIB.
- A UPS need not perform multiple tests simultaneously but may if it can.
- It is implementation specific whether configuration options cause changes.
- SNMP proxy mechanism preferred for handling multiple UPSes, rather than table with index.
- A community string or party defines an agent.
- Fielded systems are basic and advanced, MIB represents advanced, suggest organizing MIB accordingly. Agreed too much mandatory for less expensive devices. Agree with option by Group for predictable functions. This is a marketing issue. This discussion should be deferred for formal Working Group. We all want one standard.

We discussed whether we do indeed want to form a Working Group.

- Charter 1 or 2 documents to monitor and possibly control UPS, low or high end, existing and future UPS technology.
- Consider application to similar embedded systems such as power systems or power conditioners.

- Prefer not to encumber UPS needs.

- There was considerable concern about voting and influence. Process is by consensus without company dominance or simple votes. The major need is to define the problem and rally around a Standard. SNMP itself was such a compromise.

- We need an editor and email communication. MCI, Sprint, Compuserve, etc., provide mail service, and some do not charge for reading mail.

- Arbitration is informal, by consensus and compromise.

- Credit is author's on front and sometimes individuals in acknowledgements section, with name and affiliation.

- Mail is easier than news group.

- Publicity is acceptable as long as not claiming standard before complete. Internet Drafts should not be referenced in procurement or product literature, but RFCs may.

- Mailing list administrative address is upsmib-request@cs.utk.edu.

- We plan to have a document by the next IETF, final by following. The next IETF conflicts with Comdex, a big problem. Suggested Las Vegas meeting that vendors attend is a problem for Chairs.

One of the vendors (APC) presented an alternative proposal.

- MIB being implemented but needs to consider strawman proposal.

- Proprietary features were removed for presentation. Remainder was divided into basic and advanced to maximize compatibility with past and future systems.

- Extension objects point to further MIB, assuming it is similarly structure. A single object is preferred and sufficient.

Several general issues were discussed before adjournment at 10:20.

- Someone suggested a breaker group. That varies considerable across implementations. It could be handled by alarm group in strawman.

- We looked at several objects in strawman and general consensus was they are implementable.

- Although one of the proposers, Emerson, does not implement everything in the strawman, it was proposed for value to the industry.

- On the issue of credit for draft and RFC the suggestion was to limit it to SNMP Research. Suggested that Emerson receive consideration as catalyst. Deferred to private discussions and mailing list.

- Suggested September meeting central in U.S. deferred to mailing list.

Attendees

Richard Baxter	
Tom Brennan	brennan@exide.com
Jeffrey Case	case@cs.utk.edu
James Davin	jrd@ptt.lcs.mit.edu
Michael Davison	davison@cs.utk.edu
Roger Draper	rdraper@cerf.net
Bill Elliot	
David Engel	david@ods.com
David Fencl	
Owen Gallagher	oweng@jjmhome.uucp
Theodore Greene	
Kenneth Key	key@cs.utk.edu
Sharon Lewis	lewis@cs.utk.edu
Les Matheson	matheson@cerf.net
Paul Moran	Paul_Moran@3com.com
David Perkins	dperkins@synoptics.com
Marshall Rose	mrose@dbc.mtview.ca.us
Koichiro Seto	seto@hitachi-cable.co.jp
Houman Shafiezadeh	houman@exide.com
Timon Sloane	peernet!timon@uunet.uu.net
Einar Stefferud	stef@nma.com
Bob Stewart	rlstewart@eng.xyplex.com
Adam Stolinski	
Ray Wasson	
Brian Young	

2.3.1 Bridge MIB (bridge)

Charter

Chair(s):

Fred Baker, fbaker@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"

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

Minutes of the Bridge MIB Working Group (BRIDGE)

The Group met for three purposes:

1. To discuss IEEE 802.5's changes to their MIB, and its impacts on the MIB described in RFC 1286.

2. To discuss implementation experience with RFC 1286.

3. To determine whether RFC 1286 is ready to advance to Draft Standard status.

Anil Rijsinghani proposed to facilitate convergence with the Source Routing Addendum to 802.1 by including an optional group for SRT bridges, called the Source Route Bridge Port Pair Group. It contains the following objects:

```
dot1dSrPortPairTableSize
        INTEGER
        "The total number of entries in the Bridge Port Pair Database."
```

This number is n(n+1)/2, given that source routing is occurring over n bridge ports.

```
dot1dSrPortPairTable
dot1dSrPortPairEntry [dot1dSrPortPairLowPort, dot1dSrPortPairHighPort]

dot1dSrPortPairLowPort  - an Source Route Port Number
dot1dSrPortPairHighPort  - an SOURCE ROUTE Port Number
dot1dSrPortPairBridgeNum - the bridge number used in the Source
Route Descriptor tuple
dot1dSrPortPairState     - "enabled", "disabled", or "invalid"
```

Richard Sweat, IEEE 802.5's designated liaison to the Bridge MIB Working Group, then presented their view of RFC 1286. To converge with our work, IEEE 802.5 has deleted or modified a number of its managed objects and attributes. They also have some specific requests for changes in the Source Routing Group of RFC 1286.

IEEE 802.5, having already made these changes in its own MIB, suggests that we:

- Adopt Anil's Port Pair Group.

- Divide dot1dSrPortHopCountExceededDiscards and dot1dSrPortHopCount, which relate to the maximum number of routing descriptors in an All Paths Explorer (APE) or Spanning Tree Explorer (STE) frame, into two maxima and two counters: one each for APEs and STEs.

- Extend dot1dSrPortLargestFrame (which is an enumerated integer with 8 values) to have 64 possible values as described in draft 7 of the Source Route Addendum.

- Count occurrences of duplicate LAN IDs or Tree errors, in an effort to detect problems in networks containing older IBM Source Routing Bridges.

- Count LAN ID Mismatches (cases where a frame is being forwarded, but the "from" LAN ID is incorrect.

- Instead of counting frames in and frames out, count frames through a device. This applies to dot1dSrPortSpecInFrames, dot1dSrPortSpecOutFrames, dot1dSrPortApeInFrames, dot1dSrPortApeOutFrames, dot1dSrPortSteInFrames, and dot1dSrPortSteOutFrames.

- To the dot1dSrGroup, add a scalar read-write variable enumerated the same way as dot1dSrPortLargestFrame to indicate the largest frame that may pass through the bridge.

- Add a read-write LF Mode field indicating whether the bridge operates using older 3 bit length negotiation fields or the newer 6 bit length field in its RIF.

- Either change the names of objects or include text explaining the use of the path type acronyms, as IEEE 802.5 has changed their names. The mapping is:

  - Spanning Tree Explorer (STE) becomes a Spanning Tree Explorer (STE).
  - All Paths Explorer (APE) becomes an All Routes Explorer (ARE).
  - Specifically Routed Frame (Spec) becomes a Source Routed Frame (SRF).

Fred Baker applauds the efforts of IEEE 802.5 to achieve convergence. The Working Group felt that the proposals made by Anil and Richard were basically workable, and drew the following conclusions. We also felt (although there were three source routing implementations represented) that our best expertise was not present at the meeting, and so feel that the subject should be discussed on the mailing list before reaching a final conclusion.

The Group's initial conclusions were:

- Adopt Anil's Port Pair Group.

- The Group is not sure of the necessity of dividing the hop counts and hop count

discards by APEs and STEs.

- Extend dot1dSrPortLargestFrame (which is an enumerated integer with 8 values) to have 64 possible values as described in draft 7 of the Source Routing Addendum.

- Count occurrences of duplicate LAN IDs or Tree errors, in an effort to detect problems in networks containing older IBM Source Routing Bridges.

- Count LAN ID Mismatches (cases where a frame is being forwarded, but the "from" LAN ID is incorrect.

- Do not change the way we count frames, as our implementations do in fact count them (IEEE was concerned that these could not be counted), and this method is consistent with other IETF MIBs.

- To the dot1dSrGroup, add a scalar read-write variable enumerated the same way as dot1dSrPortLargestFrame to indicate the largest frame that may pass through the bridge.

- Add a read-write LF Mode field indicating whether the bridge operates using older 3 bit length negotiation fields or the newer 6 bit length field in its RIF.

- Include text explaining the use of the path type acronyms.

The Group then moved on to discuss implementation experience. Six vendors indicated that they had implemented the MIB, and were largely happy with it. The following suggestions for clarification were made:

- Anil will provide clarifying text for the DESCRIPTION of dot1dStpPortPathCost, which some have found inadequate.

- The Default Value of dot1dStaticAllowedToGoTo be specified in the DESCRIPTION as a string of ones of appropriate length.

- The Default Value of dot1dStaticStatus is "permanent".

- Port Numbers use the range 1..65535.

- Update the bibliography.

- dot1dTpPortInFrames and dot1dTpPortOutFrames be clarified by changing the last few words in the description from "processed by the local bridging function" to "processed by the bridging function, including management frames."

We will ask the list to ratify these clarifications.

One other issue was raised which affects the strategic direction of this Working Group. Some vendors are interested in proxying for IBM Source Routing Bridges, which use a modified version of the 802.1(d) BPDU. Also, IEEE 802.1(g) currently proposes that the BPDU be modified in remote bridges when sent on the lines. It is quite possible, then, that a bridge might participate in more than one spanning tree on a port by port basis.

Fred Baker proposed that the object dot1dStpProtocolSpecification, which indicates a single spanning tree protocol in use in the system, be deprecated and replaced with an INTEGER bit string indicating the spanning tree protocols that the device is capable of:

```
     1      other
     2      decLb100
     4      ieee8021d
     8      ibmTolkienRing
    16      ieee8021g
```

In addition, an object is added to the dot1dStePortEntry indicating which of those protocols is running on the indicated port. This allows for some flexibility.

The proposal of the Working Group, given ratification of the above changes on the list, is that:

- The Group do nothing now with IEEE 801.2(g)'s proposals, as it is not sufficiently close to completion.
- The Group separate the Source Routing Group into a separate document, apply the ratified subset of Anil's and Richard's proposals, and recommend that this remain at Proposed Standard status.
- The Group apply the requested clarifications and, if ratified, Fred's proposed object change, and advance the bulk of the Bridge MIB to Draft Standard Status.

Attendees

David Arneson	arneson@ctron.com
Sam Ayers	s.ayers@sybus.com
Fred Baker	fbaker@acc.com
Andy Bierman	bierman@davidsys.com
Greg Celmainis	gregc@newbridge.com
Chris Chiotasso	chris@artel.com
Cathy Cunningham	cmc@microcom.com
David Engel	david@ods.com

Pria Graves	priag@nsd.3com.com
George Kajos	kajos@coral.com
Mark Kepke	mak@cnd.hp.com
Steven Lynes	lyness@rimail.interlan.com
Cindy Mazza	
Keith McCloghrie	kzm@hls.com
David Minnich	dwm@fibercom.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
John Payne	jop@wang.com
Venkat Prasad	vsp@3com.com
Guenter Roeck	roeck@conware.de
Michael Sapich	sapich@conware.de
Koichiro Seto	seto@hitachi-cable.co.jp
Timon Sloane	peernet!timon@uunet.uu.net
Chris Sullivan	mm@gandalf.ca
Richard Sweatt	rsweatt@synoptics.com
Stephen Tsun	snt@nsd.3com.com
Luanne Waul	luanne@wwtc.timeplex.com
Gerard White	
Kiho Yum	kxy@nsd.3com.com

2.3.2 Chassis MIB (chassis)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com
Jeffrey Case, case@cs.utk.edu

Mailing Lists:

General Discussion: chassismib@cs.utk.edu
To Subscribe: chassismib-request@cs.utk.edu
Archive:

Description of Working Group:

This Working Group will produce a document describing MIB objects for use in a "chassis" — which is a collection of traditionally discrete network devices packaged in a single cabinet and power supply. A chassis may comprise, for example, combinations of layer 1 repeater elements, MAC layer bridges, or internetwork layer routers.

The Working Group is chartered to produce up to three distinct documents that define extensions to the SNMP MIB:

(1) The Working Group is chartered to define MIB objects that represent the mapping of the logical functions of traditional network devices onto particular, physical hardware resources within the chassis. These MIB definitions will not address any aspects of the network functions comprised by a chassis box that are shared with an analogous collection of discrete network devices.

(2) The Working Group is chartered, at its option, to define MIB objects that instrument the operational state of a power supply element in a chassis.

(3) The Working Group is chartered, at its option, to define MIB objects that represent aggregated information about collections of network devices (e.g., aggregate information about devices attached to a particular LAN), provided that this MIB specification is not specific to chassis implementatons of such networks and is also readily implementable for analogous collections of discrete network devices.

The MIB object definitions produced will be for use by SNMP and will be consistent with existing SNMP standards and framework.

Although the Working Group may choose to solicit input or expertise from other relevant standards bodies, no extant standards efforts or authorities are known with which alignment of this work is required.

Because the structure of chassis implementations varies widely, the Working Group shall take special care that its definitions reflect a generic and consis-

tent architectural model of chassis management rather than the structure of particular chassis implementations.

Should the Working Group elect to define objects representing aggregated information about collections of network devices, those efforts will not compromise the operational robustness of the SNMP that depends on its realization of management system function as closely as possible to centers of responsible authority.

Goals and Milestones:

Done	Discuss the Charter and define the scope of the Working Group. In particular, review all contributed MIBs and agreement on plan for producing baseline document(s).
Jul 1992	Post the first draft of the Chassis MIB specification as an Internet Draft.
Jan 1993	Submit the Chassis MIB to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Chassis MIB Working Group (CHASSIS)

The purpose of this meeting was to try to get the Chassis MIB Working Group started/restarted.

All sorts of events have conspired to detract from efforts to get it started. A major contributor from Cabletron got reassigned. Another major contributor and one of the co-Chairs has been tied up with other standardization efforts such as (SMP). The other co-Chair has been ill for nearly a month but we are pleased to hear he is better. Finally, the members of the non-Working Group seem to think that the mailing list is a read-only object.

The Group spent time reassessing the interest in and commitment to this work with the intention of pulling the plug if appropriate (though everyone hoped it would not be necessary). In addition attempts were made to put together a new time-action-plan for the work.

This was a short meeting. There was nothing to be read or said.

The Charter was reviewed as well as the work done, noting that there were only thirteen messages on the mailing list since our first meeting.

- The Group was supposed to have Chassis MIB submissions and synthesize a proposal. Nothing had been done, and it still needs doing.

- The Group was supposed to get power supply submissions. It got 1 or 2. More needs to be done.

- Nothing was done on the Aggregate MIB.

Time was spent reassessing interest, and the following conclusions were made.

- Interest remains, though perhaps at slightly lessened level. The meeting well attended.
- The Cabletron implementation has moved in the direction of Keith & Donna's proposal.

The following plans were made.

- For Chassis MIB: authors are to post original documents, or a new one within a few weeks. Cabletron proposal could be updated. Originals are in archive.

- For Power Supply MIB: MIBs are to be submitted in a week, and about ten are expected.

- The Aggregate MIB has been placed on the back burner.

Attendees

Andy Bierman	bierman@davidsys.com
Natan Borshansky	natanb@turisc.timeplex.com
Steve Bostock	steveb@novell.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Chris Chiotasso	chris@artel.com
Dave Cullerot	cullerot@ctron.com
Cathy Cunningham	cmc@microcom.com
Michael Davison	davison@cs.utk.edu
David Engel	david@ods.com
Fred Engel	engel@concord.com
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Pria Graves	priag@nsd.3com.com
Ronald Jacoby	rj@sgi.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Deidre Kostick	dck2@sabre.bellcore.com
Cheryl Krupczak	cheryl@cc.gatech.edu
David Lindemulder	dcl@mtung.att.com
Kim Mayton	mayton@wg.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@synoptics.com
Edwin Mier	mier@tigger.jvnc.net
David Minnich	dwm@fibercom.com
Lynn Monsanto	monsanto@sun.com
Paul Moran	Paul_Moran@3com.com
Patrick Mullaney	mullaney@ctron.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
Sam Roberts	sroberts@farallon.com
Guenter Roeck	roeck@conware.de
Dan Romascanu	dan@lannet.com
Michael Sapich	sapich@conware.de
Jim Scott	scott@kali.enet.dec.com
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Mark Therieau	markt@python.eng.microcom.com

Chuck Townsend	townsend@ctron.com
Stephen Tsun	snt@nsd.3com.com
David Waitzman	djw@bbn.com
Kiho Yum	kxy@nsd.3com.com
Joseph Zur	zur@fibhaifa.com

2.3.3 DS1/DS3 MIB (trunkmib)

Charter

Chair(s):

Tracy Cox, tacox@sabre.bellcore.com
Fred Baker, fbaker@acc.com

Mailing Lists:

General Discussion: trunk-mib@saffron.acc.com
To Subscribe: trunk-mib-request@saffron.acc.com
Archive:

Description of Working Group:

This Working Group will consider revisions to the DS1 and DS3 MIBs (currently published as Proposed Stds in RFC 1232 and RFC 1233) in preparation for their consideration as Draft Standards.

Consistent with the IETF standards process, the Working Group is chartered to consider only those changes to the DS1 and DS3 MIBs that are based on implementation experience or on the need to align with relevant ANSI T1M1 standards. In this context, the Working Group will thoroughly document the implementation or alignment rationale for each considered change.

All changes made by the Working Group will be consistent with the existing SNMP framework and standards — in particular, those provisions of RFC 1155 regarding addition and deprecation of objects in standard SNMP MIBs.

This Working Group will be a short-lived activity, involving a single meeting, and will conclude its business no later than June 1992.

Goals and Milestones:

Mar 1992	Submit the DS1 document for the Network Management Directorate Review.
Apr 1992	Submit the DS1 MIB to the IESG for Draft Standard Status.
Mar 1992	Submit the DS3 MIB to the Network Management Directorate for review.
Apr 1992	Submit the DS3 MIB to the IESG for approval as a Draft Standard.
Feb 1992	Post a draft version of the new DS1 MIB to the Internet-Drafts Directory.
Feb 1992	Post a revised version of the DS3 MIB to the Internet-Drafts Directory.

CURRENT MEETING REPORT

Reported by Tracy Cox/Bellcore

Minutes of the DS1/DS3 MIB Working Group (TRUNKMIB)

The decision of the Working Group, per Bill Versteeg's suggestion, was to completely obsolete RFC 1232 and RFC 1233, and remove all of the deprecated objects from this document. All tables have been given new names and new OIDs. The beginning delimiter for all objects is dsx* (* = 1 or 3). In addition, we added LESs to the near end tables, changed all Counters remaining (current table) to Gauges, and modified the dsx1LineCoding, dsx1LineStatus, dsx*SendCode Object descriptions, and added a Signaling Mode object (DS1 MIB only). We have an action item to resolve on whether to keep or remove the BPVs count.

James Watt took an action item to supply appropriate text for the failure states description. A description of Line Errored Seconds (which T1M1 has available only from the far end, but we decided to include for the near end for symmetry's sake) was added.

The DS1 and DS3 MIBs are available for anonymous FTP at fennel.acc.com.

New Internet Drafts reflecting these changes will be sent to the trunk-mib mailing list and posted in the Internet-Drafts Directories; when consensus is achieved on the mailing list, they will be forwarded to the IESG for their review and approval as new RFCs obsoleting RFC1232 and RFC1233.

Attendees

Fred Baker	fbaker@acc.com
Alan Bryenton	bryenton@bnr.ca
Steve Buchko	stevebu@newbridge.com
Greg Celmainis	gregc@newbridge.com
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Tom Easterday	tom@cic.net
John Labbe	labbe@merit.edu
Chris Sullivan	mm@gandalf.ca
Kaj Tesink	kaj@cc.bellcore.com
James Watt	jamesw@newbridge.com

2.3.4 Ethernet MIB (ethermib)

Charter

Chair(s):

Frank Kastenholz, kasten@ftp.com

Mailing Lists:

General Discussion: enet_mib@ftp.com
To Subscribe: enet_mib-request@ftp.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.

Internet Drafts:

"Implementation Notes and Experience for The Internet Ethernet MIB", 03/24/1992, Frank Kastenholz <draft-ietf-ethermib-implexp-00.txt>

"Definitions of Managed Objects for the Ethernet-like Interface Types", 07/20/1992, Frank Kastenholz <draft-ietf-ethermib-objectsv2-02.txt>

2.3.5 FDDI MIB (fddimib)

Charter

Chair(s):

Jeffrey Case, case@cs.utk.edu

Mailing Lists:

General Discussion: fddi-mib@CS.UTK.EDU
To Subscribe: fddi-mib-request@CS.UTK.EDU
Archive:

Description of Working Group:

The FDDI MIB Working Group is chartered to define a MIB for FDDI devices that is consistent with relevant FDDI specifications produced by ANSI. All definitions produced by this working group will be consistent with the SNMP network management framework and other internet-standard MIBs for SNMP.

Goals and Milestones:

Done	"Final" initial draft of required get/set variables.
Done	Initial implementations of required get/set variables.
Done	Revised "final" draft of required get/set variables.
Done	Adoption of draft of required get/set variables.
Dec 1992	Post an Internet Draft aligned with current the current ANSI document factoring in implementation experience with RFC 1285.
Nov 1992	Hold a meeting at the November IETF Plenary.
Mar 1992	Submit the FDDI MIB to the IESG for consideration as a Proposed or Draft Standard depending on the magnitude of changes to RFC 1285.

Request For Comments:

RFC 1285	"FDDI Management Information Base"

CURRENT MEETING REPORT

Reported by Jeff Case/UTenn

Minutes of the FDDI MIB Working Group (FDDIMIB)

This being the first meeting of the Group after a lengthy hiatus, the Chair invested the time in introducing the goals as well as allowing time for participants to introduce themselves to one another.

Implementation experience with RFC 1285 and FDDI-related vendor extensions were discussed, including problems encountered with RFC 1285. Although there are several products shipping today with MIBs which are precursors of RFC 1285, there is not a large body of experience with the version of the MIB found in the RFC.

ANSI participants reported on the status of ANSI SMT 7.1 and 7.2 standardization status. ANSI has made much progress since our last meeting. The 7.1 and 7.2 documents are very similar with respect to the MIB. Reported differences include such things as additional enumerations on some integer valued objects.

They discussed the next step. The consensus was that it would be far better to attempt to align with the new ANSI documents rather than attempt to advance 1285 from proposed to draft without modification(s).

The majority of the meeting was spent discussing the differences between 1285 and the 7.1 specification. This was aided by a document prepared at great labor by Judi Theg Talley and the Chair recognizes this effort.

A first level consensus was reached on what to do about most objects with respect to translation algorithms, and rules for inclusion. The Chair is to prepare a document reflecting those decisions and post it to the list for discussion.

Attendees

Michael Anello	mike@xlnt.com
David Arneson	arneson@ctron.com
Zareh Baghdasarian	zareh@tekelec.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Christopher Bucci	bucci@ngc.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Chris Chiotasso	chris@artel.com
David Engel	david@ods.com
Bill Fardy	fardy@ctron.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Shawn Gallagher	gallagher@quiver.enet.dec.com

James Halpin	halpin@turkey.sw.stratus.com
Daniel Hansen	dan@ngc.com
Jeff Hughes	jeff@col.hp.com
David Husak	dave@synnet.com
Ronald Jacoby	rj@sgi.com
George Kajos	kajos@coral.com
Nick Kawaguchi	mamster@lanai.cs.ucla.edu
Mark Kepke	mak@cnd.hp.com
Ron Mackey	rem@dsiinc.com
Kim Mayton	mayton@wg.com
David Minnich	dwm@fibercom.com
Rohit Mital	rm@protools.com
Paul Moran	Paul_Moran@3com.com
Patrick Mullaney	mullaney@ctron.com
Luc Pariseau	pariseau@quiver.enet.dec.com
David Perkins	dperkins@synoptics.com
Eric Peterson	elpeterson@eng.xyplex.com
James Reeves	jreeves@synoptics.com
Salvatore Ricci	ricci@mtunm.att.com
Anil Rijsinghani	anil@levers.enet.dec.com
Dan Romascanu	dan@lannet.com
Koichiro Seto	seto@hitachi-cable.co.jp
Timon Sloane	peernet!timon@uunet.uu.net
Lawrence Stefani	stefani@quiver.enet.dec.com
Bruce Taber	taber@interlan.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Stephen Tsun	snt@nsd.3com.com
Ahmet Tuncay	atuncay@synoptics.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
David Waitzman	djw@bbn.com
Kiho Yum	kxy@nsd.3com.com
Joseph Zur	zur@fibhaifa.com

2.3.6 Host Resources MIB (hostmib)

Charter

Chair(s):

Steven Waldbusser, waldbusser@andrew.cmu.edu

Mailing Lists:

General Discussion: hostmib@andrew.cmu.edu
To Subscribe: hostmib-request@andrew.cmu.edu
Archive:

Description of Working Group:

The Host Resources MIB Working Group is chartered to produce exactly one document that defines SNMP MIB objects that instrument characteristics common to all internet hosts. The goal of this work is to address the urgent operational need in the internet community for management of host systems. Owing to this urgency, the Working Group will focus exclusively on the alignment of existing MIB technology in order to achieve common solutions in a timely manner.

For purposes of this effort, the term "internet host" is construed to mean any computer that communicates with other similar computers attached to the internet and that is directly used by one or more human beings. Although the work of the Group does not necessarily apply to devices whose primary function is communications services (e.g., terminal servers, routers, bridges, monitoring equipment), such relevance is not explicitly precluded. The single MIB produced shall instrument attributes common to all internet hosts including, for example, both personal computers and systems that run variants of Unix.

The methodology of this Working Group is to focus entirely on the alignment of existing, enterprise-specific MIBs for SNMP that are relevant to its task. The Group will work towards its goal by distillation and generalization of these existing MIBs into a single, common MIB definition.

Owing to the urgent operational need for managing host systems, this effort will not be comprehensive in scope. Rather, the MIB produced by this Group will be confined to critical information about hardware and software configuration, processor and memory use, and data storage capacities, backup, and use.

Owing to the lack of a well-understood and accepted architecture, the Working Group will not address in any way, mechanisms that could be used to monitor or control the use of licensed software products.

All definitions produced by the Group will be consistent with the SNMP network management framework and all other internet-standard MIBs for SNMP. Wherever possible, the definitions produced will make use of or align with

relevant work in progress with chartered working groups of the IETF. Also, wherever possible, the Working Group will take into consideration pre-existing, stable work produced by other, accredited standards bodies.

Goals and Milestones:

Done	First Working Group meeting. Discuss the initial proposed document.
Sep 1992	Post an Internet Draft describing the Host Resources MIB.
Dec 1992	Submit the Host Resources MIB to the IESG as a Proposed Standard.
Oct 1992	Hold an interim meeting to discuss the current document.
Nov 1992	Meet at the IETF plenary to identify changes necessary for Working Group closure.

CURRENT MEETING REPORT

Reported by Steven Waldbusser/CMU

Minutes of the Host Resources MIB Working Group (HOSTMIB)

Review of Charter

The Charter was presented and briefly reviewed. The Charter's aggressive schedule for completion was also noted. A concern was voiced that the Charter was too narrow, but there was general agreement that due to the desire to move quickly, a short, focused effort needed to be made.

MIB Archive

An archive of vendor mibs for host management has been created and 10 vendors have placed mibs there. In addition, a proposal for the Host MIB was created by Pete Grillo and Steve Waldbusser and was placed in the archive.

The archived MIBs are available for anonymous FTP from lancaster.andrew.cmu.edu in the directory pub/hostmib. Currently the following MIBs may be found there:

```
apple           unix-Krupczak      unix-cmu        unix-sony
hostmib.txt     unix-Rose-Davison  unix-dec        unix-sun
pc-att          unix-bsd-2         unix-hp
```

Any other MIBs for host management should be sent via email to waldbusser@cmu.edu so that they may be made available in the archive.

Several of the attendees mentioned that they had other MIBs to send to the archive. In addition, somebody mentioned that MIT Project Athena has a MIB that should be included.

Host Resource MIB Proposal

Pete Grillo and Steven Waldbusser made a presentation of the MIB that they created and were submitting to the Working Group us a starting point for the Group's efforts. The MIB was explained table by table and a demonstration was given that the MIB objects defined were in fact platform independent, using Unix, PCs and Macintosh systems as examples.

After the presentation, there were many questions asked about the MIB. Several bugs was found and many helpful suggestions were noted.

One area that got a lot of attention was the software group which instruments the installed software and the running software on a host. An opinion was raised that this installed software table was hard to implement given the potentially large number of pieces of software installed on a system. A couple of implementation strategies were suggested, and at least one participant expressed plans to implement the table and report back on the feasibility. A concern was also raised that this table might not represent software packages well, but there wasn't a generally accepted principle of what a software package would look like or when it would be a general enough mechanism for a standard MIB.

Future Plans

There was general agreement that the proposed MIB should be used as a baseline for future work by the Working Group. The comments received should be incorporated into another draft which would be submitted as an Internet Draft.

The Working Group agreed with the aggressive schedule and that the Group should next meet in September at a date and time to be announced later.

Attendees

David Arneson	arneson@ctron.com
John Ballard	jballard@microsoft.com
Jonathan Biggar	jon@netlabs.com
Steve Bostock	steveb@novell.com
Lida Carrier	lida@apple.com
Tom Coppeto	tom@mit.edu
Cathy Cunningham	cmc@microcom.com
James Davin	jrd@ptt.lcs.mit.edu
Michael Davison	davison@cs.utk.edu
Raymond Edgerton	bellmf!edgerton@uwm.edu
David Engel	david@ods.com
Michael Erlinger	mike@lexcel.com
Bill Fardy	fardy@ctron.com
Rob Graham	robert_graham@protools.com
Pete Grillo	pl0143@mail.psi.net
James Halpin	halpin@turkey.sw.stratus.com
Fred Henry	fgh@sodium.att.com
Ronald Jacoby	rj@sgi.com
Mark Kepke	mak@cnd.hp.com
Kenneth Key	key@cs.utk.edu
Cheryl Krupczak	cheryl@cc.gatech.edu
Arun Mahajan	axm@proteon.com
Kim Mayton	mayton@wg.com
Mike Migliano	bellmf!mike@uwm.edu
David Minnich	dwm@fibercom.com

Greg Minshall	minshall@wc.novell.com
Rohit Mital	rm@protools.com
Lynn Monsanto	monsanto@sun.com
Paul Moran	Paul_Moran@3com.com
Patrick Mullaney	mullaney@ctron.com
Muki Murthy	muki@microcom.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Tom Nisbet	nisbet@tt.com
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
Bob Page	bob.page@eng.sun.com
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Mike Ritter	mwritter@applelink.apple.com
Kary Robertson	kr@concord.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Simon Rosenthal	simon@bos.locus.com
Jon Saperia	saperia@tcpjon.ogo.dec.com
Koichiro Seto	seto@hitachi-cable.co.jp
Timon Sloane	peernet!timon@uunet.uu.net
Einar Stefferud	stef@nma.com
Thomas Talpey	tmt@osf.org
Rodney Thayer	rodney@world.std.com
Mark Therieau	markt@python.eng.microcom.com
Chris Thomas	chris_thomas@ccm.hf.intel.com
Dean Throop	throop@dg-rtp.dg.com
Panos-Gavriil Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Stephen Tsun	snt@nsd.3com.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Kathleen Walsh	walshy@ctron.com
Steven Wong	wong@took.enet.dec.com
Honda Wu	natadm!honda@uunet.uu.net
Kiho Yum	kxy@nsd.3com.com
Joseph Zur	zur@fibhaifa.com

2.3.7 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.
Done	Distribute updated documents for more E-mail discussion.
Done	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Done	Documents available with specified changes incorporated.
Done	Submit the Repeater MIB to the IESG for consideration as a Proposed Standard.

Internet Drafts:

"Definitions of Managed Objects for IEEE 802.3 Repeater Devices", 07/23/1991, Donna McMaster, Keith McCloghrie <draft-ietf-hubmib-mib-04.txt>

CURRENT MEETING REPORT

Reported by Keith McCloghrie/Hughes

Minutes of IETF 802.3 Hub MIB Working Group (HUBMIB)

The meeting was called to order at 4:00 by co-Chairs Donna McMaster and Keith McCloghrie.

Agenda

- Introduction
- IEEE Report
- Repeater MIB
  - Status
  - Implementation Experience
- Repeater MAU (Media Attachment Unit) MIB
  - Review of Instantiation Issues
  - Technical Walkthrough
- Plans for Future

IEEE Report

Donna summarized the status of the IEEE 802.3 Management work as follows:

- Repeater Management Draft
  - Passed 2nd IEEE 802.3 confirmation ballot.
  - Forwarded to TCCC ballot.
    * Several comments.
    * All but one easily resolved.
    * Will go for TCCC confirmation ballot this month.

- MAU Management Draft
  - Few changes from Feb-July 92 (through 3 meetings).
  - Going for first 802.3 ballot.

- 802.3 Layer Management (like Ethernet-like Interfaces MIB)
  - Being rewritten in ISO GDMO format.
  - Some rearrangement of conformance packages.
  - Frank Kastenholtz has details.

Repeater MIB Status

The Repeater MIB had been updated after the last meeting, circulated to the Working Group's mailing list, updated according to the few comments received, and forwarded to the IESG Area Director for Network Management with the Working Group's recommendation for advancement to Proposed Status. It was currently awaiting consideration by the Network Management Directorate.

Implementation experience feedback from attendees brought forth the following comments:

- The nonDisruptiveSelfTest description should be clarified to allow returning "ok" after doing only a trivial test.

- The setting of rptrReset to cause the Repeater to reset should allow the agent to delay the reset (for a short period) if it so wishes (e.g., to allow the SNMP Response to be transmitted.)

- It was suggested that the enumerated values autoPartitioned(1) and notAutoPartitioned(2) of the object rptrPortAutoPartitionState be changed to have the normal condition (notAuthPartitioned) have the value (1), since the normal condition of most other "state" MIB objects has the value (1).

It was agreed that the latter was preferable, but only if the MIB could be changed without causing any further delay in the approval of the MIB. The others would be remembered for updating the MIB at a later date (e.g., when moving from Proposed to Draft status).

Discussion of MAU MIB

A first draft of an IETF MAU MIB had been distributed at the last meeting, and a second draft mailed to mailing list early June. Donna presented the following overview of MAU management status and issues:

- 802.3 Medium Attachment Unit (MAU) attaches repeater port or Ethernet-like interface to the local network medium.

- MAU types include 10BASE5 (thick coax), 10BASE2 (thin coax), 10BASE-T (twisted pair), FOIRL and 10BASE-F (fiber optic).

- MIB information includes MAU type, link status, jabbering.

At the previous meeting, the issue of how to instantiate MAUs had been briefly discussed, and a preliminary decision made to have separate MIB objects for Etherlike Interface MAUs and Repeater MAUs. Before this meeting, Frank Kastenholz, editor of the Etherlike Interfaces MIB had expressed no interest in doing an Etherlike Interface MAU MIB. Thus, this meeting decided to develop one MIB having both a table for Repeater MAUs indexed by

"group.port.mau", and a table for Etherlike Interface MAUs indexed by "interface.mau", even though these tables would be remarkably similar except for their indexing.

The meeting then proceeded with a walk-through of the objects contained in the current MAU MIB draft (containing only the table indexed for Repeater MAUs). The following issues were raised:

- The optional grouping in this MIB is slightly different to IEEE's. No one had any specific comments, but were invited to think about this issue and post any comments to the mailing-list.

- The need for having multiple MAUs on a port was discussed, and it was suggested that choice of media, and having a backup MAU were two reasons for this. Thus, it was not thought to be necessary to allow there to be 65535 MAUs on each port !! It was agreed that (1..9) was a good range, since this allowed the MAU-number to be displayed as a single digit.

- It was suggested that rpMauType should be an OBJECT IDENTIFIER, and that this MIB should define values for the various types (i.e., not use IEEE-defined values).

- The overlap between rpLostMedias with an object in the Ethernet-interface MIB was discussed. Donna agreed to check further into this.

- An elaboration of the description of Jabber counters was requested, giving further definition and/or an IEEE reference.

- A number of other minor editorial changes were suggested and agreed.

Future Plans

A branch within the experimental subtree (as defined by the Internet SMI) would be requested for use by this MIB.

The editors agreed to produce an updated draft of the MAU MIB and mail it to the list for comments. After incorporating any further comments, it would be posted as an Internet Draft. If no further comments were forthcoming, the mailing-list would be polled to determine if there was consensus approval that the MIB was ready to be forwarded with the Working Group's recommendation for advancement to Proposed Standard status.

Attendees

David Arneson	arneson@ctron.com
Sam Ayers	s.ayers@sybus.com
Andy Bierman	bierman@davidsys.com

Jonathan Biggar	jon@netlabs.com
Jack Brown	jbrown@huachuca-emh8.army.mil
David Carr	dcarr@gandalf.ca
Greg Celmainis	gregc@newbridge.com
Chris Chiotasso	chris@artel.com
Dave Cullerot	cullerot@ctron.com
Michael Davison	davison@cs.utk.edu
Michael DiGioia	digioia@interlan.interlan.com
David Engel	david@ods.com
Fred Engel	engel@concord.com
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Doron Gartner	doron@chipcom.com
James Halpin	halpin@turkey.sw.stratus.com
Mark Kepke	mak@cnd.hp.com
Kenrick Kutzler	kkutzler@synoptics.com
David Lindemulder	dcl@mtung.att.com
Steven Lynes	lyness@rimail.interlan.com
Kim Mayton	mayton@wg.com
Keith McCloghrie	kzm@hls.com
John McKenna	mckenna@ralvm12.vnet.ibm.com
Donna McMaster	mcmaster@synoptics.com
David Minnich	dwm@fibercom.com
Paul Moran	Paul_Moran@3com.com
John Payne	jop@wang.com
Anil Rijsinghani	anil@levers.enet.dec.com
Sam Roberts	sroberts@farallon.com
Dan Romascanu	dan@lannet.com
Michael Sapich	sapich@conware.de
Jim Scott	scott@kali.enet.dec.com
Koichiro Seto	seto@hitachi-cable.co.jp
Timon Sloane	peernet!timon@uunet.uu.net
Einar Stefferud	stef@nma.com
Bob Stewart	rlstewart@eng.xyplex.com
Chris Sullivan	mm@gandalf.ca
Chuck Townsend	townsend@ctron.com
Stephen Tsun	snt@nsd.3com.com
Gerard White	
Joseph Zur	zur@fibhaifa.com

2.3.8 Internet Accounting (acct)

Charter

Chair(s):

Cyndi Mills, cmills@nnsc.nsf.net
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.
Done	Architecture submission as Internet Draft.
Jul 1992	Architecture submission as RFC.
Done	Architecture working papers written.

Internet Drafts:

"Internet Accounting Meter Services MIB", 07/09/1992, C. Mills, C. Brooks, A. Owen <draft-ietf-acct-metermib-00.txt>

"INTERNET ACCOUNTING: USAGE REPORTING ARCHITECTURE", 07/09/1992, C. Mills, K. Laube, G. Ruth <draft-ietf-acct-archreport-00.txt>

Request For Comments:

RFC 1272 "Internet Accounting: Background"

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

Minutes of the Internet Accounting Working Group (ACCT)

Document Status

- Internet Accounting Background: published as an RFC in November 1991.
- Internet Accounting Architecture: needs some clarification, otherwise it is ready for security review and publication as an RFC.
- Internet Accounting MIB: simplifications needed, not ready for publication as MIB yet.

The University of Auckland is implementing the MIB and will report on their progress at the next IETF. Nevil Brownlee submitted several suggested modifications which require clarification on the mailing list. It was agreed that the MIB will be published as an Experimental MIB in the hope that network management protocols will improve sufficiently in the next year to provide a better basis for a standard accounting MIB. In particular, bulk transfer and the reporting of data in traps are desirable.

A second implementation is needed. It was suggested that we solicit support from FARNET, publish an article in the Simple Times, and give a plenary technical presentation at the next IETF.

Semyan Dukach of MIT gave a brief presentation on has work on application-level billing protocols.

Future Direction

The Internet Accounting Working Group plans to hold its last active meeting in November, since with the publication of the three documents it has fulfilled its Charter. At that meeting the final recommendation to the IESG for the status (e.g., Draft Standard or Experimental) of the architecture and the MIB will be made. Any further activity after that will take place on the mailing list.

Possible future directions for future working groups include:

- Implementor's Working Group to improve SNMP MIB.
- SMP and/or SNMP over TCP/IP implementation.
- Application-level accounting extensions.
- Exploration of Tariffs for the Internet.

Attendees

George Abe	abe@infonet.com
Steve Bostock	steveb@novell.com
Scott Cossette	sdc@concord.com
Fred Engel	engel@concord.com
Michael Erlinger	mike@lexcel.com
Kathryn Fuller	kathryn@codex.com
Owen Gallagher	oweng@jjmhome.uucp
Pria Graves	priag@nsd.3com.com
J. Scott Marcus	smarcus@bbn.com
John McKenna	mckenna@ralvm12.vnet.ibm.com
Cyndi Mills	cmills@nnsc.nsf.net
Clifford Neuman	bcn@isi.edu
Bob Page	bob.page@eng.sun.com
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
Kary Robertson	kr@concord.com
Allan Rubens	acr@merit.edu
Gregory Ruth	gruth@bbn.com
Michael Sapich	sapich@conware.de
Robert Shirey	shirey@mitre.org
Chris Thomas	chris_thomas@ccm.hf.intel.com
Curtis Villamizar	curtis@ans.net
John Vollbrecht	jrv@merit.edu

2.3.9 Token Ring Remote Monitoring (trmon)

Charter

Chair(s):

Michael Erlinger, mike@lexcel.com

Mailing Lists:

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

Description of Working Group:

The Token Ring Remote Monitoring MIB Working Group is chartered to produce a new MIB specification that extends the facilities of the existing Remote Monitoring (RMON) MIB (RFC 1271) for use in monitoring IEEE 802.5 Token Ring networks.

The Token Ring RMON MIB extensions will be developed in the same architectural framework as the existing Ethernet-based RMON MIB. The original RMON MIB architecture was designed with the intention of incorporating MIB extensions devoted to monitoring other network media types. This Token Ring activity is the first attempt at such integration.

In creating the Token Ring Extensions the Working Group will, wherever possible, conform to terminology and concepts defined by relevant IEEE standards. It may be that a MIB devoted to monitoring may need to expand on the IEEE objects and definitions. Such modifications will be accompanied by a detailed rationale.

All work produced by the Token Ring Remote Monitoring Working Group will be consistent with the existing SNMP network management framework and standards.

Goals and Milestones:

Done	Discussion and agreement on models and terminology. Comparison of RMON architecture and Token Ring requirements. Assign author and editor responsibilities.
Mar 1992	Working Group Meeting at IETF. Present and confirm results of February meeting. Develop MIB draft. Publish initial version as Internet Draft.
Jul 1992	Working Group Meeting at IETF to discuss and revise draft of Token Ring Extensions. Publish revised version as Internet Draft.
Nov 1992	Working Group meeting to discuss and reach closure on Token Ring MIB Extensions MIB. Publish agreed version MIB as Internet Draft. Make Working Group recommendation on Token Ring Extensions MIB.

CURRENT MEETING REPORT

Minutes of the Token Ring Remote Monitoring MIB Working Group (TRMON)

Report not submitted. Please refer to the Area Report for a brief summary.

Attendees

David Arneson	arneson@ctron.com
Zareh Baghdasarian	zareh@tekelec.com
Andy Bierman	bierman@davidsys.com
Jonathan Biggar	jon@netlabs.com
Steve Bostock	steveb@novell.com
Christopher Bucci	bucci@ngc.com
Isaac Choi	ichoi@synoptics.com
Dave Cullerot	cullerot@ctron.com
Cathy Cunningham	cmc@microcom.com
David Engel	david@ods.com
Fred Engel	engel@concord.com
Michael Erlinger	mike@lexcel.com
Jeff Erwin	Jeff_Erwin@protools.com
Bill Fardy	fardy@ctron.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Shawn Gallagher	gallagher@quiver.enet.dec.com
Rob Graham	robert_graham@protools.com
Martin Gray	mg@spider.co.uk
Daniel Hansen	dan@ngc.com
Jeff Hughes	jeff@col.hp.com
Robin Iddon	axon@cix.compulink.co.uk
Mark Kepke	mak@cnd.hp.com
Kenrick Kutzler	kkutzler@synoptics.com
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Kim Mayton	mayton@wg.com
John McKenna	mckenna@ralvm12.vnet.ibm.com
Rohit Mital	rm@protools.com
Patrick Mullaney	mullaney@ctron.com
Tom Nisbet	nisbet@tt.com
Roger Osmond	bytex!rfo@uunet.uu.net
John Payne	jop@wang.com
Donald Perkus	dperkus@chipcom.com
Narendra Popat	albin@frontier.com
Dan Romascanu	dan@lannet.com
Michael Sapich	sapich@conware.de

Michael Scanlon	scanlon@interlan.com
Jim Scott	scott@kali.enet.dec.com
Koichiro Seto	seto@hitachi-cable.co.jp
Bruce Taber	taber@interlan.com
Mark Therieau	markt@python.eng.microcom.com
Steve Thompson	steve@ods.com
Chuck Townsend	townsend@ctron.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Desikan Venkatrangan	venkat@geoduck.matrix.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Albin Warth	albin@frontier.com
Kiho Yum	kxy@nsd.3com.com
Uwe Zimmerman	metrix!uwe@uunet.uu.net
Joseph Zur	zur@fibhaifa.com

2.3.10 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.
Done	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 LAPB", 10/07/1991, Dean Throop, Fred Baker <draft-ietf-x25mib-lapbmib-05.txt>

"SNMP MIB extension for MultiProtocol Interconnect over X.25", 10/07/1991, Dean Throop <draft-ietf-x25mib-ipox25mib-03.txt>

"SNMP MIB extension for the X.25 Packet Layer", 10/07/1991, Dean Throop <draft-ietf-x25mib-x25packet-05.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 July 16th at the IETF meeting in Cambridge. The status of the LAPB and X.25 MIBs were discussed. The SNMP Directorate has reviewed the MIBs and they requested some changes to x25CallParmIndex. These changes should simplify the description. They also requested a number of other minor changes. Updated versions of those documents will be forthcoming.

The Working Group also discussed the MultiProtocol Interconnect over X.25 MIB. The Group approved the draft without any changes and the document has been referred to the Network Management Area Director.

Attendees

Cathy Cunningham	cmc@microcom.com
Michael Davison	
Zbigniew Kielczewski	zbig@eicon.qc.ca
Andrew Malis	malis@bbn.com
Rodney Thayer	rodney@world.std.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com

2.4 OSI Integration Area

Director(s):

- David M. Piscitello: dave@sabre.bellcore.com
- Erik Huizer: Erik.Huizer@surfnet.nl

Area Summary reported by Dave Piscitello/Bellcore and Erik Huizer/SURFnet

The OSI area contains the following working groups:

NOOP	Network Osi Operations
MPSNMP	SNMP over a Multi-protocol Innternet
OSI-DS	OSI Directory Services
MHS-DS	Message Handling Service Usage of Directory Services
X.400OPS	X.400 Operations
MIMEMHS	MIME to MHS Mapping
ODA	Office Documentation Architecture

The OSI General Working Group has been disbanded.

Related working groups:

DISI	Directory Information Services Infrastructure Working Group (report under User Services area)

BOFs in the OSI Integration Area held in Boston.

SWIP	Shared Whois Information Project
UDI	Universal Document Identifiers

Related BOF:

NIR	Networked Information Retrieval BOF (report under User Services area)

Shared Whois Information Project BOF (SWIP)

This BOF was organised by Merit to discuss the possibilities for using X.500 to set up a shared whois like service between the Major network coordination centers (currently there are 3: Ripe NCC, GSI- NIC, Merit) in the Internet. This is meant for easy access and exchange of network management data. Which ip address belongs to who, what point of contact, etc.

The goals of the SWIP BOF were to a) present the idea and project that Merit had conceived of to converge the network data stored by GSI-NIC, RIPE, and Merit. b) get general agreement on the idea and the method being used c) define requirements for a shared whois database d) get consensus on the need for a distributed whois database of networks and a consensus that the platform be X.500.

Most of these goals were achieved. There was a very clear consensus from the attendees that a distributed whois database of networks should be implemented, should be done in X.500, and that it should be done "right". It was further decided that Merit should proceed with their X.500 project to converge the network data currently available from RIPE, GSI-NIC, and Merit, and for them to put into place a procedure to keep the data converged until the distributed whois database is in place and working. There is an action item to combine the two X.500 architectural models presented in the bof pertaining to a distributed model for network data.

Universal Document Identifiers BOF (UDI)

This Group discussed naming issues intended to support the discovery and access of resources in an Internet environment. It was agreed that the term "Uniform Resource Locator" (URL) would be used to refer to standardized identifiers which specify location information for resources. The discussion of other aspects of the naming problem was deferred until a later meeting.

A document written by Tim Berners-Lee (timbl@info.cern.ch) proposing a standard for URLs was discussed and the syntax and general content of the document was accepted with some revisions. The revised document will be made available from info.cern.ch and circulated to the list below for further discussion.

The Group decided to draw up a charter and form an IETF working group on this issue. The mailing list for discussion of URL design issues will be "ietf-url@merit.edu". This list will be archived on the anonymous FTP archive on "merit.edu".

MHS-DS Working Group (MHSDS)

The MHS-DS Working Group met at JENC-3 in Innsbruck, Austria in May. A small group of technical experts met once to discuss editorial and technical revisions to the set of seven Internet Drafts being written by the Group. In addition, an open meeting of MHS-DS was held to present general concepts to a broad cross section of the European R&D

community. An open discussion followed, and valuable comments were contributed to the Internet Drafts.

The focus of the third meeting of MHS-DS (Boston) was on editing the seven Internet Drafts (listed below). We went through the documents, page by page, and contributed both simple editorial changes as well as some recommendations for minor technical revisions. As a result, three of the documents will be progressed as Experimental Standards, and the other four will be cycled through another round of review after they are revised. In addition, two new documents will be produced: a general overview of the whole set, and a document which focuses upon the subject of Content Conversion.

The document status follows:

1. Representing Tables and Subtrees in the Directory
   status: revise and progress as Experimental Standard
2. Representing the O/R Address Hierarchy in the Directory Information Tree
   status: revise and progress as Experimental Standard
3. MHS use of Directory to support MHS Routing
   status: revise and cycle as an Internet Draft
4. Use of the Directory to support mapping between X.400 and RFC 822 Addresses
   status: revise and progress as Experimental Standard
5. MHS use of the Directory to support distribution lists
   status: revise and cycle as an Internet Draft.
6. A simple profile for MHS use of Directory
   status: revise and cycle as an Internet Draft (depends upon 3)
7. Use of the Directory to support routing for RFC 822 and related protocols
   status: revise and cycle as an Internet Draft.

New documents to be produced:

1. Overview of Document Set.
2. MHS use of the Directory to support Content Conversion.

As a final note: The MHS-DS Charter will be revised to add the two new documents and also to add the following two features:

1. MHS-DS will coordinate piloting of MHS use of the Directory.
2. MHS-DS will specify requirements for tools which facilitate interworking between X.500-capable MTA's and MTA's which are not X.500-capable.

MIME-MHS Interworking Working Group (MIMEMHS)

There have been two papers produced since the last meeting:

1. X.400/MIME body equivalence Harald Tveit Alvestrand, Steven Thompson.
2. Mapping between X.400 and RFC-822 Message Bodies, Harald Alvestrand et al.

Several mappings have been defined, and for those without a clear X.400(88) equivalent there is a trapdoor/catchall External bodypart defined in X.400: EBP-mime-body-part. In the other direction the trapdoor in MIME is a new Mime subtype: application/x400-bp. The issues are:

- How to get vendors to register OIDs as well as the equivalent MIME subtype with the IANA.
- How to manage IANA registration of different versions of BPs like WP5.0 and WP5.1.
- How to handle mapping in X.400(84)?
  - Simplest case single BP IA5.
  - T.61 strings in header vs RFC 1327 needs resolving.
  - Three-party mail issue (mime-X.400(88)-X.400(84)).
- Automatic OID assignment for registered subtypes.
- Appendix with OIDs defined.
- Security: viruses will be gatewayed too, not solved in this paper.
- Criticality of header extensions.

The issues will be resolved by E-mail in the next couple of months. Both documents will be forwarded as Proposed Standard RFCs.

Network OSI Operations Working Group (NOOP)

The Group reviewed the status of RFC 1139, CLNP ping, and agreed to, (a) eliminate the short-term solution, and (b) align/revise the long-term, solution to match the ISO PDAM expected from ISO next week. A new RFC will be produced. Since this is an integral part of the tools RFC, NOOP expects to process this rapidly.

Work continues on the OSI Tools RFC. The Group also reviewed a list from RARE identifying the ten most desirable Managed Objects from the CLNP MIB, and reacted favorably to the selection of OSI connectionless transport as a means of mapping SNMP onto OSI.

The Working Group reviewed the ISO Transport MIB submitted by Russ Blaesing. Following a discussion of what and how many managed objects would be useful for network operations, Dave Piscitello agreed to evaluate this MIB against MIB-II, TCP Group. He

will post the results of this comparison to NOOP mailing list. NOOP will then discuss what MOs are required for operations, and will make this set known to vendors.

The Working Group received a presentation of TUBA from Ross Callon; of Interop '92 spring experiences from Rich Colella; and of X Window System over OSI and the "skinny OSI stack" from Jim Quigley.

OSI Directory Services Working Group (OSIDS)

Discussion topics:

- The latest1992 CCITT X.500 version is dated 25-12-1992 (The Christmas paper).
- RFC-1279 (Representing DNS in the Directory) has been revised.
- OIW established a new specification: IGOS (industry and Government OSI Specification), it requires a.o. many X.500 1992 extensions.
- NADF split the Naming Schema document into two docs:
  - Naming schema set-up for countries.
  - Specific case for the US.
- NADF security paper, protection by passwords, weak credentials. There is a defect in the replay of passwords in simple auth. (fixed in 1992 version?)
- There is a road-map paper indicating all NADF publications.
- None of (12) vendors in NADF was supporting strong auth, none had timelines for 1992 extensions.

Documents discussed:

- Naming guidelines for Directory pilots paper to be progressed to an Informational RFC.
- A string representation of distinguished names to Proposed Standard.
- User Friendly Naming to Experimental RFC.
- Strategy Document. Those who read it (ca 60document. The document will be redistributed after processing minor comments and then submitted to IESG/IAB for policy approval, and subsequent publishing as an Informational RFC.
- IP address information in the Directory. The paper was discussed and several major changes were suggested. Work ongoing.

Pilots:

- QOS no progress yet.
- JPEG ongoing.
- DIT counting.
- Char set ongoing.

A Schema group has not yet been set up. A suggestion was made to ask the IANA to take this on.

- Discussion on DUA and DSA metrics papers. Meant to set metrics for comparing DSAs and DUAs (functionality, capacity etc.) The papers will be used to describe existing implementations and results/comments on the paper will be reported back into the next meeting. Papers will then be revised and put forward as informational RFCs.

- Two papers on a lightweight access protocol for the directory were discussed. Minor comments were given, after these have been worked into the paper it will be submitted to the IESG for publication as Proposed Standard.

- DSA naming. This paper was discussed. The paper is seen as being still too much oriented towards one single implementation, to be publishable as an RFC. Therefore the Working Group will drop the issue until other manufacturers have commented or supplied their solution to DSA naming and knowledge distribution.

Office Document Architecture Working Group (ODA)

Progress of products: Six products now known to the Group, most of them not yet with full vendor support.

Progress on pilots: Few groups put up some of products. Use is mostly internal and limited.

Expectations: Over next year the int. profile FDD 26 is being ratified and will probably lead to new products and pilots. However in the next six months little progress is expected, so the IETF ODA Working Group will not meet in November 92, but will sleep until new activity pops up.

SNMP over a Multi-protocol Internet Working Group (MPSNMP)

The Group met and reviewed three Internet Drafts.

- SNMP over OSI (CL Transport Service)
- SNMP over Appletalk
- SNMP over IPX

All three were aligned with respect to the treatment and assignment of Object Identifiers for

the transportDomain All three have at least one implementation presently. It was agreed that all three would have essentially the same boilerplate recommendation with respect to multiple transport implementations; i.e., that agents are only required to implement *one* transport mapping of SNMP, and managers were expected to implement as many as necessary to allow communication to all agents within a network. Implementations are encouraged to implement UDP.

All three documents require minor rewrites; they will all be posted for a two week last call before recommending to the IESG that they be moved to Proposed RFCs.

A fourth document describing "how to write a transport mappings" was aligned with the three Internet Drafts; this will be revised and submitted as an FYI RFC.

With no further work to consider, the Working Group agreed to disband.

X.400 Operations Working Group (X400OPS)

Status report on the pilots (XNREN and Cosine MHS) was given. The amount of usage as well as the amount of connected MTAs is growing steadily.

Work on daily update tool for outing and mapping tables is still ongoing and expected to be ready by the end of 1992.

Connectivity issues:

- Internet-X.400 to public X.400;
- RFC-822 to public X.400; Various unstandardised gateways from commercial service providers (e.g., AT&T, MCI, IBM) to the Internet were discussed. These gateways cause problems like address mangling.
- In the US rather than ADMD=<space> they have proposed an ADMD=usbb, to register nationwide-multiple-carrier PRMDs. So if a prmd subscribes to e.g., ADMD=attmail, they can choose to do so under ADMD=attmail or request from AT&T to do it under ADMD=usbb.

Documents

- Proposed in this text is to use the X.400/88 GeneralText option to use extended character sets. This option is not really in X.400/88, but only in the ISO version, however it is an extended bodypart and thus can be used without modifications. The paper further describes the ISO 8859 character sets that should be used.

  NOTE: o.a. the Dutch ligature ij is missing!! Jammerlijk maar geen ramp.

  Paper will be revised to the comments made and discussed with char set experts (e.g., RARE wg-char), and then it will be put forward to RARE wg-mhs and the November

IETF meeting.

- Coordination Procedures for RFC 1327 gateways by Cosine MHS The paper documents current procedures for reference and with the purpose of making it more globally known. There is no information in the paper on how the tables should be formatted and what can and cannot go in. This will be a separate paper. However this paper should still contain some general indications. The paper will be adapted to the comments and then put forward as informational RFC.

- Operational requirements.... - Rob Hagens/Alf Hansen Minor comments were made to this document. It will be published as an Experimental Standard RFC.

- Routing coordination for X.400 Urs Eppenberger The document is almost finished. However a new perspective has been brought in by GMD (Panos G.) to allow for more X.500 oriented syntaxes in the document. Panos, Steve Hardcastle-Kille and Urs will discuss this off-line. Time pressure is high. To be progressed to Experimental Standard RFC soon.

- Using DNS to maintain RFC987 mapping tables - Claudio Allocchio This paper shows the various issues that have come forward out of the Trieste experiments with the use of DNS. The paper provides an independent way to distribute the table WITHOUT distributing necessarily the authority. It is proposed that the various alternatives will be put forward to the IETF DNS Group for advise, and that following that the paper will be progressed to an Experimental RFC

- Mapping between X.400 and Mail-11 - Claudio Allocchio There was unfortunately no time to discuss this paper. There is one implementation around, discussion will be done by E-mail.

- New document: Grandfathering of ADMD=internet in the US, to be produced.

- New document: The use of s=postmaster in X.400, to be produced.

Next meeting at November IETF.

CURRENT MEETING REPORT

Reported by Sheri Repucci/Merit

Minutes of the Shared WHOIS Project BOF (SWIP)

The goals of the SWIP BOF were to:

- Present the idea and project that Merit had conceived of to converge the network data stored by GSI-NIC, RIPE, and Merit.
- Get general agreement on the idea and the method being used.
- Define requirements for a shared WHOIS database.
- Get consensus on the need for a distributed WHOIS database of networks and a consensus that the platform be X.500.

We achieved most of these goals. There was a very clear consensus from the attendees that a distributed WHOIS database of networks should be implemented, should be done in X.500, and that it should be done "right". It was further decided that Merit should proceed with their X.500 project to converge the network data currently available from RIPE, GSI-NIC, and Merit, and for them to put into place a procedure to keep the data converged until the distributed WHOIS database is in place and working. There is an action item to combine the two X.500 architectural models presented in the BOF pertaining to a distributed model for network data.

Attendees

Cyrus Chow	cchow@orion.arc.nasa.gov
James Conklin	jbc@bitnic.educom.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Russ Hobby	rdhobby@ucdavis.edu
Tim Howes	Tim.Howes@umich.edu.
Mark Knopper	mak@merit.edu
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Thomas Lenggenhager	lenggenhager@switch.ch
Bill Manning	bmanning@rice.edu
April Marine	april@nisc.sri.com
Sheri Repucci	smr@merit.edu
Jim Romaguera	
Srinivas Sataluri	sri@qsun.att.com
Mark Smith	mcs@umich.edu
Marten Terpstra	terpstra@ripe.net

Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Russ Wright	wright@lbl.gov
Yung-Chao Yu	yy@qsun.att.com

CURRENT MEETING REPORT

Reported by Tim Berners-Lee/CERN

Minutes of the Universal Document Identifiers BOF (UDI)

The UDI BOF was held at the 24th IETF in Cambridge, MA, USA on July 14, 1992.

Introduction

Tim Berners-Lee opened the BOF with a summary of the terms used in the discussion to date. The information one quoted in a reference to an object could comprise many things, among which were possible one unique name, (Unique Resource Number, URN was one acronym), and zero or more addresses (Uniform Resource Locators or URLs) which gave instructions for retrieving the object.

The purpose of the meeting was to formalize a standard string syntax for URNs and URLs in general, and to define specific syntaxes for addresses in the namespaces of each of the existing network protocols. [There was a discussion on acronyms at various times. URL was decided upon for an address, and that is used throughout these minutes for clarity.] The result should be a standards track document (requiring a working group, which should probably be in the Directory Area but could be in Applications).

NOT to be discussed were the differences between names and addresses, URN schemes (which are not yet well enough defined), the full set of information to be given in a reference, or IPv7.

To be discussed were the overall string syntax, including allowed characters and escaping systems for unallowed characters, the order of components (little/big-endian), punctuation characters, as well as the particular prefix to be used to identify each namespace.

Specific schemes should be handled in appendices of the resulting document, and should include, Prospero, FTP, WWW, telnet, net man. db?, nntp, WAIS, gopher, finger, X.500.

Discussion

We need methods of keeping up to date the set of appendices without the same standards track procedure which applies to the full document.

It was pointed out that for WAIS one could imagine a separate name space for databases and for documents. If this was taken further, a separate prefix would be used for each type of object. It was on balance agreed that this could go too far. One prefix should be used per protocol, but it should be made clear how to determine the type of an object from the URL.

Peter Deutsch is concerned that we need a syntax for attaching a URN to a URL, but accepted that it was not for discussion at the BOF.

Cliff Lynch suggests three part structure of name, address, and other stuff. Peter Deutsch suggests that the document should talk about what the addresses aren't, in a section on Scope. This section could also provide an example of a complete reference, including other information, by way of explanation but not recommendation.

Tim Berners-Lee had submitted in the background document, the W3 implementation:

" ____scheme_specific_stuff___"

(the syntax of ____scheme_specific_stuff___ depending on the value of scheme, within certain constraints.) There was no dissent, although he noted that this is the reverse order from the WAIS proposal.

John Curran has concerns about URLs being resolvable in such a way that any two references to the same URL get the same thing. (unambiguity). It was generally felt that the system W3 uses to allow URLs to be incompletely quoted in context was an application issue and was not relevant.

The issue of what we are identifying came up "resource locator"? – a scheme for somehow identifying resources. Perhaps identifying procedure for locating a resource (Karen Sollins and Cliff Lynch). Cliff Neumann suggested Document Access Instructions as an alternative handle/name/identifier for these addresses. URL was decided on by an almost unanimous vote. (Uniform Resource Locator).

Peter Deutsch pointed out that we want to focus on interoperability, not on longevity. We ought to be able to hand URLs around in short-term, but not long-lived. URLs are not unique (in the sense that one document may have several). This should be made clear in the document.

The class of object you get back should be predictable (Cliff Lynch). W3 has a real problem with that, since everything is a "document" and handled in a similar way. Might get a pointer to a database in a piece of mail. The question of whether one gets back a file or a directory from a FTP URL arose. Archie really wants to know what it is getting back. Within a scheme, should be documented syntax that will clarify which sort of object will come back. If we go too far down this track, we fairly quickly get to full object-oriented world, with fullscale typing. Alan Emtage. suggested that simple enumeration of acceptable types. Extensions based on documented new subtypes, based on documented protocols.

A separate issue of whether human or only machine readable. Previously, included issue of printable. This is needed because don't have names now. Question arose of whether once these addresses exist will be replaceable with names - will be presented as new functionality, not replacing existing systems. Agreement on some way of specifying class of objects.

"Context" Prefix

IT WAS AGREED that the context, or namespace prefix be the first (leftmost) part of the

URL, and be separated from the rest of the URL by a colon.

The punctuation "//" was discussed. Currently, W3 URLs use "@" for login information. An extension of server the server hostname can include a port number in all current schemes.

The issue of how to manage additional schemas was discussed. Each appendix should be checked out by a particular group within the IESG, and perhaps should be an "Experimental Protocol," rather than simply an "Informational RFC." The document will describe how to write appendices.

Syntax Details

The syntax should be human typable (majority agreement).

Should one use punctuation, or attribute-value pairs? Attribute value pairs get mispelt. (note x.400 vs. Internet addresses)

It was decided to use a short string with punctuation rather than an attribute-value pair system.

We must specify the terminator (by declaring some characters as illegal inside a URL).

Is a URL nestable? If one URL can contain another, one needs nestable begin-end pairs. (Alan Emtage). Currently W3 URLs are not nested visibly although escaping allows URLs to be encapsulated within URLs, for example by gateways.

Allowed characters: characters should be disallowed if they are needed as terminators ('"', ';') or are too easily mutialtable by passage though (for example ASCII/EBCDIC/ASCII) gateways (tilde, backslash). A subset of an ISO 7-bit code should be defined, with reference to MIME work.

Future Discussion

Mailing lists: NIR list at McGill will be used by Jill Foster and George Brett's NIR BOF. ietf-url@merit.edu will be used by this Group. First of all, we should ask Mike Schwartz whether he is willing to run all mailing lists on one machine (at least nir and url) in order to cut down on multiple copies opf cross-posted messages.

Accomplishments

Things which the meeting had brought to light included:

- "File:" is too broad a description, "ftp:" would be better. If a given client knows that it can in fact access some FTP sites as local files, that is a local client issue.
- Escaping is to be defined.

- Relative naming is a client issue.
- We should look at what we call "news:" ("usenet:"?).
- We should we be able to tell what sort of an object we have (eg database or document) by simple examination of the URL
- We need a scheme for managing the addition of new schemas (cf Directory object definitions).
- The document should be an "Experimental Protocol" RFC.
- Mailing lists will be defined and a single message sent to cniarch to say what is happening.
- The time-frame for the document was: soon. Probably in the OSI Directory Services Area (Erik Huizer has suggested this). We need a working group if we want to go through the IESG. But in reality, if have big applications buying in, then will be a de facto standard.

Final Comments

We may need to separate WAIS from Z39.50. (C Lynch). We may also need SQL. Also we should include X-junk: type extension mechanism for experimental schemes.

These minutes noted on-line by Karen Sollins (Thanks!) and edited by Tim Berners-Lee.

Attendees

Vikas Aggarwal	aggarwal@nisc.jvnc.net
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Mark Baushke	mdb@cisco.com
Tim Berners-Lee	timbl@info.cern.ch
George Brett	ghb@jazz.concert.net
Mitchell Charity	mcharity@lcs.mit.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Curran	jcurran@bbn.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Ray Freiwirth	5242391@mcimail.com
Jim Fullton	jim_fullton@unc.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu

Erik Huizer	huizer@surfnet.nl
Bill Manning	bmanning@rice.edu
Clifford Neuman	bcn@isi.edu
Sam Nicholson	scion@pblx.knox.tn.us
Jon Rochlis	jon@mit.edu
Karen Roubicek	roubicek@faxon.com
Richard Schmalgemeier	rgs@merit.edu
Vincent Sgro	sgro@cs.rutgers.edu
Jane Smith	jds@jazz.concert.net
Karen Sollins	sollins@lcs.mit.edu
Simon Spero	ses@cmns.think.com
Chris Weider	clw@merit.edu

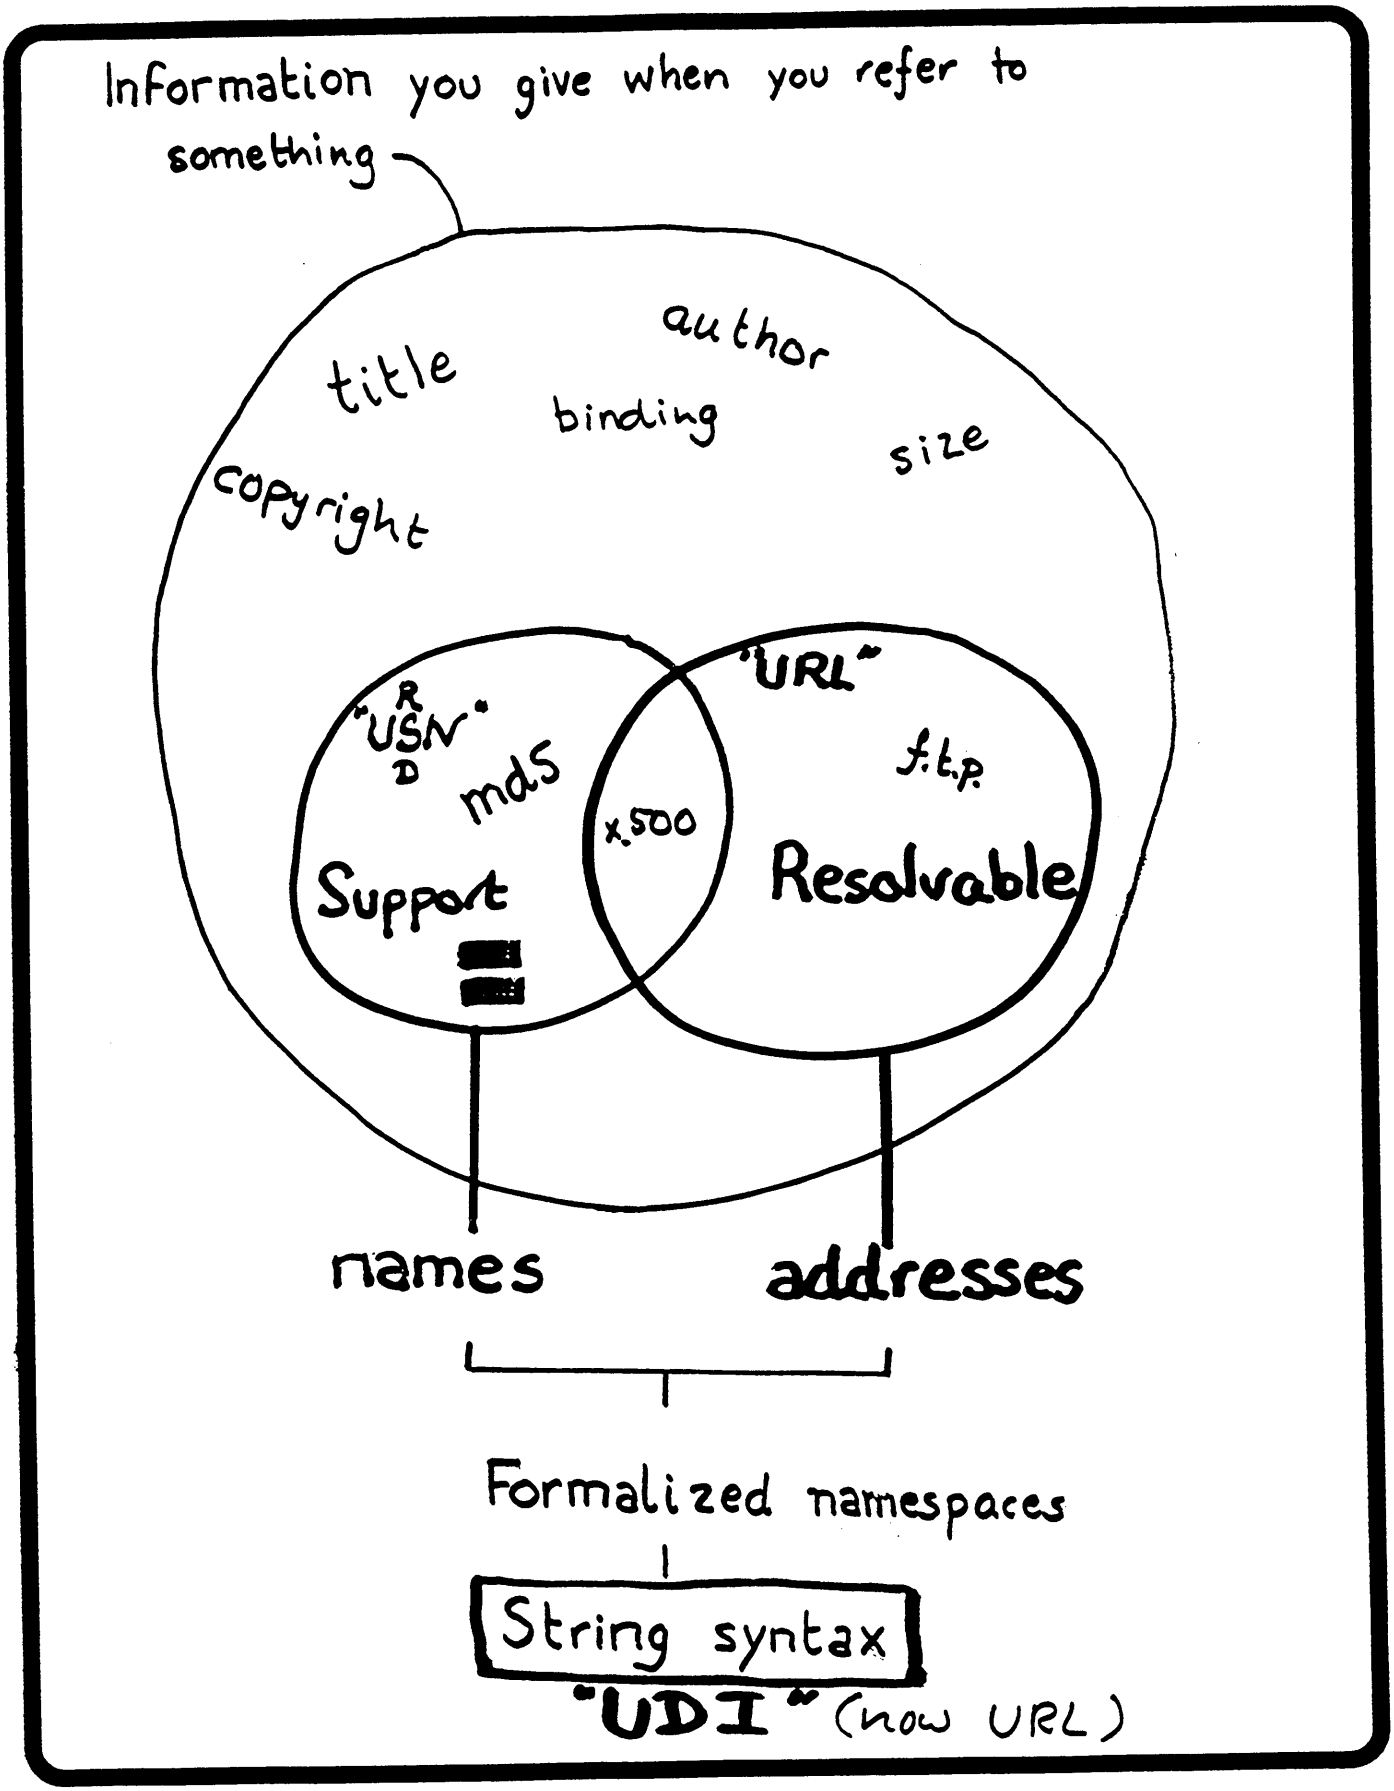
Information you give when you refer to something
author
title
binding
size
copyright
"URL"
"URN"
R
D
mdS
ftp
x.500
Support
Resolvable
names
addresses
Formalized namespaces
String syntax
"UDI" (now URL)

CURRENT MEETING REPORT

Reported by Jim Quigley/Hewlett-Packard

Minutes of the XWindows over OSI and Skinny Stack BOF (THINOSI)

Jim Quigley of Hewlett-Packard gave an hour plus presentation on the OSI Basic Communication aspects of OSI, aka "skinny stack." Quigley also discussed the University of London's prototype of X over an OSI "skinny stack" (Xosi) including performance data. Xosi is part of the ANSI standard for X that is also included in the U.S. IGOS (Industry and Government Open System – formerly GOSIP) requirements.

The group attending the BOF generally accepted Quigley's premise that most [if not all] potential OSI applications only have the basic communication requirements of connect/disconnect and send/receive. Such applications only require an OSI "skinny stack" such as that defined for Xosi. Besides the specification of OSI skinny stacks, the critical need is the specification of OSI 7-layer "skinny" APIs.

The Group agreed that Peter Furniss and Steve Hardcastle-Kille should be requested to do a Simple OSI Stack (SOS) versus OSI skinny stack comparison. The results would be published to the Skinny Stack BOF mailing list.

The consensus of the Group was to pursue the formation of a working group for the specification of APIs for an OSI skinny stack.

Jim Quigley was instructed to ask OSI Area Directors Dave Piscitello and Eric Huizer how to start such a working group on XTI and SOCKETS APIs for an OSI skinny stack.

It was agreed that Jim Quigley should get on the next IETF: technical presentation to discuss the requirements and potential of OSI skinny stacks and APIs.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
C. Allan Cargille	cargille@cs.wisc.edu
Cyrus Chow	cchow@orion.arc.nasa.gov
Richard Colella	colella@osi.ncsl.nist.gov
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Jisoo Geiter	geiter@gateway.mitre.org
Kevin Jordan	kej@udev.cdc.com
Mark Prior	mrp@itd.adelaide.edu.au
James Quigley	Jim_Quigley%YD@hp6600.desk.hp.com
David Reiser	dsr@codex.prds.cdx.mot.com
William Warner	warner@ohio.gov

2.4.1 MHS-DS (mhsds)

Charter

Chair(s):

Kevin Jordan, kej@udev.cdc.com
Harald Alvestrand, Harald.Alvestrand@delab.sintef.no

Mailing Lists:

General Discussion: mhs-ds@mercury.udev.cdc.com
To Subscribe: mhs-ds-request@mercury.udev.cdc.com
Archive: pub/archives/mhs-ds-archive:mercury.udev.cdc.com

Description of Working Group:

The MHS-DS Group works on issues relating to Message Handling Service use of Directory Services. The Message Handling Services are primarily X.400, but issues relating to RFC 822 and RFC 822 interworking, in as far as use of the Directory is concerned, are in the scope of the Group. Directory Services means the services based on X.500 as specified by the OSI-DS Group (RFCs 1274, 1275, 1276. 1277, 1278, 1297). The major aim of this Group is to define a set of specifications to enable effective large scale deployment of X.400. While this Group is not directly concerned with piloting, the focus is practical, and implementations of this work by members of the Group are expected.

Goals and Milestones:

Mar 1992 Define a set of service requirements for MHS use of Directory. This should include: support for routing; support for security services; support for user agent capabilities; support for distribution

Mar 1992 Define a work program for the Group, to write a set of RFCs to meet the service requirements. As far as possible, reference should be made to existing standards.

Dec 1992 Release RFCs meeting the service goals. This target should be refined in the light of specifying the service goals.

Ongoing Liaisons should be established with similar groups working on X.400 and X.500, i.e., RARE WG1 and RARE WG3, IETF OSI-DS and IETF X.400.

Internet Drafts:

"MHS use of the Directory to support distribution lists", 04/09/1992, S. Hardcastle-Kille <draft-ietf-mhsds-mhsuse-01.txt, .ps>

"Use of the Directory to support routing for RFC 822 and related protocols", 04/09/1992, S. Hardcastle-Kille <draft-ietf-mhsds-822dir-01.txt, .ps>

"A simple profile for MHS use of Directory", 04/09/1992, S. Hardcastle-Kille <draft-ietf-mhsds-mhsprofile-01.txt, .ps>

"Representing Tables and Subtrees in the Directory", 04/09/1992, S. Hardcastle-Kille <draft-ietf-mhsds-subtrees-01.txt, .ps>

"Representing the O/R Address hierarchy in the Directory Information Tree", 04/09/1992, S. Hardcastle-Kille

<draft-ietf-mhsds-infotree-01.txt, .ps>

"Use of the Directory to support mapping between X.400 and RFC 822 Addresses", 04/09/1992, S. Hardcastle-Kille <draft-ietf-mhsds-supmapping-01.txt, .ps>

"MHS use of Directory to support MHS Routing", 04/17/1992, Steve Hardcastle-Kille <draft-ietf-mhsds-routdirectory-01.txt, .ps>

INTERIM MEETING REPORT

Reported by Kevin Jordan/CDC

Minutes of the MHS-DS Working Group (MHSDS)

May 11, 1992

The following documents written by Steve Hardcastle-Kille are to undergo serious reviewing by the MHS-DS Group:

1. Representing Tables and Subtrees in the Directory.
2. Representing the O/R Address hierarchy in the Directory Information Tree.
3. MHS use of Directory to support MHS Routing.
4. Use of the Directory to support mapping between X.400 and RFC 822 Addresses.
5. MHS use of the Directory to support distribution lists.
6. A simple profile for MHS use of Directory.
7. Use of the Directory to support routing for RFC 822 and related protocols.

Classification by Harald Tveit Alvestrand

1. Non controversial
2. Simple.
3. This is the principal document. You need to understand 1 and 2 first.
4. Simple use of 1 and 2.
5. Simple 'repairing' of an X.400(88) bug.
6. Terse. You must understand 1, 2 and 3 first! It tells you what you need to implement as a minimum.
7. Terse. This document is about applying the algorithms of 3 to Internet-like networks.

In addition, document 3 currently covers issues related to content conversion. Harald has recommended that this be removed from document 3 into a new document (8).

Harald also gave a short tutorial on the documents. (This does not replace a serious reading of the documents by everybody but helps understanding the overall concept.)

Parallel efforts in ISO

There exists also a proposal from Robert Willmoet, UK, distributed within OSI. It has a different concept. Steve Hardcastle-Kille intends to forward his proposal also to ISO and CCITT. There is not much hope to get a stamp by CCITT since the proposal allows to bypass the ADMD infrastructure. There is at least some hope to get it accepted within

ISO.

Secretary's note: (I learned in the mean time that DEC will release a X.400(88) SW with X.500 usage for MHS routing and mapping (Dec 1992). I do not know the methods they use but it is probably different from Steve's and Robert's proposals. I even do not know how the mapping works: DEC, RFC987, RFC1148, RFC1327...)

Security

There is the still unsolved problem of the need of bilateral agreements to make the X.400 service a little bit more secure than SMTP. (There is the possibility of strong authentication in the applications.) It is proposed to spend some time on inserting more security functions into Steves documents. The hooks are already built in (see doc 3, chapter 20).

PAP: The researchers in the internet are happy with what they have now. With any new solution we should take the industrial community into account. Authentication is an important topic there.

Questions and Answers

A part of the discussions was based on questions put forth by the meeting participants, and answers were provided mainly by those who have read Steve's papers.

Will people implement it? Yes, Control Data and the ISODE Consortium will lead to two implementations. Other commercial suppliers know that they must sort out the routing problem, so there might be a chance to get more implementations. A solution for the EAN software depends on funding.

Will it work and be fast enough? Probably.

Will people keep the data up to date? Yes, as soon as it is needed for proper operation.

What happens with an email address containing ADMD='single space'? This is solved. See doc 2 fig 1.

How to use trees? How many trees are needed? This is very open. Practical experience will show what is reasonable. Doc 3 drafts several solutions. Users will probably also start to set up private subtrees, on a service provider level as within networks, or even within universities where they do not want to publish user information in the public tree.

How to control access to the trees? What info should go to the open tree? Even UA information can be stored in the open tree. It is possible to set ACLs such that UA/user information is hidden. The routing algorithm will select routing information from higher up in the tree.

What is the behaviour of the algorithm, if it gets errors back on queries? This may happen due to ACLs set, or unavailable DSAs,... We need a list of possible failure reasons and a

definition of the appropriate behaviour of the algorithm.

Where to locate the open tree? Discussion lead to the proposal to put the routing tree under @O=x400routing. This gives independence from the owners of the country entries.

How to optimise access to the DIT? Attribute inheritance was found a good solution to optimise DIT access. Reading an attribute at the end of the DIT will then also give back the routing info stored higher up. Attribute inheritance is not defined in X.500 yet.

How does a network without DSA access send mail to the DSA-documented world/MTAs? Such a network will have to have an agreement with an MTA which has access to DSAs. At a first glance this seems feasible. In addition, it will probably be necessary to develop tools which extract routing information from the directory and generate static tables like the tables currently used by the RARE X.400 community. These tools will allow existing X.400 implementations to migrate toward direct usage of X.500 over time.

How does routing through body type converters work? It is proposed to take this item out of the documents until the problems are better understood.

A pilot

Alf Hansen proposes to start a pilot project to check if the ideas really work. A very simple implementation will be included in the next public release of PP. The pilot should take into consideration:

- The co-ordination of the MTA managers which start to use it
- Mechanisms for loading the directory information tree with existing data
- Strategies for transition from current table-based routing and mapping methods to Directory-based methods
- Implementation and distribution of tools to support MTAs in both worlds, i.e., both Directory and table-based MTA's
- Time frame - date mentioned: spring 1993
- The co-ordination could also be done in the COSINE-MHS framework. This should then be included in the follow-up contract for the organisation which wants to provide these services in 1993.

The distribution list (all of the following addresses should work)

- mhs-ds@mercury.udev.cdc.com
- S=mhs-ds;OU=mercury;O=udev;P=CDC;A=ATTmail;C=US
- S=mhs-ds;OU=mercury;OU=oss;OU=arh;O=cpg;P=CDC;A=ATTmail;C=US

will be used for the co-ordination of the pilot in a first stage.

Next Meeting

The next meeting of the MHS-DS Working Group will take place at the IETF meeting in Boston, Massachusetts. Scheduled sessions are, Monday, July 13th at 4:00pm and Tuesday, July 14th at 9:30am.

Everybody is invited to attend this meeting. Further discussion will also take place on the distribution list (see above). Send request for registration to the list to:

- mhs-ds-request@mercury.udev.cdc.com
- S=mhs-ds-request;OU=mercury;O=udev;P=CDC;A=ATTmail;C=us
- S=mhs-ds-request;OU=mercury;OU=oss;OU=arh;O=cpg;P=CDC;A=ATTmail;C=US

Attendees

Harald Alvestrand	harald.alvestrand@delab.sintef.no
Denis Buffenoir	denis.buffenoir.inria.fr
Jim Craigie	craigie@jntg.ac.uk
Havard Eidnes	havard.eidnes@runit.sintef.no
Urs Eppenberger	eppenberger@switch.ch
Andrew Findlay	Andrew.Findlay@brunel.ac.uk
Ole Frendved Hansen	ole.frendved.hansen@uni-c.dk
David Goodman	d.goodman@cs.ucl.ac.uk
Alf Hansen	Alf.Hansen@delab.sintef.no
Erik Huizer	huizer@surfnet.nl
Wolfgang Jaretzki	jaretzki@dfn.dbp.de
Kevin Jordan	kej@udev.cdc.com
Peter Kaufmann	kaufmann@dfn.dbp.de
Bruno Koechlin	bruno.koechlin@inria.fr
Erik Lawaetz	uniel@uts.uni-c.dk
Jean-Paul Le Guigner	leguigner@cicb.fr
Thomas Lenggenhager	lenggenhager@switch.ch
Damanjit Mahl	damanjit.mahl@brunel.ac.uk
Michael Mueller	mueller@z1wnsv.gmd.de
Hector Nunez	h.nunez@cs.ucl.ac.uk
Paul-Andre Pays	paul-andre.pays@inria.fr
Catherine Pierre-Radenac	catherine.pierre-radenac@inria.fr
Colin Robbins	c.robbins@xtel.co.uk
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Marjo Rottschaefer	marjo.rottschaefer@surfnet.nl
Peter Sylvester	sylvester@gmd.de
Katy Treca	treca@urec.fr
Panos Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Ton Verschuren	ton.verschuren@surfnet.nl
Peter Yee	yee@ames.arc.nasa.gov

CURRENT MEETING REPORT

Reported by Urs Eppenberger/SWITCH and Kevin Jordan/CDC

Minutes of the MHS-DS Working Group (MHSDS)

Agenda

- Revision of the Charter
- Status of the Documents
- Review of the Documents
- AOB

The Minutes of the last MHS-DS meeting in Innsbruck, Austria were approved.

Charter

The revision and progression of Steve Hardcastle-Kille's documents have highest priority.

The Charter will be revised to include the following new items:

- Study the support of X.400 communities without directory access, and define requirements for tools which:
  - Extract info from the DIT into files.
  - Upload info from files into DIT.

- Identify additional information possibly needed to facilitate MHS management (beyond what is already defined in the Internet Drafts) and define where and how to place it in the DIT. (Note: it may be the case that no additional information is needed, at least in the short to medium term). Jim Romaguera volunteered to work on this item.

- The Group will coordinate a pilot and document the results. Experience from the piloting activity may necessitate changes in the MHS-DS specifications. Also, guidelines for deployment and tuning may result.

Status of the Documents

The following papers have been updated by Steve Hardcastle-Kille and were submitted to the Internet Drafts:

1. Representing Tables and Subtrees in the Directory.
2. Representing the O/R Address hierarchy in the Directory Information Tree.
3. MHS use of Directory to support MHS Routing.

4. Use of the Directory to support mapping between X.400 and RFC 822 Addresses.
5. MHS use of the Directory to support distribution lists.
6. A simple profile for MHS use of Directory.
7. Use of the Directory to support routing for RFC 822 and related protocols.

What needs still to be done:

- ASN.1 definitions in some of the documents.
- Routing algorithm in pseudo code. Perl code in beta status available from Harald Alvestrand. It's well suited for sanity checking of the algorithms since the code is executable
- A short overview paper is needed, Kevin Jordan volunteered to write it.
- Content conversion issues to be taken out of the routing document [3] and moved into a separate document.
- Notes on expected performance [3].
- Section on acknowledgements.
- References section.
- Section defining regular expressions (used in routing filters) [3].
- A table of MTA attributes and how they affect routing [3].

The next public release of PP will include the functionality defined in the minimum profile [6]. Quipu in ISODE 8.0 includes most of the objects defined in [1]-[7]. PP and Quipu will provide enough functionality to start a pilot. The restriction of supporting the open tree only may limit the scale of the pilot.

X-Tel got a contract from the European Commission for the project EXPLODE. Consequently, most functionality defined in [5] will be implemented in PP.

Review of the Documents

The documents [1]-[6] were thoroughly reviewed (page by page).

Editorial changes were pointed out for the documents [1], [2] and [4]. Steve Hardcastle-Kille will update the documents. The MHS-DS Working Group proposes to submit these documents as Experimental RFC's. Experience gained through the pilot may lead to document updates. After gaining experience via piloting, the documents will be revised, then the next goal will be to submit them as proposed standards.

[5] This document was discussed quite thoroughly at the meeting. As a result, additional needed functionality and DL policies were identified. A new version will be distributed and discussed at the next meeting.

[3] Some of the algorithms are changed.

The location of the open tree under O=Internet was discussed. The conclusion was to keep the open tree under the root of the DIT root, as originally proposed. It has been pointed out that [3] defines only a framework for using X.500 for X.400 routing. MHS communities will need to specify in more detail how the routing must be organised. It is a task of the pilot to study these issues.

Steve Hardcastle-Kille will update [3] and resubmit it as an Internet Draft.

[6] Did not get significant changes. Since it is tightly coupled with [3], it will also be updated and resubmitted as an Internet Draft.

[7] Will be discussed at the next meeting.

AOB

Next meeting: 25th IETF, November 16-20, 1992, Washington, D.C.

Action List

Kevin Jordan	To revise the Charter.
	To write an overview paper.
Urs Eppenberger	To write the Minutes.
Harald Alvestrand	To provide pseudo code.
Jim Romaguera	To write a document on possible additional objects and attributes needed in the framework of MHS management.
Steve Hardcastle-Kille	To make editorial changes to the documents [1], [2] and [4] and submit them as Experimental RFC's.
	To revise documents [3] and [5] and distribute them to the list in time for the next meeting.
	To create an additional document on content conversion and access units.

Everybody to study the revised documents and contribute comments.

Attendees

Ed Albrigo	ealbrigo@cos.com
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
C. Allan Cargille	cargille@cs.wisc.edu
James Conklin	jbc@bitnic.educom.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppenberger@switch.ch
Ray Freiwirth	5242391@mcimail.com
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@nersc.gov
Arlene Getchell	getchell@nersc.gov
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Erik Huizer	huizer@surfnet.nl
Takashi Ikemoto	tikemoto@xerox.com
Kevin Jordan	kej@udev.cdc.com
Todd Kaehler	kaehler@zk3.dec.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Thomas Lenggenhager	lenggenhager@switch.ch
Eric Nowak	nowak@ans.net
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Einar Stefferud	stef=isoc@nma.com
Panos-Gavriil Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Brien Wheeler	blw@mitre.org
Peter Williams	p.williams@uk.ac.ucl.cs
Russ Wright	wright@lbl.gov
Yung-Chao Yu	yy@qsun.att.com

2.4.2 MIME-MHS Interworking (mimemhs)

Charter

Chair(s):

Steve Thompson, sjt@gateway.ssw.com

Mailing Lists:

General Discussion: mime-mhs@surfnet.nl
To Subscribe: mime-mhs-request@surfnet.nl
Archive:

Description of Working Group:

MIME, (Multipurpose Internet Mail Extensions) currently an Internet Draft, is expected to become an Internet Proposed Standard. MIME redefines the format of message bodies to allow multi-part textual and non-textual message bodies to be represented and exchanged without loss of information. With the introduction of MIME as a Proposed Standard it is now possible to define mappings between RFC-822 content-types and X.400 body parts. The MIME-MHS Interworking Working Group is chartered to develop these mappings, providing an emphasis on both interworking between Internet and MHS mail environments and also on tunneling through these environments. These mappings will be made in the context of an RFC-1148bis environment.

Goals and Milestones:

Done	Post an Internet Draft describing MIME-MHS Interworking.
Jul 1992	Submit a completed document to the IESG describing MIME-MHS Interworking as a Proposed Standard.
Done	Post an Internet Draft describing the "core" set of Registered conversions for bodyparts.
Jul 1992	Submit the "core" bodyparts document to the IESG as a Proposed Standard.

Internet Drafts:

"Mapping between X.400 and RFC-822 Message Bodies", 07/01/1992, H. Alvestrand, S. Hardcastle-Kille, R. Miles, M. Rose, S. Thompson <draft-ietf-mimemhs-mapping-01.txt>

"Equivalences between 1988 X.400 and RFC-822 Message Bodies", 07/01/1992, H. Alvestrand, S. Thomspon <draft-ietf-mimemhs-body-equival-01.txt>

CURRENT MEETING REPORT

Reported by Steve Thompson/ODS

Minutes of the MIME-MHS Interworking Working Group (MIMEMHS)

There have been two papers produced since the last meeting:

1. X.400/MIME body equivalence Harald Tveit Alvestrand, Steven Thompson
2. Mapping between X.400 and RFC-822 Message Bodies Harald Alvestrand et al.

Several mappings have been defined, and for those without a clear X.400(88) equivalent there is a trapdoor/catchall External bodypart defined in X.400: EBP-mime-body-part. In the other direction the trapdoor in MIME is a new Mime subtype: application/x400-bp

The session consisted mostly of a presentation by S. Thompson and H. Alvestrand on the material, since the papers were released shortly before the IETF. The discussion raised the following issues for resolution:

- Registration
  - How to get vendors to register OIDs as well as the equivalent MIME subtype with the IANA. Thompson is working with the EMA (Electronic Mail Association) and Jon Postel on this issue.
  - How to manage IANA registration of different versions of BPs like WP5.0 and WP5.1.
  - Should there be automatic OID assignments for MIME Subtypes?
- '84 X.400 Support. This subject got the most discussion, by far.
  - Should the RFC address '84 explicitly or should 88-84 downgrading be enhanced to handle it.
  - Simplest case single BP IA5 - Should this be a MIME document or just an RFC822 message.
  - T.61 strings in header vs RFC 1327 needs resolving
  - Three-party mail issue (mime-X.400(88)-X.400(84)).
- The documents need appendices with OIDs defined
- Security: viruses will be gatewayed too, not solved in this paper.

- Criticality of header extensions must be defined.
- Conformance: Is there a different level of support for gateway's vs. User Agents for the newly defined X.400 and MIME constructs?

Issues will be resolved by E-mail in the next couple of months. Both documents will be forwarded as proposed standard RFCs.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Mark Bokhan	bokhan@abitok.enet.dec.com
Luc Boulianne	lucb@cs.mcgill.ca
Urs Eppenberger	eppenberger@switch.ch
Gary Gaudet	gaudet@zk3.dec.com
Tony Genovese	genovese@nersc.gov
Steve Hardcastle-Kille	s.kille@isode.com
Paul Hill	pbh@mit.edu
Christian Huitema	christian.huitema@sophia.inria.fr
Erik Huizer	huizer@surfnet.nl
Todd Kaehler	kaehler@zk3.dec.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Kent Malave	kent@chang.austin.ibm.com
Keith Moore	moore@cs.utk.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Marshall Rose	mrose@dbc.mtview.ca.us
Einar Stefferud	stef@nma.com
Steve Thompson	steve@ods.com
Linda Winkler	lwinkler@anl.gov

2.4.3 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: merit.edu:pub/noop-archive

Description of Working Group:

The Working Group is chartered to work on issues related to the deployment of CLNP in the Internet. The first area of this Group's work has been the learning necessary to start deploying OSI in internet networks. This phase includes planning for OSI deployment by creating routing plans for regional networks and education on using OSI routing protocols.

This first area of the Group's work will be on-going as we continue to deploy OSI in the Internet. This step has lead to people deploying OSI for Pilot projects and demonstrations of OSI.

The second step of deploying OSI will be the transition of OSI from a pilot service to a production service. During this phase we will work on specifying the network debugging tools and test beds. We will need to track the level of OSI support in the Internet. We will need to provide documentation for new users of OSI on the Internet.

Goals and Milestones:

Jan 1992	Post as an Internet Draft, a tutorial for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.
Jul 1992	Post as an Internet Draft, a collection of regional Routing and Addressing plans.
Ongoing	Provide a forum to discuss OSI routing plans by email or in group discussions.
Jul 1992	Post as an Internet Draft, 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.
Jul 1992	Post as an Internet Draft, a description of OSI network layer debugging methods.
Jul 1992	Post as an Internet Draft, a list of OSI Network Layer NOC tools available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.

Apr 1992 Post as an Internet Draft, a requirements document specifying what OSI network tools are needed on every host and router.

Jul 1992 Submit to the IESG for Proposed Standard, a requirements document specifying what network tools are needed on every OSI host and router.

Aug 1992 Submit to the IESG as an Informational RFC, a description of OSI network layer debugging methods.

CURRENT MEETING REPORT

Reported by Sue Hares/Merit

Minutes of the Network OSI Operations Working Group (NOOP)

Agenda

- RFC 1139 and ISO Echo
- Tools RFC
- Transport MIB
- IP v7
- TUBA - Ross Callon (DEC)
- INTEROP Issues - Richard Collela (NIST)
- Skinny Stack and X windows over OSI - Jim Quigley HP

Notes: At the request of the Chair, these Minutes should be considered "unapproved".

RFC 1139

No one thinks the short-term echo-reply option in RFC 1139 is good. The long-term option is the only desired. This long-term ping function was taken into ISO in June of 1991, and is going on PDAM ballot. The IESG needs to do something with RFC 1139. Since the ISO work is coming a long, and a little different than the RFC, the RFC will be re-written to match the ISO text. A copy of this will be forwarded to the Working Group.

Dave Piscitello, Sue Hares and Cathy Wittbrodt will re-write RFC 1139 to match the ISO text.

Tools RFC

The tools RFC will go through a final revision, and be published in August as an Internet Draft. Comments are really needed about the variables selected in the end of the text.

A long discussion ensued here about whether a few variables are enough or the whole MIB. Mostly, the people thought a few variables would be enough. Also, a discussion occurred about SNMP over connectionless transport. Why don't you just use SNMP over IP to manage your node. The problem some people face is that IP paths and CLNP paths do not go the same way.

Transport MIB

A discussion was held to see if this work was useful in the NOOP Working Group. The Group generally agreed that they would be interested in reviewing the MIB. The following general things were suggested as things to monitor: checksum errors, protocol errors, counts

on octets sent.

Dave Piscitello will do TCP MIB/Transport MIB comparison. Transport MIB will be sent to mailing group.

TUBA Presentation

Ross Callon (DEC) gave a presentation of his TUBA/TUCAN document. Please refer to that document for the full details.

INTEROP Spring '92

Richard Collela (NIST) and Doug Montgomery (NIST) spent a day at hot-stage on IS-IS, and found a couple of bugs. Due to other problems, the operational show network used static routing and cisco IGRP-OSI routing. The OSI connectivity to the rest of the world was stable.

Also, one evening, IS-IS was brought up on all the show floor as a trial and went reasonably smoothly.

The demos were: NIST routing Lab and Infrastructure. The NIST routing Lab had 7 vendor implementations if IS-IS with a network monitor showing the topology. The equipment was configured in 4 L1 areas and connected to the show network.

The Infrastructure demonstration had 8 sites in US, and 8 sites in Europe. The network monitor station showing End system connectivity and status. The tools used on the Infrastructure network monitor were: clnp ping, clnp trace, and FTAM to remote sites.

The October '92 INTEROP has no spot for an InfraStructure demonstration. IS-IS may be operational on the show network The Spring '93 INTEROP will have an OSI Technology Showcase.

Skinny Stack and OSI Presentation

Jim Quigley (HP) gave an excellent presentation of what the OSI applications need as upper layers. In fact most applications only need a basic communications. Three types of this basic communication has been defined by the OSI Implementor's workshop (NIST OIW) by the upper layer Special Interest Group (ULSIG): Migrant (non-OSI such as TCP), X windows, and kernel (OSI applications). Prototypes of X windows over skinny stack over TP-4 have been produced and demonstrated by ULCC and Dr. Peter Furniss. This code will be released into the public domain.

Watch for X-windows over OSI code announcement.

Attendees

Ed Albrigo	ealbrigo@cos.com
William Barns	barns@gateway.mitre.org
Ross Callon	callon@bigfut.enet.dec.com
Alan Clegg	abc@concert.net
Marc Cochran	mjc@vnet.ibm.com
Richard Colella	colella@osi.ncsl.nist.gov
John Curran	jcurran@bbn.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Tom Easterday	tom@cic.net
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Osten Franberg	euaokf@eua.ericsson.se
Joan Goldstein	j_goldstein@tnpubs.enet.dec.com
Susan Hares	skh@merit.edu
Paulina Knibbe	knibbe@cisco.com
Mark Knopper	mak@merit.edu
Walter Lazear	lazear@gateway.mitre.org
Bill Manning	bmanning@rice.edu
J. Scott Marcus	smarcus@bbn.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Brad Passwaters	bjp@sura.net
David Piscitello	dave@sabre.bellcore.com
Lars Poulsen	lars@cmc.com
James Quigley	Jim_Quigley%YD@hp6600.desk.hp.com
David Reiser	dsr@codex.prds.cdx.mot.com
April Richstein	abm@tycho.ncsc.mil
Erik Sherk	sherk@sura.net
William Warner	warner@ohio.gov
Luanne Waul	luanne@wwtc.timeplex.com
Linda Winkler	lwinkler@anl.gov
Steven Winnett	swinnett@bbn.com
Cathy Wittbrodt	cjw@nersc.gov

2.4.4 OSI Directory Services (osids)

Charter

Chair(s):

Steve Hardcastle-Kille, s.kille@isode.com

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
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, .ps>

"Handling QOS (Quality of service) in the Directory", 03/20/1991, S.E. Kille <draft-ietf-osids-qos-01.txt, or .ps>

"DSA Naming", 03/21/1991, S.E. Hardcastle-Kille <draft-ietf-osids-dsanaming-02.txt, or .ps>

"Naming Guidelines for Directory Pilots", 03/21/1991, P. Barker, S.E. Hardcastle-Kille <draft-ietf-osids-dirpilots-05.txt, .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>

"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-02.txt, .ps>

"The Simple OSI Stack", 03/09/1992, S. E. Hardcastle-Kille <draft-ietf-osids-simple-stack-00.txt, .ps>

"Counting the Directory Information Tree", 04/08/1992, Steve Hardcastle-Kille <draft-ietf-osids-dirtree-00.txt, .ps>

"Lightweight Directory Access Protocol", 04/17/1992, Wengyik Yeong, Tim Howes, Steve Hardcastle-Kille <draft-ietf-osids-lightdirect-01.txt>

"The String Representation of Standard Attribute Syntaxes", 05/05/1992, T. Howes, S. Hardcastle-Kille, W. Yeong, C. Robbins <draft-ietf-osids-syntaxes-01.txt>

Request For Comments:

RFC 1275 "Replication Requirements to provide an Internet Directory using X.500"

RFC 1276 "Replication and Distributed Operations extensions to provide an Internet Directory using X.500"

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"

CURRENT MEETING REPORT

Reported by Steve Hardcastle-Kille/UCL, Doug Simmons/IBM and Justin Walker/Apple

Minutes of the OSI Directory Services Working Group (OSIDS)

Comments on Agenda

Mark Knopper sent apologies for non-attendance and then turned up.

Here follow comments on the Minutes from San Diego (OSI-DS-MINUTES 7), particularly relating to action items from that meeting:

- Regarding maintenance of RFC-1274, it is Steve and Paul Barker who will be involved, not Colin as the Minutes claimed.
- Eric Huizer's strategy document is ready (comments will be made later in the Minutes).
- Chris's documents (OSI-DS 14, 16, 17, 19) have not been revised. This will be done by the next meeting. Mark Knopper has taken up the network schema.
- Documents OSI-DS-12, 23, 24 have been revised and submitted to the IESG. As suggested in the previous meeting, Steve Hardcastle-Kille read NADF 175 (which has been revised to NADF-(***Didn't get the reference***). He cleared up some misconceptions regarding he NADF position, but overall, his position did not change from the last meeting.
- Wengyik continued to work on interoperability issues. However, he had no input, since he was not present.
- There will be a NADF meeting next week (that is, the week of 7/20).
- The QOS experiments will be discussed as indicated in the Agenda.
- The JPEG schema was reviewed. However, the Schema group had not yet formed, so this item was continued. Paul Barker was to establish the schema group, but this had not done this because of overload due to Steve's departure to the ISODE Consortium. Resources were requested (volunteers were solicited; no response was heard from Russ and Mark).
- The Character Set experiment wasn't discussed because Geir Pederson wasn't present. The action was continued.

- Tim Howes brought us up to date on the DIT Counting effort. The code has been written, and will appear in the next ISODE Consortium release of QUIPU. The item was continued.

- Work on schema publishing was completed and will be discussed later.

- The action item on preferred names was continued (no one was present to speak on the subject).

- Steve H-K finished the revision of RFC-1279 , with Wengyik consulting. The paper has not actually been updated (in its electronic form). It will be circulated.

- The lightweight protocol note (LDAP) was revised and circulated.

- Steve Hardcastle-Kille looked into possible ISO alternatives to SOS (the Simple OSI Stack). There are no current ISO proposals addressing the SOS issues, but John Day (from BBN) has circulated a document (in OSI circles) on the OSI upper layers. Reviews are not complete, but this document does not seem to be an answer.

There were no Matters Arising. We therefore moved right into the liaison reports.

1. RARE (Eric Huizer)

   As reported at the last meeting, WG3 is no more. The new RARE structure has eight Working Groups, of which one, WGNAP (the Working Group for Network Application support), will undertake directory services (as well as time protocols, etc.). A major problem is that, while a Chair has been identified and wants to undertake the work, he can't get permission to do the job. WGNAP hopes to meet in November. A distribution list has been set up for other than directory service issues. WGNAP will continue to use OSI-DS for its directory service discussions. The Working Group has small budget from RARE, provided they can come up with a priority list of tasks. This could be applied to travel.

2. OSI/CCITT (Ken Rossen)

   There were two significant events to report. ISO 9594 passed to DIS. The most significant change was in the area of access control (replication and an extended information model). DISP (shadow); DOP (binding) are new protocols. An access control context is a combination of levels of access control. The US pushed successfully for simplified access control: this only allows a decision to be made at administration points (new in model); a decision isn't overridden by lower levels of the tree. As of the last editing meeting, merged text was produced. Unfortunately, the circulated stuff was a mess. There is a good copy, dated 12/25/91 (hence it is called "the Christmas text").

   The second event occurred in May. When ISO SC21 met in Ottawa, the Directory

Services Group also met, and changes to the Standard were discussed (with a 2 year target, down from the usual 4). Use of OSI management (CMIP) to manage the directory was put on hold, since the responsible party (from the US) resigned. Work on authentication could be undertaken as there is support for small changes, e.g., certificate revocation. This will wait for the next meeting to commit effort to this work. There was a feeling that there is need for closer work with (ISO) security folks for a more sophisticated security model. Given upper layer security services, there is a need for a scheme to apply to directory services. Also, there is a new edition of ASN-1 encoding rules, which could effect directory. Distinguished encoding rules were introduced that are different from those currently used by the directory. There is need to work out conflicts. This could affect digital signatures. The 1992 X.500/9594 should progress at the next editing meeting in Orlando, in the fall of '92 (this will involve serious cleanup. Rows of ducks will be set up at a US meeting in Nashua this week.).

3. OIW (Russ Wright)

   The OIW continued work on standardized profiles for DAP, replacing agreements from the OIW and EWOS. They are on schedule for results by the end of year. A joint meeting was held with the X400 SIG to look at MHS and the directory. Their desires right now are unclear, but they will provide a clearer specification.

   The IGOSS document was reviewed. This is a combined document representing input from GOSIP, the power industry, and the manufacturing industry. This requires '92 directory extensions, including replication. They were asked to review POSIX documents relating to directory services. The documents themselves are in the mail.

4. DISI (Chris Weider)

   Documents describing advanced directory usage and how to get registered in the directory have been worked on, but not circulated. A revision RFC-1292 has been worked on Four new papers have been prepared: a pilot catalog, a description of DIT setup, the directory naming philosophy, and a schema for restaurant information.

5. AARN Mark Prior

   There is not much happening at this time. AARN is not willing to commit to further work, nor are they willing to say no to further work. They are waiting for December (***Why?***). There are currently 40000 entries in their directory, and they have just added affiliates. Master and slave machines will be soon be upgraded.

6. NADF (Marshall Rose/Einar Stefferud)

   The last NADF meeting was in April, the next will be next week (7/20). Discussion of vendor plans at the last meeting was exciting (depressing?). Several documents are available. One provides a naming scheme for a country (discussing principles), and a second provides an application of these principles to the US. A third discusses

the theory and practicality of directory security.

This latter is up for more debate. There is a desire for simple authentication, but this may be difficult to protect from replay attack. The recommendation may be for protected passwords. The documents should become RFCs (but some can't even seem to be put into the politically *in*correct PostScript format). Marshall will provide copies for Steve Hardcastle-Kille. None of the twelve vendors present supported any but simple authentication. None would commit to supporting '92 extensions (except one who was planning to support the extended information model). In short, things don't seem to be going very well (according to public comment at the meeting. This is born out by Ken's observations at COS). There seems to be more positive support for simplified access control (over the basic version).

Ken noted that they think they've fixed NADF complaints. Time was spent at the Ottawa meeting on defect resolution (there is a directory implementor's guide; see Ken). There seems to be some interplay between ISO, NADF.

As no pilot project representatives were present, we continued on with the rest of the Agenda.

The Naming Guidelines Document, the UFN Document, and the Document Defining String Representation of UFN's. All three were submitted in April to the IESG. They are expected to move forward by end of this month.

The Strategy Document. The Strategy document (based on Steve's original) was much modified, based on comments received. Most of the original was retained, but with editing and restructuring. One of the main criticisms was references to other RFCs without indicating the RFC content. Eric's solution was to pull the main points from the RFCs in question, using reference only for detail. He added deployment details and requirements. Therefore, there were a lot of references to DISI papers. The ASCII version (as posted) was quite unreadable. Apologies were tendered, along with a promise to fix it. Comments were requested.

One comment at the meeting: a possible extension involving the use of large data values was questioned. The response was that this is only a *possible* extension, not a planned (or required) one. An observation was made that all items in this section (of the document) could be termed controversial. The main point is that the model is not rigid: if deployment experience indicates that a change is needed, it will be addressed.

Regarding progress to ID-hood for the strategy document, the approval of the other authors is needed. Then an informational RFC can be submitted. Steve Hardcastle-Kille wants to see this done reflecting an IAB/IESG consensus (as was done, e.g., for RFC-920). He wants the submission and publication to reflect IAB policy. It is unclear what the tradition is. It was felt that we should have OSI-DS consensus, so a sense of meeting was taken; there were no votes against the document, but there were a large number of abstentions (from those who had not read it yet). Eric will take changes, publish the new document as an RFC

(both text and PS formats), and get it into the IESG stream. The attendees seemed to favor not waiting for the next meeting, given the consensus here (all who had read approved).

Eric noted that none of the three documents mentioned earlier showed up on the IESG action list that he gets. This was deemed to be a dropped ball. Eric will follow up to determine how the ball got dropped and assure that it doesn't happen again.

Tim Howes: Some comments on the schema document, from Colin Robins (sent by email to the OSI-DS list), were distributed. Given that the schema is rapidly changing, the idea of storing (a description of) the schema in the DIT has been investigated. Tim looked first at the '92 Standard, which was very complicated. The '92 information is in his document, but comments he's received indicate that it (the '92 content) should be pared down. The document talks about representing attribute. information in the directory, although no syntaxes were defined. Although the document says this work will be a subset of '92, Tim doesn't think it really is. We must decide on compatibility with '92 vs. having something "now".

The question was asked: what are the areas of incompatibility? Among others, there is the attribute syntax, which is difficult to figure out. From Colin: how does one go from an OID to an identification of information it represents?

It was noted that an OID tree may be useful by itself, independent of other uses. There is a bootstrap problem with this. The issue is where to find a description of information, and what is the efficiency hit? Using well-known locations in the DIT may avoid a recursive upward walk of the tree. This also assumes a configuration run that tells the DSA what well-known locations to check. The directory doesn't do dynamic interpretations of OIDs. It was observed that "compatibility w. 92" and "something that works" may not be exclusive. Two actions resulted. The first was to define the OID tree. The second was to revise the schema notes in light of the discussion. Tim took both.

QOS Experiments. There was no change from the previous meeting. This work has not been a priority (although there is work "scheduled", to be done on Macintosh DUA). Sylvain noted that code that he has seen doesn't match the RFC (which may have changed since he last checked it). Tim wanted this taken off the Agenda, since it isn't a priority. He would like to surprise us with progress when it happens.

JPEG. The JPEG attribute is not in the schema, but there is code to handle it in ISODE. Russ would like this to be its end. Proposed to carry over to next time when the schema group is represented (and so it shall be).

Character Set. (Geir Pederson was not present): Again, the schema group was an issue. A discussion commenced on how to get this done. IANA was suggested as a source of help. A problem with this is that we would need to find someone with directory experience to take on some editing load. It was recommended that we talk with IANA, then worry about the short-term.

Selection of the time and place for the next meeting involved two choices: INTEROP (October in San Francisco), and the next IETF (November in Washington). A vote marginally favored the November IETF meeting, and this was agreed on.

DSA and DUA Metrics (OSI-DS33 ,OSI-DS 34)

- Measure pilot projects' success.
- Deliverables - metrics papers for:
  - DUAs.
  - DSAs.
  - Pilots' metrics.
- No absolute measure of goodness or badness of DUAs; there's SOME importance to the numbers, though.

  Comments on these papers:

- Set up an FTP ID to keep the OSI-DS documents in for easy retrieval before these meetings. SEH to address.
- DSA Document - need hands-on experience to tell if this document is really worthwhile and accurate. (comment by Eric Huizer).
- DUA document - section 12 (query resolution) not very clear what one should enter to initiate the query (comment by Time Howes).
- DUA document - 5 steps to enter a query as opposed to on line via UFN.
- BOTH - is this a Consumer Reports on DUAs/DSAs? SEH - the user endorsement section contains the necessary feedback for analysis.
- BOTH - there were comments from Paul Andre, were they being incorporate?
- DSA Document - section 5,need to discuss the environment - how can we measure implementations on different machines? (comment by Tim Howes).
- DSA Document - need more than 100 to 5000 entries for accurate testing (comment by Tim Howes).
- DSA Document - need more discussion on security aspects (unknown).
- BOTH - metrics will not be useful until they are tried out/tested against (unknown).

- BOTH - make measurements available via informational RFC.

- DSA Document - other implementations tested besides QUIPU? (comment by Sylvain Langlois) (Pissaro(sp?), ICL, Dirwiz....)

- S. Hardcastle-kille: How many of us are responsible for DUA implementations? Would it be worthwhile to make these documents publicly available? SEH to use RFC for informational test until next meeting for feedbacks.

E. Huizer	To do Siemens DSA.
T. Howes/R. Wright	Will do DUAs we'll evaluate findings at next meeting.
S. Hardcastle-Kille	To get these published as RFCs.
Everyone	To see that these get filled in when DUA's and DSA's tested.

- Comment - what's the difference between RFC 1292 and DS 33 and 34? SEH: 33 and 34 much lower level (and more work to fill out) S. Hardcastle-Kille suggested that the vendor be asked if they filled out a 33 or 34 before answering to RFC 1292.

Representing Network Infrastructure Information in X.500 (Mark Knopper) Draft circulated.

Soft Pages Project (Steve Hardcastle-Kille).

- Comment - IP name space: defining an address hierarchy. You really don't need that, what advantage over a flat design?

- Comment - Network elements diagram is a network topology. What happens if that changes? (comment by Tim Howes).

- Comment - (Mark Hopper).
  - Not sure if this resolves the problem.
  - It is too inefficient.
  - How do you get the bootstrap up and running?

- ACTION * - Mark Knopper to document how we might use this (where might the holes be).

- Comment - this tree can be kept small just by keeping the DSAs "near" you in the DIT, as they are the only ones which should interest you for cost purposes.

- Comment - need FTP address for this document (FTP.TOHOKU.AC.JP).

- Comment - do we need a Working Group to address this problem?
- ACTION* Thomas Johansen and Mark Knopper to reconsider their approaches and attempt some kind of convergence.

LDAP (OSI-DS 26, OSI-DS 27)

- Comment - kerberos and simple authentication: do we think this is worthwhile and should it be added to the document before it becomes an RFC? (Tim Howes).
- S. Hardcastle-Kille: Because it is implemented and deployed, then it should be documented.
- Comment - we should submit this to the Standards committee as soon as possible.
- Comment - suggestion that we have Christian look at it, as he has strong views on the subject.

DSA Naming (OSI-DS 13)

- Issue: Avoiding Deadlock.
- Comment - the DSA must be named higher in the tree (country level) to prevent deadlock, but you do not insure uniqueness.
- Comment - Erik seemed to remember opposition by the Pissaro group, but could not elaborate.
- Comment - using subtrees seems to be the way we fix things we can't fix via X.500.
- ACTION* S. Hardcastle-Kille: To re-write the paper to using non-QUIPU language and references.
- Comment - Erik not comfortable, seems like a way to fix a design problem in QUIPU. Need input from other DSA vendors.
- ACTION* S. Hardcastle-Kille: To drop this as an OSIDS item and take it up as a design issue with ISODE.

Action Items

Chris Weider	Update OSI-DS 14, 16, 17, 19 (carried forward)

E. Huizer	Progress Naming Guidelines, DN Syntax, UFN, and LDAP and LDAP Syntaxes as RFCs.
	Do Siemens DSA.
T. Howes/R. Wright	Will do DUAs we'll evaluate findings at next meeting.
S. Hardcastle-Kille	Get these published as RFCs.
Everyone	to see that these get filled in when DUA's and DSA's tested.
M. Knopper	To document how we might use this (where might the holes be).
T. Johansen/M. Knopper	Reconsider their approaches and attempt some kind of convergence.
S. Hardcastle-Kille	to re-write the paper to using non-QUIPU language and references.
	Drop this as an OSI-DS item and take it up as a design issue with ISODE.
	Revise Charter.
S. Hardcastle-Kille/E. Huizer	Discuss IANA support for Schema Management.
T. Howes	Write note on representation of OID Trees in Directory.
P. Barker	Publish Metric Papers as Internet Drafts.
S. Sataluri	Collect DUA survey results and publish as Internet Draft.

Attendees

Ed Albrigo	ealbrigo@cos.com
Claudio Allocchio	c.allocchio@elettra.trieste.it
C. Allan Cargille	cargille@cs.wisc.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppenberger@switch.ch
Ray Freiwirth	5242391@mcimail.com

Jisoo Geiter	geiter@gateway.mitre.org
Arlene Getchell	getchell@nersc.gov
Joan Goldstein	j_goldstein@tnpubs.enet.dec.com
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Takashi Ikemoto	tikemoto@xerox.com
Kevin Jordan	kej@udev.cdc.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Thomas Lenggenhager	lenggenhager@switch.ch
John McKenna	mckenna@ralvm12.vnet.ibm.com
John Murray	murray@premenos.sf.ca.us
Mark Needleman	mhn@stubbs.ucop.edu
William Nichols	nichols@wick.enet.dec.com
Eric Nowak	nowak@ans.net
Rakesh Patel	rapatel@hardees.rutgers.edu
Mark Prior	mrp@itd.adelaide.edu.au
Sheri Repucci	smr@merit.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Marshall Rose	mrose@dbc.mtview.ca.us
Rich Rosenbaum	rosenbaum@lkg.dec.com
Kenneth Rossen	kenr@isc.com
Srinivas Sataluri	sri@qsun.att.com
Douglas Simmons	
Mark Smith	mcs@umich.edu
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Justin Walker	justin@apple.com
William Warner	warner@ohio.gov
Chris Weider	clw@merit.edu
Brien Wheeler	blw@mitre.org
Scott Williamson	scottw@nic.ddn.mil
Steven Winnett	swinnett@bbn.com
Russ Wright	wright@lbl.gov
Yung-Chao Yu	yy@qsun.att.com

2.4.5 Office Document Architecture (oda)

Charter

Chair(s):

Peter Kirstein, P.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.
Jun 1992	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)

The Chair repeated that for the ODA Working Group, some papers are available in an electronic form on the UCL-CS info-server: "info-server@cs.ucl.ac.uk". Concern was expressed that a number of people, who had expressed previous interest, might not be on the mailing list; the Chair promised to ensure that those who had attended the IETF ODA Working Group would be on the mailing list.

The documents in the info-server are accessed by standard message systems, giving a message body of the form:

request:ietf-osi-oda

topic:xxxx

where xxxx is the name of the document required.

All the documents in the info-server are available in text form; many of them are stored also in ODA/ODIF format. The list of documents currently in the collection is listed in a document called INDEX.

The latest document defining the current status is [1] below. It is available on the info-server. It gives details about the current implementations producing Q112 ODA available for the Pilot. These are listed below:

PRODUCT	Supplier	Status	Product	Availability	Source
SLATE/ODA V1.2	BBN/UCL	OK	SLATE-yes	Now	UCL
Word-for-Windows/ODA v3	Bull	OK	Yes	Now	Bull
DECWRITE/ODA	DEC	OK	Yes	Now	DEC
Global View	Xerox	Testing	Yes	Now	Xerox
MACODA	Apple	Testing	Beta	??	Apple
WordPerfect	UPC	Dev	??	Q4 92	UPC

UCL started shipping SLATE/ODA v1.2 in February 1992; this is based on SLATE v1.2. Few have tested it. BBN is planning to bring out V2 of SLATE during Q4 92. From a user viewpoint, this has better graphics and font support, supports a fuller character set, and should be easier to configure; nevertheless, for reasons given below, we expect SLATE/ODA v2 to be ready for shipping only during Q1 1993 at the earliest - and not for this stage of the Pilot. The Bull product has been released; UCL has tested it for conformity. The DEC product is available now, and has been tested for conformity; no other site had tried a recent version of that software. The Apple and Xerox products have been delivered to UCL. They are still being tested for interoperability by UCL.

The availability situation of the WORDPERFECT/ODA system from University Politechnic of Catalonia has kept changing. There has been a contractual problem in getting that version used in the Pilot, because it included some software from ICL - which UPC had not had permission to release for the Pilot. UPC has now stated that they intend to have software entirely of their own available for the September 1992 - this will be made available to UCL for in September; UCL hopes that it will be usable on the Pilot during Q4 1992.

During an earlier discussion, it became clear that there was a requirement to use a large number of National characters; for instance support for the specific Nordic ones were requested by one participant, and of Greek ones by another. It was considered desirable to use character set switching according to ISO 8859/4 in the long term. The SLATE V2 system now uses ISO 8859/1; it was not clear who else supported it. The Bull implementation is expected to support a richer character set on the next version; we are discussing the requirements with them. No information was available on the Apple, DEC or Xerox ones.

A number of small pilots were mentioned - a Nordic one, one based on Mitre, one organised by UCL in Europe, and a small one at NASA-Ames. There was interest in broader pilots. The Chair promised to set up a list of ODA and normal mailboxes. This list will be available from ietf-osi-oda, and the interested parties will be requested to provide their details to the Chair.

So far, most use of the systems have been using X.400 for transport; interest has been expressed in working with the Multi-media Mail format (MIME). This is in accordance with the Charter, since this is a Proposed Internet Standard. The Chair mentioned that ODA was a recognised Content Type in MIME, and stated that at least the BBN/UCL SLATE now supports MIME. He promised to investigate when there could be more general integration of MIME with some of the ODA products, and report this to the Distribution List.

It had been expected that there would be substantial experience with the Pilot prior to the July IETF meeting in Boston. In practice there has been little such experience by IETF members.

Future Extensions

The current Profile used in all the ODA document activities described here is the Q112 Profile. A new Profile called FOD26 is in the process of ratification. The new profile has some advantages, but that is less important than that a number of large manufacturers have agreed to support it for products (e.g., Bull, DEC, IBM, ICL, Siemens-Nixdorf and UNISYS as part of the ODA Consortium). In the last few months, it has become clear that the manufacturers are going to provide most of their further products only according to the FOD26 version. Pending the ratification of the Standard, and the availability of Toolkits like those being produced by ODAC above, no additional new ODA products that can be expected; all enhancements will come in the FOD26 form.

In view of the above, it would seem advisable to continue pilots based on the products mentioned above, but not to envisage new products prior to Spring 1993. As far as the SLATE/ODA activity, it does not seem sensible to provide anything other than SLATE/ODA v2 - based on SLATE v2 and FOD 26.

Security Extensions

The Chair mentioned that UCL had prepared a subsystem that implemented the ODA extensions specified in the Addendum to the Standard - subject to the vital restriction that the extensions apply to whole documents, not also to the separate ODA portions in the document. The extension had been integrated into one version of their SLATE 1.2 systems, and could be adapted to the Bull and UPC ones (and possibly others); it would be used in the PASSWORD Pilot in Europe, which was piloting a number of security services (including confidentiality, integrity and authenticity for document interchange). While it was considered an interesting development, it was agreed that this system would not be used in the IETF ODA Pilot during the remainder of 1992 for two reasons:

1. It was considered undesirable to duplicate the functionality of Privacy Enhanced Mail (PEM), before PEM had been properly piloted;

2. It was considered unnecessary to introduce yet further functionality until the present Pilots were more advanced.

Future Meetings

While the attendance at this IETF ODA Working Group meeting was better than that at the previous one, we propose not to have meeting in November - but to postpone a further full meeting until the March 1993 IETF. It is hoped that by this time there would be extensive Pilots, and that the timing of FOD26 products would be clearer.

Reference

1. D. Sadok et al: The ODA Document Convertors, UCL Internal Report No 2, Version 4, July 1992.

People so far Receiving SLATE/ODA from UCL

FIRST NAME	SURNAME	ORGANISATION	CITY	COUNTRY
Velu	Sinhu	OKIDATA	Framingham, Mass	USA
Jim	Knowles	NASA-AMES	Mountain View, CA	USA
Mark	Stansfield	Paisley Coll	Paisley	Scotland
Mark	Savela	Telecom Res	Otakaari	Finland
Erik	Lillevold	NTRA	Kjeller	Norway

Encarna	Pastor	U. Pol. Madrid	Madrid	Spain
Knut	Holm	SINTEF	Trondheim	Norway
Trevor	Hales	CSIRO	Carlton, Vic	Australia
Naresh	Kumar	Touch Com	Campbell, CA	USA
Andrew	Macpherson	BNR-Europe	Harlow, Essex	UK
Paul	Kennedy	DEC	Galway	Ireland

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jisoo Geiter	geiter@gateway.mitre.org
David Katinsky	dmk@rutgers.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Thomas Lenggenhager	lenggenhager@switch.ch
Mark Needleman	mhn@stubbs.ucop.edu
Sam Nicholson	scion@pblx.knox.tn.us
Rakesh Patel	rapatel@hardees.rutgers.edu
Robert Reschly	reschly@brl.mil
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Vincent Sgro	sgro@cs.rutgers.edu
Sandro Wallach	sandro@elf.com
Peter Williams	p.williams@uk.ac.ucl.cs
Steven Winnett	swinnett@bbn.com

2.4.6 SNMP over a Multi-protocol Internet (mpsnmp)

Charter

Chair(s):

Theodore Brunner, tob@thumper.bellcore.com

Mailing Lists:

General Discussion: snmp-foo@thumper.bellcore.com
To Subscribe: snmp-foo-request@thumper.bellcore.com
Archive: thumper.bellcore.com:pub/snmp-foo/archive

Description of Working Group:

Within the SNMP management framework, the philosophy is to place the burden of management processing on managers, not on agents. As the Internet evolves to accommodate multiple protocol suites, there may be SNMP agents in the Internet that do not support the recommended method of exchanging SNMP messages using UDP/IP. In these instances, the proper model for managing a multiprotocol internet should be that agents must only be required to support one method of exchanging SNMP messages (i.e., encapsulation of SNMP messages in *one* of the protocol suites of the multi-protocol internet), and the managers support as many encapsulation methods as needed (potentially, all) to communicate with all resources it manages.

The SNMP over a Multi-protocol Internet Working Group is chartered to identify and provide solutions for communication between SNMP agents and managers in those configurations where the recommended method of exchanging SNMP messages using UDP/IP cannot be used; i.e., where a managed resource supports a single protocol suite that protocol is not UDP/IP but another protocol suite of the multi-protocol internet (for example, OSI, AppleTalk, or XNS/IPX).

Questions to be considered include: What are the appropriate protocol suites to consider? What is the appropriate method of encapsulating SNMP? What are the addressing considerations for SNMP messages What new MIB Modules are required? What (positive) effect can SNMP-based management have on resource-sharing among multiple protocols?

Goals and Milestones:

Done	Post an Internet Draft describing operation of SNMP over OSI.
Done	Post an Internet Draft describing operation of SNMP over IPX.
Done	Post an Internet Draft describing operation of SNMP over Appletalk.

Nov 1992 Submit a document describing the operation of SNMP over OSI as a Proposed Standard.

Nov 1992 Submit a document describing the operation of SNMP over IPX as a Proposed Standard.

Nov 1992 Submit a document describing the operation of SNMP over Appletalk as a Proposed Standard.

Internet Drafts:

"SNMP over AppleTalk", 12/23/1991, G. Minshall, M. Ritter <draft-ietf-mpsnmp-appletalk-02.txt>

"SNMP over OSI", 04/10/1992, Marshall Rose <draft-ietf-mpsnmp-overosi-03.txt>

"SNMP over IPX", 06/23/1992, Steve Bostock <draft-ietf-mpsnmp-overipx-01.txt>

CURRENT MEETING REPORT

Reported by Ted Brunner/Bellcore

Minutes of the SNMP over a Multi-protocol Internet Working Group (MPSNMP)

Four drafts were considered in turn:

1. "SNMP over IPX"
2. "SNMP over AppleTalk"
3. "SNMP over OSI"
4. "Guidelines for the Specification of Protocol Support for the SNMP".

The first three are descriptions of how to carry SNMP over particular transport domains, and are meant to be uniform in their directives. Thus many of the issues decided in one draft are similarly decided in the other three. In particular, much of the boilerplate at the beginning of each draft is the same. All three drafts have implementations. The issues considered in each come from a check list, which is enumerated in the final draft.snmp-lists@bir.com.

"SNMP over IPX"

Several issues had come up on the mailing list and were considered in turn. The most time consuming was the issue of OID assignments for Transport Domain and initial Party ID used by Party mechanism in rfc1351-rfc1353. The first plan had been to assign OIDs within the enterprise subtree. This proprietary technique was dropped and a second emerged: assign OIDs under the transportDomains and partyAdmin subtrees of the SNMP Party MIB (rfc1353.) Several draft had decided on this form. The problem is that independent authors may have conflicting assignments. The final choice was that the IANA assign these OIDs thus avoiding conflict. Thus the draft is written with the OID assignments under the experimental subtree (choice of leaf is left blank e.g., xxx) and the IANA will re-position them under the mib-2 subtree at the time the draft becomes RFC.

The wording of the maximum message size had been discussed on the mailing list, and the wording of the recommendation for a 546 byte packet was deemed acceptable.

The wording of the recommendation that an agent support one transport mapping and a manager support as many as necessary was deemed acceptable.

Citations will be fixed.

The Working Group voiced its support that with these minor changes the draft will be posted as Internet Draft and given the usual two week, last call before recommendation for promotion to Proposed Standard.

SNMP over AppleTalk

The same decisions regarding OID assignment were applied here.

On the mailing list was a discussion of the convention for displaying the Appletalk name, as used by the Name Binding Protocol when binding to a DDP address. There are three fields in an AppleTalk name, expressed in the form: "object:type@zone". The delimiter is the issue: either a special character, or a length field. Although special characters are commonly used in expressing an AppleTalk name, it was felt that some utility may be gained with a length field. No manager should break when displaying non-printable ASCII characters. So the convention was adopted whereby the octet preceding each of the three fields contains the length of the field (1-255). (e.g., 6object4type4zone) That length will commonly be an un-printable ASCII character.

The Working Group voiced its support that with these minor changes the draft will be posted as Internet Draft and given the usual two week, last call before recommendation for promotion to Proposed Standard.

"SNMP over OSI"

The OID assignment and language on supporting one transport mapping were brought into alignment with the other drafts.

At the last meeting of the Working Group, the decision to run the SNMP exclusively over CLTP was made. Subsequent feedback from other OSI groups has been positive, according to the OSI Integration Area Director.

A question was raised about GOSIP compliance. Some history was recalled. The original impetus for this draft came from an OSI Integration Working Group (NOOP), a router vendor, and their management needs for CLNP pilot projects. That Group is happy with the form of this draft.

The Working Group voiced its support that with these minor changes the draft will be posted as Internet Draft and given the usual two week, last call before recommendation for promotion to Proposed Standard.

"Guidelines for the Specification of Protocol Support of the SNMP"

The OID assignment will be brought into alignment.

Because of its similarity with the other three drafts, it was proposed that this draft receive similar treatment, with respect to timing and promotion, even though the Working Group had little time to read it before the IETF meeting.

The Working Group voiced its support that with these minor changes the draft will be posted as Internet Draft and given the usual two week, last call before recommendation

that it be an informational RFC.

Having completed its Charter - and in the absence of new work showing up soon to force re-examination of the decision - the Working Group agreed to disband itself.

Attendees

Jonathan Biggar	jon@netlabs.com
Steve Bostock	steveb@novell.com
Theodore Brunner	tob@thumper.bellcore.com
Philip Budne	phil@shiva.com
James Davin	jrd@ptt.lcs.mit.edu
James Halpin	halpin@turkey.sw.stratus.com
Bob Jeckell	rrj@3com.com
Greg Minshall	minshall@wc.novell.com
David Piscitello	dave@sabre.bellcore.com
Mike Ritter	mwritter@applelink.apple.com
Marshall Rose	mrose@dbc.mtview.ca.us
David Waitzman	djw@bbn.com
Steven Wong	wong@took.enet.dec.com

2.4.7 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.

Internet Drafts:

"Routing coordination for X.400 MHS services within a multi protocol / multi network environment", 03/03/1992, U. Eppenberger <draft-ietf-x400ops-mhs-service-00.txt>

"Mapping between X.400(1984/1988) and Mail-11 (DECnet mail)", 03/03/1992, Claudio Allocchio <draft-ietf-x400ops-mapmail-01.txt>

"Operational Requirements for X.400 Management Domains", 03/11/1992, Robert Hagens, Alf Hansen <draft-ietf-x400ops-mgtdomains-01.txt>

"X.400 use of extended character sets", 06/18/1992, Harald Alvestrand <draft-ietf-x400ops-charactersets-00.txt>

"Using the Internet DNS to maintain RFC1327 Address Mapping Tables and X.400 Routing Informations", 09/23/1992, C. Allocchio, A. Bonito, B. Cole, S. Giordano, R. Hagens <draft-ietf-x400ops-dnsx400-00.txt>

CURRENT MEETING REPORT

Reported by Allan Cargille/UWisc

Minutes of the X.400 Operations Working Group (X400OPS)

First Session

Alf Hansen Chaired the meeting. Allan Cargille volunteered to take Minutes. The Agenda was modified to discuss Working Group status and the status of the Wisconsin NSF X.400 Project and the Cosine MHS Project.

Alf distributed the new Charter before the meeting. It was agreed that the proposed new time schedule for the documents would be revised after discussion of the documents. **Note: this was not done in the meeting, and should be done on the mailing list. Action - Alf**

Other issues discussed during the first session included:

Change Control

The IESG (and IAB ?) agreed that change control for RFC1327 (the latest version of mapping between X.400 and RFC822 mail) was assigned to the RARE Working Group on Mail and Messaging (RARE-MSG). This prompted the following discussions:

- Is it OK for IETF RFCs to be assigned to another group?
- How will people in the x400ops Working Group be able to participate in further revisions of this document?
- How will this be publicized ?

It was clarified that RARE-MSG is an open Working Group. Members of the X400OPS Working Group are welcome to join. Here's how to join:

```
Send a message to MAILSERVER@RARE.NL with the following text in
the BODY of the message (NOT the subject).

SUBSCRIBE WG-MSG Your-given-name Your-surname

This will automatically subscribe you to the list.  An automatic
reply will be sent back to you.

The address of the mailing list itself is wg-msg@rare.nl, or
/S=wg-msg/O=rare/PRMD=surf/ADMD=400net/C=nl/.
```

Mailing Lists Dealing with X.400 Issues

There was also discussion about the number of mailing lists which deal with X.400 issues. Often messages are posted to multiple lists. It was recognized that having these multiple lists is a pain, but this Working Group is unlikely to be able to change the situation. It was recommended that when an initial message is posted to multiple lists, the message should clearly identify *one* list on which the follow-up discussion should take place.

Action Items from March 92 Meeting

a. John Sherburne (SPRINT) will work with Tony Genovese to figure out how US can provide an MTA that has X.25 connectivity.

- Tony reported that accepting ADMD = <single space> is a problem for Sprint. He did not know if that is for technical, political, or financial reasons.

  [action] Tony continue to work on a WEP which is accessible over public X.25.

- Ed Albrigo from the Corporation for Open System (COS) gave a report on their X.400 activities. They are working on the following:

  1. Establishing direct network-layer connections to the Internet. They plan to route both IP and CLNP.

  2. Establishing X.400 links which connect the OSINet X.400 community to the GO-MHS community.

  3. They are planning to go to complete "electronic-only" communication with ten COS member companies by December 1992.

  Ed confirmed that COS will comply with current RFCs and recommendations for the GO-MHS community.

  It was clarified that COS uses X.25 in their private OSINet network, but that is a private network that is not connected to public X.25.

- There was a discussion about connections to ATTmail. Internet RFC822 mail users should be able to send mail to all ATTmail users. However, the ATTmail <–> Internet mail gateway produces bad addresses, so mail is often un-replyable.

b. Urs will ask the COSINE MHS Project Team to submit the address mapping table procedures as a draft RFC. - Done.

c. Stef - Start a discussion on X.400 OPS and WG1 lists about ADMD name in the U.S. See section 3.1.2. [of March 1992 Minutes]

- Not done.
- Note that the rare-wg1 mailing list has been succeeded by the wg-msg list (see section 2 above).

[action] Stef start this discussion. [action] Someone email Stef to start this discussion. [done]

See related discussion of this in Agenda item 5.

d. Alf will send the updated Charter to the list. - Done

e. Claudio will produce a draft document that will propose a method for using DNS to store X.400 to RFC 822 mapping and routing. - Done.

f. Claudio will follow up the MAIL 11 mapping document. - Done

g. Harald will follow up the International Character set document. - Done

Status of X.400 Operations

a. Allan Cargille discussed the status and future of the NSF X.400 Project. The project has been running since August, 1990 and is now toward the end of the initial grant. The project has operated the experimental PRMD "XNREN". Fifteen to twenty sites have registered as members of this PRMD, but only approximately five are currently exchanging X.400 traffic. The project has acted as a coordination point for U.S. entries in the RFC987/1148/1327 mapping tables. The project also served as a beta site for several PP releases, and developed and contributed software to support the Fujitsu dexNET 200 fax modem in PP. The project is operating a primary MTA running PP 6.0 on a dedicated DecStation 3100/Ultrix. Some sites, including Wisconsin, are running the IBM/Wisconsin Argo X.400 software, which includes a UA. The project has also acted as a Well-Known Entry Point (WEP) to the Cosine MHS Project (see below). We are seeking an extension of the grant to continue supporting a stable U.S. WEP and to participate in the ongoing research work to develop a stable X.400 infrastructure. Without continued funding, our project will end at the end of this calendar year.

b. Jim Romaguera presented an overview of the Cosine MHS Project at SWITCH (Switzerland). That project began in (January 1991 ?) and continued work begun by the RARE MHS Project Team. They coordinate the academic and research X.400 service in Europe. They have finished 80 percent of their goals for the current project period, which ends at the

end of this calendar year. The project supports international X.400 connections between all Western European countries, as well as Greece, Slovenia, Lithuania, the United States, Canada, Australia, New Zealand, South Africa, China, India, and the Republic of Korea. Some countries have multiple networks participating in the service. Most European participants have private connections to one or more commercial ADMDs. Some are purchasing value-added services (such as fax gateways) from ADMDs. Several project participants have online services available (via telnet or over X.25) to translate between X.400 and RFC822 addresses according to the current mapping rules.

The exact future of the project is unclear, but it is expected that they will continue. It is likely that the future project will be coordinated by the RARE Operational Unit and will be contracted out.

The project team is still working on several projects. They plan to have a daily RFC1327 mapping table update tool operational by the end of this year. They are working on evaluating publicly available X.400 implementations. They plan to produce a catalog of existing X.400 implementations. They have done work on evaluating ADMDs and plan to report on this (verifying connectivity, etc). They plan to produce a tutorial and overview on RFC1327. They have done work on evaluating international X.400 connections, and are working on tools to automatically process a common statistics format. They are also working on a connectivity tool which will be based on sending mail to echo servers and evaluating the results. Lastly, they operate a file server with lots of documents. You can reach the fileserver via anonymous ftp to host "nic.switch.ch".

Discussion

- It was recommended to refer to RFC1292 (a catalog of X.500 implementations) for X.400 product evaluations.
- Will this information on implementations be released as an RFC ?
- There is a question of liability when producing such evaluations.
- It was recommended that vender and user comments about implementations be placed in separate documents.

c. Stef reported on the current work of the MHS-MD study group on ADMD/PRMD naming. By way of review, Stef covered the history of connections between the U.S. Internet and commercial email services. Vint Cerf was the founder of MCI Mail and then went to CNRI. He concluded agreements on behalf of the Internet with MCIMail, ATTmail, G.E. Information Services, and CompuServe (and possible others) that are "sender keeps all revenue" agreements.

There was also discussion about what internal protocols these services use. All operate gateways between RFC822 and their internal protocol. Several problems were discussed.

- If the service is using a poor or nonstandard gateway, then the addresses coming out of the gateway are messed up.
- People did not know of any connections between U.S. commercial ADMDs.
- There are no connections between the U.S. Internet X.400 community and commercial ADMDs.

Current MHS-MD status. Commercial ADMDs have been arbitrarily selecting their own names, and then arbitrarily naming PRMDs under their ADMD. There is strong feeling that these existing (ADMD, PRMD) name pairs must be valid in the future. Any new registration procedure must support these existing names. The Group is also working on a structure for a U.S. ADMD backbone, which does not mean a specific ADMD. Currently the string ADMD=USBB is being used to refer to such a structure. Stef cautioned us that the "USBB" name is just a placeholder and is likely to change to some other (as yet undefined) text string. PRMD names could then be registered under this "ADMD=USBB". There are still unresolved questions about how the USBB should be routed and supported.

Stef proposed that the U.S. Internet declare itself as an ADMD. This could be justified because at present, all the other ADMDs are self-declared as well. Stef argued that there is currently no regulation of US X.400 service providers, so each ADMD is more or less making up their own rules as they proceed. Many people are making lots of assumptions. One has been that the INTERNET does not qualify to be an ADMD, and the that other ADMDs would block its attempt to assert that it is an ADMD.

Discussion:

- The issue of connecting to the U.S. ADMDs is not an issue of naming, it is an issue of service agreements and charging. The routing can be worked out.
- Connections over X.25 will probably be necessary to connect to the commercial ADMDs, although many US carriers are moving to offer IP service, and to interconnect with the INTERNET.
- The Internet ADMD could offer to provide RFC1327 gateway services to the commercial ADMDs. That way the gateways would be operated according to existing agreements and recommendations and would generate "good" addresses.
- If the Internet succeeds in defining itself as an ADMD, then the other C=US ADMD service providers can no longer use the excuse that they "cannot pass ADMD-ADMD

traffic via the INTERNET PRMD".

- If the commercial services were interested, the Internet ADMD could play a role as a relay between them. [Note - this would not necessarily require commercial traffic to flow across the research Internet.]

There was a proposal to decide on the matter at this meeting. There was heated argument that the issue had not been discussed before the meeting, and should be discussed more in a wider forum and on the mailing list. It was agreed that Stef would write an internet draft proposing to create an ADMD=Internet [action]. If approved in the future, this paper could evolve into an RFC.

The Working Group recommended that each country should write an Internet Draft describing the national solution for X.400 addressing of Internet addresses. Stef's draft could be used as a template for other countries' Internet Drafts. The result will in the end be (if the drafts are approved) a series of RFCs. [This paragraph supplied by Alf Hansen.]

Future U.S. Internet X.400 Organization - not discussed beyond the above information.

Second Session

Continuation of Connections to ADMDs Discussion.

Steve Hardcastle-Kille proposed generating a document that addresses the issue of ADMDs and how they are connected to the R&D world (or "Internet" to coin a phrase). The contents of this document should be something like:

- ADMDs presently connected to the Internet (or R&D world, same thing, as I'm talking about the global Internet).
- Policy restrictions on such connections ie. are they available for free & for anyone on the Internet, can R&D people relay via a connected ADMD to 3rd party ADMDs, etc.
- Whether the ADMDs are using RFC 1327 gateways & the global mapping tables
- Which PRMDs these ADMDs support - ADMD connectivity between themselves. - anything else that fits in to the above context.

Goals are to:

- Stimulate ADMDs to deploy well run ADMD to Internet connections, preferably by using R&D operated gateways.

- Document the PRMDs reachable via ADMDs and of course the ADMD's connectivity to other ADMDs.

Jim Romaguera (wearing the hat of NetConsult AG, not the Cosine MHS Project Team) volunteered to write a draft document [action]. [notes in this (cont'd) section courtesy of Jim R.]

Document Review

In general, detailed comments are not included if a new version of the document will be released.

a. "X.400 use of extended character sets" (Harald Alvestrand). Discussion. Harald will update the document and release the updated version as an Internet draft [action]. The draft will be discussed at the upcoming RARE Character Set and RARE Messaging meetings. These comments will be presented at the next IETF meeting, and the document will be finalized.

b. "Operational Requirements for X.400 Management Domains in the GO-MHS Community" (Hansen/Hagens). Comments were taken on the document. The document will be revised and a new Internet Draft will be released [action].

There was discussion about what kind of RFC this document should be released as. People felt that it should be a requirement that X.400 domains should support the "postmaster" address in the same manner as RFC822 domains do. It was proposed that a very short RFC be drafted which explains the need for supporting "postmaster" addresses. This short postmaster RFC will then be advanced in the standards track. Allan Cargille volunteered to write the RFC [action]. It will use the recommendations from the recent Cosine MHS Managers meeting as a starting point. It was pointed out that to support the introduction of X.400(88), both S=Postmaster and CN=Postmaster must be supported.

The revised Hansen/Hagens paper cannot be progressed as an RFC until the Eppenberger routing paper and Cargille Postmaster paper are also ready to be submitted, because it references those documents. The document may also have to be modified based on the group's recommendations for C=us/ADMD=Internet.

c. "Routing coordination for X.400 MHS services within a multi-protocol/ multi-network environment" (Urs Eppenberger). Changes to this document were discussed in light of a recent submission by Panos Tsigaridas, "MHS Information Exchange Format" (MHS-IEF). Panos' paper recommended using the same basic information and routing algorithm as the

Eppenberger document, but providing a syntax and structure so that this information could be easily placed into X.500 under well-known places. Further information already stored in X.500 could easily be extracted by tools and translated into the proposed text format. These text tables could then by exchanged the "old-fashioned" way (E-mail).

The desire to support X.500 must be weighed against the fact that this new document format is needed immediately and in fact is already being introduced in the Cosine MHS Project. Changing the document format would introduce delays due to discussion and take longer to become operational. It was agreed that Urs, Panos, and Steve H-K. would meet to see if minimal changes could be made to the Eppenberger document which would make it easier to store the information in X.500. Steve reported that they agreed that Panos would propose a set of detailed "short term" change requests to Urs's document [action]. A revised document should be sent out, which should be approved via email and then submitted as an experimental RFC [action].

d. "Using the Internet DNS to maintain RFC987/RFC1148 Address Mapping Tables and X.400 Routing Informations" (Allocchio, Bonito, & Giordano). All three tables will be stored under the domain ".x400.arpa". Change control will still be centralized – the tables will still be collected and managed by the Cosine MHS Project Team. The use of the DNS tables will be described in a separate document [action]. Mapping conventions are used to represent the RFC1327 table entries in a format that is legal for the DNS. Claudio will produce a new version of the document, and distribute it to the DNS and x400ops mailing lists [action]. If consensus is reached, the document will be submitted as an Experimental RFC.

e. OSI area procedures. Erik Huizer requested that to progress a document in the OSI area as an Internet Draft, people should send email to Dave Piscitello (dave@sabre.bellcore.com), himself (huizer@surfnet.nl) and CC the IESG Secretary Greg Vaudreuil (gvaudre@nri.reston.va.us). [Note - this information should probably be sent to all the OSI area working groups. [action]] Erik also reported the following procedures for IETF OSI working groups [actions]:

- He will create a mailing list for these Working Group Chairs.
- He will distribute each message to him from higher IETF people to Working Group Chairs (Chairpersons).

There was also discussion about what classes of RFCs there are. RFCs *not* on the standards track can be classified as "Informational" or "Experimental". RFCs on the standards track proceed from "Proposed Standard" to "Draft Standard" to "Standard". [Note - is this documented in an RFC?] It was also pointed out that RFCs cannot reference Internet Drafts, but they may reference any class of RFC.

Major Operation Problem

Not discussed.

Review of Action Items

Deferred to mailing list, due to time. See below.

Other Business, and Plans for Next Meeting

Erik Huizer (OSI Area Co-Director) proposed to resume the "old" meeting schedule for the OSI area at the next IETF. Other than that, the next meeting schedule not discussed. Erik will distribute this new schedule [action].

We decided to have the next x400ops meeting at the next IETF meeting in Washington DC, U.S.A., during the week Nov. 16-20, 1992.

Revised Summary of Action Items

Allan Cargille	Distribute draft Minutes. - Done.
Alf Hansen	Revise timetable for documents on new Charter by discussion on the list.
	Update Operational Requirements document and release as an Internet Draft.
Tony Genovese	Continue to work on a WEP which is accessible over public X.25.
Einar Stefferud	Start discussion on mailing lists about U.S. ADMD naming issues. - Done.
	Write an Internet Draft proposing to create ADMD=Internet.
Someone	Email Stef to start this discussion. - Done.
Jim Romaguera	(NetConsult AG) Generate draft document that addresses the issue of ADMDs and how they are connected to the R&D world.
Harald Alvestrand	Update document on extended character sets and release as an Internet Draft.
Panos Tsigaridas	Provide a set of detailed "short term" change requests to Urs' routing document.
Urs Eppenberger	Release revised version of routing coordination document (if there are any changes). Hopefully get consensus on mailing list about the document and submit as an RFC.

Claudio Allocchio	Produce new version of the X.400 DNS paper and distribute it to the x400ops and DNS mailing lists. If consensus is reached, submit document as Experimental RFC.
	Produce new document explaining how the X.400 DNS tables should be used and distribute to x400ops list.
Erik Huizer	Distribute information on the procedure for progressing a document in the IETF OSI area to all area mailing lists.
	Create a mailing list for IETF OSI area Working Group Chairs.
	Distribute working group meeting schedule for OSI area for next IETF meeting.

Attendees

Ed Albrigo	ealbrigo@cos.com
Claudio Allocchio	c.allocchio@elettra.trieste.it
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
C. Allan Cargille	cargille@cs.wisc.edu
Chris Chiotasso	chris@artel.com
Cyrus Chow	cchow@orion.arc.nasa.gov
Alan Clegg	abc@concert.net
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Tom Easterday	tom@cic.net
Urs Eppenberger	eppenberger@switch.ch
Tom Farinelli	tcf@tyco.ncsc.mil
Osten Franberg	euaokf@eua.ericsson.se
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@nersc.gov
Arlene Getchell	getchell@nersc.gov
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@isode.com
Erik Huizer	huizer@surfnet.nl
Takashi Ikemoto	tikemoto@xerox.com
Kevin Jordan	kej@udev.cdc.com
Mark Knopper	mak@merit.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Linda Winkler	lwinkler@anl.gov
Steven Winnett	swinnett@bbn.com
Russ Wright	wright@lbl.gov

2.5 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 Bernhard Stockman/NORDUnet

Operational Requirements Area Directorate (ORAD)

The Operational Requirements Area Directorate has so far been an open Group with sessions at IETF meetings. The meeting discussed this time the formation of a closed Directorate similar to SAAG in the security area or USAC in the User Services Area. It was decided that ORAD will consist of 12-16 people with specific duties as described in the ORAD Job description. Meeting at IETF's will continue to be open. Two mailing lists will be created: one for ORAD as such and one, orad-interest@ans.net, for people interested in operational related issues.

Job Description for ORAD.

ORAD will be chartered to:

- Approve Charters of proposed working groups.
- Explicitly create necessary working groups.
- Track progress of all Operational Requirements Area working groups.
- Review *all* Internet Drafts for operational impacts.
- Liaison with other IETF areas and external groups.
- Others...

Operational Concerns.

A broad review was made over topics for ORAD. An extensive list was produced containing most operational related topics. For each topic a short description of the topic will be written down. Among the selected items some where regarded as more important and deserving immediate attention especially the routing and addressing problems.

Current routing and addressing problems. This includes very short-term issues like reuse of IP address space and effects of more thrust in routers. An address assingment plan is under preparation as part of the CIDR concept. Proposed routing and addressing plans will will be reviewed by ORAD with respect to the operational impact such as:

- Details of proposed schemes.
- Changes required in hosts and routers.

- Performance impact.
- Scaling properties.
- Support for old host and routers.
- Deployment plan.
- Security impact.
- Operational management and training impact.

The IEPG initiated specification and implementation of a Global Internet eXchange (GIX) is part of the solution of the today routing problems and will accordingly be on the ORAD agenda.

Other operational items mentioned were:

- Globally coordinated DNS.
- NIC/NOC coordination.
- Network performance measurements.
- Security.
- OSI operations.

IP Addressing Plan BOF (IPADDR)

The intention with this BOF was to discussed the possibilities of an Internet Addressing plan as part of the CIDR deployment. Two Internet Drafts have been submitted proposing similar schemes. One main difference was the proposed size of blocks to be assigned. After the BOF it was decided to make the following recommendations:

1. Blocks are allocated initially to Europe and the US with (216.*.* - 219.*.* to Europe) and (220.*.* - 223.*.* to the US).

2. Blocks will be distributed at a size of (256*256 class C networks).

3. Class B address space will only be used in rare circumstances for example to international multihomed organizations.

4. Upper half of Class A address space will be reserved for future use.

The BOF discussed issues related to the administration of delegation of authorities. It is recommended that there will be top level geographically based top authorities.

BGP Deployment and Application Working Group (BGPDEPL)

The Internet BGP topology was presented. There are about 25 administrative systems which are reachable from the NSFnet via EGP2, and about 40 which are reachable via BGP.

Peter Loethberg described the European EBONE BGP deployment. In several ways EBONE is further along in BGP deployment than US networks.

Several vendors described the status of their implementations. CISCO and Cornell gated have BGP3 running today, and are working on BGP4. BBN has BGP3 running and is working on BGP4. Other vendors including Proteon and Wellfleet are actively working on BGP3 in preparation for BGP4.

ANS intends to have BGP4 deployed January 1993 and has offered to help vendors with interoperability testing. Vendors can arrange to bring equipment into the ANS test facility. ANS is also exploring support for remote testing by "tunneling" BGP from the ANS test network through the NSFnet. Contact Jordan Becker for further information.

Benchmarking Methodology Working Group (BMWG)

Two proposed additional performance tests were discussed. The Group concluded that one of the proposals should not be adopted. The second proposal, to add application level performance test, should be discussed with the test equipment vendors to see if they would be willing to implement such tests.

The latest draft was discussed and a few editorial changes were suggested. It was strongly urged that some additions be made to the draft to add the rational behind the various tests and procedures.

Network Joint Management Working Group (NJM)

The Operational Impact of the NSF recompetition was investigated. A mailing list - recompete@nsf.gov - has been created to discuss these topics. Papers are available at expres.cise.nsf.gov.

Operational implications of address changes. How should N CIDR class C networks be routed internally.

Application and Information Services was discussed with respect to troubleshooting and maintenance. There is a need for service level maps, e.g., video topology maps.

The session continued with some networking "war stories". The BBN network 192.1.1 will be discontinued because some vendors of medical equipment use this network as their default! There are rumors that the SUN example network number was returned and is now re-issued to some unsuspecting site.

Network Status Reports (NETSTAT)

Four networks delivered reports:

1. EBONE - Bernhard Stockman

2. NASA Science Internet - Jeffry Burgan
3. ANS/Merit T1/T3 J- Jordan Becker/Elise Gerich
4. Brazil networking - Pushpendra Motha

EBONE - KTH (The Royal Institute of Technology) in Stockholm has been nominated as the EBONE Network Operations Center. The topology of EBONE is today a pentagram with core hubs in Stockholm (Sweden), London (United Kingdom), Montpellier (France), Geneva (Switzerland) and Amsterdam (Netherlands). On-line information on EBONE is available at nic.nordu.net or archive.ripe.net in directory ebone.

NSI - The Korean link has now been moved from Hawaii to NASA Ames as this solution was cheaper! The "virtual" upgrade of the link between Hawaii is done with compression.

ANS/Merit T1/T3 - T1 traffic is now falling. 75 percent of T1 AS's has been moved to the T3 net. The T1 network continue to carry CLNP traffic. The amount of configured networks are today around 5700 of which around 4500 are announced. The RFC960 FDDI cards will be upgraded in August. The target is to have a complete cutover at the end of August. CLNP is planned to be installed during the fall. Backup of the T3 net is being planned as T1 connections between diverse ENSS's and CNSS's.

Brazil Networking - Networking in Brazil is fostered by CNPK an analogous organization to the US NSF. The network form a core backbone with hubs are in Rio and Sao Paulo with bandwidths between 9.6 to 64 Kbps. The network is connected to the US via a 64 Kbps line from Rio to CERFnet and a 64 Kbps link between Sao Paulo and ESnet.

Operational Statistics Working Group (OPSTAT)

A review was made of the Internet Draft (I-D) submitted after the previous IETF. Some comments on the I-D had been received and the meeting discussed these.

The "time-section" was renamed to "label-section" and extended with stop-time and a data-file-name. Clarifications will be added on how sections could be stored, one data-section per file or multiple sections per file. It was agreed that, how files are physically arranged was outside the scope of the draft. Clarification is also need for how tags should be used. A comment field will be added. Finally some editorial remarks will be included.

After decided changes are included, the draft will be circulated for a final round of comments. When ready the paper will be submitted as an Informational RFC.

Eric Hood, Executive Director of FARNET announced the possibility of FARNET funding the development of reference tools according to the OPSTAT model. The meeting agreed that this was important. Eric Hood will make a survey of already ongoing efforts and forward the result to the OPSTAT Working Group before a decision for funding is made.

User Connectivity Problems Working Group (UCP)

Dan Long reported that Tom Sandoski of ConcertNet has released both a curses and X-windows interface to his freeware Trouble Ticket System on ftp.concert.net in the pub/tickets directory.

There are now 50 NOC's listed in the Network Service Center Phonebook, maintained at NEARnet. Finger access is available for peer NOC's and a publicly-searchable subset of the data will be announced soon.

Most of the meeting was devoted to refining the description of information that Paul Zawada has compiled for handing off tickets from one NSC to another. Several organizations have expressed interest in standardizing these handoffs. Paul will revise his ASN.1 version and Kaj Tesink will draft a specification of how to encode this information in email messages.

CURRENT MEETING REPORT

Reported by Bernhard Stockman/SUNET and Steven Winnett/BBN

Minutes of the IP Addressing Plan BOF (IPADDR)

Introduction

The BOF was Chaired by Bernhard Stockman of SUNET. The subject of the meeting was the Supernetting or CIDR (Classless Inter-Domain Routing) scheme to extend IP addressing for a few years as presented by V. Fuller, T. Li, J. Yu, and K. Varadhan in RFC 1338, "Supernetting: an Address Assignment and Aggregation Strategy".

The Chair expressed the hope that the BOF session would come to a consensus agreement as to how this problem could be solved.

Presentation

Tony Li (cisco) gave a presentation of the CIDR scheme as outlined in RFC 1337. The presentation discussed the mechanics of how the hierarchical addressing scheme would work (allowing groups of neighboring class C networks to be described as aggregates) along with examples. Blocks would initially be allocated at the continental level, with service providers (a.k.a. regionals) in turn allocating groups of networks (where applicable) to clients.

Tony Li stressed that the scheme as presented in RFC 1337 was designed as a solution for from 3 to 7 years only to the IP addressing problem in order to buy some time to come up with an acceptable longer-term solution.

The handling of unused addressed space, of multihomed sites, and of service provider changes under CIDR was presented.

As an initial allocation, North America and its service providers would be allocated the Class C networks under <220.0.0.0 - 252.0.0.0>, Europe and its service providers would be allocated the Class C networks under <216.0.0.0 - 252.0.0.0>, and the remaining continents would be allocated about 10,00 Class C networks each.

Discussions

There was some discussion of the block sizes to be allocated (i.e., would these be in 252 x 255 chunks (mask 252.0.0.0) or 255 x 255 chunks (mask 255.0.0.0).

The issue of renumbering hosts was discussed. R. Callon pointed out the desirability of this, stating that if addresses were reassigned so that routing could be done on the flat space of all regionals, then CIDR would scale to the point of every IP address assigned. Many others in the room, however, questioned whether it was realistic to expect that such

a renumbering could ever be agreed to by the Internet community.

There was some discussion of the problems which could be caused by clients giving back gaps of unallocated network numbers.

Vint Cerf stated that it would be useful to gather data with respect to the number of networks already assigned, and the number of existing service providers, on a per country basis. This might provide useful background information so that scaling choices could be made in a more informed manner.

Decisions

Consensus appeared to be reached on the general merits of the hierarchical addressing plan as presented as a short-term solution to the IP addressing problem.

There was no consensus over the issue of renumbering existing hosts and this was tabled for the next BOF.

With respect to the block size issue, the Chair stated that there appeared to be a consensus around 255 x 255 chunks.

Attendees

Steve Alexander	stevea@i88.isc.com
Nagaraj Arunkumar	nak@3com.com
Ashar Aziz	ashar.aziz@eng.sun.com
William Babson	bill@penril.com
John Ballard	jballard@microsoft.com
William Barns	barns@gateway.mitre.org
Tony Bates	tony@ean-relay.ac.uk
Jordan Becker	becker@nis.ans.net
James Beers	beers@nr-tech.cit.cornell.edu
Lou Berger	lberger@penril.com
David Bolen	db3l@nis.ans.net
Scott Brim	swb@cornell.edu
Steve Buchko	stevebu@newbridge.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Greg Celmainis	gregc@newbridge.com
Vinton Cerf	vcerf@nri.reston.va.us
John Chang	jrc@uswest.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Henry Clark	henryc@oar.net
James Cobban	jcobban@bnr.ca

Richard Colella	colella@osi.ncsl.nist.gov
Robert Collet	rcollet@icm1.icp.net
Rob Coltun	rcoltun@ni.umd.edu
Jim Comen	comenj@interlan.interlan.com
Michael Conn	0043874510@mcimail.com
John Curran	jcurran@bbn.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Peter DiCamillo	Peter_DiCamillo@brown.edu
Chas DiFatta	chas@sei.cmu.edu
Pierre Dupont	dupont@mdd.comm.mot.com
Tom Easterday	tom@cic.net
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Dennis Ferguson	dennis@mrbill.canet.ca
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Cliff Frost	cliff@cmsa.berkeley.edu
Vince Fuller	vaf@stanford.edu
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Der-Hwa Gan	dhg@nsd.3com.com
Elise Gerich	epg@merit.edu
Rob Graham	robert_graham@protools.com
Phillip Gross	pgross@nis.ans.net
Simon Hackett	simon@internode.com.au
Steve Hardcastle-Kille	s.kille@isode.com
Kathleen Huber	khuber@bbn.com
Matthew Jonson	jonson@server.af.mil
Kevin Jordan	kej@udev.cdc.com
Dan Jordt	danj@nwnet.net
Scott Kaplan	scott@ftp.com
Kenneth Key	key@cs.utk.edu
Paulina Knibbe	knibbe@cisco.com
Mark Knopper	mak@merit.edu
John Krawczyk	jkrawczy@wellfleet.com
David Kumpf	kumpf@hprnd.rose.hp.com
John Labbe	labbe@merit.edu
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Sharon Lewis	lewis@cs.utk.edu
Fong-Ching Liaw	fong.liaw@eng.sun.com
Hock-Koon Lim	lim@po.cwru.edu
Arthur Lin	yalin@srv.pacbell.com
Tony Li	tli@cisco.com
Joshua Littlefield	josh@cayman.com
Andrew Maffei	arm@aqua.whoi.edu
Carl Malamud	carl@malamud.com

Bill Manning	bmanning@rice.edu
Ed Menze	menze@cs.arizona.edu
Pushpendra Mohta	pushp@cerf.net
Paul Moran	Paul_Moran@3com.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Donald Morris	morris@ucar.edu
Mark Needleman	mhn@stubbs.ucop.edu
Erik Nordmark	nordmark@eng.sun.com
Kraig Owen	tko@merit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Jon Postel	postel@isi.edu
Stephanie Price Marasciullo	price@cmc.com
Mark Prior	mrp@itd.adelaide.edu.au
Bala Rajagopalan	braja@qsun.att.com
David Reiser	dsr@codex.prds.cdx.mot.com
Sheri Repucci	smr@merit.edu
Simon Rosenthal	simon@bos.locus.com
Anthony Rutkowski	amr@sprint.com
Henry Sanders	henrysa@microsoft.com
Hellen Sears	sears@interlan.interlan.com
Tim Seaver	tas@concert.net
Karen Seo	kseo@bbn.com
Erik Sherk	sherk@sura.net
John Shriver	jas@proton.com
William Simpson	Bill_Simpson@um.cc.umich.edu
Henry Sinnreich	hsinnreich@mcimail.com
Frank Solensky	solensky@clearpoint.com
Michael St. Johns	stjohns@umd5.umd.edu
Martha Steenstrup	msteenst@bbn.com
Bernhard Stockman	boss@sunet.se
Fumio Teraoka	tera@csl.sony.co.jp
Marten Terpstra	terpstra@ripe.net
Linda Tom	toml@interlan.interlan.com
Claudio Topolcic	topolcic@nri.reston.va.us
Panos-Gavriil Tsigaridas	tsigaridas@fokus.berlin.gmd.dbp.de
Kannan Varadhan	kannan@oar.net
Ross Veach	rrv@uiuc.edu
Curtis Villamizar	curtis@ans.net
Bernhard Volz	volz@process.com
John Wagner	jwagner@princeton.edu
Justin Walker	justin@apple.com
Carol Ward	cward@westnet.net
William Warner	warner@ohio.gov
Liming Wei	liming.wei@eng.sun.com
Evan Wetstone	evan@rice.edu

Chris Wheeler	cwheeler@cac.washington.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
Steven Winnett	swinnett@bbn.com
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com
Joseph Zur	zur@fibhaifa.com

IP Address Allocation

- Based on RFC 1237
- Provide hierarchical addressing for CIDR
- Provide hierarchical address allocation scheme
- Migration issues
- Recommend initial address allocations for US, Europe

Hierarchical addressing

- Assign sites blocks of addresses
- Masks get longer with the depth of the hierarchy
- Neighbors in the hierarchy have adjacent addresses
- A group of neighbors can be described by an aggregate

Example:

Site 1: 192.24.0.0 through 192.24.7.0 (aka 192.24.0.0 255.255.248.0)

Site 2: 192.24.16.0 255.255.240.0

Site 3: 192.24.8.0 255.255.252.0

Generate aggregate: 192.24.0.0 255.255.0.0

Levels of Hierarchy

- Root (existing) - responsible for allocation to continents, large multicontinental organizations, and intercontinental PDNs
- Continental - responsible for allocation throughout continent, including multihomed organizations, backbone networks, and service providers
- Service provider - responsible for allocation to all clients
- Individual site - responsible for allocation within the site

Any address allocation authority may recall unused address space that it has allocated to another address allocation authority.

Multihomed sites

- Take addresses from any possible provider
- Other providers must advertise aggregate
- Service provider which gives allocation may also choose to advertise site

Example:

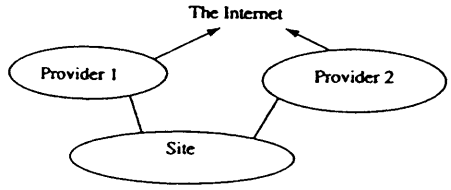

Provider 1: 192.24.0.0 255.255.0.0

Provider 2: 192.25.0.0 255.255.0.0

Site: 192.24.0.0 255.255.248.0 (advertised by provider 2 and optionally, provider 1)

Change of provider

What happens when a site changes from provider 1 to provider 2?

1. Original address space is from provider 1.
2. Provider 2 now advertises site, allocates new address space for site out of provider 2's address space.
3. Site now functions similarly to a multi-homed site (see previous slide).
4. Site SHOULD migrate to new address space. Note that the time frame for the migration is set by agreement between the site and provider 2.

Other multihomed options

- Get address allocation from authority that spans organization (e.g., continental authority, root)
- Get some address allocation from each provider. Allocate based on proximity to provider.

Migration issues

- Existing networks need not change addresses
- Establishing an address allocation authority may be slow. A higher level authority may act in proxy until the new authority is active.

Initial allocation

For North America:

- Allocate <220.0.0.0 252.0.0.0> (262144 class C nets)
- 12% of unallocated space
- Each service provider gets an allocation that should last 2 years
- Class A nets are retained by the root for intercontinental PDNs
- Class B nets are retained by the root for multicontinental, multi-homed organizations

For Europe:

- Allocate <216.0.0.0 252.0.0.0> (same)
- rest as above

For Asia, Africa, Australia, and South America:

- Allocation size TBD, but approx. 10k Class C nets each

CURRENT MEETING REPORT

Minutes of the Router Requirements Checklist BOF (RREQLIST)

Report not submitted.

Attendees

William Babson	bill@penril.com
Steve Buchko	stevebu@newbridge.com
Greg Celmainis	gregc@newbridge.com
Dean Cheng	dean@sun2.retix.com
Robert Ching	natadm!rching@uunet.uu.net
AnneMarie Freitas	afreitas@microcom.com
Vince Fuller	vaf@stanford.edu
Ian Heavens	ian@spider.co.uk
Matthew Jonson	jonson@server.af.mil
Alan Kullberg	akullber@bbn.com
Joshua Littlefield	josh@cayman.com
Pushpendra Mohta	pushp@cerf.net
Steven Onishi	sonishi@wellfleet.com
Stephanie Price Marasciullo	price@cmc.com
Michael St. Johns	stjohns@umd5.umd.edu
Brad Steinka	brad@microcom.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Fumio Teraoka	tera@csl.sony.co.jp
David Wang	wang@xylogics.com
Carol Ward	cward@westnet.net
Evan Wetstone	evan@rice.edu
Steven Winnett	swinnett@bbn.com

2.5.1 BGP Deployment and Application (bgpdepl)

Charter

Chair(s):

Jessica Yu, jyy@merit.edu

Mailing Lists:

General Discussion: bgpd@merit.edu
To Subscribe: bgpd-request@merit.edu
Archive: /pub/bgpd-archive on merit.edu

Description of Working Group:

The major purpose of this Group is to coordinate BGP deployment and application in the current Internet.

It intends to create a forum for BGP users to share BGP deployment experiences and also provide a channel for users to communicate with router vendors who implemented or who are implementing BGP. It also intends to discuss BGP policy application and coordinate policy implementation in the current internet routing environment which includes defining the usage of policy, defining a mechanism to share policy information, etc.

Goals and Milestones:

Ongoing	Facilitate the deployment of BGP as widely as possible.
TBD	Define the issues and the needs of policy routing in the current Internet architecture. Discuss how BGP policy routing capability applies to Internet policy routing needs. A document may be generate
Mar 1993	Post an Internet Draft defining a mechanism to share policy information between Administrative Domains.
Dec 1992	Post as an Internet Draft, a report of BGP deployment status.

CURRENT MEETING REPORT

Reported by Matt Mathis/PSC

Minutes of the Border Gateway Protocol Deployment and Application Working Group (BGPDEPL)

Executive Summary

The Internet BGP topology was presented. There are about 25 administrative systems which are reachable from the NSFnet via EGP2, and about 40 which are reachable via BGP.

Peter Lothberg described the European EBONE BGP deployment. In several ways EBONE is further along in BGP deployment than US networks.

Several vendors described the status of their implementations. CISCO and Cornell gated have BGP3 running today, and are working on BGP4. BBN has BGP3 running and is working on BGP4. Other vendors including Proteon and Wellfleet are actively working on BGP3 in preparation for BGP4.

ANS intends to have BGP4 deployed January 1993 and has offered to help vendors with interoperability testing. Vendors can arrange to bring equipment into the ANS test facility. ANS is also exploring support for remote testing by "tunneling" BGP from the ANS test network through the NSFnet. Contact Jordan Becker for further information.

Minutes

Two representations of the known BGP topology of the Internet were presented. The first was a map compiled by Jessica Yu, of Merit, which is compilation of several sources including Merit configuration databases and routing information extracted from the NSFnet T1 and T3 backbones. The other was BGP routing table as extracted from a router in Pittsburgh which was peering with both NSFnet backbones. The BGP topology includes about 40 administrative systems. There are only 25 administrative systems which are visible with EGP2 as the protocol of origin. Note that EGP2 does not propagate further topological detail, so there are additional EGP2 only administrative systems beyond the 25 which are visible.

One interesting artifact was noted: The T1 backbone does not run IBGP so it can not propagate BGP path information. All BGP routes learned via the T1 backbone show one administrative system beyond the backbone itself with the "incomplete" origin attribute. This is dangerous because it completely defeats BGP loop suppression.

Peter Lothberg described the European EBONE BGP deployment. They are using BGP3 in an unusual configuration where it is, in effect, their interior routing protocol. All EBONE routers are border routers. (There are no interior routers). The EBONE IGP (cisco's IGRP)

is used only to distribute the EBONE internal links and interfaces. The IGP does not carry exterior routes at all. IBGP is run as a fully interconnected mesh between all routers. Tony Li of cisco added a configuration feature to BGP to defeat the IGP alignment check. Thus every router has an IBGP route for for all exterior networks and an IGP route for EBONE routers. Packet switching takes two route lookups: one to select the exit router and then to find the interior route to that exit router.

Peter described several routing misadventures that they encountered, as well as European configuration management issues and related politics.

Several vendors described the status of their implementations. Nearly all of the major router vendors were present and actively working on various aspects of BGP.

- Proteon plans to support BGP3 in release 13 for "all platforms" by the end of the year.
- Wellfleet plans to have BGP3 in mid '93.
- BBN will be replacing the MILNET mailbridges by T20s in about six months. the T20's will do IBGP with each other, and will support BGP2, BGP3 and EGP2.
- cisco did a full rewrite of their BGP code earlier this year. Version 9.02 has many patches from early testing in Alternet and NEARnet. Check release notes on ftp.cisco.com.
- The T3 backbone (ANS) will be running BGP3 in the fall, using a gated based implementation. (The current version is routed based). They plan to have BGP4 early '93.
- CA*NET is currently using BGP3 in gated. BGP4 will require significant changes to gated, but may be available in the fall. The current Cornell gated code includes the CA*NET BGP3 code.

There was discussion about what the mid-levels could do to help the vendors. Tony Li of cisco indicated that test installations would be useful. Their primary tool has been core dumps from sites which are having problems. It was noted that it is difficult to test BGP in a laboratory, because the interesting behaviors come from its ability to represent complex events in the global internet.

Jordan Becker volunteered to provide vendors with access to the ANS test facility for BGP interoperability testing. Contact him for details.

There was also discussion about adding a feature to the ANS gated to support BGP tunneled from the ANS test facility across the Internet. This would permit remote vendors to do initial interoperability testing in their own labs. The necessary code change is trivial but there are some operational and scheduling issues to be addressed. Contact Jordan Becker.

There was a short discussion of the forms being used for exchanging Routing Policy and configurations. These forms are an initial attempt at a mechanism to detect certain classes of global policy inconsistencies. RIPE commented that the forms should be network oriented, and not AS oriented.

Attendees

Nagaraj Arunkumar	nak@3com.com
Tony Bates	tony@ean-relay.ac.uk
Jordan Becker	becker@nis.ans.net
David Bolen	db3l@nis.ans.net
Henry Clark	henryc@oar.net
Michael Craren	mjc@proteon.com
John Curran	jcurran@bbn.com
Osmund de Souza	osmund.desouza@att.com
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Dennis Ferguson	dennis@mrbill.canet.ca
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Der-Hwa Gan	dhg@nsd.3com.com
Elise Gerich	epg@merit.edu
Eugene Hastings	hastings@a.psc.edu
Kathleen Huber	khuber@bbn.com
Steven Hubert	hubert@cac.washington.edu
Paulina Knibbe	knibbe@cisco.com
Mark Knopper	mak@merit.edu
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
Daniel Long	long@nic.near.net
Matt Mathis	mathis@a.psc.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
Brad Passwaters	bjp@sura.net
Robert Reschly	reschly@brl.mil
Erik Sherk	sherk@sura.net
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Marten Terpstra	terpstra@ripe.net
Carol Ward	cward@westnet.net
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov
Jane Wojcik	jwojcik@bbn.com
Paul Zawada	Zawada@ncsa.uiuc.edu

Overview of BGP Connections

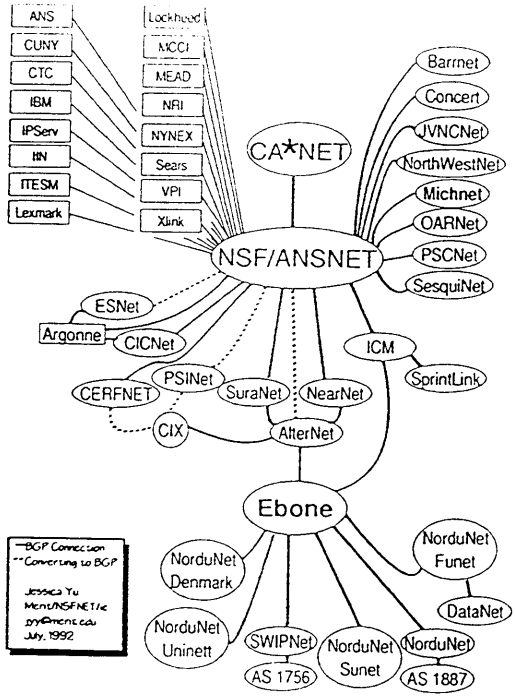

EBONE '92

- Defined fall '91 as a core network.
- Defined the RBS/EBS concept.
- Core routers in one AS
- USE BGP !!
- NO AUP on exchange points and shared resources.
- Initially building on more-or-less existing resources.
- Support "TCP/IP" and "CLNP-Pilot".
- Done without support of the X25 camps' money.

Ebone '92

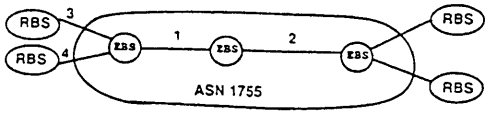

1,2 Networks internal to Ebone, IGRP

3,4 Networks belonging to Region, BGP —> Ebone

IBGP between EBS boxes carry all external Routes

IGP (Igrp) and BGP are not synchronized, thus the igp does not carry external Routes

Ebone '92

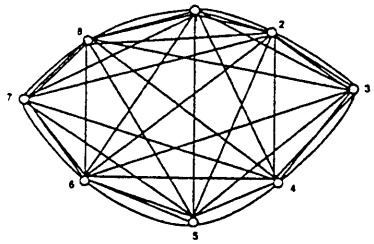

1. Stockholm 7 (STH)
2. Stockholm 2 (STH)
3. Stockholm 4 (AMS)
4. Amsterdam 1 (AMS)
5. Geneva 1 (GEN)
6. Mont Peller 1 (MNT)
7. London 1 (LON)
8. Stockholm 5 (LON)

- 15mps
- 256/512k
- IGRP 1755
- IBGP 1755

EBONE '92

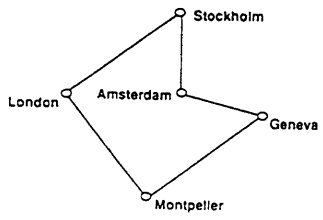

CORE NETWORK AND EBS SITES

Ebone '92

Problems:

- Teaching Regions
- IGRP --> BGP Transition
- Routing Registratory
- Multi-homed regions
- Transit networks doing packet filtering
- Playing with weights inside the IBGP can create loops.
- Everyone is not talking BGP.

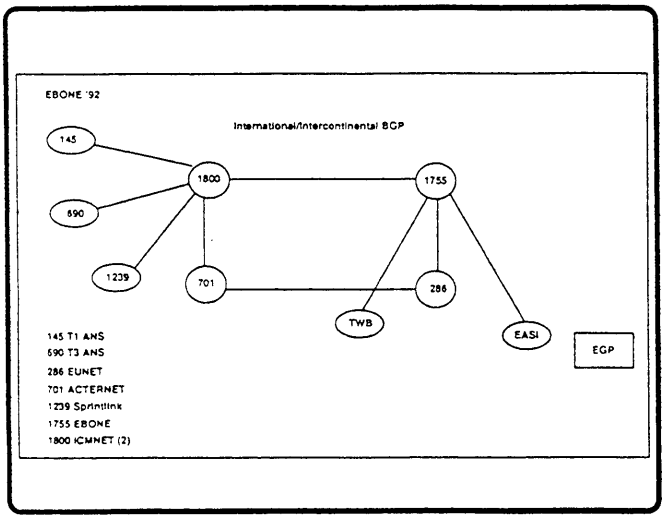
EBONE '92
International/Intercontinental BGP
145
1800
1755
690
1239
701
286
TWB
EASI
EGP
145 T1 ANS
690 T3 ANS
286 EUNET
701 ACTERNET
1239 Sprintlink
1755 EBONE
1800 ICMNET (2)

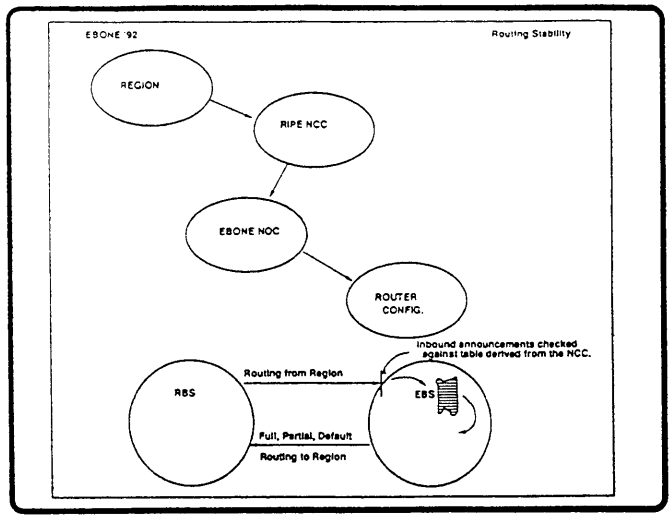
EBONE '92
Routing Stability
REGION
RIPE NCC
EBONE NOC
ROUTER CONFIG.
Inbound announcements checked against table derived from the NCC.
RBS
Routing from Region
EBS
Full, Partial, Default
Routing to Region

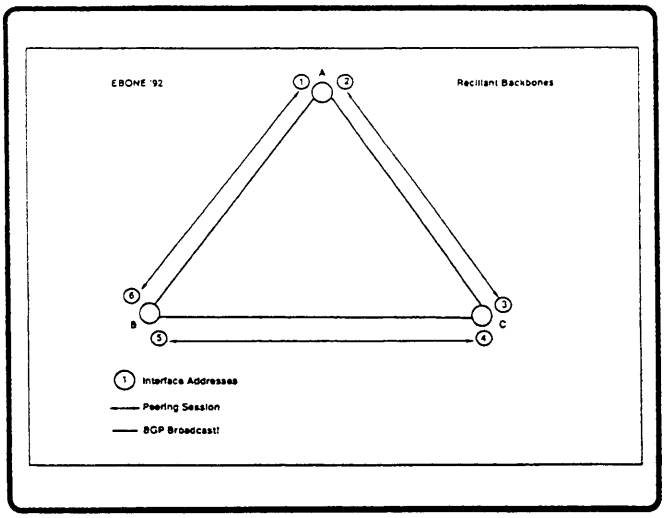
EBONE '92
Recillant Backbones
A
B
C
1 Interface Addresses
Peering Session
BGP Broadcast!

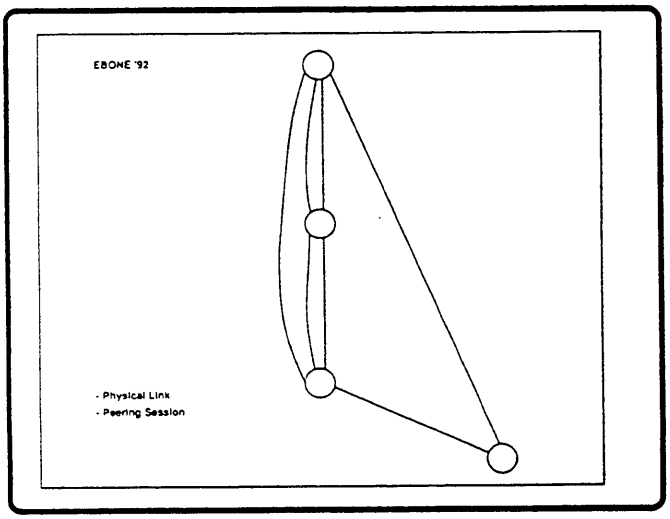
EBONE '92
- Physical Link
- Peering Session

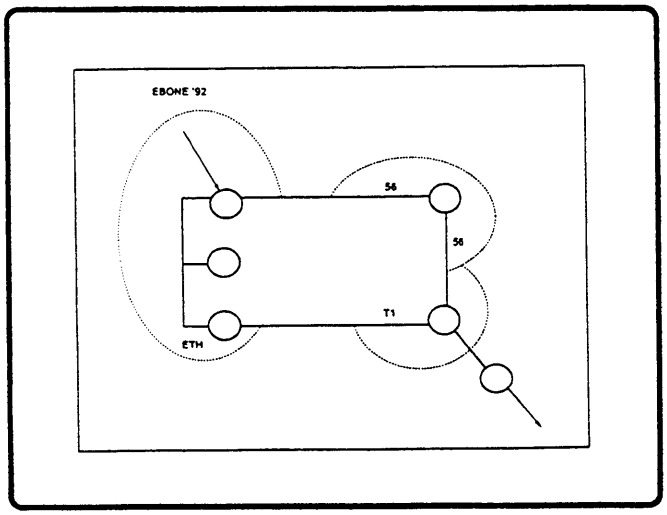
EBONE '92
56
56
T1
ETH

2.5.2 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 Benchmarking 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 concentr
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 re

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 Benchmarking Methodology Working Group met on Monday, July 13th.

Two proposed additional performance tests were discussed. The Group concluded that one of the proposals should not be adopted. The second proposal, to add an application level performance test, should be discussed with the test equipment vendors to see if they would be willing to implement such tests.

The latest draft was discussed and a few editorial changes were suggested.

It was strongly urged that some additions be made to the draft to add the rational behind the various tests and procedures.

Attendees

James Beers	beers@nr-tech.cit.cornell.edu
Scott Bradner	sob@harvard.edu
Randy Butler	rbutler@ncsa.uiuc.edu
James Carlson	carlson@xylogics.com
Constance Fleenor Lloyd	cfleenor@lloyd.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Bob Jeckell	rrj@3com.com
Alan Kullberg	akullber@bbn.com
Kim Long	klong@sura.net
Bill Manning	bmanning@rice.edu
Cindy Mazza	
Jim McQuaid	
Steve Mohr	srm@nso.3com.com
Pushpendra Mohta	pushp@cerf.net
Chris Munnelly	munnelly@xylogics.com
K. K. Ramakrishnan	rama@erlang.enet.dec.com
Martin Schulman	mas@loyola.edu
Chris Sullivan	mm@gandalf.ca
Bernhard Volz	volz@process.com
Gerard White	

2.5.3 Network Joint Management (njm)

Charter

Chair(s):

Gene Hastings, hastings@psc.edu

Mailing Lists:

General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu
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 Gene Hastings/PSC

Minutes of the Network Joint Management Working Group (NJM)

Agenda

- Follow-up on Past Actions
  - Matt Mathis' presentation on BGP is available via anonymous ftp from a.psc.edu as psc-bgp-utilization.ps
  - Many operators have adopted the canonical trouble mailboxes previously discussed in NJM (net-trouble@your.net, net- trouble@noc.your.net) but the practice is not yet ubiquitous. Part of the reason is a range of system organization.
  - If you haven't submitted an entry to Dan Long's online phonebook, DO IT!
- Operational Implications of NSF Recompete.
  - While it is not yet clear what the impact will be, it IS clear that the NSF backbone recompete has the potential for great change in the operations and practices of the Regional Network Operators. Operations and engineering personnel should be following the developments and considering ramifications.
  - Subscribe to the NSF recompete mailing list: recompete-request@nsf.gov. Related papers are available on expres.cise.nsf.gov.
  - Look at the papers on expres.cise.nsf.gov, in remcompete/ Aiken-Braun-Ford, impl.ps - Aiken presentation to FNCAC, Aiken_HPCC_NREN.ps - Draft solicitation.
- Operational Implications of Addressing Changes. Dan Jordt - small campus given N class C nets, but cannot (or will not) run 8 bit subnets. router use on same net? ARP hacks?
  - How does an operator handle N class C(s) when a site can not or will not run 8 bit subnets?
  - There are numerous tricks, such as ARP hacks/router interface hacks etc.
  - There was agreement of all assembled that there is a need to really explore ways to do this and collect experiences. Unfortunately, the IP Addressing BOF was scheduled in Parallel, so discussion was limited. A crossover attendee said that the BOF was focusing on a recent paper (Rekhter/Li), but there was little notion yet of what its consequences might be.

- OpStats - Are you doing it? What stands in your way? What help do you need? Money available for statistics software from FARNET/NSF.
  - Action Item - Ittai Hershman volunteered to speak to NYSERnet about releasing into the public domain the old SNMP software and relaxing license agreements.
  - Several folks complained that there is no statistics gathering package freely (or cheaply) available. The free SNMP libraries are only that, and complete packages, where they are affordable, tend to be highly platform dependent for displays or for other libraries.

- New Services - How to operate and debug services.
  - Do you operate any?
  - Do you support, debug them?
  - WAIS
  - ARCHIE
  - GOPHER
  - AFS
  - Do you need map of logical service topology?

  It is the case that network operators are called upon more and more to troubleshoot services above the transport layer, and a proliferation of new services to troubleshoot. (DNS, NFS, NTP, NNTP, BITNET II, WAIS, Gopher, Archie, etc.) There was agreement that maps of service topologies would be of great value in network operations, but there were no dramatic ideas on conventions or format.

- Scuttlebutt
  - Tom Easterday of CICnet announced that as of July 1 1992, ANS is running the CICnet NOC now.
  - Observations: What should we do with the T-1 network connectivity? Matt Mathis observed that the T1-T3 Interconnects seem to be stable and not critically loaded at present, but what will happen to them if sites peering with both T3 and T1 backbones stop listening to the T1 (this would send all of their T1 traffic through the interconnects). The decommissioning of the T1 entirely is contingent on the availability of CLNP in the T3 backbone, and fallback T1 links for the current T3 ENSSes.
  - 192.1.1 (a BBN Network) will be discontinued because some vendors of medical equipment ship using it as a default host network.
  - Some operators continue to experience random syslog messages from hosts that have lost a route to their loopback address. Should operators black hole traffic to 127 net to prevent wandering syslog messages.

- Caution: SURAnet has seen a CISCO applique ("335" style) that will pass even length packets but not odd length.

- There was agreement that [the attending] Network Service Providers are willing to take on the job of handing out address blocks if and when hierarchical schemes might be adopted. Some discussion followed as to whether operators should seek large block assignments and start handing them out now to get a head start on things.

Attendees

Henry Clark	henryc@oar.net
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Vince Fuller	vaf@stanford.edu
Eugene Hastings	hastings@a.psc.edu
Ittai Hershman	ittai@nis.ans.net
Wendy Huntoon	huntoon@a.psc.edu
J. Lance Jackson	jackson@noc.harvard.edu
Dan Jordt	danj@nwnet.net
John Labbe	labbe@merit.edu
Hock-Koon Lim	lim@po.cwru.edu
Kim Long	klong@sura.net
Matt Mathis	mathis@a.psc.edu
Kim Mayton	mayton@wg.com
Stephen Miller	smiller@bbn.com
Bill Norton	wbn@merit.edu
Kraig Owen	tko@merit.edu
Bob Page	bob.page@eng.sun.com
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Tim Seaver	tas@concert.net
Erik Sherk	sherk@sura.net
Carol Ward	cward@westnet.net
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

2.5.4 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, boss@ebone.net
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?

Internet Drafts:

"A Model for Common Operational Statistics", 03/24/1992, Bernhard Stockman <draft-ietf-opstat-model-00.txt>

CURRENT MEETING REPORT

Reported by Bernhard Stockman/SUNET

Minutes of the Operational Statistics Working Group (OPSTAT)

Agenda

- Administrative Items
- Review Internet Draft, is it ready to progress to RFC?
- Development of tools
- Review Charter, Goals, and Milestones

Review of Comments Received About Internet Draft

Several comments were made by various Working Group participants. Jon Boone (PSC) wants stop-time and filename in a header section so that it isn't necessary to scan the entire file to find the ending time of the data. Agreement was reached to change the time_section to label_ section as follows:

label_section ::= "BEGIN_LABEL" <FS>
<start_time> <FS>
<stop_time> <FS>
<data_file> <FS>
"END_LABEL" <FS>

There was a question about file setup... is there a need for one big file or lots of little files? Should there be multiple sections within a file? Matt (PSC) noted that important information was not mentioned and could be potentially confusing. They wanted to use very large files because of their tape storage facilities. Agreement was reached to add a sentence saying the specifics of how files are physically arranged is outside the scope of the document.

There was some confusion about the use of tags and variables in the poll-data section. Are there multiple tags for different sets of variables? It was noted that the draft is vague on time aggregation in this context and there was no clear way to do it. There is a need to provide a representation for, say, the average for 1 hour, 1 minute, and a maximum value, as well as, a need for classes of operators since aggregation is different for counters vs. an interface status variable. Agreement was reached to rewrite the section to make it more clear.

Would the addition of comments within the data files be useful? Yes, add comments, something like:

FS :== "," — <LF> — <LF> # text <LF>

Additional questions were posed by Dave S. (BNL).

- What is the use of the networkname field? Ross noted that network names were unique. Evan noted that the sharing of data among networks would be facilitated such that data wouldn't possibly become confused. Consensus was reached that the networkname field was useful :-).

- Question about the routername. This brought up a bigger discussion about how addresses are bound to each interface on a router. Usually the name of a router is tied to the interface most commonly used to access it. The name must be unique. Discussion digressed to involve yellow post-its and neon lights to name routers. Matt mentioned limitations of the DNS to name routers usefully. No action was taken.

- Questions about the linkname prompted discussion on what should be in it and what would be mandatory. The concept of a virtual link is needed to represent one or more physical links which can be grouped for statistical purposes. Should the name represent the ISO layer 2 or layer 3 name? Should there be an external name map to map the linkname field to a meaningful string? How should information be encoded in the field? Should there be resource vs. time aggregation? There are several unanswered question which could be answered in a later document covering implementation details.

Ed Reeder (IBM) suggested several editorial comments. All were approved as suggested.

In Section 5.1 there was a question about the difference between the raw data and the presented data. Agreement was reached to rewrite sentence/section to make more clear about the difference between the two.

In a separate discussion, Matt suggested adding a field to show specifically whether the data had been aggregated, versus having an implicit indication currently. Everyone agreed to this change.

There was some question about having a minimum value as well. Lengthy discussion about the minimum value always being very close to or equal to 0. Consensus had been reached at an earlier meeting to drop the minimum value.

James Barr (NIKHEH-H) asked a question about adding a comment. Since this had been agreed to previously, no further discussion was held.

Peter Fenwick (Univ. of Auckland) pointed out a syntax error in Section 6.1.3 in the data field specification where the number of "[" and "]" were unequal. Everyone agreed the error should be fixed.

Pietrak Rafal had extensive comments on the aggregation periods in the draft. Comments about the effect of extra weekend days in a month skewing data for a physical month, as well as, as questions about the period of time we wanted a peak value for raised. Everyone

agreed that two hours was too long of a period. Agreement was reached to leave the time values as they were in the Internet Draft.

Is the draft ready for RFC'ing?

Agreement was reached that the draft was getting close and needed the changes mentioned above. Once the changes are made, Bernhard will send a copy for review before forwarding it to the IESG.

Future Tool Development

Eric Hood, Executive Director of FARNET, made some comments about network statistics and the availability of financial assistance from FARNET to help fund some amount of development. Eric was going to have the FARNET staff survey networks to see what tools are currently available and what is under development. Everyone agreed that there needed to be reference implementations to flesh bugs out of the draft (and future drafts) and show the direction of future work. A Consensus was reached on the need for a common way to store data and to share common tools which are freely redistributable (in the public domain). We are currently unsure of the questions which will need to be asked at some point about what statistics are important and implementations will help answer the questions.

Eric will do the survey and forward the results to the OPSTAT Working Group. Discussions about tools will continue at the Washington, DC IETF, November 16-20, 1992.

Attendees

Vikas Aggarwal	aggarwal@nisc.jvnc.net
Tony Bates	tony@ean-relay.ac.uk
Henry Clark	henryc@oar.net
Robert Collet	rcollet@icm1.icp.net
Scott Cossette	sdc@concord.com
Osmund de Souza	osmund.desouza@att.com
Tom Easterday	tom@cic.net
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Cliff Frost	cliff@cmsa.berkeley.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Eugene Hastings	hastings@a.psc.edu
Eric Hood	ehood@nwnet.net
John Labbe	labbe@merit.edu
Walter Lazear	lazear@gateway.mitre.org
Hock-Koon Lim	lim@po.cwru.edu
Kim Long	klong@sura.net
Matt Mathis	mathis@a.psc.edu

John McKenna	mckenna@ralvm12.vnet.ibm.com
Chris Myers	chris@wugate.wustl.edu
Kraig Owen	tko@merit.edu
Sonya Reimer	sonya@brl.mil
Robert Reschly	reschly@brl.mil
Bernhard Stockman	boss@sunet.se
Ross Veach	rrv@uiuc.edu
Curtis Villamizar	curtis@ans.net
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov

2.5.5 User Connectivity (ucp)

Charter

Chair(s):

Dan Long, long@nic.near.net

Mailing Lists:

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

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 Dan Long/BBN

Minutes of the User Connectivity Working Group (UCP)

Dan Long reported that Tom Sandoski of ConcertNet has released both a curses and X-windows interface to his freeware Trouble Ticket System on ftp.concert.net in the pub/tickets directory.

There are now 50 NOCs listed in the Network Service Center Phonebook, maintained at NEARnet. Finger access is available for peer NOC's and a publicly-searchable subset of the data will be announced soon.

Most of the meeting was devoted to refining the description of information that Paul Zawada has compiled for handing off tickets from one NSC to another. Several organizations have expressed interest in standardizing these handoffs. Paul will revise his ASN.1 version and Kaj Tesink will draft a specification of how to encode this information in email messages.

Attendees

Henry Clark	henryc@oar.net
Tom Coppeto	tom@mit.edu
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Atul Garg	agarg@synoptics.com
Eugene Hastings	hastings@a.psc.edu
Alisa Hata	hata@cac.washington.edu
John Labbe	labbe@merit.edu
Walter Lazear	lazear@gateway.mitre.org
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Bill Manning	bmanning@rice.edu
Kraig Owen	tko@merit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Sonya Reimer	sonya@brl.mil
Tim Seaver	tas@concert.net
Erik Sherk	sherk@sura.net
Kaj Tesink	kaj@cc.bellcore.com
Ross Veach	rrv@uiuc.edu
John Vollbrecht	jrv@merit.edu
Carol Ward	cward@westnet.net
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu

Linda Winkler lwinkler@anl.gov
Jane Wojcik jwojcik@bbn.com
Paul Zawada Zawada@ncsa.uiuc.edu

2.6 Routing Area

Director(s):

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun

Border Gateway Protocol Working Group (BGP)

The BGP-4 Internet Draft was discussed. The Group received numerous comments which resulted in many changes to the document but did not change the resulting protocol. The editors will proceed to incorporate these changes and then post the result to the Working Group mailing list. If there are no further changes, the Working Group will present this to the IESG for Proposed Standard.

The BGP Communities' attributes was presented and discussed. This is an optional attribute which is useful for configuring BGP to help enforce Acceptable use Policies (AUP). A revised specification will be posted to the mailing list and then subsequently submitted to the IESG.

The BGP4/OSPF interaction document was discussed. A revised specification will be posted to the mailing list and the submitted to the IESG.

The BGP-4 usage document was discussed but closure was not reached. It is still unclear as to how much aggregation should happen by default with BGP-4. Discussion will proceed on the mailing list to resolve this.

The BGP-4 MIB was discussed. The changes required from the BGP3 MIB appear to be minimal.

IP Over Large Public Data Networks Working Group (IPLPDN)

The IPLPDN Working Group discussed Frame Relay <=> ATM inter-working and offered input to the IP over ATM Working Group.

They discussed "Shortcut Routing" document. A revised version of the Internet Draft will be produced.

The Group discussed the "Directed ARP" draft. After discussing the draft the Group agreed to submit the draft for Proposed Standard.

The Group discussed IP over Circuit ID proposal and agreed on a framework for an approach. An Internet Draft will be released for review.

Also discussed was PPP parameter negotiation over Large Public Data Networks. The

Group agreed on a general approach and will release a draft for review.

Inter-Domain Policy Routing Working Group (IDPR)

The Group discussed the status of the "gated" version of IDPR. The software is available form Robert Woodburn (woody@sparta.com) at SAIC. There are still things to modify and add, but the software in generally in good shape.

The Group also solicited help and received volunteer commitments for the following tasks:

- DNS Support for address/domain translation.
- MIB document update and implementation.
- Configuration Guide Update.
- Ongoing IDPR "gated" support.

In addition, the Group introduced and discussed several topics for IDPR enhancements. These include:

- Super domains.
- Multicast.
- Resource Allocation.
- Dynamic Source Policies for Hosts.

The Working Group also conducted a demo of the IDPR "gated" version. The Group also held a joint meeting with the NIMROD Group.

Mobile IP Working Group (MOBILEIP)

The Mobile IP Group met as a real Working Group for the first time. The first of two sessions consisted of presentations of three alternative proposals for handling mobile IP hosts by Charlie Perkins of IBM, John Ioannidis of Columbia University, and Fumio Teraoka of the Sony Computer Sciences Lab.

The second session reviewed and approved the Groups Charter. They added to the Charter a mention of the need to consider security and authentication issues related to IP mobility. Andrew Maffei gave a short talk on the problems of internetworking among ocean going vessels, buoys, submarines, and seabed devices.

Steve Deering identified a number of similarities and differences among the three proposed mobile host protocols and urged the Group to distinguish between fundamental difference vs. "engineering trade-off's" and to look for ways to combine the best features of each proposal.

Dave Kumpf described his experience with implementing the Columbia Mobile Host protocol and the Group reviewed and discussed a written critique of the protocol by Andrew Myles.

Multicast Extensions to OSPF Working Group (MOSPF)

The main topic discussed was interoperation of MOSPF with the Distance Vector Multicast Routing Protocol (DVMRP) which is being used to support the IETF audiocast.

John Moy reported on his experience with deploying MOSPF in some of the OARnet routers in preparation for the audiocast. He found a few unanticipated problems with mixing multicast capable and non-multicast capable routers in the same OSPF domain. This will be noted in the protocol specification. He also re-discovered the need for some "glue" protocol between MOSPF and DVMRP.

Steve Deering described DVMRP and the subset of DVMRP currently implemented in the "mrouted" routing demon for Unix. The two major shortcomings of the current implementation are:

1. The use of source-routing rather than encapsulation for multicast tunnels.
2. The lack of multicast tree pruning.

The Working Group will submit the MOSPF draft as a Proposed Standard after some minor editorial work by John Moy.

Open Shortest Path First IGP Working Group (OSPF)

The Group reviewed four documents that will soon be ready for publication before the next IETF. These are:

1. Update to the OSPF V2 Specification.
2. Update to the OSPF MIB.
3. OSPF Trap MIB.
4. OSPF Not So Stubby Area (NSSA) Option.

They also discussed two different ways of running OSPF over Frame Relay. Frame relay poses some special problems since it is a non-broadcast network that is generally not fully mush connected. Approaches to gracefully deal with OSPF database overflow was also discussed.

RIP Version II Working Group (RIPV2)

The Group made minor editorial changes to the RIP-2 Internet Draft. They also made one type change to the RIP-2 MIB Internet Draft. This is awaiting approval from the network management directorate.

The Group plans to submit the RIP-2 Internet Drafts for inclusion in the standards track.

ROUTING AREA REPORT

Robert Hinden

July 17, 1992

Sun Microsystems

July 17, 1992

BORDER GATEWAY PROTOCOL W.G.
Yakov Rekhter / IBM

- BGP-4 Internet Draft Discussed
  - After Changes Made will Submit for Proposed Standard
- BGP Communities Attributes Presented and Discussed
  - Intended to be Used to Help Enforce Acceptable Use Policies (AUP)
- BGP/OSPF Interaction Document Discussed and Revised
  - Will be Re-Submitted for Proposed Standard
- BGP-4 Usage Document Discussed
  - Major Outstanding Issue is How Much Aggregation Should Happen by Default
- BGP-4 MIB Discussed

Sun Microsystems

July 17, 1992

INTERDOMAIN POLICY ROUTING W.G.
Martha Steenstrup / BBN

- Discussed Status of IDPR *Gated* Implementation
  - Available from SAIC (woody@sparta.com)
- Discussed Ongoing Work
  - DNS Support for Address / Domain Translation
  - MIB Document Update
  - Configuration Guide Update
- Introduced and Discussed IDPR Enhancements
  - Super Domains
  - Multicast
  - Resource Allocation
  - Dynamic Source Policies for Hosts
- Held Joint Meeting with NIMROD Group

Sun Microsystems

July 17, 1992

IP OVER LARGE PUBLIC DATA NETWORKS W.G.
George Clapp / Ameritech

- Discussed Frame Relay - ATM Interworking and Presented Input to IP/ATM Working Group
- Discussed Short Cut Routing Draft
- Discussed Directed ARP Draft
  - Plan to Submit for Proposed Standard
- Discussed IP over Circuit ISDN
  - Agreed on Framework
- Discussed PPP Parameter Negotiation over LPDNs
  - Agreed on General Approach

Sun Microsystems

July 17, 1992

OPEN SHORTEST PATH FIRST IGP W.G.
John Moy / Proteon

- Reviewed Updates to Four Documents
  - OSPF V2 Specification
  - OSPF MIB
  - OSPF Trap MIB
  - OSPF Not So Stubby Area (NSSA) Option
- Discussed Two Approaches for Running OSPF over Frame Relay
  - Issues are Non-Broadcast and Not Fully Mesh Connected
- Discussed Proposal to Gracefully Deal with OSPF Database Overflow

Sun Microsystems

July 17, 1992

RIP-2 Work Group
Gary Malkin / Xylogics

- Made Minor Changes to RIP-2 Internet Draft
- Changed one Type in RIP-2 MIB
  - Discussing with NM Directorate
- Plan to Submit RIP-2 for Proposed Standard
- Completed RIP-2 Document On Schedule and MIB Ahead of Schedule
- Currently Two Independent Implementations of RIP-2

Sun Microsystems

MULTICAST OSPF WORKING GROUP
Steve Deering / Xerox PARC

- Presentation on Distance Vector Multicast Routing Protocol (DVMP)
  - Used to Support IETF Audiocast
  - Discussed Two Short Comings of DVMP
- Discussed Interoperation of MOSPF with DVMRP
  - Uncovered Problems with Mixing Multicast Capable and Non-Capable Routers in Same OSPF Domain
  - Need *Glue* between MOSPF and DVMRP
- Plan to Submit MOSPF for Proposed Standard

MOBILE IP WORKING GROUP
Steve Deering / Xerox PARC

- Presentations on Several Approaches to Mobile IP Hosts
  - Charlie Perkins / IBM
  - John Ioannidis / Columbia University
  - Fumio Teraoka / Sony Computer Science Lab
- Discussed Similarities and Differences between Approaches
- Discussed Implementation Experience with Columbia Scheme
- Reviewed and Approved Working Group Charter
  - Added Need to Consider Security and Authentication Issues
  - SAAG Member Appointed to Assist W.G.
- Talk by Andrew Maffei on Mobile Internetworking Ocean-Going Vessels, Buoys, Submarines, ...

CURRENT MEETING REPORT

Reported by Paul Tsuchiya/Bellcore

Minutes of A New Internet Protocol (PIP) BOF

The PIP BOF was held on Wednesday night, July 15, 1992. The purpose of the BOF was to educate people about PIP. PIP (for "P" Internet Protocol) is a new internet protocol designed for advanced internet applications, such as policy routing and flow, but is still able to accommodate existing internetworking techniques (that is, hop-by-hop routing with simple hierarchical addresses).

No real action items or decisions came out of the meeting. Mr. Tsuchiya indicated his intent to start a PIP mailing list (pip@thumper.bellcore.com), and to proceed as quickly as possible on a PIP specification and prototype implementation.

Otherwise, the BOF consisted of a presentation of examples of PIP operation under various circumstances including plain old hierarchical addresses, policy routing, separation of external addressing conventions from internal addressing, and mobile systems. Many lively discussions ensued, and a good time was had by all.

Attendees

J. Allard	jallard@microsoft.com
Nagaraj Arunkumar	nak@3com.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
Tony Bates	tony@ean-relay.ac.uk
Lou Berger	lberger@penril.com
David Bolen	db3l@nis.ans.net
Robert Braden	braden@isi.edu
Scott Brim	swb@cornell.edu
Alan Bryenton	bryenton@bnr.ca
Ross Callon	callon@bigfut.enet.dec.com
C. Allan Cargille	cargille@cs.wisc.edu
Stephen Casner	casner@isi.edu
Bilal Chinoy	bac@sdsc.edu
David Clark	ddc@lcs.mit.edu
Henry Clark	henryc@oar.net
James Cobban	jcobban@bnr.ca
Rob Coltun	rcoltun@ni.umd.edu
Tony DeSimone	tds@hoserve.att.com
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Vince Fuller	vaf@stanford.edu
Elise Gerich	epg@merit.edu

Rob Graham	robert_graham@protools.com
Kathleen Huber	khuber@bbn.com
Steven Hubert	hubert@cac.washington.edu
Bob Jeckell	rrj@3com.com
John Krawczyk	jkrawczy@wellfleet.com
Alan Kullberg	akullber@bbn.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Tony Li	tli@cisco.com
Gary Malkin	gmalkin@xylogics.com
Bill Manning	bmanning@rice.edu
Andy Nicholson	droid@cray.com
Lars Poulsen	lars@cmc.com
Stephanie Price Marasciullo	price@cmc.com
Tim Seaver	tas@concert.net
William Simpson	Bill_Simpson@um.cc.umich.edu
Frank Solensky	solensky@clearpoint.com
Karen Sollins	sollins@lcs.mit.edu
Marten Terpstra	terpstra@ripe.net
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Vollbrecht	jrv@merit.edu
Chris Wheeler	cwheeler@cac.washington.edu
Rick Wilder	wilder@ans.net
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com

CURRENT MEETING REPORT

Reported by Dave Crocker/TBO

Minutes of the IP Address Encapsulation BOF (IPAE)

The IP Address Encapsulation BOF was a 3-hour session on Thursday morning, July 9th. Approximately 100 people attended.

Bob Hinden gave a summary presentation of the goals and approach for solving the IP address-space limitation problem by use of an encapsulation technique which adds fields for new, globally unique addresses in a mini-layer above IP but below transport. A technical proposal was previously submitted as an Internet Draft (draft-crocker-ip-encaps-00.txt) and copies were provided. The handouts also included an initial assessment of implementation issues for BSD Unix, by Craig Partridge.

A key point in the questioning was the absence of a detailed proposal for the specific addressing scheme to be used. The current IPAE proposal constrains the last 4 octets of the larger, global addresses to be old-style IP addresses, but no other constraints or details were provided. While this was intentional, it was clear that the audience desired to have a concrete proposal. A separate effort to develop an addressing proposal has since begun as a sub-effort of the IPAE team.

The IPAE approach has as its main operational goal to limit impact on the installed base. It does this by defining the Internet to comprise a set of addressing "commonwealths" each with its own 32-bit address space. Within a commonwealth, classic IP always will be used. So, routers within the commonwealth and hosts which need to communicate only within that commonwealth will be able to conduct their business using old-style IP and require no change.

It was observed that this can turn segments of the Internet into isolated "islands". However, such islands are already part of the operational Internet, since some organizations operate that way for security purposes. Further, only hosts need to upgrade. It is intended that interior routers will not need to.

The BOF was intended to test the waters for forming a working group, as well as report on work-to-date. It was generally agreed that the working group should be formed. It intends to develop a complete specification for the core changes, a complete specification for an addressing scheme, and a specification of the procedures and mechanisms that will be needed to permit transition and coexistence for IP and IPAE.

Attendees

Nagaraj Arunkumar	nak@3com.com
William Babson	bill@penril.com
John Ballard	jballard@microsoft.com

Lou Berger	lberger@penril.com
Pravin Bhagwat	w2pravin@watson.ibm.com
David Borman	dab@cray.com
Robert Braden	braden@isi.edu
David Bridgham	dab@epilogue.com
Scott Brim	swb@cornell.edu
Geetha Brown	geetha@decvax.dec.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Greg Celmainis	gregc@newbridge.com
Vinton Cerf	vcerf@nri.reston.va.us
John Chang	jrc@uswest.com
A. Lyman Chapin	lyman@bbn.com
Chris Chiotasso	chris@artel.com
Richard Colella	colella@osi.ncsl.nist.gov
Michael Conn	0043874510@mcimail.com
David Crocker	dcrocker@mordor.stanford.edu
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Kurt Dobbins	dobbins@ctron.com
Irina Dobrina	irina@ftp.com
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Pierre Dupont	dupont@mdd.comm.mot.com
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Vince Fuller	vaf@stanford.edu
Maria Gallagher	maria@nsipo.nasa.gov
Atul Garg	agarg@synoptics.com
Robert Gilligan	Bob.Gilligan@eng.sun.com
Terry Gray	gray@cac.washington.edu
William Haggerty	haggerty@ctron.com
James Halpin	halpin@turkey.sw.stratus.com
Dimitry Haskin	dhaskin@bbn.com
Robert Hinden	hinden@eng.sun.com
Russ Hobby	rdhobby@ucdavis.edu
Tim Howes	Tim.Howes@umich.edu.
Kathleen Huber	khuber@bbn.com
Steven Hubert	hubert@cac.washington.edu
John Ioannidis	ji@cs.columbia.edu
Ronald Jacoby	rj@sgi.com
Dwight Jamieson	djamies@bnr.ca
Dan Jordt	danj@nwnet.net
Tom Kessler	kessler@eng.sun.com

Kenneth Key	key@cs.utk.edu
Stev Knowles	stev@ftp.com
John Krawczyk	jkrawczy@wellfleet.com
Alan Kullberg	akullber@bbn.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Anthony Lauck	lauck@tl.enet.dec.com
Fong-Ching Liaw	fong.liaw@eng.sun.com
Tony Li	tli@cisco.com
Carl Malamud	carl@malamud.com
Ed Menze	menze@cs.arizona.edu
Michael Millikin	mmillikin@interop.com
Lynn Monsanto	monsanto@sun.com
Dennis Morris	morrisd@imo-uvax.disa.mil
John Moy	jmoy@proteon.com
John Murray	murray@premenos.sf.ca.us
Erik Nordmark	nordmark@eng.sun.com
Charles Perkins	perk@watson.ibm.com
David Piscitello	dave@sabre.bellcore.com
Bala Rajagopalan	braja@qsun.att.com
Robert Reschly	reschly@brl.mil
April Richstein	abm@tycho.ncsc.mil
Simon Rosenthal	simon@bos.locus.com
Anthony Rutkowski	amr@sprint.com
Henry Sanders	henrysa@microsoft.com
John Sax	jsax@bbn.com
Jeremy Siegel	jzs@nsd.3com.com
Sam Sjogren	sjogren@tgv.com
Mark Smith	mcs@umich.edu
Frank Solensky	solensky@clearpoint.com
Michael St. Johns	stjohns@umd5.umd.edu
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Fumio Teraoka	tera@csl.sony.co.jp
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Walter Wimer	walter.wimer@andrew.cmu.edu
Jane Wojcik	jwojcik@bbn.com
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com

CURRENT MEETING REPORT

Minutes of the New Internet Routing and Addressing Architecture BOF (NIMROD)

Report not submitted.

Attendees

Steve Alexander	stevea@i88.isc.com
Karl Auerbach	karl@empirical.com
Robert Austein	sra@epilogue.com
Ashar Aziz	ashar.aziz@eng.sun.com
William Babson	bill@penril.com
Fred Baker	fbaker@acc.com
Lou Berger	lberger@penril.com
Andy Bierman	bierman@davidsys.com
David Bridgham	dab@epilogue.com
Steve Buchko	stevebu@newbridge.com
Randy Butler	rbutler@ncsa.uiuc.edu
John Chang	jrc@uswest.com
J. Noel Chiappa	jnc@lcs.mit.edu
Isaac Choi	ichoi@synoptics.com
Rob Coltun	rcoltun@ni.umd.edu
Tony DeSimone	tds@hoserve.att.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Chas DiFatta	chas@sei.cmu.edu
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Dennis Ferguson	dennis@mrbill.canet.ca
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Vince Fuller	vaf@stanford.edu
Robert Gilligan	Bob.Gilligan@eng.sun.com
Steven Hubert	hubert@cac.washington.edu
Christian Huitema	christian.huitema@sophia.inria.fr
John Ioannidis	ji@cs.columbia.edu
Dwight Jamieson	djamies@bnr.ca
Frank Kastenholz	kasten@ftp.com
Stev Knowles	stev@ftp.com
E. Paul Love	loveep@sdsc.edu
Carl Malamud	carl@malamud.com
Ed Menze	menze@cs.arizona.edu
Douglas Miller	dmm@telebit.com
John Moy	jmoy@proteon.com

Andy Nicholson	droid@cray.com
Erik Nordmark	nordmark@eng.sun.com
Stephanie Price Marasciullo	price@cmc.com
Bala Rajagopalan	braja@qsun.att.com
Richard Ramos	ramos@mtunm.att.com
Sam Roberts	sroberts@farallon.com
Manoel Rodrigues	manoel_rodrigues@att.com
Henry Sanders	henrysa@microsoft.com
Tim Seaver	tas@concert.net
Erik Sherk	sherk@sura.net
John Shriver	jas@proton.com
Frank Solensky	solensky@sceng.ub.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Fumio Teraoka	tera@csl.sony.co.jp
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
John Wagner	jwagner@princeton.edu
Justin Walker	justin@apple.com
James Watt	jamesw@newbridge.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com
Paul Zawada	Zawada@ncsa.uiuc.edu

2.6.1 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 var
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:

"Default Route Advertisement In The Border Gateway Protocol", 08/09/1991, Dimitry Haskin <draft-ietf-bgp-defaultroute-01.txt>

"Multicast Communications Using BGP", 08/26/1991, Scott Brim <draft-ietf-bgp-multicast-01.txt>

"A Border Gateway Protocol 4 (BGP-4)", 05/05/1992, Y. Rekhter, T. Li <draft-ietf-bgp-bgp4-02.txt>

"Definitions of Managed Objects for the Border Gateway Protocol (Version 4)", 09/01/1992, S. Willis, J. Burruss, J. Chu <draft-ietf-bgp-mibv4-00.txt>

"BGP4 OSPF Interaction", 09/15/1992, K. Varadhan <draft-ietf-bgp-bgp4ospf-interact-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)"

RFC 1364 "BGP OSPF Interaction"

CURRENT MEETING REPORT

Reported by David Bolen/ANS

Minutes of the Border Gateway Protocol Working Group (BGP)

During the first of two BGP Working Group sessions, the majority of the time was spent discussing two documents – the Internet Draft for BGP4 (Yakov Rekhter and Tony Li), and BGP4 <-> OSPF Interaction Document (Kannan Varadhan) – with a small portion of time devoted to discussing BGP Communities proposal (Yakov Rekhter and Tony Li).

BGP Communities Discussion

To start the meeting off, Tony Li presented the BGP Communities proposal (the use of a new path attribute to "color" a route), as previously posted to the BGP mailing list. The use of communities is intended to help solve the current AUP (acceptable use policy) routing problem by distributing some of the policy information (as kept in the NSFNET policy database) as a community associated with a route. The document predefines communities for research, education and commercial ASs. A community may be associated with a route by the source of that route, or may be added or augmented by any transit router (so a provider can "stamp" a route on behalf of a customer). While not a truly general solution (i.e., it does not help in cases where customers are using default routes), it may still prove beneficial in a large number of cases. The general consensus of the Group was that the proposal was worthwhile and would be useful to move forward.

BGP-4 Protocol Specifications Discussion

Next, the current BGP4 Internet Draft was discussed - driven primarily by comments from Michael Craren of Proteon, as he had questions about the document after having examined it in anticipation of implementing BGP for Proteon. The Working Group directly answered several of the simpler BGP implementation questions, while some points resulted in proposed changes to the draft, as follows:

- Section 2 Introduction
  - Revise to indicate that BGP4 no longer absolutely carries full AS path information (due to the possible use of the new ATOMIC_AGGREGATE attribute by intermediate routers).
  - Provide additional clarification as to what routes may be advertised by a BGP speaker (namely that it cannot advertise routes that it is not using).
  - Add a description of the FIB (forwarding information base).
- Section 3.2 (c) Routing Information Bases - Adj-RIBs-Out
  - Clarify that an outgoing policy (for the selection of routes to be advertised to a

neighbor) is applied only for "external" neighbors.

- Section 6.1 Message Header error handling
  - Remove the use of the "flash" qualifier to discuss update messages. Its use was thought to be a holdover from the early GATED implementation of BGP.

There were also a few simple grammatical changes pointed out. The BGP-4 document will be updated and released as a new version of the Internet Draft.

BGP-4 <-> OSPF Interaction Discussion

Kannan Varadhan then held a discussion of the updates necessary to his BGP<->OSPF interaction document to bring it in line with BGP4. For the most part the changes were to reflect the use of NLRI (Network Layer Reachability Information) within the BGP4 draft, since BGP4 carries IP prefixes rather than "class"-based network numbers.

One point brought up was that the document states that the OSPF router ID must be set to an interface address. OSPF does not require this, but the OSPF and BGP router IDs must be identical and BGP does set this requirement. It was agreed that the appropriate change to make was to update the BGP4 draft so that the router ID can be chosen as any address assigned to the router, but need not be associated with a physical interface. The BGP<->OSPF interaction document would then be updated to include the same restriction.

Kannan will also be releasing an updated version of his document.

The **second** BGP Working Group session was Chaired by Tony Li, and spent most of the time discussing the creation of a BGP4 usage document. The document is still to be done, but it was agreed that it would be very similar to the current BGP usage document, but extended to discuss BGP4 aggregation rules and requirements, and how to handle interactions with protocols that did not understand aggregated routes (such as EGP and older versions of BGP).

One issue that was left undecided (after a lengthy discussion) was what aggregation should be performed by a BGP4 implementation by default. There was no clear consensus on what option would be less likely to cause problems either for existing systems or for the site using BGP4 itself.

Some time was also spent on the BGP Communities proposal and on the BGP MIB document. The Group agreed that the BGP Communities document should proceed forward, probably with a release as an Internet Draft. The MIB document requires updating to include references to NLRI within BGP4's routes rather than networks as well as changes in the format of the AS_PATH attribute and creation of new path attributes. It was agreed to make the necessary changes.

Attendees

Nagaraj Arunkumar	nak@3com.com
Dennis Baker	dbaker@wellfleet.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
John Ballard	jballard@microsoft.com
Tony Bates	tony@ean-relay.ac.uk
Jordan Becker	becker@nis.ans.net
David Bolen	db31@nis.ans.net
Steve Buchko	stevebu@newbridge.com
Ross Callon	callon@bigfut.enet.dec.com
James Carlson	carlson@xylogics.com
William Carter	carter@ctron.com
Frank Chen	frankc@casc.com
Dean Cheng	dean@sun2.retix.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Jim Comen	comenj@interlan.interlan.com
Michael Craren	mjc@proteon.com
Steven Fancher	sfancher@ursa-major.spdcc.com
Dennis Ferguson	dennis@mrbill.canet.ca
Peter Ford	peter@lanl.gov
AnneMarie Freitas	afreitas@microcom.com
Vince Fuller	vaf@stanford.edu
Der-Hwa Gan	dhg@nsd.3com.com
Martin Gray	mg@spider.co.uk
Dimitry Haskin	dhaskin@bbn.com
Dwight Jamieson	djamies@bnr.ca
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
George Kajos	kajos@coral.com
Mark Knopper	mak@merit.edu
Mark Knutsen	knutsen@almaden.ibm.com
John Krawczyk	jkrawczy@wellfleet.com
Alan Kullberg	akullber@bbn.com
John Labbe	labbe@merit.edu
Tony Li	tli@cisco.com
Anthony Lisotta	lisotta@nas.nasa.gov
Robin Littlefield	rlittlef@wellfleet.com
Kent Malave	kent@chang.austin.ibm.com
Matt Mathis	mathis@a.psc.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
Kraig Owen	tko@merit.edu

Eric Peterson	elpeterson@eng.xyplex.com
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	abm@tycho.ncsc.mil
Martin Schulman	mas@loyola.edu
Hellen Sears	sears@interlan.interlan.com
Erik Sherk	sherk@sura.net
Marten Terpstra	terpstra@ripe.net
Linda Tom	toml@interlan.interlan.com
Kannan Varadhan	kannan@oar.net
Ross Veach	rrv@uiuc.edu
Curtis Villamizar	curtis@ans.net
John Vollbrecht	jrv@merit.edu
David Waitzman	djw@bbn.com
Luanne Waul	luanne@wwtc.timeplex.com
Honda Wu	natadm!honda@uunet.uu.net

BGP Communities

Purpose: Help solve the AUP problem.

Problem: ASs need to implement AUPs through policy filtering and weighting of routing updates. Currently, this is largely done by the NSFnet policy database or by long lists based on network number.

Solution: Distribute usage information automatically as part of BGP. Individual ASs can apply policy to implement their AUPs.

Usage

Each route may have a list of communities attached to it.

Policy filtering may be written based on the contents of the list.

Predefined communities: Research, Education, Commercial

Communities may be attached per set of destinations by the source AS.

Communities may also be attached by a transit AS.

A transit AS may not remove or modify attached communities.

Example (cont)

Example policy (at source):

For educational networks, always use the low cost educational transit. For research networks, prefer the high speed transit, but accept the low speed transit if the high speed is unavailable.

Configuration:

For the low speed neighbor: accept educational networks, accept research networks.

For the high speed neighbor: reject educational networks, accept research networks and prefer them.

Example

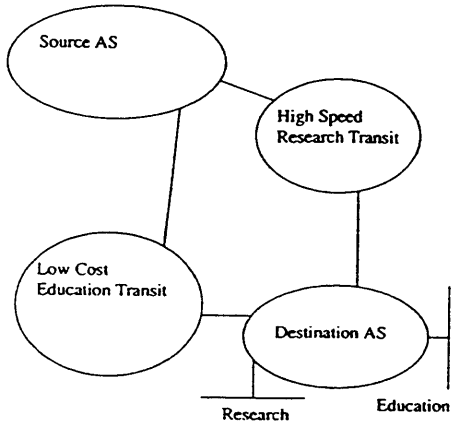

Example 2

Regional connected to NSFnet and to CIX.

Local policy: Prefer the NSFnet for educational and research networks. Prefer the CIX for commercial networks. Do not act as a transit between the two for commercial networks.

Configuration: Give high priority to educational and research networks learned from the NSFnet. Do not distribute commercial networks (including locally sourced networks) to the NSFnet.

Give high priority to commercial networks learned from the CIX. Distribute all routes to the CIX.

BGP-4 Aggregation Requirements

1. Implmenentations MUST be able to generate aggregates and other NLRIs (possibly also aggregates) that are part of that aggregate. (e.g., 192.0.0.0 255.0.0.0 and 192.1.0.0 255.255.0.0)
2. Implementations MAY allow the user to specify conditions which must be met before an aggregate is generated.
3. Deaggregation is discussed below.

BGP-4/EGP-2

A router may run BGP-4 and EGP-2 concurrently.

1. Routes from EGP-2 may be injected into BGP-4 with the ORIGIN path attribute set to 1 (EGP).

For routes from BGP-4 injected into EGP-2:

2. NLRI that denotes IP subnets requires the BGP speaker to inject the corresponding network into EGP2.
3. NLRI that denote IP network numbers may be directly injected into EGP2.
4. NLRI that denote aggregates may cause injection of
   - inject default only (0.0.0.0), no export of any other NLRI
   - allow controlled deaggregation, but only of specifically configured routes
5. A conformant implementation MUST allow export of non-aggregated NLRI.

BGP-4/BGP-3

A router may run BGP-4 and BGP-3 concurrently.

1. BGP-3 routes may be injected into BGP-4 directly. The AS_PATH becomes an AS_PATH of type AS_SEQUENCE.

For routes from BGP-4 injected into BGP-3:

2. NLRI that denotes IP subnets requires the BGP speaker to inject the corresponding network into BGP-3.
3. NLRI that denote IP network numbers may be directly injected into EGP2.
4. NLRI that denote aggregates may cause injection of
   - inject default only (0.0.0.0), no export of any other NLRI
   - allow controlled deaggregation, but only of specifically configured routes
5. A conformant implementation MUST allow export of non-aggregated NLRI.

In any of these three cases, the BGP-4 AS_PATH MUST be translated into a single AS_SEQUENCE and used as the BGP-3 AS_PATH.

6. A BGP speaker MUST NOT speak internal BGP-3 and internal BGP-4 simultaneously.

2.6.2 OSI IDRP for IP over IP (ipidrp)

Charter

Chair(s):

Sue Hares, skh@merit.edu

Mailing Lists:

General Discussion: idrp-for-ip@merit.edu
To Subscribe: idrp-for-ip-request@merit.edu
Archive: merit.edu: /pub/archive/idrp

Description of Working Group:

The IDRP for IP over IP Working Group is chartered to standardize and promote the use of IDRP (ISO Inter-Domain Routing Protocol) as a scalable inter-autonomous system routing protocol capable of supporting Policy Based Routing for TCP/IP internets. The objective is to take IDRP, as it is defined by ISO standards, and to define backward compatible extensions and/or network adaptation layers to enable this protocol to be used in the TCP/IP internets. If any ISO standardization efforts overlap this area of work, it is intended that the ISO work will supersede the standards proposed by this Group.

1) IDRP for IP over IP document (standards track)

This document contains the appropriate adaptations of the IDRP protocol definition that enables it to be used as a protocol for exchange of "inter-autonomous system information" among routers to support forwarding of IP packets across multiple autonomous systems.

2) IDRP MIB document (standards track)

This document contains the MIB Definitions for IDRP. These MIB Definitions are done in two parts; IDRP General MIB, and IDRP for IP MIB. An appendix is planned; IDRP For IP GDMO

3) IDRP - OSPF Interactions (standards track)

This document will specify the interactions between IDRP and OSPF. This document will be based on a combination of BGP-OSPF interactions document and IDRP - ISIS interaction document.

4) IDRP for IP Usage document (standards track)

Most of the IDRP for IP Usage will reference the CIDR (Supernetting document) Internet Draft. Any additional terms or protocol definitions needed for IDRP for IP will also be specified here.

Goals and Milestones:

Done	IDRP for IP submitted for Internet Draft.

Jun 1992 IDRP MIB document submitted for Internet Draft.

Jun 1992 IDRP - OSPF Interactions document submitted for Internet Draft.

Jun 1992 IDRP Usage document submitted for Internet Draft.

Nov 1992 IDRP for IP submitted to the IESG for Proposed Standard.

Nov 1992 IDRP Usage document submitted to the IESG for Proposed Standard.

Nov 1992 IDPR MIB Submitted to the IESG for Proposed Standard.

Nov 1992 IDRP - OSPF Interactions document submitted to the IESG for Proposed Standard.

CURRENT MEETING REPORT

Reported by Sue Hares/Merit and Bill Manning/SESQUINET

Minutes of the OSI IDRP for IP over IP Working Group (IPIDRP)

Agenda

- Discussion of the Working Group Charter.
- Review of IDRP for IP Document.
- ISIS-IDRP Interactions Paper.
- Next Steps.

Susan Hares (Chair) led the discussions and gave the presentations.

Working Group Charter

A great deal of discussion occurred about the Working Group Charter. The focus was what documents should be developed. The documents to be developed are:

- IDRP for IP: Merit will implement a version by Nov'92.
- MIB: Sue will derive from BGPv3. Not further discussed here.
- Usage Document. A great deal of discussion was given to what should go in the Usage document. In the usage document should go:
  - History of the protocol.
  - Typical scenarios for Usage IDRP in Internet.
  - Sample Policy description language.
  - How IDRP interacts with CIDR and C#.

A great deal of discussion was engaged in about whether interactions to other protocols belong in a usage document. It was decided the interactions should go in separate documents. Interactions that will need to be discussed are: IDRP <-> ISIS, IDRP <->EGP, IDRP <-> BGP2/BGP3/BGP4.

These documents will be new documents added to the Working Group Charter. The family tree document will tie these documents together.

Review of IDRP for IP Document

IDRP for IP can not assume correct routing if the addresses used are NOT globally unique. Mentioned because some proposal to fix the IP addressing problems do not guarantee this. This basic assumption will be placed in a section in the beginning of the document.

A great deal of discussion centered around a QOS flag for RE, and asynchronous routes. Many people thought asynchronous routes provide problems for users, but some people

challenged this discussion.

ISIS-IDRP Paper Overview

One goal of the ISIS-IDRP interactions is to allow a BIS to learn things from the IGP IS-IS. A tutorial was given about this. The paper can be obtained from merit.edu in /pub/iso/isisidrp.ps. This paper was presented to the SC6 committee during their July meeting which ran the same time as IETF. BIS can learn network layer reachability information from IS-IS in three methods:

- No Summarization:
  - Skipped over the MIB ways
  - Mash into the NRLI - use mask
- Automatic:
  - Local RDI template
- Pre-configured

Given some IS-IS changes, BIS can discover each other via IS-IS. IDRP information can either be tunneled through a domain or injected into the IGP for transmittal across the domain. A discussion occurred on how you choose when to tunnel vs. inject into the IGP. Tunneling can be based on overload bit in IS-IS. Dennis Ferguson (CA*NET) suggested we should ask for this type of feature in IP. We need a general specification on tunneling, otherwise there will be problems.

Tunneling brings up the issue of encapsulating CLNP in IP or IP in IP. Vint Cerf suggested that we look at earlier work in multiple encapsulations. Vint suggested Paul Tsuchiya's paper on encapsulation loops.

Action Items

- Discussion of ASYNC routes on this list.
- New Charter published by Chair.
- Document delivered in August.
- EGP-IDRP interactions document helpers: Scott Brim, Sue Hares
- BGP-IDRP interactions document helpers: Scott Brim, Sue Hares
- OSPF-IDRP interactions document helpers: Suggested to catch Kannan who wrote BGP-OSPF document.

Attendees

Nagaraj Arunkumar	nak@3com.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk

William Barns	barns@gateway.mitre.org
Tony Bates	tony@ean-relay.ac.uk
David Bolen	db3l@nis.ans.net
Scott Brim	swb@cornell.edu
Ross Callon	callon@bigfut.enet.dec.com
Vinton Cerf	vcerf@nri.reston.va.us
John Chang	jrc@uswest.com
Richard Colella	colella@osi.ncsl.nist.gov
Michael Conn	0043874510@mcimail.com
Michael Craren	mjc@proteon.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Pierre Dupont	dupont@mdd.comm.mot.com
Steven Fancher	sfancher@ursa-major.spdcc.com
Dennis Ferguson	dennis@mrbill.canet.ca
Vince Fuller	vaf@stanford.edu
Susan Hares	skh@merit.edu
Dimitry Haskin	dhaskin@bbn.com
Matthew Jonson	jonson@server.af.mil
James Keller	j.keller@sprint.com
Paulina Knibbe	knibbe@cisco.com
Mark Knopper	mak@merit.edu
Mark Knutsen	knutsen@almaden.ibm.com
Walter Lazear	lazear@gateway.mitre.org
Whay Lee	whay@merlin.dev.cdx.mot.com
Tony Li	tli@cisco.com
Kent Malave	kent@chang.austin.ibm.com
Bill Manning	bmanning@rice.edu
Douglas Miller	dmm@telebit.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Kraig Owen	tko@merit.edu
James Pearl	jim@microcom.com
Roy Perry	rperry@uswat.uswest.com
David Piscitello	dave@sabre.bellcore.com
Bala Rajagopalan	braja@qsun.att.com
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	abm@tycho.ncsc.mil
Erik Sherk	sherk@sura.net
Rajesh Srivastava	rajesh@flash.tandem.com
Sally Tarquinio	sallyt@gateway.mitre.org
Fumio Teraoka	tera@csl.sony.co.jp
Marten Terpstra	terpstra@ripe.net
Kannan Varadhan	kannan@oar.net
Cathy Wittbrodt	cjw@nersc.gov

2.6.3 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 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).
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:

"Shortcut Routing: Discovery and Routing over Large Public Data Networks", 06/05/1992, P. Tsuchiya, J. Lawrence <draft-ietf-iplpdn-shortcutrouting-01.txt>

"Directed ARP", 06/19/1992, John Garrett, John Hagan, Jeff Wong

<draft-ietf-iplpdn-directed_arp-00.txt>

Request For Comments:

RFC 1293 "Inverse Address Resolution Protocol"

RFC 1294 "Multiprotocol Interconnect over Frame Relay"

RFC 1315 "Management Information Base for Frame Relay DTEs"

RFC 1356 "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode"

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

Minutes of the IP over Large Public Data Networks Working Group (IPLPDN)

The IP over Large Public Data Networks (IPLPDN) Working Group met for seven sessions.

Andy Malis reported that the IAB had some questions with the Multiprotocol over X.25 document. Fortunately, these questions were easily answered, and the document should be released as an RFC soon. (It has since been released as RFC 1356.)

The Group agreed that there was sufficient proof of interoperable implementations to ask the IAB to advance the RFCs dealing with Frame Relay to the next step in the standards process. These RFCs are the following:

1315	Brown, C.; Baker, F.; Carvalho, C. Management Information Base for Frame Relay DTEs. 1992 April; 19 p.
1294	Bradley, T.; Brown, C.; A. Malis Multiprotocol Interconnect over Frame Relay. 1992 January; 28 p.
1293	Brown, C. Inverse Address Resolution Protocol. 1992 January; 6 p.

The issue concerning Frame Relay and ATM interworking arose from discussion within the IP over ATM Working Group. The IPLPDN Working Group wished to reach consensus over a solution, and they agreed that the preferred encapsulation and protocol identification scheme was that specified in RFC 1294 and in CCITT document I.355. I.355 specifies the encapsulation of Frame Relay over ATM, and RFC 1294 specifies multiprotocol encapsulation over Frame Relay. The Group agreed to take this proposal to the IP over ATM Working Group.

Tuesday afternoon was spent on Directed ARP and Shortcut Routing. Following a presentation on Directed ARP, the Group agreed to release the document to the IESG and IAB for approval as an RFC on the standards track. Release of Shortcut Routing was deferred as the authors wished to do more work and to gain implementation experience.

Multiprotocol over Circuit ISDN was discussed on Wednesday afternoon and Thursday morning. The anticipated environment is one in which the signaling protocol cannot be relied upon to inform the "called" party of the type of the encapsulation protocol, and in which Point-to-Point Protocol (PPP) implementations are expected to be common. The Group agreed that a Frame Relay encapsulation as specified in RFC 1294 would be the default encapsulation and outlined a mechanism to differentiate between Frame Relay and PPP implementations. Keith Sklower, Chris Ranch, and Cliff Frost volunteered to write the Internet Draft. They also volunteered to draft a document in which the PPP parameter

negotiation protocol is investigated as a means for parameter negotiation between Frame Relay devices.

Issues for the Next Meeting

- LAPF (Frame Relay) <-> ATM Interworking
- Directed ARP
- Shortcut Routing
- Multiprotocol over Circuit ISDN
- Point-to-Point Parameter Negotiation over Large Public Data Networks

The Group felt that a total of four sessions over two afternoons would be adequate to address these topics at the next IETF meeting.

Attendees

Sam Ayers	s.ayers@sybus.com
Douglas Bagnall	deb@proteon.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
William Barns	barns@gateway.mitre.org
Ken Benstead	kbenstead@coral.com
Caralyn Brown	cbrown@wellfleet.com
Alan Bryenton	bryenton@bnr.ca
Steve Buchko	stevebu@newbridge.com
John Burnett	jlb@adaptive.com
David Carr	dcarr@gandalf.ca
Greg Celmainis	gregc@newbridge.com
Dilip Chatwani	dilip@synoptics.com
Dean Cheng	dean@sun2.retix.com
George Clapp	clapp@ameris.center.il.ameritech.com
James Cobban	jcobban@bnr.ca
Tracy Cox	tacox@sabre.bellcore.com
Andy Davis	andy@spider.co.uk
Hans Dittler	dittler@conware.de
Cliff Frost	cliff@cmsa.berkeley.edu
Kathryn Fuller	kathryn@codex.com
John Garrett	jwg@edsel.att.com
David Getchell	getchell@cisco.com
Scott Ginsburg	ginsburg@wellfleet.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.enet.dec.com
John Hagan	hagan@dccs.upenn.edu
James Harville	jbh@unh.edu

Frank Heath	heath@cmc.com
Ian Heavens	ian@spider.co.uk
Juha Heinanen	juha.heinanen@datanet.tele.fi
Takashi Ikemoto	tikemoto@xerox.com
Dwight Jamieson	djamies@bnr.ca
James Keller	j.keller@sprint.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Joseph Lawrence	jcl@sabre.bellcore.com
Whay Lee	whay@merlin.dev.cdx.mot.com
Fong-Ching Liaw	fong.liaw@eng.sun.com
Arthur Lin	yalin@srv.pacbell.com
Andrew Malis	malis@bbn.com
Douglas Miller	dmm@telebit.com
Jack Ng	jkn@network.com
Luc Nguyen	luc@bellsouth.com
James Pearl	jim@microcom.com
Ed Perreault	ed@python.eng.microcom.com
Eric Peterson	elpeterson@eng.xyplex.com
Venkat Prasad	vsp@3com.com
Russell Pretty	pretty@bnr.ca
David Robinson	drb@relay.prime.com
Benny Rodrig	4373580@mcimail.com
Miguel Sasson	sasson@xylogics.com
Lansing Sloan	ljsloan@llnl.gov
Andrew Smith	asmith@synoptics.com
Rajesh Srivastava	rajesh@flash.tandem.com
Chris Sullivan	mm@gandalf.ca
Terrance Sullivan	terrys@newbridge.com
Sally Tarquinio	sallyt@gateway.mitre.org
Kaj Tesink	kaj@cc.bellcore.com
Sylvia Tom	tom_sylvia@tandem.com
Robert Ullmann	ullmann@process.com
Scott Wasson	sgwasson@eng.xyplex.com
James Watt	jamesw@newbridge.com
Liming Wei	liming.wei@eng.sun.com
Jehu Westmark	westmark@ds1.scri.fsu.edu
Jeff Wong	jaw@io.att.com

2.6.4 ISIS for IP Internets (isis)

Charter

Chair(s):

Ross Callon, callon@bigfut.lkg.dec.com

Mailing Lists:

General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu
Archive:

Description of Working Group:

The IETF ISIS 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 d
Done	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"

2.6.5 Inter-Domain Policy Routing (idpr)

Charter

Chair(s):

Martha Steenstrup, msteenst@bbn.com

Mailing Lists:

General Discussion: idpr-wg@bbn.com
To Subscribe: idpr-wg-request@bbn.com
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-05.txt, .ps>

"Inter-Domain Policy Routing Protocol Specification: Version 1", 03/05/1991, M. Steenstrup <draft-ietf-idpr-specv1-02.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-01.txt>

"Inter-Domain Policy Routing Configuration and Usage", 07/25/1991, H. Brown, M. Steenstrup <draft-ietf-idpr-configuration-00.txt>

"IDPR as a Proposed Standard", 04/28/1992, M. Steenstrup <draft-ietf-idpr-summary-00.txt, .ps>

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 in two sessions during the July 1992 IETF meeting in Boston. In the first session we talked about shorter-term as well as longer-term work on IDPR, and in the second session we offered a spur-of-the-moment demo and shared the remainder of the session with the NIMROD Group.

Shorter-term Topics

- Work on the Gate-d Version of IDPR. Currently, Woody Woodburn has a version of IDPR that runs as part of gated. Woody provided a detailed description of the status of the gated implementation at the first session and at the second session gave a demo of the software. To obtain a copy of the IDPR software, please contact woody@sparta.com. We plan a pilot demonstration of this software in the Internet in late summer or early fall. Moreover, we expect that, as a result of this experimentation, we will want to make changes to the software and perhaps to the protocols as well. The IDPR Working Group needs a set of people that are able and interested in working on enhancing the IDPR gated software.

- MIB Development. We need to implement the MIB, and we need to update the Internet Drafts describing both the IDPR MIB and the IDPR configuration and usage guide.

- Adding Facilities to the DNS. The DNS should be able to return domain information in response to a query giving entity name or address.

- Adding the Capability for Hosts to Request Source Policies Dynamically. In the current version of IDPR, source policies are specified as part of path agent configuration and remain active for a given host until they are reconfigured. Thus, hosts needn't do anything special to communicate this information to the path agents. We made this choice so that no host changes were necessary to reap the benefits of IDPR. However, we expect that in the future, hosts may want to have different source policies, depending upon the application active at the moment, and hence need a way to communicate this information to the path agents. SRI is working on a subset of this problem, but we may require a more general solution.

At the Working Group sessions, several of you volunteered to work on the shorter-term topics, pending approval from employers. Would those of you who have obtained such approval and who are still interested in working on these topics, please send mail to me at msteenst@bbn.com so that we can move these efforts along.

Longer-term Topics

- Multicast Support. IDPR should have the ability to construct multicast trees and establish the appropriate paths associated with these trees. Moreover, IDPR should be able to take advantage of intra-domain multicast where available, both for IDPR control information distribution and for multicast user applications. We are pursuing solutions targeted for IDPR. However, those of you interested in inter-domain multicasting in general may also wish to join the mailing list set up by Jon Crowcroft and Tony Ballardie of UCL. To subscribe to the list, send email to idmr-request@cs.ucl.ac.uk.

- Multipath Support. IDPR should have the ability to maintain as a unit multiple paths of the same service between the same source and destination. Multiple paths provide a vehicle for obtaining sufficient bandwidth, for load-balancing, and for providing backups when a primary path fails.

- Resource Allocation. IDPR should be able to interoperate with resource reservation and flow control mechanisms to provide service guarantees.

- Super Domains. As the Internet grows, one may wish to aggregate domains into super domains in order to reduce the amount of information and computation related to routing. Aggregation produces a hierarchy of domains. IDPR should provide a domain address format that permits addressing of domains at arbitrary positions in a domain hierarchy.

Any comments on any of these topics are most welcome on the IDPR mailing list, idpr-wg@bbn.com.

Attendees

Nagaraj Arunkumar	nak@3com.com
Robert Austein	sra@epilogue.com
William Babson	bill@penril.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
Tony Bates	tony@ean-relay.ac.uk
James Beers	beers@nr-tech.cit.cornell.edu
Lou Berger	lberger@penril.com
David Bridgham	dab@epilogue.com
Alan Bryenton	bryenton@bnr.ca
John Chang	jrc@uswest.com
Robert Ching	natadm!rching@uunet.uu.net
Bilal Chinoy	bac@sdsc.edu
Michael Craren	mjc@proteon.com
Osmund de Souza	osmund.desouza@att.com
Irina Dobrina	irina@ftp.com

Tom Farinelli	tcf@tyco.ncsc.mil
Elise Gerich	epg@merit.edu
William Haggerty	haggerty@ctron.com
Dimitry Haskin	dhaskin@bbn.com
Steven Hubert	hubert@cac.washington.edu
Takashi Ikemoto	tikemoto@xerox.com
Bob Jeckell	rrj@3com.com
Matthew Jonson	jonson@server.af.mil
Michael Khalandovsky	mlk@ftp.com
John Krawczyk	jkrawczy@wellfleet.com
Whay Lee	whay@merlin.dev.cdx.mot.com
Tony Li	tli@cisco.com
Anthony Lisotta	lisotta@nas.nasa.gov
Ed Menze	menze@cs.arizona.edu
Shehzad Merchant	merchant@erg.sri.com
Donald Merritt	don@brl.mil
Kraig Owen	tko@merit.edu
Lars Poulsen	lars@cmc.com
Stephanie Price Marasciullo	price@cmc.com
Manoel Rodrigues	manoel_rodrigues@att.com
Karen Seo	kseo@bbn.com
Frank Solensky	solensky@clearpoint.com
Martha Steenstrup	msteenst@bbn.com
Fumio Teraoka	tera@csl.sony.co.jp
Huyen Vu	vi@polaris.disa.mil
David Wang	wang@xylogics.com
Luanne Waul	luanne@wwtc.timeplex.com
Steven Winnett	swinnett@bbn.com
Robert Woodburn	woody@cseic.saic.com

2.6.6 Mobile IP Working Group (mobileip)

Charter

Chair(s):

Steve Deering, deering@parc.xerox.com

Mailing Lists:

General Discussion: mobile-ip@parc.xerox.com
To Subscribe: mobile-ip-request@parc.xerox.com
Archive: pub/mobile-ip/mail-archive@parcftp.xerox.com

Description of Working Group:

The Mobile IP Working Group is chartered to develop or adopt architectures and protocols to support mobility within the Internet. In the near term, protocols for supporting transparent host "roaming" among different subnetworks and different media (e.g., LANs, dial-up links, and wireless communication channels) shall be developed and entered into the Internet Standards track. The work is expected to consist mainly of new and/or revised protocols at the (inter)network layer, but may also include proposed modifications to higher-layer protocols (e.g., transport or directory). However, it shall be a requirement that the proposed solutions allow mobile hosts to interoperate with existing Internet systems.

Longer term, the Group may address, to the extent not covered by the mobile host solutions, other types of internet mobility, such as mobile subnets (e.g., a local network within a vehicle), or mobile clusters of subnets (e.g., a collection of hosts, routers, and subnets within a large vehicle, like a ship or spacecraft, or a collection of wireless, mobile routers that provide a dynamically changing internet topology).

Goals and Milestones:

Done	Review and approve the Charter, making any changes deemed necessary.
Nov 1992	Post an Internet Draft documenting the Mobile Hosts protocol.
Mar 1993	Submit the Mobile Host Protocol to the IESG as a Proposed Standard.
Mar 1993	Review the Charter of the Mobile IP Working Group for additional work required to facilitate non-host mobility.

CURRENT MEETING REPORT

Minutes of the Mobile IP Working Group (MOBILEIP)

Report not submitted. Please refer to Area Report for a brief summary.

Attendees

George Abe	abe@infonet.com
Philip Almquist	almquist@jessica.stanford.edu
Derek Atkins	warlord@thumper.bellcore.com
Robert Austein	sra@epilogue.com
Sam Ayers	s.ayers@sybus.com
Ashar Aziz	ashar.aziz@eng.sun.com
William Babson	bill@penril.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
John Ballard	jballard@microsoft.com
Lou Berger	lberger@penril.com
Pravin Bhagwat	w2pravin@watson.ibm.com
Uri Blumenthal	uri@watson.ibm.com
Luc Boulianne	lucb@cs.mcgill.ca
Alan Bryenton	bryenton@bnr.ca
William Carter	carter@ctron.com
John Chang	jrc@uswest.com
Dean Cheng	dean@sun2.retix.com
Robert Ching	natadm!rching@uunet.uu.net
Alan Clegg	abc@concert.net
Robert Clements	clements@bbn.com
James Cobban	jcobban@bnr.ca
Tom Coppeto	tom@mit.edu
Michael DeAddio	deaddio@thumper.bellcore.com
Steve Deering	deering@parc.xerox.com
Tony DeSimone	tds@hoserve.att.com
Chas DiFatta	chas@sei.cmu.edu
Irina Dobrina	irina@ftp.com
Daniel Duchamp	duchamp@liberty.cs.columbia.edu
Pierre Dupont	dupont@mdd.comm.mot.com
Hans Eriksson	hans@sics.se
Dennis Ferguson	dennis@mrbill.canet.ca
Ron Frederick	frederick@parc.xerox.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Elise Gerich	epg@merit.edu
Robert Gilligan	Bob.Gilligan@eng.sun.com
Martin Gray	mg@spider.co.uk

Chris Gunner	gunner@dsmail.enet.dec.com
William Haggerty	haggerty@ctron.com
Robert Hinden	hinden@eng.sun.com
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Steven Hubert	hubert@cac.washington.edu
John Ioannidis	ji@cs.columbia.edu
Dwight Jamieson	djamies@bnr.ca
Tom Kessler	kessler@eng.sun.com
Mark Knutsen	knutsen@almaden.ibm.com
David Kumpf	kumpf@hprnd.rose.hp.com
Mark Lewis	mlewis@telebit.com
Arthur Lin	yalin@srv.pacbell.com
Anthony Lisotta	lisotta@nas.nasa.gov
Brian Lloyd	brian@lloyd.com
Andrew Maffei	arm@aqua.whoi.edu
Glenn McGregor	ghm@merit.edu
Ed Menze	menze@cs.arizona.edu
Paul Milazzo	milazzo@bbn.com
Douglas Miller	dmm@telebit.com
John Moy	jmoy@proteon.com
Erik Nordmark	nordmark@eng.sun.com
Philippe Park	ppark@bbn.com
Michael Patton	map@lcs.mit.edu
James Pearl	jim@microcom.com
Charles Perkins	perk@watson.ibm.com
Stephanie Price Marasciullo	price@cmc.com
Bala Rajagopalan	braja@qsun.att.com
James Rees	jim.rees@umich.edu
Jon Rochlis	jon@mit.edu
Simon Rosenthal	simon@bos.locus.com
Allan Rubens	acr@merit.edu
Jonathan Sandberg	Jss@mitl.com
Henry Sanders	henrysa@microsoft.com
Richard Schmalgemeier	rgs@merit.edu
Eve Schooler	schooler@isi.edu
Karen Seo	kseo@bbn.com
Koichiro Seto	seto@hitachi-cable.co.jp
William Simpson	Bill_Simpson@um.cc.umich.edu
Henry Sinnreich	hsinnreich@mcimail.com
Rajesh Srivastava	rajesh@flash.tandem.com
Martha Steenstrup	msteenst@bbn.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Bruce Taber	taber@interlan.com
Fumio Teraoka	tera@csl.sony.co.jp
Claudio Topolcic	topolcic@nri.reston.va.us

Ross Veach rrv@uiuc.edu
Luanne Waul luanne@wwtc.timeplex.com

2.6.7 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):

Steve Deering, deering@parc.xerox.com

Mailing Lists:

General Discussion: mospf@comet.cit.cornell.edu
To Subscribe: mospf-request@comet.cit.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 Ex
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 as an example, the existing BSD OSPF code, written by Rob Coltun
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

Minutes of the Multicast Extensions to OSPF Working Group (MOSPF)

Report not submitted. Please refer to the Area Report for a brief summary.

Attendees

William Babson	bill@penril.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
Scott Brim	swb@cornell.edu
Alan Bryenton	bryenton@bnr.ca
Jeffrey Burgan	jeff@nsipo.nasa.gov
William Carter	carter@ctron.com
Frank Chen	frankc@casc.com
Robert Ching	natadm!rching@uunet.uu.net
Rob Coltun	rcoltun@ni.umd.edu
Steve Deering	deering@parc.xerox.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Dwight Jamieson	djamies@bnr.ca
John Moy	jmoy@proteon.com
Erik Nordmark	nordmark@eng.sun.com
Philippe Park	ppark@bbn.com
Bala Rajagopalan	braja@qsun.att.com
Benny Rodrig	4373580@mcimail.com
Karen Seo	kseo@bbn.com
Frank Solensky	solensky@sceng.ub.com
Martha Steenstrup	msteenst@bbn.com
Kannan Varadhan	kannan@oar.net
David Waitzman	djw@bbn.com
Luanne Waul	luanne@wwtc.timeplex.com
Liming Wei	liming.wei@eng.sun.com

2.6.8 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>

"Proposed modifications to RFC 1247", 04/17/1992, John Moy <draft-ietf-ospf-v2update-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 at the July 1992 IETF in Boston. The Minutes from that meeting follow.

The meeting began with a review of the four documents that are currently be considered for publication by the Working Group:

1. The updated OSPF V2 specification. This will supersede RFC 1247. Unfortunately, the document was not available prior to the meeting. (A limited number of paper copies of the updated specification were made available to implementors, and the specification was made available for anonymous ftp after the meeting.) An excerpt from the document briefly detailing the changes was handed out, and the changes (all backward- compatible) were discussed. It was also decided to make one additional change: it will now be possible to specify a set of area address ranges that will not be advertised in summary-LSAs. This will enable a network administrator to hide certain networks within their local areas. This change has already been implemented by some vendors.

2. The updated OSPF V2 MIB. This will supersede RFC1253. Fred Baker outlined the proposed changes. It was also decided to make additions for the multicast routing extensions and the new NSSA area option. An addition to the Area Range Group was also made for the above "hidden network" feature. An additional request for a network mask in the new external-LSA table entries was not acted upon.

3. The OSPF Trap MIB. Rob Coltun led the discussion. There was some question whether an additional error code should be added for receiving Illegal-LSAs. It was decided that this would probably already show up as retransmissions by the faulty sender, and as such was unnecessary. It was also decided to have the ospfLsdbApproachingOverflow trap occur at a configurable database size, instead of 90 percent of the maximum (as stated in the draft).

4. The OSPF NSSA option. Rob Coltun spent some time explaining how they work, spending time on the translation between type-7 and type-5 LSAs, and how you could distinguish a "local" type-7 default from one that can be translated into a global type-5 default. No changes were made to this document.

Osmund deSouza presented a proposal for running OSPF over Frame relay. There was general agreement on the problem: Frame relay is in general not full-mesh connected, and the network administrator sometimes wants to assign different costs to different PVCs. For these reasons, OSPF's non-broadcast model is not directly applicable. There was also general agreement on the solution: instead of treating the connection to Frame relay as a single OSPF interface, define an OSPF interface as some collection of PVCs. There was

a long discussion of how to represent this in terms of MIB II and the OSPF MIB. It was decided that Osmund et. al., with the help of Fred Baker, would rewrite their present document more along the lines of a usage document. With this document in hand, it would be hoped that equipment from different vendors would be able to interoperate using OSPF over Frame relay.

John Moy presented an alternative model for running OSPF over Frame relay, where there would be a single interface to the frame relay net and a) neighbors would be discovered dynamically using Inverse ARP b) OSPF Hellos would be used to build a spanning tree among Frame relay connected routers, for purpose of update distribution (database synchronization) c) by default, only these spanning tree links (adjacencies) would be included in router-LSAs and d) to get better routing across the Frame relay, more PVCs could be included in the router-LSAs or (not as good) a variant of short-cut routing could be used. John's main reason for preferring this approach is that it didn't need a human to configure it, and that it was optimal in terms of routing traffic. This proposal was not generally well received, being characterized as either too complicated or too different than current practice. John said that he would write it up anyway if he had the time.

John Moy presented a proposal for dealing with OSPF Database Overflow. In this proposal, only the number of type-5 LSAs would be limited. The reasoning being that these constitute a majority of the database in places like the NSF regionals. A limit for the number of these LSAs would be set identically in each of these routers, either via SNMP or negotiated in a new LSA type or in OSPF Hellos. Then, when the limit is reached in a router it a) won't accept any more and b) will flush all its self-originated type-5 LSAs, refusing to originate any more. The claim is that this logic produces an identical database in all routers, with less than the configured maximum number of type 5 LSAs, no continual retransmissions, and all internal routing intact. Enhancements to this scheme could involve limiting other LSA types (e.g., summaries), and to begin again to originate type-5 LSAs after a random time lag to automatically deal with temporary overflow.

John said that a similar scheme has been used in Proteon routers for several years. The proposal was characterized by some Working Group members as being like congestion control, and some desire for an additional congestion-avoidance-like mechanism was expressed. Some people also requested a way to prioritize the order in which excess advertisements are flushed (e.g., you might want to flush the default routes last). John promised to sort through the enhancements and publish a coherent Internet Draft.

Rob Coltun ended the meeting with a quick discussion on how hierarchical routing information might be injected into OSPF, in order to support any of the schemes for IPV7.

Attendees

J. Allard	jallard@microsoft.com
William Babson	bill@penril.com
Dennis Baker	dbaker@wellfleet.com

Fred Baker	fbaker@acc.com
John Ballard	jballard@microsoft.com
Ken Benstead	kbenstead@coral.com
Geetha Brown	geetha@decvax.dec.com
Steve Buchko	stevebu@newbridge.com
Greg Celmainis	gregc@newbridge.com
Frank Chen	frankc@casc.com
Dean Cheng	dean@sun2.retix.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Jim Comen	comenj@interlan.interlan.com
Michael Davison	davison@cs.utk.edu
Osmund de Souza	osmund.desouza@att.com
Dino Farinacci	dino@cisco.com
AnneMarie Freitas	afreitas@microcom.com
Vince Fuller	vaf@stanford.edu
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Der-Hwa Gan	dhg@nsd.3com.com
Ian Heavens	ian@spider.co.uk
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Steven Hubert	hubert@cac.washington.edu
Ronald Jacoby	rj@sgi.com
Dwight Jamieson	djamies@bnr.ca
Dan Jordt	danj@nwnet.net
John Krawczyk	jkrawczy@wellfleet.com
Alan Kullberg	akullber@bbn.com
Whay Lee	whay@merlin.dev.cdx.mot.com
Anthony Lisotta	lisotta@nas.nasa.gov
Robin Littlefield	rlittlef@wellfleet.com
John Moy	jmoy@proteon.com
Erik Nordmark	nordmark@eng.sun.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Henry Sanders	henrysa@microsoft.com
Hellen Sears	sears@interlan.interlan.com
Martha Steenstrup	msteenst@bbn.com
Linda Tom	toml@interlan.interlan.com
Kannan Varadhan	kannan@oar.net
Scott Wasson	sgwasson@eng.xyplex.com
Luanne Waul	luanne@wwtc.timeplex.com
Honda Wu	natadm!honda@uunet.uu.net

OSPF Trap Update

- Authentication Failure Trap
  - -- No Longer Part Of Config Error
- Ignoring Initial Intf Traps
  - -- Intf Related Traps Dependent On Dead Timer
  - -- Non-Intf Traps Use Smallest Dead Timer
- ospfLSDBOverflow and ospfLSDBApproachingOverflow Added
  - -- Dependent On ospfLSDBLimit
  - -- Never Ignored
  - -- Approaching Is Triggered At 90% Of ospfLSDBLimit
- Need Illegal LSA Received In BadPacket?

NSSA Update

- DP-Bit Is Now P-Bit
- Type-7 -> Type-5 Transalation Is Now Performed After Type-7 Routes Are Calculated (New Section)
- P-Bit (Used In Type-7 LSAs) And N-Bit (Used In Hello Packets) Are Now Multiplexed To The Same Option Bit
- Short Section On Configuration Parameters Added
- Default Route Originated By ABR Is Never Translated Into Type-5
- Default Route Originated By NSSA Internal Router May Be Translated
- Order Of Route Preference By ABR
  - -- Any Type-5 LSA
  - -- A Type-7 LSA With P-Bit Set
  - -- Any Other Type-7 LSA

OSPF Over Frame-Relay Networks

- Problem Statement
- An Example
- A Wish List
- Proposed Solution

Motivation

To allow interconnected routers running OSPF to make efficient and flexible use of a frame relay network.

Example - Cost Effective Network Design

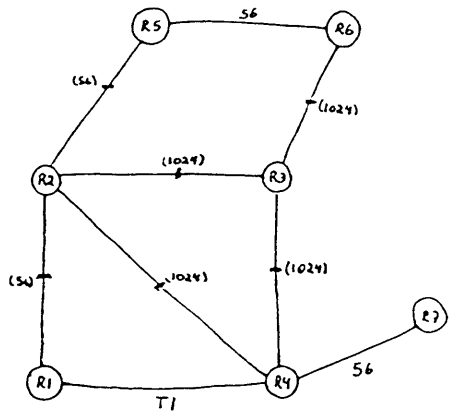

Can we use OSPF in this network?

Constraining Factors

- Non-Broadcast Multiple Access (nbma) is not a good model for this network:
  - — requires full mesh connectivity, or
  - — multiple frame-relay interfaces
- There is no mechanism to use bandwidth efficiently:
  - — different cost metrics for PVCs with different bandwidth guarantees

Requirements

The solution must:

a. Allow for non-full mesh connectivity

b. Allow different metrics for different PVCs

c. Allow for nbma as well

d. Allow a) and c) on a single frame-relay interface

e. Keep the procedure for bringing up adjacencies simple

f. Be backward compatible with minimal change to OSPF

Proposed Solution

- Fix the root cause: A generic interface should not only be a physical interface.

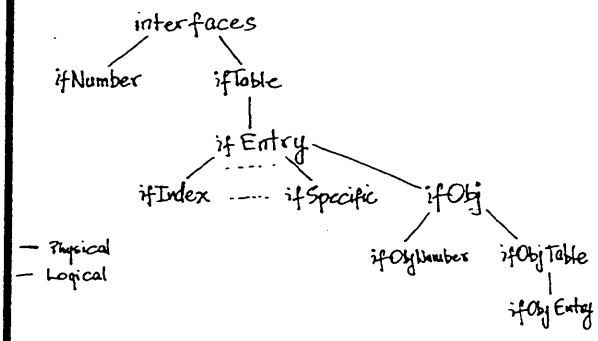

One OSPF interface for each ifObjEntry

Conclusion

- The "OSPF over FR" problem is important
- We identified key requirements
- The proposed solution meets all the requirements.
- What next?

Database Overflow

- Configuration parameter: maximum type 5 LSAs
- When limit reached
  - -- don't accept more
  - -- flush self-originated type 5 LSAs
- Claim: converge on identical database in all nodes
  - -- while still maintaining internal routing
- Limit must be the same in all routers

  How?
  - -- SNMP, or
  - -- new LSA, or
  - -- in Hellos
- Enhancements
  - -- start originating after time lag
  - -- limit other LSA types (e.g. summaries)

HOSPF

- OSPF Has 4 Levels Of Hierarchy
  - -- Intra-area
  - -- Inter-area
  - -- External Type-1
  - -- External Type-2
- Can Easily Map Into Hierarchical Address
  - -- PIP RH Fields
  - -- Domain/Area/Intra-Area
- Key Ideas
  - -- Intra-area remains the same
    - > Routes are added to routing table as first level PIP RH
    - > Routes are added to routing table as my domain/my area/this net
  - -- Inter-area Routes Advertised And Added To Routing Table As Next Level Of Hierarchical Address
    - > New Type Of LSA
    - > New type of area
      - # Supports new lsa type and type-3 for V4 and V7 coexistence
  - -- External Routes Advertised And Added To Routing Table As:
    - > Same level for type-1 externals and next level for type-2
    - > May need variable length address
    - > Need new external lsa type

2.6.9 RIP Version II (ripv2)

Charter

Chair(s):

Gary Malkin, gmalkin@xylogics.com

Mailing Lists:

General Discussion: ietf-rip@xylogics.com
To Subscribe: ietf-rip-request@xylogics.com
Archive: xylogics.com:gmalkin/rip/rip-arc

Description of Working Group:

The RIP Version 2 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:

Done	Review of RIP-II Internet Draft to ensure the additions are useful and backwards compatible. Also ensure that the additions cannot cause routing problems.
Done	Final review of RIP-II Internet Draft and submission into the standards track. First review of RIP-II MIB.
Done	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 Carrying Additional Information", 08/14/1991, Gary Malkin <draft-ietf-malkin-rip-04.txt>

"RIP Version 2 MIB Extension", 04/09/1992, Gary Malkin, Fred Baker <draft-ietf-ripv2-mibext-03.txt>

"RIP Version 2 Protocol Analysis", 08/14/1992, G. Malkin <draft-ietf-ripv2-analysis-00.txt>

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the RIP Version II Working Group (RIPV2)

We held a short meeting. There was some discussion about whether or not the routing domain field was useful. It was determined that it was useful for creating routing policies. That decision will be reflected in the draft.

It was determined, without opposition, that the drafts should be submitted for inclusion into the standards track. There are currently two independent implementations of RIP-2 and one implementation of the MIB (of which I am aware).

The Group is to be congratulated for completing the RIP-2 draft according to the original schedule and for completing the RIP-2 MIB draft ahead of the original schedule.

Attendees

Steve Alexander	stevea@i88.isc.com
David Borman	dab@cray.com
Steve Buchko	stevebu@newbridge.com
Robert Candela	candela@sceng.ub.com
Greg Celmainis	gregc@newbridge.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
AnneMarie Freitas	afreitas@microcom.com
Der-Hwa Gan	dhg@nsd.3com.com
William Haggerty	haggerty@ctron.com
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
Gary Malkin	gmalkin@xylogics.com
Glenn McGregor	ghm@merit.edu
Douglas Miller	dmm@telebit.com
Erik Nordmark	nordmark@eng.sun.com
Lars Poulsen	lars@cmc.com
Anil Rijsinghani	anil@levers.enet.dec.com
Sam Roberts	sroberts@farallon.com
Henry Sanders	henrysa@microsoft.com
James Watt	jamesw@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Honda Wu	natadm!honda@uunet.uu.net

2.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker/TIS

As usual, the Security Area was active and varied. Before surveying the activity within the IETF, there was an interesting development in the NREN program. The NREN is the National Research and Education Network. Actually, it's not really a network at all; it's a funding program within the U.S. to support a variety of important network developments and applications.

As part of the NREN program planning, the question was asked, "What are we going to do about security in the NREN?" And the next question is what does that mean and how is that different from security in the Internet as a whole? Toward this end, a small workshop was held in early July calling the question what can we do in the short run with available technology to protect the high-value portions of the NREN, which translates, roughly, into how do you keep hackers out of the super-computer centers? After a day and a half of discussion, a set of ideas emerged which are headed into the recommendations. This is not an official result yet, but what comes out should have some resemblance to this. And it is interesting because there is a lot of commonality with things of concern within the IETF Security area. The key ideas include:

- Finding a way to reduce the number of passwords which are sent in the clear.
- Using better forms of access control and better log-in mechanisms.
- Using Kerberos and successor technology for distributed authentication.
- Tightening up the configuration controls on systems, along with tools to help manage configurations.
- Using privacy enhanced mail.
- Using audit trails and logs and developing tools for analyzing and reducing the data collected in the logs.

There were also a number of people-intensive issues, which is always a tough problem because it's hard to spend money on manpower for security instead of all the other things that demand labor in a supercomputer center. Nonetheless, it is important to have people to define and implement a security policy at a center, educate people, and staff an organization to deal with security incidents when they come up.

That's the short form of what I expect will come out of this workshop by and by. These will be recommendations to the National Science Foundation in some form. The interesting thing is these people have an important problem, but it's hardly different from the problems that the rest of us have. The only difference is that they are looking at it in a smaller set of people and a smaller set of centers so that they may be able to move faster or do some things that might not be applicable across the entire spectrum. For example, if it's $100 per use to buy a device for remote log-in, it may be feasible in that context, whereas it's much harder to think about "solutions" like that for four million users around the world.

Now for the survey of where we stand in the Security Area.

Authorization and Access Control BOF (AAC)

Cliff Neuman led a BOF on authentication and access control. This is aimed at facilitating the development and deployment of applications that make use of processes running on multiple machines. Today, to do something like that, the access control mechanisms have to be jury-rigged in some ad hoc fashion, usually by setting up accounts and having processes log in as if they were users. What's needed is a common mechanism that can be used to build a wide variety of distributed applications. The question is can we identify some common tools and go down that path.

IP Security BOF (IPSEC)

An important new topic this time was an excellent BOF on IP Security, This covers authentication and encryption at the IP level so that when a packet comes in you can tell that it came from who it claims to have come from, which IP address. If necessary, it has been encrypted so that it can't be seen while in transit. There is a substantial amount of work in this area in the OSI arena and related protocols with network layer of security protocol (NLSP) and transport layer of security protocol (TLSP). The direction that we're heading in here is adopting that kind of technology into this suite. Joe Tardo from Digital give a very nice presentation on TLSP. We're going to follow up in November with a second BOF as an educational series and look at key management protocols and related matters. Then we'll choose a particular direction. My expectation is that we will kick off a formal working group in that area and put a lot of energy behind that.

Commercial IP Security Option Working Group (CIPSO)

The TSIG, the Trusted Systems Interoperability Group, is working on the commercial IP security option. This work is jointly chartered with the IETF.

The one strategic issue is whether to try to converge with the existing IP security option or whether to proceed along the original path, which was to leave it strictly alone and have a completely disjoint one. The context has changed in that there's some willingness on the part of the U.S. Government to consider some flexibility. It opens up the possibility of doing something more efficient and better. It would pay off quite a bit strategically and

it's worth the short-term pain, but the idea needs to be tested. This will be resolved very quickly, one way or another. Once that issue is resolved, I expect the document will get wrapped up rather quickly.

Common Authentication Technology Working Group (CAT)

The principal document, the GSSAPI, General Security Service Application Program Interface, document is in final review stage. I expect to submit that to the IESG for recommendation to Proposed Standard status within a very short time.

Network Access Server Requirements Workin Group (NASREQ)

Another development in the Security Area is the formation of new Working Group on Network Access Server Requirements. This is a generalization of the terminal server area focusing on the kinds of things needed in network access servers. The requirements document will be similar to a router or host requirements document. This work is taking place in the Security Area because the bulk of the concern is on access control and authorization.

A new development in another area is the development of the Mobile IP protocol. Steve Bellovin has accepted the task of working with that Group to work out the security issues. There's a definite and obvious security problem about somebody who just comes on the air and claims to be somebody from the other side of the world, but happens to have moved temporarily. That could be a little disconcerting if it's inaccurate.

Privacy Enhanced Mail Working Group (PEM)

The specifications are essentially done. There is a very small amount of tweaking to be done, and then four documents should show up ready for advancement to Proposed Standard status within the next several weeks.

One issue that does loom large is the convergence of PEM and MIME. Although it looks trivial at first glance, there are some complications. There are a number of loose ends and rough edges that are not trivial to resolve. We know what the questions are and because of the difficulty of doing anything to the PEM specification at the moment, the implementations trying to get out the door, the whole system trying to be born, we're going to leave it exactly alone except for putting in one hook to be able to bail ourselves out later and then take up as a new work item, the convergence of PEM and MIME as essentially a separate topic.

TCP Client Identity Protocol Working Group (IDENT)

Unlike the often contentious exchanges on the mailing list, the Working Group meeting was a model of restrained, productive, sensible, and efficient discourse. The Group converged on all except one detail which needs to be tracked down. This protocol should be wrapped up quickly.

Trusted Network File Systems Working Group (TNFS)

TNFS is another TSIG Working Group jointly chartered with the IETF in the Transport and Services Area. That work is proceeding quite well. Fred Glover, the Working Group Chair, presented the work and it's included in these Proceedings. For additional notes, please see TNFS under the Transport and Services Area.

TSIG Coordination BOF (TCOORD)

Of a more general nature, there were discussions about the overall relationship with the TSIG and about an overall security architecture.

With respect to TSIG interactions, we spent a considerable amount of time talking about and thinking about what is the best way to interact with them as a Group. There are some things that are working pretty well and others that could be better. The idea of taking their normal workshop and dropping in as an encapsulated activity inside of ours is beneficial, but it could be a lot better by thinking hard about how to make use of the combined environment. This would improve interactions in both directions and make better use of the time. The joint meeting this time had a very positive effect of opening eyes in both directions as to what the activities are.

The TSIG does meet more often than the IETF, so there much of their work will proceed as it always has. However, joint meetings are useful, and we will schedule more or less regularly. The details remain to be worked out.

The question of an overall security architecture keeps coming up. It is not clear to me whether or not we can in fact write down a security architecture that is going to have the effect of telling us what our choices are and making those choices and guiding us through all of the different wickets. But there is certainly some energy being applied in this direction. The Privacy Security Research Group is developing a framework document. Discussion about security architecture or ideas leading to security architecture should start to appear over the next many months.

Additional Nttes

PPP Authentication: The document has been sitting in the IAB for a while and the issues are very minor. I'm told they have been resolved.

Secure SNMP: Secure SNMP is completely done and finished. It's been published as a Proposed Standard. However, the Management area is extremely active, and there's a new version in the pipeline, SMP. One of the issues is the security model. Therefore the Secure SNMP effort will be reconstituted and will focus on SMP security.

Telnet Authentication: As is reported in the Applications Area, there is an Experimental document. We'll let that proceed for a while and build up experience before deciding what to do next.

We've had some discussions about one-time passwords and mechanisms that are possible and how to introduce them broadly in the network. That had been dormant for a while. It's been re-invigorated and announcements will be made through the usual channels.

CURRENT MEETING REPORT

Reported by Clifford Neuman/ISI

Minutes of the Authorization and Access Control BOF (AAC)

The first meeting of a new BOF on Authorization and Access Control met at the July IETF. The purpose of the BOF was to discuss authorization and access control issues for the Internet. The discussion centered around two problems: first, the need for a uniform method for specifying access control information, and second on services and mechanisms for distributing authorization in the Internet.

Agenda

- Discussion of requirements for the specification of access control information.
- Discussion of existing and evolving distributed authorization mechanisms including: DCE, DSSA, ECMA, Sesame, and Proxies.
- Discussion of the relationship between this Group and the Common Authentication Technology Working Group.
- Discussion of our goals.

Discussion

The first two items were related and discussion flowed from one into the other and back again. The need for a uniform method of specifying access control information for distributed applications was discussed.

One of the motivations for such an interface is to interact with the network authentication methods that are evolving. Such methods identify the subject accessing a service, but service specific methods are then needed to decide whether the subject is authorized access. Most existing applications do not presently maintain the information needed to make such decisions.

In the near future, application developers will have a common interface (the GSSAPI) that they can use to add strong authentication to their applications. That solves only half the problem. Ideally there would also be a common set of tools that they can use to decide whether the subject is authorized access.

The general consensus was that our work in this area should concentrate on access control lists as a conceptual model. Some of the distributed authorization mechanisms (described later) enable that model to support a full spectrum of access control methods including capabilities. Support for these mechanisms should be considered for inclusion in the model.

It was mentioned that work has gone on in POSIX and elsewhere to specify access control list mechanism for Unix. It was felt that we should consider such work, and build upon rather than replace it, but that we must support the needs of network applications.

An access control list (ACL) can be associated with objects to be protected. The ACL contains entries that identify subjects, either as individuals, or as members of groups. The entries specify the rights of the named subject to access the protected object. One point of contention was whether each distinct right for an object (read, write, execute, etc.) should be represented by a separate ACL (i.e., column in the access matrix), or an ACL should be associated with each object with the entries within the ACL specifying the rights. The two are equivalent, and consensus was that a single ACL per object was the preferred choice. It was also felt that in general the meaning of rights in an ACL entry would be application specific (interpreted by the server), but that the meaning of certain rights, in particular the ability to modify the ACL, should be common across all ACLs.

One extension to the ACL concept important for use on the Internet is that the identification of the subject should also identify the authentication method (or set of acceptable methods) to be used in identifying the subject. This is important because of the varying strengths of alternative authentication methods, and perhaps more importantly because the methods might not share a common name space for principals. There was very little discussion on this topic and any decisions here can await further work on the security model and access control abstractions.

A less straightforward extension is the addition of an optional field to each ACL entry that would allow restrictions of the authorization to be specified. Examples of restrictions include time of use, the need for additional authentication or authorization (e.g., a co-signer), etc.

Steve Crocker pointed out that the mechanism and abstraction must be simple and easy to understand in the common case, a sentiment with which everyone agreed. It was also felt, however, that in the ideal case the mechanism would also be flexible enough to support the needs of most applications, so that developers would not be forced to design their own mechanisms.

It was felt that our goals in this area should be to:

1. Identify the target applications from which to draw requirements (e.g., Mailing lists, files, login, X windows).

2. Identify the security models to be supported (e.g., We might consider discretionary access control, mandatory access control, capabilities, access control lists, etc.).

3. Work out the appropriate access control abstractions.

4. Consider an application programmer interface (API) to support those abstractions (consider POSIX, DCE, our own, etc.).

A final issue raised concerned whether our model needed to support the needs of licensing and accounting mechanisms. It was felt that we should keep such mechanism in mind, but that it was premature to consider them as an integral part of our current work.

The second phase of the meeting concentrated on the evolving mechanisms and architectures for authorization in distributed systems. This phase consisted of informal presentations of some of the mechanisms.

Joe Pato described the authorization mechanism used by OSF's Distributed Computing Environment. Authorization in DCE is based on privilege attribute certificates issued by a privilege server. These certificates are restricted Kerberos tickets (see restricted proxies described later) that specify the UID and groups to which a principal belongs. The privilege attribute certificate is then used by the principal to assert its membership in groups, and its UID. The DCE also supports an extensible ACL model for distributed systems based on and extending that in the POSIX draft. Authorization is a combination of privilege attributes and control attributes (e.g., ACLs). Support for delegation is being considered, further exploiting the ability to construct restricted proxies.

John Linn then described the authorization mechanisms that are part of the Digital Distributed System Security Architecture (DSSA). One key aspect of the DSSA is that authorization credentials are pulled by the server, rather than pushed by the client, though the client provide's hints suggesting which credentials should be pulled. A second aspect of the DSSA is that reduced authorization can be granted by establishing a new role, a new principals with reduced privileges. The DSSA supports delegation with identifiable intermediaries, but delegation is of all rights possessed by a particular role.

Piers McMahon then outlined the main features of the ECMA TC32/TG9 Security Architecture. The ECMA model is that a trusted authentication service authenticates subjects using a suitable authentication method, and that a logically separate privilege attribute server (PAS) grants privileges (e.g., identity, role, group, capability, clearance) to that subject. The privilege acquisition is constrained by the level to which the subject is authenticated - but is independent of the authentication method. The privileges are cryptographically protected by the PAS and returned in a data element called a Privilege Attribute Certificate (PAC), and are sent (pushed) by the security subject to target systems to inform access control decisions. Methods for protection of PACs, together with controls on their use and delegation are defined by current ECMA work.

Piers then described SESAME. Some background information was given to explain that SESAME is a phased project based on ECMA TC32/TG9 work. Phase 1 produced a prototype which showed that the basic model was feasible. Phase 2 is building on this to develop product-level distributed security infrastructure with support for dialogue protection and DCE-interworking. An outline of the SESAME Phase 2 architecture was given to show how it built on the ECMA architecture, and a brief walkthrough of the privilege acquisition protocol was presented. It was stated that SESAME Phase 2 supports a subset of the full ECMA privilege attributes (identity, role, group), and a profile of ECMA PAC protection and controls.

Next, Clifford Neuman described restricted proxies. A restricted proxy can be implemented on top of an authentication mechanism by issuing authentication credentials authorizing a second party (the grantee) to act as the issuer for the purpose of performing a restricted set of operations, under specific conditions. These restrictions are supported by Kerberos V5 in the authorization data field. It was then described how restricted proxies can be used to implement a range of authorization mechanisms from capabilities, to authorization and group servers, and how restricted proxies might interact with the access control list mechanisms described earlier.

The next topic of discussion was how these mechanism might be used by applications. In particular, might it be appropriate to develop a common API within which they might fit. If so, might this API become part of the common authentication technology (GSSAPI) so that the programmer only need deal with one mechanism, instead of two.

Finally, we discussed the possibility of convergence of the various approaches. Some of the approaches are still in their early stages, and it would be helpful if we could encourage, for example, a common certificate structure across mechanisms. However, some of the mechanisms, in particular DCE, are further along, and significant changes would present problems. In any event, where possible, we should try to promote fewer protocols and message formats.

It was felt that our immediate goals in the distributed authorization area should be:

1. To look for common characteristics among the mechanisms.
2. Decide on a course of action. The range of possibilities include encouraging the use of a common credential format, developing other interoperability mechanism, defining a common API, and unifying the protocols.

Finally, It was felt that work is needed in the area of authorization and access control, and that the Group should continue to meet. As a potential working group, we must:

- Decide what the product of the working group would be.
- Develop a set of goals and milestones.
- Write a Charter.

It was felt that we should refine the Group's objectives through the mailing list. If we can develop a Charter in time for the next IETF, we can form a working group. If not, we should meet again as a second BOF, part of the purpose of which will be to agree on a Charter.

A mailing list has been set up, ietf-aac@ISI.EDU. Requests for addition or deletion should be sent to ietf-aac-request@ISI.EDU.

Attendees

Derek Atkins warlord@thumper.bellcore.com

Tony Ballardie	a.ballardie@cs.ucl.ac.uk
Doug Barlow	barlow@decwet.dec.com
Mark Baushke	mdb@cisco.com
David Borman	dab@cray.com
Geetha Brown	geetha@decvax.dec.com
Stephen Crocker	crocker@tis.com
Tom Farinelli	tcf@tyco.ncsc.mil
Barbara Fraser	byf@cert.org
Maria Gallagher	maria@nsipo.nasa.gov
Neil Haller	nmh@thumper.bellcore.com
John Linn	linn@erlang.enet.dec.com
Ellen McDermott	emcd@osf.org
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Cyndi Mills	cmills@nnsc.nsf.net
Clifford Neuman	bcn@isi.edu
Brad Passwaters	bjp@sura.net
Allan Rubens	acr@merit.edu
Gregory Ruth	gruth@bbn.com
Paul Sangster	sangster@ans.net
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
Jeremy Siegel	jzs@nsd.3com.com
Sam Sjogren	sjogren@tgv.com
Simon Spero	ses@cmns.think.com
Theodore Ts'o	tytso@mit.edu
John Vollbrecht	jrv@merit.edu
Jesse Walker	walker@eider.lkg.dec.com
Peter Williams	p.williams@uk.ac.ucl.cs

CURRENT MEETING REPORT

Minutes of the IP Security BOF (IPSEC)

Minutes not received. Please refer to Area Report for a brief summary.

Attendees

Derek Atkins	warlord@thumper.bellcore.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Doug Barlow	barlow@decwet.dec.com
John Batzer	
Luc Boulianne	lucb@cs.mcgill.ca
David Bridgham	dab@epilogue.com
C. Douglas Brown	cdbrown@sandia.gov
Mark Christenson	mgc@cray.com
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Jack Drescher	drescher@concert.net
Donald Eastlake	dee@ranger.enet.dec.com
Fran Fadden	fran@decvax.dec.com
Tom Farinelli	tcf@tyco.ncsc.mil
Barbara Fraser	byf@cert.org
Shari Galitzer	shari@shari.mitre.org
Pria Graves	priag@nsd.3com.com
Neil Haller	nmh@thumper.bellcore.com
James Harville	jbh@unh.edu
Mark Hollinger	myth@ssg.lkg.dec.com
Alton Hoover	hoover@ans.net
William Huyck	
John Ioannidis	ji@cs.columbia.edu
Ronald Jacoby	rj@sgi.com
Dean Jagels	dpj@sware.com
Frank Kastenholz	kasten@ftp.com
Stephen Kent	kent@bbn.com
Stev Knowles	stev@ftp.com
James Lin	yeejang@cup.hp.com
John Linn	linn@erlang.enet.dec.com
Ellen McDermott	emcd@osf.org
Glenn McGregor	ghm@merit.edu
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
April Richstein	abm@tycho.ncsc.mil
Allan Rubens	acr@merit.edu
Paul Sangster	sangster@ans.net
Richard Schmalgemeier	rgs@merit.edu

Martin Schulman	mas@loyola.edu
Aaron Schuman	schuman@sgi.com
Ron Sharp	rls@neptune.att.com
Robert Shirey	shirey@mitre.org
Sam Sjogren	sjogren@tgv.com
Richard Slade	ricks@ssd.csd.harris.com
Charles Watt	watt@sware.com

TRANSPORT LAYER SECURITY PROTOCOL

July 13, 1992

Joseph J. Tardo
Digital Equipment Corporation (LKG 1-2/A19)
550 King Street
Littleton, MA. 01460

tardo@tardo.lkg.dec.com

Slide 1

Overview

- Concepts and References
- History and Implementations
- Some Technical Details
- Relevance to IETF

Slide 2

Motivation

- Data Protection
  - Confidentiality and Integrity
  - Access Control
- Transparency
  - Applications
  - Protocols
- Implementation
  - Software
  - Inboard or Outboard Hardware
- Placement of Devices
  - At Each Host
  - In Gateways or Routers

Slide 3

Review Concepts

- Threats
  - Passive Monitoring (Eavesdrop, Traffic Analysis)
  - Active Attack (Spoof Identity, Insert Bogus Data, or Delete, Replay, Reorder, Splice, etc. "Real" Data)
  - Denial of Service
- Lower Layer Placement Issues
  - Data Link Provides Multi-Protocol, LAN-Bounded Protection
  - Network Protects Between Hosts, Subnets or Routers
  - Transport Protects Host-Host
- Policy Granularity
  - All Traffic Between Protocol Entities
  - Selective by Connection, Packet or Field

Slide 4

History and Implementations

- Derived From Secure Data Network Study (SDNS)
  - SP4 Provided Base OSI Draft (1986 - 1989)
  - Final Sequence Number Removed
  - Cleaned Up "CO/CL" Keying Terminology
- Implementations (SP4)
  - NIST Demo Lab (DEC, Hughes, IBM - c1989)
  - DEC A/D Chip (1992)
  - Others?
- NLSP (SP3) vs. TLSP (SP4)
  - The Controversy Continues...
  - Not Quite The Same (Can Mix Clear and Ciphertext Connections Between Same Transport Entities)

Slide 5

Security Encapsulation

- Extremely Simple Concept: Wrap TPDU
- No New Entities, No Added Protocol States
- Easy to Parse (Discriminate on TPDU Type)
- Red and Black Headers and Trailers

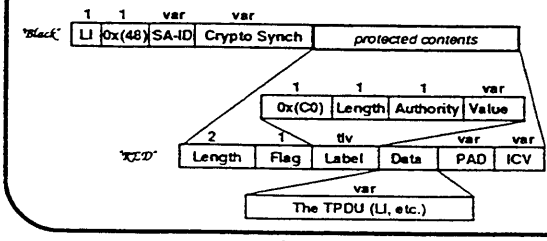

Slide 6

Transport Security Requirements

- OSI Security Architecture (ISO 7498-2)
  - Peer-Entity Authentication
  - Data Origin Authentication
  - Access Control Service
  - Connectionless Confidentiality
  - Connection Integrity with Recovery
  - Connection Integrity without Recovery
  - Connectionless Integrity
- Transport Layer Service Specification (ISO 8072)
  - None
  - Protection Against Passive Monitoring
  - Protection Against Modification, Replay, Addition or Deletion
  - Both Forms of Protection

Slide 7

Providing Security Services

- Encapsulating TPDUs Provides:
  - Peer-Entity Authentication
  - Data Origin Authentication
  - Access Control Services
  - ✔ Connectionless Confidentiality
  - ✔ Connection Integrity with Recovery (32 bit TPDU-NR, Class 4)
  - ✔ Connection Integrity without Recovery (32 bit TPDU-NR)
  - ✔ Connectionless Integrity
- Key Management Assumed To Do The Rest
  - Shared Key Between Transport Entities (CO or CL)
  - By Transport Entity Pair and Security Level (CO or CL)
  - _ Or Per Connection (CO)
  - Re-Key on TPDU-NR Overflow

Slide 8

Security Associations

- Defined In The LLSG
  - Pairwise Relationship Between [Transport] Entities
  - May Be Established Prior to Communication
  - SA-ID Names Shared Security Context (Keys, Labels, Other Attributes, e.g. ASSR)
- Use in TLSP
  - SA-ID Sufficient to Identify Incoming Key
  - Single SA Can Span Multiple Transport Connections
- Establishment
  - SP4: Relegate to Out of Band Mechanism (e.g., KMAE)
  - Could Also Do In-Band
  - No Protocol Elements Defined in TLSP or LLSG

Slide 9

Key Management PDAM

- "In-Band" SA Establishment
  - Definitions for SCI Messages
  - Includes Connection Authentication and Re-Key
  - Mechanism Independent Key Exchange (Symmetric KEK, Asymmetric, Exponential, etc.)
  - Rules, Encoding Still Being Defined
  - Some Similarities With IEEE 802.10 Key Management Proposal
- ToDo...
  - State Machine?
  - Service Primitives?
  - Option Negotiation?
  - MIB?

Slide 10

Don't Gets

- Denial of Service
- Broadcast
- Management Support (e.g., Auditing, MIB Definitions)
- Selective Invocation Beyond QOS
- Algorithms or Other Practical Details

Slide 11

Relevance For IETF

- Possible Use With TCP/IP
  - Encapsulate TCP or UDP Over IP/CLNP
  - Needs Its Own Protocol Type
  - Need Prototol Type Option In Protected Header
  - TCP Pseudo-Header Checksum
- RFC1006 Issues
  - Between TCP and IP?
  - In Convergence Layer?
- Key Management: Use PEM Certificates?
- International Export Problems

Slide 12

Abbreviations

ASSR	Agreed Set of Security Rules
TPDU	Transport Protocol Data Unit (Message)
TPDU-NR	Sequence Number (as, DT-TPDU-NR for Data)
SA-ID	Secure Association Identifier
LLSG	Lower Layer Security Guidelines
SCI	Security Control Information
ISN	Integrity Sequence Number
ICV	Integrity Check Value
KMAE	Key Management Application Entity (SDNS)

Slide 13

Referenced Documents

ISO/IEC JTC1/ SC6 ...

- N6779 DIS 10736 ... Transport Layer Security Protocol (October 16, 1991)
- N6959 ... Lower Layer Security Guidelines (August 9, 1991)
- N7053 CD 11577 ... Network Layer Security Protocol (November 11, 1991)
- N6794 DIS 10736 PDAM 1: Security Association Establishment Protocol (April 20, 1992)

Slide 14

CURRENT MEETING REPORT

Reported by Allan Rubens/Merit and John Vollbrecht/Merit

Minutes of the Network Access Server Requirements Working Group (NASREQ)

The session began with a review of the status and goals of the Network Access Server (nasreq) Working Group. There has been a BOF at each of the last two IETF meetings where the purpose of this Group was discussed. A Charter has been submitted to the IESG which is waiting final approval. This Charter proposes that the Group should write an RFC which spells out a set of standards required for Network Access Servers in the same way the router requirements RFC describes protocol requirements for internet routers. The NAS requirements are expected to be mainly in areas of Authentication, Authorization, and Accounting (AAA).

The Group met as a Working Group for the first time at the Boston IETF. There was discussion about how the Group will work and what its initial goals should be. The Security Area Director, Steve Crocker, stated that he would like this Group to only state the required standards; not define them. Steve suggested that the initial work of this Group be to refine the "Requirements" draft and submit it as an Internet Draft. It should explicitly be sent to the Chairs of the CAT Group, the Accounting Group, and the soon-to-be-formed Authentication and Access Control Group and they should be asked to discuss it at their meetings at the next IETF. We should schedule our meeting(s) later in the week so we can discuss what these other groups had to say.

A copy of a draft "NAS Requirements Document" was distributed. This document states the NAS AAA functionality requirements for which a required set of standards needs to be adopted. This document needs to be reviewed and reworked by next IETF. There was strong recommendation from Steve Crocker that this document be a first draft of the final document, but a desciption of NAS requirements. It should not include specific protocol recommendations; the protocol recommendations will come in consultation with other working groups. The "NAS Requirements Document" will be submitted as an Internet Draft about a month before the fall IETF.

At the last BOF it was suggested that the Group be divided into two subgroups; one to investigate the requirements for character-stream NAS access and the other to examine framed access. The feeling at that time was that the issues involved with these two access classes were sufficiently different that it would be better to proceed on two tracks. However, after further discussion at this meeting, the Group unanimously felt that the issues are not that different and it would be better to proceed as a single group working on both classes. The Charter will be revised to indicate this decision.

Volunteers were solicited to assist in the refining of the "Requirements" draft . Work on the draft will be shared on the auth-acct@merit.edu mail group. There was one new volunteer to help with the document. The document will be revised and sent to the mail group as

quickly as possible.

Attendees

George Abe	abe@infonet.com
Jim Barnes	barnes@xylogics.com
Charles Bazaar	bazaar@emulex.com
Charles Blauner	chazx@ctt.bellcore.com
Stephen Crocker	crocker@tis.com
Michael DeAddio	deaddio@thumper.bellcore.com
Kelly Furlong	kelly@kyle.ksc.nasa.gov
Shari Galitzer	shari@shari.mitre.org
Pria Graves	priag@nsd.3com.com
Ittai Hershman	ittai@nis.ans.net
Alton Hoover	hoover@ans.net
William Huyck	
David Katinsky	dmk@rutgers.edu
John Labbe	labbe@merit.edu
John Linn	linn@erlang.enet.dec.com
J. Scott Marcus	smarcus@bbn.com
Glenn McGregor	ghm@merit.edu
Stephanie Price Marasciullo	price@cmc.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
Allan Rubens	acr@merit.edu
Paul Sangster	sangster@ans.net
Robert Shirey	shirey@mitre.org
Jeremy Siegel	jzs@nsd.3com.com
John Vollbrecht	jrv@merit.edu
Jesse Walker	walker@eider.lkg.dec.com
David Wang	wang@xylogics.com
Peter Williams	p.williams@uk.ac.ucl.cs

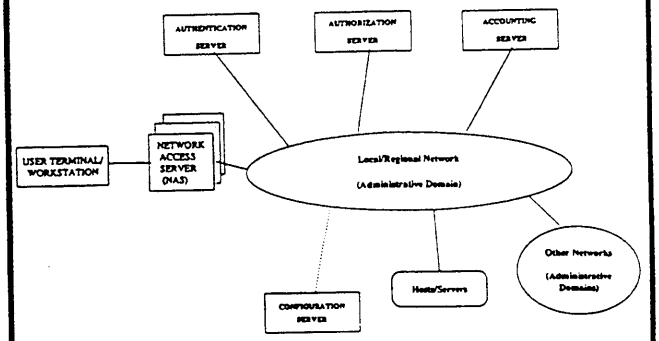

Figure 1 - The NAS AAA Model

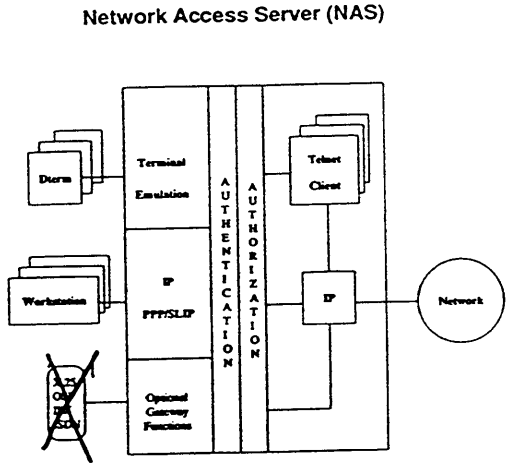

Figure 2 - NAS Functions

Authentication Requirements

- o A secure authentication mechanism is a necessity. It needs to be resistive to passive and active attacks.
- o The authentication mechanism should allow multiple NAS's in the same administrative domain to access a common authentication database. This database may be distributed. and, if so, the user should be able to indicate to the server which database entity his/her credentials are associated with.
- o Automatic authentication at the time of access to the NAS is desirable for both frame and character stream user access.
- o Some NAS ports may be configured to be implicitly authenticated (e.g., hard-wired PC in private office).
- o A user should be able to reauthenticate when desired.
- o The user's credentials should be presentable to the Authentication and Accounting servers.

Authentication Issues

- o What about user credentials associated with systems that don't share a common authentication mechanism with the NAS?
- o Is it practical to run an authentication system such as Kerberos on NAS's? Possibilities and tradeoffs of a proxy authentication agent?
- o How long is a given authentication valid for? How can use of a previous user's authentication on a shared access device (a workstation in a public area) be prevented?

Authorization Requirements

- Use authentication credentials to obtain "Access Control List" (ACL) from the authorization server; ACL might specify the following kinds of access restrictions:
  - Authorized IP addresses (ranges) and ports ;
  - Unauthorized IP addresses (ranges) and ports.
  - Ability to impose charges.
  - Ability to utilize specific NAS services.
  - Routing table to be used (e.g., separate routing tables for commercial and research or educational user).
- Authorization server must interact with Accounting server to determine if appropriate account limit restrictions are met before authorizing NAS access.

Authorization Issues

- Does authorization by destination server potentially require access to information that must be provided by the NAS? For example, how would a destination server restrict access to connections from a specific set of NAS ports?
- Can ACL be made "rich" enough to satisfy arbitrary NAS access control requirements? For example, what should be done for time-of-day restricted access or "No access allowed on the 4th Tuesday of each month?"
- Mobile hosts?
  - IP address(s) to use for PPP/SLIP link need to be authorized by NAS.
  - Need mechanism to insure that address is not multiply assigned.

Mobile users

Accounting Requirements

- The NAS relies upon the Accounting server to process the accounting information it generates and provides to the Accounting server. It also relies on the Accounting server to maintain user account balances.
- The NAS accounting mechanism must provide sufficient data to generate any desired (and "reasonable") usage statistics.
- Need to collect sufficient information for audit tracking.
- Need to insure that generated accounting information is general enough to allow for charges based upon a variety of parameters and a variety of charging units. Need to include a handle for account to be charged.
- Because charges may be involved, a reliable transport mechanism for the collection of accounting information needs to be utilized.
- Need to be able to examine accounting/statistics information on a timely basis to enable audit tracking
- The Accounting server must interact with the Authorization server to provide account limit information for access control purposes.

- Secure

Accounting Issues

- How are charges monitored for active NAS access? Does a user get disconnected when account balance goes to $0.00?
- Should the accounting policy (rates and charges) be implemented in the NAS, the Accounting server, or shared by both? Can rate determination and charging be deferred to post-processing?
- One way to do charging in the NAS would be to have the Accounting server download a charging algorithm (a shared accounting policy), in some standard format, into the NAS. Is this a reasonable approach to pursue?

CURRENT MEETING REPORT

Reported by Keith McCloghrie/Hughes

Minutes of the SNMP Security Implementors' BOF (SNMPSECI)

A BOF session for SNMP Security Implementors was held during the Boston IETF meeting on July 13, 1992. The BOF's purpose was to allow implementors to share their implementation experiences. The meeting was Chaired by Keith McCloghrie. Jim Galvin sent his apologies for not being able to attend.

The meeting began with a review of the status of SNMP Security:

- RFCs 1351, 1352, 1353 have been published with Proposed Internet Standard status,
- The RFCs have lots of editorial changes from the Internet Drafts which the Working Group had approved, but
- The only change affecting implementations was the assignment of OBJECT IDENTIFIERs under the mib-2 branch.

After reviewing the status, the meeting was opened to questions and comments from the attendees. An informal poll of the audience indicated that at least six implementations of secure SNMP existed. The discussion topics included:

- Export issues
- Clock synchronization
- Access control granularity
- MD5/DES performance overhead
- BER encoding
- Relation to SMP
- "Next steps" for the RFCs.

During the discussion of export issues, some (second-hand) information was presented on a proposal being considered by NIST for an "improved" process for U.S. export control of cryptography.

The discussion on clock synchronization raised the issue of how SNMP Security relates to the recent SMP specification, since a change to clock synchronization is proposed by the SMP specification. Thus, each of the changes to SNMP Security being proposed as part of SMP were presented. In particular, in the area of clock synchronization, SMP simplifies the algorithm by including both the destination party's clock as well as the source party's clock in the authInfo structure of a message; this removes the need for a SetRequest to be issued (in the "case 1" scenario described in RFC 1352). Another suggestion concerning clock synchronization was the use of automatic, "on the fly" synchronization of clocks whenever

an application requests a message be sent to an agent which it hasn't recently communicated with.

In other discussions, the impact on performing access control on MIB views with instance-level granularity was discussed, particularly the performance aspects of it.

Performance was also discussed in regard to the overhead of MD5 and DES. Feedback from newer implementations was compared to previously known information, and was found to be within the same ballpark. David Partain's article in the July issue of "The Simple Times" was mentioned as a source of more information.

One implementor indicated that differences in BER encodings by different implementation could cause problems. The authDigest value calculated on the SnmpAuthMessage by the receiving entity has to match the authDigest value contained in the message when these values are compared during authentication processing. In particular, ISO 8825 allows multiple valid encodings of a length field. Thus, the receiving entity must not perform an independent BER serialization/encoding, but must use the same serialized value as it received. Not only is this necessary but it can also be beneficial, since it allows implementors to minimize the number of times BER encodings are performed in their code.

Several attendees raised questions on the "next steps" for secure SNMP in light of the changes outlined in the SMP documents. There were questions on whether the SNMP Security RFCs would be updated and when. Additionally, there were questions on whether implementors should "hold off" on implementing SNMP Security until the status of SMP/SNMP II was known. Attendess were urged to participate in the SMP BOF scheduled for later in the week where these issues would be discussed.

Attendees

Steve Alexander	stevea@i88.isc.com
David Arneson	arneson@ctron.com
Jim Barnes	barnes@xylogics.com
Andy Bierman	bierman@davidsys.com
Tom Brennan	
David Bridgham	dab@epilogue.com
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Robert Ching	natadm!rching@uunet.uu.net
Chris Chiotasso	chris@artel.com
Tracy Cox	tacox@sabre.bellcore.com
Cathy Cunningham	cmc@microcom.com
James Davin	jrd@ptt.lcs.mit.edu
Michael Davison	davison@cs.utk.edu
David Engel	david@ods.com
Michael Erlinger	mike@lexcel.com

Rob Graham	robert_graham@protools.com
Pria Graves	priag@nsd.3com.com
Jeff Hughes	jeff@col.hp.com
Ronald Jacoby	rj@sgi.com
Frank Kastenholz	kasten@ftp.com
Nick Kawaguchi	mamster@lanai.cs.ucla.edu
Mark Kepke	mak@cnd.hp.com
Kenneth Key	key@cs.utk.edu
Deidre Kostick	dck2@sabre.bellcore.com
Hock-Koon Lim	lim@po.cwru.edu
John Linn	linn@erlang.enet.dec.com
Arun Mahajan	axm@proteon.com
Kent Malave	kent@chang.austin.ibm.com
Kim Mayton	mayton@wg.com
Keith McCloghrie	kzm@hls.com
Thomas McGinty	mcginty_t*corp_m@msm.cdx.mot.com
John McKenna	mckenna@ralvm12.vnet.ibm.com
David Minnich	dwm@fibercom.com
Lynn Monsanto	monsanto@sun.com
Paul Moran	Paul_Moran@3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Sam Nicholson	scion@pblx.knox.tn.us
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
Andrew Patka	apatka@wellfleet.com
John Payne	jop@wang.com
David Perkins	dperkins@synoptics.com
Richard Ramos	ramos@mtunm.att.com
Ed Reeder	EREEDER@ralvm12.vnet.ibm.com
Sam Roberts	sroberts@farallon.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Michael Sapich	sapich@conware.de
Koichiro Seto	seto@hitachi-cable.co.jp
Timon Sloane	peernet!timon@uunet.uu.net
Einar Stefferud	stefisoc@nma.com=
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Stephen Tsun	snt@nsd.3com.com
Ahmet Tuncay	atuncay@synoptics.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Huyen Vu	vi@polaris.disa.mil
David Waitzman	djw@bbn.com
Gerard White	
Steven Wong	wong@took.enet.dec.com

Honda Wu natadm!honda@uunet.uu.net
Kiho Yum kxy@nsd.3com.com
Joseph Zur zur@fibhaifa.com

2.7.1 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 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 (CIPSO 2.2)", 12/03/1991, Trusted Sys Interop. Group (TSIG)

<draft-ietf-cipso-ipsecurity-01.txt>

CURRENT MEETING REPORT

Reported by Ron Sharp/AT&T

Minutes of the Commercial Internet Protocol Security Option Working Group (CIPSO)

Due to other IETF meetings and additional TSIG plenary sessions the Working Group met for only six hours this meeting. The primary discussions involved the IETF/TSIG relationship and how to allow and encourage more participation from other IETF members. There was also much discussion concerning the Internet Draft entitled "Son of IPSO" that was submitted by Michael StJohns.

The format for this meeting changed a little. Issues were presented and discussed, however there was no voting to determine the Group's consensus. It was felt by some new attendees that this led to the idea that all work and decisions was done at the meetings and if you could not attend the meetings then you were left out. That was not the intended purpose of voting, however, it must be admitted that the result may still be the same. We encourage anyone to participate either at the meetings or electronically. Ron Sharp has been trying to push people into using the electronic media more but it has been used only a little. When an issue does come up on the mailing list and is not resolved Ron includes it in the Agenda for the next meeting. Even if there is a consensus at the meeting the issue is still alive as long as someone is willing to discuss it in any forum.

Ron will go over the issues discussed and the resolutions that were purposed. Please respond to the mailing list if you disagree with any of the proposals. If Ron hears no discussion he will make the appropriate change to the specification. Even then the issue is not dead and may be brought at a later time when some things may be clearer, though the sooner the better for everyone.

Issue 1: Changes to CIPSO Version 2.2

There were several nit changes to the CIPSO specification for accuracy and readability. These changes will be marked in the next release of the CIPSO specification. The process for releasing the specification was also changed. As editor of the specification Ron will gather the comments from the meetings and the mailing list and will make the appropriate changes. He will first put the new specification out on the mailing list for comments. After two weeks or so, depending on the comments received, Ron will send a revised version to the Internet Drafts database. He hopes to have a new draft of CIPSO 2.2 out for comments soon that will include the last two meetings and discussions between the meetings.

Issue 2: CIPSO MIBs

Tabled.

Issue 3: Router Participation

There were at least two router vendors at this meeting from cisco and 3com. It is hoped that more will be heard from them on the mailing list describing their needs and requirements. The cisco representative said that cisco is waiting for a decision to be made as to which is going to be the next IP label. She said they were about to go with CIPSO when SIPSO came out. We need to get to one specification, one label soon.

Issue 4: Test Plan

The next IETF CIPSO Working Group meeting will be in conjunction with TSIG in Minneapolis, Sept 22-24. At this meeting several vendors will bring their CIPSO implementation and test interoperability. Cray has graciously offered to host the meeting for the interoperability test. Aaron Schuman of SGI wrote a test plan to use. The plan was reviewed and several changes were made. The primary change was to use telnet as the application to test basic CIPSO functionality. Telnet was chosen since it was common to all implementations. Aaron will get a revised test plan out prior to the next meeting.

Issue 5: CIPSO, BSO Translation

Aaron presented a solution to allow a CIPSO gateway machine to translate BSO labels to CIPSO labels and CIPSO labels to BSO. The security level would mapped to the corresponding value for the other label. The BSO PAFs would map to CIPSO DOIs. Each combination of PAf flags would be a unique DOI. Mike suggested including this map directly in the CIPSO specification.

Issue 6: BSO tag type

Tabled.

Issue 7: Future of CIPSO Working Group

The Group decided to meet next time in conjunction with TSIG. A lot of electronic discussion is needed to resolve some of the remaining issues. The primary issues are described at the end of the Minutes. Steve Crocker agreed to work with us to resolve all issues between CIPSO and SIPSO prior to the next IETF meeting. The goal is to have a CIPSO specification that is acceptable to the IETF and the CIPSO vendors which incorporates the best of both specifications. Without a resolution soon, we will end up with three standards. IPSO will still be out there and included in new systems since there is no new unified label. CIPSO vendors will continue to ship CIPSO, but it will not be based on an IETF standard which they would prefer and SIPSO will be trying to get vendor participation.

Issue 8: CIPSO Option Processing

MSJ felt that the description of option processing in the specification should be split out by

end systems, intermediate systems, and routers. I will look at SIPSO and make appropriated changes to make the processing clearer.

Overall it was a good meeting. The Group did not get many issues covered but there was more dialogue as to what is expected of CIPSO to finally get to Proposed Standard stage which is long overdue. Ron feels there are still four primary issues that must be addressed and resolved between the CIPSO vendors and a few other IETF members. These are listed below:

- A. IPSO backward compatibility.

  MSJ feels that the first 4 bytes of CIPSO could look like IPSO and thus have interoperability. The PAFs would represent a unique DOI like discussed in issue 5 above. If we could truly get backward compatibility then we could more quickly move to one IP security option which is what everyone wants. There is the question of whether existing implementations like BLACKER can accept these new CIPSO options without modifications. If modifications are necessary than why not just move to a full CIPSO and get the added flexibility and interoperability a full CIPSO implementation offers. There is also concern that this would tie CIPSO to a particular security policy, that of the US DOD when the commercial market has show little interest in hierarchical labels.

- B. Number of CIPSO tags supported in this RFC

  The current draft has three tags to allow for large category sets. MSJ questions whether 3 are necessary.

- C. CIPSO currently allows for tag types above 127 to be defined by the DOI. This allows for support of new policies such as integrity and to hide classified formats and definitions. There is a concern that this could hinder interoperability. The Working Group has been working on this issue and the current draft includes words that state that implementations that support tags above 127 must be able to configure a DOI that does not require those tags. This will assure communication using standard well defined tags in the event of an emergency like the Gulf war.

- D. Inclusion of application to TCP or UDP interface processing rules.

  It is felt that, while this is a good idea, it may belong in an RFC that describes a network level security option. SIPSO includes some of these rules, however they are included as suggestions.

The above should cover the last meeting and where the Group is currently. If anything has been missed, please respond to the mailing list. Discussion of the four issues identified is needed. If anyone feels there are others than please include them. There are other issues such as options processing, however Ron has confidence that these can be worked out.

Thanks for attending the meeting and helping out. A special thanks to Aaron Schuman who presented two homework items AND recorded the minutes which were used to produce these minutes. Ok now lets hear some discussion on the remaining issues.

Attendees

George Abe	4247140@mcimail.com
J. Allard	jallard@microsoft.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Suhas Badve	badve@cup.hp.com
Uri Blumenthal	uri@watson.ibm.com
C. Douglas Brown	cdbrown@sandia.gov
Robert Ching	natadm!rching@uunet.uu.net
Mark Christenson	mgc@cray.com
Frank Coviello	
John Ioannidis	ji@cs.columbia.edu
James Keller	j.keller@sprint.com
Paulina Knibbe	knibbe@cisco.com
Kent Malave	kent@chang.austin.ibm.com
Mary Christine O'Connor	oconnor@interlan.com
Charles Perkins	perk@watson.ibm.com
Paul Sangster	sangster@ans.net
Aaron Schuman	schuman@sgi.com
Ron Sharp	rls@neptune.att.com
Jeremy Siegel	jzs@nsd.3com.com
Richard Slade	ricks@ssd.csd.harris.com
Michael St. Johns	stjohns@umd5.umd.edu
Dean Throop	throop@dg-rtp.dg.com
James Watt	jamesw@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Peter Williams	p.williams@uk.ac.ucl.cs

2.7.2 Common Authentication Technology (cat)

Charter

Chair(s):

John Linn, linn@albeit.lkg.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
Oct 1991	Update mechanism-independent Internet Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural descriptio
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 sch
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-02.txt, .ps>

"The Kerberos Network Authentication Service (V5)", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-01.txt, .ps>

"Generic Security Service API : C-bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-01.txt>

"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)

Recorded by John Linn and incorporating information from summary slides submitted by P. Rajaram.

The CAT Working Group met for one session at the July 1992 IETF. Primary discussion topics were:

- Document status review.
- Liaison requests from other standards organizations.
- Technical discussion of mechanism negotiation.

Document Status Review

The request to advance the base GSS-API Internet Draft to Proposed Standard was the first such request to be processed after the adoption of a security area policy requiring that specifications be submitted to independent reviewers before passing them on to the IESG. This pre-review process is currently in progress, with one set of comments so far received and forwarded to the CAT mailing list. Steve Crocker indicated that the pre-review will be complete by August 10th.

Some changes to the GSS-API C Bindings Internet Draft will be required, in response to accumulated comments and to track updates to the base specification. Cliff Neuman indicated that he expects to produce an updated version of the Kerberos V5 Internet Draft by the end of July.

Liaison Requests from other Standards Organizations

Subsequent to the San Diego IETF, John Linn had been approached by representatives from the X/Open Security Working Group and the POSIX Distributed Security Study Group, both of which indicated interest in adopting GSS-API within their standards processes. A short hardcopy paper, "Distributed Security Services Programme", was received from the X/Open group; copies were distributed at the meeting.

X/Open engages in activities including definition and implementation of interface test suites, and in establishment of agreements with end system vendors to encourage availability of adopted interfaces on a wide variety of platforms. These activities appear usefully complementary to IETF-CAT goals. In terms of process, Vint Cerf underscored the position that it was wholly appropriate for X/Open or other bodies to incorporate IETF specifications into their processes through reference to standards-track RFCs, but that change control to the specifications must remain within the IETF.

Piers McMahon attended the CAT session and has been a participant in the POSIX Distributed Security Study Group. He indicated his perception that POSIX would prefer to adopt GSS-API by way of X/Open rather than initiating a separate POSIX activity in this area. It was noted that authorization-oriented GSS-API extensions are under consideration in several forums, and that X/Open might also be a likely forum for standardization of such extensions.

John Linn accepted the action to coordinate with X/Open representatives to evolve a scope and action plan for liaison activities, and to report results back to the CAT Working Group. Piers plans to send a note reporting on relevant status of the POSIX study group, which was meeting in Chicago simultaneously with the Cambridge IETF.

Technical Discussion of Mechanism Negotiation

P. Rajaram indicated that he had been investigating approaches on mechanism selection and negotiation within the GSS-API framework, and led a discussion on the topic. He observed that GSS-API mechanism types correspond to groupings of authentication protocol per se, associated encryption algorithm, and associated hash function, and expressed a belief that three or four basic authentication protocols would likely exist in the marketplace but with many algorithm combinations. Further, different protocol/algorithm combinations would vary in their support for per-message confidentiality and integrity features and in their performance characteristics. It was observed, however, that divergent feature support within a single mechanism could result in cases where a given GSS-API implementation might not be able to determine the feature set supported by a desired peer. (Editor's Note: I believe that this could be reconciled by design of a mechanism in which peers declared their supported features to one another in exchanged tokens, without returning an indication of the features jointly supported until the context establishment sequence is complete.)

Raj suggested that the selection of appropriate mechanisms, and of feature sets within those mechanisms, should be refined based on application-stated requirements (e.g., integrity and confidentiality support, in addition to the service indicators already incorporated) and domain-administered policies (e.g., based on application and user names, initiator and target addresses, connection paths, time of day, ...). It was further suggested that GSS_Init_sec_context() be extended to allow an application to indicate a set of mechanism types as input to negotiation, rather than only a single type or a "default" specifier, and that a new Map_mechanisms() call accept a mechanism set and an indicator of application service requirements and return a subset of the input mechanism set suitable to satisfy the indicated requirements. Mechanism selection would be performed on an end-to-end basis, between peer applications, based on an intersection of the sets acceptable to both peers. This proposal led to an active discussion about the danger that use of negotiation to arbitrate among multiple mechanisms would generally result in the use of the weakest ("low water") alternative. It was suggested that each end system would appropriately maintain a database identifying (individually or through wildcards) the correct mechanism to use with particular peers. It was further suggested that target systems should be able to write ACLs selectively granting access based on the mechanism with which an initiator was

authenticated.

Given the level of controversy about the mechanism negotiation concept, no specification changes to aid its support were immediately adopted. Raj accepted an action to write a strawman proposal for a rendezvous scheme which would arbitrate mechanism selection in a secure fashion, and to distribute the result for mailing list discussion. Interface impacts would be revisited in the course of evaluating and evolving this proposal.

Attendees

Derek Atkins	warlord@thumper.bellcore.com
Tony Ballardie	a.ballardie@cs.ucl.ac.uk
Mark Baushke	mdb@cisco.com
Mark Bokhan	bokhan@abitok.enet.dec.com
Geetha Brown	geetha@decvax.dec.com
Vinton Cerf	vcerf@nri.reston.va.us
Stephen Crocker	crocker@tis.com
Michael DeAddio	deaddio@thumper.bellcore.com
Pierre Dupont	
Tom Farinelli	tcf@tyco.ncsc.mil
Barbara Fraser	byf@cert.org
Shari Galitzer	shari@shari.mitre.org
Gary Gaudet	gaudet@zk3.dec.com
Kenneth Grant	kgrant@oracle.com
Theodore Greene	
Neil Haller	nmh@thumper.bellcore.com
Mark Hollinger	myth@ssg.lkg.dec.com
William Huyck	
David Katinsky	dmk@rutgers.edu
Stephen Kent	kent@bbn.com
John Linn	linn@erlang.enet.dec.com
Ellen McDermott	emcd@osf.org
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Marjo Mercado	marjo@cup.hp.com
Clifford Neuman	bcn@isi.edu
Rakesh Patel	rapatel@hardees.rutgers.edu
Lars Poulsen	lars@cmc.com
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
Martin Schulman	mas@loyola.edu
Vincent Sgro	sgro@cs.rutgers.edu
Robert Shirey	shirey@mitre.org
Jeremy Siegel	jzs@nsd.3com.com
Sam Sjogren	sjogren@tgv.com
Theodore Ts'o	tytso@mit.edu

John Vollbrecht jrv@merit.edu
David Wang wang@xylogics.com
Peter Williams p.williams@uk.ac.ucl.cs

2.7.3 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.
Done	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-02.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-01.txt>

"Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", 08/22/1991, David Balenson <draft-ietf-pem-algorithms-01.txt>

"Privacy Enhancement for Internet Electronic Mail", 04/28/1992, <rfc1113-.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services", 09/01/1992, B. Kaliski <draft-ietf-pem-forms-01.txt>

Request For Comments:

RFC 1319 "The MD2 Message-Digest Algorithm"

RFC 1320 "The MD4 Message-Digest Algorithm"

RFC 1321 "The MD5 Message-Digest Algorithm"

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

Minutes of the Privacy-Enhanced Mail Working Group (PEM)

The PEM Working Group met once during the Boston IETF meeting. Various topics relative to the documents which will supersede RFCs 1113-1115 were discussed and resolved. The consensus of the attendees was that, when the changes discussed in this meeting have been executed, the resulting Internet Drafts will be ready for submission as Proposed Standard RFCs. The authors of RFCs 1113 and 1114 were present at the meeting and agreed to make the requisite changes by the end of July. The expectation is that the changes to RFC 1115 are very minor and also can be effected by the end of July. No modifications to the FORMS Internet Draft were identified, so that document also should be ready by the end of July.

The identified changes to be made to the documents are described below:

- Any certificate emitted by a PEM implementation, shall use the object identifier for RSA (see Annex G of X.509) to identify an RSA public key carried in the SubjectPublicKeyInfo field. However, PEM implementation shall accept both this object identifier and the "RSAEncryption" object identifier (from PKCS), in this field in "received" certificates, e.g., certificates in incoming PEM messages.

- The term "Internet Certificate Authority" will be changed to "Internet Policy Registration Authority" throughout RFC 1114bis.

- A new field, "Content-Domain" will be added to the PEM header. This field will be used to specify the type of content which has been protected by PEM and thus what "UA" should be invoked after PEM processing has been effected upon a received message. This provides a facility for future carriage of data type other than simple, RFC 822 mail, e.g., MIME, X.400, etc. This field must appear exactly once in the message, immediately after Proc-Type. The initial parameter value permitted for this field is "RFC-822" and will be so specified in RFC 1115bis.

The Working Group agreed to make integration of PEM with MIME the next major work item to be addressed on the PEM-DEV list and in future IETF meetings. It was agreed that this is a non-trivial task which will require careful study. There is a very strong desire from a variety of Internet community members to proceed with deployment of PEM for use with "vanilla" RFC 822 mail, hence this decision to make PEM-MIME integration a new work item rather than delaying progress of the current set of Internet Drafts. In recognition of this approach to accommodating MIME, RFC 1113bis will be revised to make explicit that it is a specification of core PEM functions plus use of PEM with RFC 822 mail, and that subsequent RFCs will address use of the core PEM functions with other mail systems, e.g., MIME, X.400, etc.

There was a discussion of issues related to deployment of PEM, summarized below:

- The PEM specification documents should all be ready for advancement by the end of July.
- TIS should be able to quickly accommodate the very minor change to the PEM header decided upon at this meeting, so availability of the reference implementation should not be substantially affected by the decisions at this meeting.
- TIS and RSADSI have executed the license agreement necessary for Internet distribution of PEM.
- The Internet Society is making preparation to instantiate its role as an Internet Policy Registration Authority. MIT has developed software that impelments the CRL service defined in FORMS and which needs to be operated by the IPRA. Steve Kent has provided a strawman algorithmic description of processing for the DN conflict resolution database, another database which the IPRA will operate.
- TIS and RSADSI have approached the IPRA about establishing PCAs, and RSADSI has recently distributed, via PEM-DEV, a candidate policy statement for a PERSONNA PCA.

It was suggested that an FYI on how prospective PEM users "get started" would be a useful document, once PEM deployment has progressed. This would augment the PCA policy statements which will be published as informational RFCs. It also was suggested that a PEM implementors' BOF might be scheduled for the next IETF, based on expectations for PEM deployment progress during the next 6 months.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Ashar Aziz	ashar.aziz@eng.sun.com
Mark Baushke	mdb@cisco.com
Uri Blumenthal	uri@watson.ibm.com
Mark Bokhan	bokhan@abitok.enet.dec.com
Luc Boulianne	lucb@cs.mcgill.ca
James Conklin	jbc@bitnic.educom.edu
Stephen Crocker	crocker@tis.com
Michael DeAddio	deaddio@thumper.bellcore.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Tom Farinelli	tcf@tyco.ncsc.mil
Barbara Fraser	byf@cert.org
Shari Galitzer	shari@shari.mitre.org
Gary Gaudet	gaudet@zk3.dec.com

Neil Haller	nmh@thumper.bellcore.com
Stephen Kent	kent@bbn.com
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
John Linn	linn@erlang.enet.dec.com
Kent Malave	kent@chang.austin.ibm.com
Ellen McDermott	emcd@osf.org
Clifford Neuman	bcn@isi.edu
Marshall Rose	mrose@dbc.mtview.ca.us
Paul Sangster	sangster@ans.net
Jeffrey Schiller	jis@mit.edu
Richard Schmalgemeier	rgs@merit.edu
Einar Stefferud	stef@nma.com
Theodore Ts'o	tytso@mit.edu
Huyen Vu	vi@polaris.disa.mil
Sandro Wallach	sandro@elf.com
David Wang	wang@xylogics.com
Charles Watt	watt@sware.com
Peter Williams	p.williams@uk.ac.ucl.cs

2.7.4 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.
Done	Submit specification to IESG for consideration as a Proposed Standard.
Dec 1991	Submit specification to IESG for consideration as a Draft Standard.
Dec 1993	Submit specification to IESG for consideration as a Standard.

Request For Comments:

RFC 1351	"SNMP Administrative Model"
RFC 1352	"SNMP Security Protocols"
RFC 1353	"Definitions of Managed Objects for Administration of SNMP Parties"

2.7.5 TCP Client Identity Protocol (ident)

Charter

Chair(s):

Mike St. Johns, stjohns@umd5.umd.edu

Mailing Lists:

General Discussion: ident@nri.reston.va.us
To Subscribe: ident-request@nri.reston.va.us
Archive: nri.reston.va.us:ietf.mailing.lists/ident.mailing.list

Description of Working Group:

The TCP Client Identity Protocol Working Group is chartered to define a protocol for returning the identity of the user initiating a TCP connection. When a client on host A initiates a TCP connection to host B, host B may query a server on host A to determine the identity of the client on host A. The primary purpose of this protocol is to record the identity of requesters initiating a connection.

This work is a clarification and standardization of the Experimental Protocol currently published as RFC 931.

Goals and Milestones:

Apr 1992	Post an Internet Draft of the revised RFC 931 Identity Server Protocol.
Jun 1992	Submit the Identity Server Protocol to the IESG for consideration as a Proposed Standard.
Ongoing	Review implementations, and resolve outstanding issues in preparation for Draft Standard.

Internet Drafts:

"Ident MIB", 04/13/1992, Michael St. Johns, Marshall Rose <draft-ietf-ident-mib-03.txt>

"Identification Server", 06/02/1992, Mike StJohns <draft-ietf-ident-idserver-02.txt>

CURRENT MEETING REPORT

Reported by Mike St. Johns/DOD and Dave Borman/Cray Research

Minutes of the TCP Client Identity Protocol Working Group (IDENT)

Discussion Items

- Security Section
- Format of User IDs
- Character Set
- Error Response
- MIB

Security Section

The security section has been argued/discussed to death on the mailing list - the current text represents what the Chair considers a reasonable compromise. The Chair did not want to reopen the section again.

After a little discussion about generalities (e.g., Is the section too big?; Is the section too small?), Marshall Rose stated that "Finger has a large security section, it's a fact of life that with documents like this a longer security section is needed." Steve Crocker said that "I read this from scratch - thought about dropping last two paragraphs, but then decided they were important." The Chair polled the room on three questions:

1. Section too strict: 0
2. Section not strict enough: 2
3. OK as is: 13
4. No opinion: 5

(Note that there was some overlap between "not strict enough" and "OK" - the "too strict" people were willing to accept as stands). The Chair said there was not enough heartburn to warrant reopening section and was not shouted down, so section is CLOSED and will stand as it is currently.

At the request of Marshall Rose, the Group discussed the MIB document out of turn.

There was a brief discussion on how MIB document relates to Ident protocol. We noted that the only critical overlap was the security section. As we had closed the security section, without objection, the Chair declined to pull the MIB document back from standards submission. Please note that objectors may send their comments to the IESG during the normal two week comment period once the document is announced.

Format of Userids

There was a lot of very good discussion here on character length and format: Why limit UNIX to 8 characters? Should we get rid of OPSYS field? The character set information is useful and the rest should be arbitrary? MIME is specifying US-ASCII instead of NVT-ASCII.

The final result of the discussion was to redesignate OCTET as a character set indicator; to remove the syntax implications from the OPSYS ID [a "real" operating system identifier implies a "real" user identifier, but does not indicate any specific syntax of the user identifier]; and to make US-ASCII the default character set vs NVT-ASCII.

Random Interjections

There was a brief discussion on the feasibility of using UDP for transport but the general consensus was that it was not a good idea. By using UDP as the transport, it would be very easy to spoof the response - even more so than is possible in TCP.

The point was make that there needs to be something in overview about the client shutting down the connection if it gets no response. The Chair accepted this and took for action.

Stever Crocker asked if the server is allowed to respond on a selected basis, and is the server allowed to respond with something not directly interpretable? The Chair deferred the first part for later discussion, but pointed Steve at the "OCTET" operating system identifier and the "OTHER" character set identifier for the second part.

Steve also indicated a problem with the "Query/Response" section; he got lost reading this. He suggested "Port on Client Machine/Port on Server Machine" vs. "local/remote". The Chair accepted this and will edit it.

Character Sets

This section remains open for one go-around on the items suggested above.

Error Responses

There were two proposals:

1. Allowing the server to return "NONE" for machines like MS/DOS or others which don't have the concept of a userid. This was rejected on the basis we'd have to reserve "NONE" and no one could use it as a userid. Systems like this should use the "NO-USER" error code instead.

2. Allowing the server to return "HIDDEN-USER" as an error code. The code would mean the system had valid user information for this query, but refused to return it at the request of the user. The consensus was to accept it.

Attendees

Steve Alexander	stevea@i88.isc.com
Mark Baushke	mdb@cisco.com
David Borman	dab@cray.com
Luc Boulianne	lucb@cs.mcgill.ca
Stephen Crocker	crocker@tis.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Chas DiFatta	chas@sei.cmu.edu
Tom Farinelli	tcf@tyco.ncsc.mil
Barbara Fraser	byf@cert.org
Ned Freed	ned@innosoft.com
Cliff Frost	cliff@cmsa.berkeley.edu
Pria Graves	priag@nsd.3com.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
April Richstein	abm@tycho.ncsc.mil
Marshall Rose	mrose@dbc.mtview.ca.us
Robert Shirey	shirey@mitre.org
Jeremy Siegel	jzs@nsd.3com.com
Michael St. Johns	stjohns@umd5.umd.edu
Theodore Ts'o	tytso@mit.edu
John Wagner	jwagner@princeton.edu
Walter Wimer	walter.wimer@andrew.cmu.edu

2.8 Transport and Services Area

Director(s):

- David Borman: dab@cray.com

Area Summary reported by Dave Borman/Cray Research)

Four working groups and one BOF met at the Boston IETF meeting.

TCP/UDP over CLNP-based Networks BOF (TUBA)

The TCP/UDP, (or TCP/UPD with Big Addresses) BOF discussed the TUBA proposal for dealing with the growth of the Internet. The proposal was published as RFC 1347. It uses a (possibly modified) CLNP header for the new IP. Some of the advantages discussed were that it would be able to support mobile hosts and would be able to use existing CLNP routing protocols.

Audio/Video Transport Working Group (AVT)

The AVT Working Group is looking at two candidate audio transport protocols: one from Steve Casner, and one from Van Jacobson. Steve's protocol is the second-cut straw-man protocol, based on discussions at the San Diego meeting on the first-cut straw-man protocol. Van's protocol is the one used by the "vat" program. The Working Group will define one audio transport protocol. It has not been decided yet, but the ground work has been laid. A new draft will be written and discussed via e-mail, with hope of having an Internet Draft available by November.

Domain Name Systems Working Group (DNS)

The main item on the Agenda of the DNS Working Group was discussion of the DNS MIB. It is moving right along, changes discussed at the meeting will be added in. Time was also spent discussing policy statements for the .COM and .EDU domains, which currently don't exist. RFC1348, NSAP Resource Records, was also briefly discussed.

Services Location Protocol Working Group (SVRLOC)

The Services Location Protocol Working Group is busy trying to put together a document. Current discussions are based on a proposal from FTP Software. It was decided to defer the RPC issue ("Should this protocol run on top of RPC, and if so, which one?"). One area that is lacking in the current discussion is a definition of a the values and format for the data that will be transmitted by this protocol. There has been an attempt to separate the protocol definition from the format/values of the data; but it was recognized that there will need to be a base set of definitions in the document.

Trusted Network File Sytem Working Group (TNFS)

The TNFS Working Group is a joint Working Group with TSIG. A plenary presentation was made to the IETF on Thursday by Fred Glover, the Chair of the Working Group. The document is on its last review. As soon as the changes agreed upon at the meeting are incorporated, the document will be submitted to the IESG for publication as a Proposed Standard.

For more information on each of these sessions, please refer to the Minutes for the individual working groups.

CURRENT MEETING REPORT

Minutes of the TCP/UDP over CLNP-addressed Networks BOF (TUBA)

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

Attendees

J. Allard	jallard@microsoft.com
William Babson	bill@penril.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
William Barns	barns@gateway.mitre.org
Lou Berger	lberger@penril.com
David Bolen	db31@nis.ans.net
David Borman	dab@cray.com
Robert Braden	braden@isi.edu
Scott Brim	swb@cornell.edu
Geetha Brown	geetha@decvax.dec.com
Theodore Brunner	tob@thumper.bellcore.com
Alan Bryenton	bryenton@bnr.ca
Jeffrey Burgan	jeff@nsipo.nasa.gov
Randy Butler	rbutler@ncsa.uiuc.edu
Ross Callon	callon@bigfut.enet.dec.com
Vinton Cerf	vcerf@nri.reston.va.us
John Chang	jrc@uswest.com
Robert Ching	natadm!rching@uunet.uu.net
James Cobban	jcobban@bnr.ca
Richard Colella	colella@osi.ncsl.nist.gov
Michael Collins	collins@ccc.nersc.gov
Michael Conn	0043874510@mcimail.com
John Curran	jcurran@bbn.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Irina Dobrina	irina@ftp.com
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@mrbill.canet.ca
Peter Ford	peter@lanl.gov
Osten Franberg	euaokf@eua.ericsson.se
Kathryn Fuller	kathryn@codex.com
Elliot Gerberg	gerberg@took.enet.dec.com
Terry Gray	gray@cac.washington.edu
Phillip Gross	pgross@nis.ans.net
William Haggerty	haggerty@ctron.com
Jack Hahn	hahn@sura.net

Tony Hain	hain@nersc.gov
James Halpin	halpin@turkey.sw.stratus.com
Robert Hinden	hinden@eng.sun.com
Kathleen Huber	khuber@bbn.com
Steven Hubert	hubert@cac.washington.edu
Bob Jeckell	rrj@3com.com
Matthew Jonson	jonson@server.af.mil
Miriam Kadansky	mckadansky@eng.xyplex.com
Scott Kaplan	scott@ftp.com
Frank Kastenholz	kasten@ftp.com
James Keller	j.keller@sprint.com
Tom Kessler	kessler@eng.sun.com
Kenneth Key	key@cs.utk.edu
Alan Kullberg	akullber@bbn.com
Anthony Lauck	lauck@tl.enet.dec.com
Tony Li	tli@cisco.com
Gary Malkin	gmalkin@xylogics.com
Allison Mankin	mankin@gateway.mitre.org
Bill Manning	bmanning@rice.edu
Ed Menze	menze@cs.arizona.edu
Douglas Miller	dmm@telebit.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Donald Morris	morris@ucar.edu
John Moy	jmoy@proteon.com
David Piscitello	dave@sabre.bellcore.com
Lars Poulsen	lars@cmc.com
Russell Pretty	pretty@bnr.ca
David Reiser	dsr@codex.prds.cdx.mot.com
April Richstein	abm@tycho.ncsc.mil
Anil Rijsinghani	anil@levers.enet.dec.com
Rich Rosenbaum	rosenbaum@lkg.dec.com
Anthony Rutkowski	amr@sprint.com
Henry Sanders	henrysa@microsoft.com
Jeremy Siegel	jzs@nsd.3com.com
Keith Sklower	sklower@cs.berkeley.edu
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@sceng.ub.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Sally Tarquinio	sallyt@gateway.mitre.org
Marten Terpstra	terpstra@ripe.net
Sylvia Tom	tom_sylvia@tandem.com
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Justin Walker	justin@apple.com

David Wang	wang@xylogics.com
William Warner	warner@ohio.gov
Liming Wei	liming.wei@eng.sun.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Steven Winnett	swinnett@bbn.com
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com
Joseph Zur	zur@fibhaifa.com

2.8.1 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: nic.es.net:[anonymous.ietf.rem-conf]/av-transport-archiv

Description of Working Group:

The Audio/Video Transport WG was formed to specify experimental protocols for real-time transmission of audio and video over UDP and IP multicast. The focus of this group is near term and its purpose is to integrate and coordinate the current AV transport efforts of existing research activities. No standards-track protocols are expected to be produced because UDP transmission of audio and video is only sufficient for small-scale experiments over fast portions of the Internet. However, the transport protocols produced by this WG should be useful on a larger scale in the future in conjunction with additional protocols to access network-level resource management mechanisms. Those mechanisms, research efforts now, will provide low-delay service and 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 also the subject of current research. It is expected that standards-track protocols integrating transport, resource management, and connection management will be the result of later working group efforts.

The AVT WG may design independent protocols specific to each medium, or a common, lightweight, real-time transport protocol may be extracted. Sequencing of packets and synchronization among streams are important functions, so one issue is the form of timestamps and/or sequence numbers to be used. The WG will not focus on compression or coding algorithms which are domain of higher layers.

Goals and Milestones:

Done	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 vide

Done	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
Done	Review contributions of existing protocols, and discuss which features should be included and tradeoffs of different methods. Make writing assignments for first-draft documents.
Oct 1992	Post an Internet Draft of the lightweight video transport protocol.
Nov 1992	Broadcast the working group meeting using implementations of the draft audio/video transport protocols.
Dec 1992	Submit to the IESG the Audio/Video Transport protocol for Proposed Standard Status.

CURRENT MEETING REPORT

Reported by Steve Casner/USC-ISI

Minutes of the Audio/Video Transport Working Group (AVT)

The AVT Working Group met during three separate sessions. The first session began with presentations of candidate protocols for real-time audio/video transport, followed by a lively discussion of the differences among the candidates and the underlying questions implied by those differences. The discussion resumed in the second session and part of the third, followed by live demonstrations of experimental packet audio and video programs.

As part of the second IETF "audiocast", live audio and video from all three sessions was transmitted via UDP and IP multicast to participants at a number of locations around the world. At least two remote participants made multiple contributions to the Working Group discussion.

1. Presentations of Candidate Protocols

Steve Casner began with a quick review of the descriptions of the Network Voice Protocol (NVP-II) data packet format and the first-cut strawman protocol from the San Diego meeting, then presented a second-cut strawman based on the discussions in San Diego. The data packet header contains the following fields:

- Timestamp (16 bits of seconds + 16-bit fraction)
- Packet Sequence Number (16 bits)
- Flow Identifier (8 bits)
- Options Length (8 bits)
- Options

Since Van Jacobsen could not attend, Steve also described the protocol used by the VAT audio program, based on a protocol description sent by Van to the rem-conf list. The data packet header format is:

- Protocol Version (2 bits)
- Number of Site Identifiers to follow header (6 bits)
- Start-of-Talkspurt Flag (1 bit)
- Audio Format/Encoding (5 bits)
- Conference Identifier (16 bits)
- Timestamp (32-bit audio sample counter)
- Site Identifiers (0 to 63; 32 bits each)

Both of these data packet formats depend on a session/control protocol to carry information that is not required in every data packet. Henning Schulzrinne described the extensions to the VAT session protocol used in his NEVOT audio program, in particular the periodic

transmission of the sender's state (the current time and how many samples have been transmitted) to enable measurement of loss at the receiver.

Simon Hackett gave an impromptu overview of his Multimedia Data Switch (MMDS) application and protocol. For purposes of experimentation, Simon chose to use large headers including a variety of fields to make the data self-describing. He also continues to send packet headers during silence as a keep-alive, but just omits the data to reduce the bandwidth.

See section 5 below for references on these protocols.

2. Discussion of Protocol Differences

The goal of the discussion was to identify the issues that must be resolved in order to produce a draft protocol. The primary ones were:

- Timestamp format, media sample clock or real time
- Sequence number versus start-of-talkspurt flag
- What multiplexing is required beyond address+port
- Whether or not to indicate encoding format in data packets

The first two issues underlie a key question for the Working Group, namely whether we should define one real-time transport protocol or multiple application-specific protocols. The rough concensus was for the former, but this may conflict with ease of implementation.

The Working Group discussed timestamp formats at the last meeting and this one, but the issue is still not finally decided. For purposes of synchronization among multiple media sources, the only practical means is to relate all streams to real time (synchronized time of day). This would be simplified if the timestamps are in real time, but the implementation of audio buffering is much easier with an audio sample clock timestamp. The timestamp format could be converted either at the sender or receiver; what's needed is a detailed analysis of the tradeoffs.

The strawman protocols propose a packet sequence number in addition to the timestamp in order to differentiate lost packets from packets not sent during silence. The VAT protocol uses a flag on the first packet of a talkspurt because packet mis-ordering makes the sequence number hard to use. On the other hand, a sequence number may be required for video applications that don't have talkspurts but require multiple packets per frame all with the same timestamp.

The Flow ID in the strawman protocol serves two purposes: it provides multiplexing of multiple streams (e.g., audio and video) from the same source on one IP multicast address and port, and it allows for different encodings to be used, with each Flow ID bound to an encoding descriptor using the session/control protocol. As defined, the VAT protocol includes an explicit encoding format field in the data packet, but the Working Group deemed 5 bits to be too small a number. The VAT encoding values could also be bound a dynamic

set of encoding descriptors using a control protocol.

The VAT Conference ID discriminates among conferences in case of a collision in random IP multicast address allocation and because many BSD derived systems don't allow discriminating on the multicast destination address. The strawman assumes a repair of the BSD deficiency (which seems feasible at this time for multicast capable systems) and assumes some other method to avoid address collisions.

3. Completeness and Compatibility with Connection Management

In addition to resolving differences among the protocol proposals, we must consider whether the protocols are sufficiently complete. Unlike the audio and video conferencing applications, distributed simulation and PBX trunking may require aggregation of multiple frames of data into a single packet. If the frames can all share the same header information, then aggregation can be consigned to the next layer up; if not then some additional encapsulating mechanism would be required. We did not consider this further.

Another extension would be flow control. In previous Working Group discussions, it has been assumed that network resource management mechanisms and protocols would be available to allow real-time applications to avoid congestion. Christian Huitema pointed out that at least over some paths we will probably need a feedback mechanism to allow adjustable codecs to accommodate congestion. The Group was unsure whether an application-independent feedback mechanism could be defined. Christian is to write a specification as a starting point.

This Working Group's low-level protocol must also be compatible with higher-level connection management protocols such as those under discussion in the Remote Conferencing Architecture BOF. Provision of encoding format selections from a conference directory server seems straightforward. However, the server must also have a means to acquire an IP multicast address. Lixia Zhang suggested (remotely!) that we really should consider a distributed system of servers to hand out globally unique IP multicast addresses; this capability will be needed by several groups considering multicast, not just ours.

4. Software Encoding and Enumeration

The real-time transport protocol should be independent of the media encoding algorithms and formats that belong to the next higher layer except that the format must be identified by the lower layer. However, in keeping with the Working Group goal to foster interoperation and experimentation with packet audio and video, it may be valuable to agree on some (perhaps low performance) software compression techniques for use until hardware is generally available. This suggests that some of the encoding formats we need to identify will be non-standard and hence not included in any standard enumeration.

The Working Group feels a strong need to pick up a task that has been deferred by others, to define an IANA-managed enumeration or naming convention for audio and video encoding

algorithms to enable interoperation. The enumeration should not be part of the protocol itself, but the protocol must provide the space to carry the encoding identification. There was substantial discussion of numeric vs text/parametric identification of formats. This issue was not resolved.

The third Working Group session was concluded with descriptions and demonstrations of the software encoding algorithms developed by Working Group participants. Paul Milazzo gave an update on the protocol for the BBN Desktop Video Conference program which was used to multicast packet video from IETF. Christian Huitema showed the INRIA H.261 video compression software. Hans Eriksson described the packet audio and video experiments at SICS.

5. Further Discussion

While several issues were not resolved, we laid out the considerations for each choice well enough to guide the design of a complete set of consistent choices as the first draft protocol from this Group. Our (revised) goal is to have an Internet Draft protocol submitted by November. Further discussion by email will be required to make this happen.

During the IETF meeting, some notes from the first session, including a description of the strawman and VAT protocols, was sent to the rem-conf list. It should be in the archive, or may be requested from casner@isi.edu. A message from last March on MMDS is also available.

An extensive summary of the issues and a protocol recommendation has been prepared by Henning Schulzrinne and is available from:

```
gaia.cs.umass.edu:~ftp/pub/rtp/rtp.ps
```

This working paper will be made an Internet Draft for wider distribution.

Thanks to Eve Schooler, Henning Schulzrinne and Christian Huitema for taking the notes from which these Minutes were prepared.

Attendees

George Abe	abe@infonet.com
J. Allard	jallard@microsoft.com
John Batzer	
Lou Berger	lberger@penril.com
James Berry	beri@sandia.llnl.gov
Luc Boulianne	lucb@cs.mcgill.ca
Scott Brim	swb@cornell.edu

Alan Bryenton	bryenton@bnr.ca
Randy Butler	rbutler@ncsa.uiuc.edu
Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@concert.net
Andrew Cherenson	arc@sgi.com
Robert Clements	clements@bbn.com
Michael Collins	collins@ccc.nersc.gov
Steve Deering	deering@parc.xerox.com
Tony DeSimone	tds@hoserve.att.com
Jack Drescher	drescher@concert.net
Hans Eriksson	hans@sics.se
Julio Escobar	jescobar@bbn.com
Roger Fajman	raf@cu.nih.gov
Margaret Forsythe	mrf@ftp.com
Osten Franberg	euaokf@eua.ericsson.se
Ron Frederick	frederick@parc.xerox.com
Jerry Friesen	jafries@sandia.llnl.gov
Robert Gilligan	Bob.Gilligan@eng.sun.com
Simon Hackett	simon@internode.com.au
Robert Hagens	hagens@ans.net
Christian Huitema	christian.huitema@sophia.inria.fr
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Padma Krishnaswamy	kri@sabre.bellcore.com
Matt Mathis	mathis@a.psc.edu
Cindy Mazza	
Donald Merritt	don@brl.mil
Paul Milazzo	milazzo@bbn.com
Robert Mines	rfm@sandia.llnl.gov
Donald Morris	morris@ucar.edu
Ari Ollikainen	ari@es.net
Roger Osmond	bytex!rfo@uunet.uu.net
Larry Palmer	lp@decvax.dec.com
Michael Powell	mdpowel@pacbell.com
Russell Pretty	pretty@bnr.ca
K. K. Ramakrishnan	rama@erlang.enet.dec.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Allan Rubens	acr@merit.edu
Henry Sanders	henrysa@microsoft.com
Eve Schooler	schooler@isi.edu
Koichiro Seto	seto@hitachi-cable.co.jp
Vincent Sgro	sgro@cs.rutgers.edu
Louis Steinberg	louiss@vnet.ibm.com
Terrance Sullivan	terrys@newbridge.com
Sally Tarquinio	sallyt@gateway.mitre.org

Claudio Topolcic	topolcic@nri.reston.va.us
Mark Uhrmacher	maui@tabasco.lcs.mit.edu
Andrew Veitch	aveitch@bbn.com
John Vollbrecht	jrv@merit.edu
David Waitzman	djw@bbn.com
Sandro Wallach	sandro@elf.com
Abel Weinrib	abel@bellcore.com
Rick Wilder	wilder@ans.net
Walter Wimer	walter.wimer@andrew.cmu.edu
Linda Winkler	lwinkler@anl.gov
Jeff Wong	jaw@io.att.com
Richard Woundy	rwoundy@rhqvm21.vnet.ibm.com
John Wroclawski	jtw@lcs.mit.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

2.8.2 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.

2.8.3 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.

Internet Drafts:

"DNS MIB Extensions", 03/05/1992, Jon Saperia <draft-ietf-dns-mibext-02.txt, .ps>

CURRENT MEETING REPORT

Reported by Win Treese/DEC

Minutes of the Domain Name System Working Group (DNS)

Agenda

- The proposed DNS MIB. (Jon Saperia, DEC)
- The NSAP resource record proposal. (RFC1348)
- Possible guidelines for the assignment of names in top-level domains. (Mike St. Johns, Dept. of Defense)

DNS MIB

Jon Saperia handed out copies of the Internet Draft DNS MIB, which had previously been sent to the namedroppers mailing list. Jon had made several changes based on comments from the meeting in San Diego. There were a few minor comments, which Jon agreed to incorporate.

There was some discussion of whether or not SNMP should be used to update authoritative date in a server. Jon Saperia and Rob Austein pointed out that this is a hard problem. Adding a new host, for example, involves adding new address and pointer records, and possibly HINFO and MX records. Should the server store the updated information in its permanent database? How are all replicas updated? Is the update protocol atomic for all replicas? Is it atomic at a single replica? How are changes authenticated and authorized?

The Group awaits a proposal on updates; the MIB will be moved forward as it is now.

NSAP Resource Records

Bill Manning presented RFC1348, which proposes new resource records for NSAPs to support RFC1347 (the TUBA proposal). There were some minor suggestions for modifications to the proposal. We also discussed having a wildcard address query for any type of addresses. The Group concluded that it was best not to create a new class for extended addresses of any type, since the class mechanisms are not well-understood in practice.

Naming Guidelines

Mike St. Johns proposed establishing some naming guidelines for the top-level domains (e.g., COM, EDU), because they are growing quite rapidly, and may eventually overwhelm the capabilities of both humans and computers to use them. There was quite a bit of discussion on this issue, but very little was resolved. A detailed proposal of guidelines is probably needed to further the discussion.

Attendees

Robert Austein	sra@epilogue.com
David Borman	dab@cray.com
Lida Carrier	lida@apple.com
Walter Lazear	lazear@gateway.mitre.org
Bill Manning	bmanning@rice.edu
Michael Patton	map@lcs.mit.edu
Mike St. Johns	stjohns@umd5.umd.edu
Jon Saperia	saperia@tcpjon.ogo.dec.com
Win Treese	treese@crl.dec.com

2.8.4 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)

RES_Type is a well defined list of attributes registered with the Assigned Numbering authority.

Decisions arrived at:

The protocol does not have to be RPC based. Scott Kaplan is reviewing the benefit of RPC as well as the performance issues.

There should be some sort of standard data representation for query data. XDR and ASN.1 were some of the proposals.

Several examples of the RES_Type and related attribute types should be given in the document to give guidance to developers of services to be located.

Some attributes are site specific (e.g., the services physical location, the organization the service belong in, etc.) these attributes should have a way to be delineated as site specific so that they will not collide with later defined types.

Scott Kaplan and John Veizades are working on a specification to be presented at the November IETF.

Attendees

J. Allard	jallard@microsoft.com
David Borman	dab@cray.com
Rich Brown	Richard.E.Brown@Dartmouth.edu
Andrew Cherenson	arc@sgi.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Margaret Forsythe	mrf@ftp.com
Atul Garg	agarg@synoptics.com
Simon Hackett	simon@internode.com.au
Michael Khalandovsky	mlk@ftp.com
Stev Knowles	stev@ftp.com
Joshua Littlefield	josh@cayman.com
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Jeremy Siegel	jzs@nsd.3com.com
John Veizades	veizades@apple.com

2.8.5 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 ezeinterest@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:

Done Review the TCP Extended Window Size proposal from the IRSG End to End Research Group and if acceptable, recommend it for standards status.

Request For Comments:

RFC 1323 "TCP Extensions for High Performance"

2.8.6 Trusted Network File Systems (tnfs)

Charter

Chair(s):

Fred Glover, fglover@zk3.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 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, auditing, documentation, and other items as identified in the Trusted Computer System Evaluation Criteria (TCSEC) and Compartmented 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 context into the appropriate server security policies which address file access. In addition, file security attributes are to be returned with each TNFS procedure call. Otherwise, the NFS V2 protocol remains essentially unchanged.

The Trusted System Interoperability Group (TSIG) has already developed a specification which defines a set of MLS extensions 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-01.txt>

INTERIM MEETING REPORT

Reported by Fred Glover/DEC

Minutes of the Trusted Network File Systems Working Group (TNFS)

May 1992

Agenda

- Reviewed the recent modifications to the TNFS document (TNFS-001.2.02)
- Reviewed the TKM document (TNFS-006.01.01)
- Reviewed implementation status and issues
- Discussed Lock Manager impact
- Discussed TNFS auditable events

TNFS Document Review

The IETF TNFS document has been available for comments in the IETF Draft Directory and TNFS archive since July, 1991. During the May meeting, the Working Group made "final" wordsmithing modifications, and edits to the document. This concludes the work planned for the TNFS specification, and the next steps are to transition the document to Proposed Standard. Trusted Systems technology providers are encouraged to commence implementation based upon the current TNFS draft.

Final updates to the TNFS document include:

- Add vendor specific token to directory response structure, set label procedure arguments.
- Change the arguments to the MLD procedure to reference operations rather than flags.
- Add explicit indication within ACCESS and file open sections to identify conditions in which a client caching security attributes must revalidate the rights of a given client application (i.e., by calling ACCESS).
- Add rationale indicating why security attribute tokens are not required to be included in the read directory response structure.

We spent some additional time discussing auditing. Conclusions include:

- Some environments may require that both client and server auditing is enabled in order to ensure full audit of events such as:
  - First read (server audits; may need to check audit id and XID in order to identify actual issuing PID if that is required).

  - Floating of objects (server only; client can't see).
  - Server override of privileges (server only; client can't see).
- The transaction ID is a possible "key" which can be used to correlate a client request with the server side procedure.

An updated document will be placed in the IETF and TNFS archives after an "email" review of the updates has been completed. In addition, Fred will contact our IETF Area Director, Dave Borman, to understand our "next steps" in the standardization process.

Token Manager Review

The TKM document was also reviewed. It now contains an attribute to token procedure. The following updates will be made to the document:

- Editorial (wordsmithing).
- Length field added to string arguments.

The updated TKM document will be placed into the IETF Draft Directory (informational RFC) and the TSIG TNFS archive.

We have been reviewing both the TKM and MAC6 token mapping proposals. Our current recommendation is for each of these to be submitted as IETF informational RFCs. One of these would be identified to become the actual standard, once all of the IETF/TSIG token mapping requirements were understood and were accommodated by at least one of these proposals. We recommended that a new working group be formed which would assume the responsibility of collecting the requirements, reviewing all of the proposals, and making recommendations.

Implementation Status, Issues

The Working Group reviewed the progress of current implementation efforts. Two implementations are very "close" to conformance with the current version of the specification. We discussed testing possibilities for the end of this year. We have already identified a test plan, a set of "non-mapped" security attributes for testing, and a modified test routine to be used with the Connectathon test suite. An update of the tnfs.h file, describing all of the TNFS procedures and data structures, will be placed into the TSIG/IETF archive to facilitate development of additional implementations.

Lock Manager Update

Charlie Watt recently completed a review of the NFS lock manager and suggests that no changes appear to be required to the lock manager to work in the TNFS environment. This confirmed the results of an earlier Working Group discussion and closes an action item. Thanks Charlie!

TNFS Auditable Events

We started the discussion of TNFS auditable events at this meeting. Mark Saake will be developing a document which describes this area. Conclusions reached at this meeting:

- The TNFS Group will focus on server side (i.e., protocol procedure) auditable events; we will expect that POSIX will identify the client side (i.e., application, API based) events and formats.
- When the server receives a TNFS request, it can identify:
  - Host address (and thus host)
  - File handle
  - Export structure
  - Procedure number
  - Version number
  - Credential information (ID info); result of subsequent authorization check (i.e., pass or fail).
  - Log entry and exit status of each called TNFS server procedure.

Next Meeting

The TNFS Group will plan to meet jointly with IETF and TSIG at the July meeting in Boston. At that meeting, we plan to:

- Present a summary to interested IETF attendees during a designated two hour time slot.
- Review the "final" version of the TNFS documents (updated documents placed into the TNFS archive and IETF Drafts Directory: Fred, Fran, Carl, Ali).
- Review the interoperability test opportunities, plans (all).
- Review NFS test suite extension for TNFS (Fran).
- Identify conforming implementations to support our request to transition our TNFS document (all).
- Review identification of auditable TNFS events (Mark).
- Place "tnfs.h", test plan, test attributes into TNFS archive (Fred).

The July meeting is planned for the 13th-17th at the Hyatt Regency in Cambridge, Massachusetts.

Attendees

Lida Carrier	lida@apple.com
Fran Fadden	fran@decvax.dec.com
Jonathon Fraser	
Fred Glover	fglover@zk3.dec.com
Ali Gohshan	
Brian Hardy	
Mark Saake	saake@netcom.netcom.com
Carl Smith	cs@eng.sun.com
T.T. Tao	
Charles Watt	watt@sware.com

CURRENT MEETING REPORT

Reported by Fred Glover/DEC

Minutes of the Trusted Network File Systems Working Group (TNFS)

The TNFS Working Group met for three sessions during the IETF July meeting. We:

- Presented an overview and status of the TNFS Working Group during the first session for those participating for the first time.
- Reviewed the final modifications to the TNFS document (TNFS-001.2.02) during the second session.
- Reviewed implementation status and issues during the third session.

In addition, Fred presented an overview of the TNFS Working Group achievements and current status during the Thursday afternoon IETF plenary.

TNFS Working Group Status and Summary

During the first TNFS session, Fred presented an overview of the TNFS Working Group effort, current status, and current set of documentation. Dave Borman, our Area Director, was present, and helped us position our documents. We are presenting the TNFS document as a candidate for proposed standard RFC. We will plan to present the token mapping document as a (new) prototype RFC, and the implementation and administrative documents as informational RFCs. The other TNFS Working Group documents, (interoperability test plan, tnfs.h, ...) will not be submitted as RFCs, but will continue to be maintained within the TNFS archive.

TNFS Document Review

The IETF TNFS document has been available for comments in the IETF Draft directory and TNFS archive since July, 1991. During the July meeting, the TNFS Working Group reviewed the current TNFS document for any final edits prior to requesting the transition of its status to proposed RFC draft:

- We agreed to add one additional flag to the access protocol operation, STAT, which would request permission to "stat" the requested file. Since the full set of attributes are returned with most all procedure calls, there is a question of which attributes a given client should be able to see from the set of cached attributes. By adding STAT to our access procedure, and by maintaining per process cached attributes, the client will be able to present an appropriate set of attributes to a requesting application from the cache.
- We agreed to change the name of the SETLABEL procedure to the SETNAMELA-

BEL procedure. This new procedure name reflects the intended use of the procedure (i.e., to modify the label of the file name only; the SETATTR procedure is used to modify the label of the file).

The document will be updated to reflect these changes, and will be placed in both the TNFS archive and IETF draft directory. Fred will then contact our IETF Area Director, Dave Borman, to indicate that the document is ready to be reviewed for Proposed Standard status.

Implementation Status, Issues

The Working Group reviewed the progress of current implementation efforts. One implementation now conforms to the current TNFS specification, and several others are very close. We reviewed TNFS test plan, and discussed several upcoming opportunities for interoperability testing:

- The September TSIG meeting hosted by Cray Research: this may be too early; we will "poll" via email during late August to determine if this is a possibility.
- A November west coast site: Carl Smith and Fred Glover will investigate possible test sites; November was identified as a more realistic timeframe for the completion of the other TNFS implementations.
- The 1993 Connectathon: no date exists for this at present, but this event is usually held during February or March; many of the TNFS Working Group member's companies already participate in Connectathon, so this is another good possibility.
- The spring TSIG meeting: it would be convenient from a testing perspective if this meeting were held during/near the time of Connectathon, but this would still be a good possibility with respect to timeframe.

Interoperability testing is an important milestone in supporting our goal to promote the TNFS specification to Draft Standard. Our test plan, a documented set of "non-mapped" security attributes for testing, and an update of the tnfs.h file, describing all of the TNFS procedures and data structures are available in the TSIG/IETF archive to facilitate the development of additional implementations.

Next Meeting

Since the TNFS document is now being proposed for promotion to Proposed Standard, the TNFS Working Group will plan to meet next in conjunction with our interoperability testing. This date will be identified via email. We will plan to review updates of the TKM specification, implementation guide, and administrative guide via email.

Attendees

David Borman	dab@cray.com
Fran Fadden	fran@decvax.dec.com
Owen Gallagher	oweng@jjmhome.uucp
Fred Glover	fglover@zk3.dec.com
Andy Lai	al@hpindda.cup.hp.com
Sharon Lewis	lewis@cs.utk.edu
William Middlecamp	wjm@cray.com
Richard Newton	rnewton@csd.harris.com
Sam Nicholson	scion@pblx.knox.tn.us
Carl Smith	cs@eng.sun.com
Thomas Talpey	tmt@osf.org
Charles Watt	watt@sware.com

2.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Eight working groups and one BOF (Birds of a Feather) in the User Services Area of the IETF met in Boston, Massachusetts. Below is a summary of the User Services IETF activities.

Directory Resources Engineering Group BOF (DREGS)

One BOF was held in the User Services area regarding a Working Group formation on WHOIS and User Lookup Services. This Group would be a joint venture between the Applications Area and the User Services Area of the IETF. The intent of this BOF was to look at recommendations to make WHOIS a generally useful service, with two specific focuses. On the one hand, to have a common name host for a WHOIS services (e.g., whois@ucdavis.edu), and secondly to develop an optional suggested inquiry and response format.

Directory Information Services (pilot) Infrastructure Working Group (DISI)

The DISI Working Group, chaired by Christopher Weider is a Working Group that provides a forum to define user requirements in X.500. It is a combined effort of the User Services Area and the OSI Integration Area of the IETF.

DISI's session focused on additional documentation they would like to produce. There was a discussion of what other documents should be written in addition to those that are listed in the DISI Charter. An advanced usages document has made some progress since the last IETF. Discussion of the revision of RFC 1292 ("A Catalog of Available X.500 Implementations") took place. Chris asked for volunteers to help with the updating of this FYI RFC. This Working Group is beginning to wind down and will probably disband within the next 4-8 months.

Internet Anonymous FTP Archives Working Group (IAFA)

Chaired by Peter Deutsch and Alan Emtage.

The IAFA Working Group, Chaired by Peter Deutsch and Alan Emtage, is chartered to define a set of recommended standard procedures for the access and administration of anonymous FTP archive sites on the Internet.

The IAFA Working Group had its final review of their working draft of "An Anonymous

FTP Site Manager's Guide", which had been worked on via electronic mail in the past few months. There was discussion of future work, including a manager's overview guide for other systems such as Gopher, WAIS, WWW, etc.

Internet School Networking Working Group (ISN)

Chaired by John Clement, Art St. George, and Connie Stout.

The ISN Working Group, Chaired by John Clement, Art St. George, and Connie Stout, 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 Working Group sessions gather educators and Internet folks together. This Boston session continued the discussion on connectivity models for schools...what networking infrastructure currently exist in schools, and to determine a pattern of growth. ISN presented various rough working drafts by various authors. The drafts represent multi-generic connectivity models for schools.

ISN has two current goals. One is to prepare an RFC-to-be that includes various examples of models for schools, with the focus primarily on IP connectivity. ISN's document and intent is not to recommend to schools any one model, but to present various suggestions and various models for school systems to look at. Then the school systems can take these models, choose one for their needs, their students' needs, and how much they have in their budgets to contribute to equipment and software. The second goal is to develop a FAQ (frequently asked questions) on resources.

Internet User Glossary Working Group (USERGLOS)

USERGLOS, Chaired by Gary Malkin and Tracy LaQuey Parker, is chartered to create an Internet specific glossary of networking terms and acronyms for the Internet community.

USERGLOS had two afternoon marathon sessions to continue to find Internet specific terms that are needed in this glossary and to weed out additional words that are not pertinent. This Group competed the final entry review of the glossary and determined that many of the definitions need more work. Volunteers were enlisted. The end of August will be the final Working Group pass and then the document will be submitted as an Internet-Draft by the end of September.

NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

The "Son of NOCTools" Working Group, Chaired by Bob Enger, is updating and revising their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

The NOCTOOL2 Working Group has been a little slow in getting the entries in. Entries are continuing to arrive, and Bob is pushing hard on final document compilation and completion

for an Internet Draft submission in the next two months.

Network Information Services Infrastructure Working Group (NISI)

The NISI Working Group is exploring 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.

NISI's session focused on their Internet Draft, "Privacy and Accuracy Issues in Network Information Center Databases". This document is ready to be submitted to the RFC Editor for publication. There were no additional comments or changes. Review of NISI's NetHelp draft included content and implementation scrutiny. What is NetHelp?? It is a new utility that is being developed within the NISI Working Group, which provides information to users. Further work on NetHelp includes setting up a template via the NISI e-mailing list, developing NetHelp, and writing documentation. NISI needs to reset their goals and milestones within their Charter to reflect this new work.

User Documents Revisions Working Group (USERDOC2)

The USERDOC2 Working Group, Chaired by Ellen Hoffman and Lenore Jackson, is preparing a revised bibliography of on-line and hardcopy 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.

USERDOC2's original bibliography was published in August of 1990. It is sorely out of date. This is the new revised Working Group. A draft document was presented at this session, "FYI on Introducing the Internet - a Short Bibliography of Introductory Internetwork Readings for the Network Novice". This is document is intended to be a short 2-3 page document, with a "companion" document to follow that will be longer in length and more comprehensive for various levels of users, not just novices. A keywords list will also be included for quick, easy reference.

User Services Working Group (USWG)

USWG, Chaired by Joyce K. Reynolds, 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.

The NOCTOOL2 session took 15 minutes of the USWG session to discuss their progress. They needed to meet in this slot, due to time constraints. Darren Kinley has resigned as co-Chair due to lack of funding money to continue to attend IETFs.

Daniel Dern presented a talk on his book in progress, "Internet Guide for New Users" (McGraw/Hill Publishers). This discussion included what Dan is doing, the perception of

some of the need for these types of books, other similar activities (e.g., Zen, EFF, SRI NISC, SIGUCCS, etc.), publishers and contracts - pointers for authors, and what he has learned from this experience. Ed Krol is also working on a book. Paul Holbrook briefly discussed the CICNet Resources Guide that has just been published.

Jill Foster provided an update on RARE activities, including a report on the RARE Information Services/User Services (ISUS) activities. Peter Deutsch led a discussion on a USWG project in development currently called, "An Internet Quick and Dirty". It is intended to be a short document on descriptions of each network service with pointers on where to obtain additional information.

Ideas and thoughts on forming a working group on training materials and having the endeavor be a joint effort between RARE & IETF was discussed. The Group decided that a BOF at the next IETF would be a good start to pursue this topic. Jill Foster and Ellen Hoffman volunteered to co-Chair the BOF.

CURRENT MEETING REPORT

Reported by Joan Gargano/UCDavis

Minutes of the Directory Resources Engineering Group BOF (DREGS)

I. Introductions

Joyce Reynolds opened the session with an overview of how this session was initiated. A last call was issued soliciting comments on the possible elevation of WHOIS to Full Standard. This prompted Russ Hobby, Area Director for Applications, to ask the Internet community for examples of working WHOIS services. A list of about 80 servers was generated. Staff at the University of California, Davis looked at each of these servers and produced a working paper describing valuable new features found in these servers and additions to the protocol that would enhance its usefulness.

II. Review of Current Whois Servers

A brief review of current WHOIS servers and valuable features was presented, based upon the discussion paper. Topics included:

- Hostnames used for WHOIS service for a domain
- Review of information content
- Review of information display
- Review of help responses
- Definition of desired information content
  - People
  - Machines
  - Domains
  - Services

III. Review/Discussion of Draft Working Group Charter

A lengthy discussion followed regarding the role of the WHOIS Working Group in relation to the X.500 Directory Services Working Group. Several members of the Group felt no further work on the WHOIS protocol should be done as it would detract from work on the X.500 Directory Service. In rebuttal, several members of the Group felt there was a role for WHOIS, especially since there are many working servers, and this effort would add value to an existing standard protocol. It was argued that it was consistent with IETF philosophy to allow the development of more than one protocol to perform similar services and to let the Internet community choose which one to use.

Joyce Reynolds brought the focus of the Group back to the issue of the Working Group Gharter and how the WHOIS work would proceed. Members of the Group originally objecting to continued work on WHOIS, agreed with the Charter presented at the meeting

if "Directory Services" was removed from the Working Group name. The Working Group name was changed to "Whois and Network Information Lookup Services" and the Charter was approved for forwarding to the User Services Area.

Attendees

George Brett	ghb@jazz.concert.net
Cyrus Chow	cchow@orion.arc.nasa.gov
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Peter Deutsch	peterd@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Cliff Frost	cliff@cmsa.berkeley.edu
Maria Gallagher	maria@nsipo.nasa.gov
Joan Gargano	jcgargano@ucdavis.edu
Jisoo Geiter	geiter@gateway.mitre.org
Arlene Getchell	getchell@nersc.gov
Steve Hardcastle-Kille	s.kille@isode.com
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Tim Howes	Tim.Howes@umich.edu.
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Hock-Koon Lim	lim@po.cwru.edu
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Mark Needleman	mhn@stubbs.ucop.edu
Eric Nowak	nowak@ans.net
Mark Prior	mrp@itd.adelaide.edu.au
Sheri Repucci	smr@merit.edu
Joyce K. Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Jane Smith	jds@jazz.concert.net
Mark Smith	mcs@umich.edu
Marten Terpstra	terpstra@ripe.net
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com

2.9.1 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 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 Chri
Done	Submit an Internet Draft on 'Catalog of available X.500 Implementations'

Done	Submit to the IESG the 'Catalog of available X.500 Implementations' as an informational document.
Done	Submit an Internet Draft on 'Executive Introduction to X.500'
Done	Submit to the IESG the 'Executive Introduction to X.500' as an informational document.
Done	Submit an Internet Draft on 'A Technical Overview of Directory ervices and X.500'.
Done	Submit to the IESG the 'Technical Overview of Directory Services and X.500' as an informational document.
Jul 1992	Submit as an Internet Draft the 'Advanced Usages' paper.
Jul 1992	Submit as an Internet Draft the 'How to get registered' paper.
Nov 1992	Submit to the IESG the 'How to get registered' paper as an informational document.
Nov 1992	Submit to the IESG the 'Advanced Usages' paper as an informational document.
Nov 1992	Submit as an Internet Draft the 'Pilot Projects Catalog' paper.
Nov 1992	Submit as an Internet Draft the 'Where do I belong in the Directory' paper.
Mar 1993	Submit to the IESG the 'Pilot Projects Catalog' as an informational document.
Mar 1993	Submit to the IESG the 'Where do I belong in the Directory' paper as an informational document.
Mar 1993	Submit as an Internet Draft the 'Guide to setting up a DSA'.
Jul 1993	Submit to the IESG the 'Guide to setting up a DSA' as an informational document.

Internet Drafts:

"Interim Schema for Network Infrastructure Information in X.500", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-disi-netinfrax500-00.txt>

Request For Comments:

RFC 1292	"A Catalog of Available X.500 Implementations"
RFC 1308	"Executive Introduction to Directory Services Using the X.500 Protocol"
RFC 1309	"Technical Overview of Directory Services Using the X.500 Protocol"

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Directory Information Services Infrastructure Working Group (DISI)

Agenda

1. Old Minutes
2. Progress on Assigned Documents
3. Assigning Revision of RFC 1292 (X.500 implementations catalog)
4. New Documents
5. Restaurant Location Scheme
6. New Business

1. No corrections were suggested for the old Minutes.

2. Two documents were discussed: the Advanced usages of X.500 document and the 'How to get registered' document. Drafts had not yet been produced, so the discussion was rather general. The template sent out to elicit responses for the Advanced Usages document was circulated, and three new fields were added: Validation Timestamp, Date of Implementation, and Version Number. It was mentioned that more than ten responses for the document had been received. On the 'How to get registered' document, it was decided that that document should also incorporate the 'Pilot Project Catalog' document.

3. The duty of revising RFC 1292 was graciously accepted by Sri Sataluri and Arlene Getchell.

4. Four new documents were discussed:
   - The Pilot Project Catalog,
   - The DSA Setup Guide,
   - The Naming Philosophy of X.500 paper, and
   - The Restaurant Location Schema paper.

   The Pilot Project Catalog had already been folded into the 'How to Get Registered' paper, and so will not be assigned as a separate document. The DSA Setup Guide was shelved as this paper was seen as essentially being a rewrite of the documentation available with each X.500 implementation. The Naming Philosophy document was seen as being in OSI-DS's territory, and the Restaurant Schema paper was seen as belonging to OSI-DS's Schema subgroup. Thus none of these four papers were assigned in Boston, and they were dropped from DISI's work list.

5. Covered in 4

6. There was no new business.

Attendees

George Brett	ghb@jazz.concert.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Peter Deutsch	peterd@cc.mcgill.ca
Joan Gargano	jcgargano@ucdavis.edu
Jisoo Geiter	geiter@gateway.mitre.org
Arlene Getchell	getchell@nersc.gov
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Tim Howes	Tim.Howes@umich.edu.
Takashi Ikemoto	tikemoto@xerox.com
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Edward Krol	e-krol@uiuc.edu
Thomas Lenggenhager	lenggenhager@switch.ch
April Marine	april@nisc.sri.com
Charlotte Mooers	mooers@nnsc.nsf.net
Eric Nowak	nowak@ans.net
Sheri Repucci	smr@merit.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	
Karen Roubicek	roubicek@faxon.com
Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@nsinic.gsfc.nasa.gov
Jane Smith	jds@jazz.concert.net
Mark Smith	mcs@umich.edu
Patricia Smith	psmith@merit.edu
Simon Spero	ses@cmns.think.com
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com

2.9.2 Internet Anonymous FTP Archives (iafa)

Charter

Chair(s):

Peter Deutsch, peterd@bunyip.com
Alan Emtage, bajan@bunyip.com

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 fo

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/Bunyip and Peter Deutsch/Bunyip

Minutes of the Internet Anonymous FTP Archives Working Group (IAFA)

The Minutes of the San Diego meeting were accepted as read.

The Site Administrator's Draft Document has now been split into two parts by the authors. It was agreed that this was appropriate and the change endorsed. Document 1 describes the procedure for setting up and maintaining an anonymous FTP archive site. Document 2 describes a list of templates which may be used by anonymous FTP archive sites to distribute information about the site itself as well as data and services served by that site.

A detailed discussion of the current drafts of the two new documents was conducted.

- Questions about the examples used in the second document were raised due a possible perception of "UNIX bias" in the current draft. The general consensus was that those individuals wishing to perform the exercise of constructing appropriate templates for other operating systems should do so and contribute them for incorporation into the final document. However, it was considered that the current draft was adequate for the purpose.

- It was agreed that a section on "Packaging for Delivery" was needed, which would explain the various compressing, archiving and other packaging techniques now in use on the Internet. This will include (but not be limited to) UNIX tar and compress, Mac BINHEX, etc. formats. A draft of this section will be inserted for the next draft.

It was agreed that the section on "ethics" and "illegal" practices such as storing copyrighted material without proper release, etc. should be strengthened and split recognizing however, that the definition of "illegal" may vary widely between different jurisdictions. Administrators should be encouraged to check for the rules, laws and regulations governing their own environment. A disclaimer on liability was also suggested.

Following up on a suggest made at the last meeting in San Diego, it was agreed that a section on "Other recommended tools for archive administrators" should be added. This will provide a brief overview of other information tools now available on the Internet with care taken not to "endorse" any specific project. This will also be added for the next draft.

It was agreed that, given the quick progress being made on Uniform Resource Locators (URL), Uniform Resources Identifiers (URI) and Uniform Resource Serial Numbers (URSN) (with work coming out of the UDI BOF and probable Working Group), that fields to accommodate these be added to the appropriate templates in Document 2.

It was agreed that more work needs to be done on the "Why you should run an archive" section of the first document.

Ellen Hoffman of Merit (ellen_hoffman@merit.edu) has volunteered her site to coordinate submissions of Services records (as defined in the second document) until services directories start to become widely available. It was agreed that an additional field "Sponsoring Organization:" should be added to support this sort of "proxy-list" service.

No agreement was reached on the the issue of developing a "cost daemon" program for ranking the distance or reachability of an archive from a particular user's host, although there was some interest in the idea. Peter Deutsch (peterd@bunyip.com) will act as coordinator for Volunteers interested in working on such a project.

Questions were asked about a recommended procedure for registering new information sites as they come on-line. Currently they are discovered through ad hoc methods such as a posting to comp.archives.admin, email to an archie site, etc. It was agreed that this should be noted in the first document, advising new sites to post their availability to comp.archives and other avenues for the distribution of this information should be pursued.

In discussing the section on security, it was pointed out that a number of sites continue to run anonymous FTP archives to exchange non-public information. Such sites function without password protection and the information so stored is publicly available via the standard anonymous FTP login procedure. It was agreed that a strengthened section on security would specifically warn against this practice, as it constitutes a form of "security through obscurity" that is not endorsed and which has already been shown to be problematic in practice.

The subject of a companion "User Guide to Anonymous FTP" was again raised and it was again agreed that such a document would be useful. Ellen Hoffman of Merit (ellen_hoffman@merit.edu) and April Marine of SRI (april@nisc.sri.com) have agreed to examine a number of current freely available documents with the intention of editing them to make them suitable for a more general audience.

In a general discussion on missing attributes in various templates, it was suggested that the following be considered:

- In templates describing textural information a field to list the ISO standard encoding for language and character set should be added.
- In templates describing documents (such as IAFA-ABSTRACT) a field listing the ISBN number (if available) should be considered.
- Where appropriate, fields to list appropriate "Distribution:" and "Copying:" fields should be added.
- Additional available contact information such as FAX numbers should be added.

New Business

It was agreed that work should begin on documenting URIs, UDIs and URSNs. It was agreed that the chairs in concert with Tim Berners-Lee of CERN (timbl@nxoc01.cern.ch) would canvas the Area Directors for User Services and Applications to determine whether IAFA should play a role in this, or new/additional Working Groups should be tasked with the job.

It was agreed that there might be interest in a document explaining tools intended to coordinate sharing of information across the Internet. These might include existing tools or look at specifications to encourage new tools. A summary of such tools will be investigated and a concrete proposal prepared for the next meeting.

Attendees

Mark Baushke	mdb@cisco.com
Pravin Bhagwat	w2pravin@watson.ibm.com
George Brett	ghb@jazz.concert.net
Mitchell Charity	mcharity@lcs.mit.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Daniel Dern	ddern@world.std.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
James Halpin	halpin@turkey.sw.stratus.com
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Thomas Lenggenhager	lenggenhager@switch.ch
Hock-Koon Lim	lim@po.cwru.edu
Kim Long	klong@sura.net
April Marine	april@nisc.sri.com
Ed Menze	menze@cs.arizona.edu
Charlotte Mooers	mooers@nnsc.nsf.net
Chris Myers	chris@wugate.wustl.edu
Clifford Neuman	bcn@isi.edu
Hank Nussbacher	hank@vm.tau.ac.il
Bob Page	bob.page@eng.sun.com
Rakesh Patel	rapatel@hardees.rutgers.edu
Michael Patton	map@lcs.mit.edu
Marsha Perrott	mlp+@andrew.cmu.edu

Mel Pleasant	pleasant@cs.rutgers.edu
Mark Prior	mrp@itd.adelaide.edu.au
David Reiser	dsr@codex.prds.cdx.mot.com
Robert Reschly	reschly@brl.mil
Joyce K. Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com
Anthony Rutkowski	amr@sprint.com
Henry Sanders	henrysa@microsoft.com
Richard Schmalgemeier	rgs@merit.edu
Tim Seaver	tas@concert.net
Vincent Sgro	sgro@cs.rutgers.edu
Erik Sherk	sherk@sura.net
Michael Slocombe	slocombe@bbn.com
Jane Smith	jds@jazz.concert.net
Karen Sollins	sollins@lcs.mit.edu
Simon Spero	ses@cmns.think.com
Marten Terpstra	terpstra@ripe.net
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Evan Wetstone	evan@rice.edu
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com

2.9.3 Internet School Networking (isn)

Charter

Chair(s):

John Clement, clement@educom.edu
Arthur St. George, stgeorge@bootes.unm.edu
Connie Stout, cstout@tenet.edu

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 1992 Release X.500 "K-12 People Directory" version in collaboration with Merit. Develop plans and milestones for K-12 Resources Directory.

Mar 1992 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 1992 Release X.500 K-12 Resource Directory version in collaboration with Merit. Present final draft guideline statement.

CURRENT MEETING REPORT

Reported by John Clement/EDUCOM/CoSN

Minutes of the Internet School Networking Working Group (ISN)

The Agenda had two items: connectivity models, and development of an FAQ (Frequently Asked Question) archive. A third item was introduced by Scott Williamson: developing a recommendation on domain naming schemes for K-12 schools and organizations.

We reviewed connectivity alternatives and growth paths. The mailing list for this subtask is connect@unmvma.unm.edu. We began by a brief review of a number of documents (of varying formality and draft status) on topics related to connectivity for schools.

Newman, Denis, Susan Bernstein and Paul A. Reese. Local Infrastructures for School Networking: Current Models and Prospects. BBN Report No. 7726, April 1992. Available in hardcopy from Bank Street School of Education and by ftp from nysernet.org.

Susan Bernstein presented major findings of this report. She made it clear that the report was not itself about connectivity alternatives themselves, but about the current situation in school networks and machine- (or LAN-) to-WAN connectivity. Many schools have LANs, and many have WANs; but the latter are almost exclusively used for administrative (not instructional) purposes; and the instructional LANs are not connected to the available WANs, and indeed often instructional LANs are dedicated to Individualized Learning Software (ILS) systems, using proprietary network protocols. The terminal-host model is almost entirely what is currently used for school-to-WAN connection.

The paper recommends a possible growth path from the point of individual dialin connections from a machine to a host resource, a path that does "..not assume the initial availability of LANs but begin(s) developing the expertise and the software to support true network connections in schools with stand-alone computers.." (p.34). The path would lead schools through connections to a remote network, to a leased line to a local Internet node, and finally to a local Internet server.

Reilly, Rob, and Kurt Lidl. A National Computer Infrastructure: The Light at the End of the Tunnel May be an Oncoming Rain! Unpublished manuscript, July 1992. Available by request from rreilly@athena.mit.edu.

Rob Reilly sketched the main points of his paper as emphasizing the need to deal with both the physical network and the logical one; a synopsis of connectivity models within one overview structure; and process suggestions for future steps.

Burns, Pat. Establishing Connections to the Internet. Unpublished manuscript, n.d. (3/92 app.?) Available by ftp as models_last_revision.ps from ariel.unm.edu (library directory).

No one was present from CSU's shop to discuss this paper.

Hastings, Eugene. Connectivity Models for Internet Access. Version 1.1. Monday, July 13, 1992. Unpublished manuscript. Available by request from hastings@psc.edu.

Gene Hastings' paper begins with the presumption that one is connecting networks together. He calls for schools to build school-specific infrastructures (up to a consortium of school districts) and then connect in bulk to the Internet.

His models address needs and constraints of connectivity situations in different situations – for example, in some areas telephone tariffs are based on metered but untimed calls. In others, the near future will bring ISDN capability – in the Bell Atlantic market, perhaps for as low as $40/month.

Pricing models are not yet set; there is a chance for user communities such as schools to affect these decisions. However, many of the connection options are currently being priced as extra-cost luxury items for small customers, which is exactly the wrong approach.

Discussion then began of the various models and designs presented. Susan Bernstein pointed out that the model outlined by Reilly and Lidl was, even in its fullest expression, not yet a "local Internet server" model, but relied on a leased-line connection to an external host machine, itself on the Internet. For a long-term solution to the problem of connecting K-12, we have to migrate Internet and other servers to the school settings.

Brian Lloyd mentioned work he was doing in the Bay Area to connect two schools, as well as preparing a paper for BARRNET on methods. He reported that he was obtaining old PCs and installing them as routers using PPP. He envisions, he said, a three-phase process: in Phase 1 a connection to a school is provided, but (the campus or regional – some already-connected group) carries the technical and administrative load; in Phase 2, the school learns to manage its own access to the network; and in Phase 3 the school provides a name and file server for a group of schools.

Gerry White of Applitek mentioned that his company is exploring providing broadband Ethernet connectivity to schools via installed cable. An unspecified number of the nation's school cable plants have "backchannel" capability, which will allow interactive uses. A number of concerns were expressed with issues such as tuning and management of broadband over cable, with ghosting and interference and their impact on packet transmission; but there was consensus that such models were well worth exploring and noting in an RFC draft.

Eric Hood of NorthWestNet and FARNET K-12 asked that any analysis of models consider costs of network management.

Michael Powell of Pacific Bell mentioned that his company has entered an agreement with California State University and the California Technology Project to provide gateways for an Internet connection to every educator in the state, under the name Knowledge Network. They are eager to participate in the process of developing connectivity models.

John Clement offered that there seemed to be enough written, and enough areas of general agreement between the different documents, that one could now attempt a draft RFC. He asked for volunteers, and said there would be a document posted for discussion before the 25th IETF in Washington, DC (November 16-20, 1992). The following individuals volunteered:

- Rob Reilly, Lanesboro, MA Schools (first draft synthesizer)
- Susan Bernstein (reviewer)
- Eugene Hastings, Pittsburgh Supercomputer Center
- Ellen Hoffman, Merit Network Inc.
- Brian Lloyd and Constance Lloyd, Cameron Park, CA
- Michael Powell, Pacific Bell
- Gerry White, Applitek

Development of an FAQ archive on school connectivity issues. The mailing list for this subtask is faq@unmvma.unm.edu. It was mentioned that Ed Vielmetti of CICnet has installed the Kidsnet archives on a WAIS server. The resource seems likely to provide useful "first answers" for an FAQ file.

John Clement offered to use the cosndisc (Consortium for School Networking) discussion forum to develop a set of questions. April Marine of SRI and Jennifer Sellers of NASA then offered to draft answers to the questions using the Kidsnet WAIS resource. It was agreed that a draft FAQ file could be made available by the time of the 25th IETF in Washington, DC (November 16-20, 1992).

Developing a recommendation on domain naming schemes for K-12 schools and organizations. A discussion was introduced by Scott Williamson of ISI. An increasing number of requests are being received for K-12 domain names, and there is no agreed-upon naming structure. The prospect of a very large number of K-12 domains and nodes raises serious concerns.

Discussion points raised: already existing are names such as xxhs.edu. Expansion of this model will provide a very flat structure with very little information, and offer a lot of chances for conflict over names (there are a large number of George Washington or Martin Luther King, Jr. high schools).

Additionally, although elementary schools and high schools are identified with relative clarity, there is considerable ambiguity in the middle range (middle schools, intermediate schools, junior high schools). Finally, the point was raised that this information is of uncertain value for a naming domain.

Considerable discussion was held on the notion of a geography- based naming convention within the .us domain. Ellen Hoffman of Merit Network, Inc. offered to prepare a draft document for discussion on this issue.

The general mailing list for this discussion is isn-wg@unmvma.unm.edu. Copies of the papers cited, can be had by asking John Clement <clement@educom.edu>. Reactions and comments are much appreciated.

Attendees

John Clement	clement@educom.edu
Daniel Dern	ddern@world.std.com
Sallie Fellows	sallie%ed@psc.plymouth.edu
Constance Fleenor Lloyd	cfleenor@lloyd.com
Maria Gallagher	maria@nsipo.nasa.gov
Neil Haller	nmh@thumper.bellcore.com
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Eugene Hastings	hastings@a.psc.edu
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
J. Paul Holbrook	holbrook@cic.net
James Keller	j.keller@sprint.com
Hock-Koon Lim	lim@po.cwru.edu
Brian Lloyd	brian@lloyd.com
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
April Marine	april@nisc.sri.com
Cindy Mazza	
Cyndi Mills	cmills@nnsc.nsf.net
Marsha Perrott	mlp+@andrew.cmu.edu
Michael Powell	mdpowel@pacbell.com
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Michael Roberts	roberts@educom.edu
Anthony Rutkowski	amr@sprint.com
Jennifer Sellers	sellers@nsinic.gsfc.nasa.gov
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Carol Ward	cward@westnet.net
Moira West	mjw@cert.org
Evan Wetstone	evan@rice.edu
Gerard White	
Scott Williamson	scottw@nic.ddn.mil

2.9.4 Internet User Glossary (userglos)

Charter

Chair(s):

Tracy LaQuey Parker, tracy@utexas.edu
Gary Malkin, gmalkin@xylogics.com

Mailing Lists:

General Discussion: usergloss@xylogics.com
To Subscribe: usergloss-request@xylogics.com
Archive: xylogics.com:gmalkin/usergloss/usergloss-arc

Description of Working Group:

The Internet User Glossary 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
Done	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.
Done	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the Internet User Glossary Working Group (USERGLOS)

The Userglos Working Group met twice at this IETF. The final include/exclude pass was completed on the entries in the Glossary. In particular, networks were excluded. Several definitions were found to be wanting and will be worked on. It was also decided that there will be no indices.

Agenda

- Discuss format.
- Discuss inclusion of networks and standards levels.
- Weeding-out Process (which definitions should be excluded?).
- Discuss indices.

The only format decision was whether or not to right-justify the document. Joyce pointed out that RFC 1111 requires "ragged edge". The column width will also be increased.

It was decided that it was impossible to draw a fair line to determine which networks (e.g., NSFnet, Ebone, AARNET, NEARnet, etc.) should be included and which should be excluded. Therefore, no networks will be included. The exception is ARPAnet, for historical reasons.

It was decided to remove the definitions for standards levels (e.g., Experimental, Historical, Recommended, etc).

Overall, the number of entries was decreased.

Given the intended audience, and the use to which they will put this Glossary, it was decided that indices based on an entry category would not be useful. That is, the document will be used by people to find the meaning of a word used by a guru. We decided that anyone interested in a list of networks, organizations or protocols would have much more definitive sources from which to work.

Attendees

George Brett	ghb@jazz.concert.net
Alan Clegg	abc@concert.net
Daniel Dern	ddern@world.std.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov

Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Hank Nussbacher	hank@vm.tau.ac.il
Marsha Perrott	mlp+@andrew.cmu.edu
Jack Pope	pope@usdcsv.acusd.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	
Jim Romaguera	
Karen Roubicek	roubicek@faxon.com
Richard Schmalgemeier	rgs@merit.edu
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Simon Spero	ses@cmns.think.com
Fred Ziegler	ziegler@aspentec.com

2.9.5 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 Joyce K. Reynolds/ISI

Minutes of the NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

The NOCTOOL2 Working Group met during the User User Services Working Group (USWG). For a brief summary, please refer to the USWG Minutes.

2.9.6 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 ser
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
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.

Request For Comments:

RFC 1302 "Building a Network Information Services Infrastructure"

RFC 1355 "Privacy and Accuracy Issues in Network Information Center Databases"

CURRENT MEETING REPORT

Reported by April Marine/SRI

Minutes of the Network Information Services Infrastructure Working Group (NISI)

The NISI Group met its primary goal, which was to come up with some concrete contents for the "nethelp" document discussed at the last meeting.

Announcements

There were no comments in the meeting regarding the current Internet Draft "Privacy and Accuracy Issues in Network Information Center Databases." As that version had been sent online as well to no comments, after the meeting we requested that Joyce as Area Director foster the document along the track to FYI RFC status.

Status of Other Work

We had two documents on our plate, one being the "nethelp" project and the other a list of services document. After a brief discussion, the Group agreed that the need for a "list of services" document was being met by the Quick and Dirty document being prepared within the USWG Group, so we punted that project over there.

Discussion/modification of Nethelp Outline

The draft outline for a document specifying a "nethelp" utility was passed out. We decided not to break into small groups, and instead the whole Group discussed the project.

There was some disagreement about whether or not we should discuss how the nethelp utility should be implemented. However, it was pointed out that until we knew what we wanted, no one could implement it anyway. So, we agreed eventually to worry less about implementation than about specifying the function and content of Nethelp. Those in the business of implementing things seemed to feel it would be pretty easy to do if they just knew what we wanted.

We came to the following goal agreements:

- We wanted something Very Simple that any site could implement.
- Nethelp should give info about the Internet, not operating system help or help regarding other programs, etc. It would let people know that they were on the Internet and tell them where to get more information.

- We would start with a basic system, then add bells and whistles after the initial concept was specified.

This set of agreements eliminated, for the moment, implementing an Internet-wide distributed help system or a central network server.

We had another discussion regarding distributing information down the pyramid from a mid-level network information center to the end-user on some connected network. It was observed that a nethelp utility would circumvent the current chain from NIC to site contact to whomever to end-user. True enough. But sometimes a local contact acts as a bottleneck for information and does not pass things along. The "nethelp" idea is designed to be easy to implement locally (so as not to be a strain on an overworked LAN administrator), and to directly inform users that they are on a network and that there is more information available about that network's resources. This will allow the end-user to seek information himself from a NIC or to prod his local contact to pass along information he may not have known his users were interested in.

So, after all this, we actually did decide what we wanted a first "nethelp" to do. We wanted a user to be able to type "nethelp" and see one screen of information. That screen would say something like:

```
--------------------
  NETHELP answers questions about ''the network''.

  You are on X-Net.  X-Net is part of the worldwide Internet,
  thus providing you with access to a myriad of resources.

  Your local contact for help is:

  For system help:                  For network information:

   Name:                              Name:
   Phone Number:                      Phone Number:
   E-mail:                            E-mail:

--------------------
```

The template would be designed to allow local sites to easily fill in information pertinent to that site. For example, the "X-net" slot could consist of a local network name, the host name, a host address, a network address, or any subset of that information. We would prefer that users see a fully qualified domain name, or at least the address of their network, so that they can eventually orient themselves in relation to the rest of the Internet.

We discussed adding a line that says, "For more network help, type nethelp X," but I'm not sure we want to do that until the next phase.

What Next

Our Action Items from the meeting were:

- Develop the nethelp template/screen on the mailing list.
- Write the document. Poll the mailing list regarding whether it should have a section discussing implementation strategies (and to get a volunteer to write it :-).
- Review the goals and milestones and revise if necessary.
- Start discussing the next phase of nethelp and what that more extensive utility will feature.

We agreed that in the interest of time and of getting something DONE, we would start with the simple nethelp recommendations above. But we are generally interested in expanding the information about the Internet that is readily available to a user, so one of the next things we want to do is talk about what information a more extensive help system would contain and how it would be deployed. So, that is a topic for our next Agenda.

Our other goal, obviously is to write up the recommendations for the first stage of nethelp and get that draft out for comments. (Feel free to comment on the plan as outlined in these Minutes.)

Attendees

Luc Boulianne	lucb@cs.mcgill.ca
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov
Joan Gargano	jcgargano@ucdavis.edu
Jack Hahn	hahn@sura.net
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
J. Paul Holbrook	holbrook@cic.net
Takashi Ikemoto	tikemoto@xerox.com

Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Thomas Lenggenhager	lenggenhager@switch.ch
Hock-Koon Lim	lim@po.cwru.edu
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Charlotte Mooers	mooers@nnsc.nsf.net
Hank Nussbacher	hank@vm.tau.ac.il
Bob Page	bob.page@eng.sun.com
Mark Prior	mrp@itd.adelaide.edu.au
Sonya Reimer	sonya@brl.mil
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@nsinic.gsfc.nasa.gov
Michael Slocombe	slocombe@bbn.com
Patricia Smith	psmith@merit.edu
Karen Sollins	sollins@lcs.mit.edu
Simon Spero	ses@cmns.think.com
Marten Terpstra	terpstra@ripe.net
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com

2.9.7 User Documents Revisions (userdoc2)

Charter

Chair(s):

Ellen Hoffman, ellen_hoffman@um.cc.umich.edu
Lenore Jackson, jackson@nsinic.gsfc.nasa.gov

Mailing Lists:

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

Description of Working Group:

The USERDOC2 Working Group will revise a bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "How to use the Internet" (RFC 1175, FYI 3). This effort will also be used to help locate documentation produced by other organizations and will make it available for distribution on the Internet. The target audience is those individuals who provide services to end users and end users themselves.

A segment of the document will:

(1) Identify and categorize useful documents/reference materials/training tools,

(2) Publish both an on-line and hard copy of this bibliography,

(3) Develop and implement procedures to maintain and update the bibliography and investigate methods to provide the information in an on-line format.

(4) As a part of the update process, identify new materials for inclusion into the active bibliography.

(5) Review procedures for periodic review of the bibliography by the User Services Working Group.

Goals and Milestones:

Mar 1992	Identify new "sources of information" (e.g., individuals, mailing lists, bulletins, etc.) Review existing document and obtain comments from others in USWG about needed revisions at the San Diego IET
Jun 1992	Publish Internet Draft for review comments at the Boston IETF.
Jul 1992	Final review, reposting as updated Internet Draft for last pass comments.
Dec 1992	Bibliography submitted as an FYI RFC to RFC Editor.

CURRENT MEETING REPORT

Reported by Ellen Hoffman/Merit and Lenore Jackson/NASA

Minutes of the User Documents Revisions Working Group (USERDOC2)

Review of Previous Session

The Group reviewed the discussion of the San Diego IETF meeting held during the USWG meeting to issue two documents. This includes producing a new, short, 2-pager aimed specifically at novice users, and a revision of the longer bibliography to update RFC 1175.

Status of Documents

Two documents were distributed. The first, a draft of the proposed short bibliography was distributed as well as a list of proposed additions to the long bibliography.

Review of New Short Bibliography

A number of changes were suggested for the draft novice user bibliography. There was some discussion that there is no single, brief (2-3 page) description of the Internet, and it was agreed that this should be further discussed at the USWG. One book was added to the list, some reorganization was proposed, and it was agreed there should be a line to encourage users who needed more information to contact their local network provider. It was also agreed not to add how-tos about applications like mail or FTP.

A focus of the discussion is how new users, who may not have a connection or may not know how to get materials on the net, could find on-line documents and whether there should be more instructions about FTP, WAIS, Gopher, etc. The Group accepted a proposal to have a single directory with all the on-line materials which several NICs would agree to mirror, and when possible, have these accessible through other tools such as WAIS, Gopher, and Archie. This directory could also be updated as new resources are added. While this was agreed to as a short-term solution that could be implemented immediately, it was also recognized that this was not an ideal longer term solution. The Group agreed to go ahead so the material would at least be available soon using existing mechanisms.

Merit volunteered to create a "master" directory on nis.nsf.net and seven other NICs volunteered to create "mirror" directories. These include ftp.nis.sri.com, nnsc.nsf.net, ftp.concert.net, paccom?, nic.ddn.mil, nic.mr.net, and ftp.jvnc.net. Merit will also set up its mail server to send out a generalized "help" file for those who want more information on getting the on-line documents and using retrieval tools to help keep the bibliography short.

Review of Proposed Additions to Longer Bibliography

The list of proposed additions was briefly discussed. It was agreed that the bibliography

needed to reflect the world-wide Internet and should not focus on NREN and other U.S. policy papers or controversies. Volunteers were recruited to help write abstracts for the new entries.

Review of Action Items

All attendees will be automatically added to the mailing list. A revised short bibliography will be sent out within two weeks, and if there are no major objections, it will be sent to the RFC editor. The long list additions will be sent out and the Group will recommend additions. The goal is a draft with all additions for review at the next IETF. Merit will put up a directory and work with the volunteer mirror sites to implement the on-line sources.

Attendees

Vikas Aggarwal	aggarwal@nisc.jvnc.net
George Brett	ghb@jazz.concert.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Alan Clegg	abc@concert.net
Daniel Dern	ddern@world.std.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Edward Krol	e-krol@uiuc.edu
Hock-Koon Lim	lim@po.cwru.edu
Andrew Maffei	arm@aqua.whoi.edu
April Marine	april@nisc.sri.com
Charlotte Mooers	mooers@nnsc.nsf.net
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	
Karen Roubicek	roubicek@faxon.com
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@nsinic.gsfc.nasa.gov
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Scott Williamson	scottw@nic.ddn.mil

2.9.8 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.

(1) 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.

(2) Create working groups or other focus groups to carry out projects deemed worthy of pursuing.

(3) 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"

RFC 1325 "FYI on Questions and Answers Answers to Commonly asked "New Internet User" Questions"

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI

Minutes of the User Services Working Group (USWG)

The NOCTools2 Working Group session took 15 minutes of the USWG session to discuss their progress. They needed to meet in this slot, due to time constraints. Darren Kinley has resigned as co-Chair of NOCTools2 due to lack of funding money to continue to attend IETFs. Robert Enger will carry on as Working Group Chair. Additional volunteers were enlisted to help Robert Enger complete the updated NOCTools catalog. Joyce Reynolds will rewrite the introduction. Charlotte Mooers will assist in updating the indices. Robert Enger will keep pursuing and collecting the vendor updates. Kim Long volunteered to assist in additional collecting of information. There are also plans to incorporate the collected information into WAIS.

Daniel Dern presented a talk on his book in progress, "Internet Guide for New Users" (McGraw/Hill Publishers). This discussion included what Dan is doing, the perception of some of the need of these types of books, and other similar activities (e.g., Zen, EFF, SRI NISC, SIGUCCS, etc.). Discussion also focused on publishers and contracts, and authors' rights. Daniel provided additional pointers for authors via what he has learned from his experience. There is an estimated six books that will be published by the end of this year focusing on users and the Internet.

Ed Krol is also working on a book and he passed around a draft copy of his work for USWG attendees to comment. Ed's book is due to be published by the end of this year. The second chapter of Ed's book, entitled, "What is the Internet?" seemed to Ellen Hoffman and other USWGers a good start for new users. Ed made mention that the User Services Area of the IETF may use this chapter for FYI RFC publication if he can obtain approval by his publisher. Ellen Hoffman and Ed Krol will work together on this endeavor.

Jill Foster provided an update on RARE activities, including a report on the RARE Information Services/User Services (ISUS) activities. She talked about the European Information Service - CONCISE, and how this effort is expanding (e.g., there are new people teaching new users to use the networks). The RARE Status Report has been published, which includes contributions from eighteen countries. This report is available in hardcopy, or on-line. Jill announced to the USWG the NSC'92, The Network Services Conference 1992, November 3-5, 1992, which is being organized in Pisa, Italy this fall by EARN in cooperation with EUnet, NORDUnet, RIPE and RARE.

Peter Deutsch led a discussion on a USWG project in development currently called, "An Internet Quick and Dirty". It is intended to be a short document on descriptions of each network service with pointers on where to obtain additional information. At the San Diego IETF, Peter volunteered to produce this document with the assistance of: Pat Smith, Cyndi Mills, Dan Matzke, Ellen Hoffman, Anders Gilner, April Marine, Jill Foster and Joyce Reynolds. This project will continue to be developed within the USWG and will interact

with the Network Information Retrieval Working Group that is currently in formation.

Ideas and thoughts on forming a Working Group on training materials and having the endeavor be a joint effort between RARE & IETF was discussed. The USWG decided that a BOF at the next IETF in Washington, D.C. would be a good start to pursue this topic. Jill Foster and Ellen Hoffman volunteered to co-Chair the BOF.

Paul Holbrook briefly discussed the CICNet Resources Guide that has just been published.

Attendees

George Abe	abe@infonet.com
George Brett	ghb@jazz.concert.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
James Conklin	jbc@bitnic.educom.edu
Daniel Dern	ddern@world.std.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Roger Fajman	raf@cu.nih.gov
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Alf Hansen	Alf.Hansen@delab.sintef.no
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
J. Paul Holbrook	holbrook@cic.net
Edward Krol	e-krol@uiuc.edu
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Thomas Lenggenhager	lenggenhager@switch.ch
Hock-Koon Lim	lim@po.cwru.edu
Kim Long	klong@sura.net
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Cyndi Mills	cmills@nnsc.nsf.net
Charlotte Mooers	mooers@nnsc.nsf.net
Hank Nussbacher	hank@vm.tau.ac.il
Marsha Perrott	mlp+@andrew.cmu.edu
Michael Powell	mdpowel@pacbell.com
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Karen Roubicek	roubicek@faxon.com

Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@nsinic.gsfc.nasa.gov
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Marten Terpstra	terpstra@ripe.net
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com
Fred Ziegler	ziegler@aspentec.com

Chapter 3

Network Status Briefings

3.1 CERFnet

Presented by Pushpendra Mohta/CERFnet

THE INTERNET IN BRAZIL

Pushpendra Mohta
pushp@cerf.net

July 13, 1992

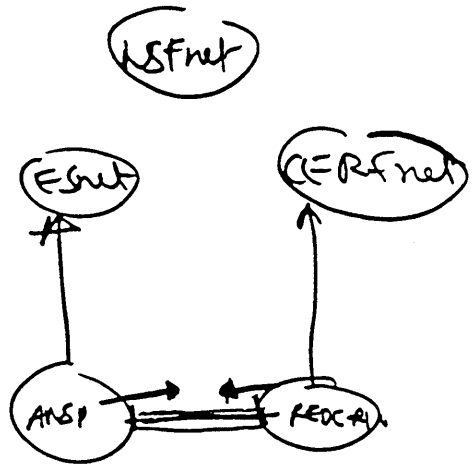

ACKNOWLEDGEMENTS

- Most material for this talk comes from Michael Stanton of PUC-RIO, BRAZIL. For information on the RNP, contact him at stanton@cerf.net or michael@inf.puc-rio.br

- Material for REDE-RIO comes from UFRJ. More information can be had from edmundo@lcca.nce.ufrj.br (Edmundo Da Silva)

REDE-RIO

- Early efforts in 1990, contacts with CERFnet
- Blessed by NSF, small grant to CERFnet
- Many expired "letters of credit" and export licences later, the network was installed in May 1992
- The backbone, the International link to CERFnet and many sites were brought on line in about 10 days

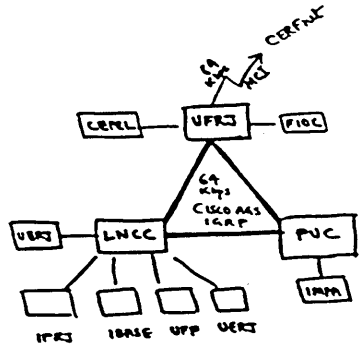

NATIONAL BACKBONE NETWORK

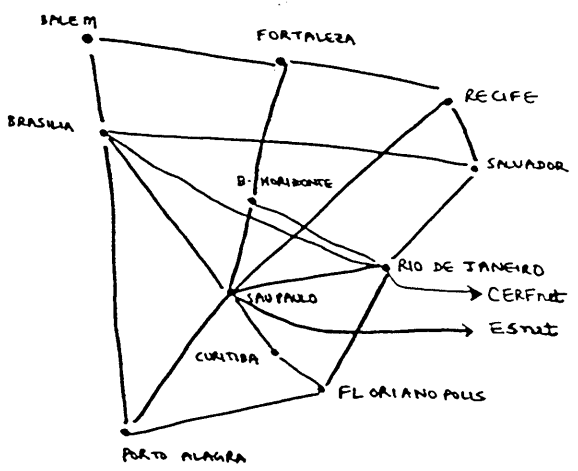

IMPLEMENTATION

RNP: National Backbone with 11 nodes
Mesh Topology
9.6 → 64 Kbps
Initially Sun workstations as routers

REDE-RIO: RIO DE JANEIRO state network
Backbone is a ring topology
64 Kbps backbone
Cisco routers
~12 sites
connected to CERFnet

ANSP: SAO Paulo state network
Backbone is a ring topology
Ports are 64Kbps
Cisco routers
~12 sites
Connected to ESnet

ANSP & REDE-RIO are also connected by a 9600 bps link

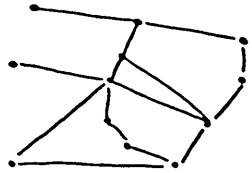

ARCHITECTURE OF THE NEW NETWORK

- Modelled after the US Internet, the National Council For Scientific and Technological Development (CNPq) adopted a hierarchical model
  1. NATIONAL
  2. REGIONAL
  3. INSTITUTIONAL
- The National Research Network (RNP) project was created to
  (a) Provide Interconnectivity between state + regional networks
  (b) Provide International Access
- State governments were encouraged to provide regional infrastructure.
- National + International routers were required to support a variety of protocol suites, in particular TCP/IP and DECnet and CLNS.

EARLY EXPERIMENTATION WITH INTERNET

- With the election of a new government in 1990, utilization of non-OSI standards and in particular, TCP/IP, was permitted
- Growth of UNIX machines led to many TCP/IP LANs
- In 1991, FAPESP in Sao Paulo established a 9600 bps link to ESNet at Fermilabs in Chicago. FAPESP offered Internet access to those who could use it

NETWORKING INFRASTRUCTURE IN BRAZIL

- Prior to 1988 Embratel did not permit sharing of links : NO BACKBONE
- Early networking : BITNET

9/88 → LNCC, RIO ⇔ Univ of Maryland, College Park 9600 bps

11/88 → FAPESP, Sao Paulo ⇔ Fermilab, Chicago 9600 bps

5/89 → UFRJ, RIO ⇔ UCLA, Los Angeles 4800 bps

→ In 1988 Embratel provisionally permitted sharing of channels for academic purposes, opening the way to setting up the national backbone.

- By 1990, most important centers were connected to either RIO or Sao Paulo. This network used standard BITNET protocols. Its major contribution was e-mail access to many.

TELECOMMUNICATIONS IN BRAZIL

- Presently a state and federal monopoly
  - → Long range carrier : EMBRATEL
  - → Regional operators : TELERJ etc.
- Embratel expects to replace current MICROWAVE and radio links with optical fiber by 1996
- International service is via submarine cables and Intelsat

DATA COMMUNICATIONS

- PSDN 1200 – 9600 bps national coverage Linked to other national PSDNs
- Leased lines predominantly at 1200 – 9600 bps
- Slow movement to 64 K bps

3.2 EBONE Report

Presented by Bernhard Stockman/SUNET

EBONE - 92

SEPTEMBER 91

INITIAL PROPOSAL

AUTUM 91

FORMATION OF THE EUROPEAN CONSORTIUM OF CONTRIBUTING ORGANIZATIONS (ECCO)

FORMATION OF THE EBONE ACTION TEAM (EAT)

EAT DEFINES THE EBONE IMPLEMENTATION PLAN

SPRING 92

PHYSICAL IMPLEMENTATION STARTED

SPECIFICATION OF ROUTING CONFIG

SPECIFICATION OF MANAGEMENT AND OPERATIONS

KTH IN STOCKHOLM SWEDEN NOMINATED AS THE EBONE NETWORK OPERATIONS CENTER (EBONE NOC)

EBONE DEVELOPMENT

INTEGRATION WITH THE EUROPEAN MULTIPROTOCOL SERVICE

EXPANSION TO CENTRAL AND EAST EUROPE

INCREASE OF BANDWIDTH WITHIN EUROPE AND BEWTEEN EUROPE AND THE US

FORMAL CONNECTIVITY AGREEMENTS FOR 1993

GENERAL PURPOSE TRANSIT CAPACITY TO OTHER CONTINENTS VIA THE GLOBAL INTERCHANGE (GIX)

EBONE TOPOLOGY

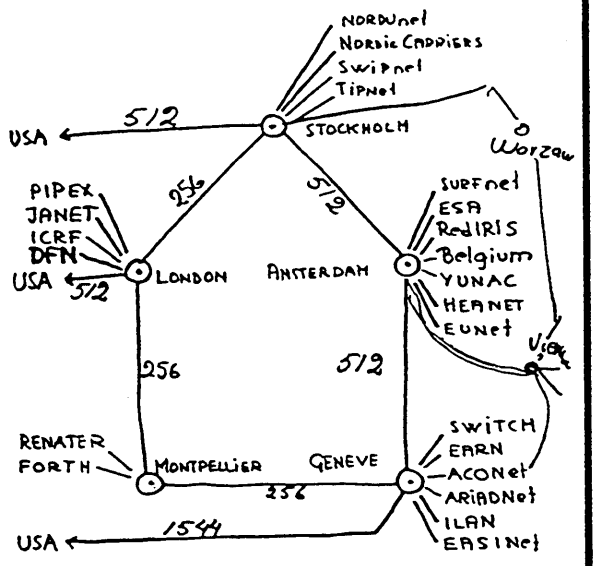

EBONE TODAY

STOCKHOLM - LONDON 256 KBPS UP AND RUNNING

STOCKHOLM - AMSTERDAM UPGRADE TO 512 KBPS SCHEDULED FOR JUNE 17

LONDON - MONTPELLIER 256 KBPS ORDERED

AMSTERDAM - GENEVA 512 KBPS UP AND RUNNING

EBONE TO IXI GATEWAYS PROPSED IN AMSTERDAM AND LONDON

ROUTING PLAN UNDER IMPLEMENTATION

EBONE NOC BEING INSTALLED

MAIN EBONE SPECIFICATION DOCUMENTS AVAILABLE VIA ANONYMOUS FTP FROM NIC.NORDU.NET OR ARCHIVE.RIPE.NET IN DIRECTORY EBONE

EBONE FUNCTIONALITY

PROVIDED SERVICES

- INTERNET IP
- ISO IP (CLNS)

ROBUST GENERAL PURPOSE BACKBONE
OPEN FOR ALL NETWORKS

CENTRALIZED MANAGEMENT

POLICIES IMPLEMENTED AT THE LOCAL
REGIONAL CONNECTIONS

GENERAL PURPOSE TRAFFIC EXCHANGE

THE NEED FOR A TOP LEVEL POLICY-FREE
INTERCHANGE

HIERARCHICAL MODEL WITH LESS
RESTRICTIONS THE CLOSER TO THE TOP

CONNECTIVITY BETWEEN NETWORKS
ON THE SAME POLICY LEVEL

ORDERED CONNECTIVITY BETWEEN
NETWORKS ON ADJACENT LEVELS

CONNECTIVITY BETWEEN NETWORKS
AT DIFFERENT POLICY LEVELS
OUTSIDE THE HIERARCHICAL STRUCTURE

EBONE

STOCKHOLM EBONE BOUNDARY SYSTEM

- NORDUnet
- Nordic Carriers
- SWIPnet
- TIPnet

LONDON EBONE BOUNDARY SYSTEM

- JANET
- ICRF

AMSTERDAM EBONE BOUNDARY SYSTEM

- SURFnet
- ESA
- RedIRIS
- K.U. Leuven
- YUNAC
- HEANET
- EUnet

MONTPELLIER EBONE BOUNDARY SYSTEM

- RENATER
- FORTH

GENEVA EBONE BOUNDARY SYSTEM

- SWITCH
- EARN
- ACOnet
- ARIADNEt
- ILAN
- EASINET

EBONE POLICIES

EBONE OPEN FOR ANY NETWORK ....

IF SIGNIFICANT BENEFITS TO THE EUROPEAN
R&D COMMUNITY

EBONE IS AIMED AT HAVING LESS
RESTRICTIONS THAN ANY OF THE REGIONAL
NETWORKS

EBONE BOUNDARY SYSTEM
PROVIDES BACKBONE ROUTING

REGIONAL BOUNDARY SYTEM
IMPLEMENTATION OF LOCAL POLICIES

EBONE TRANSIT POLICIES

EBONE TO US TRAFFIC

- GENERAL EBONE TO US CAPACITY SHOULD HAVE IDENTICAL POLICIES TO BE USEFUL
- NETWORKS PROVIDING TRANSIT CAPACITY TO SOME EBONE REGIONAL BUT NOT ALL CREATES PROBLEMS

EBONE TO PACIFIC AND ASIA TRAFFIC

- TRANSIT CAPACITY MUST NOT HAVE MORE RESTRICTIONS THAN THE EBONE ITSELF

3.3 NSFnet Report

Presented by

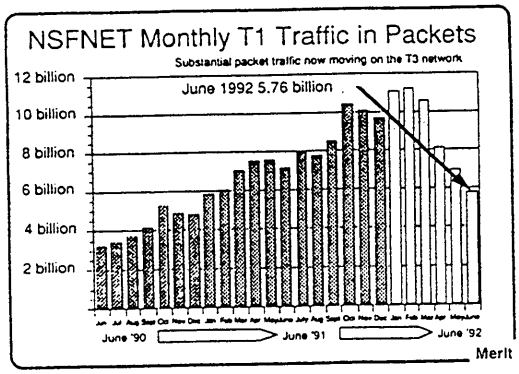
NSFNET Monthly T1 Traffic in Packets
Substantial packet traffic now moving on the T3 network
June 1992 5.76 billion
12 billion
10 billion
8 billion
6 billion
4 billion
2 billion
June '90
June '91
June '92
Merit

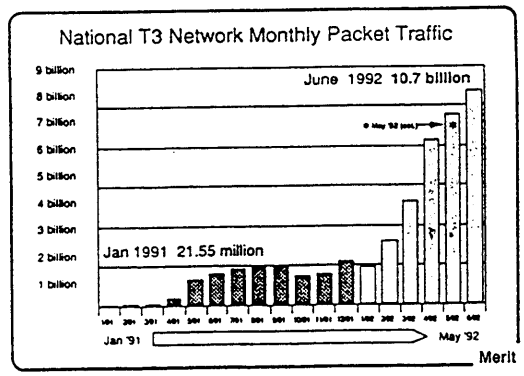
National T3 Network Monthly Packet Traffic
June 1992 10.7 billion
Jan 1991 21.55 million
9 billion
8 billion
7 billion
6 billion
5 billion
4 billion
3 billion
2 billion
1 billion
Jan '91
May '92
Merit

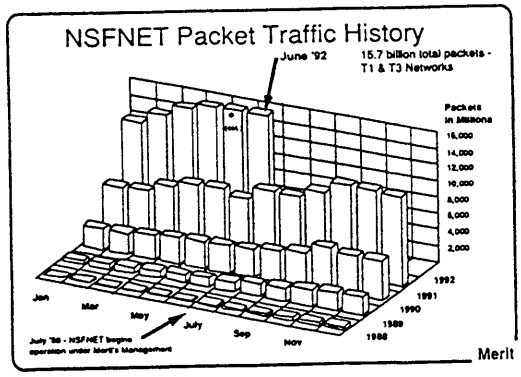
NSFNET Packet Traffic History
June '92
15.7 billion total packets - T1 & T3 Networks
Packets in Millions
Jan
Mar
May
July
Sep
Nov
1992
1991
1990
1989
1988
Merit

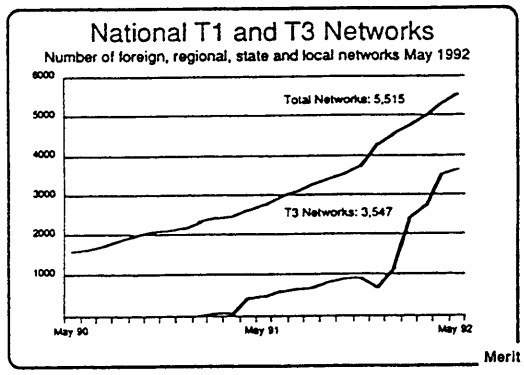
National T1 and T3 Networks
Number of foreign, regional, state and local networks May 1992
Total Networks: 5,515
T3 Networks: 3,547
6000
5000
4000
3000
2000
1000
May 90
May 91
May 92
Merit

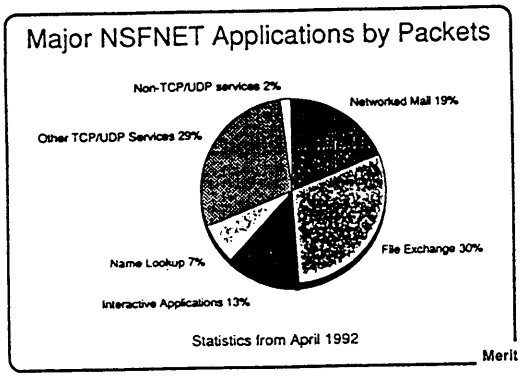
Major NSFNET Applications by Packets
Non-TCP/UDP services 2%
Networked Mail 19%
Other TCP/UDP Services 29%
File Exchange 30%
Name Lookup 7%
Interactive Applications 13%
Statistics from April 1992
Merit

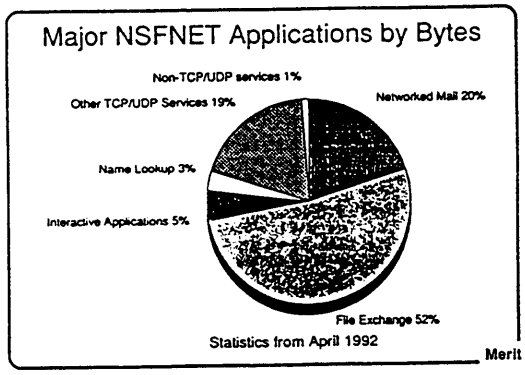
Major NSFNET Applications by Bytes
Non-TCP/UDP services 1%
Other TCP/UDP Services 19%
Networked Mail 20%
Name Lookup 3%
Interactive Applications 5%
File Exchange 52%
Statistics from April 1992
Merit

3.4 NSI Report

Presented by Jeff Burgan/NASA

NASA Science Internet Status Report

July 13, 1992
IETF Meeting, Boston

Jeffrey G. Burgan
NASA Science Internet Office
NASA Ames Research Center

Network Configuration

- 96 connected sites using Proteon routers
  - 41 running DECnet Phase IV in conjunction with TCP/IP
  - 26 running OSPF (1 area)
- NSFnet connections upgraded to T3 at FIXes
  - FIX-East (SURAnet, College Park, MD)
  - FIX-West (NASA Ames)
- International Link Status
  - Australia (512K)
  - Hong Kong (64K)
  - Korea (56K)
  - New Zealand (64K)

Upcoming Events

- Upgraded UK "fat pipe" (768Kbs)
- FIX-East connection upgraded to T3
- Backbone router upgrade
- "Virtual" upgrade of Hawaii link
- Fast packet deployment

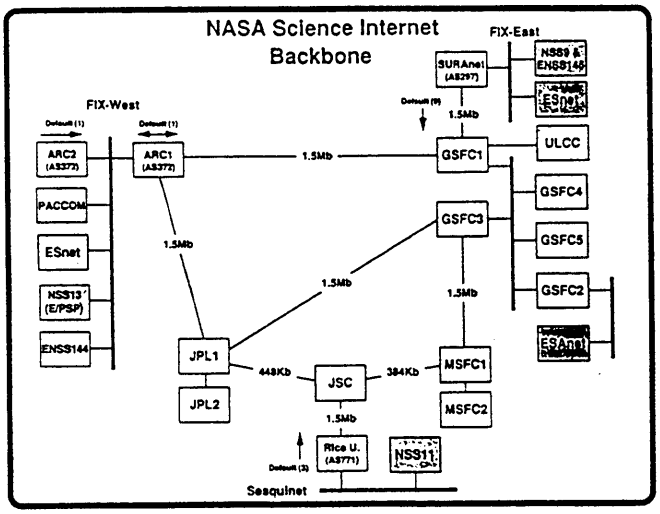

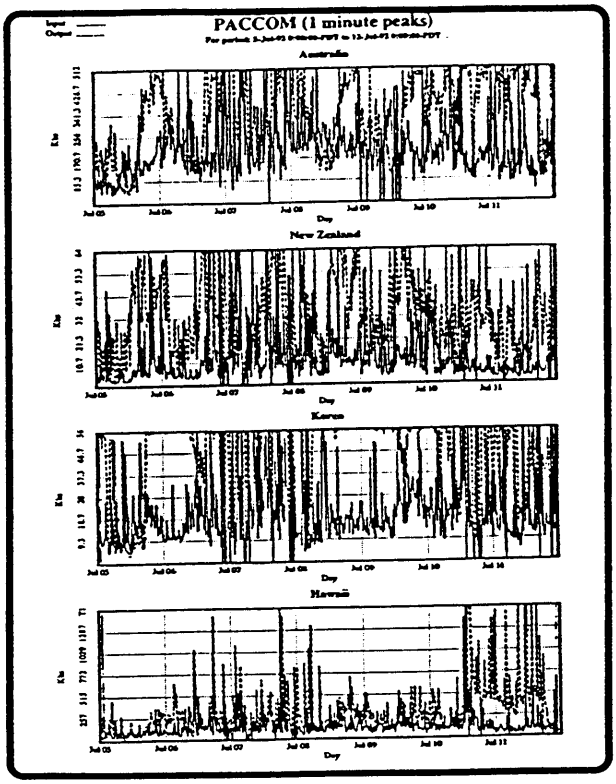

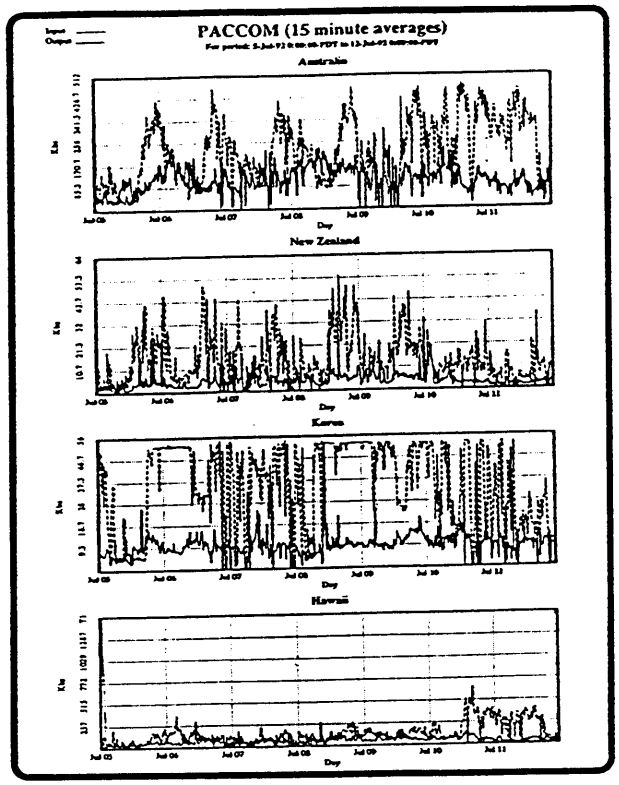

Network Status Report
July 13, 1992

T3 Network Status

- Phase-III RS/960 Upgrade Completed During 4/27-5/23
- T1 Backbone Traffic Cutover Proceeding
- Preparation for RS960 FDDI Upgrade During 8/92
- New Routing Software Enhancements

T1 Backbone Stability Improving

- Congestion Problems Reduced with T1->T3 Traffic Migration
- Route Scaling Problems Continue to Be Addressed

Other Proposed 1992 T3 System Enhancements

Phase-III T3 Network Upgrade Components

New Hardware Components

1. RS960 Serial Interface with HSSI

2. New T3 DSU Adapter
   - Serial Interface Card with C-bit Parity
   - High Speed Serial Interface (HSSI)

New Software Components

1. RS960 DS3 AIX Driver & Kernel Modifications

2. SNMP Monitoring Software
   - SNMP Daemon, DSU Proxy Subagent
   - T3 DSU Logging and Interactive Control Programs

3. RS960 Adapter Firmware

4. New RS960 Utility Programs for AIX Operating System
   - ifstat - Interface statistics
   - ccstat - On-Card Statistics

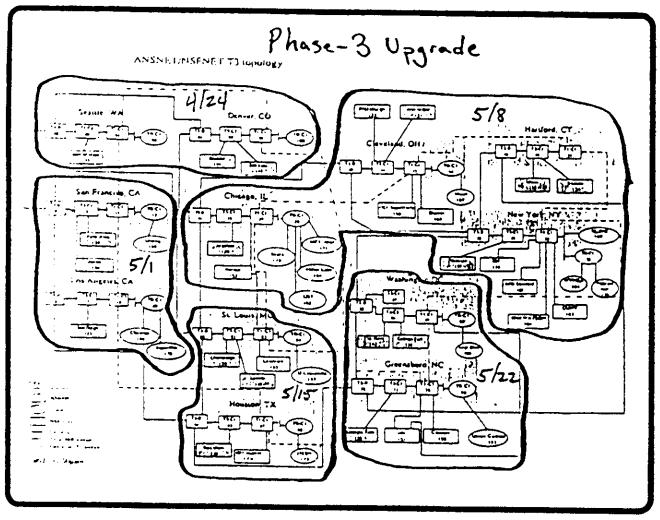

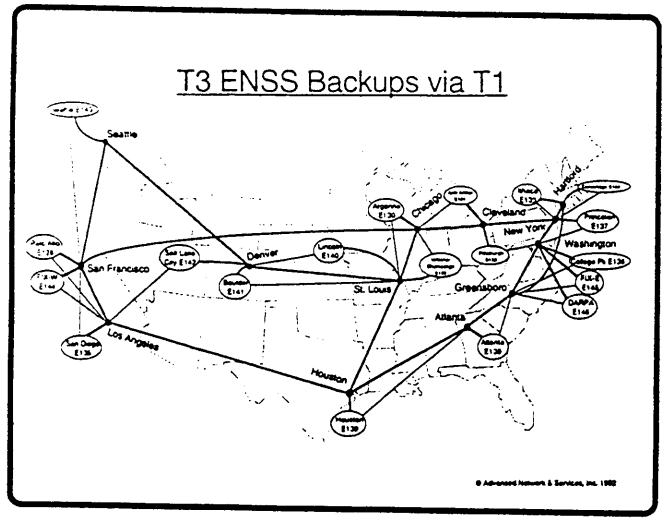

T3 Network Status

- o T3 Routing Software Enhanced
  - EGP & BGP Nets From Same AS Could Caused Black Holes
  - Unneccesary BGP Disconnects Fixed
    - Due to IGP Transients or Peer Routers Busy
    - Better Message Buffering
- o New Route Consistency Monitor Programs Installed
  - Examine Log Files For IBGP Disconnect Events
  - Examine Routing Tables For Self-Consistency
  - Used to Fix Operational Problems
  - Observed Chronic Route Flapping with Some Networks

Outstanding Problems

- o RS960 On-Card Memory Errors Due to Batch of Bad Memory
  - 7 Problems Identified So Far
  - Spare Cards To Be Upgraded with New Memory
- o ENSS T3 Link Failure Results in Black Hole
  - Due to Default Routes Pointing to 140.222
  - ICMP Redirects Ignored
  - Workaround: ENSS Will Not Advertise 140.222 if Isolated
- o FDDI Output Queue Packet Loss on ENSS128
  - Near Term: Controlled Load Splitting Across Two Interfaces
  - Long Term: Install RS960 FDDI Interface in August '92

NSF Midlevel T3 Traffic Migration Status
Networks Already Cutover

ASN	NSSN	Description
26	10	Cornell Regional Network
38	12	UIUC/NCSA Regional Network
81	17	CONCERT Gateway Network
86	9	SURANET Regional Network
114	11	Sesquinet
174	10	NYSERNET Regional Network
177	17	Merit Regional Network
181	1	Merit
195	6	SDSC Regional Network
200	13	BARRNET Regional Network
201	13	BARRNet Regional Network
232	4	Merit
233	17	Merit Regional Network
275	11	Sesquinet
279	18	SURANet out of Georgia Tech
280	11	Sesquinet Regional Network
281	8	NEARnet Regional Network
283	10	Cornell Gated Development Network
297	9	Nasa Science Network (SURANET connection)
560	8	NEARnet Regional Network
698	12	UIRnet
1206	5	PSCNET
1224	12	NCSA - Champaign, Illinois
1700	11	Sesquinet
1740	6	CERFnet
97	8	JvNCnet Regional Network
266	17	CICNET Regional Gateway
267	12	CICNET Regional Network (UIUC connection)
93	16	Midnet
372	13	NASA Science Network (West Coast) National Network
194	7	NCAR Regional Network
209	7	WestNet Regional network (East)
73	14	NorthWestNet Regional Network - U of Washington
685	14	NorthWestNet Regional Network *
701	9	Alternet Network *
702	8	Alternet at Boston *
68	7	Los Alamos Regional Network (Don Morris)
771	11	NASA Science Network at Sesquinet
1262	7	NASA Science Network at NCAR, Boulder, Colorado
210	15	Westnet Regional Network (West)
545	17	BBN-NET, Reston, VA 22091
101	14	NorthWestNet Regional Network

NSF Midlevel T3 Traffic Migration Status
Networks Still To Be Cutover

ASN	NSSN	Description
278	7	Mexican Networks
555	12	Minnesota Supercomputer Center Network (MSCNet)
590	10	EASInet Regional Network
697	10	EASINet-AS2
601	10	CA*net router in Toronto
602	14	CA*net router in Montreal
603	8	CA*net router in Quebec

FIX-E and FIX-W ASN's:

60	13	TWBNET regional network
274	9	TWBNET regional network
164	13	NSFNET-MAILBRIDGE gateway - scheduled 7/8
184	9	Second Mailbridge - scheduled 7/8
291	13	Energy Science National Network
293	9	Energy Science Network National Network
1240	13	USSPRINT - Palo Alto, CA
1238	10	ICM-NORDUNET
1800	9	ICM/Sprint (Fix/East) - partially cut over

T1 Network Stability Status

- Congestion Related Problems
  - Traffic Migration T1 -> T3
- LSP Packet Size Exceeded
  - Multiple EGP Peers at Midlevel Sites
  - Over 2000 Networks Announced at Individual Midlevel
  - T3->T1 Network Announcements During Traffic Migration
  - First Fix Did Not Work, Applying 2nd Fix
  - Will Be Fixed by Improving LSP Packet Coding Efficiency
- CPU Starvation on RCP RT/PC Nodes
  - Installing EAPC Cards In Several RT/PC Nodes
  - Optimized Routing Software for Performance

Other Proposed Network Changes in 1992

Proposed Software Enhancements

- BGP3 Support On T3 Network
- Connectionless OSI Network Protocol Support (CLNP, ES-IS)
- Migrate to Common IGP Within T3 System
  - Intra-Domain IS 10589 Routing (IS-IS)
- Future Inter-Domain Routing (BGP4, IDRP)
- Dual-mode Routing
  - Deploy Dual IS-IS and Dual IDRP (OSI and IP routing)
- Routing Software Being Implemented in GATED & Cisco

Network Infrastructure Enhancements

- IBM FDDI Adapter Upgrade
- T3 ENSS Backup Support via Redundant ENSS Connections
- Dismantling of T1 Backbone
  - Requires Full Cutover to T3, OSI Support, ENSS Backup
- Cisco T3 & FDDI Interface Support
- CNSS Reconfiguration

Current Network Status

T3 Network Stable As of 11/91

- Phase-3 T3 Upgrade Complete As of 5/23
- Reliability Continues to Be Very Good
- Performance (Throughput, Latency) Much Improved
- Majority of T1 Network Traffic Cutover to T3

T1 Backbone Stability Improved with Traffic Reduction

- Several New Routing/Performance Problems Due to Net Growth
- Route Computation Complexity Represents Ultimate Limit
- Plan to Dismantle T1 Once T3 Supports OSI, ENSS Backup

Chapter 4

Technical Presentations

4.1 Federal Network Council

Presented by Tony Villasenor/NASA

Bio: *Tony Villasenor from NASA Headquarters in Washington, is responsible for managing the NASA Science Internet and NASA's NREN; he is NASA's representative on the Federal Network Council and FCCSET subcommittees, and he chairs the FNC's Engineering and Operations Working Group.*

This presentation gives federal perspectives on IINREN/NREN status, including the government committees involved, nearterm milestones, relevant technical issues and concerns; stress on collaboration among government-industry-academia to achieve success.

Federal Network Council

Engineering & Operations Working Group

IETF

Boston
July 1992

Tony Villasenor @ NASA
EOWG Chair

Federal Network Council

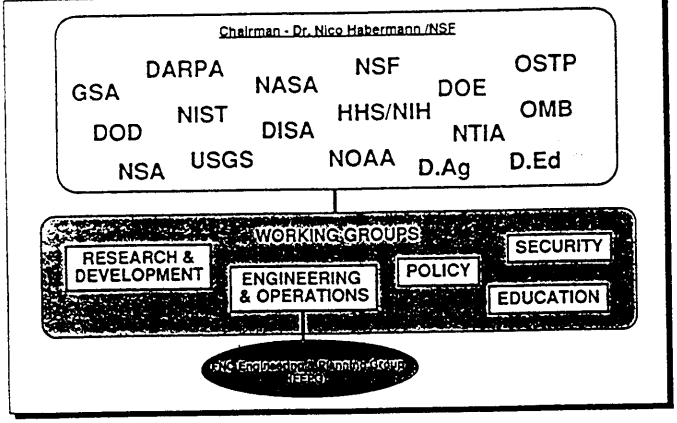

HPCC Committee Hierarchy

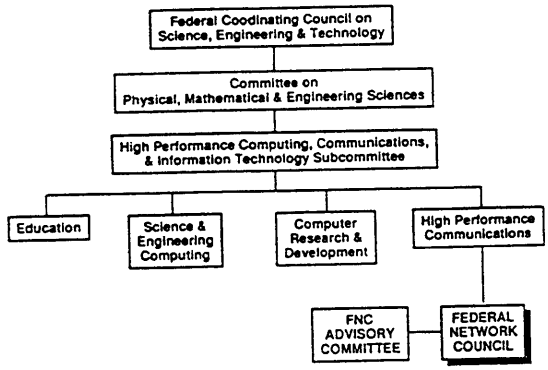

FNC Activities

Collaboration in federal networking
- NREN: requirements, management, implementations, etc.
- NSFNET Solication
- DOE & NASA solicitation
- DARPA research & development activities
- Interagency activities

Coordinated international projects
- CCIRN: Americas, Europe, Pacific
- Shared links to Europe, Pacific... Antarctica

FNC Advisory Committee inputs from academia & industry
- computer vendors & carriers

Engineering & Operations Working Group

MEMBERS

INITIAL MEMBERSHIP OF THE EOWG WOULD CONSIST OF BUT NOT BE LIMITED TO THE FOLLOWING, BECAUSE OF THEIR STRONG VESTED INTEREST IN THE ENGINEERING AND OPERATION OF THE INTERNET:

Tony Villasenor, NASA - Chairman
Steve Wolff, NSF
John Cavallini, DOE
Paul Mockapetris, DARPA
Jerry Linn, NIST

Note: Operations managers of key federal networks may be included as members. The EOWG will provide delegates to CCIRN meetings.

Inter-Agency Technical Collaborations

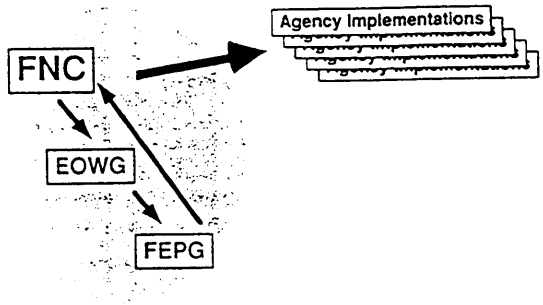

IMMINENT IINREN MILESTONES

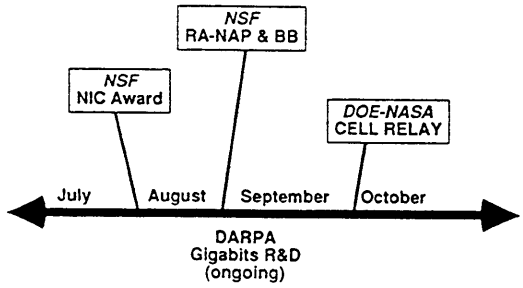

Some Concerns

"We're running out of IP address space!"

"Routing tables are exploding!"

Scoping the Address Problem

Statistics:	Total	Allocated	%
Class A	126	49	38%
Class B	16383	7354	45%
Class C	2097151	44014	2%

Registration Policies Under Consideration

Class A	No more allocations!
Class B	• Multi-national or multi-service provider • at least 15 subnets • more than 3000 hosts
Class C	As requested

LARGE ROUTING TABLES

ROUTERS!

Vision for the Future

Which Way to Go: IP or OSI ?

NREN

Answer: Maybe Neither !

Why Neither IP nor OSI ?

- Isochronous traffic (e.g. real-time video and voice)
- Gigabit performance with ATM/SONET technologies TBD!
- Policy support
  (i.e., more than current routing table implementations)
- Multicast (routing, addressing, forwarding, etc.)
- Resource reservation : 'guaranteed bandwidth' and predictable quality of service
- Accounting & Configuration Management
- Security

Note: The HPCC programs are leaning towards IP to achieve high performance networking.

Federal "Decision Making"

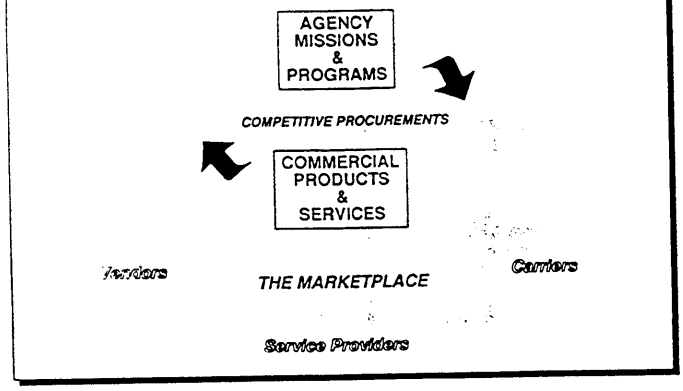

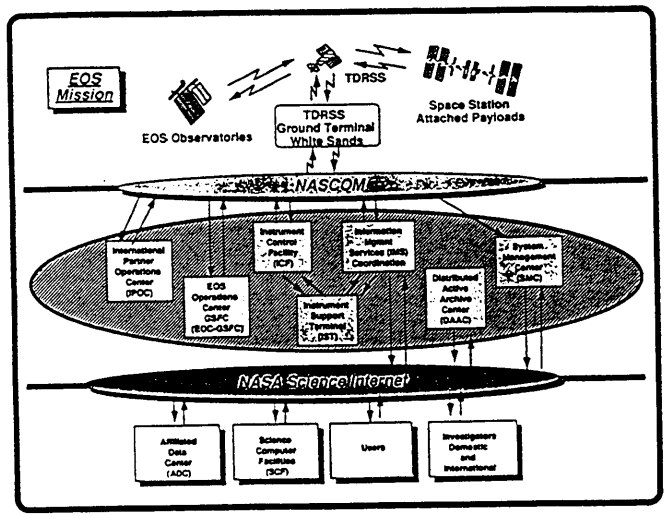

Working Together

Strategic Endeavor	*FNC Support*
NEAR TERM OPERATIONS CIDR, C#, etc.	NSF, others
RESEARCH IN NETWORK STRUCTURES Addressing, routing, etc.	DARPA, NSF, others
PRODUCTS & SERVICES Routers, network applications, etc.	ALL

4.2 The Futures of the Internet

Presented by Mitch Kapor/EFF

Bio:*Mitchell Kapor is the co-founder and President of the Electronic Frontier Foundation, an organization which works to develop and implement public policies to promote openness, diversity, and innovation in emerging electronic social environments. Mr. Kapor is also currently Chairman of the Commercial Internet Exchange Association (CIX). Previously, Mr. Kapor founded Lotus Development Corporation and served as its Chief Executive Officer, President and Chairman.*

THE ELECTRONIC FRONTIER FOUNDATION'S OPEN PLATFORM PROPOSAL

Until now the nation's telecommunications policy debate has largely been perceived as a struggle among entrenched commercial interests over who will control and dominate markets such as information services, manufacturing, and long distance service. We believe it is time to refocus the debate by seeking near-term technological, economic, legislative and regulatory solutions which will encourage the rapid development of a diverse information services market and help realize the democratic potential of new information media.

In the Fall of 1991, the Electronic Frontier Foundation was invited by Representative Edward Markey to testify before the House Subcommittee on Telecommunications and Finance on the subject of Bell company entry into the information services market. To address concerns that Bell entry into this market would reduce the diversity of information through anti-competitive behavior, EFF proposed the rapid deployment of a digital information platform, using existing technology and facilities, which could be made available to all on a ubiquitous, affordable, equitable basis. Our testimony suggested that narrowband Integrated Services Digital Network (ISDN) could be such a platform.

Narrowband ISDN, if offered nation-wide, and tariffed at affordable, mass- market rates, can offer end-to-end digital service without major infrastructure investments. This narrowband technology can also serve as a transitional telecommunications platform until national switched broadband access options become available early in the 21st century. With an ISDN platform in place, information entrepreneurs will soon be able to reach an expanded market in which to offer text, video, and interactive multimedia services. Public agencies, private communications, computer, and publishing firms, and even individuals will be able to access an inexpensive, widely available medium in which to publish and communicate electronically. Other technologies from outside the public telephone network may also come to play an important role in providing digital access, but because of the importance of the public switched telephone network, ISDN has a key role to play.

EFF believes that ISDN deployment and other developments in the public telecommunications infrastructure should proceed with the following goals in mind:

- make end-to-end digital service widely available at affordable rates;
- promote First Amendment free expression by reaffirming the principles of common carriage;
- ensure competition in local exchange services;
- foster innovations that make networks and information services easy to use;
- protect personal privacy; and
- preserve and enhance equitable access to communications media for all segments of society.

A robust, open telecommunications infrastructure is certainly important for the international competitiveness and economic health of our nation. But also, as people become more dependent on telecommunications services in their daily lives, the character of the evolving infrastructure and the laws which govern its operation will come to have a profound impact on politics, culture, education, and entertainment. Therefore, the steps that we take at this critical moment in the development of telecommunications technologies must be carefully considered.

ISDN is a platform which could stimulate innovation in information services in a way that will benefit much of the American public that currently has no access to electronic information services. Lessons from the personal computer industry can help guide telecommunications policy makers in the development of an information infrastructure. The desktop personal computer represented a revolutionary platform for innovation of the 1980's because it was affordable, and was designed according to the principle of open architecture, allowing numerous hardware and software entrepreneurs to enter the computer industry.

To bring the benefits of the information age to the American public in the 1990's, we need to build an open, ubiquitous digital communications platform for information services. Just as the personal computer brought access to computing power beyond large organizations, widely available ISDN can enable the citizen's access into the Information Age.

What is ISDN?

ISDN (Integrated Digital Services Network) is a technology designed for the public switched telephone network which allows low-cost communication in data, voice, video, and graphic media over the existing copper telephone network. ISDN is not an information service, but a transmission medium – a platform – for delivering and receiving information in a variety of formats. Crude data communication is possible over standard analog telephone lines now, but the fact that the existing transmission system was designed for voice, not for data, means that transmission rates are very slow, error rates are high, and equipment (modems) are difficult to use. Basic Rate ISDN offers transmission speeds fifteen to sixty times faster than most data transmission schemes now used on voice grade lines. More

the just the increased speed, what is important about ISDN is that it offers the minimum capacity necessary to carry full multi-media – voice, text, image, and video – transmissions.

ISDN is not a "field of dreams" technology. It is a fully-developed international standard that has been extensively tested in the United States and has already been implemented in the public switched telephone networks of other countries. Real applications have been demonstrated over ISDN lines. Major communications carriers have field-tested distance learning applications which allow students in classrooms all across a city to participate in multimedia presentations run by a teacher in a remote location. Inexpensive desktop and home video conferencing systems are now being introduced which run over ISDN lines. These applications have real value, but are only a small sample of what entrepreneurs will inevitably produce if ISDN were widely available. Yet, the promise of this service can only be realized if the local phone companies tariff and deploy the service.

Prospects for Near Term ISDN Deployment

EFF's Open Platform proposal for ISDN is a work-in-progress. We have received valuable comments and support from key players among the Regional Bell Operating Companies (RBOCs), interexchange carriers, information providers, and state public service commissions, all of whom believe that ISDN can play a crucial role in developing the information arena for the benefit of all today. To date, we have reached the following conclusions:

1. ISDN deserves a *second look* because it can meet many of the information needs of residential and commercial users long before a public, switched broadband network will be available.
2. ISDN can be made widely available within the next three to five years, without massive infrastructure investment or new technology development.
3. ISDN can and must be tariffed as a *basic service* at affordable rates.
4. ISDN is a critical and even necessary *transitional* technology on the path toward the future broadband national public network.
5. The benefits of other networks that are already important information distribution media can be enhanced by interconnection with ISDN.

More investigation of many issues is still required, especially the regulatory economics of deployment. Still, we are optimistic that ISDN is an important step along the path to the development of a telecommunications infrastructure that meets the diverse needs of the nation.

The Second Look

ISDN can meet many of the critical information needs of both residential and commercial users even without broadband capacity. ISDN is the only switched, digital technology

available today in the public switched network that can be deployed widely in the near term. For text-based data users and publishers, ISDN offers a dramatic advantage over data transmission technology currently used by individuals and small organizations. One of the two 64kbits/sec data channels available in the ISDN Basic Rate Interface can fax 30 typewritten pages of text in one minute, and send a 1000-word newspaper article in less than one second. Dramatic advances in video compression make transmission of videoconference images possible today, and all indications are that new compression algorithms will allow real-time transmission of VCR-quality video images in the near future. The Massachusetts Department of Public Utilities found, in the course of its recent investigation of ISDN, that "residential customers will benefit from the availability of significant enhancements to services such as home banking, library access, work at home, home health care monitoring, home shopping, and information access."(1)

Some telecommunications cognoscenti view the promise of narrowband ISDN as quite limited, because they are aware that ISDN has languished unimplemented for over ten years, and because they know that other copper- based transmission technologies offering much higher bandwidth are available. We are fully supportive of implementing higher capacity narrow band and broadband networks in the future, when technology and user demand make it possible.

The personal computer industry shows that raw power is not all that matters in a new technology. By about 1980, corporations already had good access to massive computational facilities at the institutional level through their mainframes and minicomputers. But individual workers had no effective direct access to those facilities. In practice, all the computing power didn't directly help the white-collar worker get her job done. Personal computers made a difference in the office and in the home because they were directly under the control of the individual, despite the fact that they were anemically under- powered. Similarly, there may be high data capacity at the institutional data network level already, but if individuals and small organizations can't connect with it, its value is limited. We must make tapping into the digital, switched network as easy as ordering a phone line for a fax. Just as PCs enhanced individual productivity, ISDN can enhance individual connectivity.

In this regard, we are encouraged by the fact that the computer industry has recently joined the debate on telecommunication infrastructure. With the growing recognition that the hardware and software they design will be severely limited by the lack of a nationwide switched, digital communications infrastructure, key players in the computer industry have lent their support to EFF's Open Platform Proposal as a transitional infrastructure strategy.

ISDN Means an Available Infrastructure Soon

ISDN can be made widely available in the near future without massive new infrastructure investment or new technology development. In sharp contrast to fiber optic-based broadband technologies, only modest infrastructure investment is required. Digital central office switches are required for ISDN (2), but with the Bell companies aggressive deployment

of a fully-digital switching and signaling system (Signaling System Seven), the bulk of the infrastructure necessary to support ISDN is already installed or planned.(3) Some Bell companies such as Bell Atlantic and Ameritech plan to have over 70ISDN-ready by the end of 1994. Other companies, however, project deployment rates as low as 21are expected to be capable of carrying ISDN calls by 1994.(4) (See Appendix A)

Many segments of the telecommunications industry are engaged in a concerted effort to make nation-wide ISDN deployment a reality. Problems that haunted ISDN in the past, such as lack of standard hardware and software protocols and corresponding gaps in interoperability, are being addressed by National ISDN-1. This a joint effort by Bell companies, interexchange carriers, and switch manufactures, and Bellcore, is solving major outstanding standards problems. By the end of 1992, a single hardware

standard will make ISDN central office switches and customer premises equipment interoperable, regardless of which vendor made the equipment. Following National ISDN- 1, National ISDN-2 will address standards problems associated with ISDN Primary Rate Interface (PRI), a switched 1.5Mbit/sec service with 23 separate 64kbit/sec data channels and one 64kbit/sec signaling channel.

Led by Bellcore, the communications industry has a nationwide demonstration of real, off-the-shelf, ISDN services planned for November 1992, called TRIP'92. A variety of local and national ISDN services will be demonstrated on a working ISDN network covering twenty cities around the country. TRIP'92 will show that Bell companies, long distance carriers, and information providers can work together to provide the kind of ubiquitous, standards-based service that is critical to the overall success of ISDN.

Additional interconnection problems do remain to be solved before ISDN is truly ubiquitous. Among other things, business arrangements between local Bell companies and interexchange carriers must be finalized before ISDN calls can be passed seamlessly from the local exchange to long distance networks.

ISDN Must be Priced Affordably

ISDN can and must be tariffed as a basic service at affordable, mass-market rates. If ISDN is to be a platform that spurs growth and innovation in the information services market, it must be priced affordably for the average home and small business user. Here, the telephone industry has a valuable lesson to learn from the computer industry. The most valuable contribution of the computer industry in the past generation is not a machine, but an idea–the principle of open architecture. Typically, a hardware company (an Apple or IBM, for instance) neither designs its own applications software nor requires licenses of its application vendors. Both practices were the norm in the mainframe era of computing. Instead, in the personal computer market, the hardware company creates a "platform"–a common set of specifications, published openly so that other, often smaller, independent firms can develop their own products (like the spreadsheet program) to work with it. In this way, the host company takes advantage of the smaller companies' ingenuity and creativity.

Platform services, even if they are ubiquitous, are useless unless they are also affordable to American consumers. Just as the voice telephone network would be of little value if only a small fraction of the country could afford to have a telephone in their home, a national information platform will only achieve its full potential when a large majority of Americans can buy access to it. Therefore, the tariffs adopted by state public utility commissions are critical to the success or failure of ISDN.

Since few states have adopted single-line business and residential ISDN tariffs, there is a window of opportunity to establish pricing principles for ISDN which make it viable as a mass-market service. The Massachusetts Department of Public Utilities (DPU) recently completed proceeding should serve as a valuable example to other states. The Massachusetts regulators found that ISDN is a "monopoly, basic service that has a potentially far- reaching and significant role in the telecommunications infrastructure of the Commonwealth."(5) The DPU also recognized that the "risks of pricing the service too high are of much greater concern... [because] high rates could discourage the development of new ISDN-dependent technologies and their applications."(6) The final tariff approved has a monthly access charge of $13.00 for single line residential service and usage sensitive fees of 2.6 cents for the first minute and 1.6 cents for each additional minute. After much dispute, New England Telephone (NET) based the usage sensitive component of the tariff on measured voice rates already in place in Massachusetts. We believe that NET's decision to link prices to existing basic voice rates is an important signal to other LECs and other state commissions that low-priced ISDN service is indeed possible.

Studies by experts in the field of regulatory economics indicate that ISDN can be priced affordably. Dr. Lee Selwyn found, based on data from the Massachusetts proceeding, that the average monthly price for ISDN service should be approximately $10.(7) An analysis of ISDN deployment by a leading consumer advocate also indicates that ISDN can be offered at a relatively low cost to consumers. Dr. Mark Cooper, Research Director of the Consumer Federation of America, found that average ISDN monthly costs are now at roughly $7.50, and can be expected to decline to $4.50 in the near future.(8)

To encourage widespread use of ISDN, it must be priced at or near the price levels already in place for basic voice services. ISDN line charges will be somewhat higher than analog voice services because there are some additional one-time capital costs associated with offering ISDN service, but basing prices on voice telephone rates is possible and rational from a regulatory standpoint.

The digital switches which carry ISDN calls treat voice and data calls in exactly the same manner. A five minute data call uses no more or less switching resources than a five minute voice call, so their pricing should be equivalent. Some states may chose to tariff ISDN only with measured (usage sensitive) rates, while others may also want to adopt a flat rate scheme similar to that which exists for residential voice services. The economics of this issue need more study, but we believe that both options have arguments in their favor.(9)

Current prices for ISDN telephones, data links, and in-home network terminators are high. An ISDN telephone with voice and data interfaces costs roughly $1000. If these price

levels persist, many small scale users will never enter the market. However, with increased demand, ISDN terminal appliance prices can be expected to follow the steep downward curve of VCRs and PCs prices. When first introduced, VCRs cost well over $1000, but now sell below $200 for a basic unit.

Ill-considered pricing policy could, alone, cripple ISDN's chances for success. We are hopeful that Bell companies with more aggressive deployment plans will file such residential tariffs and set a precedent for progressive, mass-market pricing that will make ISDN affordable. In any event, legislative or regulatory action may be necessary to guaranty affordable rates and widespread availability of ISDN around the country.

ISDN: A Critical Transitional Technology

ISDN is a critical transitional technology on the road to a nation-wide public broadband network. ISDN is not a permanent substitute for a broadband network, but it is a necessary transitional technology on the way to public switched broadband networking. Though some might like to leap directly to a broadband network, the entire telecommunications and information industry still has much to learn about designing a broadband digital network before it can be implemented.(10) Though a first generation of broadband switches are now being introduced, many basic questions still remain about the most appropriate design for a broadband network that can replace or be built on top of the analog telephone network. These questions are impossible to answer without experience in the ways that people will use a public, digital switched network.

Some are reluctant to make any investment in ISDN because it is perceived as old technology. But this is not an either/or choice If implemented at prices that encourage diverse usage, ISDN will provide important new services to all segments of society, and offer vital perspectives on how to design the next generation of public, switched broadband networks.

ISDN Enchances Extant Networks

The benefits of other networks that are already important information distribution media can be enhanced by interconnection with ISDN. The public switched telephone network is a critical, central part of the nation's telecommunications infrastructure, so ISDN has a vital role to play in the overall information infrastructure. In addition to being an information platform itself, ISDN can interconnect with other networks that offer a variety of information resources. Cable television systems, which already provide broadband connections to 60pass by 90ISDN, cable systems could develop interactive video applications. The Internet, an international packet network that serves universities, government organizations, and an increasing number of commercial enterprise, has over two million users and access to vast archives of information. Wireless transmission systems such as PCS (Personal Communications Systems) could also serve as open platforms for information services.

Guiding Communications Policy Principles

The public switched telephone network is just one part of what we call the National Public Network, a vibrant web of information links that will come to serve as the main channels for commerce learning, education, politics, social welfare, and entertainment in the future. With or without ISDN, the telephone network is undergoing dramatic changes in structure, scope, and in its growing interrelationship with other communications media. These changes should be guided by a public policy vision based on the following principles.

- Create an Open Platform for Innovation in Information Services by Speedily Deploying a Nation-wide, Affordable ISDN.

  To achieve the information diversity currently available in print and broadcast media in the new digital forum, we must guaranty widespread accessibility to a platform of basic services necessary for creating information services of all kinds. Such a platform offers the dual benefit of helping to creating a level playing field for competition in the information services market, and stimulating the development of new services beneficial to consumers. An open platform for information services will enable individuals and small organizations, as well as established information distributors, to be electronic publishers on a local, national, and international level.

- Promote First Amendment Free Expression by Affirming the Principles of Common Carriage.

  In a society which relies more and more on electronic communications media as its primary conduit for expression, full support for First Amendment values requires extension of the common carrier principle to all of these new media. Common carriers are companies which provide conduit services for the general public. The common carrier's duties have evolved over hundreds of years in the common law and later in statutory provisions.

  The rules governing their conduct can be roughly distilled in a few basic principles. Common carriers have a duty to:
  - provide services in a non-discriminatory manner at a fair price,
  - interconnect with other carriers, and
  - provide adequate services.

  The public must have access to digital data transport services, such as ISDN, which are regulated by the principles of common carriage.

  Unlike arrangements found in many countries, our communications infrastructure is owned by private corporations instead of by the government. Therefore, a legislatively imposed expanded duty of common carriage on public switched telephone carriers is necessary to protect free expression effectively. A telecommunications provider under a common carrier obligation would have to carry any legal message regardless of its

content whether it is voice, data, images, or sound. For example, if full common-carrier protections were in place for all of the conduit services offered by the phone company, the terminations of "controversial" 900 services such as political fundraising would not be allowed, just as the phone company is now prohibited by the Communications Act from discriminating in the provision of basic voice telephone services. As a matter of law and policy, the common carriage protections should be extended from basic voice service to cover basic data service as well.

- Ensure Competition in Local Exchange Services.

  The divestiture of AT&T in the early 1980s brought with it various restrictions on the kinds of markets in which the newly created local Bell companies were allowed to compete. Many consumer and industry groups are now concerned that as these judicially-imposed restrictions are lifted (know as the MFJ), the Bell companies will come to dominate the design of the emerging National Public Network, shaping it more to accommodate their business goals than the public interest. The bottleneck that Bell companies have on local exchange services critical to information providers can be minimized by unbundling these services and allowing non-Bell company providers to offer them in competition with Bell companies.

  The post-divestiture pattern of providing long distance service offers us a valuable lesson: a telecommunications network can be managed effectively by separate companies–even including bitter opponents like AT&T and MCI–as long as they can connect equitably and seamlessly from the user's standpoint. Together with the open platform offered by ISDN, unbundling and expanded competition is a key to ensuring equitable access to Bell company facilities needed for information service delivery.

- Protect Personal Privacy.

  As the telecommunications infrastructure evolves, there are increasing threats to both communications privacy and information privacy. Strong government intervention will, at times, be necessary to protect people's constitutional right to privacy. Careful thought must also be given to the appropriate use of search warrants and wiretap authorizations in the realm of new electronic media. While new technologies may pose some difficult challenges to law enforcement, we must protect people's constitutionally- guaranteed right to be free from "unreasonable searches and seizures." Fundamental civil liberties tenets are at stake as long-standing constitutional doctrine is applied to new technologies.

  The privacy of telephone conversations and electronic mail is already protected by the Electronic Communications Privacy Act. However, communications in other media, such a cellular phone conversations, can be intercepted using readily available technology by private third parties without the knowledge or consent of the conversants. In addition to this, however, we believe that technological advances should be used to help people protect their own privacy and exercise more control over infor-

mation about themselves. In general, citizens should be given greater control over information collected, stored, and disseminated by telephone companies and information providers. As the public outcry over Caller ID demonstrates, citizens want and deserve to have adequate notice about what information is being collected and disseminated by communications firms and must be able to exercise informed consent before information collected for one purpose can be used for any other purpose.

- Make the Network Simple to Use.

  One of the great virtues of today's public switched telephone network, from a user's perspective, is that it operates according to patterns and principles that are now intuitively obvious to almost everyone. As this network grows beyond just voice services, information services that become part of this network should reflect this same ease- of-use and accessibility. The development of such standards and patterns for information services is vital, not just because it helps makes the network easier to use, but also because it ensures an open platform for information providers. However, standards development will be ad hoc and even chaotic at first. Numerous standards may be tried and found inadequate by users before a mature set of standards emerges. Congress and government regulatory bodies may need to set out the ground rules for standards planning in order to ensure that all interested parties have an equal voice, and the resulting standards should be closely analyzed to make sure that they reflect public needs. But, direct government involvement in the process should be avoided if possible.

- Preserve and Enhance Socially Equitable Access to Communications Media.

  The principle of equitable access to basic services is an integral part of nation's public switched telephone network. From the early history of the telephone network, both government and commercial actors have taken steps to ensure that access to basic voice telephone services is affordable and accessible to all segments of society. Since the divestiture of AT&T, many of the constituent parts of the "social contract" for universal service have fallen away. Re- creation of old patterns of subsidy may no longer be possible nor necessarily desirable, but serious thought must be given to sources of funds that will guaranty that the economically disadvantaged will still have access to basic communications services.

  The universal service guaranty in the Communications Act of 1934 (11) has, until now, been interpreted to mean access to "plain old telephone service" (POTS). In the information age, we must extend this guaranty to include "plain old digital service." Extending this guaranty means ensuring that new basic digital services are affordable and ubiquitously available. Equity and the democratic imperative also demand that these services meet the needs of people with disabilities, the elderly, and other groups with special needs.

  Failure to do so is sure to create a society of "information haves and havenots."(12)

Conclusion

The path toward ISDN deployment requires that cooperation of numerous public and private sector organizations and political constituencies. National policy direction is needed to ensure that the necessary ubiquity and interconnection of service providers is achieved. Federal policy makers in Congress and the Federal Communications Commission will also have to consider the appropriate regulatory role for guidance of a new national resource: the information infrastructure. State public service commissions will be at the forefront of establishing pricing policy for ISDN service. The success of residential applications for ISDN will depend heavily on the PUCs' approach to ISDN pricing.

The communications industry – including the Bell Companies, the interexchange carriers, equipment manufacturers – all have cooperative roles to play in making ubiquitous ISDN a reality. The computer industry is a new, but critical player in telecommunications policy. Many of the innovative products and services to take advantage of ISDN will likely come from the computer community.

In the policy arena and in relations with industry, many public interest advocacy organizations have a vital role to play in ensuring that new technologies are implemented and regulated in a way that promotes wide- spread access to new media and preserves the fundamental guarantees of affordable, universal service.

The Electronic Frontier Foundation is working to solicit comments, support, and criticism from all of these constituencies. This version of the Open Platform Proposal has been much improved with the help thoughts and reactions from many concerned parties. We welcome more comments from all who are concerned about the development of the telecommunications infrastructure.

For More Information Please Contact:

Mitchell Kapor
President
Electronic Frontier Foundation
155 Second St.
Cambridge, MA 02141
617-864-0665
mkapor@eff.org

Daniel J. Weitzner
Communications Policy Analyst
Electronic Frontier Foundation
666 Pennsylvania Ave, SE
Washington, DC 20003
202-544-9237
djw@eff.org

Appendix A: ISDN Deployment Data

Regional Bell Operating Company ISDN Deployment Plans Through 1994 (Numbers in Thousands)

Regional Bell Operating Co.	Total Lines	Lines Access w/ ISDN access	ISDN %
Ameritech	16,410	11,400	70%
Bell Atlantic	18,600	16,200	87%
BellSouth	20,000	10,500	52%
NYNEX	16,360	5,100	31%
Pac Telesis	15,900	10,900	69%
SW Bell	13,600	2,900	21%
US West	14,100	8,300	59%
TOTAL	114,970	65,300	56%

Source: Bellcore Report SR-NWT-002102, ISDN Deployment Data, Issue 2, June 1992.

Note: This table does not include deployment data for independent telephone companies.

Notes

- (1) Mass. D.P.U. 91-63-B, p. 86-7. See Appendix B for an overview of the Massachusetts proceeding.
- (2) In central offices where digital switches have not yet been installed, ISDN can still be provided at lower cost than by installation of special "switch adjuncts".
- (3) Though the Bell companies are not required to install Signaling System Seven, it is the only practical way that they can meet new FCC requirements for 800 number portability. See Memorandum Opinion and Order on Reconsideration and Second Supplemental Notice of Proposed Rulemaking, FCC Docket 86-10, Released September 4, 1991.
- (4) See FCC Docket 89-624 and Bellcore Special Report SR_NWT-002102, ISDN Deployment Data, Issue 2, June 1992.
- (5) ISDN Basic Service, Mass. D.P.U. 91-63-B, p. 34 (February 7, 1992).

- (6) Id. at 86.

- (7) L. Selwyn, A Migration Plan For Residential ISDN Deployment, April 20, 1992 (Prepared for the Communications Policy Forum and the Electronic Frontier Foundation).

- (8) M. Cooper, Developing the Information Age in the 1990s: A Pragmatic Consumer View, June 8, 1992. See p. 52.

- (9) Since the average length of a data call may be longer than the average voice call, the flat rate for ISDN would have to be adjusted upward to reflect added load on central office switching systems. However, the mere fact that data lines may remain open longer does not preclude a flat rate, non-usage- sensitive tariff.

- (10) The most optimistic BOC estimates on fiber deployment promise ubiquitous fiber optic cable in roughly 20 years.

- (11) 47 USC 151, et seq.

- (12) Modified Final Judgment: Hearings Before the Subcommittee on Telecommunications and Finance of the House Committee on Energy and Commerce, 101st Cong., 1st Sess. 2 (1989) (Opening Statement of Chairman Markey). Chairman Markey set the following goal for the development of new information services: to make [information services] available swiftly to the largest number of Americans at costs which don't divide the society into information haves and have nots and in a manner which does not compromise our adherence to the long-cherished principles of diversity, competition and common carriage.

4.3 Pip: The 'P' Internet Protocol

Presented by Paul Tsuchiya/Bellcore

PIP is an internet protocol that scales, encodes policy, is high speed, allows multiple defaults routing, has a compact header, allows strong firewalls between domains, makes address administration easier, makes mobility easier, and does multicast.

PIP has several significant differences from conventional internet protocols. This talk gives a brief high-level overview of the basic concepts of PIP. A more detailed talk was given at the PIP Birds of a Feather Session.

I am proposing PIP as an alternative to the two "medium term" proposals that emerged from the ROAD (Routing and Addressing) Group, to deal with the dual IP problems of scaling and address depletion.

PIP:
because
Only Surviving Is
not enough

PAUL F. TSUCHIYA
BELLCORE

Pip: 'P' Internet Protocol

"Next generation" IP protocol

Proposed replacement for IP

- IP Version 7

Current status:

- Basic ideas established
  - Overview paper
  - Drafty paper with additional detail
- Lots of work still to do

Pip

Does everything IP currently does

- Except fragmentation
  - Can be added

Therefore, can take advantage of much existing technology

- Routing algorithms .
- Host algorithms

This necessary for transition

Makes early fielding of Pip possible

Pip

But, can do much more than IP

- Advanced features can be evolved
  - Because basic packet format doesn't change even when new features added
- This not true (or MUCH less true) with CLNP
- Perhaps also true for IPAE
  - Which hasn't specified new address

Regarding IAP Decision for CLNP

Melt-down can be delayed more that IAB is designing for

- CIDR
- Address Reuse
  - Ugly, but will allow almost indefinate delay of melt-down

Therefore, worth waiting a bit and choosing/designing high quality solution

- Allow for advanced routing such as Nimrod or Unified

Pip is worth waiting for

- And, don't need to wait that long

Pip Feature Set

Scales to virtually unlimited number of systems

- Hierarchical addresses

Multiple simultaneous "address" formats

- Not just multiple branches of a single address tree, as in NSAPs, but:
- Hierarchical addresses, VCIs, multicast, source routes, policy routes...

Fast routing table lookup

- Routing lookup mechanics are independent of "address" format
- Routing lookup time bounded by hierarchy depth

Pip Feature Set (cont.)

Compact header

- Less than CLNP (for equivalent semantics)

General "tagging" function

- QOS, packet "coloring", etc.
- Routers as well as hosts can tag packets

Separation of "routing" and "identification" information

- Simplifies mobility
- Allows all kinds of "tricks" to be played with routing
  - For instance, isolation from inter-domain addressing inside stubs

Pip Design Philosophy

Reduce host and router functions to their most basic elements

- That is, the real ***mechanics*** of handling an IP packet, not some high level semantics that humans are comfortable with

Design IP header whose fields directly trigger those basic functional elements

- That is, don't design IP header, and then design IP engine around it...
- Instead, design ***basic*** IP engine, then design IP header that drives that engine

Basic Internet functions:

Identification

- Which "entity" sends and receives,
- independent of the entity's location

Routing

- How to get from sending to receiving entity

Other

- I call "handling"

IP "mismatch" of header to function

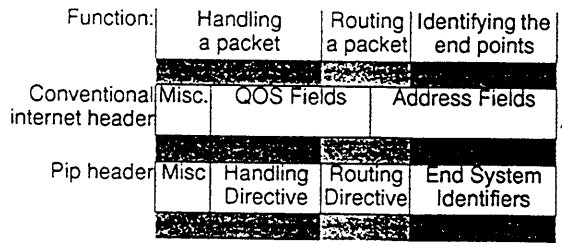

Some Practical Consequences of IP/CLNP mismatch

Routing scope limited

- Hard to change semantics of routing information
  - for instance, adding policy information to header
- Hard to introduce QOS
  - Because routers don't look at QOS field

Mobility awkward

- Because address does both routing and identification

Problem is, IP and CLNP contain a certain limited set of semantics, rather than basic host and router functional parts

Pip Header Again

Misc	Handling Directive (HD)	Routing Directive (RD)	End System Identifiers (ESI)

End System Identifiers (ESI)

- ***Only*** identify source and destination (flat identifiers)

Routing Directive (RD)

- ***Only*** part of packet that influences routing (next hop) decision

Handling Directive (HD)

- Various handling functions that do not influence routing decision

Some Routing Fundamentals

All routing information in an internet header is basically a:
Loose Source Route

For instance, IP address:

IP Address		
Network	Subnet	Host

This nothing more than a Loose Source Route

- o First route to Network, then to Subnet, then to Host

Some Routing Fundamentals (cont)

Difference between IP address and real Loose Source Route is lack of "Active Field Indicator" in IP address

- o To parse IP address, must work from left to determine "active field"

Alternative IP Address format:

	Loose Source Route		
Active LSR Field Indicator	Field 1 (Network)	Field 2 (Subnet)	Field 3 (Host)

This serves same function of IP address, but is more general

Routing Directive (RD)

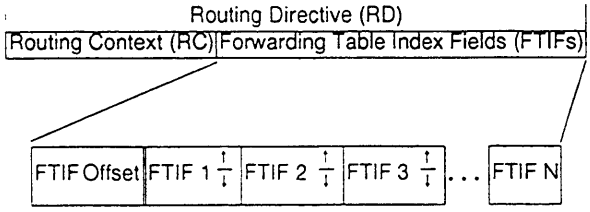

Routing Context (RC)
(formally Logical Router)

Single field

- o treated as flat by forwarding function
- o control algorithms understand content

Logically speaking, determines context within which router is acting

- o QOS type, Packet color, Hierarchy level, Address type, etc.

Mechanistically speaking, chooses one of multiple forwarding tables

All routing contexts encoded in this single field

- o Unlike CLNP, where multiple QOS-type fields required to determine context

Forwarding Table Index Fields
(formally Routing Hint Fields)

Once Routing Context (and therefore correct forwarding table) determined;

- o The FTIF Offset is used to isolate correct FTIF,
- o The FTIF is used as direct index into forwarding table

Each FTIF has two parts:

- o Value
- o Relator (up, down, none)
  - o Used to indicate whether next FTIF is hierarchically up, down, or unrelated to current FTIF

Value and Relator concatenated is used as index into forwarding table

Example: Plain Old Hierarchical Addresses

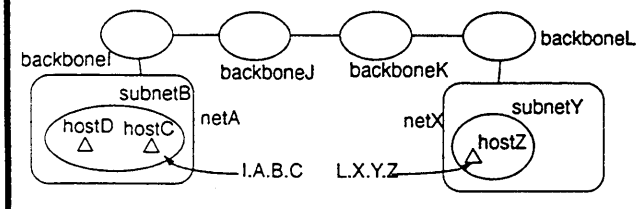

FTIF Offset	Source Address				Destination Address			
5	C↑	B↑	A↑	I-	L↓	X↓	Y↓	Z

Plain Old Hierarchical Addresses (cont)

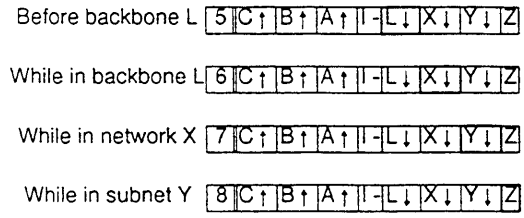

Example: Policy Route

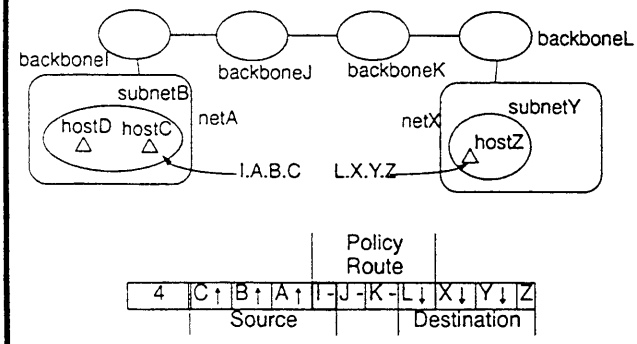

Example: hostC to hostD

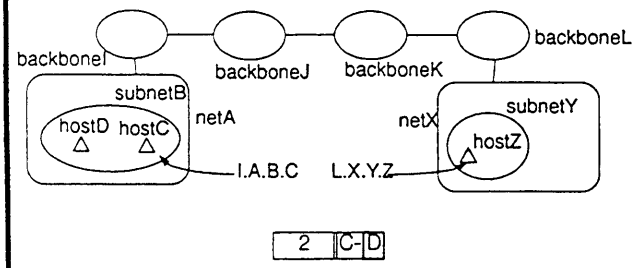

Where to go from here

First, design and install "Basic" Pip

- Minimum number of modifications to existing algorithms to make Pip work
- This gets us on-line with new header format as soon as possible
- Only differences from current operation:
  - Extra layer of hierarchy (for scaling)
  - Multiple defaults routing
  - Stub isolation from inter-domain address conventions

At same time, work on advanced features

- Nimrod
- Unified

Near-term Requirements for Basic Pip

Need following specs:

- Pip protocol
- Router Operation
  - PipBGP, PipOSPF
- Host Operation
  - Pip Configuration Protocol
- DNS Operation
- Transition (IP <—> Pip) Gateway

Need (public domain code) prototypes corresponding to four boxes above

Basic-Pip Working Group

- Pip BOF this week
  - Wednesday night session
  - Will give many detailed examples of Pip operation

Hope to establish "basic Pip" working group

- Looking for small but enthusiastic membership
- Goal to specify, prototype, and experiment with basic Pip
  - Basic means not much more sophisticated than current IP
  - Use existing routing protocols
- Should be able to establish basic Pip fairly quickly

Later Working Groups Evolve Advanced Pip Features Over Time

Policy

Mobility

Flow Setup

Multicast

Other

- o Byzantine Routing
- o Alternate Path or Multi-path Routing
- o Etc.

4.4 DARTnet: A Progress Report

Presented by Bob Braden/ISI

Bio:*Bob Braden started to design ARPANET host software in 1970, and beginning in 1978 he developed one of the research prototype TCP/IP implementations. He is a charter member of the IAB and chairs the End-to-End Research Group of the IRTF. Bob is a project leader at the USC Information Sciences Institute in Marina del Rey, CA.*

DARPA RESEARCH TESTBED NETWORK:
A PROGRESS REPORT

Bob Braden
USC Information Sciences Institute

Boston IETF Meeting
July 16, 1992

DARTnet Progress Report — July 1992

Topics:

- Overview of DARTnet
- The Research Program
- Recent Progress
- Future plans

WHAT IS DARTnet?

DARPA Research Testbed Network

DARTnet is a DARPA-funded research network, to support experiments that require dedicated facilities, e.g., the ability to change the packet switching software.

- "A NETWORK WE'RE ALLOWED TO BREAK".

WHY IS DARTnet IMPORTANT
TO IETF AND THE INTERNET?

Research is needed to steer the future technical evolution of the Internet.

There are many important gaps to be filled, especially:

- Resource reservation for real-time service
- Advanced routing

DARTnet is an essential experimental facility for research related to the future of the Internet.

ORGANIZATIONS

SPONSORSHIP & SUPPORT:

- DARPA – *primary funding*
- Sun Microsystems – *hardware*
- NSF, DOE, Xerox, Sun, Bellcore – *research support*

RESEARCH AGENDA:

- End-to-End Research Group and
- Autonomous Networks Research Group of the IRTF.

DARTnet Research Community

ROUTER SITES:

BBN	Bolt Beranek & Newman, Inc.
Bellcore	Bell Communications Research [future]
ISI	USC Information Sciences Institute
LBL	Lawrence Berkeley Laboratories
MIT	MIT Laboratory for Computer Science
PARC	Xerox Palo Alto Research Center
SRI	SRI International
Sun	Sun Microsystems, Inc. [future]
UDel	Univ of Delaware

OTHERS:

Mitre	Mitre Corporation
UMass	Univ of Massachusetts
USC	Univ of Southern California

DARTnet Router Topology 7/92

T1 (1.5Mbps) Lines: Cross-country spine + tail circuits.

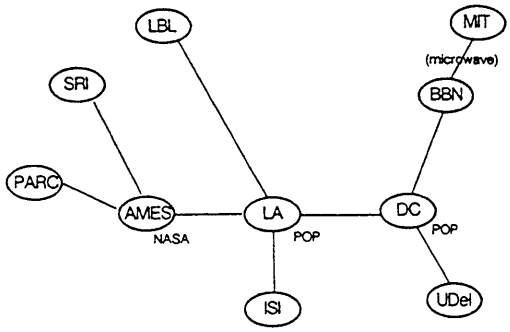

DARTnet Router Topology: Projected

T1 (1.5Mbps) Lines: Cross-country spine + tail circuits.

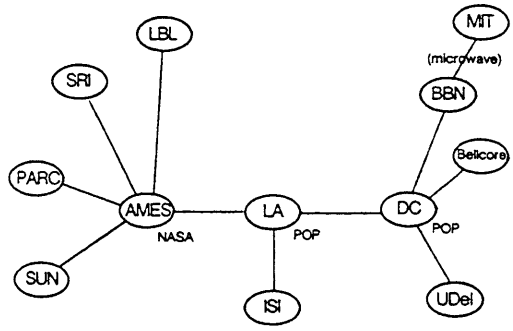

TYPICAL EXPERIMENTAL SITE

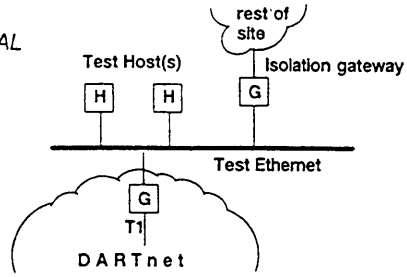

ROUTERS:

===> Requirement: open and programmable hardware

- **Sun Sparcstations** w/o displays

ROUTER OPERATING SYSTEM

- Requirement: an open, widely known system,
  to facilitate academic research and technology transfer.

  Choice: **BSD4.x**

  Interim: **Sun OS 4.1** *plus:*
  - 4.4BSD network code
  - NTP clock synchronization
  - IP multicasting
  - Berkeley Packet Filter

NETWORK OPERATION CENTER:

DARTnoc at ISI

RESEARCH-ORIENTED, NOT SERVICE-ORIENTED

- Hardware maintenance (8x5)
- Remote booting, power-cycling
  - *-> Machines in phone company POPs: dial in.*
  - *-> Machines in research sites: "buddy" hosts.*
- Maintain "baseline" router system

DARTnoc

- Information directories
  - *Master system configuration*
  - *Master kernel build files*
  - *Documentation, maps, etc.*
- Scheduling and Coordination
  - *Master schedule:* Signups for exclusive access to specified routers, in 3-hour blocks.
  - *Experiment coordination*

EXPERIMENTAL PROCEDURES

1. Build experimental kernel on home machine.
2. Boot experimental kernel into all relevant routers and perhaps test hosts
3. Make measurements.
4. Transfer data from routers to home machine for analysis.

TOOLS:

- Accurate **Time-Keeping** (Dave Mills and NTP).
- Traffic generation tool **TG** (SRI).
- Measurement and display tools (LBL and MIT).

DARTnet RESEARCH AREAS

(1) Resource Reservation (*)

(2) Wide-area Multicasting (*)

(3) Collaboration Technology (*)

(4) New Routing Paradigms

(5) Network Dynamics

---

* "DARTnet conferencing technology"

RESOURCE RESERVATION

The pieces of the problem are:

(1) **Resource / Service Model**
- Service parameters: "Flow Spec"

(2) **Traffic Control (Scheduling) Mechanism**
- Packet queueing and forwarding

(3) **Resource Setup Protocol**
- Set up real-time sessions
- Recover from failures

RESOURCE RESERVATION IN IP ROUTER

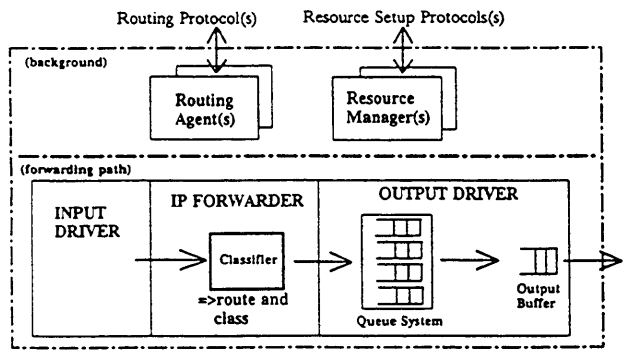

RESOURCE SETUP PROTOCOLS

ISSUES:

- Building multicast delivery paths *(complex problem)*
- CO ('stateful') or CL ('stateless') approach?
- Routing vs. resource allocation
- Failure recovery
- Internet problem: partial coverage and diversity

COLLABORATION TECHNOLOGY:

DARTnet Packet Audio & Video Programs

PACKET AUDIO:

- VT — ISI Voice Terminal (ST-II or UDP)
- VAT — LBL Voice program (UDP)
- MIT 386-based voice program (UDP)
- NEVOT — UMass Voice program (UDP)

PACKET VIDEO:

- PVP — ISI Packet Video Processor (ST-II or UDP)
- MIT 386-based video program (UDP)

EXPERIMENTAL PROGRAM: July 1992

EXPERIMENTS COMPLETED:

- IP multicasting (PARC)
  - Wide-area multicasting
  - Extension through IP tunnels
  - Widearea multicast Mazewar
  - Software DES encryption for multicast voice
- ST-II (BBN, ISI)
  - Video and audio on SparcStations (BBN, ISI)
  - ST-II traffic control using Virtual Clock
- InterDomain Policy Routing experiments (BBN)

EXPERIMENTAL PROGRAM: July 1992

ONGOING EXPERIMENTS ...

- Baseline DARTnet performance measurements (SRI)
- Accurate time for 1-way delay measurements (UDel)
- Traffic Control algorithms:
  - WFQ/FIFO+ scheduling (MIT, PARC)
  - Hierarchical allocation (LBL)
  - Stochastic Fair Queueing (SRI)
  - Deadline scheduling (UMass)

EXPERIMENTAL PROGRAM: July 1992

MORE ONGOING EXPERIMENTS...

- Resource Setup Protocols
  - ST-II: 'stateful' setup (BBN)
  - RSVP: 'stateless' multicast setup (PARC, USC)
- Multicast Routing (PARC)

EXPERIMENTAL PROGRAM

FUTURE EXPERIMENTS:

- Congestion Control Experiments (ISI, Mitre, LBL, PARC)
- Bellcore Touring Machine for conference control (Bellcore)
- Loss Preference experiments (UMass, MIT)
- Multicast congestion control (LBL, PARC)
- Synchronization protocol (BBN)
- Verify TCP simulations and extensions (LBL)

First DARTnet Research Demo: April 13, 1992

- Demonstrated three different **traffic control mechanisms** to support internet integrated service — a mixture of real-time and best-effort traffic.
- Used **wide-area IP multicasting** and **collaboration technology:** teleconferencing with packet video and voice.

Demo'd Three Traffic Control Mechanisms [Three different kernels]

- Virtual Clock under ST-II [BBN, ISI]
  ST-II is a connection-oriented setup and forwarding protocol.
  Virtual Clock was invented by Lixia Zhang (PARC).
- Predicted and Guaranteed Service [MIT, PARC]
  Dave Clark, Scott Shenker, and Lixia Zhang.
- Hierarchical Resource Allocation [LBL]
  Van Jacobson and Sally Floyd

FUTURE RESEARCH DEMOS

- Resource setup protocols

  Both stateful and stateless, plus admission control.

- Controlled link-sharing
- Loss Preference
- Advanced collaboration technology

DARTnet FUTURE

General goals:

A. Develop and use 'DARTnet conferencing technology'

B. Prototype 'Internet integrated service'

Traffic control algorithm and resource setup protocol.

C. Develop personal teleconferencing.

D. Test and verify new IP version.

E. Extend fundamental research into new domains.

F. Continue collaboration among network research groups.

4.5 Trusted NFS: Protocol Extensions for Multi-Level Security

Presented by Fred Glover/DEC

Bio: *Fred Glover is the Chair of the Trusted Network File System Working Group. He is employed as a consulting engineer within the UNIX Engineering Group at Digital Equipment Corporation.*

The Trusted Network File System (TNFS) Working Group has developed a specification describing extensions to the NFS V2 protocol which supports network file access between Multilevel Security (MLS) systems. This presentation describes the general approach used to transport additional security attributes between a Trusted NSF client and server, the current status of the TNFS Working Group effort, and the set of documents which support the development of TNFS implementations.

TRUSTED NFS:
Protocol Extensions for MultiLevel Security

Fred Glover
UNIX* Engineering Group
Digital Equipment Corporation

*UNIX is a registered trademark of UNIX Systems Laboratories (U.S.L.)

Trusted NFS: Working Group

History:

- TSIG TNFS working group formed in July 1989
- IETF TNFS working group formed in July 1991
- Meeting bimonthly since July 1989, 2 days/meeting

Initial Objectives:

- review requirements identified in the TCSEC, CMW, POSIX 1003.6 documents; others identified by members of the security community
- identify set of NFS V2 protocol extensions which support trusted, MultiLevel Security (MLS) environments
  - —discretionary access control (DAC)
  - —subject and object labeling
  - —mandatory access control (MAC)
  - —auditing
- extend NFS V2 for security attributes only

Trusted NFS: Working Group

General Approach:

- transport additional *user context* in the authentication parameter of each TNFS RPC request
- return additional *security extended file attributes* in each TNFS RPC response
- accommodate multiple policies and attribute formats through use of translated attributes: tokens
- document the protocol extensions, token mapping scheme, implementation details, interoperability test plan for both TSIG and IETF archives

Trusted NFS: Protocol Extensions

USER CONTEXT EXTENSION:

- AUTH_MLS credential:
  - —NFS V2 AUTH_UNIX information,
  - —audit ID, privileges, sensitivity, information, integrity

FILE ATTRIBUTE EXTENSIONS:

- modified *fattr*, *sattr* structures:
  - —ACLs, privileges, sensitivity, information, integrity

FILE NAME EXTENSIONS:

- sensitivity, information labeled file names

Trusted NFS: Protocol Extensions

MULTILEVEL DIRECTORY EXTENSIONS:

- file name attributes, diversion directories

NFS V2 PROCEDURE EXTENSIONS:

- NFSPROC_ACCESS=18: file open enhancement
- NFSPROC_SETLABEL=19: file name security attributes
- NFSPROC_MLD=20: diversion directories

EXPLICIT POLICY EXTENSIONS:

- Access Control
- Auditing

Trusted NFS: Working Group Status

- Seven TNFS documents completed and archived
- TNFS Audit Guide being drafted
- TNFS Protocol Specification is complete
- 2-3 implementations "close" to final specification
- planning for interoperability testing by EOY '92
- working group continuing to meet every 2 months

Trusted NFS: Documentation, Mail

Documents:

- TNFS Protocol Specification INTERNET-DRAFT
- TNFS Implementation Guide
- TNFS TKM Specification
- TNFS Administration Guide
- TNFS Interoperability Test Plan
- TNFS Test Attributes
- TNFS tnfs.h
- TNFS Meeting Minutes

Mail Lists:

- General Discussion: tnfs@wdl1.wdl.loral.com
- To Subscribe: tnfs-request@wdl1.wdl.loral.com
- Archives: archive-server@wdl1.wdl.loral.com

4.6 IP Address Encapsulation

Presented by Bob Hinden/Sun and Dave Crocker/TBO

Bio: *Bob Hinden is the Manager of Internet Engineering at Sun Microsystems. He has been involved in the Internet community since 1980 and has been the IESG Routing Area Director since 1989. He is currently involved in work in internet routing and addressing, and the issues relating to internetworking using Asynchronous Transfer Mode (ATM).*

Bio: *Dave Crocker is a co-Chair of the IP Address Encapsulation (IPAE) BOF (soon working group). He is a princap with The Branch Office, an open networking systems consultancy. Dave is also the IETF Area Director for Standards Management.*

IP ADDRESS ENCAPSULATION

(IPAE)

Robert Hinden
Dave Crocker

July 16, 1992

INTRODUCTION

- Motivation for IP Address Encapsulation
- How It Works
- Transition Plan
- Implementation Status
- Next Steps (or why we need your help)
- Benefits and Differences

GROWTH OF THE INTERNET

- Medium Term Solution
  - Solve Routing Table and Computation Limits
  - Support Larger IP Addresses
  - Minimum Changes
  - Keep Internet Growing within Current Architecture
  - Keep Details of Internet Intact

IPAE FOCUS

- Solve the Internet Routing and Addressing Problems
  - Routing Explosion
  - IP Address Exhaustion
- Change as Few Things as Possible
  - Protocol Modules
  - Router and Host Devices
  - People and Operations
- Strive for the Best Cost / Benefit Ratio

IP ADDRESSING COMMONWEALTH

- Extension of Current Internet Model of Subnets and Networks
  - No Changes Required for Subnet and Network Routing/Addressing Mechanisms
- Adds New Level Called *IP Addressing Commonwealth*
  - Commonwealths are connected by Commonwealth Routers
- Definition: *Area where 32-Bit IP Addresses are Unique*
  - Example: Today's Internet is a Commonwealth

HOW IT WORKS

- Add New Extension Header after IP Header
  - Identified by PROT field = IPAE
- New Header Contains Global IP Source and Destination Addresses
  - Format and Size Not Yet Defined except for inclusion of 32-bit IP address
  - Global IP Addresses are Globally Unique
- Protocol Field in Extended Header indicates Next Protocol

GLOBAL IP ADDRESSES

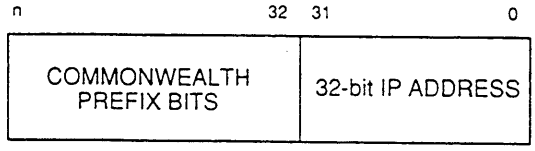

- Global IP Addresses are Globally Unique
- Prefix Bits identify Commonwealth
- Global IP Addresses Uniquely Identify End Points of Connections
  - Same as 32-Bit IP Addresses do Today
- 32-Bit IP Address used for Compatibility with Current Protocols

PACKET FORMAT

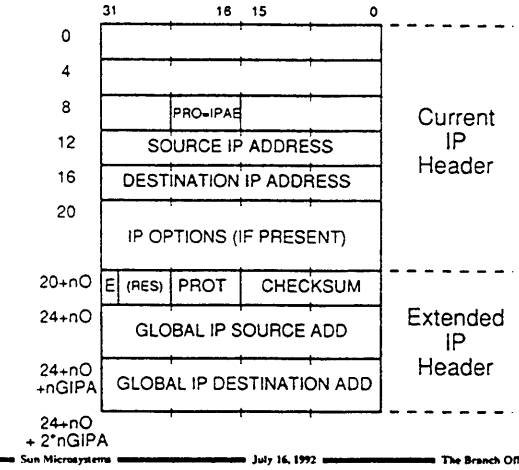

ROUTING MODEL

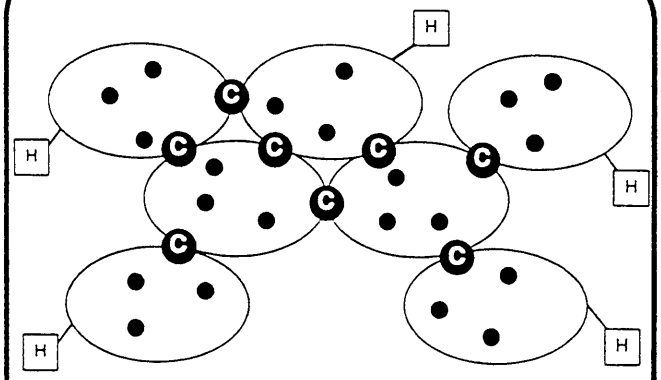

PACKET FORWARDING

- Routing and Packet Forwarding inside of Commonwealth Unchanged
- Hosts Perform Additional Lookup for Traffic Exiting Commonwealth
  - Add Extended IP Header
  - Forward Packet to Commonwealth Router
- Current IP Header Source and Designation Addresses used to Address Next Commonwealth Hop
  - Just Like IP Uses Network MAC Layer Addresses
- Commonwealth Router Forward Datagrams based on Global IP Addresses

MOBILE HOSTS USING IPAE

- Compatible with Current IETF Work on Mobile IP Hosts
- Mobility Supported inside of a Commonwealth without any Extra Headers
  - IP Destination Address in IP Header Serves as Current Location of Mobile Hosts
  - Global IP Address serves as Identifier of Host.

TRANSITION PLAN OVERVIEW

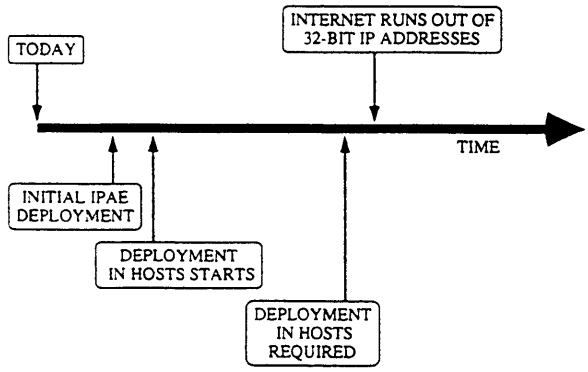

IMPLEMENTATION ANALYSIS

- Find out Effort to Implement IPAE on BSD Networking Release 2
- New Code Size
  - 400 Lines New Code
  - 100 Diffs to Existing Code
- Performance Impact Looks Small
- General Conclusion
  - Changes were Easy
  - Completed in Less than One Week

Sun Microsystems — July 16, 1992 — The Branch Office

SUMMARY OF CHANGES

- Changes
  - *in_addr* structure to contain Global IP Address
  - New Layer above IP for IPAE
  - Global Addresses in Routing Code and Expand Size of Entries in IP Routing Table
- Unchanged
  - IP Addresses in IP
  - TCP/UDP Pseudo-Headers
- Multihomed Hosts (different Commonwealths) require interface to stamp incoming traffic with Commonwealth

Sun Microsystems — July 16, 1992 — The Branch Office

NEXT STEPS

- Resolve Remaining Technical Issues
- Finish Documents
  - Protocol Description
  - Transition Plan
  - Addressing and Routing Plan
- Get IP Protocol identifier assigned for IPAE
- Develop and Test Implementations
  - 4.X BSD (Host and Commonwealth Router)
  - GATED (Commonwealth Router)
  - Others Implementations
- Write IESG Evaluation Criteria Report

Sun Microsystems — July 16, 1992 — The Branch Office

ISSUES TO RESOLVE

- Is E-Bit Necessary?
- Are Extended Options Needed?
- Reduction in IP Max Datagram Size Serious Problem?
- ICMP Mechanisms?
- Commonwealth Router Discovery Mechanism?
- Others?

Sun Microsystems — July 16, 1992 — The Branch Office

COMPONENT CHANGES

- Application Level
  - Support Larger Addresses
- Transport Level
  - Support Larger Addresses
- Address Resolution and Framing
  - No Changes
- Internet Layer
  - Retain Current IP Header
  - Encapsulate Extended Header
  - Add Commonwealth Routing Decision
  - Configure Default Commonwealth Route
  - Minor ICMP Changes
- Routing
  - Commonwealth Routing
- DNS
  - Add Larger Addresses

Sun Microsystems — July 16, 1992 — The Branch Office

BENEFITS

- ROAD Problems
  - Routing Table Explosion Problem Solved in Short Term
  - IP Address Exhaustion Problem Solved in Medium Term
- Protocols
  - Current IP Layer Protocols Retained
  - Multicast Retained
  - IP-over-<MEDIA> Retained

Sun Microsystems — July 16, 1992 — The Branch Office

BENEFITS (CONTINUED)

- Devices
  - Most Interior and Exterior Routers Not Required to Change
  - Hosts do not change current 32-Bit IP Addresses
  - Hosts Not Required to Implement IPAE Until Second Transition Step
- People and Operations
  - Existing Formats and Terminology Retained
  - Operational Tools Continue to Function
  - Current Investment in Training, Procedures, Documentation Retained

HOW THIS IS DIFFERENT FROM TUBA

- ROAD Problems
  - Deals with Routing Table Size and Routing Computation Problems when Initially Deployed
- Protocols
  - Retains Current IP Layer Infrastructure
    - Routing, Address Resolution, Network Management, Multicast, Training, Documentation, Operations Tools, ...
  - Uses Existing IP Checksum Algorithm
  - No Changes to Transport Protocols Required
    - Pseudo-Header Checksum uses Embedded 32-Bit IP Addresses
  - Contains PROT Identifier in Header (Like IPv4)

HOW THIS IS DIFFERENT FROM TUBA (CONTINUED)

- Devices
  - No Hosts Changes Required at Initial Step
- People and Operations
  - Graceful Transition
  - Supports Old Hosts Communicating with New Hosts until 32-Bit IP Addresses Run Out
  - Extends Current IP Technology Base
  - No Issues of Protocol Ownership
- All of the Benefits of TUBA but at a Much Lower Cost

IPAE WORKING GROUP INFORMATION

- Mailing Lists

  Mailing List: ip-encaps@sunroof.eng.sun.com

  Requests: ip-encaps-request@sunroof.eng.sun.com
- Anonymous FTP Site Draft Documents and Mail Archive

  /pub/ip-encaps/ parcftp.xerox.com
- Meeting Schedule (Video Conferences in SF Bay Area and East Coast)

  July 30

  August 27

  September 24

  October 15

  November IETF

4.7 A Cloudy Crystal Ball - Visions of the Future

Presented by Dave Clark/MIT

Bio: *David Clark is a Senior Research Scientist at the MIT Laboratory for Computer Science. He has been involved in the Internet community since 1976, and was Chair of the IAB from 1981 to 1989. He is involved in research on high-speed networks, support of real-time services, and networking for the information age.*

A Cloudy Crystal Ball

—

Visions of the Future

David D. Clark

M.I.T. Laboratory for Computer Science

IETF, July 1992

Alternate title: Apocalypse Now

Guessing the future

Identify the major external forces.

Consider each separately: the future if each dominates.

Speculate on what happens when we mix the stories together.

Forces that shape us

New services:
- Real time (video)
- Information access

Commercial network offerings:
- SMDS -> B-ISDN -> ubiquitous ATM access
- A new "kid" on the block?

Cyber-terrorists:
- "Security" gateways (Mail relays...)

Us:
- We have met the enemy and he is ...

Video and Real-Time

Our best success was not computing, but hooking people together.

Video and related services might be even more powerful.
- Do not use the phone analogy to speculate.

Small technical problems:
- Figure out how to do it. (MIT research here)
- Change all the routers.
- Charge for service.
- Make it affordable.

Work-stations are "almost there".

Computer mediated video interaction, a.k.a. games.

The network as an Information Mesh

An old goal, not yet achieved.

Recently, some neat stuff.
- WAIS, W3, Archie, Gopher, Prospero, etc.
- IETF, IRTF activities.

Does it require changes in the infrastructure?
- Scale, not speed, is the issue.
- Infrastructure must know about information objects. (MIT research here)
  - Names
  - Types

New services: charging for services, security (MIT research)

Video as an information interface.

Commercial network services

What are the issues?

Policy:
- How do we charge?
- Is there a role for monopoly?
- Business vs. ubiquitous access?

Technical:
- Control of routing.
- Support for accounting.
- Security.
- ...

ATM – A really big elephant

Myths from New Jersey:
- "They" will supply the scalable address space.
- "They" will solve the routing problem.
- ATM will solve the problem of real-time and QOS.
- "They" will be here real soon.

What are the real issues here?
- The network designers with telephony background do not understand multi-application networks.
- The phone companies have no history or approach to rapid deployment.
- They do not know how to do QOS either.

An example: why ATM LANs.
- My personal research: Everyone -> Sun-> standard.
- WHEN will the standard come? Mismatch possible.

The 90's – the decade of the cyber-terrorist

What I hated about the Morris worm:
- I found out about it on the Today show.

A worked example of a painless act of terrorism.
- The hacks of today are the commonplace of tomorrow. (True for good stuff, why not bad stuff?)

A digression: my Internet security talk.

SECURITY

Security is a CRITICAL problem.

Lack of security means the END OF LIFE AS WE KNOW IT!!

A time for ACTION!!!

(Can I be more explicit?)

WHAT'S THE PROBLEM?

Large networks and poor security don't mix.

Users will less and less tolerate the risk of being attacked from anywhere in the universe.

Look at the Internet worm.
- Check out the level of publicity.
- Consider the potential for damage.
- Consider who else has noticed the above.

Will this be the decade of the cyber-terrorist?

WHAT WILL HAPPEN?

Without better levels of protection, people will not be willing to attach to the Internet.

The "GREAT UNPLUGGING"?
- Too dramatic...

The decade of firewalls?
- Already happening.

MAIL RELAYS (Yuck!).

WHY ARE APPLICATION RELAYS SO BAD?

Application level relays have two problems:

- The signal the end of flexible service introduction.
- They don't work very well (consider mail today).

The end of the open road....
The fencing of the West....
The Italian telephone system....

WHY DO APPLICATION LEVEL RELAYS HELP?

Why do they help?
- Most security bugs are not in the specification, but in an Implementation of the specification.
- To penetrate a system protected behind an application level relay, it is necessary to break two implementations.
- Lower level attacks (tunneling attacks) cannot get past the relay.
- Insecure services can be blocked.

WHAT CAN WE DO?

Option 1: Make system security better.
- Not "our" problem".
- We must band together and make demands.
- Fix insecure services.

Option 2: Accept the inevitable; make it work.

Why doesn't it work well?

RELAYS ARE NOT CONSISTENT WITH THE BASIC ASPECTS OF THE PROTOCOL ARCHITECTURE!!

THE ARCHITECTURE AND THE RELAY

The protocol architecture assumes universal connectivity at the network layer.

Relays break that assumption. Things stop working.

Some examples:
- Names, addresses, routes.
- Fault isolation.

Recreation in *ad hoc* manner of the whole network functionality at application level.
- Consider X.500 and X.400.

WHAT "WE" SHOULD DO

Lobby for better system security.

Fix insecure services.
- PASSWORDS!!!

Push for "open domains".
- Better security = larger domains.

Develop a new protocol reference model for application level relay networks. Make it work. Accept it.

Don't just sit there and think it does not matter.
- Security is the problem we love to ignore.

Some lessons

Bad things do not happen all at once.
- AIDS, crime, routing collapse

Things get worse slowly. People adjust.

The problem is assigning the correct degree of fear to distant elephants.
- When should we (have) declared panic about:
  - Addressing, security,

Always ask: What will happen if I do nothing?
- Use these rules.

No security -> mail gateways.
No addressing -> ?? MAIL GATEWAYS and X.400.

Walking among the wild elephants

Plan of today:
- Fix addressing and routing.
- Leave security at end point. Pray.
- See if new services stamp out mail gateways.

An alternative plan (just for fun!!!)
- Build application-independent boarder crossing boxes.
- Ignore the addressing problem.
- Build a new network based on application, not IP connectivity. Routing and addressing at this level.

VIEWS OF THE FUTURE

The last force on us – us

The standards elephant of yesterday – OSI.

The standards elephant of today – its right here.

As the Internet and its community grows, how do we manage the process of change and growth?
- Open process – let all voices be heard.
- Closed process – make progress.
- Quick process – keep up with reality.
- Slow process – leave time to think
- Market driven process – the future is commercial.
- Scaling driven process – the future is the Internet.

We reject: kings, presidents and voting.
We believe in: rough consensus and running code.

SLIDE 19

VIEWS OF THE FUTURE

A look at us

What are we good at?
- Responding to short term reality.
- Building stuff that works.
- Calling bad stuff bad.

What are we bad at?
- Growing our processes to match our size.
- Setting long-term direction.

SLIDE 20

VIEWS OF THE FUTURE

An example – making standards.

What is the correct model?
- I am trying to ask this in a constructive way, please.

Today: IESG proposes, with IAB advice and consent.
- Sort of like the House of Lords.

IESG alone is enough?
- I think some "checks and balances" are good.

Supreme court model?
- Life appointments!!! No...
- Arbitration? TANNSAAFL judging?

What is the community (meta-)process that will create the acceptable process?

SLIDE 21

VIEWS OF THE FUTURE

An example – long term planning

Consider the addressing/routing situation.

Consider (just for fun) my security elephant.

How could we as a group decide what to do about security?
- Can we converge on an assessment of the peril?
- Can we rank this with other perils?
- Can we direct the funds to do research?
- Can we hold a steady course in the storm?

I offer these questions for your deliberation?
- Think positive thoughts.
- Remember: If we have a problem it is due to too much success.

SLIDE 22

Chapter 5

Trusted Sytems Interoperability Group

The Trusted Systems Interoperability Group (TSIG) is a consortium of trusted systems developers, system integrators and government users. The TSIG began in 1989 as an integration forum for Compartmented Mode Workstation (CMW) vendors. The TSIG's goal is to foster the interoperability of multilevel secure systems.

Technical activity on any specific topic within the TSIG is addressed within working groups. There are currently five chartered TSIG working groups: Commercial Internet Protocol Security Option (CIPSO), Trusted Network File System (TNFS), Trusted Sessions (TSESS), Trusted Administration (TADMIN) and Trusted X-Windows (TXWIN). Two of the working groups, CIPSO and TNFS, are co-chartered with the IETF. The working groups conduct business during 2.5 day meetings held about 4 to 5 times per year. Meeting reports, charters and general information on current TSIG activities are available on-line by an archive server (archive-server@wdl1.wdl.loral.com).

Several TSIG mailing lists exist. The master list is tsig@wdl1.wdl.loral.com. Each of the working groups (e.g., CIPSO, TNFS, TSESS) also have a related mailing list at wdl1.wdl.loral.com. To join a mailing list, send a request to the associated request list. All TSIG mailing lists have a companion "-request" list. Send requests to join a list to <listname>-requests@wdl1.wdl.loral.com.

CURRENT MEETING REPORT

Minutes of the Trusted Administration Working Group (TADMIN)

Summary not provided. A detailed listing of issues covered can be found in the remote directories under tadmin-minutes-92jul.txt.

Attendees

David Belote	
Lee Benzinger	lab@wdl1.wdl.loral.com
Charles Blauner	chazx@ctt.bellcore.com
Luc Boulianne	lucb@cs.mcgill.ca
Jeffrey Case	case@cs.utk.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Jeffrey Edelheit	edelheit@mitre.org
William Huyck	
Nina Lewis	nina@cam.unisys.com
Vern McGeorge	vern@hpda.cup.hp.com
Wally Ramsey	wbr@mitre.org
Bradley Rhoades	bdrhoades@mmc.mmmg.com
Paul Vazquez	vazquez@dockmaster.ncsc.mil
Moira West	mjw@cert.org
Peter Williams	p.williams@uk.ac.ucl.cs

CURRENT MEETING REPORT

Reported by Doug Barlow/DEC

Minutes of the Trusted Sessions Working Group (TSESS)

The Group had drafted Julie LaMoine (MITRE) to Chair the Group at this meeting, but she couldn't make it, so Doug Barlow filled in for her.

Mike Matthews (Addamax) presented an overview of the Addamax token mapping service. Addamax is also planning on writing an Addamax ATN Profile to compliment the TSWG framework.

We reviewed the status of outstanding homework. Available for progress were the following:

- The Framework Document
- The Commercial Multi-level Distributed Security (CMDS) Profile
- The MaxSix V2.0 Profile
- The DNSIX V4.0 Profile

The Framework document was approved for submission to the TSIG plenary, with the following edits:

- The order of sections 4.1 and 4.2 are to be reversed, to conform to the order presented in the diagram.

- In section 5, the phrase "Addamax plans to present this" is to be changed to read, "Addamax has presented this".

The CMDS Profile was approved for submission to the TSIG plenary, with the following edits:

- Section 2.4, last bullet, the phrase, "supplying he local" is to be changed to read, "supplying the local".

- Section 4, the incorrect ASN.1 syntax in the first line of the Commercial Label Exchange protocol is to be corrected to read, "COMMERCIAL-LABEL DEFINITIONS ::=".

The MaxSix Profile was approved for submission to th TSIG plenary this decision was rescinded later – keep reading), with the following edits:

- The version number specified in the title is to be changed from "3.0" to "2.0".

- In Section 2.3, paragraph 4, the phrase "the MaxSix Security" is to be changed to read, "the MaxSix proposal for DNSIX Security".

- A paragraph will be added explaining how to obtain the referenced MaxSix documents, since they are not in the TSIG archives.

The Trusted Realm Environment Exchange Service (TREES) document was approved for submission to the TSIG plenary without modification.

The DNSIX V4.0 Profile was provided as status information, but is not yet complete, and was not considered for progression at this time.

John Batzer (ITT) told us about work he is doing on a project named "Dragonfly". It is a hardware-assisted session layer security protocol which uses RSA to validate packets. As there were several newcomers, we also provided an overview of the work the Trusted Sessions Working Group has done.

We examined possible future paths for the TSWG. Suggested alternatives were:

- Help other TSIG working groups utilize trusted sessions.
- Move existing applications (telnet, ftp, rcmd, etc.) to trusted sessions.
- Work on the token mapping problem.
- Agree on a common API for operating trusted sessions.
- Work on a TSIG Security Architecture Framework.
- Provide consistent management of trusted sessions (a la MIBs).

Paul Vasquez (DIA) was invited to attend our Group and give us an update on DIA's plans for DNSIX V3.0. Paul called for any and all TSIG attendees to comment on the MaxSix proposal for DNSIX V3.0. Comments must be received by the end of July. So far, out of the 22 vendors to which DIA has made the MaxSix documents available, only IBM and Digital have returned comments. Two other proposals for DNSIX V3.0 have been received by DIA, the one from Addamax, and one from Digital. However, DIA does not plan on distributing those proposals. Paul recommended that people contact the submitters directly to obtain them. DIA would entertain comments on the other proposals as well.

Paul went on to describe what he felt were requirements that any proposal for DNSIX V3.0 must meet:

- IPSO (nee RIPSO) is required.
- An API specification is desirable, but not required.
- A token mapping capability is desirable, but not required.
- Backwards compatibility with DNSIX V2.1, which was originally stated to be a requirement, isn't really a requirement, since there are no installed DNSIX V2.1 sites in DIA to be backwards compatible with.

The Group felt that the current TSWG method of providing a profile for every possible DNSIX V3.0 submission did not meet the goal of standardizing on a single solution. Hence the previous decision to submit all completed documents was rescinded, and a vote to forward each individual document to the TSIG plenary was taken. The results of the

voting was:

- Framework for Trusted Session Protocol – Yes: 6, No: 0, Abstaining: 3.
- CMDS Profile – Yes: 6, No: 0, Abstaining: 3.
- MaxSix V2.0 Profile – Yes: 0, No: 5, Abstaining: 4.
- TREES Document – Yes: 5, No: 0, Abstaining: 4.

The "No" vote on the MaxSix V2.0 Profile is taken to be an indication that the Group wishes to wait and see the progress of the DNSIX V3.0 specification. The Group reserves the right to reconsider this document for submission to the TSIG plenary at a later time.

NOTE: In the closing TSIG plenary, TSIG voted to accept the submitted documents – Yes: 14, No: 1, Abstaining: 7. Concerns were expressed that the profile mechanism still does not guarantee interoperability between ALL secure systems, and that some newer people were not familiar with the TSWG work. Doug Barlow (Digital) volunteered to present an overview of the adopted papers at the next TSIG meeting in Minneapolis.

Attendees

Doug Barlow	barlow@decwet.dec.com
John Batzer	
Luc Boulianne	lucb@cs.mcgill.ca
Dean Jagels	dpj@sware.com
James Lin	yeejang@cup.hp.com
Clifford Neuman	bcn@isi.edu
Richard Newton	rnewton@csd.harris.com
Paul Sangster	sangster@ans.net
Paul Vazquez	vazquez@dockmaster.ncsc.mil
Charles Watt	watt@sware.com
W. Stan Wisseman	swissema@oracle.com

CURRENT MEETING REPORT

Reported by Mark Smith/AT&T

Minutes of the Trusted X Working Group (TRUSTEDX)

The current thinking on Secure X is divided into two separate approaches:

1. The establishment of a core security policy derived from prior vendor efforts.

   Vendors are still reluctant to publish their policies, although there are signs that some will soon be published. Inasmuch as there are no current proposals in this area, no further discussion on this approach was offered at the meeting.

2. The abstraction of the security policy via a policy-free protocol.

   The key here is the construction of a mechanism for security-cognizant applications to determine what the security policy is. We briefly discussed the "RequestPolicy" proposal (distributed via email shortly before the meeting), which allows a client to probe specific points of the policy, and agreed that the approach is promising but that a proof of concept is needed.

The Boston TSIG X Working Group was not well attended. For that reason little progress was made other than the discussion on "RequestPolicy" above. We need vendor support, especially in the form of new proposals for (1) above, although more work in area (2) is very welcome also.

We need to have an idea of the attendees for the next TSIG meeting in Minneapolis so that we can judge whether another "cooling off period" is required. Please let me or Mark Christianson know fairly soon whether you'll be attending the next meeting. I would like to know whether this low attendance was an aberration or not.

Attendees

Charles Blauner	chazx@ctt.bellcore.com
Edward Cande	cande@zk3.dec.com
Alton Hoover	hoover@ans.net
Richard Newton	rnewton@csd.harris.com
Mark Smith	mark@neptune.att.com

Appendix A

Attendees

George Abe
Infonet Services Corporation
(W)310-335-2867
(WF)310-335-2876
abe@infonet.com

Vikas Aggarwal
JvNCnet
(W)609-258-2403
(WF)609-258-2424
aggarwal@nisc.jvnc.net

Masuma Ahmed
Bell Communications Research
(W)908-758-2515
(WF)908-758-4192
mxa@sabre.bellcore.com

Ed Albrigo
Corporation for Open Systems
International
(W)703-883-2723
(WF)703-848-4572
ealbrigo@cos.com

Steve Alexander
INTERACTIVE Systems
Corporation
(W)708-505-9100
(WF)708-505-9133
stevea@i88.isc.com

J. Allard
Microsoft Corporation
(W)206-936-9031
(WF)206-936-7329
jallard@microsoft.com

Lynn Alley
Dayna Communications, Inc.
(W)801-531-0600
dayna@utah.edu

Claudio Allocchio
National Institute for
Nuclear Physics
(W)+39 40 3758523
(WF)+39 40 226338
c.allocchio@elettra.trieste.it

Guy Almes
Advanced Network and
Services, Inc.
(W)914-789-5333
(WF)914-789-5310
almes@ans.net

Philip Almquist
Consultant
(W)415-752-2427
almquist@jessica.stanford.edu

Harald Alvestrand
SINTEF DELAB
(W)+47 7 59 70 94
(WF)+47 7 53 25 86
Harald.Alvestrand@delab.sintef.no

Michael Anello
XLNT Designs, Inc.
(W)619-487-9320
(WF)619-487-9768
mike@xlnt.com

David Arneson
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
arneson@ctron.com

Nagaraj Arunkumar
3Com Corporation
(W)408-764-5104
(WF)408-764-5003
nak@3com.com

Derek Atkins
Bell Communications Research
(W)201-829-4233
warlord@thumper.bellcore.com

Randall Atkinson
Naval Research Laboratory
(W)202-404-8590
(WF)202-404-7942
atkinson@itd.nrl.navy.mil

Karen Auerbach
Epilogue Technology
Corporation
(W)805-650-7107
(WF)805-650-7108
karen@epilogue.com

Karl Auerbach
Empirical Tools and
Technologies
(W)408-427-5280
(WF)408-427-5281
karl@empirical.com

Robert Austein
Epilogue Technology
Corporation
(W)617-942-0915
(WF)617-942-7846
sra@epilogue.com

Sam Ayers
Sybus Corporation
(W)813-535-6999
(WF)813-536-4773
s.ayers@sybus.com

Ashar Aziz
Sun Microsystems, Inc.
(W)415-336-2735
(WF)415-336-3156
ashar.aziz@eng.sun.com

William Babson
Penril Datacomm
Networks, Inc.
(W)301-921-8600
(WF)301-921-0161
bill@penril.com

Suhas Badve
Hewlett-Packard
(W)408-447-3434
(WF)408-447-3660
badve@cup.hp.com

Zareh Baghdasarian
Tekelec, Inc.
(W)818-880-7817
(WF)818-880-6993
zareh@tekelec.com

Douglas Bagnall
Proteon, Inc.
(W)508-898-2800
deb@proteon.com

Cynthia Bagwell
The MITRE Corporation
(W)703-883-5438
(WF)703-883-7142
cbagwell@gateway.mitre.org

Dennis Baker
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
dbaker@wellfleet.com

Fred Baker
Advanced Computer
Communications
(W)805-685-4455
(WF)805-685-4465
fbaker@acc.com

Tony Ballardie
University College London
(W)+44 71 387 7050
a.ballardie@cs.ucl.ac.uk

John Ballard
Microsoft Corporation
(W)206-936-4688
(WF)206-883-8101
jballard@microsoft.com

Atul Bansal
Digital Equipment Corporation
(W)508-486-5355
(WF)508-486-7417
bansal@wile.enet.dec.com

Doug Barlow
Digital Equipment Corporation
(W)206-562-8552
(WF)206-865-8890
barlow@decwet.dec.com

Jim Barnes
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
barnes@xylogics.com

William Barns
The MITRE Corporation
(W)703-883-6832
(WF)703-883-7142
barns@gateway.mitre.org

Richard Basch
Massachusetts Institute of
Technology
(W)617-253-0100
(WF)617-258-8736
basch@mit.edu

Tony Bates
University of London
Computer Center
(W)+44 71 405 8400
(WF)+44 71 242 1845
tony@ean-relay.ac.uk

John Batzer
ITT A/CD
(W)201-284-3895

Mark Baushke
cisco Systems
(W)415-688-7849
mdb@cisco.com

Richard Baxter
Pensar Corporation
(W)414-739-4355
(WF)414-739-9615

Charles Bazaar
Emulex Corporation
(W)714-668-5416
(WF)714-668-6818
bazaar@emulex.com

Jordan Becker
Advanced Network and
Services, Inc.
(W)914-789-5308
(WF)914-789-5310
becker@nis.ans.net

James Beers
Cornell University
(W)607-255-2627
(WF)607-255-5771
beers@nr-tech.cit.cornell.edu

David Belote
McDonnell Douglas
(W)714-566-4153
(WF)714-566-4100

Ken Benstead
Coral Network Corporation
(W)508-460-6010
(WF)508-481-6258
kbenstead@coral.com

Lee Benzinger
Loral WDL
(W)408-473-4668
lab@wdl1.wdl.loral.com

Lou Berger
Penril Datacomm
Networks, Inc.
(W)301-921-8600
(WF)301-921-0161
lberger@penril.com

Tim Berners-Lee
CERN
(W)+41 22 767 3755
(WF)+41 22 767 7155
timbl@info.cern.ch

James Berry
Sandia National Laboratories
(W)510-294-2914
(WF)510-294-1225
beri@sandia.llnl.gov

Joel Berson
Digital Equipment Corporation
(W)603-884-6215
berson@brake.enet.dec.com

Pravin Bhagwat
IBM Corporation
(W)914-784-6866
w2pravin@watson.ibm.com

Andy Bierman
David Systems, Inc.
(W)408-720-6854
(WF)408-720-1337
bierman@davidsys.com

Jonathan Biggar
NetLabs, Inc.
(W)415-961-9500
(WF)415-961-9300
jon@netlabs.com

Charles Blauner
Bell Communications Research
(W)908-699-8395
(WF)908-336-2943
chazx@ctt.bellcore.com

Uri Blumenthal
IBM Corporation
(W)914-945-1267
uri@watson.ibm.com

Mark Bokhan
Digital Equipment Corporation
(W)603-884-4479
(WF)603-884-2607
bokhan@abitok.enet.dec.com

David Bolen
Advanced Network and
Services, Inc.
(W)914-789-5327
(WF)914-789-5310
db3l@nis.ans.net

David Borman
Cray Research, Inc.
(W)612-683-5571
(WF)612-683-5599
dab@cray.com

Natan Borshansky
Ascom Timeplex
(W)201-391-6000
(WF)201-391-2279
natanb@turisc.timeplex.com

Steve Bostock
Novell, Inc.
(W)408-473-8203
(WF)408-435-1706
steveb@novell.com

Luc Boulianne
McGill University
(W)514-398-6698
(WF)514-398-3883
lucb@cs.mcgill.ca

Robert Braden
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
braden@isi.edu

Scott Bradner
Harvard University
(W)617-495-3864
(WF)617-495-0914
sob@harvard.edu

Laura Breeden
FARNET, Inc.
(W)617-890-5120
(WF)617-890-5117
breeden@farnet.org

Tom Brennan
Exide Electronics
(W)919-872-3030
(WF)919-878-2385
brennan@exide.com

George Brett
MCNC
(W)919-248-1886
(WF)919-248-1405
ghb@jazz.concert.net

David Bridgham
Epilogue Technology
Corporation
(W)617-942-0915
(WF)617-942-7846
dab@epilogue.com

Scott Brim
Cornell University
(W)607-255-9392
(WF)607-255-5771
swb@cornell.edu

C. Douglas Brown
Sandia National Laboratories
(W)505-845-8699
(WF)505-844-2067
cdbrown@sandia.gov

Caralyn Brown
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
cbrown@wellfleet.com

Geetha Brown
Digital Equipment Corporation
(W)603-881-2126
(WF)603-881-2257
geetha@decvax.dec.com

Jack Brown
United States Army
Computer Engineer Center
(W)602-533-2834
(WF)602-533-5676
jbrown@huachuca-emh8.army.mil

Rich Brown
Dartmouth College
(W)603-646-3648
(WF)603-646-2810
Richard.E.Brown@Dartmouth.edu

Gregory Bruell
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
gob@wellfleet.com

Theodore Brunner
Bell Communications Research
(W)201-829-4678
(WF)201-267-9298
tob@thumper.bellcore.com

Alan Bryenton
Bell Northern Research
(W)613-763-2791
(WF)613-763-2626
bryenton@bnr.ca

Christopher Bucci
Network General Corporation
(W)201-616-8307
(WF)201-616-8311
bucci@ngc.com

Steve Buchko
Newbridge Networks Corporation
(W)603-591-3600
(WF)603-591-3680
stevebu@newbridge.com

Philip Budne
Shiva Corporation
(W)617-252-6356
(WF)617-252-6852
phil@shiva.com

Jeffrey Burgan
NASA Ames Research Center
(W)415-604-5705
(WF)415-604-0063
jeff@nsipo.nasa.gov

John Burnett
Adaptive
(W)415-780-7768
(WF)415-364-0123
jlb@adaptive.com

Randy Butler
National Center for
Supercomputing Applications
(W)217-244-8285
(WF)217-244-7396
rbutler@ncsa.uiuc.edu

Ross Callon
Digital Equipment Corporation
(W)508-486-5009
(WF)508-486-5279
callon@bigfut.enet.dec.com

Edward Cande
Digital Equipment Corporation
(W)603-881-0460
cande@zk3.dec.com

Robert Candela
Ungermann-Bass
(W)617-851-0198
candela@sceng.ub.com

C. Allan Cargille
University of Wisconsin-Madison
(W)608-262-5084
(WF)608-262-9777
cargille@cs.wisc.edu

James Carlson
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
carlson@xylogics.com

David Carr
Gandalf Data Limited
(W)613-723-6500
(WF)613-226-1717
dcarr@gandalf.ca

Lida Carrier
Apple Computer, Inc.
(W)408-974-8325
(WF)408-974-0892
lida@apple.com

William Carter
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
carter@ctron.com

Jeffrey Case
University of Tennessee
(W)615-573-5931
(WF)615-573-9197
case@cs.utk.edu

Stephen Casner
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
casner@isi.edu

Greg Celmainis
Newbridge Networks Corporation
(W)613-591-3600
(WF)613-591-3680
gregc@newbridge.com

Vinton Cerf
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
vcerf@nri.reston.va.us

John Chang
US West Advanced Technologies
(W)303-541-6255
(WF)303-541-6250
jrc@uswest.com

Yee-Hsiang Chang
MCNC
(W)919-248-1169
(WF)919-248-1405
yhc@concert.net

A. Lyman Chapin
Bolt Beranek and Newman
(W)617-873-3133
(WF)617-873-4086
lyman@bbn.com

Mitchell Charity
Massachusetts Institute of Technology
mcharity@lcs.mit.edu

Dilip Chatwani
SynOptics Communications, Inc.
(W)408-764-1624
dilip@synoptics.com

Cheng Chen
Teleos Communications, Inc.
(W)908-544-6444
(WF)908-544-9890
ctc@huey.bellcore.com

Frank Chen
Cascade Communications Corporation
(W)508-692-2400
(WF)508-692-9214
frankc@casc.com

Dean Cheng
Retix
(W)310-828-3400
(WF)310-828-2255
dean@sun2.retix.com

Andrew Cherenson
Silicon Graphics, Inc.
(W)415-960-1980
arc@sgi.com

J. Noel Chiappa
(H)804-898-8183
(HF)804-898-7663
jnc@lcs.mit.edu

Robert Ching
Network Application Technology
(W)408-370-4270
(WF)408-370-4222
natadm!rching@uunet.uu.net

Bilal Chinoy
San Diego Supercomputer Center
(W)619-534-8305
(WF)619-534-5113
bac@sdsc.edu

Chris Chiotasso
Artel Communications Corporation
(W)508-562-2100
(WF)508-562-6942
chris@artel.com

Isaac Choi
SynOptics Communications, Inc.
(W)408-764-1546
(WF)408-988-5525
ichoi@synoptics.com

Cyrus Chow
NASA Ames Research Center
(W)415-604-6843
(WF)415-604-4003
cchow@orion.arc.nasa.gov

Mark Christenson
Cray Research, Inc.
(W)612-683-5208
(WF)612-683-5599
mgc@cray.com

Jodi-Ann Chu
University of Hawaii
(W)808-956-2400
(WF)808-956-2412
jodi@uhunix.uhcc.hawaii.edu

George Clapp
Ameritech Services, Inc.
(W)708-248-5462
(WF)708-248-3775
clapp@ameris.center.il.ameritech.com

Cynthia Clark
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
cclark@nri.reston.va.us

David Clark
Massachusetts Institute of
Technology
(W)617-253-6003
(WF)617-253-8682
ddc@lcs.mit.edu

Henry Clark
OARnet
(W)614-292-6483
(WF)614-292-7168
henryc@oar.net

Alan Clegg
MCNC
(W)919-248-1120
(WF)919-248-1405
abc@concert.net

John Clement
EDUCOM
(W)202-466-6296
(WF)202-872-4318
clement@educom.edu

Robert Clements
Bolt Beranek and Newman
(W)617-873-3612
(WF)617-873-3776
clements@bbn.com

Tim Cleveland
Digital Equipment Corporation
(W)508-486-5806
tac@lkg.dec.com

James Cobban
Bell Northern Research
(W)613-763-8013
(WF)613-763-8581
jcobban@bnr.ca

Marc Cochran
IBM Corporation
(W)203-783-4308
(WF)203-783-5046
mjc@vnet.ibm.com

Richard Colella
National Institute of
Standards and Technology
(W)301-975-3627
(WF)301-590-0932
colella@osi.ncsl.nist.gov

Robert Collet
US Sprint
(W)703-904-2230
(WF)703-904-2119
rcollet@icm1.icp.net

Michael Collins
Lawrence Livermore
National Laboratory
(W)510-422-4018
(WF)510-422-0435
collins@ccc.nersc.gov

Rob Coltun
Consultant
(W)301-340-9416
rcoltun@ni.umd.edu

Jim Comen
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
comenj@interlan.interlan.com

James Conklin
EDUCOM
(W)202-872-4200
(WF)202-872-4318
jbc@bitnic.educom.edu

Michael Conn
MCI Communications
(W)214-918-5710
(WF)215-918-6019
0043874510mcimail.com

Robert Cooney
Naval Computer and
Telecommunications Station
(W)202-433-2448
(WF)202-433-3188
cooney@wnyose.nctsw.navy.mil

Eric Cooper
Fore Systems, Inc.
(W)412-967-4040
(WF)412-967-4044
ecc@fore.com

Tom Coppeto
Massachusetts Institute of
Technology
(W)617-253-8400
(WF)617-258-8736
tom@mit.edu

Scott Cossette
Concord Communications, Inc.
(W)508-460-4646
(WF)508-481-9772
sdc@concord.com

Frank Coviello
ITT
(W)201-284-4874
(WF)201-284-2849

Curtis Cox
Naval Computer and
Telecommunications Station
(W)202-433-4026
(WF)202-433-3188
ccox@wnyose.nctsw.navy.mil

Tracy Cox
Bell Communications Research
(W)908-758-2107
(WF)908-758-4177
tacox@sabre.bellcore.com

Stephen Coya
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
scoya@nri.reston.va.us

Michael Craren
Proteon, Inc.
(W)508-429-7114
mjc@proteon.com

David Crocker
The Branch Office
(W)408-246-8253
(WF)408-249-6205
dcrocker@mordor.stanford.edu

Stephen Crocker
Trusted Information Systems
(W)301-854-6889
(WF)301-854-5363
crocker@tis.com

Dave Cullerot
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-335-3467
cullerot@ctron.com

Paul Cummings
Digital Equipment Corporation
(W)508-486-6152
(WF)508-486-6100

Cathy Cunningham
Microcom, Inc.
(W)617-551-1131
(WF)617-551-1015
cmc@microcom.com

John Curran
NEARnet
(W)617-873-4398
(WF)617-873-5620
jcurran@bbn.com

Megan Davies
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
mdavies@nri.reston.va.us

James Davin
Massachusetts Institute of Technology
(W)617-253-6020
(WF)617-258-8682
jrd@ptt.lcs.mit.edu

Andy Davis
Spider Systems, Ltd.
(W)+44 31 554 9424
andy@spider.co.uk

Michael Davison
SNMP Research, Inc.
(W)615-573-1434
(WF)615-573-9197
davison@cs.utk.edu

Osmund de Souza
AT&T Bell Laboratories
(W)908-949-1393
(WF)908-949-4673
osmund.desouza@att.com

Michael DeAddio
Bell Communications Research
(W)201-829-5024
(WF)201-829-2504
deaddio@thumper.bellcore.com

Steve Deering
Xerox Corporation
(W)415-812-4839
(WF)415-812-4471
deering@parc.xerox.com

Marty Del Vecchio
Shiva Corporation
(W)617-252-6300
(WF)617-252-6852
marty@shiva.com

Daniel Dern
Dern Associates
(W)617-926-8743
(WF)617-926-3048
ddern@world.std.com

Tony DeSimone
AT&T Bell Laboratories
(W)908-949-4521
(WF)908-834-5906
tds@hoserve.att.com

Richard desJardins
NASA
(W)301-286-8554
(WF)301-286-5269
desjardi@boa.gsfc.nasa.gov

Peter Deutsch
Bunyip Information Systems
(W)514-398-3709
(WF)514-398-6876
peterd@cc.mcgill.ca

Peter DiCamillo
Brown University
(W)401-863-7582
(WF)401-863-7329
Peter_DiCamillo@brown.edu

Jeffrey Dietz
Proteon, Inc.
(W)508-898-2800
(WF)508-898-3176
jbd@proteon.com

Chas DiFatta
Software Engineering
Institute
(W)412-268-7788
(WF)412-268-5758
chas@sei.cmu.edu

Michael DiGioia
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
digioia@interlan.interlan.com

Hans Dittler
Conware Computer Consulting
GmbH
(W)+49 721 94950
(WF)+49 721 9495130
dittler@conware.de

Kurt Dobbins
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
dobbins@ctron.com

Irina Dobrina
FTP Software, Inc.
(W)617-246-0900
irina@ftp.com

Roger Draper
Emerson Computer Power
(W)714-457-3638
(WF)714-457-3633
rdraper@cerf.net

Jack Drescher
MCNC
(W)919-248-1412
(WF)919-248-1405
drescher@concert.net

Daniel Duchamp
Columbia University
(W)212-854-8115
(WF)212-666-0140
duchamp@liberty.cs.columbia.edu

Pierre Dupont
Motorola
(W)604-241-6243
(WF)604-241-6042
dupont@mdd.comm.mot.com

Tom Easterday
CICNet, Inc.
(WF)313-998-6105
tom@cic.net

Donald Eastlake
Digital Equipment Corporation
(W)508-486-2358
(WF)508-486-2311
dee@ranger.enet.dec.com

Jeffrey Edelheit
The MITRE Corporation
(W)703-883-7586
(WF)703-883-1397
edelheit@mitre.org

Raymond Edgerton
Bell Atlantic
(W)215-963-6740
(WF)215-963-6776
bellmf!edgerton@uwm.edu

Bill Elliot
ONEAC Corporation
(W)708-816-6000
(WF)708-680-5124

Alan Emtage
Bunyip Information Systems
(W)514-398-8117
(WF)514-398-6876
bajan@cc.mcgill.ca

David Engel
Optical Data Systems, Inc.
(W)214-234-6400
(WF)214-234-4057
david@ods.com

Fred Engel
Concord Communications, Inc.
(W)508-460-4646
(WF)508-481-9772
engel@concord.com

Robert Enger
Advanced Network and
Services, Inc.
(W)914-789-5355
(WF)914-789-5310
enger@ans.net

Urs Eppenberger
SWITCH
(W)+41 1 261 81 12
(WF)+41 1 261 81 33
eppenberger@switch.ch

Hans Eriksson
Swedish Institute of
Computer Science
(W)+46 8 752 1527
(WF)+46 8 751 7230
hans@sics.se

Michael Erlinger
Lexcel
(W)714-870-2117
(WF)714-680-0279
mike@lexcel.com

Jeff Erwin
ProTools, Inc.
(W)503-645-5400
(WF)503-645-3577
Jeff_Erwin@protools.com

Julio Escobar
Bolt Beranek and Newman
(W)617-873-4579
(WF)617-873-3776
jescobar@bbn.com

Susan Estrada
CERFnet
(W)619-455-3902
(WF)619-455-3990
estradas@cerf.net

Fran Fadden
Digital Equipment Corporation
(W)603-881-2115
fran@decvax.dec.com

Roger Fajman
National Institutes of Health
(W)301-402-1246
(WF)301-480-6245
raf@cu.nih.gov

Steven Fancher
Engineering On Demand, Inc.
(W)617-884-1876
sfancher@ursa-major.spdcc.com

Bill Fardy
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
fardy@ctron.com

Dino Farinacci
cisco Systems
(W)415-688-4696
(WF)415-688-8282
dino@cisco.com

Tom Farinelli
Department of Defense
(W)410-859-4491
tcf@tyco.ncsc.mil

Stefan Fassbender
GMD
(W)+49 2241 14 3180
(WF)+49 2241 14 3002
stf@easi.net

Sallie Fellows
New Hampshire State
Department of Education
(W)603-271-2778
(WF)603-271-1953
sallie%ed@psc.plymouth.edu

David Fencl
ONEAC Corporation
(W)708-816-6000
(WF)708-680-5124

Dennis Ferguson
CA*net
(W)416-978-3328
(WF)416-978-6620
dennis@mrbill.canet.ca

Robert Finn
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
rfinn@wellfleet.com

Constance Fleenor Lloyd
Lloyd & Associates
(W)916-676-1147
(WF)916-676-3442
cfleenor@lloyd.com

Peter Ford
Los Alamos National Laboratory
(W)505-665-0058
(WF)505-665-7793
peter@lanl.gov

James Forster
cisco Systems
(W)415-688-7854
(WF)415-688-8245
forster@cisco.com

Margaret Forsythe
FTP Software, Inc.
(W)617-246-2920
(WF)617-245-7943
mrf@ftp.com

Jill Foster
University of
Newcastle-Upon-Tyne
(W)+44 91 222 8250
(WF)+44 91 222 8232
jill.foster@newcastle.ac.uk

Craig Fox
Network Systems Corporation
(W)612-424-1703
foxcj@network.com

Paul Fox
Cayman Systems, Inc.
(W)617-494-1999
(WF)617-494-9270
pgf@cayman.com

Osten Franberg
Ellemtel
(W)+46 8 7274121
(WF)+46 8 6478276
euaokf@eua.ericsson.se

Barbara Fraser
CERT
(W)412-268-7090
(WF)412-268-6989
byf@cert.org

Ron Frederick
Xerox Corporation
(W)415-812-4459
(WF)415-812-4471
frederick@parc.xerox.com

Ned Freed
Innosoft International, Inc.
(W)714-624-7907
(WF)714-621-5319
ned@innosoft.com

AnneMarie Freitas
Microcom, Inc.
(W)617-551-1210
(WF)617-551-1015
afreitas@microcom.com

Ray Freiwirth
RCI
(W)508-453-7857
(WF)508-934-9066
5242391@mcimail.com

Jerry Friesen
Sandia National Laboratories
(W)510-294-3144
(WF)510-294-1225
jafries@sandia.llnl.gov

Karen Frisa
Carnegie Mellon University
(W)412-268-8576
(WF)412-268-4987
karen.frisa@andrew.cmu.edu

Cliff Frost
University of California,
Berkeley
(W)510-642-5360
(WF)510-643-5385
cliff@cmsa.berkeley.edu

Kathryn Fuller
Motorola Codex
(W)508-261-4000
kathryn@codex.com

Vince Fuller
BARRNet
(W)415-723-6860
(WF)415-723-0010
vaf@stanford.edu

Jim Fullton
University of North Carolina
(W)919-962-9107
(WF)919-962-5604
jim_fullton@unc.edu

Kelly Furlong
Boeing Aerospace Operations
(W)407-867-4847
(WF)407-867-4079
kelly@kyle.ksc.nasa.gov

Shari Galitzer
The MITRE Corporation
(W)703-883-6229
(WF)703-883-5200
shari@shari.mitre.org

Maria Gallagher
NASA Science Internet
(W)415-604-3601
(WF)415-604-0063
maria@nsipo.nasa.gov

Owen Gallagher
Stratus Computer, Inc.
(W)508-560-2336
oweng@jjmhome.uucp

Shawn Gallagher
Digital Equipment Corporation
(W)508-952-3519
gallagher@quiver.enet.dec.com

Der-Hwa Gan
3Com Corporation
(W)415-764-5152
dhg@nsd.3com.com

Joan Gargano
University of California,
Davis
(W)916-752-2591
(WF)916-752-9154
jcgargano@ucdavis.edu

Atul Garg
Synoptics Communications, Inc.
(W)408-764-1318
(WF)408-988-5525
agarg@synoptics.com

John Garrett
AT&T Bell Laboratories
(W)908-580-4719
(WF)908-580-4721
jwg@edsel.att.com

Doron Gartner
LANart Corporation
(W)617-444-1994
(WF)617-444-3692
doron@chipcom.com

Gary Gaudet
Digital Equipment Corporation
(W)603-881-2504
(WF)603-881-2257
gaudet@zk3.dec.com

Jisoo Geiter
The MITRE Corporation
(W)703-883-6571
(WF)703-883-7142
geiter@gateway.mitre.org

Tony Genovese
Lawrence Livermore
National Laboratory
(W)415-423-2471
(WF)415-422-0435
genovese@nersc.gov

Elliot Gerberg
Digital Equipment Corporation
(W)508-486-7587
gerberg@took.enet.dec.com

Elise Gerich
University of Michigan
(W)313-936-2120
(WF)313-747-3745
epg@merit.edu

Arlene Getchell
Lawrence Livermore
National Laboratory
(W)510-423-6349
(WF)510-422-0435
getchell@nersc.gov

David Getchell
cisco Systems
(W)415-688-7803
(WF)415-688-8282
getchell@cisco.com

Robert Gilligan
Sun Microsystems, Inc.
(W)415-336-1012
(WF)415-968-4298
Bob.Gilligan@eng.sun.com

Scott Ginsburg
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
ginsburg@wellfleet.com

Fred Glover
Digital Equipment Corporation
(W)603-881-0388
fglover@zk3.dec.com

Joan Goldstein
Digital Equipment Corporation
(W)508-486-7806
(WF)508-486-5788
j_goldstein@tnpubs.enet.dec.com

Rob Graham
ProTools, Inc.
(W)503-645-5400
(WF)503-645-3577
robert_graham@protools.com

Kenneth Grant
Oracle Corporation
(W)415-506-2908
(WF)415-506-7226
kgrant@oracle.com

Pria Graves
3Com Corporation
(W)408-764-6204
priag@nsd.3com.com

Martin Gray
Spider Systems, Ltd.
(W)+44 31 5549424
(WF)+44 31 5540649
mg@spider.co.uk

Terry Gray
University of Washington
(W)206-543-3880
(WF)206-543-3909
gray@cac.washington.edu

Theodore Greene
Exide Electronics
(W)919-870-3353
(WF)919-870-3111

Pete Grillo
Network Innovations
(W)503-526-9766
(WF)503-526-9766
pl0143@mail.psi.net

Daniel Grossman
Motorola Codex
(W)617-821-7333
(WF)617-821-4218
dan@merlin.dev.cdx.mot.com

Phillip Gross
Advanced Network and
Services, Inc.
(W)914-789-5341
(WF)914-789-5310
pgross@nis.ans.net

Chris Gunner
Digital Equipment Corporation
(W)508-486-7792
(WF)508-486-5279
gunner@dsmail.enet.dec.com

Simon Hackett
Internode Systems Pty Ltd
(W)+61 8 373 1339
(WF)+61 8 373 4911
simon@internode.com.au

John Hagan
University of Pennsylvania
(W)215-898-9192
(WF)215-898-9348
hagan@dccs.upenn.edu

Robert Hagens
Advanced Network and
Services, Inc.
(W)914-789-5300
(WF)914-789-5310
hagens@ans.net

William Haggerty
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
haggerty@ctron.com

Jack Hahn
SURAnet
(W)301-982-4600
(WF)301-982-4605
hahn@sura.net

Tony Hain
Lawrence Livermore
National Laboratory
(W)510-422-4200
(WF)510-422-1482
hain@nersc.gov

Neil Haller
Bell Communications Research
(W)201-829-4478
(WF)201-829-2504
nmh@thumper.bellcore.com

Martyne Hallgren
Cornell University
(W)607-254-8767
(WF)607-255-5771
martyne@mitchell.cit.cornell.edu

James Halpin
Stratus Computer, Inc.
(W)508-460-6863
halpin@turkey.sw.stratus.com

Alf Hansen
SINTEF DELAB
(W)+47 7 592982
(WF)+47 7 532586
Alf.Hansen@delab.sintef.no

Daniel Hansen
Network General Corporation
(W)201-616-0986
(WF)201-616-8311
dan@ngc.com

Steve Hardcastle-Kille
ISODE Consortium
(W)+44 71 223 4062
(WF)+44 71 223 3846
s.kille@isode.com

Susan Hares
Merit Network, Inc.
(W)313-936-2095
(WF)313-747-3745
skh@merit.edu

Douglas Hartman
Open Software Foundation
(W)617-621-8818
(WF)617-225-2943
hartman@osf.org

James Harville
University of New Hampshire
(W)603-862-0205
jbh@unh.edu

Dimitry Haskin
Bolt Beranek and Newman
(W)617-873-8609
(WF)617-873-4086
dhaskin@bbn.com

Eugene Hastings
Pittsburgh Supercomputing Center
(W)412-268-4960
(WF)412-268-5832
hastings@a.psc.edu

Alisa Hata
University of Washington
(W)206-543-7732
(WF)206-543-3909
hata@cac.washington.edu

William Hawe
Digital Equipment Corporation
(W)508-486-7666
(WF)508-486-5279
hawe@erlang.enet.dec.com

John Hawthorne
Rome Research Corporation
(W)315-330-7928
(WF)315-330-2396
johnh@tigger.rl.af.mil

Frank Heath
Rockwell International
(W)714-833-6864
(WF)714-833-6863
heath@cmc.com

Ian Heavens
Spider Systems, Ltd.
(W)+44 31 555 5166
(WF)+44 31 555 0664
ian@spider.co.uk

Juha Heinanen
Telecom Finland
(W)+358 49 500958
(WF)+358 31 2432211
juha.heinanen@datanet.tele.fi

Arturo Heinsen
Motorola Codex
(W)508-261-4000
heinsen_a@msm.cdx.mot.com

Fred Henry
AT&T
(W)908-576-3763
(WF)908-576-2993
fgh@sodium.att.com

Ittai Hershman
Advanced Network and Services, Inc.
(W)914-789-5337
(WF)914-789-5310
ittai@nis.ans.net

Paul Hill
Massachusetts Institute of Technology
(W)617-253-0124
(WF)617-258-8736
pbh@mit.edu

Robert Hinden
Sun Microsystems, Inc.
(W)415-336-2082
(WF)415-336-6015
hinden@eng.sun.com

Russ Hobby
University of California,
Davis
(W)916-752-0236
(WF)916-752-9154
rdhobby@ucdavis.edu

Ellen Hoffman
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
ellen_hoffman@um.cc.umich.edu

Ron Hoffmann
Massachusetts Institute of
Technology
hoffmann@mit.edu

J. Paul Holbrook
CICNet, Inc.
(W)313-998-7680
(WF)313-998-6105
holbrook@cic.net

Mark Hollinger
Digital Equipment Corporation
(W)508-486-7936
myth@ssg.lkg.dec.com

Jeffrey Honig
Cornell University
(W)607-255-6460
(WF)607-255-5771
jch@nr-tech.cit.cornell.edu

Eric Hood
NorthWestNet
(W)206-562-3000
(WF)206-562-4822
ehood@nwnet.net

Alton Hoover
Advanced Network and
Services, Inc.
(W)914-789-5325
(WF)914-789-5310
hoover@ans.net

Tim Howes
University of Michigan
(W)313-747-4454
(WF)313-764-5140
Tim.Howes@umich.edu.

Kathleen Huber
Bolt Beranek and Newman
(W)617-873-2520
(WF)617-873-4086
khuber@bbn.com

Steven Hubert
University of Washington
(W)206-543-6384
(WF)206-543-3909
hubert@cac.washington.edu

Jeff Hughes
Hewlett-Packard
(W)719-531-4777
(WF)719-531-4505
jeff@col.hp.com

Christian Huitema
INRIA
(W)+33-93657715
(WF)+33-93657766
christian.huitema@sophia.inria.fr

Erik Huizer
SURFnet
(W)+31 30 310290
(WF)+31 30 340903
huizer@surfnet.nl

Doug Hunt
Bolt Beranek and Newman
(W)617-873-4261
(WF)617-873-6263
dhunt@bbn.com

Wendy Huntoon
Pittsburgh Supercomputing Center
(W)412-268-4960
(WF)412-268-5832
huntoon@a.psc.edu

David Husak
Synernetics, Inc.
(W)508-670-9009
(WF)508-670-9015
dave@synnet.com

William Huyck
ITT Aerospace
(W)201-284-2548

Robin Iddon
AXON Networks, Inc.
(W)+44 31 654 0409
(WF)+44 31 654 0409
axon@cix.compulink.co.uk

Takashi Ikemoto
Fuji Xerox
(W)415-813-7599
(WF)415-813-7393
tikemoto@xerox.com

John Ioannidis
Columbia University
(W)212-854-8120
(WF)212-666-0140
ji@cs.columbia.edu

J. Lance Jackson
Harvard University
(W)617-495-7632
(WF)617-495-0914
jackson@noc.harvard.edu

Lenore Jackson
NASA Science Internet
(W)301-286-4574
(WF)301-286-5152
jackson@nsinic.gsfc.nasa.gov

Ole Jacobsen
Interop Company
(W)415-962-2515
(WF)415-949-1779
ole@interop.com

Ronald Jacoby
Silicon Graphics, Inc.
(W)415-390-1371
(WF)415-969-2314
rj@sgi.com

Dean Jagels
SecureWare, Inc.
(W)404-315-6296
(WF)404-315-0293
dpj@sware.com

Dwight Jamieson
Bell Northern Research
(W)613-763-2791
(WF)613-763-2626
djamies@bnr.ca

Bob Jeckell
3Com Corporation
(W)408-764-6049
(WF)408-764-5003
rrj@3com.com

Carl Jensen
Motorola Codex
(W)508-261-4000
jensen_c@msm.cdx.mot.com

Matthew Jonson
United States Air Force
(W)205-416-4075
(WF)205-416-4738
jonson@server.af.mil

Kevin Jordan
Control Data Systems, Inc.
(W)612-482-6835
(WF)612-482-4455
kej@udev.cdc.com

Dan Jordt
NorthWestNet
(W)206-562-3000
(WF)206-562-4822
danj@nwnet.net

Miriam Kadansky
Xyplex, Inc.
(W)508-264-9900
(WF)508-264-9930
mckadansky@eng.xyplex.com

Todd Kaehler
Digital Equipment Corporation
(W)603-881-1915
(WF)603-881-2257
kaehler@zk3.dec.com

George Kajos
Coral Network Corporation
(W)508-460-6010
(WF)508-481-6258
kajos@coral.com

Scott Kaplan
FTP Software, Inc.
(W)415-543-9001
(WF)415-543-9002
scott@ftp.com

Frank Kastenholz
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
kasten@ftp.com

Neil Katin
SunSoft
(W)415-336-3718
(WF)415-961-0992
katin@eng.sun.com

David Katinsky
Rutgers University
(W)908-932-4783
(WF)908-932-2968
dmk@rutgers.edu

David Kaufman
Magnalink Communications
(W)617-255-9400
(WF)617-255-5885

Nick Kawaguchi
Trillium Digital Systems
(W)310-479-0500
(WF)310-575-0172
mamster@lanai.cs.ucla.edu

James Keller
Sprint
(W)703-689-7016
(WF)703-689-6998
j.keller@sprint.com

Stephen Kent
Bolt Beranek and Newman
(W)617-873-3988
(WF)617-873-4086
kent@bbn.com

Mark Kepke
Hewlett-Packard
(W)303-229-4201
(WF)303-229-3526
mak@cnd.hp.com

Tom Kessler
Sun Microsystems, Inc.
(W)415-336-3145
(WF)415-969-9131
kessler@eng.sun.com

Kenneth Key
University of Tennessee
(W)615-974-5067
(WF)615-974-4404
key@cs.utk.edu

Michael Khalandovsky
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
mlk@ftp.com

Zbigniew Kielczewski
Eicon Technology, Inc.
(W)514-631-2592
(WF)514-631-3092
zbig@eicon.qc.ca

Peter Kirstein
University College London
(W)+44 71 380 7286
(WF)+44 71 387 1397
P.Kirstein@cs.ucl.ac.uk

John Klensin
Massachusetts Institute of
Technology
(W)617-253-1355
(WF)617-491-6266
klensin@infoods.mit.edu

Paulina Knibbe
cisco Systems
(W)415-961-8228
knibbe@cisco.com

Mark Knopper
Merit Network, Inc.
(W)313-763-6061
(WF)313-747-3745
mak@merit.edu

Jim Knowles
Sterling Software
Federal Systems
(W)415-604-3056
(WF)415-604-6999
jknowles@trident.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
stev@ftp.com

Mark Knutsen
IBM Corporation
(W)203-783-7276
(WF)203-783-5046
knutsen@almaden.ibm.com

Deidre Kostick
Bell Communications Research
(W)908-758-2642
(WF)908-758-4177
dck2@sabre.bellcore.com

John Krawczyk
Wellfleet Communications, Inc.
(W)617-275-7750
(WF)617-275-5001
jkrawczy@wellfleet.com

Padma Krishnaswamy
Bell Communications Research
(W)908-758-2362
(WF)908-758-4177
kri@sabre.bellcore.com

Edward Krol
University of Illinois
(W)217-333-7886
(WF)217-244-7089
e-krol@uiuc.edu

Cheryl Krupczak
Georgia Institute of Technology
cheryl@cc.gatech.edu

Alan Kullberg
Bolt Beranek and Newman
(W)617-873-3955
(WF)617-873-4086
akullber@bbn.com

David Kumpf
Hewlett-Packard
(W)916-785-5162
(WF)916-786-9185
kumpf@hprnd.rose.hp.com

Kenrick Kutzler
SynOptics Communications, Inc.
(W)408-764-1256
(WF)408-988-5525
kkutzler@synoptics.com

John Labbe
Merit Network, Inc.
(W)313-936-3335
labbe@merit.edu

Andy Lai
Hewlett-Packard
(W)408-447-2448
(WF)408-447-3660
al@hpindda.cup.hp.com

Lawerence Lang
cisco Systems
(W)415-688-4638
(WF)415-326-1989
llang@cisco.com

Sylvain Langlois
Electricite de France
(W)+33 1 47 65 44 02
(WF)+33 1 47 65 30 01
Sylvain.Langlois@der.edf.fr

Anthony Lauck
Digital Equipment Corporation
(W)508-486-7644
(WF)508-486-5279
lauck@tl.enet.dec.com

Joseph Lawrence
Bell Communications Research
(W)908-758-4146
jcl@sabre.bellcore.com

Walter Lazear
The MITRE Corporation
(W)703-883-6515
(WF)703-883-7142
lazear@gateway.mitre.org

Gordon Lee
FTP Software, Inc.
(W)617-246-0900
gordon@ftp.com

Whay Lee
Motorola Codex
(W)617-821-7244
(WF)617-821-4218
whay@merlin.dev.cdx.mot.com

Barry Leiner
Universities Space
Research Association
(W)415-390-0317
(WF)415-390-0318
leiner@nsipo.nasa.gov

Thomas Lenggenhager
SWITCH
(W)+41 1 261 81 78
(WF)+41 1 261 81 33
lenggenhager@switch.ch

Debra Legare
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
dlegare@nri.reston.va.us

Mark Lewis
Telebit Corporation
(W)408-745-3232
(WF)408-745-3810
mlewis@telebit.com

Nina Lewis
Paramax Systems Corporation
(W)805-987-6811
(WF)805-388-7792
nina@cam.unisys.com

Sharon Lewis
University of Tennessee
(W)615-974-5067
(WF)615-974-4404
lewis@cs.utk.edu

Fong-Ching Liaw
Sun Microsystems, Inc.
(W)415-336-5415
(WF)415-336-3156
fong.liaw@eng.sun.com

Tony Li
cisco Systems
(W)415-688-8186
(WF)415-688-8282
tli@cisco.com

Hock-Koon Lim
Case Western Reserve University
(W)216-368-2982
(WF)216-368-4693
lim@po.cwru.edu

Arthur Lin
Pacific Bell
(W)510-823-1480
(WF)510-867-1405
yalin@srv.pacbell.com

David Lindemulder
NCR
(W)908-576-2276
(WF)908-576-4613
dcl@mtung.att.com

James Lin
Hewlett-Packard
(W)408-447-3482
(WF)408-447-3625
yeejang@cup.hp.com

John Linn
Digital Equipment Corporation
(W)508-486-7384
(WF)508-486-5279
linn@erlang.enet.dec.com

Anthony Lisotta
NASA Ames Research Center
(W)415-604-4634
(WF)415-604-4377
lisotta@nas.nasa.gov

Joshua Littlefield
Cayman Systems, Inc.
(W)617-494-1999
(WF)617-494-5167
josh@cayman.com

Robin Littlefield
Wellfleet Communications, Inc.
(W)617-275-2400
rlittlef@wellfleet.com

Brian Lloyd
Lloyd & Associates
(W)916-676-1147
(WF)916-676-3442
brian@lloyd.com

Daniel Long
Bolt Beranek and Newman
(W)617-873-2766
(WF)617-873-5660
long@nic.near.net

Kim Long
SURAnet
(W)302-982-3214
(WF)302-982-4605
klong@sura.net

E. Paul Love
San Diego Supercomputer Center
(W)619-534-5043
(WF)619-534-5152
loveep@sdsc.edu

Robert Lutz
Stac Electronics
(W)619-431-7474
(WF)619-431-0880

Steven Lynes
Racal-Datacom, Inc.
(W)508-263-9929
lyness@rimail.interlan.com

Ron Mackey
Distributed Systems
International, Inc.
(W)708-665-4693
(WF)708-665-4706
rem@dsiinc.com

Keith Mader
Microcom, Inc.
(W)617-551-1211
(WF)617-511-1015
keith@microcom.com

Carl Madison
Star-Tek, Inc.
(W)508-393-9393
(WF)508-393-6934
carl@startek.com

Andrew Maffei
Woods Hole
Oceanographic Institution
(W)508-457-2000
(WF)508-457-2174
arm@aqua.whoi.edu

Arun Mahajan
Proteon, Inc.
(W)508-898-2800
(WF)508-898-3176
axm@proteon.com

Carl Malamud
carl@malamud.com

Kent Malave
IBM Corporation
(W)512-838-3486
(WF)512-838-3509
kent@chang.austin.ibm.com

Andrew Malis
Bolt Beranek and Newman
(W)617-873-3419
(WF)617-873-4086
malis@bbn.com

Gary Malkin
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
gmalkin@xylogics.com

Allison Mankin
The MITRE Corporation
(W)703-883-7907
(WF)703-883-7142
mankin@gateway.mitre.org

Bill Manning
SESQUINET
(W)713-285-5415
(WF)713-527-6099
bmanning@rice.edu

J. Scott Marcus
Bolt Beranek and Newman
(W)617-873-3075
(WF)617-873-4086
smarcus@bbn.com

April Marine
SRI International
(W)415-859-5318
(WF)415-859-6028
april@nisc.sri.com

Les Matheson
Emerson Computer Power
(W)714-457-3840
(WF)714-457-3633
matheson@cerf.net

Matt Mathis
Pittsburgh Supercomputing
Center
(W)412-268-4960
(WF)412-268-5832
mathis@a.psc.edu

Kim Mayton
Wandel & Goltermann
Technologies, Inc.
(W)919-941-5730
(WF)919-941-5751
mayton@wg.com

Cindy Mazza
Applelink Corporation
(W)508-475-4050
(WF)508-475-0550

Keith McCloghrie
Hughes LAN Systems
(W)415-966-7934
(WF)415-960-3738
kzm@hls.com

Ellen McDermott
Digital Equipment Corporation
(W)617-621-7266
(WF)617-225-2782
emcd@osf.org

Vern McGeorge
Hewlett-Packard
(W)408-447-2768
(WF)408-447-3660
vern@hpda.cup.hp.com

Thomas McGinty
Motorola Codex
(W)617-821-7194
(WF)617-821-4956
mcginty_t*corp_m@msm.cdx.mot.com

Glenn McGregor
Merit Network, Inc.
(W)313-936-3000
(WF)313-747-3745
ghm@merit.edu

John McKenna
IBM Corporation
(W)919-543-7464
(WF)919-543-0723
mckenna@ralvm12.vnet.ibm.com

P.V. McMahon
ICL
(W)+44 734 586211
(WF)+44 734 855106
p.v.mcmahon@rea0803.wins.icl.co.uk

Donna McMaster
SynOptics Communications, Inc.
(W)408-764-1206
(WF)408-764-1299
mcmaster@synoptics.com

Jim McQuaid
Wandel & Goltermann
Technologies, Inc.
(W)919-941-5730
(WF)919-941-5751

Ed Menze
University of Arizona
(W)602-621-4889
(WF)602-621-4246
menze@cs.arizona.edu

Marjo Mercado
Hewlett-Packard
(W)408-447-2826
(WF)408-447-3660
marjo@cup.hp.com

Shehzad Merchant
SRI International
(W)415-859-3594
(WF)415-859-4812
merchant@erg.sri.com

Greg Merrell
Digital Equipment Corporation
(W)415-691-4658
(WF)415-691-4589
merrell@caldec.enet.dec.com

Donald Merritt
Ballistic Research Laboratory
(W)410-278-6808
(WF)410-278-5075
don@brl.mil

Beth Miaoulis
Cayman Systems, Inc.
(W)612-494-1999
(WF)612-494-5167
beth@cayman.com

William Middlecamp
Cray Research, Inc.
(W)612-683-5518
(WF)612-683-5599
wjm@cray.com

Edwin Mier
Meir Communications, Inc.
(W)609-275-7311
(WF)609-275-8813
mier@tigger.jvnc.net

Mike Migliano
Bell Atlantic
(W)414-253-3310
(WF)414-255-7767
bellmf!mike@uwm.edu

Paul Milazzo
Bolt Beranek and Newman
(W)617-873-4325
(WF)617-873-3776
milazzo@bbn.com

Douglas Miller
Telebit Corporation
(W)408-734-3124
dmm@telebit.com

Kenneth Miller
Motorola Codex
(W)617-821-7026
(WF)617-821-4426

Stephen Miller
NEARnet
(W)617-873-4153
smiller@bbn.com

Michael Millikin
Interop Company
(W)415-962-2519
(WF)415-941-2913
mmillikin@interop.com

Cyndi Mills
Bolt Beranek and Newman
(W)617-873-4143
(WF)617-873-5620
cmills@nnsc.nsf.net

Robert Mines
Sandia National Laboratories
(W)510-294-2548
rfm@sandia.llnl.gov

David Minnich
FiberCom, Inc.
(W)703-342-6700
(WF)703-342-5961
dwm@fibercom.com

Greg Minshall
Novell, Inc.
(W)510-975-4507
(WF)510-947-1238
minshall@wc.novell.com

Rohit Mital
ProTools, Inc.
(W)503-645-5400
(WF)503-645-3577
rm@protools.com

Steve Mohr
3Com Corporation
(W)408-764-5850
(WF)408-764-5090
srm@nso.3com.com

Pushpendra Mohta
CERFnet
(W)619-455-3908
(WF)619-455-3990
pushp@cerf.net

Lynn Monsanto
Sun Microsystems, Inc.
(W)415-336-2946
(WF)415-336-3156
monsanto@sun.com

Charlotte Mooers
Bolt Beranek and Newman
(W)617-873-4143
(WF)617-873-5620
mooers@nnsc.nsf.net

Keith Moore
University of Tennessee
(W)615-974-3735
(WF)615-974-8296
moore@cs.utk.edu

Paul Moran
3Com Corporation
(W)+44 2231000
(WF)+44 2254701
Paul_Moran@3com.com

Dennis Morris
Defense Information
Systems Agency
(W)703-487-3329
(WF)703-487-3190
morrisd@imo-uvax.disa.mil

Donald Morris
National Center for
Atmospheric Research
(W)303-497-1282
(WF)303-497-1137
morris@ucar.edu

John Moy
Proteon, Inc.
(W)508-898-2800
(WF)508-366-8901
jmoy@proteon.com

Patrick Mullaney
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
mullaney@ctron.com

Chris Munnelly
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
munnelly@xylogics.com

John Murray
Premenos Corporation
(W)510-602-2000
(WF)510-602-2133
murray@premenos.sf.ca.us

Muki Murthy
Microcom, Inc.
(W)914-968-2300
(WF)914-968-7100
muki@microcom.com

Chris Myers
Washington University
(W)314-935-7390
(WF)314-935-7302
chris@wugate.wustl.edu

Rina Nathaniel
RAD Network Devices
(W)+972 3 6458555
(WF)+972 3 5447368
rina!rnd!rndi@uunet.uu.net

Mark Needleman
University of California
(W)510-987-0530
(WF)510-763-9673
mhn@stubbs.ucop.edu

Clifford Neuman
Information Sciences Institute
(W)310-822-1511
bcn@isi.edu

Scott Newman
Digital Equipment Corporation
(W)603-884-6123
(WF)603-884-0829
newman@broke.enet.dec.com

Richard Newton
Harris Computer Systems
(W)703-448-5425
(WF)703-827-0335
rnewton@csd.harris.com

Jack Ng
Network Systems Corporation
(W)612-424-4888
jkn@network.com

Luc Nguyen
BellSouth
(W)404-420-8424
(WF)404-880-9011
luc@bellsouth.com

Andy Nicholson
Cray Research, Inc.
(W)612-683-5473
(WF)612-683-5599
droid@cray.com

Sam Nicholson
Network Imaging Corporation
scion@pblx.knox.tn.us

William Nichols
Digital Equipment Corporation
(W)508-486-5266
(WF)508-486-7417
nichols@wick.enet.dec.com

Tom Nisbet
Telecommunications Techniques
Corporation
(W)301-353-1550
(WF)301-353-1536
nisbet@tt.com

Erik Nordmark
Sun Microsystems, Inc.
(W)415-336-2788
(WF)415-969-9131
nordmark@eng.sun.com

Bill Norton
Merit Network, Inc.
(W)313-936-2656
(WF)313-747-3745
wbn@merit.edu

Eric Nowak
Advanced Network and
Services, Inc.
(W)703-758-7705
nowak@ans.net

Hank Nussbacher
Tel Aviv University
(W)+972-3-6408309
hank@vm.tau.ac.il

Mary Christine O'Connor
Racal-Datacom, Inc.
(W)508-263-9929
oconnor@interlan.com

Ari Ollikainen
Lawrence Livermore
National Laboratory
(W)510-423-5962
(WF)510-423-8744
ari@es.net

Steven Onishi
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
sonishi@wellfleet.com

Alan Oppenheimer
Apple Computer, Inc.
(W)408-974-4744
(WF)408-862-1159
oppenheime1@applelink.apple.com

Roger Osmond
Bytex Corporation
(W)508-366-8000
(WF)508-366-0244
bytex!rfo@uunet.uu.net

Kraig Owen
Advanced Network and
Services, Inc.
(W)313-764-3600
(WF)313-747-3745
tko@merit.edu

Bob Page
SunConnect
(W)415-336-2745
(WF)415-336-3156
bob.page@eng.sun.com

Larry Palmer
Active Software, Inc.
(W)508-263-7651
(WF)508-263-7651
larry%acts@uunet.uu.net

Larry Palmer
Digital Equipment Corporation
(W)603-881-0329
lp@decvax.dec.com

Luc Pariseau
Digital Equipment Corporation
(W)508-952-3563
(WF)508-952-3023
pariseau@quiver.enet.dec.com

J. Bradford Parker
Cayman Systems, Inc.
(W)617-494-1999
(WF)617-494-5167
brad@cayman.com

Philippe Park
Bolt Beranek and Newman
(W)617-873-2892
ppark@bbn.com

Jay Pasco-Anderson
Motorola Codex
(W)617-821-7324
jay@merlin.dev.cdx.mot.com

Brad Passwaters
SURAnet
(W)301-982-4600
(WF)301-982-4605
bjp@sura.net

Rakesh Patel
Rutgers University
(W)908-932-3436
(WF)908-932-2968
rapatel@hardees.rutgers.edu

Andrew Patka
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
apatka@wellfleet.com

Michael Patton
Massachusetts Institute of
Technology
(W)617-253-6061
(WF)617-258-8682
map@lcs.mit.edu

John Payne
Wang Laboratories
(W)508-967-1526
jop@wang.com

James Pearl
Microcom, Inc.
(W)617-551-1215
(WF)617-551-1015
jim@microcom.com

Charles Perkins
IBM Corporation
(W)914-784-7350
(WF)914-784-7007
perk@watson.ibm.com

David Perkins
SynOptics Communications, Inc.
(W)408-764-1516
(WF)408-988-5525
dperkins@synoptics.com

Donald Perkus
Chipcom Corporation
(W)508-460-8922
(WF)508-460-8952
dperkus@chipcom.com

Ed Perreault
Microcom, Inc.
(W)617-551-1221
(WF)617-551-1015
ed@python.eng.microcom.com

Marsha Perrott
PREPnet
(W)412-268-7870
(WF)412-268-7875
mlp+@andrew.cmu.edu

Roy Perry
US West Advanced Technologies
(W)303-541-6245
(WF)303-541-6276
rperry@uswat.uswest.com

Eric Peterson
Xyplex, Inc.
(W)508-264-9900
(WF)508-264-9930
elpeterson@eng.xyplex.com

David Piscitello
Bell Communications Research
(W)908-758-2286
(WF)908-758-4177
dave@sabre.bellcore.com

Mel Pleasant
Rutgers University
(W)908-932-2023
(WF)908-932-2968
pleasant@cs.rutgers.edu

Narendra Popat
Frontier Software Development
(W)508-851-8872
(WF)508-851-6956
albin@frontier.com

Bill Pope
Open Software Foundation
(W)617-621-8889
zpope@osf.org

Jack Pope
University of San Diego
(W)619-260-4810
(WF)619-260-4669
pope@usdcsv.acusd.edu

Lars Poulsen
Rockwell CMC
(W)805-968-4262
(WF)805-968-8256
lars@cmc.com

Jon Postel
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
postel@isi.edu

Michael Powell
Pacific Bell
(W)209-454-4267
(WF)209-323-5638
mdpowel@pacbell.com

Venkat Prasad
3Com Corporation
(W)408-764-5226
(WF)408-764-5003
vsp@3com.com

Russell Pretty
Bell Northern Research
(W)613-763-9400
(WF)613-763-9225
pretty@bnr.ca

Mark Prior
The University of Adelaide
(W)+61 8 228 5680
(WF)+61 8 223 6245
mrp@itd.adelaide.edu.au

Stephanie Price Marasciullo
Rockwell CMC
(W)805-968-4262
(WF)805-968-8256
price@cmc.com

James Quigley
Hewlett-Packard
(W)408-447-3505
(WF)408-447-3660
jim_quigley@hp6600.desk.hp.com

Mary Quinn
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
quinn@maui.InterLan.Com

Scott Quinn
Bell Communications Research
(W)201-829-4147
(WF)201-829-5976
smq@aries.bellcore.com

Joe Ragland
MCNC
(W)919-248-1404
(WF)919-248-1405
jrr@concert.net

Bala Rajagopalan
AT&T Bell Laboratories
(W)908-949-8017
(WF)908-949-1726
braja@qsun.att.com

K. K. Ramakrishnan
Digital Equipment Corporation
(W)508-486-7267
(WF)508-486-5279
rama@erlang.enet.dec.com

Richard Ramos
AT&T Bell Laboratories
(W)908-576-6267
ramos@mtunm.att.com

Wally Ramsey
The MITRE Corporation
(W)402-292-5889
(WF)402-292-1792
wbr@mitre.org

Christopher Ranch
Novell, Inc.
(W)408-473-8667
(WF)408-946-3447
cranch@novell.com

Ed Reeder
IBM Corporation
(W)919-543-2154
(WF)919-543-0723
EREEDER@ralvm12.vnet.ibm.com

James Rees
University of Michigan
(W)313-763-4174
(WF)313-763-4434
jim.rees@umich.edu

James Reeves
SynOptics Communications, Inc.
(W)408-764-1440
(WF)408-988-5525
jreeves@synoptics.com

Robert Reilly
K12NET
(W)413-443-0027
rreilly@athena.mit.edu

Sonya Reimer
Ballistic Research Laboratory
(W)410-278-8676
(WF)410-278-5075
sonya@brl.mil

David Reiser
Motorola Codex
(W)508-261-4000
dsr@codex.prds.cdx.mot.com

Yakov Rekhter
IBM Corporation
(W)914-945-3896
yakov@watson.ibm.com

Sheri Repucci
Merit Network, Inc.
(W)313-936-3335
smr@merit.edu

Robert Reschly
U.S. Army Ballistic
Research Laboratory
(W)410-278-6808
(WF)410-278-5075
reschly@brl.mil

Tom Reuther
Digital Equipment Corporation
(W)508-486-5517
(WF)508-486-5554
reuther@nac.dec.com

Joyce K. Reynolds
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
jkrey@isi.edu

Bradley Rhoades
3M
(W)612-736-2874
(WF)612-736-0431
bdrhoades@mmc.mmmg.com

Salvatore Ricci
AT&T Bell Laboratories
(W)908-576-6337
(WF)908-576-4613
ricci@mtunm.att.com

April Richstein
Department of Defense
(W)410-684-7302
(WF)410-850-7166
abm@tycho.ncsc.mil

Anil Rijsinghani
Digital Equipment Corporation
(W)508-952-3520
anil@levers.enet.dec.com

Mike Ritter
Apple Computer, Inc.
(W)408-862-8088
(WF)408-257-7982
mwritter@applelink.apple.com

Michael Roberts
EDUCOM
(W)202-872-4200
(WF)202-872-4318
roberts@educom.edu

Kary Robertson
Concord Communications, Inc.
(W)508-460-4646
(WF)508-481-9772
kr@concord.com

Sam Roberts
Farallon Computing, Inc.
(W)510-596-9215
(WF)510-596-9024
sroberts@farallon.com

David Robinson
Computervision Corporation
(W)617-275-1800
drb@relay.prime.com

Jon Rochlis
Massachusetts Institute of
Technology
(W)617-253-4222
(WF)617-258-8736
jon@mit.edu

Benny Rodrig
RAD Network Devices
(W)+972 3 6458555
(WF)+972 3 5447368
4373580@mcimail.com

Manoel Rodrigues
AT&T Bell Laboratories
(W)908-949-4655
(WF)908-949-1726
manoel_rodrigues@att.com

Guenter Roeck
Conware Computer Consulting GmbH
(W)+49 721 9495 0
(WF)+49 721 9495 146
roeck@conware.de

Jim Romaguera
NetConsult AG
(W)+41 31 235750
(WF)+41 31 235792
romaguera@cosine-mhs.switch.ch

Dan Romascanu
Lannet Data Communications
(W)+972 3 6458414
(WF)+972 3 5447146
dan@lannet.com

Marshall Rose
Dover Beach Consulting, Inc.
(W)415-968-1052
(WF)415-968-2510
mrose@dbc.mtview.ca.us

Rich Rosenbaum
Digital Equipment Corporation
(W)508-486-5927
(WF)508-486-7417
rosenbaum@lkg.dec.com

Simon Rosenthal
Locus Computing Corporation
(W)617-229-4980
(WF)617-229-2969
simon@bos.locus.com

Christine Ross
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
ross@warp.interlan.com

Kenneth Rossen
INTERACTIVE Systems Corporation
(W)617-661-7474
(WF)617-661-2070
kenr@isc.com

Karen Roubicek
The Faxon Company
(W)617-329-3350
(WF)617-326-5484
roubicek@faxon.com

Allan Rubens
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
acr@merit.edu

Gregory Ruth
Bolt Beranek and Newman
(W)617-873-3150
(WF)617-873-8202
gruth@bbn.com

Anthony Rutkowski
Sprint International
(W)703-689-5080
(WF)703-689-6715
amr@sprint.com

Salvatore Salamone
Data Communications Magazine
(W)212-512-6020
(WF)212-512-6833
0005379565@mcimail.com

Timothy Salo
Minnesota Supercomputer Center
(W)612-626-0347
(WF)612-624-6550
tjs@msc.edu

Jonathan Sandberg
Matsushita Information
Technology Laboratory
(W)609-497-4600
(WF)609-497-4013
Jss@mitl.com

Henry Sanders
Microsoft Corporation
(W)206-882-8080
henrysa@microsoft.com

Paul Sangster
Advanced Network and
Services, Inc.
(W)703-758-7706
sangster@ans.net

Jon Saperia
Digital Equipment Corporation
(W)508-496-8333
(WF)508-496-9929
saperia@tcpjon.ogo.dec.com

Michael Sapich
Conware Computer Consulting
GmbH
(W)+49 721 9495 0
(WF)+49 721 9495 146
sapich@conware.de

Miguel Sasson
Xylogics, Inc.
(W)617-272-0924
(WF)617-272-3159
sasson@xylogics.com

Srinivas Sataluri
AT&T Bell Laboratories
(W)908-949-7782
(WF)908-949-1726
sri@qsun.att.com

John Sax
Bolt Beranek and Newman
(W)617-873-3419
(WF)617-873-4086
jsax@bbn.com

Michael Scanlon
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
scanlon@interlan.com

Jeffrey Schiller
Massachusetts Institute of
Technology
(W)617-253-0161
(WF)617-258-8736
jis@mit.edu

Richard Schmalgemeier
Merit Network, Inc.
(W)313-936-3000
rgs@merit.edu

Eve Schooler
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
schooler@isi.edu

Martin Schulman
Sprint International
(W)703-689-7171
(WF)703-689-5633
mas@loyola.edu

Aaron Schuman
Silicon Graphics, Inc.
(W)415-335-1901
(WF)415-969-2314
schuman@sgi.com

Jim Scott
Digital Equipment Corporation
(W)508-493-5514
scott@kali.enet.dec.com

Hellen Sears
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
sears@interlan.interlan.com

Tim Seaver
MCNC
(W)919-248-1973
(WF)919-248-1455
tas@concert.net

Jennifer Sellers
NASA Science Internet
(W)202-434-8954
sellers@nsipo.arc.nasa.gov

Karen Seo
Bolt Beranek and Newman
(W)617-873-3152
(WF)617-873-3776
kseo@bbn.com

Allwyn Sequeira
Ungermann-Bass
(W)408-987-6501
(WF)408-970-3789
allwyn@nmu.ub.com

Koichiro Seto
Hitachi Cable, Ltd.
(W)+81 294 25 3190
(WF)+81 294 42 2787
seto@hitachi-cable.co.jp

Vincent Sgro
Rutgers University
(W)908-932-3433
(WF)908-932-0537
sgro@cs.rutgers.edu

Houman Shafiezadeh
Exide Electronics
(W)919-872-3030
(WF)919-878-2385
houman@exide.com

Ron Sharp
AT&T Bell Laboratories
(W)201-386-4479
(WF)201-386-5449
rls@neptune.att.com

Daisy Shen
IBM Corporation
(W)914-945-3632
(WF)914-945-4439
daisy@watson.ibm.com

Erik Sherk
SURAnet
(W)301-982-4600
(WF)301-982-4605
sherk@sura.net

Robert Shirey
The MITRE Corporation
(W)703-883-7210
(WF)703-883-1397
shirey@mitre.org

John Shriver
Proteon, Inc.
(W)508-898-2800
(WF)508-898-3176
jas@proton.com

Jeremy Siegel
3Com Corporation
(W)408-764-5247
(WF)408-764-5003
jzs@nsd.3com.com

Douglas Simmons
IBM Corporation
(W)408-452-4198
(WF)408-452-4153

William Simpson
Computer Systems
Consulting Services
(W)313-543-3269
Bill_Simpson@um.cc.umich.edu

Henry Sinnreich
MCI Telecommunications
Corporation
(W)214-918-5231
(WF)214-918-6024
hsinnreich@mcimail.com

Sam Sjogren
TGV, Inc.
(W)408-427-4366
(WF)408-427-4365
sjogren@tgv.com

Keith Sklower
University of California,
Berkeley
(W)510-642-9587
sklower@cs.berkeley.edu

Richard Slade
Harris Computer Systems
(W)305-974-1700
(WF)305-977-5580
ricks@ssd.csd.harris.com

Frank Slaughter
Shiva Corporation
(W)617-252-6302
(WF)617-252-6852
fgs@shiva.com

Timon Sloane
Timon Sloane Consulting
Services
(W)415-941-9726
(WF)415-941-4679
peernet!timon@uunet.uu.net

Lansing Sloan
Lawrence Livermore
National Laboratory
(W)510-422-4356
(WF)510-423-8534
ljsloan@llnl.gov

Michael Slocombe
Bolt Beranek and Newman
(W)617-873-6014
(WF)617-491-1496
slocombe@bbn.com

Andrew Smith
SynOptics Communications, Inc.
(W)408-764-1574
(WF)408-988-5525
asmith@synoptics.com

Carl Smith
Sun Microsystems, Inc.
(W)415-336-1050
cs@eng.sun.com

Jane Smith
MCNC
(W)919-248-1886
(WF)919-248-1455
jds@jazz.concert.net

Mark Smith
University of Michigan
(W)313-764-2277
(WF)313-764-5140
mcs@umich.edu

Mark Smith
AT&T Bell Laboratories
(W)919-279-6629
(WF)919-279-5873
mark@neptune.att.com

Patricia Smith
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
psmith@merit.edu

Frank Solensky
Clearpoint Research
Corporation

Evan Solley
Farallon Computing, Inc.
(W)510-596-9201
(WF)510-596-9024
evan@farallon.com

Karen Sollins
Massachusetts Institute of Technology
(W)617-253-6006
(WF)617-258-8682
sollins@lcs.mit.edu

Simon Spero
Wide Area Information Services
(W)+44 81 500 0067
(WF)+44 81 500 3000
ses@cmns.think.com

Rajesh Srivastava
Tandem Computers, Inc.
(W)408-285-4103
(WF)408-285-4324
rajesh@flash.tandem.com

Michael St. Johns
Department of Defense
(W)301-688-6741
stjohns@umd5.umd.edu

Martha Steenstrup
Bolt Beranek and Newman
(W)617-873-3192
(WF)617-873-3776
msteenst@bbn.com

Lawrence Stefani
Digital Equipment Corporation
(W)508-952-3769
(WF)508-952-3023
stefani@quiver.enet.dec.com

Einar Stefferud
Network Management Associates, Inc.
(W)714-842-3711
(WF)714-848-2091
stef@nma.com

Louis Steinberg
IBM Corporation
(W)203-783-7175
louiss@vnet.ibm.com

Brad Steinka
Microcom, Inc.
(W)617-551-1209
(WF)617-551-1015
brad@microcom.com

Scott Stein
Apple Computer, Inc.
(W)408-974-3124
(WF)408-974-0781
scotts@apple.com

Maureen Stetson
Motorola Codex
(W)508-261-4000
stetson_m@msm.cdx.mot.com

Bob Stewart
Xyplex, Inc.
(W)508-264-9900
(WF)508-264-9930
rlstewart@eng.xyplex.com

David Stine
cisco Systems
(W)415-688-4681
(WF)415-688-8282
dstine@cisco.com

Bernhard Stockman
SUNET
(W)+46 8 790 6519
(WF)+46 8 241 179
boss@sunet.se

Adam Stolinski
Liebert Corporation
(W)714-457-3600
(WF)714-457-3787

Roxanne Streeter
NASA Science Internet
(W)415-604-4845
(WF)415-604-0063
streeter@nsipo.arc.nasa.gov

Chris Sullivan
Gandalf Data Limited
(W)613-723-6500
(WF)613-226-1717
mm@gandalf.ca

Terrance Sullivan
Newbridge Networks Corporation
(W)613-591-3600
(WF)613-591-3680
terrys@newbridge.com

George Swallow
Bolt Beranek and Newman
(W)617-873-3102
(WF)617-873-6263
gswallow@bbn.com

Richard Sweatt
SynOptics Communications, Inc.
(W)408-764-1260
(WF)408-764-1299
rsweatt@synoptics.com

Bruce Taber
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-635-9820
taber@interlan.com

Wayne Tackabury
Pacer Software, Inc.
(W)508-898-3300
(WF)508-366-1356
wft@pacersoft.com

Thomas Talpey
Open Software Foundation
(W)617-621-8923
(WF)617-225-2782
tmt@osf.org

Sally Tarquinio
The MITRE Corporation
(W)703-883-7987
(WF)703-883-7142
sallyt@gateway.mitre.org

Fumio Teraoka
Sony Computer Science
Laboratory, Inc.
(W)+81 3 3448 4380
(WF)+81 3 3448 4273
tera@csl.sony.co.jp

Marten Terpstra
RIPE
(W)+31 20 592 5065
(WF)+31 20 592 5155
terpstra@ripe.net

Kaj Tesink
Bell Communications Research
(W)908-758-5254
(WF)908-758-4196
kaj@cc.bellcore.com

Rodney Thayer
Sable Technology Corporation
(W)617-332-7292
(WF)617-332-7970
rodney@world.std.com

Mark Therieau
Microcom, Inc.
(W)617-551-1195
(WF)617-551-1015
markt@python.eng.microcom.com

Chris Thomas
Intel Corporation
(W)801-379-2251
(WF)801-373-1599
chris_thomas@ccm.hf.intel.com

Steven Thompson
Soft-Switch, Inc.
(W)215-640-9600
(WF)215-651-5358
sjt@gateway.ssw.com

Steve Thompson
Optical Data Systems, Inc.
(W)214-234-6400
(WF)214-234-4059
steve@ods.com

Dean Throop
Data General Corporation
(W)919-549-8421
(WF)919-248-5942
throop@dg-rtp.dg.com

Linda Tom
Racal-Datacom, Inc.
(W)508-263-9929
(WF)508-263-8655
toml@interlan.interlan.com

Sylvia Tom
Tandem Computers, Inc.
(W)408-285-4119
(WF)408-285-4324
tom_sylvia@tandem.com

Claudio Topolcic
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
topolcic@nri.reston.va.us

Chuck Townsend
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-332-1019
townsend@ctron.com

Win Treese
Digital Equipment Corporation
(W)617-621-6615
(WF)617-621-6650
treese@crl.dec.com

Panos-Gavriil Tsigaridas
Research Center of
Open Communications Systems
(W)+49 30 25499232
(WF)+49 30 25499202
tsigaridas@fokus.berlin.gmd.dbp.de

Theodore Ts'o
Massachusetts Institute of
Technology
(W)617-253-8091
(WF)617-258-8736
tytso@mit.edu

Paul Tsuchiya
Bell Communications Research
(W)201-829-4484
(WF)201-829-2504
tsuchiya@thumper.bellcore.com

Stephen Tsun
3Com Corporation
(W)408-764-6978
(WF)408-764-5002
snt@nsd.3com.com

Ahmet Tuncay
SynOptics Communications, Inc.
(W)408-764-1279
(WF)408-764-1299
atuncay@synoptics.com

Mark Uhrmacher
Massachusetts Institute of
Technology
maui@tabasco.lcs.mit.edu

Robert Ullmann
Process Software Corporation
(W)508-879-6994
(WF)508-879-0042
ullmann@process.com

Dono van-Mierop
3Com Corporation
(W)408-764-5930
(WF)408-764-5003
dono_van_mierop@3mail.3com.com

Kannan Varadhan
OARnet
(W)614-292-4137
(WF)614-292-7168
kannan@oar.net

Gregory Vaudreuil
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
gvaudre@nri.reston.va.us

Paul Vazquez
DIA DSO-3A
(W)703-284-0795
(WF)703-284-0853
vazquez@dockmaster.ncsc.mil

Ross Veach
University of Illinois
at Urbana-Champaign
(W)217-244-4274
(WF)217-244-7089
rrv@uiuc.edu

Andrew Veitch
Bolt Beranek and Newman
(W)617-873-2588
(WF)617-873-3776
aveitch@bbn.com

John Veizades
Apple Computer, Inc.
(W)408-974-2672
(WF)408-974-9083
veizades@apple.com

Desikan Venkatrangan
Metrix, Inc.
(W)603-888-7000
(WF)603-891-2796
venkat@geoduck.matrix.com

Curtis Villamizar
Advanced Network and
Services, Inc.
(W)914-789-5300
(WF)914-789-5310
curtis@ans.net

John Vollbrecht
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
jrv@merit.edu

Bernhard Volz
Process Software Corporation
(W)508-879-6694
(WF)508-879-0042
volz@process.com

Huyen Vu
Defense Information
Systems Agency
(W)703-487-3038
(WF)703-487-3190
vi@polaris.disa.mil

Daniel Vukelich
The MITRE Corporation
(W)617-271-2943
vukelich@mbunix.mitre.org

John Wagner
Princeton University
(W)609-258-6043
(WF)609-258-3943
jwagner@princeton.edu

David Waitzman
Bolt Beranek and Newman
(W)617-873-4323
(WF)617-873-3776
djw@bbn.com

Steven Waldbusser
Carnegie Mellon University
(W)412-268-6628
(WF)412-268-4987
waldbusser@andrew.cmu.edu

Jesse Walker
Digital Equipment Corporation
(W)508-486-7326
walker@eider.lkg.dec.com

Justin Walker
Apple Computer, Inc.
(W)408-974-4427
(WF)408-974-0892
justin@apple.com

Sandro Wallach
ELF Communications
(W)617-776-0204
sandro@elf.com

Kathleen Walsh
Cabletron Systems, Inc.
(W)603-332-9400
walshy@ctron.com

David Wang
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
wang@xylogics.com

Carol Ward
Westnet
(W)303-492-5860
(WF)303-492-4198
cward@westnet.net

William Warner
Department of Administrative
Services, State of Ohio
(W)614-466-6683
(WF)614-752-6108
warner@ohio.gov

Albin Warth
Frontier Software Development
(W)508-851-8872
(WF)508-851-6956
albin@frontier.com

Ray Wasson
Liebert Corporation
(W)714-457-3600
(WF)714-457-3787

Scott Wasson
Xyplex, Inc.
(W)508-264-9900
(WF)508-264-9930
sgwasson@eng.xyplex.com

Charles Watt
SecureWare, Inc.
(W)404-315-6296
(WF)404-315-0293
watt@sware.com

James Watt
Newbridge Networks Corporation
(W)613-591-3600
(WF)613-591-3680
jamesw@newbridge.com

Luanne Waul
Ascom Timeplex
(W)310-443-5651
(WF)310-443-5582
luanne@wwtc.timeplex.com

Chris Weider
Merit Network, Inc.
(W)313-936-3000
clw@merit.edu

Liming Wei
Sun Microsystems, Inc.
(W)415-336-1085
(WF)415-336-3156
liming.wei@eng.sun.com

Abel Weinrib
Bell Communications Research
(W)201-829-4454
(WF)201-829-5889
abel@bellcore.com

Jonathan Wenocur
Shiva Corporation
(W)617-252-6347
(WF)617-252-6852
jhw@shiva.com

Jehu Westmark
Centel Corporation
(W)904-599-1843
(WF)904-599-1033
westmark@ds1.scri.fsu.edu

Moira West
Carnegie Mellon University
(W)412-268-7090
(WF)412-268-6989
mjw@cert.org

Evan Wetstone
SESQUINET
(W)713-285-5173
(WF)713-527-6099
evan@rice.edu

Brien Wheeler
The MITRE Corporation
(W)617-271-2567
(WF)617-271-2352
blw@mitre.org

Chris Wheeler
University of Washington
(W)206-543-7741
(WF)206-543-3909
cwheeler@cac.washington.edu

Gerard White
Applitek Corporation
(W)508-475-4050
(WF)508-475-0550

Kathleen Wilde
Digital Equipment Corporation
(W)603-881-0387
(WF)508-486-5279
wilde@decvax.dec.com

Rick Wilder
Advanced Network and
Services, Inc.
(W)703-758-7703
wilder@ans.net

Jim Williams
Merit Network, Inc.
(W)313-764-9423
jimw@merit.edu

Scott Williamson
Network Solutions, Inc.
(W)703-802-8405
(WF)703-802-8376
scottw@nic.ddn.mil

Peter Williams
University College London
(W)+44 71 387 6050
p.williams@uk.ac.ucl.cs

Walter Wimer
Carnegie Mellon University
(W)412-268-6252
(WF)412-268-4987
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
(W)708-252-7236
(WF)708-252-5983
lwinkler@anl.gov

Steven Winnett
Bolt Beranek and Newman
(W)617-873-6101
(WF)617-873-4086
swinnett@bbn.com

W. Stan Wisseman
Oracle Corporation
(W)415-506-2621
(WF)415-506-7200
swissema@oracle.com

Cathy Wittbrodt
Lawrence Livermore
National Laboratory
(W)510-466-4016
(WF)510-422-0435
cjw@nersc.gov

Jane Wojcik
Bolt Beranek and Newman
(W)617-873-6247
(WF)617-873-4086
jwojcik@bbn.com

Jeff Wong
AT&T Bell Laboratories
(W)908-580-5361
(WF)908-580-4721
jaw@io.att.com

Steven Wong
Digital Equipment Corporation
(W)508-486-5344
(WF)508-486-7417
wong@took.enet.dec.com

Robert Woodburn
SAIC
(W)703-448-1683
(WF)703-734-3323
woody@cseic.saic.com

Richard Woundy
IBM Corporation
(W)203-783-7996
(WF)203-783-5046
rwoundy@rhqvm21.vnet.ibm.com

Russ Wright
Lawrence Berkeley Laboratory
(W)510-486-6965
(WF)510-486-6363
wright@lbl.gov

John Wroclawski
Massachusetts Institute of
Technology
(W)617-253-7885
(WF)617-253-2673
jtw@lcs.mit.edu

Honda Wu
Network Application Technology
(W)408-370-4365
(WF)408-370-4222
natadm!honda@uunet.uu.net

Brian Young
Best Power Technology, Inc.
(W)1-800-356-5794
(WF)608-565-2221

Kiho Yum
3Com Corporation
(W)408-764-6406
kxy@nsd.3com.com

Yung-Chao Yu
AT&T Bell Laboratories
(W)908-949-0555
(WF)908-949-4673
yy@qsun.att.com

Paul Zawada
National Center for
Supercomputing Applications
(W)217-244-4728
(WF)217-244-7396
Zawada@ncsa.uiuc.edu

Fred Ziegler
Aspen Technology, Inc.
(W)617-497-9010
(WF)617-497-7806
ziegler@aspentec.com

Uwe Zimmerman
Metrix, Inc.
(W)603-888-7000
(WF)603-891-2796
metrix!uwe@uunet.uu.net

Joseph Zur
Fibronics International
(W)+972 4 313679
(WF)+972 4 550266
zur@fibhaifa.com

Appendix B

Audiocast Participants

Below is the list of IETF audiocast participants as seen by an instance of the vat program running at ISI. This list is cumulative from late on July 10th, when the IP multicast tunnels were being set up, until 6:00 p.m. EDT on July 18th. There were 170 hosts from ten countries:

1. Australia
2. Canada
3. China
4. France
5. Japan
6. Netherlands
7. Norway
8. Sweden
9. United Kingdom
10. United States

Somewhere between 3.5 - 4 million packets were transmitted from the IETF (there is some uncertainty about the "lost" counter).

This file is available by anonymous ftp from venera.isi.edu in: pub/ietf-audiocast.txt. The list of participating hosts is shown twice, in numerical order by IP address and in alphabetical order on the domain name with the fields reversed.

```
melfyn@conger.mel.dit.CSIRO.AU
worsley@guppy.mel.dit.CSIRO.AU
ajw@squid.mel.dit.CSIRO.AU
Geoff Huston (AARNet, Australia)
Mark Turner (Australian National University)
Simon Coppins (coppins@arch.adelaide.edu.au)
chrisc@blizzard.mpce.mq.edu.au
```

kre@munnari.OZ.AU
Andy Hooper, Queen's U, Canada
John Demco, UBC, Vancouver, Canada
John Roth, University of Toronto
Eric Carroll, University of Toronto, Canada
Canadian Rate Adaptation Gateway
Lee Oattes, UofToronto, Canada
Adam Feigin, ETH Zuerich
Erik Fair (Apple Computer)
Paul Milazzo @ BBN
CLynn (BBN) Boston Ma USA
dartbbn@DART5.BBN.COM
rrosales@HAIN.BBN.COM
Bob Clements @ BBN
rmf@quartz.bbn.com
schroder@QUILL.BBN.COM
chowe@TOPAZ.BBN.COM
root@hampton-ss2.cisco.com
Paul Traina (pst@cisco.com)
denny@mm6.erg.sri.com
LUO @ USL
PARC 32kbps gateway
jlarson@parc.xerox.com
Mark Verber (PARC)
Daniel Zappala
lixia@dartvader.parc.xerox.com
juhlig@dollar.parc.xerox.com
jchow@kiwi.parc.xerox.com
Ron Frederick
Steve Deering (Xerox PARC)
nichols@osprey.parc.xerox.com
swinehar@ptarmigan.parc.xerox.com
xiong@ravern.parc.xerox.com
jchow@thyron.parc.xerox.com
dalfonso@weatherby.parc.xerox.com
Sugih Jamin (USC@PARC)
parker@sparky.fac.cs.cmu.edu
sw@sun.fac.cs.cmu.edu
hsm@cz.SEI.CMU.EDU
Marcus Augustus Alzona (Software Engineering Institute, CMU, Pittsburgh, PA, USA)
Scott Brim (swb@cornell.edu)
Dick Cogger - Cornell
Jeffrey C Honig (Cornell University)
root@OITSUN.CIT.CORNELL.EDU
edguer@guardian.CES.CWRU.Edu

limpach@elvis.INS.CWRU.Edu
lim@merlion.INS.CWRU.Edu
Stephen Trier (CWRU) Cleveland OH USA
Dan Brown (IRIS/INS-Telecom)
Dan Brown CWRU
davidc@sirius.net.Hawaii.Edu
ISI, Los Angeles, CA
Statistics Accumulator at ISI
Jeff Bailey (Kent State Univ)
IETF Multimedia Multicast
IETF Terminal Room
IETF Listener
Ira Fuchs (Princeton)
Larry Rogers (Princeton)
CIT Systems Group (Princeton)
hardware@darwin-146.psc.edu
John Deuel @ Rice U.
Listening in @ Rice U.
Scanner (RPI)
Jon Finke (RPI)
Thomas Brisco @ Rutgers
moreland@happy.sdsc.edu
hutton@opus.sdsc.edu
Lobby @ San Diego Supercomputer Center
root@pravda.sdsc.edu
hwb@serendip.sdsc.edu
Mark J. Steiglitz (steig@cs.stanford.edu)
Networking Systems @ Stanford
Edwin Allum, University of Toronto, Canada
Dave Galloway drg@csri.toronto.edu
sandra@csri.toronto.edu
mart@genie.csri.toronto.edu
evan@chop.ele.toronto.edu
Erik Perkins (Udel)
mills@malarky.udel.edu
Henning Schulzrinne@demo8
hgschulz@erlang.cs.umass.edu
dave@metro.citi.umich.edu
dave@sylvan.citi.umich.edu
breslau@santorini.usc.edu
turletti@jerry.inria.fr
dabbous@mars.inria.fr
bolot@pax.inria.fr
Phil Wood (LANL)
Steven McCanne (LBL)

Sally Floyd (LBL)
Van Jacobson (LBL)
Milo Medin (NASA Science Internet)
A. Lee Wade (NASA Science Internet)
Dan McKernan (NASA Science Internet)
Robert Michael Gutierrez (NSI Net Ops)
Warren Van Camp (NASA Science Internet)
feinler@wonderland.arc.nasa.gov
Bill Fink (NASA Goddard)
Dave Hayes (JPL/NASA) Pasadena CA USA
root@fddi_tx.ksc.nasa.gov
Shingo (uec@Japan)
asaba@turing.isr.recruit.co.jp
John Shirron (NRL, DC)
hoffman@tesuji.cmf.nrl.navy.mil
Ron Broersma (NOSC, San Diego)
yhc@happy.concert.net
Vikas @ JvNCnet
Testing
ops@noc.near.net
John Wieronski - OARnet
Kannan (OARnet)
Daniel Karrenberg (RIPE NCC)
dfk@rijp.ripe.net
Jack Jansen (CWI, Amsterdam, Holland)
lon@havik.cwi.nl
Jack Jansen@schelvis.cwi.nl
a38@frodo.nikhef.nl
Rob Blokzijl (RIPE Europe)
Geir Pedersen, Univ of Oslo, Norway
Maryann et al (MITRE, VA)
Chris Perry (MITRE, VA)
krej@julium.electrum.kth.se
perl@clotho.tds.kth.se
larsh@cupido.tds.kth.se
Patrik Ernberg@sics.se
steve@garuda.sics.se
Bengt Ahlgren, SICS
Tommy Wallo (SICS)
cc@nancarrow.sics.se
SICS Multimedia lab
Steve, SICS Sweden
Anders Klemets, SICS, Sweden
George Savvides@BT Labs, UK
Stuart @ BT Labs, UK 32kbit/s

Andie Ness (CSTR, Edinburgh)
Graeme Wood (EUCS, Univ. of Edinburgh)
Mark Handley, UC London
Ian Wakeman (UCL)
titcombe@laphroaig.cs.ucl.ac.uk
Saleem Bhatti (UCL)
sbaydere@magicflute.cs.ucl.ac.uk
nismail@mercedes.cs.ucl.ac.uk
Shaw Chuang, UCL
par@occasion.cs.ucl.ac.uk
gjoly@porridge.cs.ucl.ac.uk
Jon Crowcroft@UC London
bkumar@tamdhu.cs.ucl.ac.uk
gjoly@tomato.cs.ucl.ac.uk