Contents

Director's Message

The 26th meeting of the Internet Engineering Task Force was held in Columbus, Ohio from March 29 - April 2, 1993. The meeting was hosted by OARnet and The Ohio State University, and our thanks and appreciation go out to Kannan Varadhan and Henry Clark, and to all the others that helped with the social event and the terminal room.

As I write this message, I find myself thinking of 1993 and the changes and challenges in front of us, individually and as a group. Indeed, some changes have already or are about to occur: new members of the IESG and IAB, chosen for the first time by a selection committee; a terminal room at an IETF meeting which included working demonstrations of next generation of IP proposals, not to mention a relabeling of the effort to IP-The Next Generation (later shortened to IP-Next Generation or IPng); and this year the IETF will meet for the first time in Europe... in fact, the first time outside of North America.

The year began with the efforts of the Selection Committee whose three month mission was to implement the recommendations proposed by the Poised Working Group and adopted by the IETF at the Washington meeting. This committee, chaired by Jeff Case, launched an effort to identify and nominate new members for the IESG and IAB. Armed with only a general guideline, they went through an extensive period of identifying potential members, conducting interviews, and producing a list of recommendations by the time we met in Columbus... all this in spite of a late start and exacerbated by snow storms that delayed critical meetings and needed conversations. And in typical IETF fashion, late night meetings were held wherever space could be found in order to produce a slate of nominees and obtain the approvals in time for the plenary meeting Thursday night.

The audio-video multicasting effort of Steve Deering and Steve Casner (and others) is becoming an integral part of IETF meetings. The InterNIC launched services during the week, and a presentation was given to the IETF by representatives from the three organizations comprising the InterNIC. Part of the presentation included remote participants who spoke to the IETF via the audio-video link.

A visit to the terminal room at the Columbus Hyatt found not only a number of devices available for attendee use, but working demonstrations of PIP, SIP, and TUBA. Another new event was the first Working Group Chair Workshop conducted by Dave Crocker. Though the workshop for working group Chairs was announced relatively late, there were approximately 30 attendees at the two 8:00 a.m. sessions.

As for new challenges, the sanity of the IETF Secretariat and the patience of the IETF was challenged when the over-zealous hotel staff stored and compressed our attendee packets in a trash compactor, not to mention the other groups meeting (vocal clubs, drill team competitions, etc.). As attendance at IETF meetings remains high, meetings are being held at larger hotels and we are no longer able to "take over" the facility. This is especially true when we use hotels associated with or connected to large conference facilities.

IP - The Next Generation

The IETF meeting began Monday morning with technical presentations on four of the alternatives under consideration to address the problems of growth within the Internet. These presentations were to provide status updates and information on what had changed since the D.C. IETF meeting. There was also an increase in the interaction between the presenters and the IETF, particularly in the number of questions from the floor. The IPng effort is still one of the primary challenges facing the IETF (and the Internet) during 1993, and will undoubtably be the focus of attention throughout the year. In fact, preparations have already begun to put together the third joint presentation of the IPng candidates at the next IETF meeting in Amsterdam, the theme of which is to address what changes are required to implement each of the solutions.

In addition to the technical presentations, subsequent working group meetings were held during the week. This time, though, there were more than just working group meetings. Three of the proposed groups ran demonstrations of their solutions in the terminal room, and made informal presentations while displaying what had been accomplished since the D.C. IETF meeting. Special thanks go out to Kannan and his group for being able to accommodate the demonstrations in addition to providing devices for general IETF network access.

Meeting Statistics

As has been the trend for the past year, each meeting boasts a higher number of pre-registered attendees than the previous meeting.

It's not a fluke! After the jump in IETF meeting attendance for San Diego and Cambridge, there was some doubt as to whether this level of attendance would hold for the Washington, D.C. meeting (it did). Many IETFers were convinced that there would be a significant decrease in attendance at the Columbus meeting... not so! The number of actual attendees was slightly higher at the Columbus IETF meeting, increasing from 634 (actual D.C. attendance) to 638. There can be no doubt. Attendance at IETF meetings is consistently over 600. More and more organizations are sending people to these meetings, and for good reason; the work and activities of the IETF are appearing in the trade press more often and there is a great deal of attention being focused on the Internet and internetworking.

The number of first time attendees remained above the 200 mark. There were 214 first time attendees at the Columbus meeting, down slightly from the 225 first timers at the Washington meeting. Over 150 people showed up for the Newcomers' Orientation Sunday afternoon.

Future Meetings

The next plenary meeting of the IETF will be held in Amsterdam from July 12-16, 1993, co-hosted by RARE and SURFnet. This is the first meeting of the IETF to be held outside of North America. Please note that the meeting fee for Amsterdam will be $200.00 U.S. dollars.

The November IETF meeting will be in Houston, Texas the first week of November (November 1-5, 1993). This meeting is being hosted by SESQUINET and Rice University.

The IETF Secretariat is busy working on the 1994 meetings, and details will be announced as they firm up. At this time, it looks like the first 1994 meeting will be held in Seattle, Washington the last week of March, followed by Toronto in July. Note that information on future IETF meetings can be always be found in the file 0mtg-sites.txt which is located on the IETF shadow directories.

Stephen J. Coya
Executive Director, IETF

IETF Progress Report

The IESG and IETF have been very active since the Washington, D.C. IETF meeting last November; over 125 Internet-Draft actions, 33 IESG Protocol Actions, and over 50 RFCs.

Between the IETF meetings in Washington, D.C. and Columbus Ohio, there were five new working groups created:

1. Minimal OSI Upper-Layers (thinosi)
2. Network Training Materials (trainmat)
3. Source Demand Routing (sdr)
4. Simple Internet Protocol (sip)
5. Integrated Directory Services (ids)

and nine working groups that were concluded:

1. IP over FDDI (fddi)
2. Multi-Media Bridging (mmb)
3. Internet Mail Extensions (smtpext)
4. Internet User Glossary (userglos)
5. Directory Information Services Infrastructure (disi)
6. Ethernet MIB (ethermib)
7. DS1/DS3 MIB (trunkmib)
8. SNMP over a Multi-protocol Internet (mpsnmp)
9. TCP Client Identity Protocol (ident)

Additionally, there were 53 RFCs published since the D.C. IETF meeting in November, 1992:

RFC	Status	Title
RFC1383	I	An Experiment in DNS Based IP Routing
RFC1384	I	Naming Guidelines for Directory Pilots
RFC1386	I	The US Domain
RFC1387	I	RIP Version 2 Protocol Analysis
RFC1388	PS	RIP Version 2 Carrying Additional Information
RFC1389	PS	RIP Version 2 MIB Extension
RFC1390	S	Transmission of IP and ARP over FDDI Networks
RFC1391	I	The Tao of IETF: A Guide for New Attendees of the Internet Engineering Task Force
RFC1392	I	Internet Users' Glossary
RFC1393	E	Traceroute Using an IP Option
RFC1394	I	Relationship of Telex Answerback Codes to Internet Domains

RFC1395	I	BOOTP Vendor Information Extensions
RFC1396	I	The Process for Organization of Internet Standards Working Group (POISED)
RFC1397	PS	Default Route Advertisement In BGP2 And BGP3 Versions Of The Border Gateway Protocol
RFC1398	DS	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC1400	I	Transition and Modernization of the Internet Registration Service
RFC1401	I	Correspondence between the IAB and DISA on the use of DNS throughout the Internet
RFC1402	I	There's Gold in them thar Networks! Searching for Treasure in all the Wrong Places
RFC1403	PS	BGP OSPF Interaction
RFC1404	I	A Model for Common Operational Statistics
RFC1405	E	Mapping between X.400(1984/1988)and Mail-11(DECnet mail)
RFC1406	PS	Definitions of Managed Objects for the DS1 and E1 Interface Types
RFC1407	PS	Definitions of Managed Objects for the DS3/E3 Interface Type
RFC1408	PS	Telnet Environment Option
RFC1409	E	Telnet Authentication Option
RFC1410	S	IAB OFFICIAL PROTOCOL STANDARDS
RFC1411	E	Telnet Authentication: Kerberos Version 4
RFC1412	E	Telnet Authentication : SPX
RFC1413	PS	Identification Server
RFC1414	PS	Ident MIB
RFC1415	PS	FTP-FTAM Gateway Specification
RFC1416	E	Telnet Authentication Option
RFC1417	I	NADF Standing Documents: A Brief Overview
RFC1418	PS	SNMP over OSI
RFC1419	PS	SNMP over AppleTalk
RFC1420	PS	SNMP over IPX
RFC1421	PS	Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures
RFC1422	PS	Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management
RFC1423	PS	Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers
RFC1424	PS	Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services
RFC1425	PS	SMTP Service Extensions

RFC1426	PS	SMTP Service Extension for 8bit-MIMEtransport
RFC1427	PS	SMTP Service Extension for Message Size Declaration
RFC1428	I	Transition of Internet Mail from Just-Send-8 to 8Bit-SMTP/MIME
RFC1429	I	Listserv Distribute Protocol
RFC1430	I	A Strategic Plan for Deploying an Internet X.500 Directory Service
RFC1431	I	DUA Metrics
RFC1432	I	Recent Internet Books
RFC1433	E	Directed ARP
RFC1434	I	Data Link Switching: Switch-to-Switch Protocol
RFC1435	I	IESG Advice from Experience with Path MTU Discovery
RFC1436	I	The Internet Gopher Protocol(a distributed document search and retrieval protocol)
RFC1439	I	The Uniqueness of Unique Identifiers

Agenda of the Twenty-Sixth IETF

(March 29 - April 2, 1993)

MONDAY, March 29, 1993

8:00-9:00 am	IETF Registration and Continental Breakfast
9:00-9:30 am	Introductions
9:30-12:00 noon	Technical Presentations
	• "Next Generation of IP"
Breaks	Coffee available throughout morning.
1:30-3:30 pm	Afternoon Sessions I

APP	OSI Directory Services WG (osids) (Steve Kille/ISODE)
INT	IP over ATM WG (atm) (Bob Hinden/Sun)
MGT	SNA Systems Management BOF (snamib) (Baktha Murali/DEC)
OPS	Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)
RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)*
RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
RTG	OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)*
SEC	Security Area Advisory Group (saag) (Steve Crocker/TIS)
USV	Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/CNIDR)

3:30-4:00 pm	Break (Refreshments provided)
4:00-6:00 pm	Afternoon Sessions II

APP	OSI Directory Services WG (osids) (Steve Kille/ISODE)
INT	IP Address Encapsulation WG (ipae) (Dave Crocker/Silicon Graphics)
INT	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
INT	TCP/UDP over CLNP-addressed Networks WG (tuba) (Peter Ford/LANL and Mark Knopper/Merit)

4:00-6:00 pm	Monday, March 29, 1993 - Afternoon Sessions II (cont'd.)	
	OPS	Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)
	RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
	SEC	Common Authentication Technology WG (cat) (John Linn)
	USV	WHOIS and Network Information Lookup Service WG (wnils) (Joan Gargano/UCDavis)
7:30-10:00 pm	Evening Sessions	
	APP	Internet Message Extension WG (822ext) (Greg Vaudreuil/CNRI)
	MGT	SNA Systems Management BOF (snamib) (Baktha Murali/DEC)
	OPS	BGP Deployment and Application WG (bgpdepl) (Matt Mathis/PSC)
	RTG	IPv7 Addressing BOF (bigaddr) (Paul Tsuchiya/Bellcore)
	RTG	Virtual Circuit Routing BOF (vcrout) (Rob Coltun/Consultant and Marco Sosa/Bellcore)
	USV	Integration of Internet Information Resources WG (iiir) (Chris Weider/Merit)

* BGP and IPIDRP will be meeting in joint session.

TUESDAY, March 30, 1993

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	IETF Technical Presentations	
	• "An Architecture for Resource Management Networks" (David Clark/MIT)	
9:30-12:00 noon	Morning Sessions	
	APP	Conferencing Control BOF (confctrl) (Eve Schooler/ISI)
	APP	Office Document Architecture WG (oda) (Peter Kirstein/UCL)
	APP	TELNET WG (telnet) (Steve Alexander/Lachman Technology)
	INT	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
	INT	Simple Internet Protocol WG (sip) (Steve Deering/Xerox PARC and Christian Huitema/INRIA)
	OPS	Network OSI Operations WG (noop) (Sue Hares/Merit and Cathy Wittbrodt/BARRNet)
	RTG	Open Shortest Path First IGP WG (ospf) (John Moy/Proteon and Mike Petry/UMD)
	USV	Integrated Directory Services WG (ids) (Tim Howes/UMich and Chris Weider/Merit)
	USV	User Services WG (uswg) (Joyce K. Reynolds/ISI)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	MHS-DS WG (mhsds) (Kevin Jordan/CDS and Harald Alvestrand/SINTEF DELAB)
	INT	IP over ATM WG (atm) (Bob Hinden/Sun)
	INT	P. Internet Protocol WG (pip) (Paul Tsuchiya/Bellcore)
	MGT	SNA Systems Management BOF (snamib) (Baktha Murali/DEC)
	SEC	Network Access Server Requirements WG (nasreq) (Allan Rubens/Merit and John Vollbrecht/Merit)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)

1:30-3:30 pm	Tuesday, March 30, 1993 - Afternoon Sessions I (cont'd.)	
	USV	Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/CNIDR)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	APP	MIME-MHS Interworking WG (mimemhs) (Steve Thompson/Soft-Switch)
	INT	Dynamic Host Configuration WG (dhc) (Ralph Droms/Bucknell)
	INT	IP over AppleTalk WG (appleip) (John Veizades/Apple)
	INT	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)*
	INT	P. Internet Protocol WG (pip) (Paul Tsuchiya/Bellcore)
	INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant) *
	RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)*
	RTG	OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)*
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
	TSV	Domain Name System WG (dns) (Rob Austein/Epilogue)
	USV	User Documents WG (userdoc2) (Ellen Hoffman/Merit and Lenore Jackson/NASA)
7:30-10:00 pm	Tuesday, March 30, 1993 - Evening Sessions	
	APP	TELNET WG (telnet) (Steve Alexander/Lachman Technology)
	INT	Net Support for QOS and Real-Time Traffic BOF (rtqos) (David Clark/MIT)
	MGT	Frame Relay Network MIB BOF (frnetmib) (Tracy Cox/Bellcore)
	MGT	Mail and Directory Management BOF (madman) (Steve Kille/ISODE)
	MGT	Modem Management BOF (modemmgt) (Mark Lewis/Telebit)
	OPS	Generic Internet Service Specification BOF (giss) (Daniel Karrenberg/RIPE)
	SEC	Common Authentication Technology WG (cat) (John Linn)

7:30-10:00 pm — Tuesday, March 30, 1993 - Evening Sessions (cont'd.)

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

* BGP and IPIDRP will be meeting in joint session.
* IPLPDN and PPPEXT will be meeting in joint session.

WEDNESDAY, March 31, 1993

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	• Amsterdam IETF Update (Erik Huizer/SURFnet)	
9:30-12:00 noon	Morning Sessions	
	APP	Internet Message Extension WG (822ext) (Greg Vaudreuil/CNRI)
	APP	X.400 Operations WG (x400ops) (Allan Cargille/UWisc)
	INT	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)*
	INT	Dynamic Host Configuration WG (dhc) (Ralph Droms/Bucknell)
	INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant) *
	MGT	IEEE 802.3 Hub MIB WG (hubmib) (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
	OPS	MBONE Engineering and Operations BOF (mbone) (Matt Mathis/PSC)
	RTG	Source Demand Routing BOF (sdr) (Deborah Estrin/USC and Tony Li/cisco)
	SEC	Internet Protocol Security Protocol WG (ipsec) (Al Hoover/ANS and Paul Lambert/Motorola)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
	USV	Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/CNIDR)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	Conferencing Control BOF (confctrl) (Eve Schooler/ISI)
	APP	Interactive Mail Access Protocol BOF (imap) (Terry Gray/UWash)
	APP	X.400 Operations WG (x400ops) (Allan Cargille/UWisc)
	INT	IP over ATM WG (atm) (Bob Hinden/Sun)

1:30-3:30 pm	Wednesday, March 31, 1993 - Afternoon Sessions I (cont'd.)	
	INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant)
	INT	Simple Internet Protocol WG (sip) (Steve Deering/Xerox PARC and Christian Huitema/INRIA)
	MGT	IEEE 802.3 Hub MIB WG (hubmib) (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
	SEC	Commercial Internet Protocol Security Option WG (cipso) (Ron Sharp/AT&T)
	USV	Networked Information Retrieval WG (nir) (Jill Foster/UNewcastle-Upon-Tyne and George Brett/MCNC)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	MGT	Token Ring Remote Monitoring WG (trmon) (Steve Waldbusser/CMU)
	OPS	Operational Area Directorate (orad) (Phill Gross/ANS and Bernhard Stockman/SUNET)
	RTG	IP Routing for Wireless/Mobile Hosts WG (mobileip) (Steve Deering/Xerox PARC)
	RTG	Virtual Circuit Routing BOF (vcrout) (Rob Coltun/Consultant and Marco Sosa/Bellcore)
	SEC	Authorization and Access Control BOF (aac) (Cliff Neuman/ISI)
	TSV	Service Location Protocol WG (svrloc) (John Veizades/Apple and Scott Kaplan/FTP)
	USV	GOPHER BOF (gopher) (Jim Fullton/CNIDR and Joyce K. Reynolds/ISI)
7:30-10:00 pm	Wednesday, March 31, 1993 - Evening Session	
	USV	Low Cost IP Hardware Wish List BOF (loip) (Laura Breeden/FARNET, Inc.)

* IPLPDN and PPPEXT will be meeting in joint session.

THURSDAY, April 1, 1993

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	•	"Internet Talk Radio" (Carl Malamud)
9:30-12:00 noon	Morning Sessions	
	APP	Minimal OSI Upper-Layers WG (thinosi) (Peter Furniss/Consultant)
	MGT	Managing ATM with SNMP BOF (atmmib) (Kaj Tesink/Bellcore)
	MGT	Uninterruptible Power Supply WG (upsmib) (Jeff Case/UTenn)
	OPS	Network Joint Management WG (njm) (Gene Hastings/PSC)*
	OPS	Network Status Reports WG (netstat) (Gene Hastings/PSC)*
	RTG	IP Routing for Wireless/Mobile Hosts WG (mobileip) (Steve Deering/Xerox PARC)
	RTG	ISIS for IP Internets WG (isis) (Ross Callon/Wellfleet and Chris Gunner/DEC)
	SEC	Privacy-Enhanced Electronic Mail WG (pem) (Steve Kent/BBN)
	USV	Network Training Materials WG (trainmat) (Ellen Hoffman/Merit and Jill Foster/UNewcastle-Upon-Tyne)
Breaks	Coffee available throughout the morning.	
1:30-3:30 pm	Afternoon Sessions I	
	MGT	Chassis MIB WG (chassis) (Jeff Case/UTenn and Bob Stewart/Xyplex)
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	INT	SNA Peer-to-Peer Networking BOF (snapper) (Wayne Clark/cisco)
	INT	TCP/UDP over CLNP-addressed Networks WG (tuba) (Peter Ford/LANL and Mark Knopper/Merit)

1:30-3:30 pm	Thursday, April 1, 1993 - Afternoon Sessions I (cont'd.)	
	RTG	Multicast Extensions to OSPF WG (mospf) (Steve Deering/Xerox PARC)
	SEC	Security Area Advisory Group (saag) (Steve Crocker/TIS)
	TSV	Service Location Protocol WG (svrloc) (John Veizades/Apple and Scott Kaplan/FTP)
	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 — Technical Presentations

- "Introducing the INTERNIC"
  - General Atomics/CERFnet (Susan Calcari)
  - Network Solutions (Scott Williamson)
  - AT&T (Rick Huber)
- "Qualcomm Wireless" (Phil Karn/Qualcomm)

7:30-10:00 pm — Open Plenary and IESG

* NJM and NETSTAT will be meeting in joint session.

FRIDAY, April 2, 1993

8:30-9:00 am	Continental Breakfast	
9:00-12:00 noon	Morning Sessions	
	INT	Inter-Domain Multicast Routing BOF (idmr) (Tony Ballardie/UCL)
	MGT	FDDI MIB WG (fddimib) (Jeff Case/UTenn)
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	OPS	Network Joint Management WG (njm) (Gene Hastings/PSC)*
	OPS	Network Status Reports WG (netstat) (Gene Hastings/PSC)*

* NJM and NETSTAT will be meeting in joint session.

Key to Abbreviations

APP	Applications	Russ Hobby/UCDavis and Erik Huizer/SURFnet
INT	Internet	Philip Almquist, Stev Knowles/FTP Software and Dave Piscitello/Bellcore
MGT	Network Management	TBD
OPS	Operational Requirements	Bernhard Stockman/SUNET and Phill Gross/ANS
RTG	Routing	Bob Hinden/Sun
SEC	Security	Steve Crocker/TIS
TSV	Transport and Services	Dave Borman/Cray Research
USV	User Services	Joyce K. Reynolds/ISI

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 one or more Area Director who has primary responsibility for that one area of IETF activity. Together with the Chair of the IETF, these 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:	Erik Huizer/SURFnet
	Brewster Kahle/WAIS
Internet:	Stev Knowles/FTP Software
	Dave Piscitello/Bellcore
Network Management:	Marshall Rose/DBC
Operational Requirements:	Scott Bradner/Harvard
Routing:	Robert Hinden/Sun
Security:	Steve Crocker/TIS
Transport:	Allison Mankin/NRL
Service Applications:	Dave Crocker/SGI
User Services:	Joyce K. Reynolds/ISI
Standards Management:	A. Lyman Chapin/BBN

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 Meeting Coordinator:	Megan Davies Walnut
IETF Meeting Registrar:	Debra Legare
Administrative Support:	Cynthia Clark
	Britt Jackman

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 ds.internic.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 an IETF general discussion list and an IETF announcement 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 IETF announcement mailing list. For general inquiries about the IETF, requests should be sent to `ietf-info@cnri.reston.va.us`. An archive of mail sent to the IETF list is available for anonymous ftp from the directory `ietf-mail-archive/ietf` on `cnri.reston.va.us`.

1.1 Future IETF Meeting Sites

Summer 1993

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

Fall 1993

Houston, Texas
SESQUINET and Rice University
Host: Bill Manning
November 1-5, 1993
Status: CONFIRMED

Spring 1994

Seattle, Washington
NorthWestNet
Host: Dan Jordt
March 28 - April 1, 1994
Status: TENTATIVE

Summer 1994

Toronto, Canada
University of Toronto
Host: Warren Jackson
July 25-29, 1994
Status: TENTATIVE

1.2 On Line IETF Information

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 and Email file servers. Procedures for retrieving the information are described below.

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. Working Group charters and minutes are in sub-directories under the working group acronym. Retrieve and view the 1wg-summary.txt file for a list of working groups and their acronyms.

FILE NAME	
0tao.txt	This file contains "A Guide for New Attendees of the Internet Engineering Task Force", RFC1391.
0mtg-agenda.txt	The current Agenda for the upcoming IETF meeting, containing scheduled working group meetings, Technical Presentations and Network Status Reports.
0mtg-at-a-glance.txt	The announcement for the upcoming IETF meeting, containing specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and meeting costs.
0mtg-rsvp.txt	A standardized RSVP form to notify the Secretariat of your plans to attend the upcoming IETF meeting.
0mtg-sites.txt	Current and future meeting dates and sites for IETF meetings.
1id-guidelines.txt	Instructions for authors of Internet-Drafts.
1ietf-description.txt	A short description of the IETF, the IESG and how to participate.
1wg-summary.txt	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.

1wg-charters.txt	A single file containing an abbreviated version of all the current working group charters.

Working groups have individual directories dedicated to their particular activities. The directories contain the charters and meeting minutes for the group.

Minutes of Birds of a Feather (BOF) sessions and area summaries of the IETF meetings are grouped into directories by meeting. The directory names are of the form YYmmm, i.e. 92mar for the reports of the March 1992 meeting. These directories do not include the minutes of the working group meetings.

When using FTP, the "cd" and "dir" commands will permit you to review what working group files are available and the specific naming scheme to use for a successful anonymous ftp request.

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 IESG 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(s) whose name and email address are listed on the first page of the respective draft.

FILE NAME	
1id-abstracts.txt	This file lists the current Internet-Drafts and their pathnames.
1id-index.txt	This file contains an abbreviated listing of Internet-Drafts. This contains only the document title, the filename and the posting date.

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

The IESG Directory

The IESG directory contains the minutes of IESG meetings and regularly updates status report on protocols in the standards track.

FILE NAME	
1protocol_actions.txt	This file contains a list of protocols currently under consideration by the IESG.

1old_standards.txt	This file contains a list of Proposed and Draft Standards eligible for advancement.

The minutes are contained in files named with the pattern:

iesg.YY-MM-DD

i.e.,

iesg.92-11-10

for the minutes of the meeting held on November 10, 1992.

FTP Access

IETF Information is available by anonymous FTP from several sites.

East Coast (US) Address: ds.internic.net (198.49.45.10)

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

Europe Address: nic.nordu.net (192.36.148.17)

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

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

To retrieve this information via FTP, establish an anonymous FTP connection, then login with username "anonymous". Use your email address as the password. When logged in, change to the directory of your choice with one of the following commands:

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

Individual files can then be retrieved using the GET command:

```
get 1wg-summary.txt
get 822ext/822ext-charter.txt
```

Email Access

Internet-Drafts are available by mail server from ds.internic.net. To retrieve a file, mail a request to mailserver@ds.internic.net with a subject of anything you want. In the body, put a command of the form:

```
send internet-drafts/1id-abstracts.txt
send ietf/1wg-summary.txt
send ietf/822ext/822ext-minutes-91jul.txt
```

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@cnri.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 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-Drafts 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 ds.internic.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@cnri.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
- Erik Huizer: huizer@surfnet.nl

Area Summary reported by Russ Hobby/UC Davis

At the end of the Columbus meeting it was announced by the IAB that Brewster Kahle from Wais, Inc. will replace Russ Hobby as a co-Director of the Applications Area.

Applications Area Directorate (APPLES)

The Applications Area Directorate met for the first time at the Columbus IETF. The Directorate will help the Area Directors on architectural matters and reviews. Members of the Directorate are appointed by the Area Directors. Nominations can be made by the Application Area working group Chairs. The Directorate can be reached at `<apples@surfnet.nl>` and currently consists of the following individuals:

- Ned Freed
- John Klensin
- Steve Kille
- Christian Huitema
- Russ Hobby

The first task of the directorate is to produce a document on an email architecture. This document will be used as a basis for discussion on this topic in the Applications Area. After the document has evolved to a state of maximum consensus, working groups will be created to focus on specific issues indicated by the architecture document.

The directorate also discussed the general problem of character sets and noted that it will be a recurring problem in many applications. The directorate will develop an initial plan for dealing with character sets in applications and start a working group to address this problem in detail.

The directorate noted the increasing difficulty for working groups to make forward progress. This appears to be due to the increasing size and interest in the IETF and the Internet in general. More people, more discussions, more time. In the future, the Applications Area desires an initial draft document be written by interested parties before a working group is formed. While the final result of the working group may look nothing like the initial document, the initial document will provide focus for discussion.

There are four working groups jointly chartered under the User Services Area and Applications Area. They are:

- Integrated Directory Services (IDS)
- Integration of Internet Information Resources (IIIR)
- Uniform Resource Identifiers (URI)
- Whois and Network Information Lookup Service (WNILS)

For a report on these Working Groups see the User Services Area Report.

There were two BOFs held that reside under the Network Management Area, but are strongly related to the Applications Area. They are:

- Mail and Directory Management (MADMAN)
- IFIP Electronic Mail Management (EMAILMGT)

For a report on these BOFs see the Network Management Area Report.

Conference Control BOF (CONFCTRL)

Now that video, audio and shared applications are starting to flow over the network, there is a need for the setup and management of conference sessions. This BOF focused on various aspects of controlling distributed network conferences. Several people related their current work, and plans were made for coordinating work through an IETF working group.

Interactive Mail Access Protocol BOF (IMAP)

The BOF discussed efforts to update and standardize IMAP. Mark Crispin has a new draft of IMAP that will be submitted as an Internet-Draft. A sample working group Charter was reviewed.

Internet Message Extensions Working Group (822EXT)

The RFC822 Message Extensions Working Group met for two sessions to review and approve the revised MIME protocol for Draft Standard. With several clarifications and with the removal of several optional features, agreed to previously on the ietf-822 mailing list, MIME was so approved.

The Working Group has completed its Charter as currently written and is ready to conclude. There is significant MIME related work which still needs to be addressed and for which new working groups should be formed. Among the work are MIME extensions such as the definition of a content-integrity-check and content-disposition body headers to add general functionality to MIME. There are expected to be a large number of new content-types defined, most of which should be developed in specific single-topic working groups.

MHS-DS Working Group (MHSDS)

The MHS-DS Working Group focused on its Long Bud pilot project at this IETF meeting. Since the last meeting, some basic infrastructure has been established for supporting X.400

routing via the Internet X.500 directory service. By the next IETF, the Group plans to expand this infrastructure and generate some tools such that it can be demonstrated that the pilot project is functional and that the directory is actually being used by some MTAs to support message routing. To achieve this goal, specific action items were assigned to Working Group members. Specifically, two important documents will be written and circulated, and specific individuals will begin implementing important software tools. The documents will clearly define the purpose of the pilot project, outline its short and long-term goals, specify its relationship to the existing Internet X.400 community, and indicate how to participate in the pilot. The tools will facilitate the integration of the pilot with the existing Internet X.400 infrastructure.

In addition to working on issues relating to the Long Bud pilot project, action items were assigned for progressing three Internet-Drafts as RFCs. In addition, one or two minor technical issues were resolved which will be reflected in the next revision of the Internet-Drafts.

MIME-MHS Interworking Working Group (MIMEMHS)

There are three draft RFCs in progress:

1. "Mapping between X.400 and RFC-822 Message Bodies"
2. "Equivalences between 1988 X.400 and RFC-822 Message Bodies"
3. "HARPOON: Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages"

The first two have been stable for some time with no outstanding issues. The third (HARPOON) had some open issues and, until now, had never been discussed at an IETF meeting.

During the meeting, the HARPOON proposal was presented, the issues were resolved, and it was agreed that all three documents would be forwarded to the IESG for approval as Proposed Standards.

Minimal OSI Upper-Layers Working Group (THINOSI)

The THINOSI Working Group met for the first time as a working group. Nearly all the time was spent reviewing the first draft of the "bytestream cookbook." Various changes were agreed upon, generally applying a principle of keeping things simple (and thin) for this first case, but ensuring that interworking with "full" OSI implementations would be feasible. It will be highly desirable to achieve alignment with the "minimal OSI" profile being developed in OIW and EWOS. Identifying the range of applications to be supported is central to achieving this alignment - this should include at least DAP and X.400 P7 if at all possible.

Office Document Architecture Working Group (ODA)

Over the course of 1992 an international profile FOD26 was being approved. An industrial consortium was preparing an ODA toolkit which will become available in the second quarter of 1993. Pending the availability of this toolkit and the new profile, there has been little availability of new ODA implementations, though this will change during the third quarter of 1993.

The Working Group had previously expressed interest only in piloting with real products. In view of their non-availability at present, there was little interest in the Group.

It is recommended that the Working Group conclude. If there is further interest when products become available it can be revived, though this is unlikely to happen before November 1993.

OSI Directory Services Working Group (OSIDS)

- The Charter was discussed and several work items were defined.
- There is strong consensus on the need for continuation of this Group.
- Volunteers for editing papers are hard to find.
- The schema management issue is still not resolved. This remains a major worry.
- A new approach to presenting quality of data in the Directory was discussed. It will be put on paper and aligned with earlier ideas of the Group.
- Representation of registration, IP-addressing and Network Information was discussed. A series of Internet-Drafts will be produced on this issue.
- Representation of documents and related information in the Directory was discussed based on four draft inputs.

TELNET Working Group (TELNET)

The Working Group continued work on the Environment Option, Authentication and Encryption.

HP's Telnet MPX proposal for session multiplexing was discussed. Most people were impressed with the results but felt that, in general, session multiplexing did not belong in the Telnet layer. Perhaps this should be addressed as a TCP extension. In the meantime, the Working Group suggested that HP submit the protocol to be an Experimental Protocol.

There was enthusiastic discussion by a group of people who want to work on improving TN3270 to better match the current SNA environment. The TELNET Working Group felt that the TN3270 work would be outside the scope of their Group and work should be done within a separate working group.

X.400 Operations Working Group (X400OPS)

- Finalized Documents:
  - "Operational Requirements for X.400 Management Domains in the GO-MHS Community" as an Informational RFC.
  - "Routing coordination for X.400 MHS services within a multi protocol multi network environment Table Format V3 for static routing" as an Experimental Protocol.
  - "Evaluation of ADMDs and Integration aspects with respect to the R&D messaging community" as an Informational RFC.
  - "Assertion of C=US; A=IMX" as a Proposed Standard.
  - "X.400 use of extended character sets" as a Proposed Standard.

- Work Left To Do:
  - Automatic email distribution of tables.
  - X.400 - RFC822 mapping authorities.

CURRENT MEETING REPORT

Reported by Eve Schooler/Information Sciences Institute

Minutes of the Conferencing Control BOF (CONFCTRL)

These Minutes were prepared by Eve Schooler from notes provided by Abel Weinrib of Bellcore and Deborah Estrin of USC/ISI.

Introduction and Presentations

Two CONFCTRL sessions were held at the Columbus IETF. The first meeting was used to provide an overview of Conference Control efforts both within and outside of the IETF. Inside the IETF, the CONFCTRL Group was spawned by the Remote Conferencing Architecture BOF (REMCONF). Outside the IETF, interest in conference control, sometimes referred to as connection management, has been ongoing for some time. Thus far, the CONFCTRL mailing list has collected a sizable bibliography containing references to many of the early and ongoing research projects in this area.

Most of the first session was used for presentations on different CONFCTRL schemes (see slides). The intent of the presentations was to flesh out design assumptions, tradeoffs, complexity, scalability, etc. The systems were classified according to several parameters: whether they (1) concentrate more on groupware conferencing (shared editors, whiteboards) than on real-time audio/video conferencing, (2) provide session control of packet-based real-time media versus analog real-time media, (3) rely on centralized versus distributed session management, and/or (4) observe loose versus tight session control.

Synthesis of CONFCTRL approaches

The second session was used to identify pervasive CONFCTRL themes, and to question the applicability of the various solutions to the Internet. The main objective was to narrow the scope of the problem en route to the design of a generic CONFCTRL protocol. Observations were culled not only from the presentations at the IETF but also from templates that were filled out prior to the meeting. The templates included Dave Lewis' write up of the UCL PREPARE project, a description of the ZAPT project by Joe Touch, a contribution from Jack Jansen about the Meeting project, and Fengmin Gong's template on the MCNC CONCERT Video Network Migration effort.

Of particular interest were implementors' comments about the aspects of their approaches which were hard, easy, or warranted change. Except for a lone comment about the ease of implementation of floor control, there were several recurrent themes regarding implementation difficulties:

- It is difficult to design a CONFCTRL protocol that balances simplicity with a high degree of semantic flexibility, e.g., Jack Jansen concluded that different conferencing styles require entirely separate CONFCTRL protocols.

- A distributed model comes with distributed system complexities:
  - Support for causality of multiway message exchanges.
  - Recovery from temporary network failures.
  - Propagation of consistent state information.

  The solutions proved to be cumbersome, unexpectedly hard and often times "tricky".

- The underlying transport (that carries session control information) comes at a price, e.g., the overhead of one RPC implementation led the PREPARE project to shift to a different, lighter implementation.

- There is room for improved media integration, e.g., asymmetric flows are difficult to characterize at setup, there is a need for more powerful control over presentation of media streams.

Most experimental systems either are or began as LAN-based conferencing systems. However, it is clear that many, if not all, are aiming for WAN operation. Although the tools that currently populate the MBONE rely on loose-style session control, in the past most experimentation has taken place with tightly controlled session models – though this is clearly changing. The Group speculated that the predominance of tight-control systems may be a function of the interest in supporting "coordinated" telecollaborations, which are readily modeled using a tight-control framework, whereas the emergence of loose-control systems may be a reflection of the relative ease with which they are implemented.

Systems were clearly differentiated in their approaches to interconnectivity among participants, both for session and for media topology. In certain cases, symmetry exists for N-way communication capabilities, while in other cases conferees are asymmetrically interconnected, relying on an initiator, moderator, filter/reflector or a privileged set of designees to coordinate communication on behalf of others. Explicit versus implicit communication is another distinguishing feature; this relates to whether or not the session has policies attached to it, such as who dictates membership rules, the extent to which session information is disseminated or if participant information is meant to be kept globally coherent. Finally, it was observed that the decision to model the system in a centralized or distributed fashion influenced the degree of messaging synchrony and causality.

Group Scope, Framework and Functional Taxonomy

There was rough consensus on the definition of conference control as the management and coordination of multiple sessions and their multiple users in multiple media. It was also agreed that the focus of the Group is to design a "session layer" protocol to perform these functions. However, the Group debated the utility of designing a "teleconferencing" session protocol specifically for the coordination of users' "media" versus designing a group negotiation protocol that is extensible to act as a conduit for media details.

The Group recognizes that it cannot set out to support all conferencing scenarios. However, it proposes to support one loose style protocol (a la Xerox PARC's nv, INRIA's ivs, BBN's dvc, LBL's vat, UMass' nevot) and one tight style protocol (for negotiated and potentially private sessions). How loose and how tight? To answer this, the list of conversation styles must be mapped (from the last IETF Minutes) into their underlying CONFCTRL session protocols.

As an example of how a tight-control approach to session management might integrate with already existing MBONE tools, an X-based version of ISI's MMCC conference control tool was demonstrated at the IETF. MMCC was used to explicitly invite a specific set of participants (versus having a wide-open session), to distribute multicast addresses and a shared encryption key among those participants, and to initiate as well as tear down sessions comprised of nv, vat and/or BBN's newly released PictureWindow.

Although it was emphasized that the goal of the Group is to design a session protocol, the Group conceded that there is a need for a common framework within which it can talk about conferencing control. The framework that arose from discussion, looked as follows:

```
          User A                                           User B

        +-------------+                                 +-------------+
        |             |                                 |             |
        | Application |                                 | Application |
        |             |                                 |             |
        +------+------+                                 +------+------+
               |                                               |
        +------+------+                                 +------+------+
        |             |                                 |             |
        |   Session   |<------------------------------->|   Session   |
        |             |       "Session Protocol"        |             |
        +---+--+--+---+                                 +---+--+--+---+
           /   |   \                                       /   |   \
          /   ...   \                                     /   ...   \
    +-------+     +-------+                         +-------+    +-------+
    | Media | ... | Media |<----------------------->| Media |    | Media |
    | Agent |     | Agent |     "Media Stream"      | Agent |    | Agent |
    +-------+     +-------+                         +-------+    +-------+
```

The premise is that the session protocol would be distributed in nature, and would accommodate multiple user sessions (even though the diagram depicts only two conferees). There is a firm separation between the session protocol and media transport protocols. Thus, it is immaterial whether the media transport is packet-based or analog. Generic session state would include membership and policy information. Application-domain specific state might include media interconnectivity (topology) and media configuration (data encodings, rates). Although needing further refinment, the list of session functionality provided to the end systems and reflected in the session protocol would encompass:

- Create/Destroy Session
- Add/Delete Member
- Set Policy
  - Who may join
  - Who may invite
  - Who may set policies
  - Etc.
- Add/Change Application-Domain Specific State
  - Media interconnectivity
  - Media configuration
- Floor Control?
- Prescheduling?

Polling the interest of the BOF participants, it was found that 75% were interested in solving the session protocol problem, 40% also would be interested in defining or standardizing the media-agent-to-session-entity interface, and 30% were interested in configuration management issues.

Terminology

It became evident that there are no set definitions for terms such as conference, connection, session, media agents, etc. Many of the systems presented during the BOF and described in the templates used these terms differently. Thus, a CONFCTRL terminology reference guide needs to be developed.

The Group had been interchanging the phrases session control, session management, connection control and connection management, but later agreed that "connection" is too ambiguous since it is used at any number of levels in the protocol stack. Connection was replaced by the term "session", and was broadly defined as an association of members for control purposes. However, it was later argued that session looks too much like an OSI term. The term "conference" was also felt to be too application specific. Therefore, the Group is open to suggestions for a better name.

It was suggested (although not entirely resolved) that "media agents" handle the media specifics associated with a session. "Media" could be considered any data streams that involve communication. It was also suggested that floor control is deemed the responsibility of a media agent when it concerns a single media agent, but the responsibility of the session entity when it requires coordination across different media agents (e.g., video to follow audio).

The Group also differentiated between two meanings of configuration; the static end-system description, including hardware and software capabilities, and the per-session description.

Liaisons

The CONFCTRL Group is committed to tracking the progress of related efforts, both within and outside the IETF. An important IETF linkage is to leverage off ongoing work in the Audio/Video Transport Working Group (AVT), which is nearing completion of the Real-time Transport Protocol (RTP) specification. During the first two AVT sessions, there was considerable discussion about RTCP, the control protocol associated with RTP. Certain functions in RTCP were felt to violate "layering"; they do not belong in the transport layer, but would live comfortably within the session level, e.g., text strings of session participants. The Group will need to follow closely the outcome of these developments, especially if certain services are assumed to percolate into the session layer.

The MBONE is another strategic testing ground for a CONFCTRL solution, although its use should not preclude use of these ideas elsewhere, nor should these ideas be tailored specifically to the MBONE. By mentioning MBONE it is really meant that the Group expects, in the long term, to have access to networks that support multicast and in the longer term to support real-time services. The general Internet should suffice for now.

Individuals who volunteered to track developments in related areas include:

Ruth Lang	Directory Services
Hans Eriksson	Multicast Developments
Fengmin Gong	Resource Management/QoS
Steve Casner	Audio/Video Transport
Eve Schooler	Audio/Video Transport
Paul Lambert	Security
Stuart Stubblebine	Security
Yee-Hsiang Chang	ATM
Peter Kirstein	MIBs

Action Items

- Make CONFCTRL bibliography available.
- Documentation:
  - Terminology reference guide.
  - Refinement of functional taxonomy.

- Turn Minutes into issues/framework document.
- Mapping of conversation styles into session protocols.
- Collect suggestions for a Group name change.

Attendees

Lou Berger	lberger@bbn.com
Monroe Bridges	monroe@cup.hp.com
Al Broscius	broscius@bellcore.com
Randy Butler	rbutler@ncsa.uiuc.edu
Yee-Hsiang Chang	yhc@hpl.hp.com
Brian Coan	coan@faline.bellcore.com
Richard Cogger	R.Cogger@cornell.edu
Simon Coppins	coppins@arch.adelaide.edu.au
Dave Cullerot	cullerot@ctron.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Hans Eriksson	hans@sics.se
Deborah Estrin	estrin@isi.edu
Francois Fluckiger	fluckiger@vxcern.cern.ch
Jerry Friesen	jafries@sandia.llnl.gov
Fengmin Gong	gong@concert.net
Kenneth Goodwin	goodwin@a.psc.edu
Mark Green	markg@apple.com
Russ Hobby	rdhobby@ucdavis.edu
Don Hoffman	hoffman@eng.sun.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Michael Khalandovsky	mlk@ftp.com
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@binky.arc.nasa.gov
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Giri Kuthethoor	giri@ms.uky.edu
Paul Lambert	paul_lambert@email.mot.com
Ruth Lang	rlang@nisc.sri.com
Patrick Leung	patrickl@eicon.qc.ca
Allison Mankin	mankin@cmf.nrl.navy.mil
Donald Merritt	Don@brl.mil
Paul Milazzo	milazzo@bbn.com
Robert Mines	rfm@sandia.llnl.gov
Joseph Pang	pang@bodega.stanford.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
John Penners	jpenners@advtech.uswest.com
Bala Rajagopalan	braja@qsun.att.com

Michael Safly	saf@tank1.msfc.nasa.gov
Eve Schooler	schooler@isi.edu
Michael St. Johns	stjohns@darpa.mil
Stuart Stubblebine	stubblebine@isi.edu
Sally Tarquinio	sallyt@gateway.mitre.org
Claudio Topolcic	topolcic@cnri.reston.va.us
Mario Vecchi	mpv@thumper.bellcore.com
Abel Weinrib	abel@bellcore.com
John Wroclawski	jtw@lcs.mit.edu
Yow-Wei Yao	yao@chang.austin.ibm.com

Presentations of Existing Confctrl Schemes

- Shared Workspace
  - CECED
- Central Connection Management:
  - Analog Media: Touring Machine
  - Real-time Packet Media: CoDesk, CoCo
- Distributed Connection Management
  - Tight Control: VideoTeam, MMCC
  - Loose Control: IVS

Flexible Confctrl Protocol:

- MCP

COLLABORATIVE ENVIRONMENTS FOR CONCURRENT ENGINEERING DESIGN (CECED)

IETF
Columbus OH
March 29 - April 2, 1993

Ruth Lang
SRI International

COLLABORATIVE ENVIRONMENTS FOR CONCURRENT ENGINEERING

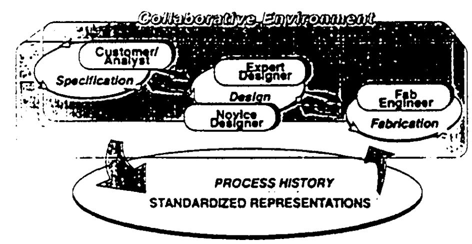

Informal Design Environments

Collaboration Should Support:

- Users in Existing Workplaces with Their Tools
- Informal Design Sessions
- Expert/Novice Dialogs

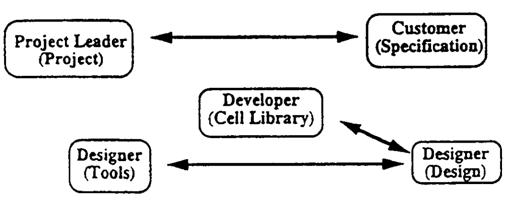

ELEMENTS OF THE DESIGN HISTORY

In Ceced

Completed Design plus Documentation

What is the design?

Design History

How was the design created?

Why were the design choices made? (Design Rationale)

Collaborative Dialogue

CECED Collaboration Services

Collaboration Services	Novice/Expert Consultation	Peer Designer	Designer/Customer Specification
Shared Presentation	Yes	Yes	Yes
Shared Tool Control	Yes	No	No
Shared Remote Pointer	Yes	Yes	Yes
Data Access	All Modify	Owner modifies, other read	Owner modifies, others read
Same Tools	Yes	No	No
Same Views	Yes	Yes	Yes

CECED System Architecture

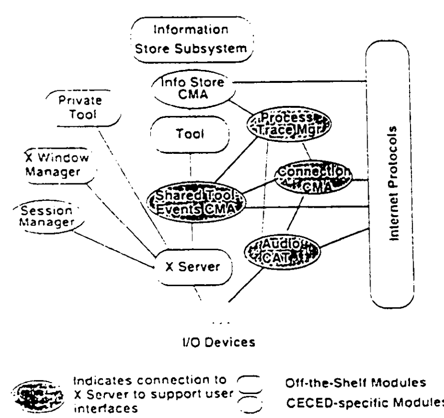

The TOURING MACHINE™ Session Abstraction

contact:
Abel Weinrib
abel@bellcore.com

Bellcore
Bell Communications Research

Touring Machine Infrastructure

Applications Programming Interface (API)
- "language" for writing multimedia communications applications
- makes available abstract network capabilities and core services

Flexible session management
- separation of application *policy* from network *mechanism* provided by infrastructure
- separate control of session from transport
- multiple active sessions
- separate control of media
  - audio
  - video
  - data

Fully integrated name server
- name and access transient as well as static objects (e.g., users, client applications, communications sessions)

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

Other services
- authentication
- negotiation

AW 3/27/93

Multimedia Telecommunications Service

One, of many, Touring Machine applications

- list of users
- separate control of media
- control of multiple sessions
- suspend/resume
- add/delete parties
- change configuration

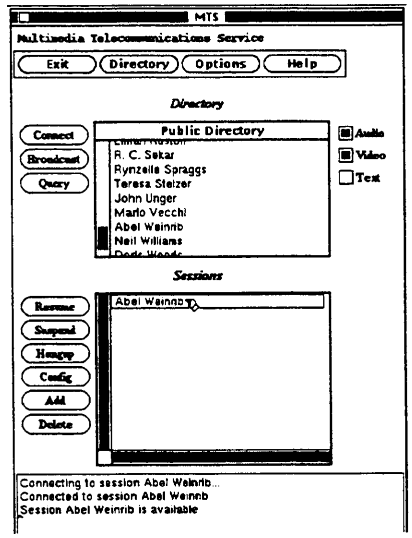

AW 3/27/93

Application Programming Interface

Client registration
(registerClient <token> <clientName> <regAction>+)

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

Session establishment and modification
(sessionCreate...<sessionName> <clientID> <sessionAction>+)
"session" separates *control* from *transport* for communicating applications

- negotiate and set privacy, billing, and other policies
- negotiate and define transport topology (as connectors)

Network-access control
(endpointMap, endpointUnmap, portCreate, endpointAssign)
map and unmap give applications control over the network-access resources separate from the network resources allocated for transport

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

Name server queries (nsQuery... <keys> <attributes>)
allows applications to access system information

Inter-client message forwarding (messageSend...)
provides facility for clients to exchange control messages

AW 3/27/93

Session Protocol

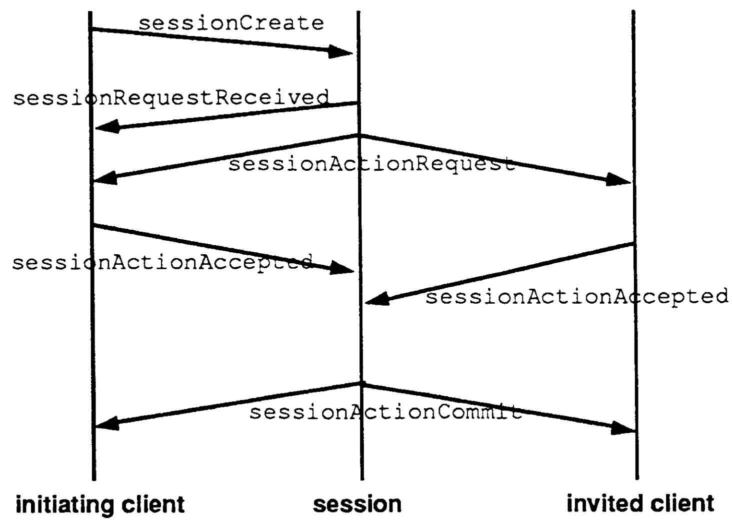

AW 3/27/93

sessionCreate

session initiator sends

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

where `<sessionAction>` **can be any of the following:**

```
(addClient <clientName>+)
(delClient <clientName>+)
(addCon <conName> <conType>)
(delCon <conName>)
(addSource <conName> <clientName:endpointName>+)
(delSource <conName> <clientName:endpointName>+)
(addSink <conName> <clientName:endpointName>+)
(delSink <conName> <clientName:endpointName>+)
(moveEndpoint <fromConName> <toConName>
 <clientName:endpointName>+)
(setPrivacy <privacy>)
(setPermission <permission>)
(setUDA <userDefinedAttributes>+)

<privacy> = {self, group, all}
<permission> = {private, protected, public}
```

AW 3/27/93

sessionRequestReceived

session responds with

```
(sessionRequestReceived <token> <sessionName>
  <clientName> <reqNum>)
(sessionRequestDenied <token> <sessionName>
  <clientName> <reason>)
```

AW 3/27/93

Example of sessionCreate

```
(sessionCreate 9876 "bob:app:s1" 34523667
    (addCon "vCon" "video") (addCon "aCon" "audio")
    (addSource "vCon" "bob:app:camera"
      "joe:app:camera" "ann:app:camera")
    (addSink "vCon" "bob:app:monitor"
      "joe:app:monitor" "ann:app:monitor")
    (addSource "aCon" "bob:app:mic" "joe:app:mic"
      "ann:app:mic")
    (addSink "aCon" "bob:app:speaker"
      "joe:app:speaker" "ann:app:speaker")
    (setPrivacy "all") (setPermission "protected")
    )
```

AW 3/27/93

Connectors

a connector is an association of sources and sinks for transport

f **describes "presentation control"**

- **eventually, general presentation control language**
- **currently, typed connector**
  - **video bridge, video PIP, audio bridge, audio mixer**

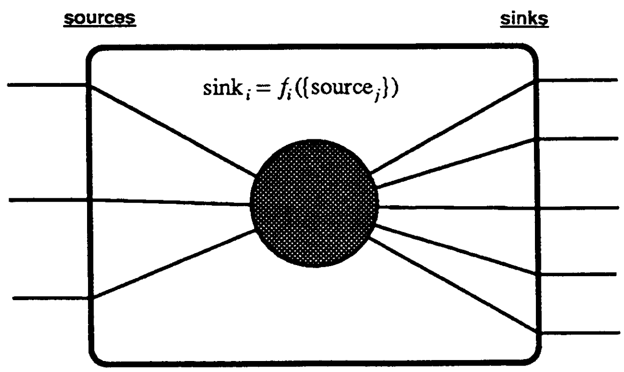

AW 3/27/93

sessionActionRequest

session sends to all involved clients

```
(sessionActionRequest <reqNum> <sessionName>
  <recipient> <originator> (memberList
  <clientName>+) [(forwardingPath <toClient>
  <fromClient>)]+ <requestAction>+)
```

where `<requestAction>` **can be any of the following:**

```
(addClient)
(delClient)
(addEndpoint <endpointName>+)
(delEndpoint <endpointName>+)
```

clients respond with

```
(sessionActionAccepted <reqNum> <sessionName>
  <clientID> [<endpointName:portName>]+)
(sessionActionDenied <reqNum> <sessionName>
  <clientID> [<reason>])
(sessionActionForward <reqNum> <sessionName>
  <clientID> <newClient>)
```

AW 3/27/93

sessionActionCommit/Abort

session sends to *all* clients

```
(sessionActionCommit <reqNum> <sessionName>
  <clientName> (client [<clientName>]+)
  [(forwardingPath <toClient> <fromClient>)]+
  [<commitAction>]+)
```

where `<commitAction>` **can be any of the following:**

```
(con <conName> <conType>
  [(source <endpointName>+)]
  [(sink <endpointName>+)])
```

or

```
(sessionActionAbort <reqNum> <sessionName>
  <clientName> <reason>)
```

AW 3/27/93

sessionChange

once a session is created, a client may initiate a change

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

a session ends when all of its members are deleted.

AW 3/27/93

VideoTeam™

Video Conferencing for the Defense Simulation Internet

BBN VideoTeam™

The Defense Simulation Internet

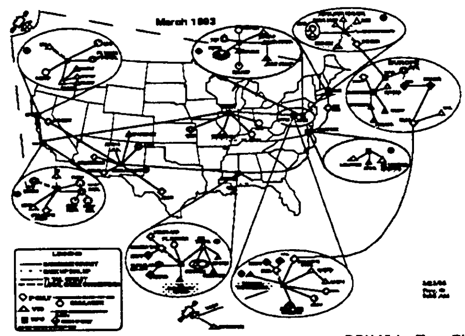

BBN VideoTeam™

Design Goals

Support "Real" Video Conferences
- Conferences are Important Meetings
- Users are not Computer Experts
- High-Quality Audio/Video Is Required
- Users May Arrive Late

Robustness
- Must survive temporary outages (net, equipment)

Flexibility and Scalability
- Different codecs, interoperability, workstation?
- Extend to Larger Meetings or Lectures
- Adapt to Other Media, e.g. Simulations

BBN VideoTeam™

"Sticky" Protocol Overview

Participant Management
`RING, REFUSE, HELLO, BYE`
Resource (Media) Management
`WANT, DONE, WATCH`
Floor Management
`REQ, UNREQ, FLOORLIST`
Utility
`ACK, NAK`

No Distinction Between "Start Up" and "Steady State"

All Interactions are Strictly Bilateral

Resource Negotiation uses "Late Binding" Approach

BBN VideoTeam™

Two-Site Example

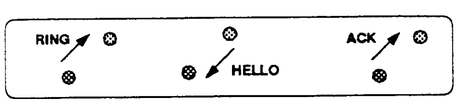

RING and HELLO convey <u>site list</u> and <u>local resource list</u>

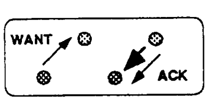

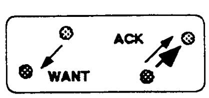

WANT requests access to media stream and involves negotiation

(Note that Media Stream is one-way.)

BBN VideoTeam™

Three-Site Example

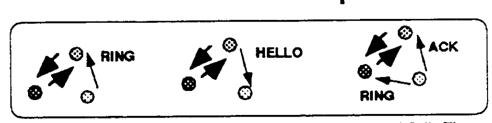

New Site Connects to Any Site, Learns of Other Sites, and Calls Them

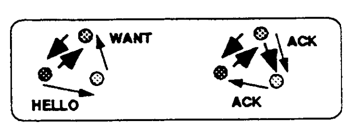

In Parallel, It Requests Media Streams

(Multicast Streams are Handy, Negotiation is Necessary)

BBN VideoTeam™

A Few Reflections

- Fully Decentralized, even at Start-Up
- No Distinction between Start-up and Late Arrival
- Existing Conference continues without "Resynch" Period
- Multiple Sites can simultaneously start the Same Conference
- Distinct Conferences Automatically Merge on Connect
- Site Lists will be Inconsistent during Transitions
- *Media Requests may Require Renegotiation of Existing Streams*

BBN VideoTeam™

Switching a View

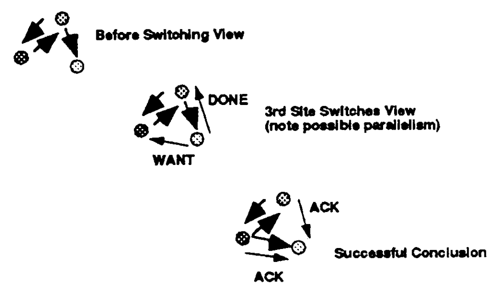

BBN VideoTeam™

Floor Management

Is Easy to Implement

Admissions Control

Is Hard to Define without a Moderator

BBN VideoTeam™

Scalability / Flexibility

Large Lectures

- Lecturing Site "trims" site list that it sends in HELLO
- Lecturing Site never issues WANT
- Results in one outbound multicast stream from lecturer
- And a connection from lecturer to each viewer (no traffic)

Flexibility of Resources

- Protocol accomodates any resource type
- "Late binding" allows just-in-time bandwidth reservation
- Sources are free to issue media in multiple formats/streams

BBN VideoTeam™

Questions . . .

1. Why distinguish "startup phase" from "late arrivals"?
2. Can resource negotiation be handled by an algorithm, or must the participants help decide?
3. How does "early negotiating" accomodate late arrivals?

BBN VideoTeam™

MCP: Multiflow Conversation Protocol*

Raj Yavatkar

Giri Kuthethoor, Lakshman K

Department of Computer Science
University of Kentucky,
Lexington, KY

email: < *firstname* >@ms.uky.edu

*Supported in part by the National Science Foundation Grant No. NCR-9111323 and STI 9108764.

Goals of MCP

- Communication support for Multimedia collaborative applications.
  Example: Application with a group editor, image display, video, and voice
- Intra floor coordination - Floor control.
- Inter flow coordination - Coordination and synchronization of traffic over related streams.
- Multipoint communication over a WAN or LAN

Research Issues

- What are the appropriate communication abstractions?
- Degree of synchronization and concurrency control
- OS support?

Approach

- Policy and mechanism are orthogonal.
- MCP will enforce and provide mechanisms.

  Applications use MCP abstractions to implement policy.

Token-Based Control

- Token is assigned when a flow is created
- Sender must hold a token
- Token management primitives are provided

  - request, transfer, replicate, delete, distribute, status.

Various Floor Control Styles

- Strict Floor-control
- Activity-sensing
- Brainstorming
- Chalkboard interaction
- Discussion groups

Multi-Flow Conversations

- Assume a network-layer *flow* abstraction
- First create individual flows with appropriate performance requirements
- Create a conversation and add related flows
- Causal synchronization among messages sent over constituent flows
- Causal Synchronization is based on a notion of Δ causality

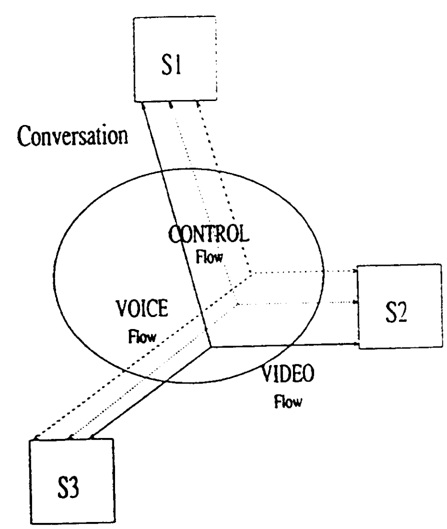

Implementation

- Underlying network layer with performance guarantees
- Two approaches
  - Library implementation
    - Upcalls from the library to indicate events.
  - Kernel
    - Applications use socket layer.
    - Upcalls are difficult.
    - Currently MCP uses signals and exception flags

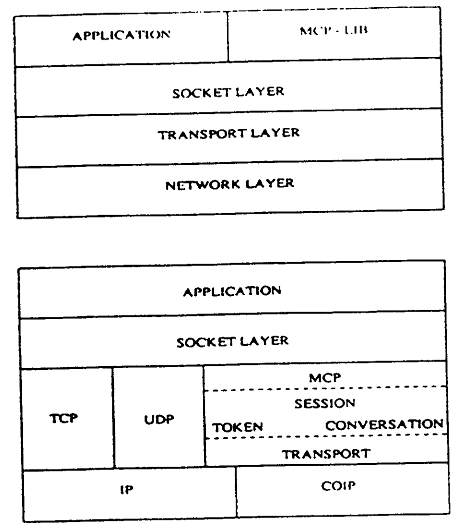

Applications and Test programs

- Mshell

  Command level interface to MCP.

- Mtalk

  Voice and text application.

- Nevot over MCP

  Provides various floor control styles.

Summary

- Both Kernel and library based implementation.
- Some experience with applications
- Still evaluating design
- Looking for input / feedback

An Architecture for Multimedia Connection Management

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

A Software Architecture for Packet Teleconferencing

- A modularized and layered design: salient components
  - *Connection Manager* coordinates multi-user, multimedia sessions
  - *Media Agents* handle media-specific details
  - *Configuration management* of end-system heterogeneity
    - Facilitates interoperation among different teleconferencing implementations
- A distributed connection control protocol (CCP)
  - Targeted for WAN operations: reliability, efficiency, robustness features
  - Conduit for control information both locally and remotely

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

Flow of Control Information

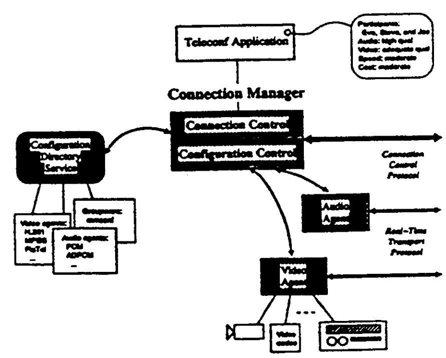

Functions Supported by CCP

- Membership: connect, invite, join, disconnect
- Media configuration: media type(s), parameters
- Floor control: receiver selected, sender selected
- Remote control: camera selection, autopilot
- Pre-schedule
- General mechanism for info exchange

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

A Scalable Architecture for Personal Teleconferencing

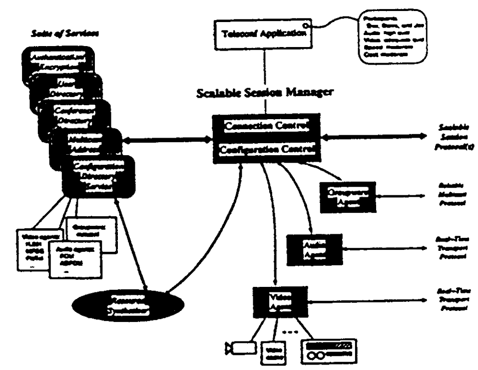

CURRENT MEETING REPORT

Reported by Terry Gray/UWash

Minutes of the Interactive Mail Access Protocol BOF (IMAP)

Agenda

- Introductions
- Draft Working Group Charter
- Discussion

Results

There was general agreement that the focus of the proposed working group should be to refine and extend the existing IMAP2/IMAP2bis protocol, hopefully without breaking the installed base of IMAP2-capable software.

The wording of the proposed working group Charter was deemed generally acceptable.

A list of desired IMAP extensions was made. Many of the proposed features appear to be within the scope of the CMU IMSP (Interactive Mail Support Protocol) project, rather than IMAP. Also, some are likely to be incompatible with current software.

It was clear that disconnected operation, ala DMSP/PC-Mail, was a very high priority.

Wish List

- Support for disconnected operation
- Off-line sorting of mailbox
- Background server searching and sorting
- Shared mailbox per-user state (like a .newsrc, but for mailboxes)
- Function to determine where to submit messages
- Storage/retrieval of MUA configuration data
- Minimal non-plain text authentication
- Minimal confidentiality (XOR with shared secret)
- Test assertion that PEM does not affect IMAP
- Remote printing (from IMAP server's copy of msg)
- Improved searching

Next Steps

The proposed Charter, with minor modifications, will soon be submitted to the IETF Application Area Directors for endorsement and forwarding to the IESG for approval as a working group.

The wish list will be combined with other input and categorized according to scope, compatibility, etc.

References

To subscribe to the IMAP Mailing list, send your request to:

<imap-request@cac.washington.edu>

Attendees

Karl Auerbach	karl@empirical.com
Sandy Bryant	slb@virginia.edu
William Chung	whchung@watson.ibm.com
James Conklin	jbc@bitnic.educom.edu
Mark Davis-Craig	mad@merit.edu
Roger Fajman	raf@cu.nih.gov
Ned Freed	ned@innosoft.com
Terry Gray	gray@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Steven Hubert	hubert@cac.washington.edu
David Katinsky	dmk@pilot.njin.net
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Bob Morgan	morgan@networking.stanford.edu
Robert Reschly	reschly@brl.mil
Marshall Rose	mrose@dbc.mtview.ca.us
Carl Schoeneberger	70410.3563@Compuserve.com

2.1.1 Internet Message Extensions (822ext)

Charter

Chair(s):

Gregory Vaudreuil, gvaudre@cnri.reston.va.us

Mailing Lists:

General Discussion: ietf-822@dimacs.rutgers.edu
To Subscribe: ietf-822-request@dimacs.rutgers.edu
Archive: ietf.cnri.reston.va.us:~/ietf-mail-archive/822ext/*

Description of Working Group:

This Working Group was chartered to extend the RFC 822 Message format to facilitate multi-media mail and alternate character sets. RFCs 1341 and RFC 1342 document the Multi-Media Extensions for Internet Mail.

The Working Group will work to progress MIME to Draft Standard status and provide a forum for the review of standards track content-type specifications and the review of character set extensions to MIME.

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.
Done	Post an Internet-Draft for the use of Japanese Characters for Internet Mail.
Done	Post a revised version of the MIME document as an Internet-Draft.
Done	Submit the revised MIME document to the IESG for consideration as a Draft Standard.
Done	Submit the Japanese Character set specification as an Informational document.

Internet-Drafts:

"MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies", 02/08/1993, N. Borenstein, N. Freed <draft-ietf-822ext-mime2-03.txt, .ps>

"MIME (Multipurpose Internet Mail Extensions) Part Two: Message Header Extensions for Non-ASCII Text", 03/22/1993, K. Moore <draft-ietf-822ext-mime-part2-01.txt>

"The text/enriched MIME Content-type", 03/23/1993, N. Borenstein <draft-ietf-822ext-text-enriched-02.txt, .ps>

"The Content-MD5 Header", 04/05/1993, M. Rose <draft-ietf-822ext-md5-02.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"

RFC 1437 "The Extension of MIME Content-Types to a New Medium"

RFC 1468 "Japanese Character Encoding for Internet Messages"

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

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

The Internet Message Extensions Working Group met twice in Columbus to review and approve the MIME protocol for submission as a Draft Standard. With a handful of changes, the MIME protocol was approved. The Working Group agreed to disband after publication of the revised document. Several new working groups will be formed to define extensions to the MIME protocol and additional content-types.

The solutions listed in the latest MIME issues list, as distributed periodically to the IETF-822 mailing list, were accepted with the following additions and changes:

- **Encoding of Content-Type Message**

  RFC1421 prohibits the use of a content-transfer-encoding other than 7bit, 8bit, and binary on the message type. This was designed to ensure that both the structure of a MIME message is visible without decoding and that nested encodings were not generated. Implementation experience has uncovered several problems with the use of message/partial and message/external-body when conversion is required in a gateway. In particular, using a non-null encoding of a partial 8-bit message for 7-bit transport is prohibited. Even if it was allowed, re-encoding the message into a 7-bit encoding would be likely to cause message size growth, defeating the intent of using message partial in the first case.

  The question for the Group was whether to limit encoding of any message type to 7-bit or only message/partial. The Group agreed to modify the prohibition to allow only content-transfer-encoding of 7bit for the message/partial content-type.

- **Representation of Filenames in Message/External-body**

  The inclusion of filenames in the content-type headers has the effect of requiring that all filenames be 7-bit ASCII. The Working Group discussed the likelihood that new operating systems will require a richer character set for filenames and the possibility that when this occurs the current filename mechanism may not be adequate. After lengthy discussion, during which the Group considered the possibility of using an encoded word from RFC1342, it was agreed that no changes should be made at this time, and that when needed, a new content-type could be defined with an enhanced mechanism.

- **Definition of Charset**

  The Working Group agreed to significantly trim the definition of a character set and to eliminate specific wording about specific unregistered character sets. The discussion of specific character sets not currently listed with IANA was eliminated (see the revised document for the new wording). Agreement was reached to remove Appendix F.2, the procedure for IANA registration, in favor of a statement pointing to the IANA for the procedure. It is expected that this procedure will evolve independently of MIME.

Issues related to the application of the general principle of a "charset" to specific current and future character sets is not part of the Charter of this Working Group and will be the subject of a new working group chartered to address the character set issues in a more general IETF context.

- **MIME-Version: 1.0 Header Semantics**

  The Working Group discovered that the MIME-Version header was insufficiently defined to be used for true versioning and that the interpretation of this header was not uniform across current implementations. Understanding that backward compatible changes to MIME were unlikely and that changing the version in the current header will cause at least one implementation to fail to recognize the message as valid MIME, the Working Group agreed that this header should now be considered a string constant; any version specific notes should be encoded as an RFC822 comment in the MIME-version header line, a feature available in all other RFC822 headers.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Gabe Beged-Dov	gabe@cv.hp.com
Nathaniel Borenstein	nsb@bellcore.com
Kevin Carosso	kvc@innosoft.com
George Chang	gkc@ctt.bellcore.com
William Chung	whchung@watson.ibm.com
James Conklin	jbc@bitnic.educom.edu
Al Costanzo	al@akc.com
Urs Eppenberger	eppenberger@switch.ch
Erik Fair	fair@apple.com
Roger Fajman	raf@cu.nih.gov
Ned Freed	ned@innosoft.com
James Galvin	galvin@tis.com
Christine Garland	garland@ihspa.att.com
Terry Gray	gray@cac.washington.edu
Alton Hoover	hoover@ans.net
Jeroen Houttuin	houttuin@rare.nl
Marko Kaittola	Marko.Kaittola@funet.fi
Neil Katin	katin@eng.sun.com
John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Mary La Roche	maryl@cos.com

Keith Moore	moore@cs.utk.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Francois Robitaille	francois.robitaille@crim.ca
Marshall Rose	mrose@dbc.mtview.ca.us
Carl Schoeneberger	70410.3563@Compuserve.com
Gregory Sheehan	gregory.c.sheehan@att.com
Einar Stefferud	stef@nma.com
Gregory Vaudreuil	gvaudre@cnri.reston.va.us

2.1.2 MHS-DS (mhsds)

Charter

Chair(s):

Kevin Jordan, Kevin.E.Jordan@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: mercury.udev.cdc.com:~/pub/archives/mhs-ds-archive

Description of Working Group:

The MHS-DS Working Group works on issues relating to Message Handling Services use of Directory Services. The Message Handling Services are primarily X.400, but issues relating to RFC822 use of Directory and Directory support for RFC822 and X.400 interworking are in the scope of the Group. Directory and Directory Services refer to the services based upon the CCITT X.500 recommendations and additional ISO standards, stable implementation agreements, and RFCs, as specified by the OSI-DS Working Group. The major aims of the MHS-DS Working Group are:

1. Define a set of specifications to enable effective, large-scale deployment of X.400.

2. Study issues associated with supporting X.400 communities which lack access to X.500 Directory, and define requirements for tools which: a) extract information from the X.500 Directory for use by non-X.500 applications, b) upload information into the X.500 Directory.

3. Coordinate a pilot project which deploys MHS information into the X.500 Directory and uses it to facilitate mail routing and address mapping. The results of this pilot will be documented, and experience gained from the project will be fed back into the Internet specifications created by the Working Group.

Goals and Milestones:

Ongoing	Provide a forum to discuss Directory support of Message Handling Services including the operational aspects of X.500 based routing in the Internet community and issues of migration from non-X.500 to X.500 based routing.
Ongoing	Establish and maintain liaison relationships with similar groups working on X.400 and X.500, e.g., RRE Mail and Messaging Group, IETF OSI-DS Working Group, IETF X.400 Operations Working Group and the IETF MIME-MHS Working Group.

Jan 1993 Post an overview of MHS use of Directory as an Internet-Draft.

Done Post a document on representing tables and subtees in the directory as an Internet-Draft.

Done Post an Internet-Draft on representing the O/R Address hierarchy in the Directory Information Tree.

Done Post an Internet-Draft on MHS use of Directory to support MHS Routing.

Done Post as an Internet-Draft a document on the use of the directory to support mapping between X.400 and RFC822 addresses.

Done Post as an Internet-Draft a document describing a simple profile for MHS use of Directory.

Done Post as an Internet-Draft a document on the use of the Directory to support routing for RFC822 and related protocols.

Done Submit as an Internet-Draft a document on MHS use of Directory to support MHS Context Conversion.

Done Post as an Internet-Draft a document describing the use of the Directory to support distribution lists.

Aug 1993 Submit the set of MHS-DS documents to the IESG for consideration as Experimental and Informational documents.

Internet-Drafts:

"Use of the Directory to support routing for RFC 822 and related protocols", 04/09/1992, S. Kille <draft-ietf-mhsds-822dir-02.txt, .ps>

"A simple profile for MHS use of Directory", 04/09/1992, S. Kille <draft-ietf-mhsds-mhsprofile-02.txt, .ps>

"Representing Tables and Subtrees in the Directory", 04/09/1992, S. Kille <draft-ietf-mhsds-subtrees-02.txt, .ps>

"Representing the O/R Address hierarchy in the Directory Information Tree", 04/09/1992, S. Kille <draft-ietf-mhsds-infotree-02.txt, .ps>

"Use of the Directory to support mapping between X.400 and RFC 822 Addresses", 04/09/1992, S. Kille <draft-ietf-mhsds-supmapping-02.txt, .ps>

"MHS use of the Directory to support distribution lists", 04/09/1992, S. Kille <draft-ietf-mhsds-mhsuse-02.txt, .ps>

"MHS use of Directory to support MHS Routing", 04/17/1992, Steve Kille <draft-ietf-mhsds-routdirectory-02.txt, .ps>

"MHS use of Directory to support MHS Content Conversion", 11/10/1992, S. Kille <draft-ietf-mhsds-convert-00.txt, .ps>

CURRENT MEETING REPORT

Reported by Urs Eppenberger/SWITCH and Keven Jordan/CDC

Minutes of MHS-DS Working Group (MHSDS)

Minutes from the previous meeting were approved. The updated Working Group Charter was reviewed and approved, and action items from the previous meeting were reviewed. The status of these action items follows:

Kevin Jordan	To write an Internet-Draft providing an overview of the main set of MHS-DS RFC's. Status: not yet done.
	To populate the DIT with US ADMD's. Status: <SPACE> and ATTMail have been added.
	To send mail to Erik Huizer when documents are ready for progression. Status: no documents are ready for progression yet.
Harald Alvestrand	To write pseudo-code for the Routing Document. Status: not done, but a contribution from Australia has been submitted, and Harald will consider it as a basis for the pseudo-code to be added to the Routing Document.
	To populate the DIT with Norwegian ADMD's. Status: <SPACE>, TELEMAX, and UNINETT added.
Jim Romaguera	To produce a document on "other" MHS-DS issues. Status: action item dropped.
	To populate the DIT with Swiss ADMD's. Status: ARCOM has been added.
Sylvain Langlois	To populate the DIT with French ADMD's. Status: not yet done.
Panos Tsigaridas	To populate the DIT with German ADMD's. Status: <SPACE> and DBP added.

In addition to the ADMD entries described above, the DIT has also been populated with the following ADMD entries:

```
C=GB; A=' '
C=GB; A=Gold 400
C=GB; A=MARK400

C=ES; A=' '
C=ES; A=MENSATEX

C=US; A=TELEMAIL
```

Thus, the DIT has already been populated with basic information about many X.400 ADMD's. This establishes the initial infrastructure needed to allow the MHS-DS pilot project to move forward.

Discussion of New Issues

Kevin Jordan raised the issue of the need for defining a new object identifier to identify the X.400/88 X.410 mode application context. X.400/88 allows MTA's to establish X.410 mode connections using the RTS protocol identifier 12. There currently is no OID which identifies this application context. The Routing Document defines an OID for the X.400/84 X.410 mode application context, but it doesn't define one for the X.400/88 X.410 mode context (and neither does the standard). Kevin recommended that a new OID should be created for the X.400/88 X.410 mode context, and that this new OID should be added to the Routing Document. The recommendation was accepted. Steve Kille will update the document accordingly.

Discussion of the Pilot Project

Most of this MHS-DS meeting was devoted to discussion of the MHS-DS pilot project, Project Long Bud. Populating the DIT with ADMD entries was an important accomplishment since the last MHS-DS meeting because it establishes the infrastructure which allows Long Bud to progress. A productive discussion took place regarding how best to populate the DIT with X.400 routing information and how best to progress the pilot project. The conclusions were:

1. Steve Kille will update the following three documents and submit them for progression by May 1st:
   - "Representing Tables and Subtrees in the Directory"
   - "Representing the O/R Address Hierarchy in the Directory Information Tree"
   - "Use of the Directory to Support Mapping Between X.400 and RFC 822 Addresses"

2. Initially, participation in Long Bud should be restricted to the GO-MHS community (the GO-MHS community is the Internet X.400 community). To simplify initial routing requirements, only MTA's willing to accept connections from all other MTA's will be eligible to participate. In addition, all participating MTA's must be able to reach all of the existing GO-MHS community. We will start by adding a small collection of MTA's which are interconnected by RFC1006 and/or ISO CLNS. We will also configure the open routing tree such that an MTA at SWITCH is defined as a default X.400 route.

   After the initial routing infrastructure is established and is being used successfully by the initial set of participating MTA's, eligibility restrictions will be relaxed.

3. Some documents are needed to define the pilot project formally and provide guidance on how to participate. Jim Romaguera agreed to coordinate and to help write these documents. In fact, he assembled a team and generated rough drafts within hours after the MHS-DS meeting ended. Kevin Jordan will edit the documents and distribute them as Internet-Drafts as soon as possible.

The need for MHS-DS tools was discussed. The conclusions reached were that the following tools need to be built:

1. Routing information browser and verifier. This tool would perform functions such as:
   (a) Verify that specified O/R addresses are reachable using routing information specified in the DIT.
   (b) Verify MTA connection information such as presentation addresses and credentials.
   (c) Chase references from one X.500 entry to another to ensure that there are no dangling pointers.

   Urs Eppenberger agreed to write the specifications for this tool.

2. Route tracer.

   This tool would accept an MTA name and destination O/R address and use the DIT to produce the set of possible routes from the MTA to the indicated destination.

In addition, the following tools are also desirable:

1. X.500 Perl

   This tool would be a version of Perl which would include built-in DUA functions. This would greatly facilitate the development of other X.500-capable tools.

2. Routing information DUA

   This would be a specialized tool facilitating the creation, maintenance, and management of X.400 routing information in the DIT.

3. Migration tools

   These tools would facilitate interworking between existing MTA's which are not X.500-capable and ones that are. Panos Tsigaridas is already planning to build these tools. His goal is to build a tool which accepts existing RARE routing documents as input (with possible enhancements for defining X.500 information) and generates X.500 entries. He is also planning to build a tool which pulls routing information from the DIT and creates RARE routing documents from it. Panos estimates that these tools can be completed by June.

MTA Support for Long Bud

The PP 6.4 beta release contains support for X.500-based routing, so this software can be used to participate in the pilot project. PP 7.0 will also support X.500-based routing, as does ISODE Consortium Release 1.

Next Meeting

The next MHS-DS meeting will take place at the IETF meeting in Amsterdam in July. Erik Huizer suggested that an MHS-DS demonstration be organized for that meeting. The audience for the demonstration would be MTA and DSA managers. The demonstration would not need to be very fancy. It could simply be based upon a utility which browses the DIT to show the routing information, and it would show one or two operational MTA's which would actually be using the DIT to make their routing decisions.

Attendees

Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
George Chang	gkc@ctt.bellcore.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Urs Eppenberger	eppenberger@switch.ch
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Jeroen Houttuin	houttuin@rare.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	Kevin.E.Jordan@cdc.com
Marko Kaittola	Marko.Kaittola@funet.fi
Mary La Roche	maryl@cos.com
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr

Ignacio Martinez	martinez@rediris.es
Edward Reed	eer@cinops.xerox.com
Francois Robitaille	francois.robitaille@crim.ca
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Einar Stefferud	stef@nma.com
Catherine Summers	cfs@cos.com
Louisa Thomson	louisa@whitney.hac.com

2.1.3 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) is currently a 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.
Done	Post an Internet-Draft describing the "core" set of Registered conversions for bodyparts.
Jul 1992	Submit a completed document to the IESG describing MIME-MHS Interworking as a Proposed Standard.
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. Kille, R. Miles <draft-ietf-mimemhs-mapping-02.txt>

"Equivalences between 1988 X.400 and RFC-822 Message Bodies", 07/01/1992, H. Alvestrand, S. Thompson <draft-ietf-mimemhs-body-equival-02.txt>

"HARPOON: Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages", 09/28/1992, H. Alvestrand, J. Romaguera, K. Jordan <draft-ietf-mimemhs-harpoon-02.txt>

CURRENT MEETING REPORT

Reported by Steve Thompson/Soft-Switch

Minutes of the MIME-MHS Interworking Working Group (MIMEMHS)

Document Review

The Group reviewed the status of the three MIMEMHS documents in progress.

- "Mapping Between X.400 and RFC-822 Message Bodies - Stable"
- "Equivalences Between 1988 X.400 and RFC-822 Message Bodies - Stable"
- "Harpoon: Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages - Issues Still Outstanding"

The consensus was that the Mapping and the Equivalence documents could be forwarded for approval as Proposed Standards, but that HARPOON needs to be discussed further.

HARPOON Discussion

The remainder of the meeting was spent discussing HARPOON, which covers MIME-X.400(84) interworking. Some alternative MIME to X.400(84) approaches were discussed at length and discarded. The only HARPOON issue that had not reached closure was the body part format of MIME messages in an X.400 message. This issue had been discussed on the list, but no consensus was reached. There were good arguments for each solution, but everybody agreed that one solution was better than supporting both. Consensus was reached on going with the IA5 body part rather than a binary one. Harald Alvestrand will update the HARPOON document and reissue a final draft within the next week or so.

The Group decided to forward all three documents for approval as Proposed Standards, once the HARPOON editing is complete.

Attendees

Gabe Beged-Dov	gabe@cv.hp.com
Nathaniel Borenstein	nsb@bellcore.com
Kevin Carosso	kvc@innosoft.com
George Chang	gkc@ctt.bellcore.com
William Chung	whchung@watson.ibm.com
Kishan Dudkikar	kishan@icm1.icp.net
Urs Eppenberger	eppenberger@switch.ch
Francois Fluckiger	fluckiger@vxcern.cern.ch
Ned Freed	ned@innosoft.com
Jeroen Houttuin	houttuin@rare.nl

Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	Kevin.E.Jordan@cdc.com
Marko Kaittola	Marko.Kaittola@funet.fi
Mary La Roche	maryl@cos.com
Ignacio Martinez	martinez@rediris.es
Francois Robitaille	francois.robitaille@crim.ca
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Marshall Rose	mrose@dbc.mtview.ca.us
Gregory Sheehan	gregory.c.sheehan@att.com
Einar Stefferud	stef@nma.com
Catherine Summers	cfs@cos.com
Steven Thompson	sjt@gateway.ssw.com
Louisa Thomson	louisa@whitney.hac.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Gregory Vaudreuil	gvaudre@cnri.reston.va.us

2.1.4 Minimal OSI Upper-Layers (thinosi)

Charter

Chair(s):
Peter Furniss, p.furniss@ulcc.ac.uk

Mailing Lists:
General Discussion: thinosi@ulcc.ac.uk
To Subscribe: thinosi-request@ulcc.ac.uk
Archive: pluto.ulcc.ac.uk:/ulcc/thinosi/thinosi-mail-archive.txt

Description of Working Group:

The OSI upper-layer protocols (above Transport) are rich in function and specified in large, complex and numerous documents. However, in supporting a particular application, the protocol actually used is only a subset of the whole. An implementation is not required to support features it never uses, and it is, or should be, possible to have relatively lightweight implementations specialized for a particular application or group of applications with similar requirements. The application protocol could be an OSI application-layer standard or a protocol originally defined for TCP/IP or other environment. It will be easier to produce such implementations if the necessary protocol is described concisely in a single document.

An implementation, of the mapping of X Window System protocol over OSI upper-layers, is based on this principle.

The Working Group is chartered to produce two documents:

"Skinny bits for byte-stream" a specification of the bit (octet) sequences that implement the OSI upper-layer protocols (Session, Presentation and ACSE) as needed to support an application that requires simple connection, and byte-stream read and write. This will be based on the octet sequences needed to support X. This will not be expected to be provide a full equivalent of TCP, nor to cover specific standardized protocols.

"Skinny bits for Directory" a specification of the bit sequences needed for the Directory Access Protocol - in the same style as the byte-stream specification, but to include DAP. The level of functionality of this is to be determined.

An important aspect of the Group's work is to find out if it is possible to produce useful and concise specifications of this kind. A minor part is to think of better names.

The Group will also encourage the deployment of X/OSI implementations and interworking experiments with it.

Goals and Milestones:

May 1993 Post an Internet-Draft for "Skinny bits for Byte-Stream".

Aug 1993 Post an Internet-Draft for "Skinny Bits for Directory".

Dec 1993 Submit the "Skinny Bits for Byte-Stream" specification to the IESG for consideration as a Proposed Standard.

Mar 1994 Submit the "Skinny Bits for Directory" specification to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Peter Furniss/Consultant

Minutes of the Minimal OSI Upper-Layers Working Group (THINOSI)

This was our first meeting as a Working Group. Most of the time was spent reviewing the first draft of the bytestream cookbook and especially the technical choices of what OSI upper-layers features should be included.

Some guiding principles emerged:

- Interworking with "full" OSI implementations was a key purpose of the thinosi approach and must be maintained, even at the cost of some complexity.

- Inclusion or exclusion of features is determined by (and determines) the set of application protocols that can be supported.

- t will be little trouble to add additional features (thus widening the application range) once the basic format and style is worked out - these could be alternatives in the same document or in separate ones.

- If at all possible the set of supported application protocols should include the DAP and X.400 P7 protocol. (These are not actually byte-stream in the sense we have now, but at least DAP should be possible with only a two octet change. P7 to be investigated. PRF)

A number of conclusions were reached (most of these are confirmations of what is in the first draft). *Editor's Note (md): A listing of these conclusions is available via ftp under thinosi-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

The Group was uncertain on whether the lengths of constructed elements should all be indefinite, definite, or follow a pattern of whatever is convenient. For interworking, all possibilities must be understood on receipt, but various performance payoffs are possible. The views of the mailing list will be sought.

There was also some discussion of related work in other circles.

- A profile for "minimal OSI functionality" (mOSI) is being developed in the OIW and EWOS regional workshops (these are the OSI profiling bodies for North America and Europe). This is intended to be part three of Common Upper-Layer Requirements. The mOSI profile and thinosi work have a common origin. At present, the upper-layer facilities supported in mOSI and in thinosi are not quite the same. One single selection of features for both would be highly desirable. Attempts will be made to align these - in particular the new draft of the cookbook should be made available to the EWOS technical liaison group (TLG) in time for their next meeting in April.

- Within X/Open, a new appendix for the XTI interface is being developed that defines the use of this interface for minimal OSI.

The Group will meet in Amsterdam.

Attendees

John Boatright	bryan_boatright@ksc.nasa.gov
Kevin Carosso	kvc@innosoft.com
George Chang	gkc@ctt.bellcore.com
Peter Furniss	p.furniss@ulcc.ac.uk
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Mark Needleman	mhn@stubbs.ucop.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Keith Sklower	sklower@cs.berkeley.edu
Klaus Truoel	truoel@gmd.de
William Warner	warner@ohio.gov
Les Wibberley	lhw24@cas.org
Charles Young	Charles.E.Young@att.com

2.1.5 Network Database (netdata)

Charter

Chair(s):

Daisy Rose, `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 the solution. Assign writing assignments for first draft of the document.
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 an 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-04.txt>

"Network Database Implementation Information Internet Draft", 12/16/1991, Daisy Shen <draft-ietf-netdata-implement-03.txt>

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 widely implemented. 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-01.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	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.
Done	Review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.

Request For Comments:

RFC 1179	"Line Printer Daemon Protocol"

2.1.8 OSI Directory Services (osids)

Charter

Chair(s):
Steve 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:

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.
Done	Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.
Done	Study the relationship of the OSI Directory to the Domain Name Service.

Internet-Drafts:

"Using the OSI Directory to Achieve User Friendly Naming", 11/26/1990, S. Hardcastle-Kille <draft-ietf-osids-friendlynaming-05.txt, .ps>

"A String Representation of Distinguished Names", 01/30/1992, S. E. Hardcastle-Kille <draft-ietf-osids-distnames-05.txt, .ps>

"Lightweight Directory Access Protocol", 04/17/1992, Wengyik Yeong, Tim Howes, Steve Hardcastle-Kille <draft-ietf-osids-lightdirect-03.txt>

"The String Representation of Standard Attribute Syntaxes", 05/05/1992, T. Howes, S. Hardcastle-Kille, W. Yeong <draft-ietf-osids-syntaxes-01.txt>

"DSA Metrics", 09/23/1992, P. Barker, R. Hedberg <draft-ietf-osids-dsa-metrics-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"

RFC 1384 "Naming Guidelines for Directory Pilots"

RFC 1430 "A Strategic Plan for Deploying an Internet X.500 Directory Service"

RFC 1431 "DUA Metrics"

CURRENT MEETING REPORT

Reported by Ed Reed/Xerox

Minutes of the OSI Directory Services Working Group (OSIDS)

The session began with introductions and was followed by a review of the Agenda. The Minutes of the November 1992 meeting were accepted without comment. The session continued with the Liaison Reports.

1. WG-NAP (Erik Huizer)

   There are three task forces working:

   - Work that Panos-Gavriil Tsigaridas is doing.
   - Data Management - how to get data into the directory, and keep it up to date and accurate.
   - Legal and Privacy issues - going to publish first results as an analysis of Dutch regulations and other regulations recently published.

2. NADF (Tim Howes)

   Decided to open to users, but as non-voting members. Piloting continues.

3. DISI (Tim Howes)

   Reformed under the Integrated Directory Services Working Group, with the addition of Whois++, want to address general directory issues.

4. AARNet (Mark Prior)

   One of the divisions of Telecom has joined the pilot. Currently doing a Whois++ startup. Trying to get a new binary distribution of ISODE 8.0 ready.

5. Paradise

   The First project ended with 1992. Next Project (transition) will run sixteen months. Now also has INRIA as a partner. Using Quipu and Pizzaro implementations. DE now will perform very broad (c=??/cn=name) searches. China, Slovenia (means ISODE 8.0 required), Croatia, Slovakia have joined. A question was raised about the issue of intellectual property rights associated with DE and other Paradise tools.

6. NREN-NIS (Sri Sataluri/Mark Kosters)

   InterNIC will begin providing services April 1. - InterNIC personal listings will be available. Will provide DE access, and other X.500 services and WAIS. Plan to make WAIS information available via X.500, too, but not immediately. Plan to make registration data available via X.500.

7. DOD

   Defense messaging system will take over the old switches including Autodin. Will be X.400 based, and are building an X.500 directory support.

8. Integrated Services Panel (US/GSA)

   There's a newsletter describing efforts available. Directory services now are flat file, with X.500 direction.

Progression of Standards (Erik Huizer)

- Published the Strategy Document.

  There was a long discussion which had been prompted by John Curran but he was unable to be here to discuss.

- LDAP

  It was not clear whether the Document was published, but it had left the IESG. There may be a block of some sort - Tim Howes took an action to see if it's being held up in the IAB.

No other documents are pending immediate progression. DSA Metrics will be discussed later.

Experiment Progress

Editor's Note (md): The details of this experiment are available via ftp under osids-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

The Group has not concluded that the draft OSI-DS 15, "Handling Quality of Service in the Directory," should be abandoned, if more of the values specified there are in fact implemented. The question is whether DSA and DUA implementors will build OSI-DS 15 approaches, or not.

There is some overlap between this experiment and the MADMAN efforts. Gavriil Tsigaridas reported some of their efforts have raised an issue with DS-15's use, or lack of, object type data in the QOS database. One approach is to just record information about times to find people.

Information is probably only of interest at the local level - views are too different from other places via other access mechanisms.

There is some value to share the implementation approach taken, but this is an experiment, still.

It seems as though the Group has gone far enough on DS-15, and should look in some other direction. DS-15 is complementary, but should be extended with interface native information which doesn't belong in the directory. Paul will see if he can make what he's done available for others to implement. A poll of the Group indicated continued interest in developing OSI-DS 15, but there were no volunteers to be the editor.

Schema Working Group

A previous meeting chartered a small subgroup to look at this. That Group never got together. There have been other issues (JPEG, etc.,) which have come up needing help, too, but still no volunteers to edit.

Panos Gavriil Tsigaridas' Document

Panos asked people to please read his document. Applications need the ability to use a common repository for information about management information, there would be a valuable synergy.

Charter Review

Steve and Erik each published Draft Charters. Steve doesn't think it makes sense to put things into the Charter which need to be done if there's not support from the Working Group members to do them. He proposes four:

1. Liaisons
2. Schema Coordination
3. DSA/DUA Metrics
4. IP Address Representation

Erik pointed out that the Group needs to be stricter in its procedures and resources as the IETF grows - specifically with regard to Charter and time schedules. Only if there are concrete objectives and times will the Charter be renewed. Erik's list includes:

1. Non-white pages use of the directory
2. Test strategies
3. Schema management
4. Guidelines for technical implementation, migration to 1993, and database coupling.

Paradise has an objective including interoperation of directory services. Interoperation is more properly a target of pilot projects, with which the Group wants to liaison, but that should not be part of this Group's Charter.

Perhaps if the Group defined where the holes in the standard exist which preclude interoperability, and publishes RFCs to fill the holes, then at least there would be a unified face to the implementors. For instance, Siemans has delivered an RFC based product, which goes beyond the OSI Standard, when pressed by pilot managers.

To some extent, this seems to be a necessary activity, in spite of the continuing claims by vendors that the extensions are non-standard, and will be obsoleted by the next standard.

(Erik) - IDS will focus on general problems relating to directory services, while OSIDS will focus on X.500 specific issues.

Metrics - Roland Hedberg

Roland has been testing the Siemens DSA and has also looked at interoperability testing. There are holes in the standards - schema handling, access control, etc. There appears to be holes in the metrics, too - they report good results, when you know there are problems.

Paradise - Paul Barker

Paul discussed new data management tools in more detail - which will be available shortly. Archie-like service based on X.500...begins with a leap of faith that it makes sense to record information about documents in the directory. The presentation provided an overview of the approach to be taken.

The sense of the Group was that the it should take the project being done as a work item. Paul will edit the papers he's done towards an RFC.

Representing WHOIS data in the X.500 Directory - Sri Sataluri

The objective is to provide access to information about network entities and to define a schema for representing that data. A concern was expressed that that approach may simply be replicating a centralized database, and not really distributing it - but there was disagreement with that concern.

Charting Networks in the Directory (OSI-DS 37-39) - Glenn Mansfield and Thomas Johannsen

Editor's Note (md): Notes to accompany the presentation are available via ftp under osids-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

The presentation included background, problem discussion and a description of a proposed solution. The objective is to provide a distributed map of the network. Not only topology, but the policies, costs, services, properties, administration and management attributes, and contacts. Many kinds of applications can use the information, but network management is the main thrust of the effort.

The CONMAN Project is addressing configuration management. The SOFTPAGES Project is addressing cost computation, using the configuration information from CONMAN, etc. In addition, file server contents is indexed in the directory.

The consensus of the Group was that it should be dealing with the problems described here. A subgroup of volunteers agreed to meet over dinner and plan work (Paul Barker, Tim Howes, Thomas Johannsen, Mark Knopper (silent volunteer) (missed dinner), Mark Kosters, Ruth Lang, Sylvain Langlois, Bruce Mackey, Glen Mansfield, Ed Reed, Sheri Repucci, Sri Sataluri, Mark Smith and Scott Williamson).

This Group identified a list of documents to be published, and accepted volunteers to edit the them.

- Roadmap (Steve Kille)
- IP Addressing Schema (Glenn, Thomas, Mark Ko, Sri)
- Network Objects Schema (Thomas, Sri, Ed, Mark Ko.)
- RFC1279 Revision (Mark Ko.)
- Naming Layout (Sri)
- Transition Plan for Existing Services and Deployment (Scott, Glenn)
- Business Process Model (Operations Guidelines) - Glenn
- Security and Privacy (Tim)
- OSI Addressing (to be determined)
- XNS Addressing (Ed)

Abstract: Charting Networks in the Directory. Work in progress at AIC, WIDE, Tohoku University.

Editor's Note (md): The Abstract is available via ftp under osids-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Attendees

Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Jules Aronson	aronson@nlm.nih.gov
Paul Barker	p.barker@cs.ucl.ac.uk
Russell Blaesing	rrb@one.com
John Boatright	bryan_boatright@ksc.nasa.gov
George Chang	gkc@ctt.bellcore.com
Wayne Clark	wclark@cisco.com
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Simon Coppins	coppins@arch.adelaide.edu.au
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Marcello Frutig	frutig@rnp.impa.br
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Marco Hernandez	marco@mh-slip.cren.edu

Gerd Holzhauer	holzhauer1@applelink.apple.com
Jeroen Houttuin	houttuin@rare.nl
Tim Howes	tim@umich.edu
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Kevin Jordan	Kevin.E.Jordan@cdc.com
David Katinsky	dmk@pilot.njin.net
Steve Kille	S.Kille@isode.com
Mark Knopper	mak@merit.edu
Mark Kosters	markk@internic.net
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Mary La Roche	maryl@cos.com
Ruth Lang	rlang@nisc.sri.com
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Bruce Mackey	brucem@cinops.xerox.com
Bill Manning	bmanning@sesqui.net
Glenn Mansfield	glenn@aic.co.jp
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
Mark Prior	mrp@itd.adelaide.edu.au
Edward Reed	eer@cinops.xerox.com
Sheri Repucci	smr@merit.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Yzhak Ronen	y.ronen@homxa.att.com
Marshall Rose	mrose@dbc.mtview.ca.us
Srinivas Sataluri	sri@qsun.att.com
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgrass@bitnic.educom.edu
Catherine Summers	cfs@cos.com
Louisa Thomson	louisa@whitney.hac.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Alan Williamson	scottw@nic.ddn.mil
Russ Wright	wright@lbl.gov

2.1.9 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:

Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.
Done	Inaugural meeting.
Done	Produce a paper stating what ODA standards or profiles still need completing.
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.

CURRENT MEETING REPORT

Reported by Peter Kirstein/UCL

Minutes of the Office Document Architecture Working Group (ODA)

There had been less progress in implementations of the newer FOD-26 version of the ODA Profile than had been hoped by the time of this meeting. Moreover, while early implementations may be available during the third quarter of 1993, it is unlikely that there can be enough to justify any pilot by that time.

In view of the above, and of the very poor attendance at the meeting, it was agreed to terminate the Working Group. If there is sufficient interest at a later stage - e.g., the November 1993 meeting, then the Working Group may be reinstated. The current status is outlined below.

Documents Available

Some papers are available electronically form on the UCL-CS info-server: "info-server@cs.ucl.ac.uk". 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 Status of Q112 Implementations

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 or V2	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	??	Q2 93	UPC

UCL started shipping SLATE/ODA v1.2 in February 1992; this is based on SLATE v1.2. Few have tested it. BBN brought out V2 of SLATE during the fourth quarter of 1992. From a user viewpoint, this has better graphics and font support, supports a fuller character set, and is easier to configure; the UCL SLATE/ODA v2 supports only the same functionality as the SLATE/ODA v1.2 - because the extra facilities in the SLATE v2 are not supported in the Convertor. The Bull product has not been changed recently; 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 were delivered to UCL. They have not really been tested for interoperability by UCL; little interest in them was expressed in the last ODA WG meeting. The WordPerfect/ODA system from University Polytechnic of Catalonia has been delivered recently to UCL. Some problems still need to be ironed out, but the software should be available by June 1993.

There was also interest expressed in ODA vs SGML. UCL has produced an SGML to ODA convertor, which works with a specific DTD. This is available to interested parties.

The Status of FOD-26 Implementations

The Profile used in all the ODA document activities described hitherto is the Q112 Profile. A new Profile called FOD26 has now been ratified. 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 under ODAC (Bull, DEC, IBM, ICL, Siemens-Nixdorf and UNISYS as part of the ODA Consortium). The manufacturers are providing their further products only according to the FOD26 version. The ODAC Toolkits is now available in binary form for $100 U.S. from the Consortium.

In view of the above, there has been little interest in working with the Q112 versions of the software, and for that reason the Group chose not to have a meeting at the November 1992 IETF. There were a number of FOD26 developments announced; these include the following:

PRODUCT	Supplier	Status	Product	Availability	Source
SLATE/ODA V2	BBN/UCL	Testing	SLATE-yes	Q3 1993	UCL
Word-for-Windows/ODA	Bull	Testing	Yes	Q3 1993	Bull
WordPerfect	UPC	Testing	Yes	Q3 1993	UPC
ODAC Toolkit	ICL	Available	Yes	Now	ICL

Reference

1. S. Bayderee et al: The ODA Document Convertors, UCL Internal Report, Version 5, March 1993.

2.1.10 TELNET (telnet)

Charter

Chair(s):

Steve Alexander, steveaQlachman.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 RFC854, "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.
Done	Post an Internet-Draft describing the authentication option.
Done	Post an Internet-Draft describing the encryption option.
Mar 1991	Rewrite RFC854.
Done	Submit the authentication option to the IESG as an Experimental Protocol.
Jul 1993	Submit the encryption option to the IESG as an Experimental Protocol.

Internet-Drafts:

"Telnet Authentication and Encryption Option", 04/01/1990, Dave Borman
<draft-ietf-telnet-encryption-02.txt>

"Telnet Environment Option", 04/05/1993, S. Alexander <draft-ietf-telnet-envmnt-option-01.txt>

"Telnet Environment Option Interoperability Issues", 04/08/1993, D. Borman <draft-ietf-telnet-interoperability-00.txt>

Request For Comments:

RFC 1116	"Telnet Linemode option"
RFC 1184	"Telnet Linemode Option"
RFC 1372	"Telnet Remote Flow Control Option"
RFC 1408	"Telnet Environment Option"
RFC 1409	"Telnet Authentication Option"
RFC 1411	"Telnet Authentication: Kerberos Version 4"
RFC 1412	"Telnet Authentication : SPX"
RFC 1416	"Telnet Authentication Option"

CURRENT MEETING REPORT

Reported by Steve Alexander/Lachman Technology

Minutes of the TELNET Working Group (TELNET)

The Telnet Working Group met in two sessions on Tuesday, March 30, 1993. In the morning the Group first discussed the Environment Specification, RFC1408. This document contains an error in two sub-option values. It appears that more implementations do what the BSD/Cray implementation does than actually follow RFC1408. The Group agreed to update RFC1408 with the correct values and some text discussing interoperability. Steve Alexander will update the document and send it out for review.

The Working Group then discussed the latest draft of the Kerberos V authentication mechanism. This has been updated with support for delegation of credentials. Ted Ts'o from MIT indicated that this had been implemented and works. Ted also mentioned that the general authentication option needs a way to indicate that credentials will be delegated. The Group agreed to this. With support in the general mechanism, some changes will be needed to the Kerberos V mechanism document as well.

The next topic was Dave Borman's new merged authentication/encryption document. The document was only briefly discussed since the most recent version came out the Friday before the meeting. It was agreed that Steve Bellovin's suggestion that the document explain why encryption cannot be turned on once it has been turned off should be adopted. This recommendation and the credential delegation suggestion will be forwarded to Dave.

Jeff Schiller recommended that the Working Group Chair be more zealous in getting new drafts out as Internet-Drafts. There was general agreement on this point.

The Group spent the remainder of the morning discussing the HP proposal for Telnet multiplexing (Telnet MPX). It was generally agreed that HP's performance numbers were very impressive, but the consensus was that Telnet was not the best place to implement this. The Group recommended that a general solution, possibly at the transport or network layers, would be a better approach. It was also recommended that HP move the multiplexing mechanism to a TCP port other than 23 for experimentation, and that the revised document be issued as an Experimental RFC. There was also a brief discussion on enhancements to the MPX packet format.

The evening session was devoted to a discussion of Telnet 3270 emulation (TN 3270). Robert Moskowitz from Chrysler gave a synopsis of defects in the current TN3270 implementations. There was general agreement on these problems and the proposed solutions. In discussion with Russ Hobby, it was decided that TN3270 was somewhat tangential to the issues that the Telnet Working Group is pursuing. TN3270 implementors and users will form their own working group in the Applications Area. TN3270 will remain a telnet-based service so that features like Authentication and the Environment Option can be utilized if desired.

Attendees

Steve Alexander	stevea@lachman.com
Jim Barnes	barnes@xylogics.com
Fred Bohle	fab@interlink.com
Larry Campbell	larry@psycho.psy.ohio-state.edu
Roger Fajman	raf@cu.nih.gov
Antonio Fernandez	afa@thumper.bellcore.com
Eric Fleischman	ericf@act.boeing.com
Jisoo Geiter	geiter@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Cleve Graves	cvgpc@oc.com
Jeff Hilgeman	jeffh@apertus.com
Russ Hobby	rdhobby@ucdavis.edu
Scott Kaplan	scott@wco.ftp.com
Frank Kastenholz	kasten@ftp.com
Charles Kaufman	kaufman@zk3.dec.com
Bill Kelly	kellywh@mail.auburn.edu
Zbigniew Kielczewski	zbig@eicon.qc.ca
William Kwan	kwan@rabbit.com
John Linn	linn@gza.com
Kent Malave	kent@bach.austin.ibm.com
Marjo Mercado	marjo@cup.hp.com
Greg Minshall	minshall@wc.novell.com
Robert Moskowitz	3858921@mcimail.com
Baktha Muralidharan	murali@smaug.enet.dec.com
Clifford Neuman	bcn@isi.edu
Jon Penner	jjp@bscs.uucp
Christopher Provenzano	proven@csi.compuserve.com
Marcus Ranum	mjr@tis.com
Owen Reddecliffe	owen%wrq@mcimail.com
Steven Richardson	sjr@merit.edu
Rick Royston	rick@lsumvs.sncc.lsu.edu
Jeffrey Schiller	jis@mit.edu
Kitty Shih	kmshih@novell.com
Mark Silverman	mark@dss.com
Sam Sjogren	sjogren@tgv.com
Barbara Sterling	bjs@mcdata.com
Theodore Ts'o	tytso@mit.edu
Richard Warwick	richard@dss.com

2.1.11 Telnet TN3270 Enhancements (tn3270e)

Charter

Chair(s):

Robert Moskowitz, 3858921@mcimail.com

Mailing Lists:

General Discussion: tn2370e@list.nih.gov
To Subscribe: listserv@list.nih.gov
In Body: sub tn3270e <first_name> <last_name>
Archive: listserv@list.nih.gov

Description of Working Group:

The TN3270 Enhancements Working Group will document the current practices that provide limited support for 3270 devices over TELNET and will develop a specification that allows the 3270 family of devices, including printers, to function properly over TCP via TELNET. Topics such as authentication, which are being addressed by other working groups, are recognized as important to TN3270, but are beyond the scope of this effort.

The specification will draw on work already done by the Internet community for supporting 3270 devices through TELNET. It will be based on appropriate portions of IBM's published documentation on 3270 display and printer data streams and LU function management. Finally, it will make use of existing TELNET facilities where possible.

The Working Group will produce: An informational RFC documenting current TN3270 terminal practices, an experimental RFC that describes an interim approach to printing and LU name selection (this will address the work that is already under way and implementations of this partial solution that are already in place), and a standards track RFC specifying the TELNET protocols that support a fully functional 3270 display and printing environment. This RFC will supersede RFC1041 and the experimental RFC describing the interim approach to printing and LU name selection.

Goals and Milestones:

Apr 1930	Post an Internet-Draft describing the Interim approach to printing and LU name selection.
Apr 1993	Submit an Internet-Draft documenting current TN3270 terminal emulation practices.
May 1993	Submit the interim printing and LU name selection document to the IESG for consideration as an Experimental Protocol.

May 1993	Post as an Internet-Draft a protocol to support a fully functional 3270 display and printing environment over TELNET.
May 1993	Sumbit the document describing current TN2370 terminal practices to the IESG for consideration as an Informational Protocol.
Sep 1993	Submit the TN3270 TELNET specification to the IESG for consideration as a Proposed Standard.

2.1.12 X.400 Operations (x400ops)

Charter

Chair(s):

Alf Hansen, Alf.Hansen@delab.sintef.no
Tony Genovese, genovese@es.net

Mailing Lists:

General Discussion: ietf-osi-x400ops@cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@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.
Done	Internet-Draft available.
Dec 1991	Document ready for publication.

Internet-Drafts:

"Operational Requirements for X.400 Management Domains in the GO-MHS Community", 03/11/1992, Robert Hagens, Alf Hansen <draft-ietf-x400ops-mgtdomains-ops-05.txt>

"X.400 use of extended character sets", 06/18/1992, Harald Alvestrand <draft-ietf-x400ops-charactersets-02.txt>

"Postmaster Convention for X.400 Operations", 11/23/1992, C. A. Cargille <draft-ietf-x400ops-postmaster-01.txt>

"Assertion of C=US; A=IMX", 12/11/1992, E. Stefferud <draft-ietf-x400ops-admd-02.txt>

"Using the Internet DNS to maintain RFC1327 Address Mapping Tables", 01/05/1993, C. Allocchio, A. Bonito, B. Cole <draft-ietf-x400ops-dnsx400maps-02.txt>

"Using the Internet DNS to maintain X.400 MHS Routing Informations", 02/01/1993, C. Allocchio, A. Bonito, B. Cole <draft-ietf-x400ops-dnsx400rout-02.txt>

"Evaluation of ADMDs and Integration aspects with respect to the R&D messaging community", 02/25/1993, J. Romaguera, P. Klarenberg <draft-ietf-x400ops-evaluation-admd-00.txt>

"Table distribution", 03/09/1993, M. Kaittola <draft-ietf-x400ops-tbl-dist-00.txt>

Request For Comments:

RFC 1405 "Mapping between X.400(1984/1988) and Mail-11 (DECnet mail)"

RFC 1465 "Routing coordination for X.400 MHS services within a multi protocol / multi network environment Table Format V3 for static routing"

CURRENT MEETING REPORT

Reported by Urs Eppenberger/SWITCH

Minutes of the X.400 Operations Working Group (X400OPS)

Agenda

- Introduction
- Action List Review
- Liaison Reports
- Review Charter
- GO-MHS Community Document
- Review X.400 Operations Documents
- Presentation by Jeroen Houttuin on "Mail based servers"
- Presentation by Marko Kaittola on "Table Distribution"
- AOB and Plan for next meeting in Amsterdam

The Minutes do not follow the Agenda but provide two main sections, liaison reports and document status. All Agenda items have been grouped into these two sections. A list of actions is appended to make it easier for the Working Group Chairs to track Action Items. The Minutes of the last X400OPS meeting in Washington, DC were approved.

Review Charter

The Goals in the Charter of the X400OPS Group need to be reviewed. The references to OPS-* should be replaced by the names of the documents in the Internet-Drafts Directory. The title and time scale of draft-ietf-x400ops-tbl-dist-00.txt (OPS-6) is not realistic.

Liaison Reports

- IETF MHSDS Working Group.

  The main focus is on the establishment of a pilot by adding routing information to the directory to gain operational experience. The Group hopes to have something to demonstrate at the next IETF in Amsterdam.

- RARE WG-MSG Working Group on Mail and Messaging.

  A task force works on a number of documents on X.400(88) deployment.

  A lot of other documents have been sent for comments to both X400OPS and WG-MSG.

WG-MSG is not a working group in the IETF sense, it is more like an IETF Area since it covers all mail protocols and is an ongoing Group which may split off chartered task forces.

WG-MSG has been approached by Allan to see if the postmaster document would be appropriate to be published as RARE Technical Report (RTR). The chairman of WG-MSG, Harald T. Alvestrand, is positive about it. It might be bundled with the GO-MHS requirement.

- Electronic Mail Association (EMA).

  Nothing was reported.

- European Electronic Mail Association (EEMA).

  Jim Romaguera reported from EEMA. He has been approached by some ADMD service providers with questions concerning the connection of their services to the Internet. He will forward the questions to the X400-OPS list.

  The most interesting activities are in the EEMA PRMD operators group and in the EEMA ADMD operators group. The problem is that they do not have electronic mailing lists for their discussions but rather use paper mail. Perhaps the X400OPS Group could offer to host distribution lists and archives for them?

- COSINE-MHS Project.

  The new service is called MHS Coordination Service. SWITCH got a contract by RARE to provide the staff and the required facilities. RARE is the umbrella organisation for research and academic networks in Europe. The contract will be moved to the Operational Unit, once they have started operation. The European members contribute to cost based on a cost key which is not fully decided yet, but key elements are the number of served organisations within a network and if the network uses the mapping tables or has registered a mapping. 50% of the total cost will be covered by the Commission of the European Countries. It is possible that networks which use the MHS Coordination Service outside Europe will be asked for a contribution as well.

  There was consensus that the Group support migrating the Cosine MHS community to the GO-MHS Community, in principle.

  Jim pointed out, and others agreed, that this migration must be thought out. Urs has to revise the Cosine MHS documents anyway, so he volunteered to write up a list of issues that are involved in the transition and send it to the list.

  There was a discussion about the relationship between the GO-MHS community and the MHS Coordination Service. People felt that a community is not "open" in the Internet sense if one must pay to join it. Erik Huizer pointed out that a GO-MHS coordination service must provide global coordination, regardless of financial contributions. He pointed out that there are U.S. Internet coordination services which provide international services but are funded solely by the U.S., including the NIC and the IANA (naming authority). There was a discussion that there is a difference between providing coordination and providing consulting and technical support. Erik was going to work to see that the MHS Coordination service is able to provide coordination for the GO-MHS community regardless of financial contributions. However, consulting and technical assistance may be limited to contributing members of the service.

- Internet Engineering Steering Group (IESG).

  Erik Huizer, IESG Application Area Director, reported on the decision to create an Application Area Directorate which will investigate overall concepts for networked applications. Priority is set for email and the character set issue.

  The document on Mail Based Servers is considered of broader interest and will be taken out of the X400OPS Group to be worked on by email experts for the SMTP and X.400 protocols. However, the current version will be used as a requirements document by the MHS Coordination Service.

Document Status

Editor's Note (md): A detailed listing of the status of the documents of the X400OPS Working Group is available via ftp under x400ops-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Action List

Jim Romaguera	To forward the questions from the ADMD operators to the X400OPS list.
	To update his document after a final call for comments and send it to the RFC Editor for consideration as an Informational RFC.
Erik Huizer	To work out the possibility of basing the coordination point for the GO-MHS community on a sound political basis.
	To make sure that documents from other groups (esp., RARE working groups) will have the appropriate group name in the ID filename instead of the author's surname.
MHS Coordination Service	To work out procedures for new members of the GO-MHS community to join the coordination service.

Alf Hansen	To send draft-ietf-x400ops-mgtdomains-ops-05.txt to the IESG secretariat with a copy to the Area Director, Erik Huizer, for consideration of the document as an Informational RFC.
Claudio Allocchio	To report to the X400OPS Group on the results of the discussion with IAB/IESG on the usage of DNS for mapping/routing tables.
	To update the documents draft-ietf-x400ops-dnsx400maps-02.txt and draft-ietf-x400ops-dnsx400rout-02.txt according to the results of the preceding action.
	To write a document on the implementation of the two proposals from the preceding action.
Urs Eppenberger	To wait for the decision of the IESG on his document, do the final modifications and send it to the RFC Editor.
Harald T. Alvestrand	To send draft-ietf-x400ops-charactersets-01.txt to the IESG with a copy to the Area Director, Erik Huizer for consideration of the document as a Proposed Standard.
Marko Kaittola	To update his document according to the comments received, choose a new title and prepare a new version for the next IETF meeting in Amsterdam.
Einar Stefferud	To integrate the final comments to the document draft-ietf-x400ops-admd-01.txt and then send it to the IESG for consideration as an Informational RFC.
Allan Cargille	To send draft-ietf-x400ops-postmaster-01.txt to the IESG with a copy to the Area Director, Erik Huizer for consideration of the document as a Proposed Standard.
All Group Members	To send comments to the mapping authority paper in a timely fashion to the list or to Jeroen Houttuin to allow him to create a new version for the next RARE WG-MSG meeting at JENC 93 in Trondheim.
Tony Genovese/Alf Hansen	To propose a revised version of the Charter.

Attendees

Claudio Allocchio	`Claudio.Allocchio@elettra.trieste.it`
Harald Alvestrand	`Harald.Alvestrand@delab.sintef.no`
C. Allan Cargille	`allan.cargille@cs.wisc.edu`

Kevin Carosso	kvc@innosoft.com
George Chang	gkc@ctt.bellcore.com
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Urs Eppenberger	eppenberger@switch.ch
Jeroen Houttuin	houttuin@rare.nl
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	Kevin.E.Jordan@cdc.com
Marko Kaittola	Marko.Kaittola@funet.fi
John Klensin	klensin@infoods.unu.edu
Mary La Roche	maryl@cos.com
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Bruce Mackey	brucem@cinops.xerox.com
Ignacio Martinez	martinez@rediris.es
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Edward Reed	eer@cinops.xerox.com
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Yzhak Ronen	y.ronen@homxa.att.com
Gary Rowe	gjrowe@attmail.com
Sue Smith	smiths.es.net
Einar Stefferud	stef@nma.com
Catherine Summers	cfs@cos.com
Louisa Thomson	louisa@whitney.hac.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Russ Wright	wright@lbl.gov

2.2 Internet Area

Directors (s):

- Stev Knowles: stev@ftp.com
- Dave Piscitello: dave@mail.bellcore.com

Area Summary reported by Stev Knowles/FTP and Dave Piscitello/Bellcore

Working Groups in the Internet Area are actively involved in the development of Internet standards for:

1. IP and multi-protocol operation over emerging wide area technologies (ATM, SMDS, Frame relay) and point-to-point technologies (including narrowband ISDN).

2. Expanded use of the IP backbone by tunneling other widely used network protocols (Appletalk, SNA).

3. Development of a "next generation" IP; i.e., a replacement protocol and addressing/routing architecture for IPv4.

4. Miscellany (Dynamic host discovery, and multicast interdomain routing).

The following Groups in the Internet Area met during the Columbus IETF:

- Birds-of-a-Feather (BOFs)
  - Net Support for Quality Of Service (QOS) and Real-Time Traffic
  - SNA Peer-to-Peer Networking

- Working Groups
  - Dynamic Host Configuration
  - IP Address Encapsulation
  - IP over AppleTalk
  - IP over Asynchronous Transfer Mode
  - IP over Large Public Data Networks
  - P. Internet Protocol
  - Point-to-Point Protocol Extensions
  - Simple Internet Protocol
  - TCP/UCP over CLNP-addressed Networks

Four candidates for IPng ("next generation", we are no longer referring to this as IPv7)–PIP, SIP, and TUBA/CLNP, ULLMAN/IPv7–provided plenary status reports and three provided demonstrations throughout the week from the terminal room provided by OARnet. Many thanks to the participants for their efforts and to OARnet for their support.

SNA Peer-to-Peer Networking BOF (SNAPPER)

The SNAPPER BOF was chaired by Wayne Clark. Three alternatives for handling SNA peer-to-peer (i.e., APPN) traffic across TCP/IP networks were presented:

1. APPN over TCP/IP
2. APPN/DLS using connection networks, and
3. APPI

Pros and cons of the three alternatives were discussed and it was decided that (a) the topic is too political at this point; (b) neither IBM or the APPI Forum is ready to change at this point in time, and (c) the APPN market at present is too small to worry about.

The conclusions of the BOF were:

1. A data link switching (DLS) working group would be very desirable..
2. Willingness to participate in the Working Group would be investigated.
3. If the conditions of (2) are met, the snapper@cisco.com mailing list would be used to develop a working group Charter and a more applicable name.

Dynamic Host Configuration Working Group (DHC)

The DHC Working Group discussed the interface between DHCP and the Domain Name System (DNS). Rob Austein explained that the DNS Working Group is currently developing an interface to DNS that considers the needs of DHCP.

Additionally, the Group discussed the possible use of SNMP with DHCP, to serve as a "second-level" bootstrap mechanism to transmit additional configuration parameters to a client. SNMP is not likely to be as useful as an implementation-specific interface for server management.

A technical problem was discovered with the current definition of the 'chaddr' field, which allows 'chaddr' to be used as either a hardware address or other unique identifier. As the 'chaddr' value must be used to return DHCP reply messages to the client, that field will be reserved for use strictly as a hardware address, and the client will be required to supply a unique identifier in a 'client identifier' option. This identifier will be a typed value with the same structure as defined for the 'chaddr' field.

Mike Carney and Jon Dreyer submitted a new definition for encapsulating vendor-specific options that the Working Group accepted with minor modifications.

Finally, the mechanism for determining the parameters returned to a particular client was discussed at length. No conclusion was reached at the meeting; an interim solution will be incorporated into the DHCP specification Internet-Draft to allow the protocol to move forward to Proposed Standard.

IP Address Encapsulation Working Group (IPAE)

The IP Address Encapsulation (IPAE) Working Group met once during the Columbus IETF. Since IPAE re-aligned itself to provide transition technology for SIP, the foci of IPAE discussions have been:

1. Modifications to SIP that are likely to facilitate transition, and
2. Implementation experiences with SIP/IPAE.

(Reference to IPAE usage is specifically to SIP-over-IP and SIP-mapped-to-IP. The first is to tunnel through the Internet and the second is to gateway to IPv4 hosts.)

The Working Group held a brief discussion about the SIP/IPAE demonstration slated for later in the week, discussing the implementations being shown and their use.

Bob Gilligan and Ron Jacoby discussed their implementation work. The Beame & Whiteside, Intercon, and Network General implementations, for DOS, Windows, MacIntosh, and network monitoring, respectively, were not available to make presentations.

Steve Deering raised the issue about whether SIP fragmentation is end-to-end only, or can occur en-route. He is interpreting the Group's response as supporting the position that SIP fragmentation need *not* occur en-route.

IP Over Asynchronous Transfer Mode Working Group (ATM)

The IP over ATM Working Group met for three sessions at the March IETF meeting.

Four presentations were given on ATM Addressing and ATM/Internet Address Resolution. Keith McCloghrie presented an overview of the ATM Forum addressing work. Subbu Subramaniam presented his ideas on how ATM address resolution should work. Fong-Ching Liaw presented pros and cons of subnet and peer model of address resolution in ATM networks, and Juha Heinanen presented an overview of his NBMA Address Resolution Protocol (NBMA ARP) proposal.

There was considerable discussion about how address resolution should work, and pros and cons of the subnet versus peer model. There were strong views on both sides. The session ended with suggesting that neither one approach would prevail and proposed mechanisms for combining the two approaches. Mark Leabach presented slides with his thoughts on which areas the Working Group should focus on first. He saw two approaches in the Working Group: Functional layerists (ATM as a wire-replacement) versus ATM IP Morphist. He made a strong argument that the Working Group should first specify how IP over ATM should work in the "classical IP" mode where ATM networks are connected by routers. Mark went on to present a list of problems which need to be solved.

Tim Salo presented a talk on the Gigabit Testbed. The goal of the project is to create a seamless connection between ATM LAN's and ATM WAN's. His preliminary observations were that there are no complete proposals available today, some are only slightly explored or identified, much new functionality must be implemented, and most of the problems are in the wide area.

Tim's final observations were that we need a complete solution if ATM is going to be the solution. It is important to avoid making the customer the system integrator. Significantly more implementation experience is needed.

Ramon Caceres presented the results of his simulation of different approaches for virtual circuit (VC) multiplexing. His conclusions were that one VC per connection gives the best performance, followed by one VC per type of traffic (e.g., telnet, ftp, mail, etc.). One VC per router pair gives poor performance. The overall goal should be to separate traffic as much as possible.

After a final discussion of subnet versus peer addressing models, the Working Group moved on to a discussion of important areas to pursue and the assignment of action items to complete.

Joint IP over Large Public Data Networks Working Group (IPLPDN) and Point-to-Point Protocol Extensions Working Group (PPPEXT)

The IP over Large Public Data Networks (IPLPDN) and PPP Extensions (PPPEXT) Working Groups met in joint session to discuss protocol specifications common to both. Since the objectives and requirements of the two Working Groups differ in some key respects, there was considerable difference of opinion at the outset.

Two subjects were discussed: how to share load among a set of parallel links to increase apparent bandwidth and potentially reduce latency between two sites, and how the IPLPDN group might best avail itself of the facilities found in PPP negotiation. Both Fred Baker and Keith Sklower had proposals though the consensus was that Keith's approach was preferred, but required some modifications. Keith will appropriately edit his proposal for further discussion on the IPLPDN mailing list.

The two Groups also discussed PPP Parameter Negotiation for Frame Relay. Consensus was reached on several issues. Keith Sklower and Bill Simpson have agreed to merge their efforts and their current proposals to implement this consensus and will co-author a "Parameter Negotiation over Frame Relay and X.25" document. The output will be discussed on the IPLPDN list.

IP over Large Public Data Networks Working Group (IPLPDN)

The revised draft of RFC1294, "Multiprotocol over Frame Relay," was approved for submittal to the IESG for release in a new RFC and for advancement from Proposed to Draft standard.

RFC1356, "Multiprotocol over X.25," was reviewed for advancement in status. It was agreed to perform tests of interoperability of implementations. The revised RFC1315, the Frame Relay MIB document, was discussed. It was agreed to keep this document as the "DTE MIB" and that the new work on a Frame Relay MIB would become the "DCE MIB."

The Charter was discussed and the Chair agreed to talk with the Routing and Internet Area Directors and working group Chairs about the transfer of open issues to other groups.

Work progressed on the "Multiprotocol over Circuit ISDN" document. The draft was re-titled "Encapsulation Determination for Circuit Switched Services." Keith Sklower will incorporate comments and will distribute the revised document by email for Working Group approval for submittal to the IESG.

The P. Internet Protocol Working Group (PIP)

The PIP Meeting was tutorial in nature. Paul Francis (formerly Tsuchiya) covered various aspects of the PIP protocol, starting with the header format and going through addressing. No serious problems were uncovered in the discussions, though Steve Deering did uncover a bug in the proposed caching mechanism.

Attendees were invited to see a demonstration of PIP during the meeting.

Simple Internet Protocol Working Group (SIP)

The SIP Group discussed clarifications and changes to the SIP specification that have occured since last November. The most significant changes were the addition of hop-by-hop options and the specification of "local-use" SIP addresses that hosts can fabricate from 802.11 addresses for the purpose of auto-configuration. The Group also discussed a tentative SIP Sensitivity Label Option and a SIP End-to-End Security Option, both based on recent work on IPv4 security.

The SIP Group also discussed proposed modifications to RIP-2, OSPF, and IDRP to support SIP routing, as documented in recent Working Group drafts. Finally, the Group received a status report on SIP "host routing" work-in-progress.

TCP/UDP over CLNP-addressed Networks Working Group (TUBA)

The TUBA Working Group met to discuss the following topics:

- Implementation Status and Demonstration.
- Document Status.
- Priortization of TUBA Work.
  - Questions asked at Opening Plenary
  - Dynamic Host Address Assignment
  - Mobile Hosts

  - Routing and Addressing Plan
  - Transition Strategies
  - Discussion of Technical Advantages of CLNP
- Demo and Implementation Targets.

Editor's Note (md): A detailed summary of each topic is provided in the Minutes which follow this Area Report.

CURRENT MEETING REPORT

Minutes of the Net Support For QOS and Real-Time Traffic Working Group (RTQOS)

Report not submitted.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Jules Aronson	aronson@nlm.nih.gov
Randall Atkinson	atkinson@itd.nrl.navy.mil
Karl Auerbach	karl@empirical.com
William Barns	barns@gateway.mitre.org
Nutan Behki	Nutan_Behki@qmail.newbridge.com
Tom Benkart	teb@acc.com
Monroe Bridges	monroe@cup.hp.com
David Bridgham	dab@epilogue.com
Scott Brim	Scott_Brim@cornell.edu
Al Broscius	broscius@bellcore.com
Theodore Brunner	tob@thumper.bellcore.com
Randy Butler	rbutler@ncsa.uiuc.edu
Ramon Caceres	ramon@mitl.com
Jeff Carman	tcarman@bnr.ca
Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@hpl.hp.com
William Chimiak	chim@relito.medeng.wfu.edu
George Clapp	clapp@ameris.center.il.ameritech.com
Wayne Clark	wclark@cisco.com
Richard Cogger	R.Cogger@cornell.edu
Robert Cole	rgc@qsun.att.com
Simon Coppins	coppins@arch.adelaide.edu.au
Dave Cullerot	cullerot@ctron.com
James Davin	davin@thumper.bellcore.com
Steve Deering	deering@parc.xerox.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Tony DeSimone	tds@hoserve.att.com
Taso Devetzis	devetzis@bellcore.com
Chas DiFatta	chas@cmu.edu
Kurt Dobbins	kurtdob@ctron.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Robert Enger	enger@reston.ans.net
Deborah Estrin	estrin@isi.edu

Antonio Fernandez	afa@thumper.bellcore.com
Michael Fidler	fidler@mitre.org
Ron Frederick	frederick@parc.xerox.com
Eugene Geer	ewg@cc.bellcore.com
Robert Gilligan	gilligan@Eng.Sun.Com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Fengmin Gong	gong@concert.net
Kenneth Goodwin	goodwin@a.psc.edu
Ramesh Govindan	rxg@thumper.bellcore.com
Mark Green	markg@apple.com
John Hascall	john@iastate.edu
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Don Hoffman	hoffman@eng.sun.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Kathy Huber	khuber@bbn.com
Steven Hubert	hubert@cac.washington.edu
John Ioannidis	ji@cs.columbia.edu
Phil Irey	pirey@relay.nswc.navy.mil
David Jacobson	dnjake@vnet.ibm.com
Ronald Jacoby	rj@sgi.com
David Johnson	dbj@cs.cmu.edu
Matthew Jonson	jonson@server.af.mil
Phil Karn	karn@qualcomm.com
Frank Kastenholz	kasten@ftp.com
Michael Khalandovsky	mlk@ftp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
Stev Knowles	stev@ftp.com
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Giri Kuthethoor	giri@ms.uky.edu
Mark Laubach	laubach@hpl.hp.com
Tony Li	tli@cisco.com
Hock-Koon Lim	lim@po.cwru.edu
Frank Liu	fcliu@pacbell.com
Charles Lynn	clynn@bbn.com
Glenn Mackintosh	glenn@canet.ca
Carl Madison	carl@startek.com
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@cmf.nrl.navy.mil
Matt Mathis	mathis@a.psc.edu
Keith McCloghrie	kzm@hls.com
Donald Merritt	don@arl.army.mil
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk

Paul Milazzo	milazzo@bbn.com
Robert Mines	rfm@sandia.llnl.gov
Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Erik Nordmark	nordmark@eng.sun.com
Zbigniew Opalka	zopalka@agile.com
Ayal Opher	aopher@synoptics.com
Joe Pagan	jrp@afterlife.ncsc.mil
Joseph Pang	pang@bodega.stanford.edu
Brad Parker	brad@fcr.com
Michael Patton	map@bbn.com
Geir Pedersen	Geir.Pedersen@usit.uio.no
John Penners	jpenners@advtech.uswest.com
Maryann Perez	perez@cmf.nrl.navy.mil
Charles Perkins	perk@watson.ibm.com
Drew Perkins	ddp@fore.com
James Philippou	japhilippou@eng.xyplex.com
Robert Reschly	reschly@brl.mil
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Shawn Routhier	sar@epilogue.com
Allan Rubens	acr@merit.edu
Michael Safly	saf@tank1.msfc.nasa.gov
Hal Sandick	sandick@vnet.ibm.com
Shiva Sawant	shiva@synoptics.com
Andrew Schmidt	ags@uius.edu
Eve Schooler	schooler@isi.edu
Tim Seaver	tas@concert.net
Kanan Shah	kshag@cmf.nrl.navy.mil
William Simpson	Bill.Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Karen Sollins	sollins@lcs.mit.edu
Marco Sosa	mxs@sabre.bellcore.com
Martha Steenstrup	msteenst@bbn.com
Stuart Stubblebine	stubblebine@isi.edu
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Wayne Tackabury	wayne@cayman.com
Sally Tarquinio	sallyt@gateway.mitre.org
John Tavs	tavs@vnet.ibm.com
Susan Thomson	set@bellcore.com
Claudio Topolcic	topolcic@cnri.reston.va.us
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Mario Vecchi	mpv@thumper.bellcore.com

Dinesh Verma	verma@watson.ibm.com
Warren Vik	wmv@lachman.com
Curtis Villamizar	curtis@ans.net
John Vollbrecht	jrv@merit.edu
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Abel Weinrib	abel@bellcore.com
Von Welch	vwelch@ncsa.uiuc.edu
Rick Wilder	wilder@ans.net
Douglas Williams	dougw@ralvmg.vnet.ibm.com
Steven Willis	steve@wellfleet.com
Linda Winkler	lwinkler@anl.gov
Jane Wojcik	jwojcik@bbn.com
Richard Woundy	rwoundy@vnet.ibm.com
Russ Wright	wright@lbl.gov
John Wroclawski	jtw@lcs.mit.edu
Yow-Wei Yao	yao@chang.austin.ibm.com
Chin Yuan	cxyuan@pacbell.com
Paul Zawada	Zawada@ncsa.uiuc.edu

ReSerVation Protocol (RSVP)

Xerox PARC: Lixia Zhang, Steve Deering, Scott Shenker
And
USC/ISI: Bob Braden, Deborah Estrin, Sugih Jamin, Danny Mitzel, Daniel Zappala

Reservation setup protocol needed to create and maintain resource reservations in switches.

Motivated by: Demands of multipoint, real-time applications (e.g., vat, nv).

Design goal is to accommodate:

1. **heterogeneous receivers.**
2. **large and dynamic groups.**
3. **different reservation styles** for various applications (e.g., mixing audio sources vs. distinct video sources).
4. **channel selection** among different sources during a session.
5. **adapt with dynamic** underlying routing, without replicating routing function.
6. **contain protocol overhead** for large groups.
7. **heterogeneous underlying** technologies (e.g., routing and scheduling algorithms).

These goals are not met by existing (conventional) setup protocols such as ST-II, IDPR, X.25...

Design Principles

RSVP designed according to the following principles:

1. **Receiver initiated reservation.**
2. **Separate reservation from packet filtering**
3. Support different reservation styles
4. Maintain soft-state in network switches
5. Separate reservation from routing
6. Adjustable protocol overhead
7. Modularity

1. Receiver Initiated Reservation

- Reservation requirement depends on constraints specific to receiver: host, bandwidth capacity, price sensitivity.
- Receiver "experiences" and "pays for" reservation...should control it.
- RSVP Receiver chooses reservation parameters (a flowspec less than or equal to source's)
- RSVP Receiver keeps reservation active as long as desired.

2. Separate Reservation from Packet Filtering

Resource reservation specifies what amount of resources is reserved for whom (receiver(s) that own/control reservation).

Packet filters specify which packets may use the reserved resource.

Filter may be changed without changing the reservation; called dynamic filtering.

The decoupling of reservation and filter enables us to offer different reservation styles.

Protocol Overview

Each data source periodically sends Path message to establish or update Path state in switches.

- Path message carries source's flowspec and filter flag.

- Path state contains
  a) incoming link upstream to source,
  b) outgoing links to downstream receivers (as indicated by routing table), and
  c) source ID and previous hop kept only for filtered reservations.

- Forwarding of messages is based on routing table entry for destination address (unicast or multicast).

Protocol Overview (cont.)

Each receiver periodically sends Reservation message to establish or update Reservation state in switches.

- Reservation message carries flowspec, reservation style, and optional packet filter.

- Reservation state contains
  a) amount of resource reservation,
  b) source filter,
  c) reservaton style, and
  d) receiver ID kept only for dynamic-filter reservations.

- Forwarding of reservation messages uses inlink information (if no-filter) or previous-hop (if filtered); i.e., reverse paths used so that reservation occurs on path used by sources' packets; since routes may not be symmetric.

When underlying routing or membership changes, next path messages will follow new route (or branches) and next reservation messages will travel up the new route.

Path and Reservation message types carry timeout values and intermediate switches set local timers accordingly.

All RSVP messages sent as datagrams.

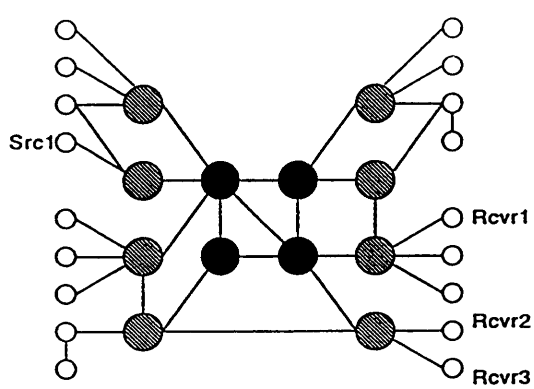

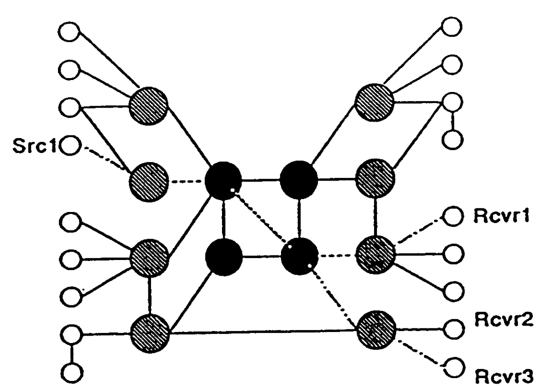

--------- Multicast distribution tree installed by routing from Src1 to Rcvr 1,2,3 (in routing tables of nodes)

Src1's Path message sent as datagram along multicast tree.

mcast_addr: outlinks, inlinks, source (if filtered) saved in reservation state.

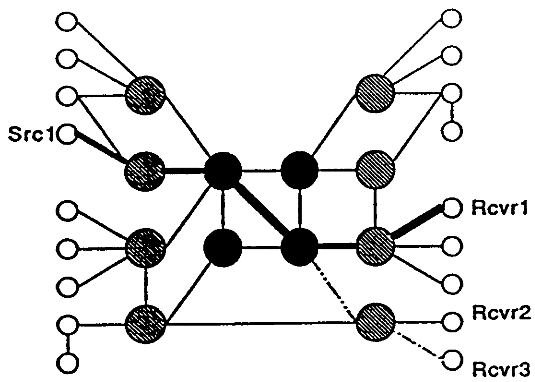

Rcvr1 initiates Reservation.
Message traverses route to Src1 using mcast_addr to inlinks mapping.

Processed by admission control.

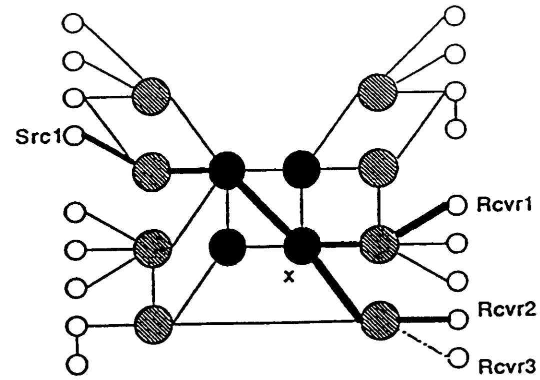

Rcvr2 initiates Reservation.
Message traverses route to Src1 using mcast_addr to inlinks mapping.

At node X, reservation is merged with existing reservation and message is not propagated further.

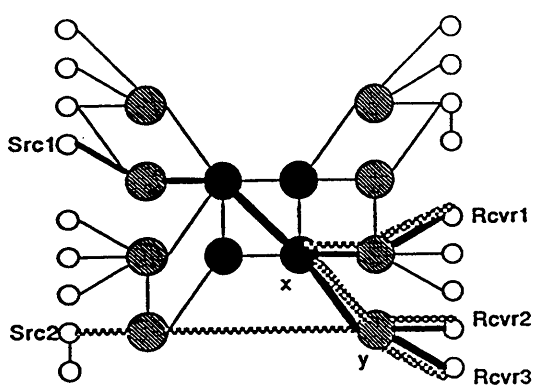

Src1, Src2 send periodic Path messages to capture change in topology/group

Rcvr1,Rcvr2,Rcvr3 send periodic Reservation messages to keep reservation active. Reservations merged at branch point

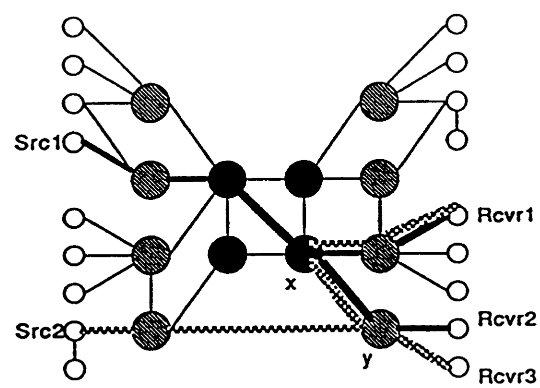

If Rcvr2 and Rcvr3 can each only handle a single live flow at any one time, they each reserve for 1 flow, indicating Dynamic Filter style. Each sends Reservation messages with desired Filter to node y when they want to select a different source.

Status

Simulator written by Lixia Zhang. Modules imitate actual behavior of network hosts, links, routers, protocols.

Zhang verified RSVP design by implementing protocol in simulator and observing, step-by-step, how the protocol handles various dynamic events.

RSVP prototype currently being implemented in Dartnet kernel by Sugih Jamin (USC).

Development will continue at ISI and Xerox, including integration and testing with real time protocol support developed by MIT and Xerox.

Also conducting larger, call-level, simulations to verify and quantify scaling properties, and to investigate design tradeoffs.

Open issues include: interaction/interface to routing, consideration of other reservation styles, caching/aggregation strategies.

3. Support different reservation styles

Some applications can share a reservation among multiple (all) sources.

- Audio sources can be mixed; small upper bound on number of simultaneous speakers
- No filter needed for these reservations.

Other applications require distinct reservation per source.

- Each video stream carried requires its own reservation.
- Requires filter for each reservation.
- Filter may be shared by downstream receivers.

Possible limitations (e.g., BW, $) on number of live, simultaneous, sources carried on a link/route to a receiver.

- Allow receivers to select which subset of sources to receive.
- Support dynamic channel selection during a call: fast and low cost and without risk of call blocking.
- Requires dynamic filter for each reservation to specify different sources over time; filters not shared by downstream receivers.

4. Avoid Hard-State Management in Routers

Avoid router state that cannot be recreated transparently to higher level protocol.

Leave maintenance (although not storage) of state to endpoints.

Endpoints periodically refresh reservation state

- Source initiated refreshes catch routing distribution tree changes (after failure or new membership)
- Receiver initiated refreshes reconfirm reservation; also reestablish reservation along new branches of tree.
- Switches timeout reservations to reclaim resources.

5. Decouple Reservation from Routing

Establish reservation over distribution tree independent of how tree was created

- IP unicast route determined by Intra- or Inter-domain routing.
- Multicast tree established by IP Multicast Protocol, distributed membership-control, or by CBT or some other inter-domain multicast routing protocol.

RSVP provides common reservation mechanism for both unicast and multicast, large and small groups, intra- and inter-domain.

Common routing mechanisms for reservation and non-reservation traffic

- Routing decision made only once per flow.
- But packet and flow routing can still share the product (RIB) of a single underlying routing protocol.
- Make use of existing adaptive routing capabilities.

6. Adjustable Protocol Overhead

Adjust refresh interval of control messages as group size grows; tradeoff between protocol overhead and state overhead.

Merge control messages as they travel up and down tree.

Tradeoff also exists between protocol/state overhead and resource overhead.

7. Modularity

RSVP interfaces with applications (or session protocols), admission control, and routing.

- Applications give flowspecs to RSVP.
- Admission control makes acceptance decision based on flowspec carried by RSVP.
- Routing forwards RSVP messages.

An Architecture for Resource Management in Networks

David D. Clark
M.I.T. Laboratory for Computer Science

Joint work with
lots of others...

Who...

Scott Shenker, Lixia Zhang and Steve Deering
- Xerox PARC

Deborah Estrin, Sugih Jamin
- USC

Craig Partridge
- BBN

Chuck Davin, Andrew Heybey and others
- MIT

Members of the End-2-End Research Group:
Bob Braden, Van Jacobson, Sally Floyd

Abhay Parekh
- MIT and IBM

What are the key questions?

What set of services can we implement?
- Technical question: theory and practice.

What set of services are needed?
- What do applications actually require?
- What can be marketed?

How do we write standards?
- Specify mechanism.
- Specify objective.

The parts of the solution

- Queue management – the per-packet processing to support the services.
- Resource setup protocol – propagate a request across the network.
- Route establishment – selecting suitable links.
- The "flow spec" – the description of the needed service.
- Admission control – the algorithm to run at each switch to determine if the request can be met.
- Packet classifier – map packets to resource class.

What are the service requirements?

We identify two kinds of service requirements:

Performance requirements of individual applications.
- Relate to delay objectives.
- Real time, interactive, bulk data.
- Ergonomic requirements.

Controlled sharing among traffic aggregations.
- Relates to bandwidth.
- Institutions, protocol families, applications.
- Economic requirements.

WHAT IS REAL-TIME?

Real-time traffic is characterized by a "playback" point.

- Packets "on-time" or early are useful.
- Late packets are not useful.

Note: not the same as interactive traffic.

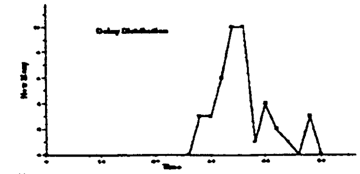

Playback time?

Loss % ?

SUPPORTING REAL-TIME APPLICATIONS

THE PLAYBACK POINT

Where does it come from?

Simple model (guaranteed service):
- Negotiated at service setup time.
- Network assures it.

Alternate model (predictive service):
- Receiver observes arriving traffic.
- Receiver adapts to observed arrival time.
- Network provided stable service.

A receiver that adapts must tolerate some late packets. Predictive service will give lower bounds.

SUPPORTING REAL-TIME APPLICATIONS

BOUNDS ON DELAY

Is there a guaranteed bound for queueing delay?

A recent result due to Parekh: yes!

If:
- Traffic conforms to "Token Bucket" at source.
- Switches use Weighted Fair Queueing service.

Then there is an upper bound to delay.
- Since we can achieve the bound in practice, it is the best bound we will get.

SUPPORTING REAL-TIME APPLICATIONS

DEFINITIONS

Token bucket traffic source:

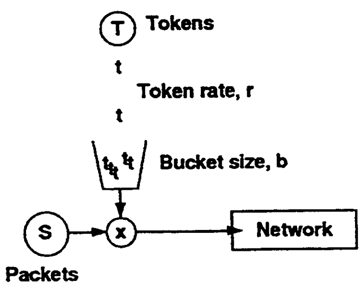

SUPPORTING REAL-TIME APPLICATIONS

DEFINITIONS

Weighted Fair Queueing service:

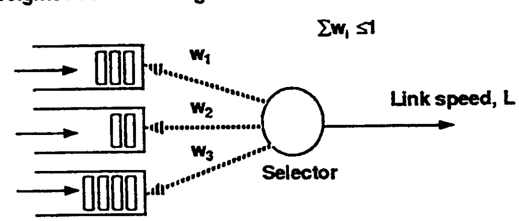

Selector: for each queue, compute the departure time of head packet at rate $w_i \cdot L$. Choose the packet that would leave first and send it.

SUPPORTING REAL-TIME APPLICATIONS

WHY IS THERE A GUARANTEED BOUND?

This scheme produced a bound because:

- WFQ isolates one flow from another.
- The bucket size limits the burst that can enter the network.

SUPPORTING REAL-TIME APPLICATIONS

CAN WE REDUCE DELAY FURTHER?

The key to the guaranteed bound was isolation.

**Isolation provides a guarantee.
It does not minimize delay.**

Consider the result of putting a burst into the network.
- My burst causes me delay. Others are unaffected.

What if several sources shared the delay.
- Higher average delay. Lower maximum delay.

**Isolation bounds delay and provides protection.
Sharing allocates delay.**

SUPPORTING REAL-TIME APPLICATIONS

AN EXAMPLE

Consider a number of sources with equal delay requirements.

If each packet has the same target delay bound in the switch, what queueing discipline results?

– FIFO! (Some find this surprising...)

FIFO is poor at isolation, good for sharing.

Service	Mean	99.9%tile
WFQ	3.16	53.86
FIFO	3.17	34.72

Times in ms. (packet transmission time). 83.5% loading.

R/T 13

SUPPORTING REAL-TIME APPLICATIONS

OUR SIMULATION MODEL

Two-state (on-off) sources. Peak rate twice average.

Shaped by token bucket (A,50). Excess was dropped.

Link bandwidth 1 Mb/s.

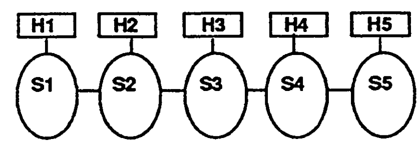

10 Flows per link. 12 of length 1, 4 of length 2, 4 of length 3, and 2 of length 4. 22 total.

R/T 14

SUPPORTING REAL-TIME APPLICATIONS

ADAPTIVE RECEIVERS

FIFO produces lower delays than WFQ.

The delay is not known precisely in advance.

The receiver must track the arriving traffic and adapt.

This is the <u>predictive</u> real-time service.

FIFO is only one method for achieving predictive service.
- Focus on service objective, not just implementation.
- Consider other sharing methods – priority.

R/T 15

SUPPORTING REAL-TIME APPLICATIONS

Admission control

If we are to meet specified service requirements, then:

- We must have a characterization of the offered load.
  - Token bucket, etc.
- We must have an algorithm to decide if the request can be met.
  - Admission control.
- We must be able to refuse the request.

R/T 16

SUPPORTING REAL-TIME APPLICATIONS

Other application requirements

Real-time traffic has playback point.

<u>Elastic</u> traffic can tolerate variation in delay.
- Continuous curve of user satisfaction.
- Wide range of delay objectives.
  - Can be as demanding as real-time.
- Much harder to characterize offered load.
  - Most systems today do not try.
  - No admission control.
- No attempt to characterize the service.
  - Best effort delivery.
  - Manage aggregates of sources.
- Management of QOS in background.
  - Purchase larger share, install more fiber.
  - Like highway system.

R/T 17

SUPPORTING REAL-TIME APPLICATIONS

Link Sharing

Controlled sharing:
- Among customers
- Among protocols

Want overload assurance <u>and</u> statistical sharing.

Worst-case objective: nested sharing.

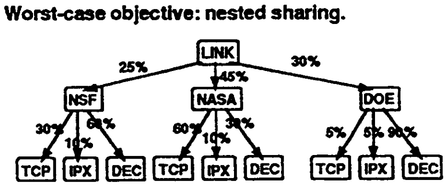

R/T 18

Summary: Service Requirements

Ergonomic:
- Providing real-time service
  - Guaranteed and predictive
- Providing lower delays for more demanding elastic services.

Economic:
- Provide isolation among protocol families.
  - Not just economic, but technical
- Provide isolation among customers.
- Achieve effective link utilization.
  - Aggregation of traffic.
  - Statistical sharing or real low loading.

Per-packet scheduling

We need one scheme that accomplishes all the goals:
- Ergonomic :guaranteed real time, predictive real time, and a range of elastic objectives.
- Economic: nested link sharing.

How can these be combined?
- Need an ordering on the requirements.

Real time "comes first".
- Delay of each individual packet matters.
- Admission control can achieve link sharing.

Guaranteed service takes precedence over predictive.
- Need not do better than guarantee.

Link sharing takes precedence over elastic objectives.
- No admission control.

The mechanism

Guaranteed R.T. requires isolation:
- -> WFQ or some similar scheme.
  - Our current scheme is "Stalled Virtual Clock".

Predictive R.T. requires a "delay sharing" method:
- -> FIFO mixing of flows.
- -> Simple priority based on delay objective.

Elastic service requires another isolation mechanism:
- -> Need not be precise for each packet.
- -> Needs to be nested.
- -> Our scheme: Nested Approximate WFQ.

Real time packets are scheduled without regard to link sharing objectives.
- -> Appropriate node in link sharing structure is "charged" for R.T. packets.

Final queueing picture

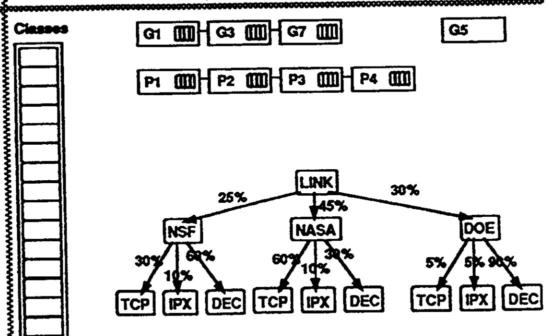

A further improvement to predictive service

Multiple switches seem to make the delay variance worse.
- But multiple switches should increase the chance for sharing.

Can we control this to reduce delay variance?
- FIFO+

For each class, compute the average delay at each switch.

For each packet, track the actual total delay.

Schedule each packet sooner or later based on actual delay.

Some simulation results

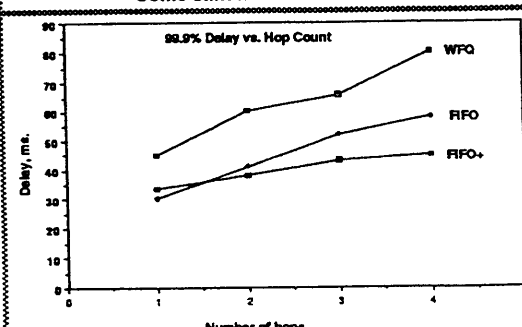

SUPPORTING REAL-TIME APPLICATIONS

The parts of the solution

- Queue management -- the per-packet processing to support the services.
- Resource setup protocol -- propagate a request across the network.
- Route establishment -- selecting suitable links.
- The "flow spec" -- the description of the needed service.
- Admission control -- the algorithm to run at each switch to determine if the request can be met.
- Packet classifier -- map packets to resource class.

R/T 25

SUPPORTING REAL-TIME APPLICATIONS

The FLOW SPEC

Flow Spec -- a specification of the needed service.
- Internet RFC by Craig Partridge

Bandwidth:
- Average rate.
- Token bucket size.
- Peak rate.

Delay:
- Target max delay. (Not a hard limit.)
- Max jitter. (We could argue about this one.)

Sort of service:
- Guaranteed or Best Effort.

Errors:
- Approx. maximum tolerable.

R/T 26

SUPPORTING REAL-TIME APPLICATIONS

Admission Control

An algorithm run at each node to see if the requested service can be accepted.

Key issue: how to characterize the existing traffic.
- From the flow spec.
- From actual behavior.

We propose to measure actual usage.
- MUCH better link utilization.
- Risk of being misled.
- Presume "well-behaved" sources.
- Also sell the guaranteed service.

Any model like token bucket will (must be) a very loose bound. A tight model is much to complex.
- Consider video.

R/T 27

SUPPORTING REAL-TIME APPLICATIONS

The Algorithm, Roughly

When a new guaranteed flow is requested:
- Check to see if there is enough bandwidth.

When a new predictive flow is requested:
- Find the proper delay class.
- See if the sum of measured bandwidth plus proposed addition will fit.
- See if the measured delay bound plus the proposed bucket size will fit in delay target of class.
- See if any other predictive delay class will be disrupted.

Results: (VERY preliminary). For a mix of delay targets, link loadings of >85% are possible.

R/T 28

SUPPORTING REAL-TIME APPLICATIONS

Packet Classifier

Two issues: efficiency and functionality.

Efficiency:
A flow id might be better but an IP header is good enough. (!)
- look at source and destination IP address plus the port numbers
- fewer bits than a CLNP header....

Functionality:
- Support end-node and network management setup.
  - Current battle: ports vs. TOS field.
- Aggregate traffic into flows.

Important issue for next IP.
- I like concept of flow ID.

R/T 29

SUPPORTING REAL-TIME APPLICATIONS

Summary -- project status

Per-packet scheduling: second generation implementation exists, and has been simulated. Must demonstrate efficiency.

Flow spec: Internet RFC published.

Admission control: preliminary paper published, more testing and evolution needed.

Resource setup: Algorithm has been implemented and tested in simulation. Currently being integrated into scheduling code.

Packet classifier: Waiting for code from Van Jacobson.

Overall IP test: DARTnet demo in spring/summer 1993.

R/T 30

What service is needed?

An application example: guaranteed vs. predicted real-time bounds.

- Are we selling a characterized service or a good one?
  - Predicted RT, delay targets, aggregated classes.

A marketing example: commercial services.

- This discussion of real-time suggests fine grained reservations.
- Customers want predictable rates.
  - -> Aggregation for charging?

Two key questions?
- Can one create a cost-effective switched service?
- How does one add QOS to this service?

How will this be standardized?

Traditional model: standardize the formats and algorithms.
- Our queue scheme.

But...
There may be several valid ways to achieve the same results.
- Low link loading.
- Alternative realizations in different switch designs.

Alternative model: functional or performance spec.
- Building standards have both.
- Networks: routing standards?

Performance specs are HARD TO WRITE.

An approach to standards

Queue management -- functional requirement.

Explicit delay management (FIFO+) -- format/semantics definition.

RSVP -- traditional protocol specification.

Flow spec -- format definition.

Packet classifier -- global agreement on format.

Admission control -- private matter.

Route setup -- (speculative)
- Link state parameters -- format/semantics def.
- Route selection algorithm -- private matter.

CURRENT MEETING REPORT

Reported by Zbigniew Kielczewski/Eicon Technology

Minutes of the SNA Peer-to-Peer Networking BOF (SNAPPER)

The SNAPPER BOF was held on April 1st in Columbus (Ohio) during the 26th IETF meeting and was chaired by Wayne Clark.

Agenda

- Introduction
- Presentations:
  - Connection Networks
  - Data Link Switching
  - APPN TCP/IP
  - APPN, DLS and Connection Networks
  - APPI
- Discussion

The suggestion was made that it could be important and interesting to work on the standard way of transporting APPN over IP. One of the propositions was the DLS specification which could become the object of future IETF work.

The issue of coordination of efforts to merge APPN and APPI technologies was raised. It was clear from statements of interested parties that this effort involves not only technical issues, but also political ones, and the IETF would not be able to solve those complicated problems.

A second proposal was that all interested people could work on the standardization of the Data Link Switching which presents more promising and realistic targets. The APPN versus APPI issues should not be directly addressed by this work but both involved parties should be present and contribute to the work. The consensus was that this proposition could be used to form a working group.

Wayne will submit a draft Charter to the SNAPPER mailing list for discussion.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Fred Bohle	fab@interlink.com
John Chang	jrc@uswest.com
George Clapp	clapp@ameris.center.il.ameritech.com
Wayne Clark	wclark@cisco.com

Kurt Dobbins	kurtdob@ctron.com
Roger Fajman	raf@cu.nih.gov
Cleve Graves	cvgpc@oc.com
Ken Hayward	Ken.Hayward@bnr.ca
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
David Jacobson	dnjake@vnet.ibm.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Deirdre Kostick	dck2@sabre.bellcore.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Tony Li	tli@cisco.com
Baktha Muralidharan	murali@smaug.enet.dec.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
John Tavs	tavs@vnet.ibm.com
Stephen Tsun	snt@3com.com
Hung Vu	hungv@fonorola.com
James Watt	james@newbridge.com
Les Wibberley	lhw24@cas.org
Douglas Williams	dougw@ralvmg.vnet.ibm.com

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 to investigate network configuration and reconfiguration management and 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 — Write a bootp extensions document.

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

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

Done — Suggest new mechanisms to manage the information identified by Objective 1.

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

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-06.txt, .ps>

"DHCP Options and BOOTP Vendor Extensions", 06/30/1992, S. Alexander, R. Droms <draft-ietf-dhc-options-03.txt>

"Interoperation Between DHCP and BOOTP", 06/30/1992, R. Droms <draft-ietf-dhc-between-bootp-03.txt>

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell

Minutes of the Dynamic Host Configuration Working Group (DHC)

The Working Group discussed the interface between DHCP and the Domain Name System (DNS). DHCP needs an interface that can allow dynamic updates to DNS entries in response to dynamic allocation of DNS names to DHCP clients. Rob Austein explained that the DNS Working Group is currently developing such an interface to DNS that considers the needs of DHCP.

The Working Group discussed the possible use of SNMP with DHCP. SNMP may be useful as a "second-level" bootstrap mechanism to transmit additional configuration parameters to a client. SNMP is not likely to be as useful as an implementation-specific interface for server management. SNMP is an interesting candidate for the server-server protocol, as it may provide the semantics and data representation tools required for exchange of DHCP binding information between servers.

The Working Group discovered a technical problem with the current definition of the 'chaddr' field, which provides for use of 'chaddr' as either a hardware address or other unique identifier. As the 'chaddr' value must be used to return DHCP reply messages to the client, that field will be reserved for use strictly as a hardware address, and the client will be required to supply a unique identifier in a 'client identifier' option. This identifier will be a typed value with the same structure as defined for the 'chaddr' field.

Mike Carney and Jon Dreyer submitted a new definition for encapsulating vendor-specific options that the Working Group accepted with minor modifications. In the accepted definition, the 'vendor- specific information' option will include an initial value that identifies how to interpret the contents of the option, and other DHCP options, encoded in the same format as the current variable- length DHCP options. The initial identifying values will be centrally administered to avoid conflicts. One identifying value will be reserved for local use.

The mechanism for determining the parameters returned to a particular client was discussed at length. The focal points of the discussion were the ways in which a client can identify its characteristics ('client type' option) and the rules by which a server can use those characteristics to choose the information to be returned to a host. No conclusion was reached at the meeting; an interim solution will be incorporated into the DHCP specification Internet-Draft to allow the protocol to move forward to Proposed Standard.

Attendees

Kannan Alagappan	kannan@dsmail.lkg.dec.com
Steve Alexander	stevea@lachman.com

Philip Almquist	almquist@jessica.stanford.edu
Robert Austein	sra@epilogue.com
John Boatright	bryan_boatright@ksc.nasa.gov
Gregory Bruell	gob@wellfleet.com
Ralph Droms	droms@bucknell.edu
Robert Gilligan	gilligan@Eng.Sun.Com
Richard Harris	rharris@atc.boeing.com
John Hascall	john@iastate.edu
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@wco.ftp.com
Frank Kastenholz	kasten@ftp.com
Mark Kepke	mak@fc.hp.com
Andrew Knutsen	andrewk@sco.com
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Paul Lustgraaf	grpjl@iastate.edu
Kent Malave	kent@bach.austin.ibm.com
Glenn Mansfield	glenn@aic.co.jp
Evan McGinnis	bem@3com.com
William Nowicki	nowicki@legato.com
William Owens	owens@acsu.buffalo.edu
Charles Perkins	perk@watson.ibm.com
Steven Richardson	sjr@merit.edu
Shawn Routhier	sar@epilogue.com
Jon Saperia	saperia@lkg.dec.com
Chris Shaw	cshaw@banyan.com
Marek Tomaszewski	marek@net.com
Walter Wimer	walter.wimer@andrew.cmu.edu

2.2.2 IP Address Encapsulation (ipae)

Charter

Chair(s):

Dave Crocker, dcrocker@mordor.stanford.edu

Mailing Lists:

General Discussion: ip-encaps@sunroof.eng.sun.com
To Subscribe: ip-encaps-request@sunroof.eng.sun.com
Archive: parcftp.xerox.com:~/pub/ip-encaps/

Description of Working Group:

The IPAE Working Group seeks to develop a capability for extending IP to support larger addresses while minimizing impact on the installed base of IP users. An enhancement to the current system is mandatory due to the limitations of the current 32 bit IP addresses. IPAE seeks to upgrade the current system, rather than to replace the Internet Protocol. The approach taken will be to sandwhich a small addressing layer, above IP but below TCP or UDP, with the new layer having its own IP Protocol-ID. This special layer will thereby encapsulate new, larger, globally-unique addresses for source and destination, as well as any other fields of information that are considered essential.

The specificaton effort will attend to issues of transition and coexistance, among unmodified "IP" hosts and hosts which support "IPAE" hosts The IPAE approach will develop a framework to organize the Internet into areas called "IP Addressing Commonwealths" within which 32-bit IP addresses are unique and are part of a larger, globally-unique Internet addressing scheme. It is a goal of this effort to avoid requiring any router within a Commonwealth to be modified, but any host wishing full Internet connectivity will need to support IPAE eventually. Further, any system wishing to support full IPAE addresses will need to be modified, including network management software.

Goals and Milestones:

Done	Review and approve the Charter at the first Working Group meeting.
Done	Post the initial IPAE specification as an Internet-Draft.
Aug 1992	Post the initial "Addressing" specification as an Internet-Draft.
Done	Post the report to the IESG as an Internet-Draft.
Done	Present work of the IPAE Working Group to the IETF.
Jun 1993	Post the "Implementation and Transition" specification as an Internet-Draft.

Internet-Drafts:

"IPv7 Criteria Analysis for IP Address Encapsulation (IPAE) and the Simple Internet Protocol (SIP)", 11/11/1992, R. Hinden, S. Deering, D. Crocker <draft-ietf-ipae-ipv7-criteria-00.txt>

"IP Address Encapsulation (IPAE): A Mechanism for Introducing a New IP", 11/11/1992, D. Crocker, R. Hinden <draft-ietf-ipae-new-ip-00.txt>

INTERIM MEETING REPORT

Reported by Bob Hinden/Sun

Minutes of the IP Address Encapsulation Working Group (IPAE)

An interim meeting of the IPAE Working Group took place on March 3, 1993. Participants via the MCI Phone Bridge included, Mike Conn, Yakov Rekhter, Christina Huitema and Carl Beame. On-site participants included, Bob Hinden, Dave Crocker, Erik Nordmark, Steve Deering, Chuck Desostoa, and Ron Jacoby.

Review Implementations Status

Christian reported that Werner Vogels/IENSC is doing a Mach implementation.

Christian is able to spend about twenty percent of his time on SIP implementation. Steve Deering suggested looking at the latest Berkeley release on FTP.UU.NET for sources of the latest BSD implementations that support variable length network addresses (called NET2). Christian would like to have a module that could be added to existing kernels. A SIP route module like MRoute module needs to be incorporated.

Chuck Desostoa reported that he is interested in doing an implementation but that he does not have the resources to do it now. Dave Crocker suggested that Hewlett-Packard work with INRIA and TGV because these implementations are based on BSD code.

The Silicon Graphics (SGI) implementation is underway and it is expected that ping will be working within the next few weeks.

Telnet and FTP are working between Beame & Whiteside (B&W) and Sun Microsystems. There is still ICMP testing being done. There are some problems with ICMP checksums. When ICMP testing is finished, complete transport will be done; then on to applications. Expect a few more hours work. B&W are currently only doing IPAE.

SIP packets are flowing from Geoff Mulligan's revision of KA9Q. Testing is needed with other implementations.

Gary Malkin is doing extensions to SIP stuff for RIPv2. Some discussion on merits of Distance Vector..... [Note: Draft of SIP RIP now out as an Internet-Draft]

Bob Gilligan met with James Davidson of Network General. Network General wants to do SIP/IPAE sniffer implementation. They collected some SIP/IPAE traffic traces.

Planning for IETF Meeting

The next IPAE meeting will be held during the March IETF on Monday at 4:30-6:00 p.m. The Group discussed what should be done for the IETF? Give implementation status reports? Discuss ease of implementations, etc. Carl Beame said that he can attend, and might have SIP working by then. Sun is to get Solaris 2.x to Beame.

The Group decided to give a demonstration of the following Protocols at the March IETF:

- Telnet
- FTP
- Ping
- TFTP (?)
- Email (?)

The Group decided to setup testing for the IETF demonstration on Wednesday, at SGI. There will be a mixture of on-site and Internet testing.

Christian suggested that it was important to have working implementations on the Internet to test against. Sun and SGI can do this. Christian has also been getting inquiries about doing a SIP pilot. He will facilitate this in Europe.

SGI will put B&W implementation on the Internet. Carl Beame will be at INTEROP and will meet with Dave Crocker and Bob Hinden.

The Group needs to put together a test scenario before the 24th. Steve Deering suggested that a short demonstration document be put together.

There was some discussion about BIND and DNS name services. Erik Nordmark will look into adding this to the Sun implementation. Christian expects to have it running soon too. Will need to put AE records in to make it useful.

IPAE Protocol Issues

Steve Deering proposed there be a 4-byte PAD between IP and SIP. This would help 64-bit machines. He wants comments from implementors and proposed always putting 4 bytes of zero PADs between headers. This will allow more efficient processing of SIP packets on 64-bit machines. Only for encapsulating SIP in IPv4. Christian didn't like the original proposal because of implementation and performance costs. The new proposal still requires extra overhead which will be useful in the short to medium term. A long discussion followed on this topic.

The Group decided to leave it the way it is now (e.g., no padding between IP and SIP headers) their view being that the complexity was not worth the gain. It needs to be proven that this is really a win in the IPAE case as the gain is not clear given the added complexity to support it.

Action Items

Steve Deering/Erik Nordmark	Determine the requirements for an audio/video demonstration at the IETF meeting.
Dave Crocker	To get a slot for the demonstration.
Bob Gilligan	To document the demonstration scenario for testing.
Ron Jacoby	To put B&W implementation up on the Internet.

CURRENT MEETING REPORT

Reported by Dave Crocker/Silicon Graphics

Minutes of the IP Address Encapsulation Working Group (IPAE)

The IPAE Working Group met once during the March IETF. Since IPAE re-aligned itself to provide transition technology for SIP, the foci of IPAE discussions have been.

1. Modifications to SIP that are likely to facilitate transition, and
2. Implementation experiences with SIP/IPAE.

(Reference to IPAE usage is specifically to SIP-over-IP and SIP-mapped-to-IP. The first is to tunnel through the Internet and the second is to gateway to IPv4 hosts.)

The Working Group held a brief discussion about the SIP/IPAE demonstration slated for later in the week, discussing the implementations being shown and their use.

Bob Gilligan and Ron Jacoby discussed their implementation work. The Beame & Whiteside, Intercon, TGV and Network General implementations, for DOS, Windows, MacIntosh, VMS and network monitoring, respectively, were not available to make presentations.

Steve Deering then raised the issue about whether SIP fragmentation is end-to-end only, or can occur en-route. He is interpreting the Group's response as supporting the position that SIP fragmentation need *not* occur en-route.

The rest of the week was used for a demonstration of SIP/IPAE in action. All of the listed vendors had implementations participating. Within the demo facility, hosts interacted with pure SIP (SIP to SIP) and with pure IP (of course) and went through the unmodified Internet to interact with other SIP hosts, using IPAE tunneling (SIP-over-IP). While the Sun workstation can perform IPAE translation, operational constraints limited the demonstration so that SIP-mapped-to-IP translation was not included.

Attendees

Randall Atkinson	atkinson@itd.nrl.navy.mil
Dennis Baker	dbaker@wellfleet.com
David Crocker	dcrocker@mordor.stanford.edu
Steve Deering	deering@parc.xerox.com
Christine Fredenburg	cfredenburg@dsac.dla.mil
Danny Hanson	hanson.tic@commlan.safb.af.mil
John Hascall	john@iastate.edu
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@wco.ftp.com
Kenneth Key	key@cs.utk.edu

John Krawczyk	jkrawczy@wellfleet.com
Charles Lynn	clynn@bbn.com
Gary Malkin	gmalkin@xylogics.com
Michael Marcinkevicz	mdm@csu.net
Erik Nordmark	nordmark@eng.sun.com
John Penners	jpenners@advtech.uswest.com
David Reese	dave@csu.net
Terry Sullivan	terrys@newbridge.com
L. Stuart Vance	vance@tgv.com
Curtis Villamizar	curtis@ans.net
Hung Vu	hungv@fonorola.com
Fred Whiteside	fred@bws.com
Jane Wojcik	jwojcik@bbn.com

Teaching Applications to SIP

Bob Gilligan

Sun Microsystems

bob.gilligan@eng.sun.com

Background

Solaris 2.1 - Sun IPAE/SIP prototype

Services modified: telnet, ftp, tftp, "ping"

Changes: user-level client programs, server programs

New library functions for SIP address parsing/printing

int ascii2addr(long af, char *cp, char *ap);

char *addr2ascii(long af, char*ap, char *buf, int buflen);

Name to address translation functions not implemented yet

Programs are ambidextrous --

clients parse SIP and IPv4 addresses, use SIP, IPAE or IPv4

servers accept connections via SIP, IPAE or IPv4

Code Changes -- Client programs

telnet

parse SIP addresses

use AF_SIP form sockaddr_in structures

ftp

parse SIP addresses

use AF_SIP form sockaddr_in structures

use low-order 32 bits of local SIP address in PORT command

tftp

parse SIP addresses

use AF_SIP form sockaddr_in structures

ping

extensive changes to form IPAE header

Code Changes -- Server Programs

inetd -- no change

in.telnetd -- no change

in.tftpd

use SIP_ADDRFORM setsockopt to have kernel return SIP addrs

use AF_SIP form sockaddr_in structures

in.ftpd

use SIP_ADDRFORM setsockopt to have kernel return SIP addrs

use AF_SIP form sockaddr_in structures

compose destination SIP address for data connection using high-order 32 bits of client's SIP address and 32 bit IPv4 addr from the PORT command (if given)

Summary

Applications easy to modify for SIP

Please read and comment on:

"SIP Program Interfaces for BSD Systems" internet draft

SIP / IPAE DEMONSTRATION

March 29, 1993

Internet Engineering Task Force Meeting
Columbus, OH

SIP / IPAE OVERVIEW

- SIP: Simple Internet Protocol
  - Evolution of IPv4
  - 64-bit Addresses
  - Header Simplification
  - Options moved to Separate Headers
- IPAE: IP Address Encapsulation
  - Transition Scheme for SIP
  - Uses Encapsulation and Translation
  - Flexible Deployment Scheme
  - Self Configuring IPv4 Compatibility

SIP & IPAE Working Groups — March 30, 1993

ACTIVE PARTICIPANTS

• Beame & Whiteside	Implementation (PC)
• INRIA	Implementation (BSD, BIND), DNS & OSPF Specifications
• Intercon	Implementation (MAC)
• MCI	Phone Conferences
• Merit	IDRP for SIP Specification
• Network General	Implementation (Sniffer)
• SGI	Implementation (IRIX, NetVisulizer)
• Sun	Implementation (Solaris 2.x, Snoop, KA9Q)
• TGV	Implementation (VMS)
• Xerox PARC	*Steve Deering*
• Bill Simpson	Address Resolution / Router Discovery

SIP & IPAE Working Groups — March 30, 1993

NEW DOCUMENTS

- *SIP: A Simple Internet Protocol*, S. Deering, to be published in May '93 IEEE Network.
- *SIP-RIP*, G. Malkin, C. Huitema, Internet Draft, draft-ietf-sip-rip-00.txt
- *IDRP for SIP*, S. Hares, Internet Draft, draft-ietf-ipidrp-sip-00.txt
- *OSPF for SIP*, C. Huitema, to be published as an Internet Draft.
- *SIP Addresses in the Domain Name Service Specifications*, C. Huitema, to be published as an Internet Draft.
- *SIP Program Interfaces for BSD Systems*, R. Gilligan, to be published as an Internet Draft.

SIP & IPAE Working Groups — March 30, 1993

IMPLEMENTATIONS

OS	Organization	Status
BSD 4.3	INRIA	In progress.
DOS &Windows	Beame & Whiteside	Completed (telnet, ftp, tftp, ping)
IRIX	Silicon Graphics	In progress (ping)
KA9Q	Sun	In progress (ping)
Mac OS	Intercon	Completed (telnet, ftp, finger, ping)
Solaris	Sun	Completed (telnet, ftp, tftp, ping)
VMS	TGV	In Progress (ping)

TOOLS	Organization	Status
NetVisualizer	Silicon Graphics	Completed (SIP & IPAE)
Sniffer	Network General	Completed (SIP & IPAE)
Snoop	Sun	Completed (SIP & IPAE)

MISC	Organization	Status
Bind	INRIA	Code done

SIP & IPAE Working Groups — March 30, 1993

DEMONSTRATION CONFIGURATION

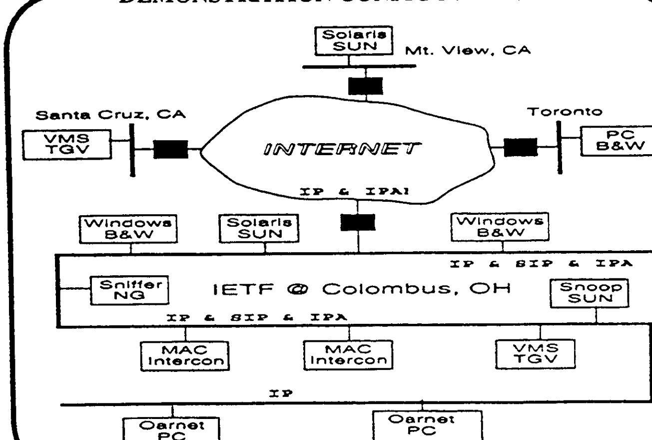

SIP & IPAE Working Groups — March 30, 1993

2.2.3 IP over AppleTalk (appleip)

Charter

Chair(s):

John Veizades, veizades@apple.com

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done	Post an Internet-Draft the current set of protocols used to connect Macintoshes to IP internets.
Done	Submit the AppleTalk MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"A Method for the Transmission of Internet Packets Over AppleTalk Networks [MacIP]", 03/08/1991, T. Evans, C. Ranch <draft-ietf-appleip-MacIP-02.txt>

"AppleTalk Management Information Base II", 12/21/1992, S. Waldbusser, K. Frisa <draft-ietf-appleip-mib2-01.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

CURRENT MEETING REPORT

Reported by John Veizades/Apple Computer

Minutes of the IP over AppleTalk Working Group (APPLEIP)

The meeting opened with discussion of the intended conclusion of the Working Group. The work of the Group will continue under the AppleTalk Network Forum (ANF). At the conclusion of the meeting the Group decided to disband, but will reactivate, if necessary, to complete the work currently on the standards track.

MIB+

Four issues were brought up in regard to the MIB+.

- Zone Name Changing

  Zone name changing will not be part of MIB+ and will be moved to a working group of the ANF for further study and work.

- Per Port Counters

  Both per port counters and aggregate counters will be kept. The per port counters can be queried by watching for changes in the aggregate counters.

- Zone Name matching of Get, Get Next and Set

  Lexical comparison of Zone names can be an issue in MIB+ implementations since the zone name equivalence comparison is different from the lexical order used in Get Next. This will be elucidated in the MIB+ document.

- Index of Zone Name

  Zone names will be used as the index into several tables.

The MIB+ will move forward and be submitted as an Internet-Draft by May 5th. It will then be submitted to the IESG for consideration as a Proposed Standard.

MacIP

The MacIP specification will be submitted as an Internet-Draft under the control of the APPLEIP Working Group if it is submitted by May 5th. After that date parties may submit the specification as an Informational RFC.

ANF

The ANF will take over all of the work that is currently done by the APPLEIP Working

Group. The ANF is still in its formation phase but it will try to have its first meeting at the MacTivity conference.

AFP MIB

The AFP MIB will soon be available on ftp.cayman.com. An announcement will be sent to the mailing list shortly. This work will be done under the ANF.

Other Issues

How does the work on IP next generation effect MacIP boxes?

Attendees

Jim Beers	Jim.Beers@cornell.edu
John Boatright	bryan_boatright@ksc.nasa.gov
Gregory Bruell	gob@wellfleet.com
Cyrus Chow	cchow@ames.arc.nasa.gov
David Dubois	dad@pacersoft.com
Karen Frisa	karen.frisa@andrew.cmu.edu
John Gawf	gawf@compatible.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
Bob Morgan	morgan@networking.stanford.edu
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Brad Parker	brad@fcr.com
Radia Perlman	perlman@dsmail.enet.dec.com
Sam Roberts	sroberts@farallon.com
Michael Safly	saf@tank1.msfc.nasa.gov
Bob Stewart	rlstewart@eng.xyplex.com
Wayne Tackabury	wayne@cayman.com
John Veizades	veizades@apple.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com

2.2.4 IP over Asynchronous Transfer Mode (atm)

Charter

Chair(s):

Robert Hinden, hinden@eng.sun.com

Mailing Lists:

General Discussion: atm@sun.com
To Subscribe: atm-request@sun.com
Archive: Send message to atm-request@sun.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 standards 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 standards 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) 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 Post an Internet-Draft for a mechanism for IP over ATM. (Multi-Protocol Interconnect over ATM AAL5)

Done Submit the Multi-Protocol Interconnect over ATM AAL5 to the IESG as a Proposed Standard.

Mar 1993 Post Internet-Draft for "Internet Requirements for ATM Signaling".

Jul 1993 Submit "Internet Requirements for ATM Signaling" to the IESG for consideration as an Informational Document.

Internet-Drafts:

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

"Partial Address Resolution in ATM Networks", 03/03/1993, S. Subramaniam <draft-ietf-atm-address-resolve-00.txt>

"IP over ATM : architecture, address translation, and call control", 03/22/1993, F. Liaw <draft-ietf-atm-address-translation-00.txt>

"NBMA Address Resolution Protocol (NBMA ARP)", 03/24/1993, J. Heinanen <draft-ietf-atm-nbma-00.txt>

CURRENT MEETING REPORT

Reported by Bob Hinden/Sun

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

Agenda

- Overview of ATM Forum Addressing Work, Keith McCloghrie
- Address Resolution in ATM Networks, Subbu Subramaniam
- Architecture, Address Translation, and Call Control, Fong-Ching Liaw
- NBMA Address Resolution Protocol (NBMA ARP), Juha Heinanen
- General Discussion of ATM Address Resolution
- Issues in the Interconnection of LANs and ATM Networks, Tim Salo

The ATM Working Group met for three sessions at the March IETF meeting. Over 129 people attended the Working Group sessions.

The first session consisted of four presentations on ATM addressing and ATM/Internet address resolution. Keith McCloghrie presented an overview of the ATM Forum addressing work. Subbu Subramaniam presented his ideas on how ATM address resolution should work. Fong-Ching Liaw presented pros and cons of subnet and peer model of address resolution in ATM networks and Juha Heinanen presented an overview of his NBMA Address Resolution Protocol (NBMA ARP) proposal.

There was considerable discussion about how address resolution should work, and the pros and cons of the subnet versus peer model. There were strong views on both sides. Subnet model is a conventual internet model where each network has it own addressing, and internetworking routing and addressing runs on top of this. In the peer model the ATM network is knowledgeable about and participates in Internet routing. Call connection (signaling) requests would use internet addresses to specify destination. The session ended with the suggestion that neither one approach would prevail and that there should be proposed mechanisms for combining the two approaches.

Mark Leabach presented slides with his thoughts on which areas the Working Group should focus on first. He saw two approaches in the Working Group: Functional layerists (ATM as a wire-replacement) versus ATM IP Morphist. He made a strong argument that the Working Group should first specify how IP over ATM should work in the "classical IP" mode where ATM networks are connected by routers. He presented the following problems which need to be solved:

- IP encapsulation in ATM (Done).

- How to "roll your own" wires over ATM between IP elements in the network.

- How to dynamically find other IP hosts on an ATM LAN.

- How to do connection-less service.

- Knowing that you have an all ATM path between IP elements, what optimizations can be made?

- Authentication of virtual circuit (VC) call setup.

- Management tools such as traceroute.

- Resource setup/quality of service (QOS).

- IP multicast to ATM multicast.

Tim Salo presented a talk on the Gigabit Testbed. The goal of the project is to create a seamless connection between ATM LAN's and ATM WAN's. His preliminary observations were that there are no complete proposals available today, some are only slightly explored or identified, much new functionality must be implemented, and most of the problems are in the wide area. He grouped issues into several problem areas:

- ATM address assignment
- IP address assignment
- Address resolution
- IP over ATM Internets
- IP route optimization
- Locality (IP only talks to local entities)

His final observations were that we need a complete solution if ATM is going to be the solution. It is important to avoid making the customer the system integrator. Significantly more implementation experience is needed.

Ramon Caceres presented the results of his simulation of different approaches for VC multiplexing. His conclusions were that one VC per connection gives the best performance, followed by one VC per type of traffic (e.g., telnet, ftp, mail, etc.). One VC per router pair gives poor performance. The overall goal should be to separate traffic as much as possible.

After a final discussion of subnet versus peer addressing models, the Working Group agreed that the following issues were important to pursue:

- Focus on subnet model first. Work on the peer model should continue, but it is important to develop a solution for the subnet model first.

- Prefer strongly typed addresses with peer model versus mechanical address translation.

- Simulation of subnet model is important. Need to look at performance issues based on data for real IP traffic.
- ARP servers will need mechanisms for updating and synchronizing databases.
- Charging is an issue with ARP servers.
- Auto-configuration and discovery of ARP servers is necessary.

Action Items

The following actions items were assigned:

Steve Willis:	Write contribution to be submitted to the ATM Forum on encapsulation (LLC) negotiation. Due in 1 week.
Mark Leabach:	Write a framework document describing the use of ATM in the classical IP model (i.e., ATM networks connected by routers). Due prior to next IETF meeting.
Robert Cole:	Write a framework document describing the use of ATM in the end-to-end model (i.e., seamless ATM from initial source to final destination). Due prior to next IETF meeting.
Juha Heinanen:	Complete draft on NBMA address resolution. Will include both general media independent specification and ATM specific usage guidelines. Due prior to next IETF meeting.
Bob Hinden:	Develop a list of IPoverATM issues. Due in May.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
Robert Aiken	aiken@es.net
N. Akiko Aizawa	akiko@nacsis.ac.jp
Kannan Alagappan	kannan@dsmail.lkg.dec.com
David Arneson	arneson@ctron.com
Fred Baker	fbaker@acc.com
Serpil Bayraktar	sbb@noc.ans.net
Nutan Behki	Nutan_Behki@qmail.newbridge.com
Tom Benkart	teb@acc.com
Lou Berger	lberger@bbn.com
B. Sue Blair	bsb@noc.ans.net
John Boatright	bryan_boatright@ksc.nasa.gov

Rich Bowen	rkb@ralvm11.vnet.ibm.com
Deborah Boyer	djb@noc.ans.net
Monroe Bridges	monroe@cup.hp.com
Thomas Brisco	brisco@pilot.njin.net
Caralyn Brown	cbrown@wellfleet.com
Theodore Brunner	tob@thumper.bellcore.com
Randy Butler	rbutler@ncsa.uiuc.edu
Ramon Caceres	ramon@mitl.com
Ross Callon	rcallon@wellfleet.com
Jeff Carman	tcarman@bnr.ca
John Chang	jrc@uswest.com
Yee-Hsiang Chang	yhc@hpl.hp.com
Anthony Chow	chow_a@wwtc.timeplex.com
Thomas Christie	tfc@aa.ans.net
George Clapp	clapp@ameris.center.il.ameritech.com
Richard Cogger	R.Cogger@cornell.edu
Robert Cole	rgc@qsun.att.com
Michael Collins	collinsms@es.net
Rob Coltun	rcoltun@ni.umd.edu
Tracy Cox	tacox@sabre.bellcore.com
Dave Cullerot	cullerot@ctron.com
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
James Davin	davin@thumper.bellcore.com
Shane Dawalt	sdawalt@desire.wright.edu
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Osmund DeSouza	osmund.desouza@att.com
M.J. Dixon	mjd@att.com
Kurt Dobbins	kurtdob@ctron.com
Bob Downs	bdowns@combinet.com
Kishan Dudkikar	kishan@icm1.icp.net
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Robert Enger	enger@reston.ans.net
Michael Fidler	fidler@mitre.org
Francois Fluckiger	fluckiger@vxcern.cern.ch
Craig Fox	foxcj@network.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Dan Frommer	dan@jeremy.enet.dec.com
John Gawf	gawf@compatible.com
Eugene Geer	ewg@cc.bellcore.com
Kenneth Giusti	kgiusti.chipcom.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Mike Goguen	goguen@synoptics.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.enet.dec.com

Tony Hain	alh@es.net
Joel Halpern	jmh@network.com
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Danny Hanson	hanson.tic@commlan.safb.af.mil
Eugene Hastings	hastings@psc.edu
Ken Hayward	Ken.Hayward@bnr.ca
Frank Heath	heath@cmc.com
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
John Heimaster	jwh@mps.ohio-state.edu
Juha Heinanen	juha.heinanen@datanet.tele.fi
Robert Hinden	hinden@eng.sun.com
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
Don Hoffman	hoffman@eng.sun.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
John Hopprich	hopprich@davidsys.com
Kathy Huber	khuber@bbn.com
Phil Irey	pirey@relay.nswc.navy.mil
Keisuke Ito	itoh@msen.com
David Jacobson	dnjake@vnet.ibm.com
David Johnson	dbj@cs.cmu.edu
Merike Kaeo	merike@alw.nih.gov
Kenneth Key	key@cs.utk.edu
Glenn Kime	itoh@msen.com
Charley Kline	cvk@uiuc.edu
Stev Knowles	stev@ftp.com
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Giri Kuthethoor	giri@ms.uky.edu
Mark Laubach	laubach@hpl.hp.com
Fong-Ching Liaw	fong@eng.sun.com
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Paul Lustgraaf	grpjl@iastate.edu
Charles Lynn	clynn@bbn.com
Carl Madison	carl@startek.com
Jamshid Mahdavi	Mahdavi@a.psi.edu
James Mahon	mahonj@honfsi1.att.com
Andrew Malis	malis_a@timeplex.com
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@cmf.nrl.navy.mil
Michael Marcinkevicz	mdm@csu.net
Cynthia Martin	martin@spica.disa.mil
Jerry Martin	jerry+@osu.edu
Matt Mathis	mathis@a.psc.edu
Jun Matsukata	jm@eng.isas.ac.jp
Keith McCloghrie	kzm@hls.com
Donald Merritt	Don@brl.mil

David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
James Miner	jjm@fibercom.com
Robert Mines	rfm@sandia.llnl.gov
Greg Minshall	minshall@wc.novell.com
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Steve Neruda	neruda@morningstar.com
Tom Nisbet	nisbet@tt.com
David O'Leary	doleary@cisco.com
Zbigniew Opalka	zopalka@agile.com
Ayal Opher	aopher@synoptics.com
Joe Pagan	jrp@afterlife.ncsc.mil
Laura Pate	pate@gateway.mitre.org
Maryann Perez	perez@cmf.nrl.navy.mil
Drew Perkins	ddp@fore.com
Radia Perlman	perlman@dsmail.enet.dec.com
Roy Perry	rperry@advtech.uswest.com
James Philippou	japhilippou@eng.xyplex.com
David Piscitello	dave@mail.bellcore.com
Narendra Popat	albin@frontier.com
Dave Rand	dave_rand@novell.com
Venkat Rangan	venkat@.metrix.com
David Reese	dave@csu.net
Bradley Rhoades	bdrhoades@mail.mmmg.com
Tony Richards	richards@icm1.icp.net
Michael Roberts	mmr@drwin.ptvy.ca.us
Ben Robinson	ben_robinson@vnet.ibm.com
Kenneth Rodemann	krr@qsun.att.com
Benny Rodrig	4373580@mcimail.com
Steve Romig	romig@cis.ohio-state.edu
Shawn Routhier	sar@epilogue.com
Timothy Salo	tjs@msc.edu
Hal Sandick	sandick@vnet.ibm.com
Vilson Sarto	vilson@fapq.fapesp.br
Shiva Sawant	shiva@synoptics.com
Andrew Schmidt	ags@uius.edu
Tim Seaver	tas@concert.net
Paul Serice	serice@cos.com
Kanan Shah	kshag@cmf.nrl.navy.mil
Karen Shay	kshay@cs.cmu.edu
Allyson Showalter	allyson@nsipo.arc.nasa.gov
Keith Sklower	sklower@cs.berkeley.edu
Andrew Smith	asmith@synoptics.com
Marco Sosa	mxs@sabre.bellcore.com

Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Richard Stuart	dickstuart@attmail.com
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Wayne Tackabury	wayne@cayman.com
Sally Tarquinio	sallyt@gateway.mitre.org
John Tavs	tavs@vnet.ibm.com
Kaj Tesink	kaj@cc.bellcore.com
Kamlesh Tewani	ktt@arch2.att.com
Susan Thomson	set@bellcore.com
Marek Tomaszewski	marek@net.com
Mario Vecchi	mpv@thumper.bellcore.com
Huyen Vu	vu@polaris.disa.mil
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Scott Wasson	sgwasson@eng.xyplex.com
Jack Waters	waters@sura.net
James Watt	james@newbridge.com
Von Welch	vwelch@ncsa.uiuc.edu
Rick Wilder	wilder@ans.net
Kirk Williams	kirk@sbctri.sbc.com
Steven Willis	steve@wellfleet.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Linda Winkler	lwinkler@anl.gov
Richard Woundy	rwoundy@vnet.ibm.com
Liang Wu	ltw99@bellcore.com
Yow-Wei Yao	yao@chang.austin.ibm.com
Chin Yuan	cxyuan@pacbell.com
Paul Zawada	Zawada@ncsa.uiuc.edu

Overview of ATM Forum Addressing Work

Keith McCloghrie
Hughes LAN Systems

March 1993

Various Proposals Considered by the ATM Forum

- use of E.164 everywhere
  - (most) public networks want to use E.164
  - thus, (I thought) cleanest solution
  - not practical, due to limited access to E.164 numbers
- use of new addressing plan
  - too difficult to setup new administration
- use of NSAPs
  - Big enough to provide routing hierarchy, and a structure for network/end-point split
  - ANSI/NIST/etc. provide address assignment
  - Q93B has codepoint for NSAPs
  - but, confusion between layer-2/layer-3 addressing

ATM Forum Agreements

[Caveat: not final until Signaling spec is balloted/published]

At a Private UNI:

- RFC 1237-structured NSAPs

At a Public UNI:

- use of one or more of:
  - RFC 1237-structured NSAPs
  - E.164 addresses (not NSAP encoded)
  - E.164 addresses (NSAP encoded)

Support of the above is mandatory – other formats and types may also be supported.

RFC 1237 Structure

```
<--------------------- NSAP ------------------------>
              <------ DSP (Domain Specific Part) ----->
! AFI ! IDI ! DFI ! AA  ! Rsvd ! RD ! Area ! ID ! Sel !
 <--Prefix assigned to AA----->
 <----Prefix assigned by AA-------->

where:  AFI   Authority and Format Identifier (1 octet)
        IDI   Initial Domain Identifier (2 octets)
        DFI   DSP Format Identifier (1 octet)
        AA    Administrative Authority (3 octets)
        Rsvd  Reserved (2 octets)
        RD    Routing Domain Identifier (2 octets)
        Area  Area Identifier (2 octets)
        ID    IEEE 802 MAC address (6 octets)
        SEL   Selector (1 octet)

using ICD (for NIST): AFI = 47, IDI = x'0005', DFI = x'80'
using DCC (for ANSI): AFI = 39, IDI = x'840f', DFI = x'dd'
```

ATM Address is a SNPA

The first sentence of the Addressing section of the Forum's current Signaling draft says:

```
An ATM address denotes a subnetwork point of attachment
to an ATM network
```

OSI's IONL defines a SNPA as a (media) address on a particular subnetwork, where subnetworks are connected via IS's to form an internet.

Thus, the ATM Forum's spec says:

- a (concatenation of) ATM networks is a subnetwork,
- an ATM address is the "media" address on such a subnetwork.

This agreement was reached after long discussions of the "subnet" model versus the "peer" model. (See references in Fong's I-D.)

Layer-2/Layer 3 Confusion

So, this addressing scheme

- uses an addressing structure, an address encoding, and an address-assignment administration defined for OSI's network layer,
- for ATM-layer "media" addresses.

There is danger of confusion here, but no other way is known to satisfy the requirements.

Address Registration

- to exchange addressing information across the UNI (e.g., at initialization/reset)
- network-side supplies "network part" of address
- user-side supplies "user part"
- NO configuration of local ATM-address in ATM end-points
- through symmetry, can be used at an NNI
- two mechanisms currently under discussion:
  - extension of ILMI
  - extension of meta-signaling

Address Resolution

- With ATM addresses as SNPAs, a means of resolving network-layer addresses (e.g., IPv4 or IPv7 addresses) to ATM-layer addresses is required.
- Contributions on each of the following three proposals were submitted to the ATM Forum's February meeting where no decisions were made pending discussions at this IETF meeting.

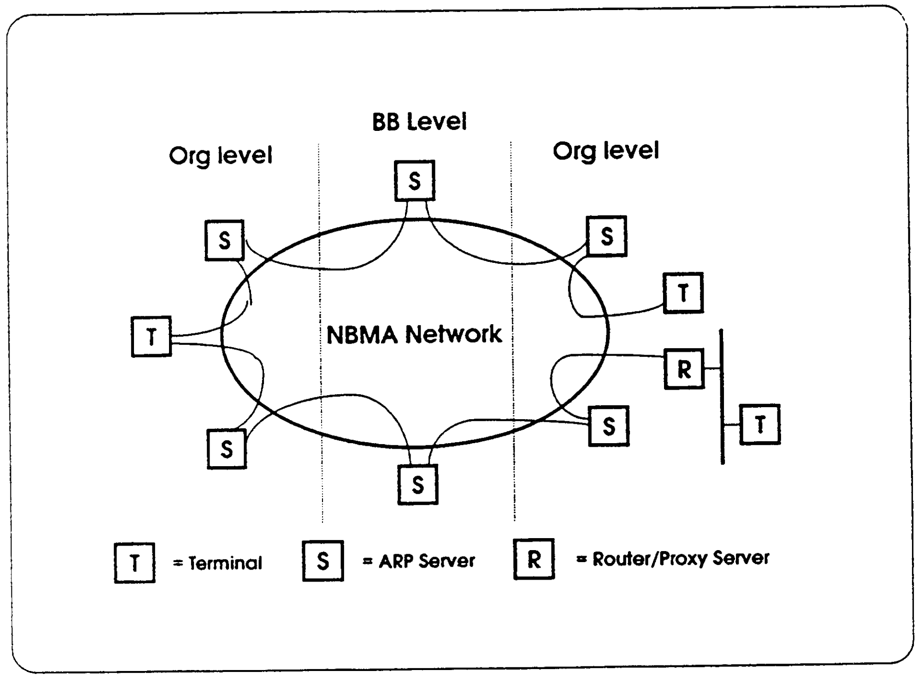
BB Level
Org level
Org level
S
S
S
T
T
NBMA Network
R
T
S
S
S
T
= Terminal
S
= ARP Server
R
= Router/Proxy Server

Address Resolution

o Before exchanging network Layer packets, hosts have to resolve a Network Layer Address (NLA) into an ATM Layer Address (ALA).

IP, CLNP, etc. →

AFI	IDI	......	IEEE MAC

as agreed upon earlier
Note: Hosts may change MAC-ID.

o Mechanisms:

- Mechanical Translation
- Translation through a protocol
  * Broadcast - based
  * Server - based

Server-based Address Resolution

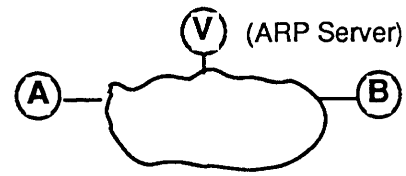

o A sends an 'ARP Request' containing B's NLA to the ARP Server V.

o A waits for 'ARP Reply' containing B's ALA.

o A initiates a connection to B's ALA.

Note: Time elapsed before exchanging network layer packets.

= Time taken for address resolution + Time taken for connection establishment.

Typically there are higher layer protocols as well.

Server-based Address Resolution

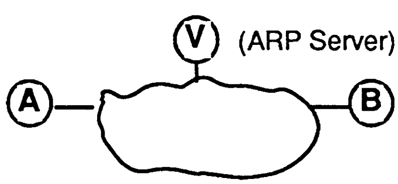

When 'B' is not "near":

o ARP Server in A's area does not know B's ALA.

o A's ARP Request is forwarded to an ARP server in B's area.

o ARP Server in B's area returns ARP Reply to A (establishes a connection?)

o A then sends out a SETUP packet.

Note: Time taken before exchanging network layer packets. = 2 times RTT

Server-based Address Resolution

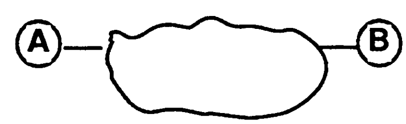

Example:

Distance btw A&B	=	3,000 miles
	~	4,800 km
Speed of light	=	300,000 km/sec
Round trip time	~	32 ms
Total time taken	~	64 ms

plus.
- Queuing
- Processing in switch
- etc.

Note: Approaching the requirement limit.

Partial Resolution Method

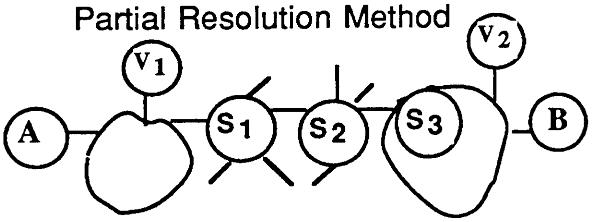

o A's ARP Request is partially resolved by V_1 (all fields except B's MAC-ID).

o A sends a SETUP packet with the partially resolved address. A includes a NLA information element.

o Switches S_1, S_2 pass NLA IE transparently, unaware that destination address is partially resolved.

o S_3 discovers a partially resolved address, completes the resolution by sending an ARP Request to V_2, and forwards the SETUP packet to B.

Partial Resolution Method

Extensions:

o Resolution can be completed in more than 2 stages.

o In the extreme, case, hosts can leave the entire ALA unresolved. (network can provide the ARP service)

Advantages:

o Meets the delay requirement set by the Forum.

o Switches continue to route based on ALA.

o Hosts can (in the extended method) use Network Layer Addresses and leave the ALA completely unresolved.

o Will work with future network layer protocols.

Partial Resolution Method

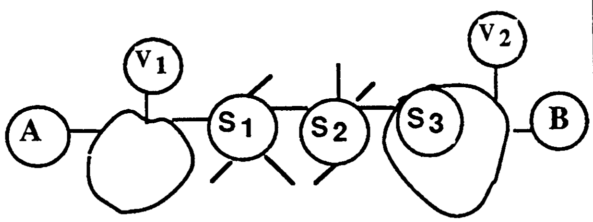

Assumptions:

o V_1 can provide the hierarchical part of B's ALA.

o S_3 implements the ARP protocol.

o New IE to carry NLA.

o A special 48-bit MAC-ID indicating "MAC Address Not Available".

o A and S_3 recognize this special MAC-ID.

It is important that:

o Address Resolution should take minimum time. In particular, a second round trip should be avoided.

o Address Resolution mechanisms should take advantage of ATM network features (e.g., signalling).

o Address Resolution methods should easily accommodate future Internet Protocols (e.g., IPv7).

Partial Resolution Method meets the above goals.

2.2.5 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

George Clapp, clapp@ameris.center.il.ameritech.com

Mailing Lists:

General Discussion: iplpdn@cnri.reston.va.us
To Subscribe: iplpdn-request@cnri.reston.va.us
Archive: ietf.cnri.reston.va.us:~/ietf-mail-archive/iplpdn/*

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:

TBD Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.

TBD Routing of IP datagrams across very large public data networks such as SMDS and Frame Relay.

Done Establish priorities and dates of completion for documents.

Internet-Drafts:

"Shortcut Routing: Discovery and Routing over Large Public Data Networks", 06/05/1992, P. Tsuchiya <draft-ietf-iplpdn-shortcutrouting-02.txt>

"Multiprotocol Interconnect over Frame Relay", 01/26/1993, T. Bradley, C. Brown, A. Malis <draft-ietf-iplpdn-framerelay-04.txt>

"The Transmission of Multi-protocol Datagrams over Circuit-mode ISDN", 02/17/1993, K. Sklower <draft-ietf-iplpdn-multi-isdn-00.txt>

"Parameter Negotiation for the Multiprotocol Interconnect", 02/17/1993, K. Sklower, C. Frost <draft-ietf-iplpdn-para-negotiation-00.txt>

"Management Information Base for Frame Relay DTEs", 03/24/1993, C. Brown, F. Baker, C. Carvalho <draft-ietf-iplpdn-frmib-dte-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"

RFC 1433 "Directed ARP"

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

Minutes of the IP over Large Public Data Networks Working Group (IPLPDN)

The following actions were taken during the week:

- The revised draft of RFC1294, "Multiprotocol Interconnect over Frame Relay," was approved for submittal to the IESG for release in a new RFC and for advancement from Proposed to Draft Standard.

- RFC1356, "Multiprotocol Interconnect over X.25," was reviewed for advancement in status. It was agreed to perform tests of interoperability of implementations.

- The revised RFC1315, the Frame Relay MIB document, was discussed. The Group agreed to keep this document as the "DTE MIB" and that the new work on a Frame Relay MIB would become the "DCE MIB".

- Earlier email had proposed that the IPLPDN Working Group complete the work on "multiprotocol over circuit ISDN," transfer open issues to other working groups, and then disband. Open issues to be transferred are address resolution, which might be taken to the IP over Asynchronous Transfer Mode Working Group (ATM), and routing issues, which may go to the Virtual Circuit Routing BOF (VCROUT), soon to be a working group. The Chair agreed to talk with the Area Directors and with working group chairs and report to the Group.

- Work progressed on the "Multiprotocol over Circuit ISDN" document. The Draft was re-titled "Encapsulation Determination over Circuit Switched Services." Keith Sklower will incorporate comments and will distribute the revised document by email for Working Group approval for submittal to the IESG.

Two joint sessions were held with the PPPEXT Working Group. Readers are referred to the Minutes written by Fred Baker.

Attendees

Nagaraj Arunkumar	nak@3com.com
Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Nutan Behki	Nutan_Behki@qmail.newbridge.com
Monroe Bridges	monroe@cup.hp.com
Caralyn Brown	cbrown@wellfleet.com
George Clapp	clapp@ameris.center.il.ameritech.com
Brian Coan	coan@faline.bellcore.com

Tracy Cox	tacox@sabre.bellcore.com
Wayne Cullen	wnc@netlink.com
M.J. Dixon	mjd@att.com
Bob Downs	bdowns@combinet.com
Kishan Dudkikar	kishan@icm1.icp.net
Michael Fidler	fidler@mitre.org
Craig Fox	foxcj@network.com
Paul Franchois	paulf@bldrdoc.gov
Anita Freeman	
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.enet.dec.com
Joel Halpern	jmh@network.com
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Ken Hayward	Ken.Hayward@bnr.ca
Frank Heath	heath@cmc.com
John Heimaster	jwh@mps.ohio-state.edu
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
Keisuke Ito	itoh@msen.com
Glenn Kime	itoh@msen.com
Giri Kuthethoor	giri@ms.uky.edu
William Kwan	kwan@rabbit.com
Mark Laubach	laubach@hpl.hp.com
Andrew Malis	malis_a@timeplex.com
Gerry Meyer	gerry@spider.co.uk
James Miner	jjm@fibercom.com
Greg Minshall	minshall@wc.novell.com
Shannon Nix	sdn@netlink.com
David O'Leary	doleary@cisco.com
Brad Parker	brad@fcr.com
Drew Perkins	ddp@fore.com
James Philippou	japhilippou@eng.xyplex.com
Kelly Quaal	quaalk@network.com
Jim Rees	jim.rees@umich.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Timothy Salo	tjs@msc.edu
Hal Sandick	sandick@vnet.ibm.com
Andrew Schmidt	ags@uius.edu
William Simpson	Bill.Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Marek Tomaszewski	marek@net.com
Stephen Tsun	snt@3com.com
Robert Ullmann	Ariel@Process.com
Scott Wasson	sgwasson@eng.xyplex.com

James Watt	james@newbridge.com
Kirk Williams	kirk@sbctri.sbc.com
Linda Winkler	lwinkler@anl.gov

2.2.6 P. Internet Protocol (pip)

<u>Charter</u>

Chair(s):

Paul Francis, `Francis@thumper.bellcore.com`

Mailing Lists:

General Discussion: `pip@thumper.bellcore.com`
To Subscribe: `pip-request@thumper.bellcore.com`
Archive: `thumper.bellcore.com:~/pub/tsuchiya/pip-archive`

Description of Working Group:

The PIP Working Group is chartered to develop an IPv7 proposal using the basic ideas of PIP as described in the PIP overview.

PIP is designed on one hand to be very general, being able to handle many routing/addressing/flow paradigms, but on the other hand to allow for relatively fast forwarding. PIP has the potential to allow for better evolution of the Internet. In particular, it is hoped that we will be able to advance routing, addressing, and flow techniques without necessarily having to change hosts (once hosts are running PIP).

While the PIP overview demonstrates a number of powerful mechanisms, much work remains to be done to bring PIP to a full specification. This work includes, but is not limited to, specifying the header format; specifying a basic set of error messages (PCMP messages); specifying the PIP forwarding rules; specifying host interface messages (particularly the directory service query response); specifying rules for host PIP header construction; specifying modifications to existing protocols for use with PIP (BGP IV, OSPF, ARP, DNS, etc.); specifying PIP maximum MTU discovery techniques; and specifying a transition strategy for PIP.

Over the near-term, the goal of the PIP Working Group will be to produce these specifications and supporting documentation. Over the long-term, up to the point where PIP is definitively rejected as IPv7, it is expected that the PIP Working Group will oversee implementations and testing of the PIP specifications.

Except to the extent that the PIP Working Group modifies existing protocols for operation with PIP, and to the extent that the PIP Working Group must be aware of routing/addressing/flow architectures to really make PIP general, the PIP Working Group will not work on routing/addresing/flow architectures.

Goals and Milestones:

Done	Review and approval of the Charter for the PIP Working Group.

Done	Post as an Internet-Draft a description of the PIP Packet Format and Forwarding Engine, the PIP Control Message Protocol (PCMP), the PIP Host Interface Message Protocol, and the PIP MTU Discovery Protocol.
Oct 1992	Post as an Internet-Draft a description of the modifications to BGP IV for PIP, the Modifications to OSPF for PIP, and the modifications to ARP for PIP.
Done	Presentation and review of the PIP specification by the IESG. If acceptable, the first Working Group meeting will be held.
Done	Post as an Internet-Draft the modifications to DNS for PIP, the Address assignment in PIP, and the PIP transition strategy.

Internet-Drafts:

"Pip Header Processing", 10/30/1992, P. Tsuchiya <draft-ietf-pip-processing-01.txt>

"The EIPIP Protocol: a Pip engine with an EIP shell", 11/03/1992, Z. Wang, P. Tsuchiya <draft-ietf-pip-eip-shell-00.txt>

"Transition to the Future Internet Protocol a comparison of three transition schemes", 11/03/1992, Z. Wang <draft-wang-transition-00.txt>

"Pip Identifiers", 11/03/1992, P. Tsuchiya <draft-ietf-pip-identifiers-01.txt>

"IPv7 Criteria Analysis for EIPIP", 11/13/1992, P. Tsuchiya, Z. Wang <draft-ietf-pip-ipv7-analysis-00.txt>

"Use of DNS with Pip", 01/29/1993, P. Tsuchiya, S. Thomson <draft-ietf-pip-dns-00.txt>

"Pip Near-term Architecture", 02/22/1993, P. Tsuchiya <draft-ietf-pip-architecture-00.txt>

"On the Assignment of Provider Rooted Addresses", 03/22/1993, P. Tsuchiya <draft-ietf-pip-provider-addr-00.txt>

"The Multi-Level Path Vector Routing Scheme", 04/08/1993, B. Rajagopalan, P. Francis <draft-ietf-pip-vector-00.txt>

CURRENT MEETING REPORT

Reported by Paul Francis/Bellcore

Minutes of the P. Internet Protocol Working Group (PIP)

The PIP Working Group met during both afternoon sessions of Tuesday, March 30th.

By and large, the meeting could be described as tutorial in nature. Paul Francis (formerly Tsuchiya) covered various aspects of the PIP protocol.

Paul first went over the changes in the header since Washington, D.C. The major changes include modification of how the Routing Context and Handling Directive are interpreted, and in how options are handled. During this discussion, Steve Deering uncovered a bug in the use of Source and Dest ID + Protocol + Packet SubID as a caching mechanism.

Next, Paul went over the structure of PIP Addresses and how they are used in the PIP header. Though many questions were asked, no particular problems were uncovered.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
Tom Benkart	teb@acc.com
Ross Callon	rcallon@wellfleet.com
Jeff Carman	tcarman@bnr.ca
Steve Deering	deering@parc.xerox.com
Eric Fleischman	ericf@act.boeing.com
Eugene Geer	ewg@cc.bellcore.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Fengmin Gong	gong@concert.net
John Hascall	john@iastate.edu
John Ioannidis	ji@cs.columbia.edu
Phil Irey	pirey@relay.nswc.navy.mil
Moshe Kochinski	moshek@FibHaifa.com
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Charles Lynn	clynn@bbn.com
Carl Madison	carl@startek.com
Marjo Mercado	marjo@cup.hp.com
David Meyer	meyer@ns.uoregon.edu
Manoel Rodrigues	manoel_rodrigues@att.com
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Richard Thomas	rjthomas@bnr.ca
Warren Vik	wmv@lachman.com

2.2.7 Point-to-Point Protocol Extensions (pppext)

Charter

Chair(s):

Fred Baker, fbaker@acc.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 layer 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:

"Requirements for an Internet Standard Point-to-Point Protocol", 07/30/1988, D. Perkins <draft-ietf-ppp-requirements-02.txt>

"The PPP Internetwork Packet Exchange Control Protocol (IPXCP)", 06/10/1992, W. Simpson <draft-ietf-pppext-ipxcp-03.txt>

"Compressing IPX Headers Over WAN Media (CIPX)", 12/08/1992, S. Mathur, M. Lewis <draft-ietf-pppext-cipx-03.txt>

"PPP LCP Extensions", 01/08/1993, W. Simpson <draft-ietf-pppext-lcpext-01.txt>

"PPP over ISDN", 03/10/1993, W. Simpson <draft-ietf-pppext-isdn-00.txt>

"PPP over X.25", 03/10/1993, W. Simpson <draft-ietf-pppext-x25-00.txt>

"PPP over Frame Relay", 03/10/1993, W. Simpson <draft-ietf-pppext-fr-00.txt>

"PPP over SONET", 03/10/1993, W. Simpson <draft-ietf-pppext-sonet-00.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"

RFC 1334 "PPP Authentication Protocols"

RFC 1376 "The PPP DECnet Phase IV Control Protocol (DNCP)"

RFC 1377 "The PPP OSI Network Layer Control Protocol (OSINLCP)"

RFC 1378 "The PPP AppleTalk Control Protocol (ATCP)"

RFC 1471 "The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol"

RFC 1472 "The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol"

RFC 1473 "The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol"

RFC 1474 "The Definitions of Managed Objects for the Bridge Network Control Protocol of the Point-to-Point Protocol"

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

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

Status of work in progress:

Proposed Standards

RFC1331 - PPP	In considering PPP for other HDLC encapsulations such as LAPB and Frame Relay, it became clear that the document does not clarify the difference between the transport (the protocol in the address and control fields) and the PPP negotiation and data services (everything else). Bill Simpson will reorganize the document, possibly into two documents, and when we concur on the output we expect to advance this to Draft Standard. We agreed on one FSM change: when a code reject occurs, the sender should swallow it EVEN IN THE OPEN STATE.
RFC1332 - IPCP RFC1333 - LQM	Several tested interoperable implementations of these exist, and we have data supporting that. Mark Lewis will write an "Experience Document" supporting the advancement of these and PPP to Draft Standard.
RFC1334 - Authentication	There is work happening on the nas-req@merit.edu list to include a version of CHAP based on recent PEM work, and a Kerberos option.
RFC1220 - Bridging	Rich Bowen of IBM will make a backward compatible update noting experience and adding an IBM SR Bridge BPDU. We will probably want to advance this to Draft Standard when the updates are ready.

The following Internet-Drafts have been advanced:

ppp-requirements-02	Bill Simpson will update and submit as an Informational RFC.
pppext-bridgemib-01	==> Proposed Standard
pppext-cipx-02	==> Proposed Standard
pppext-ipcpmib-01	==> Proposed Standard

pppext-ipxcp	One sentence change, ==> Proposed Standard. Bill will change the sentence and post the draft for IESG consideration.
pppext-lcpext-00	Some Updates Required, which Bill will do.

Other Work in Progress

- Keith Sklower will revise his load sharing draft (see minutes of the joint meetings of IPLPDN and PPPEXT) for further review. Fred Baker is withdrawing his proposal in favor of Keith's.

- Novell has prototyped the LAPB and Compression proposals, and has donated the time of Dave Rand to document their work and turn it into an interoperable consensus specification. Especially in view of earlier poor relations between the Working Group and Novell, we welcome this effort and support. Fred Baker will put the current proposal into a form that can be handed off, and Dave will revise it as appropriate.

- AppleTalk Smart Buffering and EDDP are under consideration by Brad Parker. He will post updated documents as he is able to complete them. Brad will also record experience in an updated AppleTalk/PPP document.

PPPEXT plans to have one meeting at the Amsterdam IETF, and two joint meetings with IPLPDN.

Attendees

David Arneson	arneson@ctron.com
Jim Barnes	barnes@xylogics.com
Rich Bowen	rkb@ralvm11.vnet.ibm.com
Caralyn Brown	cbrown@wellfleet.com
Anthony Chow	chow_a@wwtc.timeplex.com
George Clapp	clapp@ameris.center.il.ameritech.com
Shane Dawalt	sdawalt@desire.wright.edu
Bob Downs	bdowns@combinet.com
Avi Elenko	avi@dss.com
Craig Fox	foxcj@network.com
Karl Fox	karl@morningstar.com
Anita Freeman	
Chris Gunner	gunner@dsmail.enet.dec.com
Joel Halpern	jmh@network.com
Frank Heath	heath@cmc.com
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
Keisuke Ito	itoh@msen.com
Glenn Kime	itoh@msen.com

Mark Lewis	Mark.S.Lewis@telebit.com
Gerry Meyer	gerry@spider.co.uk
David O'Leary	doleary@cisco.com
Drew Perkins	ddp@fore.com
James Philippou	japhilippou@eng.xyplex.com
Kelly Quaal	quaalk@network.com
Dave Rand	dave_rand@novell.com
Tony Richards	richards@icm1.icp.net
Benny Rodrig	4373580@mcimail.com
Hal Sandick	sandick@vnet.ibm.com
Vilson Sarto	vilson@fapq.fapesp.br
Mark Silverman	mark@dss.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Bob Sutterfield	bob@morningstar.com
Steve Suzuki	suzu@fet.com
Steven Sweeney	steves@farallon.com
Kim Toms	kim@morningstar.com
Stephen Tsun	snt@3com.com
Richard Warwick	richard@dss.com
Scott Wasson	sgwasson@eng.xyplex.com
Kirk Williams	kirk@sbctri.sbc.com
Steve Wilson	steve@morningstar.com

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

Minutes of the joint sessions of IPLPDN and PPPEXT Working Groups

The two Working Groups met in joint session to discuss protocol specifications common to both. Since the objectives and requirements of the two Working Groups differ in some key respects, there was considerable difference of opinion at the outset. The Chairs wish to congratulate the various parties in the discussions on the level of personal restraint and professionalism displayed during the discussions. Would that there were even more of both.

Two subjects were discussed: how to share load among a set of parallel links to increase apparent bandwidth and potentially reduce latency between two sites, and how the IPLPDN Group might best avail itself of the facilities found in PPP negotiation.

Load Sharing

Two proposals for load sharing were outlined. Fred Baker briefly reminded the Group of his previous proposal to use ISO Multilink Procedures as described in ISO 7776 and ISO 7428. Keith Sklower discussed his proposal to use the existing RFC1294 segmentation encapsulation to achieve traffic ordering in much the same way that Multilink does, and provide the option of data fragmentation and reassembly.

The consensus of the Group preferred Keith's approach, but recommended that two currently unused bits be assigned the purposes of the ISO Multilink Protocol's RESET and RESET ACKNOWLEDGE flags to facilitate synchronization of links when bringing them into and out of service. Concerns were raised about the size of the resequencing buffer, most especially when the link speeds are mismatched. Joel Halpern and Craig Fox will provide input to Keith concerning a solution to this that they worked out; Keith will appropriately edit his proposal for further discussion on the IPLPDN mailing list.

PPP Parameter Negotiation for Frame Relay

The consensus we reached is encapsulated in the following points.

- By default, Frame Relay services will conform to RFC1294. This implies that if two systems attempt to communicate, one using RFC1294 and the other using PPP services, the system desiring PPP services will use RFC1294 instead.
- There may be a requirement to negotiate services in both PVC and SVC environments.
- Negotiation uses PPP negotiation frames encapsulated in a manner conforming to RFC1294.

- The system receiving a PPP negotiation frame may choose to ignore and discard it, either because the system is old or because it is configured to do so. Once both systems have decided to negotiate, the full PPP negotiation FSM takes effect.

- There may be LCP configuration options to modify the encapsulation on a virtual circuit.

- We jointly agree to specify PPP encapsulation. The PPP encapsulations of choice are:

  First choice:

```
 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Q.922 Address Field   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Control Field | NLPID (TBD)   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         PPP PID               |
```

  Second choice:

```
 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Q.922 Address Field   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Control Field | PAD (00)      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    NLPID=80   |      OUI = 00 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    OUI = 00-00                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Ethernet Packet Type TBD   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         PPP PID               |
```

  Use of the first encapsulation is subject to assignment of an NLPID value by X3S3. Bill Simpson reports that Lyman Chapin feels ANSI has assigned a block of NLPID values to the IAB for IETF purposes.

- By default, data which has an existing RFC1294 encapsulation should be encapsulated as in RFC1294, unless the two systems agree, using LCP negotiation, to use the above encapsulation for data. Data which has no defined RFC1294 encapsulation but has a PPP encapsulation must use the above for data regardless of the outcome of that negotiation. Data which has no defined PPP encapsulation but has an RFC1294 encapsulation must use the latter regardless of the outcome of the negotiation.

- Because it is specified in the PPP FSM, all data flow stops during LCP negotiation. Data streams having PPP NCPs do not restart until their PPP NCP has reached the open state. Data streams not having PPP NCPs may restart upon the LCP reaching the open state. As in PPP, subsequent renegotiation of an NCP affects only its data.

Keith Sklower and Bill Simpson have agreed to merge their efforts and their current proposals to implement this consensus. The output will be discussed on the IPLPDN list.

2.2.8 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 RFCs, 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 document 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 IP Routers Volume 1: Introduction", 09/17/1990, Philip Almquist <draft-ietf-rreq-iprouters-04.txt>

Request For Comments:

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

RFC 1354 "IP Forwarding Table MIB"

2.2.9 Simple Internet Protocol (sip)

Charter

Chair(s):
Steve Deering, deering@parc.xerox.com

Mailing Lists:
General Discussion: sip@caldera.usc.edu
To Subscribe: sip-request@caldera.usc.edu
Archive:

Description of Working Group:

SIP is another candidate for IPv7. The purpose of the Working Group is to finalize the SIP family of protocols, and to foster the early development and experimentation of this protocol.

There are two major characteristics of the SIP proposal: it is very much a continuation of IP, and it aims at maximum simplicity. A short hand definition of SIP could be "64 bits IP with useless overhead removed".

Following the IP model, SIP uses globally-unique addresses, hierarchically structured for efficient routing. SIP addresses are 64 bits long, which is believed to be adequate to scale the Internet up to, say, thousands of internet-addressable devices in every office, every residence, and every vehicle in the world.

The quest of simplicity in SIP has been described as parallel to the RISC philosophy. The minimal SIP header contains only those fields which are necessary to achieve our goal: routing packets efficiently in a very large internet. As a result of this design philosophy, the SIP header is much simpler than the IP header. Simplicity facilitates high-performance implementation and increases the likelihood of correct implementation.

Contrary to several other IPv7 candidates, the SIP effort is focused mostly on the description of the final state, not on the description of the transition. This is due to a coordination with the IPAE Working Group, which has already engaged an intensive study of transition problems, with SIP in mind as a final state.

Goals and Milestones:

Done Post the complete SIP specification as an Internet-Draft. This specification shall include the header format, the address format, ICMP and IGMP, the fragmentation protocol, the source route protocol, and the the requirements SIP imposes on higher layer protocols and lower later protocols, e.g., ARP.

Done Post an Internet-Draft specifing the SIP addressing and routing architecture. Include discussion of multicast and mobile host support as well as a discussion of how policy routing can be supported. Detail the changes required to OSPF, BGP, and RIP.

Jan 1993 Post as an Internet-Draft a specification for the SIP MIB. Detail the operation of SNMP over SIP.

Jan 1993 Make available a public domain implementation of SIP for the UNIX-BSD socket environment.

Jan 1993 Make available a public domain version of modified TCP and UDP for the UNIX-BSD socket environment.

Mar 1993 Post as an Internet-Draft a report on the initial implementation and experience with SIP.

Jun 1993 Incorporate security into SIP.

Done Post an Internet-Draft specifying changes to RIP needed for SIP.

Internet-Drafts:

"Simple Internet Protocol (SIP) Specification", 11/11/1992, S. Deering <draft-deering-sip-00.txt>

"SIP-RIP", 03/11/1993, G. Malkin, C. Huitema <draft-ietf-sip-rip-00.txt>

"SIP Program Interfaces for BSD Systems", 04/05/1993, R. Gilligan <draft-ietf-sip-bsd-api-00.txt>

"Administrative Allocation of the 64-bit Number Space", 04/19/1993, W. Simpson <draft-ietf-sip-64bit-plan-00.txt>

"SIP System Discovery", 04/21/1993, W. Simpson <draft-ietf-sip-discovery-02.txt>

CURRENT MEETING REPORT

Reported by Christian Huitema/INRIA

Minutes of the Simple Internet Protocol Working Group (SIP)

The first session of the SIP Working Group was devoted to the finalization of the SIP specification, especially in light of the first interoperability experiences. The second session was devoted to the analysis of routing protocols.

SIP Specification

Steve Deering reviewed the recent "precisions" to the SIP specification:

- The first 32 flow-ids values have been "reserved" for IP "TOS" compatibility.

- The formats for the LSR and reassembly payloads have been slightly modified, so that the "payload type" arrives first.

- The payload type 0 has been reserved for hop-by-hop options. Routers are supposed to inspect the payload type of every packet. If this type is set to zero, then hop-by-hop options should be examined.

A generic format for hop-by-hop has been defined: after a generic "option header" expressing the embedded payload type and the number of 64-bit words used to carry the option, each hop-by-hop option is represented by a set of 32-bit words; the first octet includes the option type and the option length, expressed as a number of 32-bit words. There was a discussion on the adequacy of using the 32-bit words units, and a proposal to use a byte count instead; the Working Group turned this down, and reached consensus on the 32-bit word count. Steve then presented the requirements for "option ordering". To sum up, it is required to have the following order in the packets:

- Sip header
- Hop-by-hop option, if present
- Source route option, if present
- Fragmentation header, if present
- Data

Steve then presented a change in terminology: "Cluster address" instead of "anyone" address. The Working Group noted the requirement for two precisions in the specifications:

1. Multicast addresses cannot be used in source routes.
2. Cluster addresses should never be used as source addresses. This implies that cluster addresses should not be used for TCP connections, as they cannot be used as "reply" addresses.

Steve also mentioned that the new specification will include precisions on the "pseudo-checksum computation" (e.g., for ICMP and IGMP), in particular when the "C" bit is set.

IEEE 802 Address Format

The discussion switched then to the definition of an IEEE 802 address format. Steve Deering and Bob Hinden presented a format where a SIP address can be build by prepending a 16-bit header to an IEEE 802 48-bit address.

```
-------------------------------------
| prfx | Subnet | IEEE 802 address |
-------------------------------------
```

The 16-bit prefix, in this proposition, is divided in two fields, a 6-bit prefix for recognizing this address as an IEEE 802 address, and a 10-bit "subnet identifier" to identify the "local network" to which the station is connected. This triggered a discussion on the usage of this addresses and the proper length for the prefix and subnet field, with the following conclusions:

- These addresses should only be used as long as no "real" SIP address has been assigned to the station.
- They should never be advertized outside of the autonomous system, i.e. through IDRP.
- The prefix should be 8 bits long, and the subnet also 8 bits. The need to study the interaction with the DNS was also mentioned.

Security Labelling

The last speaker in the first session was Randy Atkinson, who had submitted a "SIPSO" draft, describing a "security labeling" option, and also presented the various possibilities for designing an end-to-end security option based on SP-3 and NLSP. The discussion of the labeling option was marked by the following comments:

- The main rationale for the labeling proposal was compatibility with IPv4. Some IPv4 routers used in a "secure environment" use the labeling service, and would suffer from "reduced capabilities" if the function was not available.
- It was pointed out that the security label alone was not very useful and that security should be addressed "globally", e.g. also in the routing architecture.
- Randy Atkinson mentioned that one rationale for the labeling option was also to prevent negative feelings, e.g. propagation of the false impression that SIP would be less secure than IPv4 because it would be missing the option.

All members agreed that end-to-end security was much more important. Randy continued his presentation by explaining how SP-3 is used to encapsulate and protect IP packets. It was decided to track the IPv4 efforts on the subject, and to come out with a SIP version of their proposal as soon as the IPv4 drafts stabilize.

Session 2

The first point on the Agenda of the second session was the review of the "DNS for SIP" Draft. The Draft defines an "AA" type record for storing the 64-bit SIP addresses, and a reverse tree under "sip-addr.arpa". It was decided after discussions to align the format of the reverse addresses with the "standard" external format of SIP addresses. The name corresponding to the SIP address:

    0abc:f120:138.96.24.84

will be obtained through the inverse domain name:

    84.24.96.138.f120.abc.sip-addr.arpa

The SIP-DNS Draft will be updated accordingly.

SIP Version of Three Routing Protocols

The next part of the session was devoted to the study of the "SIP" version of three routing protocols, RIP, OSPF and IDRP.

- The SIP-RIP Draft was presented by Gary Malkin. The Draft is derived from RIP-2, with the following modifications:
  - Addresses are 64 bits.
  - Bit masks are contiguous.
  - The metric is designed to converge on the best throughput, instead of the lesser number of hops.

  Gary presented three possible amendments to SIP-RIP, proposed by himself and Yakov Rekhter:

  1. As the SIP routers can easily know the MTU of the interface on which a RIP packet is transmitted, there is no need for a standard 576 length limit – they can as well make packets as big as the MTU allows.

2. By using an algorithm suggested in the 1989 SIGCOM conference, it is possible to remove the "bouncing effect" and to compute loop free routes. The cost is to add one extra address information per routing entry, and one computation step before propagating routes.

3. Use the "route on demand" improvement.

After discussion, it was decided that modification (1) and (2) should be incorporated in the current draft, and that (3), which represents a substantial, although compatible improvement, should be specified in a separate document.

- The SIP-OSPF Draft was presented by Christian Huitema. The main features of the Draft are:
  - Regular OSPF, running over IPv4.
  - Two additional LSAs to import SIP information and an additional bit in the router-LSA to indicate SIP capability.
  - "Integrated routing" in order to ease the migration from IPv4 to SIP, after which a native OSPF for SIP would be defined.

  A number of points were raised:

  - The IPv4 address to use when tunneling should be that of the selected interface to the dual SIP/IPv4 host. There should be a way to identify this interface address.
  - The translation between SIP and IPv4 formats should only take place in border routers.
  - "Double transition" should be avoided and the specification should be definitive.
  - We should specify clearly whether "SIP" areas are requested to have the same boundaries as "IPv4" areas.
  - When defining new LSAs, we should perhaps be able to specify a "multiprotocol" format.

  It was decided that all detailed discussions of OSPF in SIP would be carried on in the SIP Working Group.

- The IDRP for SIP draft written by Sue Hares was presented by Yakov Rekhter. Options for representing SIP addresses and for identifying SIP "autonomous systems" were discussed, as well as the need to define a document for multicast routing.

Bill Simpson presented the provisional results of the working party on ES/IS, ARP and autoconfiguration.

Attendees

Kannan Alagappan	kannan@dsmail.lkg.dec.com
Guy Almes	almes@ans.net
Michael Anello	mike@xlnt.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
David Battle	battle@cs.utk.edu
Tom Benkart	teb@acc.com
Fred Bohle	fab@interlink.com
David Bolen	db31@ans.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Robert Braden	braden@isi.edu
Monroe Bridges	monroe@cup.hp.com
David Bridgham	dab@epilogue.com
Al Broscius	broscius@bellcore.com
Jeff Carman	tcarman@bnr.ca
Kevin Carosso	kvc@innosoft.com
David Carr	Carr@acsu.buffalo.edu
John Chang	jrc@uswest.com
Rob Coltun	rcoltun@ni.umd.edu
Steve Deering	deering@parc.xerox.com
Osmund DeSouza	osmund.desouza@att.com
Chas DiFatta	chas@cmu.edu
Ralph Droms	droms@bucknell.edu
David Dubois	dad@pacersoft.com
Pierre Dupont	dupont@mdd.comm.mot.com
Tom Easterday	tom@cic.net
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Eric Fleischman	ericf@act.boeing.com
Paul Francis	Francis@thumper.bellcore.com
Peter Furniss	p.furniss@ulcc.ac.uk
John Gawf	gawf@compatible.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Ramesh Govindan	rxg@thumper.bellcore.com
Danny Hanson	hanson.tic@commlan.safb.af.mil
John Hascall	john@iastate.edu
Robert Hinden	hinden@eng.sun.com
Jeffrey Honig	Jeffrey_C_Honig@Cornell.edu
Steven Hubert	hubert@cac.washington.edu
Christian Huitema	christian.huitema@sophia.inria.fr
John Ioannidis	ji@cs.columbia.edu
Phil Irey	pirey@relay.nswc.navy.mil
Ronald Jacoby	rj@sgi.com
David Johnson	dbj@cs.cmu.edu

Laurent Joncheray	lpj@merit.edu
Dan Jordt	danj@nwnet.net
Doug Karl	dkarl@osu.edu
Kenneth Key	key@cs.utk.edu
Charley Kline	cvk@uiuc.edu
Moshe Kochinski	moshek@FibHaifa.com
Mark Laubach	laubach@hpl.hp.com
Mark Lewis	Mark.S.Lewis@telebit.com
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Paul Lustgraaf	grpjl@iastate.edu
Charles Lynn	clynn@bbn.com
Bruce Mackey	brucem@cinops.xerox.com
Carl Madison	carl@startek.com
Gary Malkin	gmalkin@xylogics.com
David Meyer	meyer@ns.uoregon.edu
Greg Minshall	minshall@wc.novell.com
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
John Moy	jmoy@proteon.com
Erik Nordmark	nordmark@eng.sun.com
William Owens	owens@acsu.buffalo.edu
Michael Patton	map@bbn.com
Gaige Paulsen	gaige@intercon.com
John Penners	jpenners@advtech.uswest.com
Willi Porten	porten@gmd.de
David Reese	dave@csu.net
Yakov Rekhter	yakov@watson.ibm.com
Manoel Rodrigues	manoel_rodrigues@att.com
Yzhak Ronen	y.ronen@homxa.att.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Fumio Teraoka	tera@csl.sony.co.jp
Kamlesh Tewani	ktt@arch2.att.com
Richard Thomas	rjthomas@bnr.ca
Susan Thomson	set@bellcore.com
L. Stuart Vance	vance@tgv.com
John Veizades	veizades@apple.com
Dinesh Verma	verma@watson.ibm.com
Warren Vik	wmv@lachman.com
Ruediger Volk	rv@informatik.uni-dortmund.de
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Scott Wasson	sgwasson@eng.xyplex.com
Von Welch	vwelch@ncsa.uiuc.edu
Fred Whiteside	fred@bws.com
Steven Willis	steve@wellfleet.com

Richard Woundy	rwoundy@vnet.ibm.com
Charles Young	Charles.E.Young@att.com

2.2.10 TCP/UDP over CLNP-addressed Networks (tuba)

Charter

Chair(s):

Mark Knopper, mak@merit.edu
Peter Ford, peter@goshawk.lanl.gov

Mailing Lists:

General Discussion: tuba@lanl.gov
To Subscribe: tuba-request@lanl.gov
Archive:

Description of Working Group:

The TUBA Working Group will work on extending the Internet Protocol suite and architecture by increasing the number of end systems which can be effectively addressed and routed. The TUBA effort will expand the ability to route Internet packets by using addresses which support more hierarchy than the current Internet Protocol (IP) address space. TUBA specifies the continued use of Internet transport protocols, in particular TCP and UDP, but encapsulated in ISO 8473 (CLNP) packets. This will allow the continued use of Internet application protocols such as FTP, SMTP, Telnet, etc. An enhancement to the current system is mandatory due to the limitations of the current 32-bit IP addresses. TUBA seeks to upgrade the current system by a transition from the use of the Internet Protocol version 4 to ISO/IEC 8473 (CLNP) and the corresponding large Network Service Access Point address space.

In addition to protocol layering issues and "proof of concept" work, the TUBA approach will place significant emphasis on the engineering and operational requirements of a large, global, multilateral public data network. TUBA will work to maximize interoperatability with the routing and addressing architecture of the global CLNP infrastructure. The TUBA Working Group will work closely with the IETF NOOP and IPRP-for-IP Working Groups to coordinate a viable CLNP-based Internet which supports the applications which Internet users depend on such as Telnet, FTP, SMTP, NFS, X, etc. The TUBA Working Group will also work collaboratively with communities which are also using CLNP, and will consider issues such as interoperability, applications coexisting on top of multiple transports, and the evolution of global public connectionless datagram networks, network management and instrumentation using CLNP and TUBA, and impact on routing architecture and protocols given the TUBA transition.

The TUBA Working Group will consider how the TUBA scheme will support transition from the current IP address space to the future NSAP address space without discontinuity of service, although different manufacturers, service providers, and sites will make the transition at different times. In particular, the way in which implementations relying on current 32-bit IP addresses will migrate must be considered. TUBA will ensure that IP addresses can be assigned, for as long as they are used, independently of geographical and routing considerations. One option is to embed IP addresses in NSAP addresses, possibly as the NSAP end-system identifier. Whatever scheme is chosen must run in a majority of *-GOSIPs and other NSAP spaces. The TUBA strategy will require a new mapping in the DNS from NAMEs to NSAP addresses.

The rationale RFC (RFC-1347) documents issues of transition and coexistence, among unmodified "IP" hosts and hosts which support "TUBA" hosts. Hosts wishing full Internet connectivity will need to support TUBA.

Goals and Milestones:

Done	Post Initial TUBA rational and discussion as an RFC. (RFC 1347)
Done	Post the Initial TUBA DNS specification. (RFC 1348)
Done	Review and approve the Charter.
Done	Post the TUBA CLNP profile as an Internet-Draft.
Done	Post a Routing and Addressing specification as an Internet-Draft, coordinated with the Network OSI Operations Working Group and the IDRP for IP Working Group.
Nov 1992	Post a summary report on TUBA deployment in the Internet.
Done	Present the results of Working Group deliberations at the November IETF meeting.
Nov 1992	Post an Internet-Draft on the changes required to Internet applications affected by the deployment of TUBA.
Nov 1992	Post an Internet-Draft covering the methodologies, instrumentation, address administration, routing coordination and related topics.
Done	Post as an Internet-Draft a revision to RFC1347 reflecting lessons learned in the Working Group deliberation.

Internet-Drafts:

"Use of ISO CLNP in TUBA Environments", 09/04/1992, David Piscitello <draft-ietf-tuba-clnp-02.txt>

"Addressing and End Point Identification, For Use with TUBA", 11/06/1992, R. Callon <draft-ietf-tuba-address-00.txt, .ps>

"Assignment of System Identifiers for TUBA/CLNP Hosts", 04/30/1993, D. Piscitello <draft-ietf-tuba-sysids-00.txt>

CURRENT MEETING REPORT

Reported by Mark Knopper/Merit

Minutes of the TCP/UDP over CLNP-addressed Networks Working Group (TUBA)

Agenda

- Implementation Status and Demonstration.
- Document Status.
- Prioritization of TUBA Work.
  - Questions asked at Opening Plenary
  - Dynamic Host Address Assignment
  - Mobile Hosts
  - Routing and Addressing Plan
  - Transition Strategies
  - Discussion of Technical Advantages of CLNP
- Demo and Implementation Targets.

Implementation Status and Demonstration

The current status of TUBA implementations is:

cisco	Has telnet and finger initiators and responders, tftp initiator, and SNMP agent. The effort took a long weekend, the hardest part being getting the TCP checksum right. Paul Traina indicated that cisco intends to modify tftpd to operate over UDP/CLNP as soon as operating system support is available.
3Com	Has telnet initiator and responder. This work took about one week.
BSDi	Has telnet and SMTP initiators and responders; currently a bit buggy. This implementation is the BSDi distribution with Keith Sklower's modified 4.4 BSD network code.
NCSA Telnet	Has telnet and finger initiators; ftp responder works for command connection (support for data connection is a future work item).
SunOS	Francis Dupont (at INRIA) has grafted the 4.4 BSD modified network code onto SunOS 4.1.2, and has added support for UDP over CLNP. No application information was available (Francis was not at the TUBA Working Group meeting). Francis has also modified tcpdump to understand TUBA; contact Francis.Dupont@inria.fr for details.

AIX 3.2 IBM ported the 4.4 BSD modified network code to AIX 3.2. Merit will be testing the port. Yakov Rekhter will modify ftp for TUBA after Merit completes the kernel work. It wasn't clear what the status is for other applications.

The cisco, 3Com, BSDi, and NCSA Telnet implementations were running in the IETF terminal room. CLNP connectivity was available from the terminal room via an NSFNET EON encapsulator to other TUBA hosts at:

- cisco via Barrnet
- 3Com via SURANet and COS
- NIST via SURANet
- Merit via the NSFNET
- LANL via ESNet
- NORDUNET and other sites in Europe

Existing Document Status

- RFC 1347 (the original TUBA proposal): No identified changes.

- "CLNP for TUBA" Internet-Draft (draft-ietf-tuba-clnp-02.txt): Dave Piscitello will polish the pseudo–header checksum calculation description.

  Dino Faranacci suggested that the Group needs to think about MTU discovery. The Group might want to use the ER PDU to return the MTU size.

  The idea of padding the CLNP header to obtain word alignment for the TCP header was reopened briefly. It was decided that this had already been discussed in the past and the Group would stick to the conclusion that this is not something that can be guaranteed, given the number of different subnet services that CLNP operates over.

  Given the implementation experience, the Group decided that it would ask for this document to be moved to Proposed Standard. Dave Piscitello will take this as an action.

- "Addressing and End Point Identification, For Use with TUBA" (draft-ietf-tuba-address-00.ps): Everyone should go back and (re)read this and send comments to the mailing list.

- "DNS NSAP RRs" Internet-Draft (draft-manning-dns-nsap-01.txt): This Internet-Draft is the successor to RFC 1348. It contains a better treatment of the inverse mapping for NSAPs than was in 1348, but this aspect is still subject to change. [Note: Bill Manning has posted this Internet-Draft already.]

New Documents

- Catalog of TUBA implementations: We decided that it would be useful to collect the information about what implementations are available and who to contact. Mark agreed to take this as an action.

- CLNP changes from London ISO meeting: There was a document describing possible changes for CLNP that was distributed in a recent SC6 meeting in London. Mark took the action of getting a copy on-line.

- TUBA Frequently Asked Questions: In keeping with the theme of needing better organization of the TUBA documentation, Mark suggested we write a FAQ. Mark will produce a first draft.

- CLNP Multicast work: SC6 is working on multicast extensions for CLNP and related routing protocols. Radia Perlman said she will ask Dave Oran to post a summary status of this work on the mailing list.

Prioritization of TUBA Work

Several questions were asked during the Opening Plenary.

1. What upper layer changes are necessary?

   The core applications – including FTP, SMTP, TELNET, and DNS – were mentioned. It was decided that we should create a single document that catalogues what changes, if any, need to be made to these for TUBA. In most cases, the required changes are minimal. Mark agreed to take a first cut at this document. Dave Piscitello agreed to provide the FTP-specific section. Peter Ford, Yakov Rekhter, and Richard Colella agreed to modify FTP from this specification.

   Keith Sklower mentioned a draft description of a replacement for gethostbyname that he and Eric Allman had devised. Called getconninfo, it is more general than gethostbyname, accommodating address families other than AF_INET. This will make TUBA (and other IPng proposals) more transparent to the applications. Keith agreed to post the write-up as an Internet-Draft.

2. What is the transition scheme?

   Most of this discussion focused on a problem that John Veizades sees: there is a community of users that does not generally have the resources necessary to upgrade their small, older routers to accommodate CLNP to support TUBA (e.g., universities). After some discussion it became clear that, whereas some thought that this was not a serious issue, John was not convinced. Dino Faranacci and John agreed to take this particular issue off-line. In any case, it was clear that the TUBA work needs a transition document to answer just this kind of question. Peter Ford and John Curran agreed to draft a transition plan.

3. Address assignments – how do we get them?

   This question is fully answered by the NSAP allocation scheme outlined in RFC 1237, Guidelines for OSI NSAP Allocation in the Internet, July, 1991. There is already a well-defined method of obtaining and assigning NSAP addresses. In the U.S., address space can be obtained from either the US GSA or from ANSI.

4. How does TUBA address mobile hosts?

   Deferred due to lack of time.

5. Are there any known boundary conditions?

   There were no known boundary conditions involving TUBA.

6. What about scaling?

   In response, reference was made to a seminal paper from 1971 by Kleinrock.

   Stev Knowles asked, "What if you have one million networks? How does CLNP and its routing protocols handle this?" A lively discussion ensued; there was not a specific response as it's a complex question.

It was agreed that the Group should discuss the topics of scaling and mobile hosts.

Discussion of Technical Advantages of CLNP

Radia Perlman wanted to make the point that we need to recognize the technical strengths of CLNP. She enumerated three in particular.

1. Auto-configuration – By using a unique System ID in the NSAP, it is relatively easy to do address auto-configuration. This would greatly reduce administrative overhead in assigning and changing addresses, and allow for easier portability of systems.

2. Infinite scaling property – Given the size and flexibility of NSAP addresses, address prefix routing provides a large number of potential levels in the routing hierarchy, assuming that prefixes are based on nibble boundaries.

3. "Free" routing across WANs – Embedded subnet addressing can be used to simplify routing in environments that make use of WANs for interconnection. This entails assigning NSAPs with a WAN-based subnet address in the high-order part of the NSAP. The WAN-based part of the subnet address would then be used to perform the

cross-WAN routing hop (e.g., from one routing domain to another, both connected to the same WAN). Note that domains not connected to the same WAN would continue to route using the normal routing protocols (i.e., ISIS and IDRP).

Dynamic Host Address Assignment

One part of the solution to dynamic host address assignment is ES-IS, which is reasonably straightforward. Bill Warner agreed to draft text that describes how ES-IS is used to do dynamic address assignment.

Another part of dynamic host address assignment is how to get the information into DNS. This is not so obvious. John Curran agreed to write some text for this.

Routing and Addressing Plan

Ross Callon wrote a routing and addressing Internet-Draft for TUBA in October. Everyone was assigned to (re)read this and send comments to the list (see Internet-Draft draft-ietf-tuba-address-00.[txt,ps]).

The subject of globally unique EIDs was raised once more. There was violent agreement that we should do this in the NSAP System ID field. However, there was some disagreement on the mechanics. Ross suggested mandating that the System ID field be taken from a single, globally-coordinated 48-bit number space (*not* synonymous with IEEE MAC addresses). Keith had a somewhat different idea, allowing variable size EIDs and, hence, variable sized System IDs. Each proponent was asked to write a short description of their proposal and post it to the mailing list. Dave Piscitello agreed to write up Ross's proposal.

Demonstration and Implementation Targets

It was recognized that the TUBA demonstrations could benefit from better planning and coordination. George Chang agreed to take the lead in this area.

Summary of Action Items

Dave Piscitello	Update the CLNP for TUBA document and submit as a Proposed Standard.
	Write the FTP for alternative network layers specification. The implementation will be done by Peter Ford, Yakov Rekhter and Richard Colella.
	Write the EID administration text.
Ross Callon	Update the Addressing document with comments solicited from Group.

Manning and Colella	Update the DNS extensions for NSAPs Internet-Draft (RFC1348 update).
Mark Knopper	Write a catalog of TUBA implementations.
	Make the updated CLNP document from the London ISO meeting available.
	Write a TUBA Frequently Asked Questions document.
	Write a document on what changes need to be made to each application.
	Update the Group weekly on the status of each action item. (This item refers to the offer Mark made to remind the Group periodically on the status of each action item and what is left to be done.)
Radia Perlman	Report on the status of the CLNP Multicast work.
Paul Traina	Implement the Tftpd.
Keith Sklower	Write the getconninfo document as a replacement for gethostbyname.
Ford and Curran	Write the transition document.
Bill Warner	Write the auto-config (dynamic host address assignment using ES-IS) specification.
John Curran	Write the NSAP insertion into DNS text. The implementation will be handled by Dave Piscitello.
George Chang	Provide PR and coordination for the demonstration.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Jim Barnes	barnes@xylogics.com
Russell Blaesing	rrb@one.com
Rebecca Bostwick	bostwick@es.net
George Chang	gkc@ctt.bellcore.com
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
William Chimiak	chim@relito.medeng.wfu.edu
Richard Colella	colella@nist.gov
Michael Collins	collinsms@es.net
John Curran	jcurran@nic.near.net

Dino Farinacci	dino@cisco.com
Eric Fleischman	ericf@act.boeing.com
Francois Fluckiger	fluckiger@vxcern.cern.ch
Peter Ford	peter@goshawk.lanl.gov
Vince Fuller	vaf@stanford.edu
Peter Furniss	p.furniss@ulcc.ac.uk
John Gawf	gawf@compatible.com
Eugene Geer	ewg@cc.bellcore.com
Tony Hain	alh@es.net
Susan Hares	skh@merit.edu
Woody Huang	yuh@merit.edu
David Jacobson	dnjake@vnet.ibm.com
Laurent Joncheray	lpj@merit.edu
Mark Knopper	mak@merit.edu
Paul Lustgraaf	grpjl@iastate.edu
Carl Madison	carl@startek.com
Tracy Mallory	tracym@3com.com
Bill Manning	bmanning@sesqui.net
Jun Matsukata	jm@eng.isas.ac.jp
David Meyer	meyer@ns.uoregon.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
Peder Chr. Noergaard	pcn@tbit.dk
Laura Pate	pate@gateway.mitre.org
Maryann Perez	perez@cmf.nrl.navy.mil
Radia Perlman	perlman@dsmail.enet.dec.com
David Piscitello	dave@mail.bellcore.com
Willi Porten	porten@gmd.de
Yakov Rekhter	yakov@watson.ibm.com
Ben Robinson	ben_robinson@vnet.ibm.com
Yzhak Ronen	y.ronen@homxa.att.com
Michael Safly	saf@tank1.msfc.nasa.gov
Paul Serice	serice@cos.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Steve Suzuki	suzu@fet.com
Wayne Tackabury	wayne@cayman.com
John Tavs	tavs@vnet.ibm.com
Kamlesh Tewani	ktt@arch2.att.com
Richard Thomas	rjthomas@bnr.ca
Paul Traina	pst@cisco.com
John Veizades	veizades@apple.com
William Warner	warner@ohio.gov
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net
Charles Young	Charles.E.Young@att.com

2.2.11 TP/IX (tpix)

Charter

Chair(s):
Vladimir Sukonnik, sukonnik@process.com

Mailing Lists:
General Discussion: tpix@world.std.com
To Subscribe: tpix-request@world.std.com
Archive: world.std.com:~/pub/tpix/*

Description of Working Group:

TP/IX is a new version of the IP and TCP/UDP protocols, to advance the Internet technology to the scale and performance of the next generation of internetwork technology. TP/IX has been assigned the IP Version Number 7.

The Working Group is chartered to review the TP/IX and RAP protocols, evaluate issues arising during product development and deployment planning, and to document problems and explanations for any parts of the coexistance with IPv4 not covered directly in the TP/IX-IPv4 interoperation design.

The Group will also be the initial forum for development of the RAP protocol while it is experimental; this work will need to be moved to the Routing Area when it is to be advanced.

Goals and Milestones:

Done	Present the TP/IX (formerly IPv7) and the RAP protocols to the IETF Plenary.
May 1993	Post the TP/IX Protocol and the RAP protocol as Experimental RFCs.
Jul 1993	Hold Working Group meeting to discuss additional definitions. Prepare criteria to be met prior to standardization.
Nov 1993	Hold Working Group meeting to evaluate the TP/IX and RAP protocols for Proposed Standard.
Dec 1993	Submit the TP/IX and RAP Protocols to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"TCP/IP: Internet Version 7", 08/12/1992, R. L. Ullmann <draft-ullmann-ipv7-03.txt>

"RAP: Internet Route Access Protocol", 02/05/1993, R. Ullmann <draft-ullmann-rap-01.txt>

"Initial AD Assignment Plan", 06/07/1993, R. Ullmann <draft-ietf-tpix-adplan-00.txt>

2.3 Network Management Area

Director(s):

- Marshall Rose: mrose@dbc.mtview.ca.us

Area Summary reported by Marshall Rose/Dover Beach Consulting

The Working Groups and BOFs which met in Columbus are as follows:

Frame Relay Network MIB BOF (FRNETMIB)

The Frame Relay Network MIB BOF met to discuss whether there was interest in a standard set of objects for Frame Relay CNM. Work in this area started in the Frame Relay Forum and is continuing there. There was consensus that a working group should be formed to align with this work and to ultimately produce a MIB in this area.

IFIP Electronic Mail Management BOF (EMAILMGT)

The EMailMgt BOF met in three sessions with a cumulative participation of thirty-two people. Two documents (Requirements & Model) were reviewed and revised to align terminology and concepts. Work between now and the next meeting in June will be done via email in the EMailMgt mailing list. The next face-to-face meetings are planned for NIST OIW in June (7-11) and Amsterdam IETF in July (12-16). Our objective is to complete and publish our Requirements and Model documents (as Informational RFCs) by the end of the July IETF meeting. Work will then refocus on support of coordinated development of appropriate MO/MIB specifications. In the interim, we will coordinate MO/MIB development efforts with cross-participation in the involved working groups.

Mail and Directory Management BOF (MADMAN)

The issue of a Charter was discussed. Four documents were suggested for development: (1) a network application MIB model, (2) a MIB module for MTAs, (3) a MI module DSAs, and possibly, (4) a MIB module for message stores. Draft documents describing the first three MIBs were available (although two of them showed up "just in time"); these documents were reviewed and minor changes were made. Finally there was a long discussion concerning what might go into a Message Store MIB module.

Managing ATM with SNMP BOF (ATMMIB)

The ATMMIB BOF was held to discuss the need to standardize managed objects for ATM management. The need to standardize managed objects for SONET equipment was also discussed. The BOF was chaired by Kaj Tesink, and featured presentations by Masuma Ahmed on a strawman ATM MIB, Keith McCloghrie on the ILMI of the ATM Forum, and

Kaj Tesink on the Internet-Draft for a SONET MIB. The Group recommended unanimously to form a working group (the AToMMIB Working Group) that is chartered to produce MIBs on these topics.

Modem Management BOF (MODEMMGT)

The Modem Management BOF discussed monitoring and control issues for modem devices and also examined what currently defined MIB modules could be used in support of this. There was also discussion of alignment with other organizations. There was consensus that a working group be chartered to complete this work.

SNA Systems Management BOF (SNAMIB)

The three SNAMIB BOF sessions were well attended. The primary objective was to assess the community interest in developing standard MIBs for SNMP management of SNA systems, protocols and data links. Strong interest was expressed by all the attendees; in addition, thirteen parties came forward with resource commitments for the work needed. Hence, it was decided that working groups should be formed. Attention was then focused on identifying the SNA systems, protocols and data links that should be worked on at this time. The considerations that were applied to this discussion included resource availability, establishing a track record (with IETF) by taking on (and succeeding in) manageable amount of work and priority of the system/device/data links in question. The Group then identified two prospective working groups, one to focus on the two most important SNA NAU services protocols and the other to focus on the two most important SNA data links. Editors and Chairs were identified for the two working groups. Finally, as an initial milestone, dates were set for vendors to contribute their MIBs.

Chassis MIB Working Group (CHASSIS)

Progress continues but remains slow. The Group discussed a presentation on generalization of the MIB model as introduced at the previous meeting and partly integrated into the draft. There was fair consensus, but some unsureness of understanding. The new model will be integrated into the draft and Group members must check their chassis implementations against it. A developer reported mostly positive experience implementing the proposed Chassis MIB, but provided a few suggestions.

FDDI MIB Working Group (FDDIMIB)

The Working Group reviewed the current Internet-Draft. Several minor revisions were made. There was consensus that a new Internet-Draft be posted, briefly reviewed by the mailing list, and then submitted for consideration as a Proposed Standard. There was some interest in a traps document. A strawman will be developed.

IEEE 802.3 Hub MIB Working Group (HUBMIB)

The HUB MIB Working Group met to discuss minor issues concerning both the MAU MIB Internet-Draft (3/22/92) and the Repeater MIB Proposed Standard (RFC1368). All issues were resolved in principle, with exact working of changes to be done by the editors. The editors will then mail new drafts of both documents to the Working Group mailing list for three weeks review. If no unresolved issues surface during that time, the MAU MIB will be submitted to the IESG for consideration as a Proposed Standard, and the Repeater MIB will be submitted to the IESG for consideration as a Draft Standard. (Note that nine implementors of the Repeater MIB were represented at the meeting, and the Working Group agreed that the implementation and operational experience with the Repeater MIB was more than sufficient to warrant its forwarding to Draft.)

Token Ring Remote Monitoring Working Group (TRMON)

The TRMON Group met once to identify and resolve the final outstanding technical issues for the draft. There was consensus that the resulting draft should be submitted to the Network Management Directorate for eventual publication as a Proposed Standard. The Group then discussed priorities for future work and where a next meeting might take place. There was no clear resolution on these issues. Finally, in the remaining minutes, a few implementation issues for RFC1271 were discussed.

Uninterruptible Power Supply Working Group (UPSMIB)

A review of the strawman document published just before the meeting became a discussion of a single counter for line failures, which took most of the meeting before deciding that "failures" counted would be vendor specific, as there was no common ground (not speaking electrically). The discussion turned briefly to the model for input lines, with consensus and instructions to the editor. One vendor reproposed the idea of a very small MIB for the simplest UPS, again requesting multiple MIBs. The Chair pointed out that SNMPv2 compliance groups are the answer, and are motivated by exactly such problems. In the interest of moving faster, individuals volunteered to be spokesmen to push progress for each of the MIB groups.

CURRENT MEETING REPORT

Reported by Einar Stefferud/NMA and Ray Freiwirth/RCI

Minutes of the IFIP Electronic Mail Management BOF (EMAILMGT)

Three EMAILMGT BOF Sessions were held at the Columbus IETF, under a special arrangement for the IFIP WG 6.5/6.6 Chartered EMailMgt Working Group to meet with the IETF to both draw on IETF attendance for participation, and to bring additional IFIP participation into IETF.

IFIP WG 6.5/6.6 and the EMAILMGT participants are very pleased and grateful for the opportunity to serve and support IETF interests, and to draw on IETF resources.

First Session

Several Documents were distributed during the first BOF session, and the mailing lists for the various EMailMgt Task Teams and working groups were announced. *Editor's Note (md): A complete listing of the documents and mailing lists is available via ftp under emailmgt-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Ray Freiwirth led a review of the current EMailMgt Draft Requirements Document (EMGT 93-006). The emphasis of the meeting was on terminology and making sure that everybody understood the terms as used, and everyone understood the functions associated with the terms.

One major area of discussion was how to identify a user that becomes an "email manager" for some functions that are allowed for that user by the real manager. It was decided that no special term is needed for such a user.

The following sentence was added to the definition of "user":

> "Has capability to monitor its own mailbox, local environment and remote logs, files, etc. as may be allowed".

Some current definitions will be further modified with regard to minor spelling/phrasing problems. The concept of a Relative Domain is still being discussed. It is not clear if the concept and definition of Manager Responsibility Area (MRA) needs to be expanded. The Group is trying to avoid the confusion that would follow from using the term "domain" in yet another context with yet another meaning.

It was noted that it is important for the document to state clearly that the EMailMgt requirements are not dictating a new kind of management, but rather calls for use of existing methods and tools to meet EMailMgt requirements.

Second Session

Harald Alvestrand lead a discussion of his Draft EMailMgt Modeling Document (EMGT 93-002).

The concept of a gateway needs to be clearly defined in the diagrams to show how gateways fit into the email infrastructure (e.g., when they sit astride two different email environments). Omission of gateways would imply that email gateways are outside the scope of EMailMgt! They obviously are not.

It was generally agreed that most of the model diagrams need revision. The dataflow diagram needs more work to indicate interchange between Management Responsibility Areas.

Some managers will use information to reconfigure systems, which implies that there are different time-frames for different data flows.

The following questions were raised regarding general terminology.

- Why not define "customer" as defined in the English language?

  Answer: Because a clear distinction needs to exist between a user and a customer. A customer makes value judgments. Users do not, unless they also happen to be customers at the same time.

- Why not define mail service as just a mail transfer service?

  Answer: Inside the mail system there are many objects to be managed. Some of these objects might be managed by a single email manager, by a group of email managers or by a "user/manager" or just by a user.

- Where does Message Store fit into this model?

  Answer: There is a need to identify split User Agent and Mail Box functionality, and call out that a user can manage part of the mailbox. The model needs better definition of users having some of the capabilities of a manager, and of users having some manager roles.

In Section 3.2, the diagram needs to be expanded to show all the services that were identified in the requirements document (i.e., security, routing, etc.). This relates directly to the ability to manage the email portion of the data that resides in any of the services, and the ability to use their services: directory; network management; logging; etc.

Detailing of both MTA and UA model diagrams should be modified to show more dimension with respect to sources, queues and sinks (flow detail).

Third Session

This session was dedicated to reconciling all differences between the use of terms in the requirements and modeling documents. This was determined to be the key high priority objective of this set of EMailMgt meetings.

Ray Freiwirth lead the discussion, following the requirements document. With the work of the previous two days and a better understanding of the model document, great strides were made in reconciliation.

Requirements Section 2 needs to be better aligned with the model document, especially the fourth paragraph.

For requirements Section 3.3.2 (Remote Email Service), there was a general discussion about IMAP. Somebody should cross-participate in the IMAP Group to make sure and that both Groups are aligned.

The number one goal for both documents is to achieve alignment on the one hand, and comprehensiveness on the other. The Group is working to identify all the relevant elements and entities that require management and show how they relate to to each other in the overall model.

Final Observations

Goals for this meeting could not have been achieved without holding three separate sessions on three separate days. The final session on Friday was critical to pull everything back together in the end.

It is noted that the next EMailMgt meetings are scheduled for June (OIW at NIST) and July (IETF at Amsterdam). Other meetings of EMailMgt task teams are also planned, and will be announced on the EMailMgt mailing list.

Since the next meetings are several months into the future, plans have been made to complete edits of the EMailMgt requirements and modeling documents and publish them as Internet-Drafts to obtain wider distribution and to facilitate more robust discussions on the main mailing list. The Internet-Drafts should be published by the end of April at which time the process of review, comment, revision, and adoption of these documents using consensus methods in the EMailMgt mailing list will begin, based on the published Internet-Drafts.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jules Aronson	aronson@nlm.nih.gov
Robert Beer	r-beer@onu.edu
Richard Bjers	rich.bjers@uc.edu
Cyrus Chow	cchow@ames.arc.nasa.gov

Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Urs Eppenberger	eppenberger@switch.ch
Erik Fair	fair@apple.com
Francois Fluckiger	fluckiger@vxcern.cern.ch
Ned Freed	ned@innosoft.com
Raphael Freiwirth	5242391@mcimail.com
Marcello Frutig	frutig@rnp.impa.br
Christine Garland	garland@ihspa.att.com
Kenneth Goodwin	goodwin@a.psc.edu
Jeroen Houttuin	houttuin@rare.nl
Barbara Jennings	bjjenni@sandia.gov
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Kenneth Key	key@cs.utk.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Bruce Mackey	brucem@cinops.xerox.com
Ignacio Martinez	martinez@rediris.es
Keith Moore	moore@cs.utk.edu
Brad Passwaters	bjp@sura.net
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Yzhak Ronen	y.ronen@homxa.att.com
Gary Rowe	gjrowe@attmail.com
Chris Shaw	cshaw@banyan.com
Sue Smith	smiths.es.net
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de

CURRENT MEETING REPORT

Reported by Barbara Sterling/McDATA

Minutes of the SNA Systems Management BOF (SNAMIB)

The meeting opened with an introduction of all attendees. A roster, as well as an interest grid, was circulated. Baktha Muralidharan went over the Agenda which was then passed without dissent. The criteria for forming a working group for a MIB were discussed. Basically at least three or more people (one Chair, one author, one editor) who are willing to participate actively are needed before a working group can be formed. In addition, a working group's charter preferably should include milestone dates.

A poll was taken to show interests in the various areas of SNA. *Editor's Note (md): The results of the poll are available via ftp under snamib-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Note: This poll only reflects the interests of the attendees. It was pointed out that for some areas, such as APPI, there may be other vendors who are interested but not present.

The following issues were brought up during the meeting:

- How well will the MIBs reflect the client's point of view?
- The relationship between SNMP and NetView:
  - How do the two relate to each other: does it mean sending SNMP information to NetView?
  - What is the scope of management by each protocol?
- Splitting the host management aspect from PU 5. Some felt that there are wider interests in managing SNA topology networks than SNA hosts.
- Placement of related areas into one or more separate MIBs. For example, should APPN End Node be a separate MIB from the APPN Network Node? Should PU 2.1 belong to the APPN MIB (since it describes LEN) or to the PU MIBs (with PU 2.0)? It was decided to defer detailed MIB organization discussion until later.

Marshall Rose pointed out the differences between proprietary and standard MIBs. A standard MIB should never contain vendor-specific details but should be a core set of information common to all. In addition, a standard MIB loses meaning if it is not widely adopted and implemented by the community.

Session Two

The meeting was opened by Baktha, who recapped session one. Baktha proposed a working group structure composed of four distinct working groups, based upon the discussions of session one:

WG1	PUT2.0, PUT2.1, APPN End Node, APPN Network Node, APPI
WG2	SDLC, LLC-2, Channel, QLLC
WG3	APPC, LUs
WG4	Data Link Switching (DLSw)

A series of discussions ensued, involving:

- The number of working groups needed or desired – issues included:
  - There is a need for overall architectural control to ensure all MIBs complement each other.
  - Considerable overlap of participation is expected if there are multiple working groups.
  - How to minimize the size and number of mailing lists, including having one mailing list even if multiple working groups are formed.
  - To ensure communication between multiple working groups, each working group would have other working groups review and sign-off documents prior to publishing them.
  - Relevance of number of vendors/users present at BOF versus other interested vendors/users in determining priority of certain MIB structures.
  - Moving DLSw to WG2.
  - Moving PUT2.0 to WG3.
  - LUs topic should include LU types 0,1,2,3,4,6.2,7.

- IBM indicated that, in approximately one month, it intends to submit its current APPN MIB as a Draft Standard in order to aid network management vendors in preparing to support IBM's APPN functions for the 6611 router. Discussion followed, including:
  - Is there a real need to propose this MIB as a standard now, assuming a working group is to be formed that addresses this issue?
  - Should the APPN MIB be published as an informational draft instead?

- The IESG and/or IAB is likely to not accept IBM's submission as a Draft Standard and will refer the matter to the established working group.
- IBM also has an APPC MIB that they plan to submit to the IETF in the future.

- The Group developed consensus in these areas:
  - There will be two working groups:
    - WG1 – Logical SNA Protocols (includes PUT2.0, PUT2.1, APPN EN, APPN NN, APPI, APPC, all LUs).
    - WG2 – Data Link Layer (includes SDLC, LLC-2, Channel, QLLC, DLSw). The Charter of the group will be to examine the issues relating to and publish the MIBs required to enable management of logical SNA protocols and their data link layers by SNMP, defining capabilities that are similar to those provided to IBM's NetView network management product.

- Milestones for this Group include:
  - The working groups will publish a draft document to the mailing list identifying how many MIBs are to be defined within the scope of this effort. This draft will be published within two months of the formation of the working groups.
  - During July 1993, first drafts of the MIBs will be reviewed by the mailing list.
  - The working groups will meet at the July 1993 IETF meeting.
  - One mailing list will be defined that will include both working groups.

A suggestion was made that the Chair post a "request for participation" for these working groups to the IETF mailing list, encouraging both vendors and users to participate.

Discussions relating to staffing the two working groups were deferred to BOF Session Three.

Session Three

It was decided to make the Charters more specific and focused than had previously been decided. Proposed WG1 will work on MIBs for PUs 2 and 2.1 and LUs 1,2 and 3. Proposed WG2 will work on MIBs for SDLC and LLC-2 data link protocols. The Charters for these first two groups are not meant to preclude the creation of future working groups to develop "SNA MIBs" in other areas of common interest.

It was mentioned that the MIB for LLC-2 will need to be coordinated with the similar IEEE MIB definition effort.

Several individuals volunteered to edit and/or author the documents and other volunteered to chair the working groups. *Editor's Note (md): A list of the individuals who volunteered* *for these tasks is available via ftp under snamib-minutes-93mar.txt. Refer to Section 1.2 for retrieval instructions.*

Action Items

- Baktha Muralidharan and Deirdre Kostick are to work on appropriate wording of the Charters and distribute drafts to the mailing list.

- Authors are to post their MIBs for WG's 1 and 2 by April 30th.

Attendees

Michael Allen	moallen@ralvmg.vnet.ibm.com
David Arneson	arneson@ctron.com
David Battle	battle@cs.utk.edu
Mahesh Bhatia	bhatia@ctron.com
Fred Bohle	fab@interlink.com
Michael Bowman	meb@netlink.com
Jeff Case	case@cs.utk.edu
Jia-bing Cheng	cheng@ralvm6.vnet.ibm.com
Anthony Chow	chow_a@wwtc.timeplex.com
Wayne Clark	wclark@cisco.com
Tracy Cox	tacox@sabre.bellcore.com
Wayne Cullen	wnc@netlink.com
Kishan Dudkikar	kishan@icm1.icp.net
Eric Fleischman	ericf@act.boeing.com
Cleve Graves	cvgpc@oc.com
Jeff Hilgeman	jeffh@apertus.com
Bill Kelly	kellywh@mail.auburn.edu
Mark Kepke	mak@fc.hp.com
Kenneth Key	key@cs.utk.edu
Zbigniew Kielczewski	zbig@eicon.qc.ca
Moshe Kochinski	moshek@FibHaifa.com
Deirdre Kostick	dck2@sabre.bellcore.com
William Kwan	kwan@rabbit.com
Patrick Leung	patrickl@eicon.qc.ca
William McKenzie	mckenzie@ralvma.vnet.ibm.com
Robert Moskowitz	3858921@mcimail.com
Satinder Mundra	mundra@ctron.com
Baktha Muralidharan	murali@smaug.enet.dec.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Tom Nisbet	nisbet@tt.com
Shannon Nix	sdn@netlink.com
Bill Norton	wbn@merit.edu

Eric Olinger	eric@peregrine.com
Jon Penner	jjp@bscs.uucp
David Perkins	dperkins@synoptics.com
Thomas Pusateri	pusateri@cs.duke.edu
Owen Reddecliffe	owen%wrq@mcimail.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Joseph Rumolo	attmail!jrumolo
Chris Shaw	cshaw@banyan.com
Kitty Shih	kmshih@novell.com
Timon Sloane	timon@timon.com
Stuart Stanley	stuarts@apertus.com
Barbara Sterling	bjs@mcdata.com
Stephen Tsun	snt@3com.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
James Watt	james@newbridge.com
Kiho Yum	kxy@nsd.3com.com

SNA Peer-to-Peer Networking BOF

Internet Engineering Task Force
Twenty-sixth Meeting
Columbus, OH

Wayne Clark

cisco Systems, Inc.
Engineering Department
1525 O'Brien Dr.
Menlo Park, CA 94025
wclark@cisco.com

Thursday, April 1, 1993
1:30 pm

Agenda

1:30 - 1:35	Introductions
1:35 - 1:45	Connection Networks
1:45 - 2:00	Data Link Switching
2:00 - 2:15	APPN TCP/IP
2:15 - 2:30	APPN, DLS and Connection Networks
2:30 - 3:15	APPI
3:15 - 3:30	Question and Answer

APPN Intermediate Session Routing

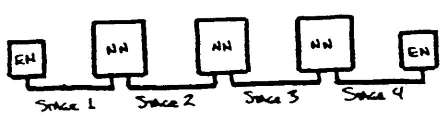

- End-to-end connection is a series of concatenated session stages
- Each stage is:
  - paced, for flow control
  - segmented, for link considerations
- Routing and re-routing is done at the SNA session level

APPN Connection Networks

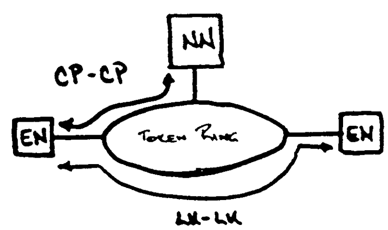

- Optimization for Shared Access Transport Facilities (SATF)
- Result of Locate Request is: (MAC Address, SAP) pair
- LU-LU session established with a direct LLC2 session EN-to-EN
- LU-LU session establishment sequence begins with TEST to destination = MAC Address

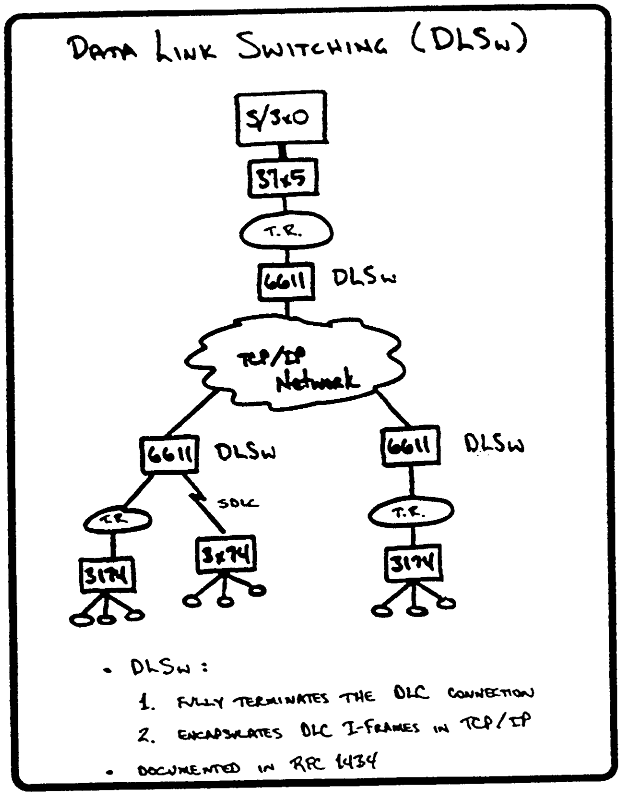
DATA LINK SWITCHING (DLSw)
S/3x0
37x5
T.R.
6611
DLSw
TCP/IP Network
6611
DLSw
6611
DLSw
SDLC
T.R.
T.R.
3174
3x74
3174
• DLSw:
1. FULLY TERMINATES THE DLC CONNECTION
2. ENCAPSULATES DLC I-FRAMES IN TCP/IP
• DOCUMENTED IN RFC 1434

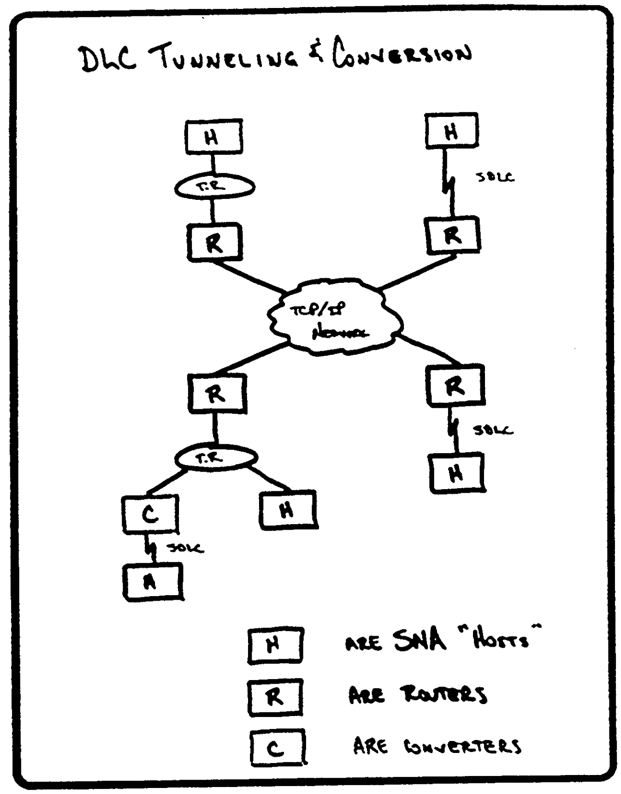
DLC TUNNELING & CONVERSION
H
H
SDLC
T.R.
R
R
TCP/IP Network
R
R
SDLC
T.R.
H
C
H
SDLC
A
H ARE SNA "HOSTS"
R ARE ROUTERS
C ARE CONVERTERS

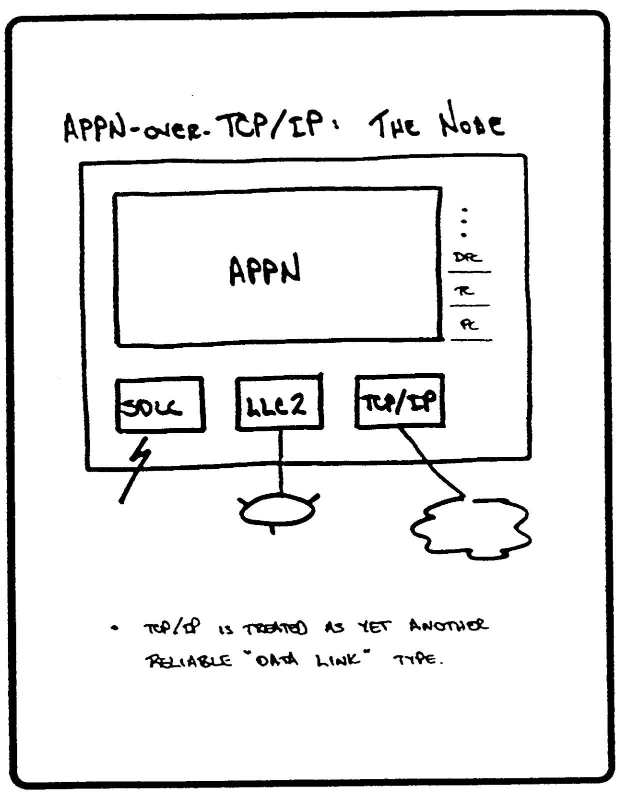
APPN-over-TCP/IP: THE NODE
APPN
DR
TR
FR
SDLC
LLC2
TCP/IP
• TCP/IP IS TREATED AS YET ANOTHER RELIABLE "DATA LINK" TYPE.

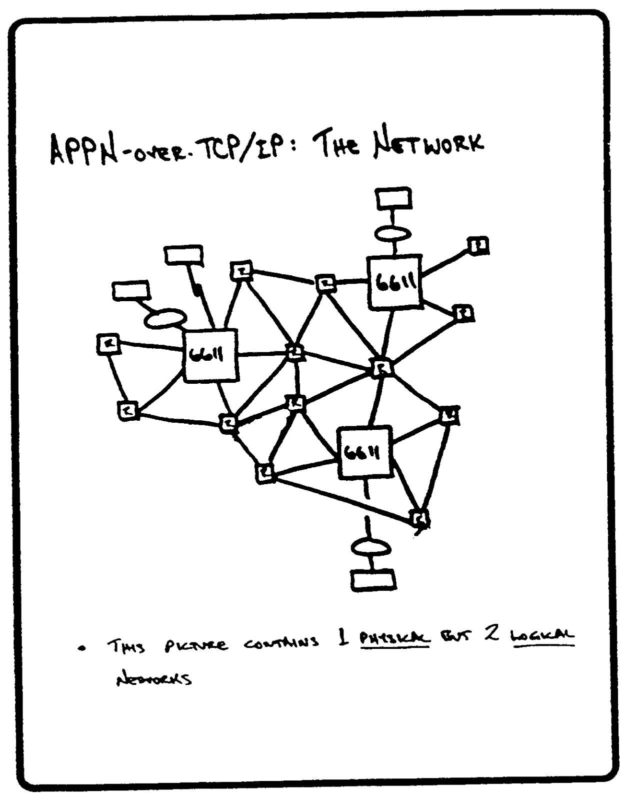
APPN-over-TCP/IP: THE NETWORK
6611
6611
6611
• THIS PICTURE CONTAINS 1 PHYSICAL BUT 2 LOGICAL NETWORKS

APPN/DLS AND CONNECTION NETWORKS

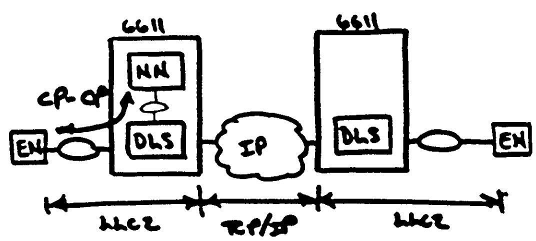

Why APPI?

- industry issues
  - proprietary APPN NN specifications
  - Is the APPN Implementor's Workshop really open?
  - Is APPN a level playing field?

- APPN technical limitations:
  - no adaptive routing/rerouting (yet)
  - limited media support
  - limited track record

- APPN user / marketplace issues:
  - integration with current router networks
  - divergence of SNA and LAN internets
  - yet another backbone protocol
  - yet another transport protocol
  - vendor dependency

The APPI Proposition

- industry perspective
  - open and, therefore, a level playing field

- technical perspective
  - dynamic adaptive rerouting (session integrity)
  - full media support
  - no new routing protocol
  - no new transport protocol (proven track record)

- user / marketplace perspective
  - well integrated into existing router networks
  - consistent with SNA / IP user directions
  - no need to deploy another backbone protocol
  - available from multiple vendors
  - available on a variety of platforms

The APPI Forum Charter

- produce APPI interoperability specifications
  - based upon open standards (where possible)
  - deviations from standards will be documented
  - available to all interested parties

- promote industry cooperation
  - achieve wide availability of APPN and APPI products and services
  - plan for interoperability through testing
  - reach consensus on design and implementation details
  - coordinate input to standards bodies

- provide venue for APPI-related standards development

The APPI Forum is OPEN

The APPI Forum is open to all parties willing to cooperate on the development of the APPI interoperability specification.

The interoperability specification will be available to all interested parties.

APPI Forum Current Membership

- **Voting Members**

Alcatel
Arkhon Technologies
Ascom/Timplex
British Telecom
Cabletron
Cascade Communications
CompuServe
Cisco Systems
Data Connection, Ltd.
Digital Equipment Corp.
Hewlett-Packard
Infonet
McData
MCI
NCR / AT&T
Netrix
Proteon
Rabbit
Sprint
SunConnect
SynOptics
Unisys

- **Auditing Members**

Computer Communications, Inc
Eicon
Proginet Corp.
SourceCom Corp.

State of APPN

Implementation	LEN	EN	NN
IBM			
AS/400	√	√	√
OS/2 2.0 ES 1.0	√	√	√
3174 w/ Config Sppt C	no	why not?	√
VTAM/NCP	√	1993	1993
6611	n/a	n/a	1993
RS/6000	√	1993	1993
OEM	no	yes?	1993
NS/DOS	√	n/a	n/a
Non-IBM			
System Strategies, Inc. (OEM)	√	1993	n/a
Data Connections Ltd. (OEM)	√	1993	1993
Novell - Netware for SAA	√	no	1993?
Eicon - SNA LAN Gateway	√	n/a	n/a
DCA - Select Comm Server	√	n/a	n/a
3Com Corporation	no	no	1993
N.E.T.	no	no	?
Apple Computer	√	SOD	?

APPI Network Topology

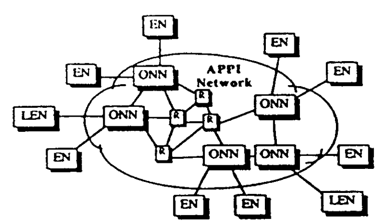

- an APPI network is a set of interconnected Open Network Nodes (ONNs) and multiprotocol routers
- ONNs are merely intermediate systems with APPN enabling software
- ONNs look like APPN NNs to the ENs and LENs
- APPI offers an open alternative to APPN:
  - APPN NN Link State –> existing routing protocol
  - APPN Distributed Directory Service –> APPI Distributed Directory Service
- APPI's data transport services will be open systems transport services

EN-to-ONN Flows

(CP-CP Session)

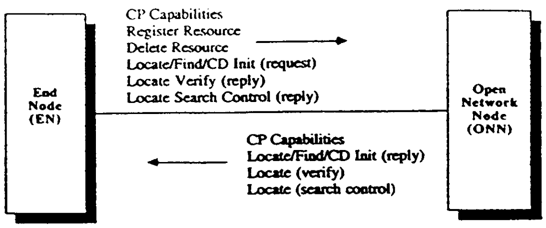

- the ONN presents an APPN NN appearance to the EN
- EN-to-ONN connection is LU 6.2 over NT 2.1
- all elements shown are standard GDS Variables
- formats and protocols are documented in NT 2.1 Node Reference (SC30-3422-2)

APPI Distributed Directory Service

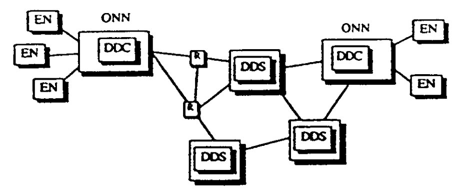

- DDC = Distributed Directory Client
- DDS = Distributed Directory Server
- each ONN contains a single DDC and may optionally contain a DDS
- DDCs maintain a cache of resources (for "locality of reference" optimizations)
- Σ(DDS) = distributed database

Building the APPI Distributed Database

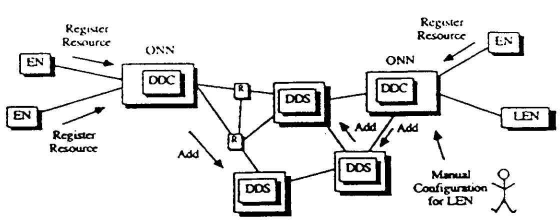

- ENs issue Register Resource GDS Variables
- LENs require static resource configuration at the ONN
- both of the above entities are intercepted by the DDC
- SNA resources are added to the designated DDS
  - an APPI DDS can accommodate a backup DDS

Referencing the APPI Distributed Database

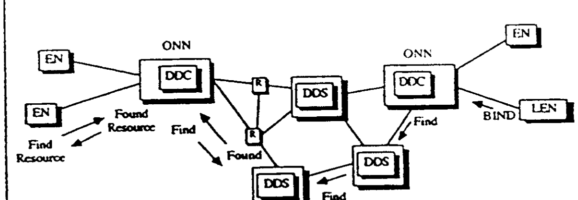

- ENs issue Find Resource GDS Variables
  - receive a Found Resource GDS Variable Reply
- LENs merely issue a BIND Request Unit to the ONN
  - Secondary LU Name must be found by ONN
  - ONN becomes a search proxy for LENs
- each DDS know about up to "n" other DDSs
- results of successful Found are cached at the DDC
  - to exploit "locality of reference"

APPI Data Transport Services

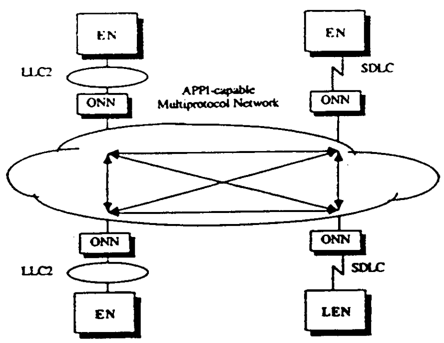

- SNA visibility is at the LU-level
  - LFSID used for making routing decisions
- any LU can get to any other LU in the APPI network
- LU transport established automatically by the participating ONNs

APPI Early Prototype

(Interop Spring '93)

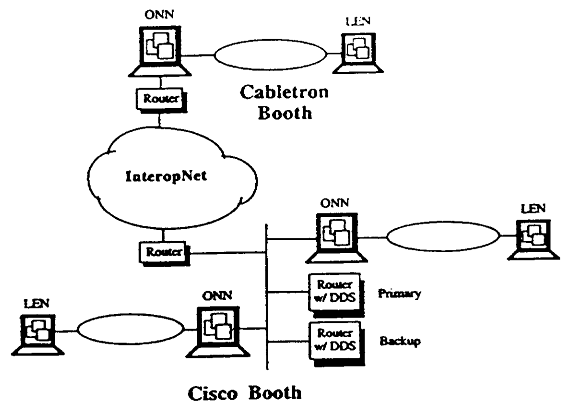

- ONNs are RS/6000 running AIX 3.2
  - SNA Subsystem w/ LU 6.2 and NT 2.1
  - DDC
- LENs are PS/2 running OS/2 ES 1.0
- DDS implemented on cisco router

LU-LU Session Establishment

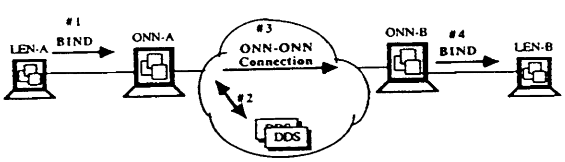

- #1: BIND received at ONN-A from LEN-A with an unknown SLU Name
- #2: LU Name query sent by ONN-A to DDS
- #3: ONN-to-ONN Connection established (socket)
- #4: BIND sent from ONN-B to LEN-B

BIND Data Path

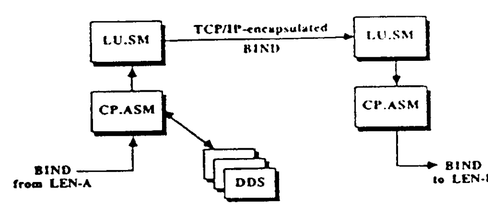

- LEN implementation of NT 2.1 used as the base
- originating Address Space Manager (ASM) enhanced to query DDC/DDS for unknown SLU Names
- originating Session Manager (SM) enhanced to equate local LFSID to socket handle
- destination SM enhanced to equate remote LFSID to socket handle
- destination ASM oblivious to socket connection
- from the APPN perspective, TCP/IP network is effectively a 1-hop ISR network

LU-LU Data Path

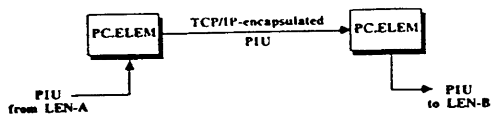

- originating Path Control Element (PC.ELEM) translates local LFSID to socket handle
- destination PC.ELEM translates socket handle to remote LFSID
- socket handle = (IP address of ONN,interface descriptor,unique no.)
- socket handle is effectively a short FQPCID
- minimal code path through PC.ELEM

APPN-over-TCP/IP vs. APPI

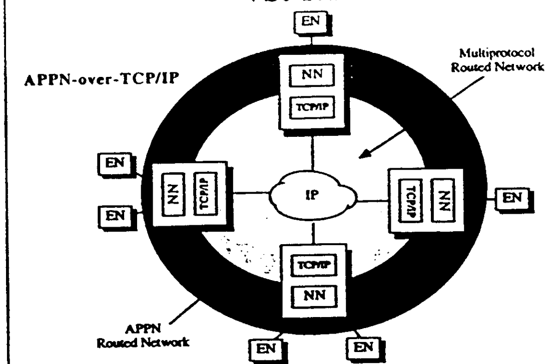

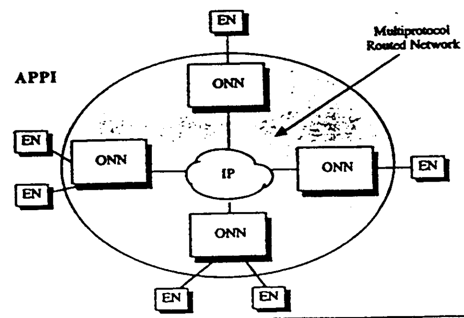

APPN/DLS vs. APPI

Example: n parallel sessions between 2 ENs

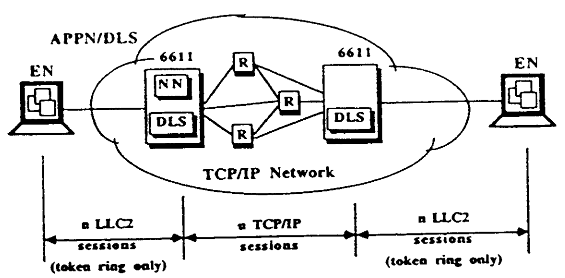

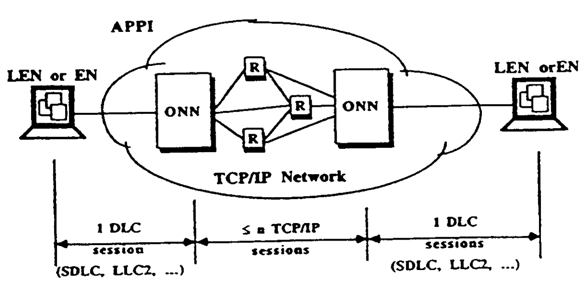

APPI Technical Benefits

- end systems can remain SNA-based
- supports both LEN and EN
- one DLC connection per SNA end system
- APPI direction established by APPI Forum
- reduces the number of protocols in the backbone
  - fewer network protocols
  - fewer routing protocols
  - SNA nodes kept on the network boundary
- media independence in the backbone
- better performance
- SNA sessions unaffected by intermediate link failure
- adaptive routing
- SNA implementation in ONN is based upon LEN
  - much smaller than APPN Network Node
  - greater availability of LEN offerings
- leverages existing transport technology
- open network management

SNA - TCP/IP Integration Issues

- there are a set of indigenous issues that ALL solutions face
- TCP/IP networks have different properties than SNA
  - connectionless vs. connection-oriented
  - reliable transport vs. reliable data link
  - reactive flow control vs. preventative flow control
- TCP/IP networks need:
  - better class of service
  - better congestion control
  - more deterministic delivery properties

2.3.1 ATM MIB (atommib)

Charter

Chair(s):

Kaj Tesink, kaj@cc.bellcore.com

Mailing Lists:

General Discussion: atommib@thumper.bellcore.com
To Subscribe: atommib-request@thumper.bellcore.com
Archive:

Description of Working Group:

The AToM MIB Working Group is chartered to define sets of managed objects which will be useful in the management of ATM and SONET equipment, interfaces, networks, and/or services that conform to the relevant ATM and SONET specifications. The initial sets defined will be:

- An interface-specific MIB for ATM interfaces, which is aligned with the managed objects for interface layering being defined by the Interfaces MIB Working Group. The Working Group should consider the ATM Forum's ILMI MIB for its suitability in this respect, plus any extensions necessary to instrument the layers between the ATM layer and the IP layer (e.g., AAL5). The latter should take into account the work of the IP over ATM Working Group (e.g., the "Multi-Protocol over AAL5" specification).

- Managed objects for the monitoring and control of ATM PVCs and SVCs, both in ATM end-points and in ATM switches or networks. (Objects for ATM SVCs will be considered after completion of the work on ATM PVCs.)

- Managed objects that instrument devices with SONET interfaces that conform with the relevant SONET specifications. This work should closely align to other trunk MIBs (DS1/E1 MIB, DS3/E3 MIB). The Working Group should consider the existing Internet-Draft SONET MIB for its suitability in this respect.

Goals and Milestones:

Oct 1993 Post an Internet-Draft of the ATM and SONET MIB.

Dec 1993 Submit the ATM and SONET MIB to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Kaj Tesink/Bellcore

Minutes of the Managing ATM with SNMP BOF (ATMMIB)

The goal of the ATMMIB BOF was to determine if there was sufficient interest to create a working group to develop managed objects for ATM. *Editor's Note (md): Following the IETF meeting this Group became a working group and took the name AToM MIB Working Group (ATOMMIB).*

Agenda

- Overview, Proposed Charter.
- Proposals.
  - ATM MIB Strawperson Proposal (Masuma Ahmed)
  - ATM ILMI (Keith McCloghrie)
  - SONET MIB (Kaj Tesink)

Overview - Kaj Tesink

Following unanimous agreement that a working group is necessary to develop managed objects related to ATM, Kaj identified the following "terms of reference" as the basis for the proposed Charter of the ATM MIB work effort:

- Priority on ATM PVCs.

- Keep the MIB small.

- Priority is creation of an SNMPv1 managed objects. Support of SNMPv2 may be a future effort.

- Difference between interface and end-to-end view of NM information for ATM.

Kaj discussed the relationship of the proposed ATM MIB with physical layer MIBs. He suggested that the Working Group include work on the SONET MIB which was posted by Tracy Cox. The Trunk MIB Working Group which had developed MIBs for DS1/E1 and DS3/E3 no longer exists. Ayal Opher cautioned against assuming that the SONET MIB is "done". Since there are more implementations of SONET equipment, folks may pay more attention and have more comments on the proposed SONET MIB module.

There was discussion on the different perspectives on ATM-related network management information:

- Local DTE view of ATM-based communication over the local interface.
- Network view of the local interface.

- End-to-end view of the communication path.
- ATM switch management - both private and public.
- ATM-based service information - both private and public.

These different perspectives drive identification of different sets of managed objects.

Based on this discussion, a Charter was agreed upon that directs the proposed Working Group to undertake work on ATM and SONET management.

Proposals

- ATMMIB Proposal - Masuma Ahmed

  Masuma gave an overview of a Bellcore strawperson proposal for an ATM MIB. See the copy of the viewgraphs for a detailed list of proposed management information. The strawperson MIB is based on the ILMI, and includes configuration information about the ATM interface, ATM layer statistics, but (unlike the ILMI) also VPL/VCL level performance information. Virtual Channel Links are the interfaces between two adjacent ATM devices. A Virtual Channel Connection is the set of VCLs used end-point to end-point. A VCL is identified by a Virtual Channel Identifier (VCI). Similarly, a Virtual Path Connection consists of VPLs, identified by VPIs.

  Issues identified during the discussion included:

  - The relation to the existing ILMI MIB (see Keith's talk).
  - The need for end-to-end PVC configuration information.
  - The modelling of ATM in relation to MIB-II interfaces table.
  - The need to keep the MIB small.
  - The need to coordinate modeling of PVCs with FR MIB effort.
  - The need for AAL performance statistics.
  - The need for VPC/VCC level performance statistics.

- ATM Forum ILMI - Keith McCloghrie

  Keith gave a presentation on the ATM Forum's Interim Local Management Interface ILMI, stressing that even though the ILMI uses SNMP packet formats, the purpose of the ILMI is not network management, but rather "interface management". However, since the management information needed for interface management is a subset of the information needed for network management, the ILMI's MIB should be appropriate for use by network management as well as for interface management.

  The ATM Forum's ILMI MIB was designed to be symmetric – that is, to have the same MIB on each side of the interface. It currently contains read-only objects. The ILMI MIB contains information on the physical port (e.g.,operStatus), ATM layer, ATM cell, and VPC/VCCs. There are no VCL or VPL level statistics in the ILMI.

Keith recommended that the proposed IETF Working Group use the ILMI MIB as a base for the development of IETF ATM MIB(s). The management of the local interface should use a common MIB module/table/group.

- SONET MIB - Kaj Tesink

  Kaj gave a short presentation of the SONET MIB. The SONET MIB has been posted as an Internet-Draft for a long time. It is consistent with other trunk MIBs. A minor change may be needed on table indexing. Interested people are encouraged to review this MIB.

  Internet-Draft: "Definitions of Managed Objects for the SONET Interface Type", Jan.1993, Tracy A. Cox, Kaj Tesink, <draft-ietf-cox-sonetmib-01.txt>.

Conclusions of the BOF

- There is interest in forming a Working Group to develop standard managed objects for ATM and SONET.

- Deirdre Kostick suggested calling it the AToMMIB Working Group, with the "o" signifying SONET, and the "AToM" as a reminder to keep MIBs small(!).

- It is likely that there will be multiple MIB modules for ATM developed to support the different management perspectives identified during the BOF.

- The MIB module(s) for management for the local interface should build on the ATM Forum ILMI MIB.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
David Arneson	arneson@ctron.com
David Battle	battle@cs.utk.edu
Nutan Behki	Nutan_Behki@qmail.newbridge.com
Caralyn Brown	cbrown@wellfleet.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
John Chang	changj@ralvm6.vnet.ibm.com
Anthony Chow	chow_a@wwtc.timeplex.com
George Clapp	clapp@ameris.center.il.ameritech.com
Michael Collins	collinsms@es.net
Tracy Cox	tacox@sabre.bellcore.com
James Davin	davin@thumper.bellcore.com
M.J. Dixon	mjd@att.com

Kurt Dobbins	kurtdob@ctron.com
David Engel	david@ods.com
Kenneth Giusti	kgiusti.chipcom.com
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
Mark Kepke	mak@fc.hp.com
Moshe Kochinski	moshek@FibHaifa.com
Deirdre Kostick	dck2@sabre.bellcore.com
David Lovering	lovering@bldrdoc.gov
Andrew Malis	malis_a@timeplex.com
Jun Matsukata	jm@eng.isas.ac.jp
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
George Mouradian	gvm@arch3.att.com
Tom Nisbet	nisbet@tt.com
Bill Norton	wbn@merit.edu
Zbigniew Opalka	zopalka@agile.com
Ayal Opher	aopher@synoptics.com
Joe Pagan	jrp@afterlife.ncsc.mil
Michael Patton	map@bbn.com
Drew Perkins	ddp@fore.com
Venkat Rangan	venkat@.metrix.com
Kenneth Rodemann	krr@qsun.att.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jon Saperia	saperia@lkg.dec.com
Andrew Schmidt	ags@uius.edu
Kanan Shah	kshag@cmf.nrl.navy.mil
Andrew Smith	asmith@synoptics.com
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Sally Tarquinio	sallyt@gateway.mitre.org
Kaj Tesink	kaj@cc.bellcore.com
Stephen Tsun	snt@3com.com
Mario Vecchi	mpv@thumper.bellcore.com
James Watt	james@newbridge.com
Peter Wilson	peter_wilson@3com.com
Yow-Wei Yao	yao@chang.austin.ibm.com
Kiho Yum	kxy@nsd.3com.com

SONET MIB

Kaj Tesink

(908) 758-5254

kaj@cc.bellcore.com

Bellcore

1

SONET MIB

- Can be used to manage SONET circuits
- Managed in the same way as DS1, E1, DS3, E3 (RFC1406, RFC1407)
- Applies DS1/DS3 MIB approach.
- Uses ANSI terminology (T1M1.3/92-005R1)
- Based on SONET specifications T1.105-1988/1990, T1.106-1988
- Internet-Draft draft-ietf-cox-sonetmib-01.txt
- One minor adjustment needed

2

SONET Layers

Manages all SONET layers separately

VT Layer
Path Layer
Line Layer
Section Layer
Photonic Layer

3

SONET MIB Specifics

Feature	Photonic	Section	Line	Path	VT
Config	Coding LineType Circuit ID			Path Width	VT Width
Status		LOS LOF	AIS RDI	AIS RDI LOP	AIS RDI LOP RFI
Statistics (15 minute Intervals)		ESs, SESs, CVs, UASs			

4

ATM MIB BOF

Kaj Tesink

(908) 758-5254

kaj@cc.bellcore.com

Bellcore

1

ATM MIB BOF AGENDA

- Terms of reference / Charter
- Presentation of a strawman proposal (M. Ahmed)
- Background ATM Forum ILMI (K. McCloghrie)
- SONET MIB, and Relationship with other MIBs
- Discussion

2

Items on terms of reference

- Priority on ATM PVCs
- Keep it small
- Priority for SNMPv1
- Difference of "Local" interface aspects and "End-to-End" aspects
- Other MIBs - in particular SONET MIB
- Relationship with other groups
- Separate WG

3

Local and End-to-End aspects

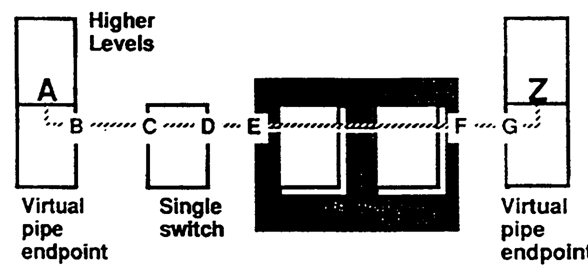

4

Bellcore Strawman Proposal of ATM MIB

Masuma Ahmed

Bellcore

Scope of Current Proposal

- Develop a standard ATM MIB that can be used to manage ATM interfaces and can be applied irrespective of public or private environments.

ATM Virtual Channel

VCC - Virtual Channel Connection

VCL - Virtual Channel Link

End-point - Equipment such as a router or workstation terminating the VCC

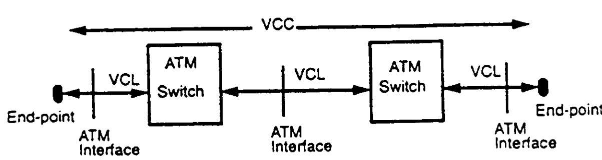

ATM Virtual Path

VPC - Virtual Path Connection

VPL - Virtual Path Link

End-point - Equipment such as a router terminating the VPC

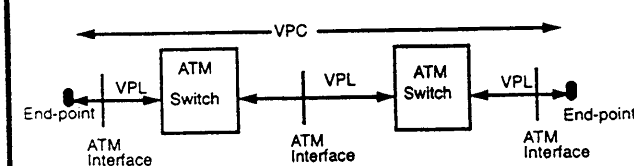

ATM VPL and VCL

VPL - Virtual Path Link

VCL - Virtual Channel Link

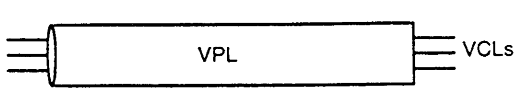

ATM Management Information

- Configuration information about the ATM interface and VCL/VPL
- Performance information about the ATM interface and ATM VCL/VPL at different protocol layers:
  - ATM physical layer and PLCP layer
  - ATM cell layer

Configuration information-ATM Interface

- ATM address
- Physical transmission type
- Physical media type
- Maximum VP and VC links
- Configured VP and VC links
- Maximum configured VPI and VCI fields
- Available ingress and egress bandwidth

Configuration Information - VCLs and VPLs

- VPI/VCI values
- Operational status
- Administrative status
- Ingress and egress QOS class
- Ingress and egress traffic parameters

ATM Statistics - Cell Layer

- ATM cells transmitted/received across an ATM interface
- ATM cell transmission errors
- ATM cell protocol and VPI/VCI violation errors

ATM Statistics - VCL or VPL Level

- ATM cells transmitted/received across an ATM VCL or VPL
- ATM cell peak rate violation errors across a VCL or VPL
- ATM statistics are kept either at VCL or VPL level and not both.

Open Issues

- End-to-end PVC configuration information
- Multi-point configuration information
- ATM level loopback and performance monitoring
- AAL performance statistics
- Indexing of information flow
- Pointer information about other MIBs, e.g.,
  - Physical layer MIBs (e.g., SONET, DS3)

2.3.2 Bridge MIB (bridge)

Charter

Chair(s):

Fred Baker, fbaker@acc.com

Mailing Lists:

General Discussion: bridge-mib@pa.dec.com
To Subscribe: bridge-mib-request@pa.dec.com
Archive:

Description of Working Group:

The Bridge MIB Working Group is chartered to define a set of managed objects that instrument devices that conform to the IEEE 802.1 standard for MAC-layer bridges.

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.

The MIB object definitions produced will be for use by SNMP and will be consistent with other SNMP objects, standards, and conventions.

Goals and Milestones:

Done	Publish initial proposal.
Done	Submit an Internet-Draft.
Done	Submit draft for RFC publication.
Done	Publish a draft revision to RFC 1286 that reflects implementation experience and the result of alignments with IEEE work as an Internet-Draft.
Done	Publish a draft SNMP MIB that instruments functions specific to source routed bridges as an Internet-Draft.
Done	Submit a draft MIB for source routing bridge functions to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Definitions of Managed Objects for Bridges", 10/15/1992, E. Decker, P. Langille, A. Rijsinghani <draft-ietf-bridge-objects-02.txt>

"Definitions of Managed Objects for Source Routing Bridges", 05/05/1993, E. Decker, K. McCloghrie, P. Langille <draft-ietf-bridge-sr-objects-01.txt>

Request For Comments:

RFC 1286 "Definitions of Managed Objects for Bridges"

2.3.3 Character MIB (charmib)

Charter

Chair(s):

Bob Stewart, `rlstewart@eng.xyplex.com`

Mailing Lists:

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

Description of Working Group:

The Character MIB Working Group is chartered to prepare a recommendation to the IESG evaluating RFCs 1316-1318 (the Character MIBs) with respect to the standards track.

The recommendation will document implementation, interoperability, and deployment experience. If these experiences suggest that changes should be made to the documents, new drafts may be prepared. The recommendation will report one of four outcomes for each RFC:

- That the RFC should be advanced from Proposed to Draft status, without changes (if no problems are found);

- That a draft prepared by the Working Group should replace the RFC, and be designated a Draft Standard (if only minor changes are made);

- That a draft prepared by the Working Group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or,

- That the RFC should be designated as historic (if this technology is problematic).

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discussion and final approval of Charter; discussion on models and terminology. Make writing assignments.
Done	First draft document, discussion, additional drafts, special meeting?
Done	Review latest draft and if OK, give to IESG for publication as RFC.
Apr 1993	Reactivation of Working Group to prepare the Character MIBs for Draft Standard.

Jun 1993 Post an Internet-Draft with the results of the survey of implementation and operational experiences with the Character MIBs. Post revised MIB documents if necessary.

Aug 1993 Submit the Character MIBs to the IESG for consideration as Draft Standards.

Request For Comments:

RFC 1316 "Definitions of Managed Objects for Character Stream Devices"

RFC 1317 "Definitions of Managed Objects for RS-232-like Hardware Devices"

RFC 1318 "Definitions of Managed Objects for Parallel-printer-like Hardware Devices"

2.3.4 Chassis MIB (chassis)

Charter

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

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 implementations 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 consistent 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).
Done	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.

Internet-Drafts:

"Definitions of Managed Objects for a Chassis Containing Multiple Logical Network Devices", 01/13/1993, K McCloghrie, D Arneson, M Kaycee <draft-ietf-chassis-mib-01.txt>

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Chassis MIB Working Group (CHASSIS)

The Group reviewed and approved the Agenda. The current draft document is dated January 27, 1993.

Revised Model for the Chassis MIB

Dave Arneson presented a revised model for the Chassis MIB. It consists of three kinds of things within a chassis - modules, resources, and entities. Each resource is used to build up a single module; independently each resource is used to build up a single entity. The first relationship is used to describe the physical location of a resource. The second relationship is used to describe the logical functional unit for management access.

There are three primary tables in the model:

1. A module table indexed by location type and slot number.
2. A resource table indexed by loc-type, slot-#, and resource-#.
3. An entity table indexed by entity number.

There was some discussion that the resource table should be two logical tables: one indexed by module information, and the other indexed by entity information. There was general consensus that the model is an improvement, and reasonably general.

There was a discussion of whether we are modeling chassis configuration or network topology. The general consensus seemed to be that network topology is a general problem and is not part of our requirement.

A remaining question was what access information is necessary for v1 and v2.

Implementation Experience

Rick Royston discussed implementation experience. He was concerned that indexing was difficult, but thought that the new model would help. He was also concerned about keeping the management application informed when the chassis configuration changed. Finally, there was concern about dynamic rules for what card is allowed to go in which slots.

Actions

Group members	Apply the model to their own devices and report back.
Pete Wilson	Provide a set of examples.

Dave Arneson	Revise draft based on model.
Jeff Case, Marshall Rose	Investigate entity access information.

Attendees

David Arneson	arneson@ctron.com
Jim Barnes	barnes@xylogics.com
David Battle	battle@cs.utk.edu
Andy Bierman	abierman@synoptics.com
Jeff Case	case@cs.utk.edu
Anthony Chow	chow_a@wwtc.timeplex.com
Shane Dawalt	sdawalt@desire.wright.edu
Manuel Diaz	diaz@davidsys.com
Sandra Durham	sdurham@synoptics.com
David Engel	david@ods.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
John Hopprich	hopprich@davidsys.com
Mark Kepke	mak@fc.hp.com
Kenneth Key	key@cs.utk.edu
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Tom Nisbet	nisbet@tt.com
Bill Norton	wbn@merit.edu
Venkat Rangan	venkat@.metrix.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Timon Sloane	timon@timon.com
Ira Steckler	isteckle@chipcom.com
Bob Stewart	rlstewart@eng.xyplex.com
Geoffrey Thompson	thompson@synoptics.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Peter Wilson	peter_wilson@3com.com
Kiho Yum	kxy@nsd.3com.com

3Com Physical Chassis

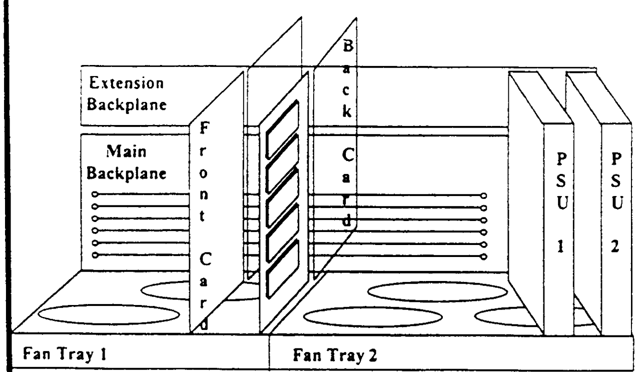

3Com Proprietary and Confidential

3Com Logical Chassis

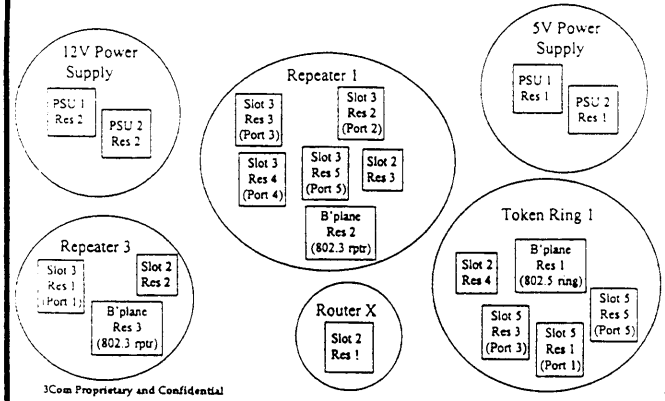

3Com Proprietary and Confidential

3Com Mappings

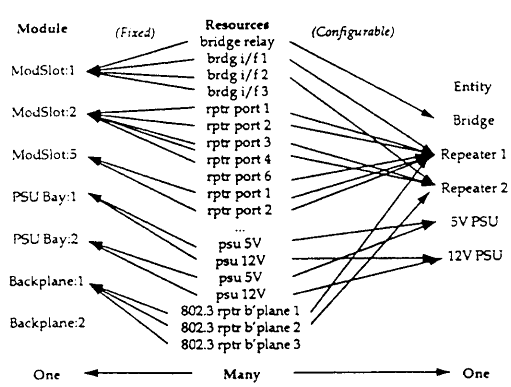

3Com Proprietary and Confidential

3Com Definitions (1)

- Chassis
  A generic 'container'. Any size or shape.A chassis can be viewed either from a physical or a logical perspective. Physically a chassis contains a number of modules. Logically it contains a collection of functional 'entities'.
- Module
  A module is a physical thing contained within a chassis. A particular chassis can contain any number of different types of module and any number of distinct instances of the same type of module. A module is distinguished by its type and its position. Only one module can occupy a particular position within the chassis. A module is itself a container containing resources.

3Com Proprietary and Confidential

3Com Definitions 2

- Resource

  Each physical module comprises a number of resources. The resources can be of any kind, for example a repeater port, bridge relay etc.
- Entity

  An entity is a functional unit wihin the chassis. An 'entity' is anologous to a seperate, stand-alone device. Examples of chassis entities are: Repeaters, Bridges, Routers, Roken Ring Concentrators.

3Com Proprietary and Confidential

3Com Physical View (1)

Module Table

A module is distinguished by its physical location in the chassis.

LocType	Position	Description	SW Ver	Oper Stat
PSU Bay	1	Std PSU		OK
PSU Bay	2	ExtPSU		Failed
ModSlot	1	12xUTP Rptr		OK
ModSlot	2	12xUTP Rptr		OK
ModSlot	3	12xUTP Rptr		OK
ModSlot	4	3x3 Bridge		Self-Test
ModSlot	5	12xUTP Rptr		Self-Test
...	...	...		...
B'plane	1	Std B'plane		OK
B'plane	2	Ext B'plane		OK

3Com Proprietary and Confidential

Physical View (2)

Resource Table

Resources are distinguished by the module on which they reside. The index on this table establishes the resource to module mapping

LocType	Position	Resource	Type	Assignment
...	...	...	...	...
ModSlot	1	1	.3 rptr pt	rptr 1
ModSlot	1	2	.3 rptr pt	rptr 1
ModSlot	1	3	.3 rptr pt	rptr 2
ModSlot	1	4	.3 rptr pt	rptr 2
...	...	...	...	...

The assignment column establishes the mapping from resource to entity.

Logical View

Entity Table

Id	Type	Description	Version	Access
1	repeater8023	802.3 Repeater 1		
2	repeater8023	802.3 Repeater 2		
3	tokenRing8025	802.5 Concentrator		
4	bridge	802.3/802.5 bridge		
5	powerSupply	5V PSU		
6	powerSupply	12V PSU		

Rows in this table may be statically created by the agent or may be dynamically created through the MIB depending on the implementation of the chassis. If the agent allows the dymamic creation of new entities then this must be achieved through the creation of rows in this table followed by the assignment of resources to the entity.

Note that a particular box may restrict the types of entities that may be created. It is generally straightforward to create a new repeater, not so a new router!

Benefits of this representation

- Simple Relationship between modules, resource and entities.
  - Simple implementaition for read-only form
  - Simple (MIB) implementation for writable form
- No seperate triply indexed mapping table
- All relationships explicit and understandable
  - Reassigning a resource moves that resource to another entity. No need to create a new entry in the module/entity/segment table which has the side-effect of deleting another relationship
- Allows a single module to connect multiple resources to an entity
  - Not possible in existing MIB without changes because the seperate relationship table only allows one relationship between a module/segment/entity triplet.
  - Required to support per-port switching of token ring and Ethernet.

Existing 'Relationship' model.

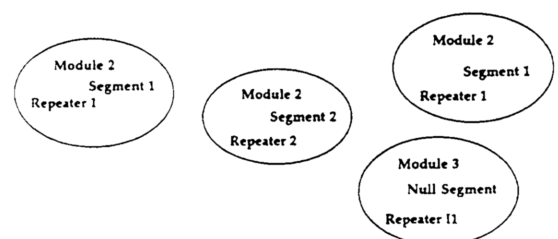

- Extra concept (4 in total).
- Finest physical granularity is a module.
- Only deals with one type of physical entity.
- Makes 'Segment' and exception.
- Creating a relationship tends to remove another by side-effect.

2.3.5 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):

Jonathan Saperia, saperia@tay.dec.com

Mailing Lists:

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

Description of Working Group:

The DECnet Phase IV MIB Working Group is chartered to prepare a recommendation to the IESG evaluating the standards track status of RFC1289 (the DECnet Phase IV MIB).

The recommendation will document implementation, interoperability, and deployment experience. If this experience suggests that changes should be made to the document, a new draft may be prepared. The recommendation will report one of four outcomes:

- That RFC 1289 should be advanced from proposed to draft status, without changes (if no problems are found).

- That a draft prepared by the Working Group, should replace RFC 1289, and be designated a draft standard (if only minor changes are made).

- That a draft prepared by the Working Group, should replace RFC 1289, and be designated a Proposed Standard (if major changes or feature enhancements are made).

- That RFC1289 should be designated as historic (if this technology is problematic).

Goals and Milestones:

May 1993 Re-activate the Working Group to advance the Decnet Phase IV MIB to Draft Standard.

Jun 1993 Post an Internet-Draft of the results from the survey of implemention and operational experience. Post a revised version of the MIB if necessary.

Oct 1993 Submit the DECnet Phase IV MIB to the IESG for consideration as a Draft Standard.

Request For Comments:

RFC 1289 "DECnet Phase IV MIB Extensions"

2.3.6 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.
Mar 1992	Submit the FDDI MIB to the IESG for consideration as a Proposed or Draft Standard depending on the magnitude of changes to RFC1285.
Done	Hold a meeting at the November IETF Plenary.
Dec 1992	Post an Internet-Draft aligned with current the current ANSI document factoring in implementation experience with RFC1285.

Internet-Drafts:

"FDDI Management Information Base", 03/08/1993, J. Case, A. Rijsinghani
<draft-ietf-fddimib-objects-02.txt>

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)

Twenty members, plus the Chair were present. The Working Group endorsed the selection of a new Editor, Anil Rijsinghami, who will assist Jeff Case. After introductions the proposed Agenda was approved and a Secretary, Jack Brown, was selected. There were no new issues on the mailing list.

7.3 SMT Prefix

Jim Reeves brought up the problem of the 7.3 SMT prefix.

- No one in the foreign community would understand.
- A new prefix needed to be selected that would not conflict with the present naming of RFC1285 and ANSI X3T9.5 .

Action Item: The Editor is to replace SMT 7.3 with FDDI MIB.

Ron Macken had three editorial issues on the draft. *Editor's Note (md): A detailed listing of the editorial issues and other proposed modifications is available via ftp under fddimib-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Action Item: By Group consensus the Editor is to make the changes to the document as agreed by the Group and will place the new draft on the Internet for comment after which it will be submitted forward. It was also agreed that Version 2 issues be separate.

Traps

Five conditions are implementable. Four events suspect are implementable. EV overflow - might be important, but not enough information is available to implement.

There was a discussion on whether there should be separate trap tables. Marshall Rose thought the tables should be second MIB in the original document. There was no final decision but a strawman indicated that five members definitely want a trap document. No one indicated that they definitely did not want a trap document.

Action Item: Ron Macken is to produce a strawman document with the following instructions:

- A simpler, more concise strawman document with traps stands a better chance.
- Leave traps out.
- Leave the MIB the way it is now so the MIB will work with SNMP Version 1 besides SNMP Version 2.

The next Working Group meeting will be planned for the Amsterdam IETF Meeting to handle the trap document and V2 compliance issues unless everything gets pushed back on the mailing list. Friday would be the best day. NOTE: Ron Macken has a working implementation of the MIB with the 6.2 draft.

Attendees

Michael Anello	mike@xlnt.com
David Battle	battle@cs.utk.edu
Jeffrey Berk	berk@ctron.com
John Boatright	bryan_boatright@ksc.nasa.gov
Caralyn Brown	cbrown@wellfleet.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Jeff Case	case@cs.utk.edu
Anthony Chow	chow_a@wwtc.timeplex.com
David Engel	david@ods.com
Wayne Foco	foco@ralvm6.vnet.ibm.com
Paul Franchois	paulf@bldrdoc.gov
Kenneth Giusti	kgiusti.chipcom.com
John Hopprich	hopprich@davidsys.com
Ronald Jacoby	rj@sgi.com
Merike Kaeo	merike@alw.nih.gov
Kenneth Key	key@cs.utk.edu
Evan McGinnis	bem@3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Steve Suzuki	suzu@fet.com

2.3.7 Frame Relay Service MIB (frnetmib)

Charter

Chair(s):
James Watt, james@newbridge.com

Mailing Lists:
General Discussion: frftc@nsco.network.com
To Subscribe: frftc-request@nsco.network.com
Archive:

Description of Working Group:

The Frame Relay Service MIB Working Group is chartered to define an initial set of managed objects which will be useful for customer network management of a provider's Frame Relay Service. The Working Group will consider existing definitions, including the Frame Relay Forum's work in this area. The objects defined by the Working Group will be consistent with the SNMP framework.

The Working Group will coordinate with both the Frame Relay Forum and the ATM MIB Working Group.

Goals and Milestones:

May 1993 Post the initial Internet-Draft for discussion.

Dec 1993 Submit the Frame Relay Service MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Definitions of Managed Objects for Frame Relay Service", 05/13/1993, T. Cox
<draft-ietf-frnetmib-fr-00.txt>

CURRENT MEETING REPORT

Reported by Deirdre Kostick/Bellcore

Minutes of the Frame Relay Service MIB BOF (FRNETMIB)

The purpose of the Frame Relay Service MIB BOF was to determine if there was interest in writing a standard Frame Relay Network MIB and to determine if the MIB should be developed in the IETF.

Tracy Cox presented the purpose of the BOF, the proposed scope of the MIB, and discussed the relation to the Frame Relay Forum Technical Committee activities. The proposed MIB has the following scope/features:

- The MIB will be an SNMPv1 MIB.
- It will contain read-only objects.
- It is intended for use by end-customers (versus service providers) to manage their portion of a Frame Relay network.
- It is intended to support fault detection, performance monitoring, and configuration for Frame Relay interfaces.
- It is NOT intended to be a switch MIB, and will NOT include managed objects for switching elements and related internal aspects of the network supporting Frame Relay.

Tracy discussed the relation with the existing Frame Relay DTE MIB (RFC1315). Based on discussion with RFC1315 authors and others, Tracy determined that the Frame Relay DTE MIB was not sufficient to manage the Frame Relay interface from the network perspective. For example, a Frame Relay Network MIB would need bi-directional information on Frame Relay parameters (CIR, Be, Bc), and an end-to-end view of the network, neither of which are supported in RFC1315.

The Frame Relay Network MIB would not include managed objects for the physical layer. Existing physical layer MIBs (e.g., DS1 and DS3) would be used.

There was agreement with the scope of the MIB. The BOF attendees also agreed that a) there was interest in writing a standard MIB, and b) that the IETF was the appropriate body for the development of standard MIBs.

There was discussion on the relationship of the efforts in the Frame Relay Forum and the proposed ATOMMIB Working Group. These issues were discussed at length during George Mouradian's presentation.

George presented ideas on service management – Architecture Principles for Service MIBs. George indicated that standards for service management are necessary, and that Frame Relay and ATM services are good candidates for standardized MIBs. Both the user and the vendor communities benefit. George raised questions and issues related to the proliferation of MIBs and of different groups working on the same management issues.

George indicated that he felt it was premature to start a separate effort in the IETF and that it was necessary to give the Frame Relay Forum more time to complete their efforts. There was lively discussion on this topic. Caralyn Brown indicated that the intent was not to discontinue the work in the Frame Relay Forum, rather it was to follow the process used for development of RFC1294. Andy Malis indicated that for RFC1294 the work in the IETF was brought into the Frame Relay Forum, and that there were ongoing efforts to keep the Frame Relay Forum informed of the IETF work and to gain their input and consensus. Both Caralyn Brown and Andy Malis were involved in this coordination effort.

Ken Rodemann indicated that he agreed that it was unwise to assume that the IETF would "rubber-stamp" a MIB brought in from an outside group. However, he felt that it would be wiser to let the work continue in the Frame Relay Forum before bringing it into the IETF.

Doug Kay supported continuing the work in the Frame Relay Forum since there was expertise on Frame Relay; however, Doug also indicated that he felt that the IETF offered the network management and SNMP-related expertise that would be necessary to develop a quality MIB. Doug indicated that he would feel comfortable if it was clear that the Frame Relay Forum was responsible for defining the managed objects and that the IETF Group would be responsible for "MIBification".

There was discussion on the timing of the establishment of the proposed Working Group and the process for coordination with the Frame Relay Forum work. Deirdre Kostick suggested that the Forum continue to work on the proposed list of managed objects via email and at their June meeting, develop a consensus on the set of objects. This set of objects would be used by the proposed IETF Working Group to begin MIB definition at the July IETF meeting. James Watt suggested that an interim meeting should be held in conjunction with the June Frame Relay Forum meeting.

Results of the BOF

- There is interest in writing a Frame Relay Network Service MIB. The scope of the MIB is consistent with the scope identified during Tracy Cox's presentation. That is, the MIB will be an SNMPv1 MIB intended to support end-customer network management for their Frame Relay interfaces. It is not intended to be a switching system MIB.

- There is agreement that a working group should be created to develop the Frame Relay Network Service MIB.

- The first meeting (assuming approval of the Working Group by the yet-to-be-named Network Management Area Director) will be June 28-30 with the Frame Relay Forum Technical Committee in Chicago. Meeting logistics will be posted on the mailing list.
- The Frame Relay Network Service MIB will be based on the managed objects identified by the Frame Relay Forum Technical Committee. There will be ongoing coordination efforts between the two Groups.
- The proposed schedule for deliverables from the Working Group are indicated below in the proposed Charter.
- There will be coordination with the proposed ATOMMIB Working Group to insure that common elements are consistently modeled.

Proposed Charter of Working Group

- Tracy Cox will Chair the Group.
- Messages for Group discussion can be sent to frftc@nsco.network.com. Subscription requests for the discussion list should be sent to frftc-request@nsco.network.com.
- To write a standard Frame Relay Network Service MIB based on the set of managed objects identified by the Frame Relay Forum Technical Committee. Close coordination with the Frame Relay Forum is essential. Once chartered, this Working Group will also coordinate their efforts with the proposed ATOMMIB Working Group.
- The goals of the Group are to complete the first draft of an Internet-Draft by July of 1993 and to submit the final draft of the document for approval as an RFC by November of 1993.

Attendees

Masuma Ahmed	mxa@sabre.bellcore.com
Rich Bowen	rkb@ralvm11.vnet.ibm.com
Caralyn Brown	cbrown@wellfleet.com
Theodore Brunner	tob@thumper.bellcore.com
John Chang	changj@ralvm6.vnet.ibm.com
Anthony Chow	chow_a@wwtc.timeplex.com
Tracy Cox	tacox@sabre.bellcore.com
Manuel Diaz	diaz@davidsys.com
Ken Hayward	Ken.Hayward@bnr.ca
Don Hofacker	hofacker@dtedi.hq.aelc.af.mil
Doug Kay	doub.kay@sprintintl.sprint.com

Kenneth Key	key@cs.utk.edu
Zbigniew Kielczewski	zbig@eicon.qc.ca
Moshe Kochinski	moshek@FibHaifa.com
Deirdre Kostick	dck2@sabre.bellcore.com
Patrick Leung	patrickl@eicon.qc.ca
Andrew Malis	malis_a@timeplex.com
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
George Mouradian	gvm@arch3.att.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Louise Reingold	l.reingold@sprint.sprint.com
Bradley Rhoades	bdrhoades@mail.mmmg.com
Kenneth Rodemann	krr@qsun.att.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Kaj Tesink	kaj@cc.bellcore.com
James Watt	james@newbridge.com
Kiho Yum	kxy@nsd.3com.com

Service MIBs
for
Frame Relay and ATM

GEORGE MOURADIAN

- Important to have Standards for Service Management
  - Frame Relay and ATM are good candidates
  - Both user and vendors benefit.
- There are some Questions/Issues
  - Service Management vs. Device Management
  - Service MIB Proliferation (already!)
- Frame Relay Forum Status
  - List of 55+ possible attributes
  - What is a service MIB?
  - Critical mass of expertise
  - Track record of success
  - Collaboration with ATM Forum

Service MIBs
for
Frame Relay and ATM

- MIB Design Principles Apply
  - Start Small and Simple,
  - Minimize impact on managed resources,
  - Emphasize Usefulness,
  - Based on Experience and Use.
- Service MIB Design
  - Consistent with DTE MIBs
  - Model the service, not the devices
  - Don't replicate carrier functions
  - Focus on current practice, not emerging standards
  - Allow for extensions, enterprise MIBs
  - Expect a phased approach
  - Support Frame Relay and ATM interworking.

Frame Relay Network MIB BOF

Tracy A. Cox

(908) 758-2107

tacox@mail.bellcore.com

Frame Relay Network MIB BOF

Agenda

- Purpose of BOF
- Purpose of MIB
- What is Customer Network Management?
- Background and Assumptions
- Frame Relay CNM Information
  - Assumptions
  - MIB Architecture
  - ATM MIB coordination
  - Standards Development
- Architecture Principles for Service MIBs -- George Mouradian

Frame Relay Network MIB BOF

Purpose of BOF

- To see if there is interest in writing a "standard" Frame Relay Network MIB.

Frame Relay Network MIB

Purpose of MIB

- Access to a Frame Relay Network MIB via SNMPv1.
- Provides read-only information that allows end-customers to obtain performance monitoring, fault detection, and configuration information from a Frame Relay Network -- MIB is used for CNM.
- By using this and other related MIBs, an end-customers' NMSs can monitor their PVCs and UNI/NNI access channels.
- Internal aspects of the network (e.g., switching elements, line cards, and network routing tables) are outside the scope of this MIB.

What is Frame Relay PVC CNM Service?

- To the Frame Relay PVC Customers:
  - All activities that end-customers perform to manage their Frame Relay UNIs, PVCs, and Frame Relay service features; should be consistent with managing their whole data communications network
- To the Service/Network Provider:
  - Providing customers with the capabilities to manage directly their portion of the Frame Relay network provided by the public service provider

What is Frame Relay PVC CNM Service?

Background and Assumptions

- Frame Relay DTE MIB (RFC1315) is not sufficient to manage Frame Relay Network interfaces
  - DTE MIB is for end-systems (e.g., routers and bridges)
  - Need network view of the interface to manage PVCs through a network
  - Need Forward/Backward CIR, Bc, and Be information
  - Need to identify the interface (e.g., UNI or NNI)

What is Frame Relay PVC CNM Service?

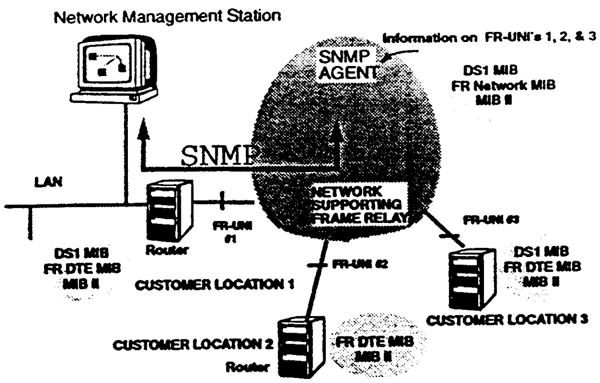

Frame Relay PVC CNM Addressing

Access to CNM service via FR-UNI

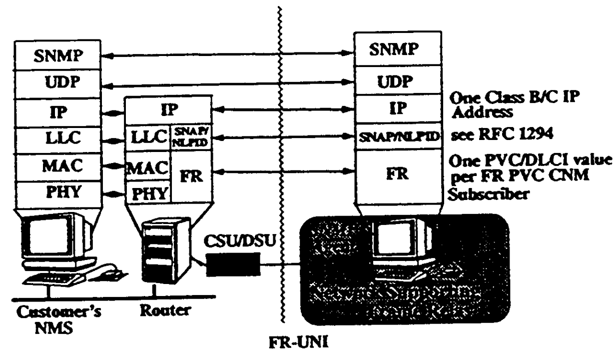

Frame Relay PVC CNM Information

Assumptions

- IETF has specified Frame Relay DTE MIB (RFC1315)
  - Not applicable to Frame Relay network interfaces (FR-UNI or FR-NNI)
- Proposing that a standardized Frame Relay Network MIB be written
- DS1, DS3, and SONET MIBs are defined in the IETF
  - There is an RS232-like MIB defined within the IETF
  - May need other physical layer MIBs

Frame Relay PVC CNM Information

MIB Architecture (example)

MIB II		
Frame Relay X.25 MIB for IW/Encap.	SIP Relay MIB	RFC1294 encapsulation no MIB needed
Frame Relay Network SNMP MIB for CNM		
Physical Layer MIB (e.g., DS1/E1 MIB)	ATM MIB	
	Physical Layer MIB	

Frame Relay PVC CNM Information

Relationship with ATM MIB BOF

- Both ATM and Frame Relay are based on virtual connections
- Frame Relay can be transported across an ATM network and can interoperate with ATM end-systems
- ATM and Frame Relay are distinct technologies and different layers of a protocol stack
- Coordination will benefit both efforts

Frame Relay PVC CNM Standards Developments

Frame Relay Forum Technical Committee

- Writing a draft of the Frame Relay Network SNMP MIB for CNM:
  - Read-only information (NOT to be used for Provisioning)
  - SNMPv1 MIB
  - Purpose is for Customer Network Management
  - CNM service capabilities are fault detection, performance monitoring and retrieving configuration information

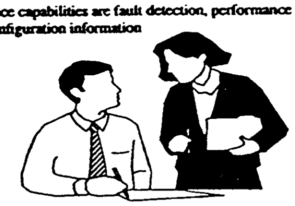

2.3.8 Host Resources MIB (hostmib)

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

Internet-Drafts:

"Host Resources MIB", 10/07/1992, Pete Grillo, Steven Waldbusser <draft-ietf-hostmib-resources-01.txt>

2.3.9 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: sweetwater.synoptics.com:~/pub/hubmib

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 (i.e., 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.

The Working Group is further chartered at its discretion to define an SNMP MIB for management of IEEE 802.3 Medium Access Units (MAUs). An 802.3 Medium Attachment Unit (MAU) attaches a repeater port or Ethernet-like interface to the local network medium. The scope of this work may include several types of MAU units: 10BASE5 (thick coax), 10BASE2 (thin coax), 10BASE-T (twisted pair), FOIRL and 10BASE-F (fiber optic). Managed objects defined as part of the MAU MIB task may, for example, represent such information as MAU type, link status, and jabbering indications.

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.
Done	Post the Media Access Unit MIB Definition as an Internet-Draft.
Apr 1993	Submit the Media Access Unit MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs)", 03/08/1993, D. McMaster, K. McCloghrie, S. Roberts <draft-ietf-hubmib-mau-02.txt>

"Definitions of Managed Objects for IEEE 802.3 Repeater Devices", 06/07/1993, D. McMaster, K. McCloghrie <draft-ietf-hubmib-objects-00.txt>

Request For Comments:

RFC 1368 "Definitions of Managed Objects for IEEE 802.3 Repeater Devices"

CURRENT MEETING REPORT

Reported by Donna McMaster/SynOptics

Minutes of IEEE 802.3 Hub MIB Working Group (HUBMIB)

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

IEEE Report

Geoff Thompson, Executive Secretary of IEEE 802.3 Repeater Management Task Force, reported on the Task Force's status. The IEEE's Repeater MIB was approved last September and published last November, and has been submitted to ISO where it is undergoing a 30-day Committee Draft (CD) ballot. A few editorial changes are being submitted as comments from the United States. The IEEE's MAU MIB has undergone two rounds of balloting and is expected to be approved and published by July 1993, and be submitted to ISO soon after. The organization of the specification has been changed to be protocol-independent with the GDMO-specification in a normative Annex. This allows, for example, the sizes of counters to be made protocol-specific.

MAU MIB

A message to the mailing-list had questioned the value of mauJabberState because that state was so short-lived. The Group agreed that this was not the case since the "jabber" state is not exited until reset, and thus decided to leave the document as-is.

The question of if and how to represent an interface/port/MAU used only to manage a repeater was discussed. Normally, these are internal to a device and thus often proprietary, and in fact such a MAU might effectively be null, in which case there would be no need to have MIB objects for it. Even if the MAU wasn't null, rpMauType could have an enterprise-specific Object Identifier value. It was agreed to add a sentence or two to the Overview section of the MAU MIB to explain this. The suggestion to add a diagram to the document was rejected, since it was thought the issues were too vendor-specific to be able to reach agreement on a diagram.

A suggestion to change the name of mauJabberStateChanges was accepted in order to better reflect the behavior of the object, since it only counts the times the "jabber" state is entered, not all state changes.

Repeater MIB

There was lengthy discussion on rptrAddrTrackLastSourceAddress. The MIB editors had made a suggestion to the mailing-list prior to the meeting to specify that a noSuch error/exception should be returned prior to the first frame being received on a port. Responses on the mailing-list had preferred other approaches. All the possible solutions discussed at the previous meeting were again listed and discussed at this meeting, with the addition of having the object be initialized with the well-known MAC address defined for use in FDDI. By a process of elimination, an agreement was reached on the solution of deprecating the object and defining a replacement which would have a zero-length value until the first frame was received on the port.

Other minor changes were agreed to, including:

- 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).

At this point, the scheduled time for the Working Group meeting expired. Some of the participants left to meet other scheduled commitments, while others continued to discuss items informally until 12:30 p.m. In addition, a second informal meeting was held later to continue discussion of open issues.

First Informal Meeting

Noone remaining in the meeting had much to say on the topic of having a "repeater index" in the MIB. A few implementors thought it might make it easier to manage multiple repeaters, but noone wanted to change the MIB.

The requirements for progression to Draft Standard status were reviewed. There were at least nine implementations of the MIB represented at the meeting. Donna asked the participants if they felt that there was enough implementation experience. It was agreed that there was enough implementation experience, but perhaps not enough interoperability experience.

Bob Stewart observed that all of the implementations of the MIB are of agents, but that agents don't interoperate with each other; manager implementations are required for interoperability experience. All of the agents have interoperated with "MIB browser" applications, but no known MIB-specific management applications had been written.

The participants agreed that a call should be issued on the mailing list for NMS implementors to let the Group know what kind of applications they're working on, and what implementation and/or interoperability experience they have. Donna and Keith will consider talking to the press to publicize the status of the MIB and encourage implementors to write applications that utilize the Repeater MIB information.

One person observed that the Group had no multiple instantiation implementation experience. It was pointed out that this wasn't part of the standard.

Dave Perkins questioned whether there was enough operational experience with the objects in the MIB. Donna observed that there is considerable operational experience with similar objects in enterprise MIBs.

The participants concluded that there was enough experience to move the MIB forward as a Draft Standard. Therefore it was decided that the editors will make the few agreed-upon changes to the draft and submit the new MIB to the mailing list for three weeks review. If no unresolved issues arise on the mailing list in that time, the Draft will be forwarded to the IESG.

The same actions and schedule are to apply to the MAU MIB.

Second Informal Meeting

About eight Working Group members met informally to discuss informative text that could be added to the Repeater MIB and/or MAU MIB documents to help readers understand the implementation options for the repeater port(s) through which management packets are transmitted and received. The text generated below by this Group will be included in the next Repeater MIB draft, to be reviewed by the Working Group.

- Describe ports as sources of traffic into the repeater, with examples such as:
  - Externally connected devices such as 10BASE-T or AUI.
  - Internal management ports.
  - Backplane internal to implementation.

- Some implementations may not manage all of the ports. For managed ports, there must be entries in the port table.

- It is the decision of the implementor to select the appropriate group(s) in which to place internal ports. GroupCapacity for a given group always reflects the number of managed ports in that group.

- If some ports are unmanaged such that not all packet sources are represented by managed ports, then the sum of the input counters for the repeater will not equal the actual output of the repeater.

Next Meeting

It was agreed not to hold a meeting during the next IETF meeting in Amsterdam.

Attendees

David Arneson	`arneson@ctron.com`
David Battle	`battle@cs.utk.edu`

Andy Bierman	abierman@synoptics.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Jeff Case	case@cs.utk.edu
John Chang	changj@ralvm6.vnet.ibm.com
Shane Dawalt	sdawalt@desire.wright.edu
Manuel Diaz	diaz@davidsys.com
Sandra Durham	sdurham@synoptics.com
David Engel	david@ods.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
Kenneth Key	key@cs.utk.edu
Moshe Kochinski	moshek@FibHaifa.com
Duane Kuang	duanek@kalpana.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
David Perkins	dperkins@synoptics.com
Sam Roberts	sroberts@farallon.com
Dan Romascanu	dan@lannet.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
Paul Serice	serice@cos.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@timon.com
Ira Steckler	isteckle@chipcom.com
Bob Stewart	rlstewart@eng.xyplex.com
Steve Suzuki	suzu@fet.com
Geoffrey Thompson	thompson@synoptics.com
Stephen Tsun	snt@3com.com
Peter Wilson	peter_wilson@3com.com
Kiho Yum	kxy@nsd.3com.com

2.3.10 Interfaces MIB (ifmib)

Charter

Chair(s):

Ted Brunner, tob@thumper.bellcore.com

Mailing Lists:

General Discussion: if-mib@thumper.bellcore.com
To Subscribe: if-mib-request@thumper.bellcore.com
Archive:

Description of Working Group:

The Interfaces MIB Working Group is chartered to accomplish two tasks.

First, to develop a collection of managed objects which model the relation between different entities in the data link and physical layers. The Working Group will explore different modeling approaches in order to develop a collection of objects which is both correct in the modeling sense and has an acceptable impact (if any) on the interfaces table from MIB-II and all media MIB modules on the standards track or under development by a working group. The objects defined by the Working Group will be consistent with the SNMP framework.

Second, to prepare a recommendation to the IESG evaluating RFC1229 (the interface-extensions MIB), RFC1231 (the token-ring MIB), RFC1304 (the SMDS MIB), and RFC1398 (the ethernet-like MIB) with respect to the standards track.

The recommendation will document implementation, interoperability, and deployment experience. If these experiences suggest that changes should be made to the documents, new drafts may be prepared.

For RFCs 1229, 1231, and 1304, the recommendation will report one of four outcomes for each RFC: that the RFC should be advanced from Proposed to Draft status, without changes (if no problems are found); that a draft prepared by the Working Group should replace the RFC, and be designated a Draft Standard (if only minor changes are made); that a draft prepared by the Working Group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or, that the RFC should be designated as historic (if this technology is problematic).

For RFC1398, the recommendation will report one of five outcomes: that the RFC should be advanced from Draft to Full status, without changes (if no problems are found); that a draft prepared by the Working Group should replace the RFC, and be designated a full Standard (if only editorial changes are made); that a draft prepared by the Working Group should replace the RFCs, and be designated a Draft Standard (if only minor changes are made);

that a draft prepared by the Working Group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or, that the RFC should be designated as historic (if this technology is problematic).

Goals and Milestones:

Jul 1993	Post the if layering document as an Internet-Draft.
Sep 1993	Submit the if layering document to the IESG for consideration as a Proposed Standard.
Sep 1993	Issue a call for implementation and operations experience with RFCs 1229, 1231, 1304, and 1398.
Oct 1993	Evaluate experience and if necessary post revised MIBs as Internet-Drafts.
Dec 1993	Submit recommendations on the various MIBs to the IESG.

Internet-Drafts:

"Evolution of the Interfaces Group of MIB-II", 06/04/1993, K. McCloghrie, F. Kastenholz <draft-ietf-ifmib-evolution-00.txt>

2.3.11 Mail and Directory Management (madman)

Charter

Chair(s):

Steve Kille, S.Kille@isode.com

Mailing Lists:

General Discussion: madman@innosoft.com
To Subscribe: mailserv@innosoft.com
In Body: subscribe ietf-madman <email address>
Archive: innosoft.com:~/ietf-madman/archive.txt

Description of Working Group:

The Mail and Directory Management Working Group is chartered to define four MIB modules: one for generic application monitoring, one for message relays (either SMTP or X.400 based), one for OSI Directory service (X.500), and a fourth for message stores. The MIB modules will provide basic monitoring capabilities, and will be consistent with the SNMP framework and existing SNMP standards.

Goals and Milestones:

May 1993 Post an Internet-Draft of the generic application monitoring MIB.

May 1993 Post an Internet-Draft of the message relay monitoring MIB.

May 1993 Post an Internet-Draft of the OSI X.500 Directory Service MIB.

Jul 1993 Submit the directory monitoring, message relay, and generic application MIBs to the IESG for consideration as Proposed Standards.

Jul 1993 Post an Internet Draft of the message store monitoring MIB.

Nov 1993 Submit the message store monitoring MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Network Services Monitoring MIB", 06/01/1993, N. Freed, S. Kille <draft-ietf-madman-networkmib-00.txt>

CURRENT MEETING REPORT

Reported by Urs Eppenberger/SWITCH

Minutes of the Mail and Directory Management BOF (MADMAN)

Agenda

- Minutes of the SAM BOF
- Review of the Charter
- Review of the Documents
  - Application monitoring MIB
  - MTA monitoring MIB
  - DSA monitoring MIB

The Minutes of the SNMP Application Monitoring BOF (SAM), which met during the November 1992 IETF, were accepted without any comments. The Group then moved on to a review of the Charter for the MADMAN BOF. The following decisions were reached:

- Monitoring of Message Store is added to the Charter. Message Store does include X.400(88) Message Store, POP, IMAP...

- The management of the applications is not a working item for the Group. The MIB enables MTA and DSA managers to detect when things go wild by monitoring the most important system variables.

- The application monitoring MIB is planned to be finalised for the Amsterdam IETF.

- The DSA and MTA monitoring MIBs are planned for the Fall 1993 IETF.

- Work on the Message Store MIB will start after the Fall 1993 IETF.

- The IFIP Working Group on email management has a more general approach. Good relations are ensured by cross participation of Working Group members and the sharing of documents.

The remainder of the session was spent in reviewing various documents.

- Application Monitoring MIB
  - It is based on Steve's original proposal and should form a framework for all network applications.

  - A name string will be added.

  - A version indication per application is needed.

  - The concept on representation of processes in the host MIB will be checked.
  - There are already two other MIBs defined, a very short one from HP and one from CDC, concentrating on Mail and Directory.

- DSA monitoring MIB
  - Additional variables: operations pending, active chained operations to other DSAs, current number of requests from clients in process, counters for replications, fail indication.

- MTA Monitoring MIB
  - Support now for multiple MTAs on a single system is added.
  - A new abstraction has been added: channel. This is to get some figures out of very simple MTAs. A problem is deferred delivered messages and alarms set on 'oldest message per channel'.
  - A counter for rejected messages is added.
  - A separate document is needed, also an RFC, to define object identifiers as a basis for identifying TCP services based on their port number.
  - Glenn volunteers to work on a layered model of the MTA monitoring approach. An association layer which receives and sends the messages and a higher level which routes and modifies (in case of gateway functions) the messages.

- Message Store MIB
  - Monitor UAs connected to the message store, submission, retrieval, messages locked in the store.
  - The MIB should support capacity planning.
  - Ed Reed volunteered to work on a document to study the issues involved.

Attendees

Harald Alvestrand	`Harald.Alvestrand@delab.sintef.no`
Robert Austein	`sra@epilogue.com`
Cyrus Chow	`cchow@ames.arc.nasa.gov`
Robert Cooney	`cooney@wnyose.nctsw.navy.mil`
Urs Eppenberger	`eppenberger@switch.ch`
Ned Freed	`ned@innosoft.com`
Terry Gray	`gray@cac.washington.edu`
Roland Hedberg	`Roland.Hedberg@rc.tudelft.nl`
Barbara Jennings	`bjjenni@sandia.gov`
Thomas Johannsen	`Thomas.Johannsen@ebzaw1.et.tu-dresden.de`

Kevin Jordan	Kevin.E.Jordan@cdc.com
Marko Kaittola	Marko.Kaittola@funet.fi
Steve Kille	S.Kille@isode.com
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Bruce Mackey	brucem@cinops.xerox.com
Glenn Mansfield	glenn@aic.co.jp
Ignacio Martinez	martinez@rediris.es
Bob Morgan	morgan@networking.stanford.edu
William Nowicki	nowicki@legato.com
Geir Pedersen	Geir.Pedersen@usit.uio.no
David Perkins	dperkins@synoptics.com
Edward Reed	eer@cinops.xerox.com
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Yzhak Ronen	y.ronen@homxa.att.com
Chris Shaw	cshaw@banyan.com
Sue Smith	smiths.es.net
Einar Stefferud	stef@nma.com
Catherine Summers	cfs@cos.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de

2.3.12 Modem Management (modemmgt)

Charter

Chair(s):

Mark Lewis, `Mark.S.Lewis@telebit.com`

Mailing Lists:

General Discussion: `modemmgt@Telebit.com`
To Subscribe: `majordomo@Telebit.com`
In Body: **subscribe modemmgt <email address>**
Archive: `ftp.telebit.com:~/pub/modemmgt`

Description of Working Group:

The Modem Management Working Group is chartered to define a MIB module for dial-up modems and similar dial-up devices. This MIB module will provide a set of objects that are the minimum necessary to provide the ability to monitor and control those devices, and will be consistent with the SNMP framework and existing SNMP standards.

The Working Group will consider existing specifications including the RS-232-like, Character, PPP and other related MIB modules. It will consider enterprise-specific MIB modules which support modem-like devices. The Working Group will also consider the TSB Study Group 14s work on an OSI CMIS/CMIP object definition for V series DCEs entitled "Managed Object Template for V-Series DCE's".

Goals and Milestones:

Jun 1993 Post an Internet-Draft of the Modem Management MIB.

Oct 1993 Submit the Modem Management MIB to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Mark S. Lewis/Telebit

Minutes of the Modem Management BOF (MODEMMGT)

There were approximately 25 attendees made up of modem manufacturers, SNMP developers, and modem users/service providers. The BOF was begun with a discussion of the need to manage modems by users and service providers. The Group explored the need to monitor and control modems, as well as account for their usage. It was clear that the ability to change the configuration of modems was important. It was suggested that other devices like fax, DSU/CSUs, and perhaps terminal adapters should be treated similarly. Most agreed that the Group should be as inclusive as possible in writing a 'modem' MIB.

```
+-----------+-----------+
| Character |    PPP    |
|-----------+-----------|
|      RS-232-Like      |
|-----------------------|
|         Modem         |
+-----------------------+

          MIB Hierarchy
```

Relevant MIBs including the RS-232-Like, Character, PPP, and DS1/E1 (RFC1406) were briefly discussed. Enterprise MIB extensions were considered including those from Telebit and US Robotics. A short presentation was made by Les Brown (TIA TR-30.4) on a Draft Recommendation V.im; Management Information Model for V-Series DCE's. It was noted that some vendors had looked at this Draft and had implemented a subset of it. This may be the essential portion of what seems to be a very comprehensive definition in the Draft.

Some raised issues about which problems the Group was trying to solve. Others suggested that a clear statement about the problems would provide helpful direction. Some questioned the model being considered: whether it was geared for a stand-alone device or a rack installation. It was agreed that a rack model should be used which supports multiple devices through tables in a MIB rather than single objects.

There was a question about whether or not it was important to produce a general modem MIB. It was suggested that some want extensive enterprise MIB support from their particular vendor. Some of the SNMP management station developers said they would not be willing to add support for each vendors enterprise MIB extensions, but would be willing to add support for a general modem MIB. It was agreed a general modem MIB was necessary.

There was discussion about how many objects might be included in a modem MIB. The V.im Draft seemed to define 200-250 objects. Some thought it would be good to limit the modem MIB to around 100 objects. As a comparison, the DS1/E1 MIB (RFC1406) defines approximately that many objects.

Given the base of relevant work, the BOF Chair proposed that an IETF working group be formed to work in conjunction with TR-30.4 to produce a joint standard modem MIB. A proposed Charter was presented and discussed. There was discussion about the different cultures of the two Groups. There was some skepticism that the two Groups could produce a joint document. The Chair proposed that discussions should be carried-out through e-mail and periodic face-to-face meetings.

A proposed schedule was discussed which set certain goals. Basically, the joint document would be completed in 1993. This means that the joint Group would produce at least two drafts, which would be reviewed face-to-face at least twice. Presumably, there would be extensive e-mail discussion in the interim. Most considered this an aggressive schedule but thought it was a good objective.

The BOF concluded with agreement that an IETF working group should be organized to pursue these goals. The Chair agreed to update the proposed Charter and submit it to the acting Management Area Director. He also agreed to setup the required mailing list, and ftp repository of relevant MIBs and documents.

Attendees

Jim Barnes	barnes@xylogics.com
David Battle	battle@cs.utk.edu
Larry Blunk	ljb@merit.edu
Thomas Brisco	brisco@pilot.njin.net
Douglas Carson	carson@utcc.utoronto.ca
Jeff Case	case@cs.utk.edu
Naomi Courter	naomi@concert.net
Mark Davis-Craig	mad@merit.edu
Robert Gutierrez	gutierre@nsipo.nasa.gov
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Ittai Hershman	ittai@ans.net
Mark Kepke	mak@fc.hp.com
Mark Lewis	Mark.S.Lewis@telebit.com
Ray Mlejnek	r.mlejnek@pdn.paradyne.com
Chris Payson	!sharps!bar!c_payson@telebit.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
Jon Saperia	saperia@lkg.dec.com
Timon Sloane	timon@timon.com
Bob Stewart	rlstewart@eng.xyplex.com
Richard Stuart	dickstuart@attmail.com
Steven Waldbusser	waldbusser@andrew.cmu.edu

2.3.13 SNA DLC Services MIB (snadlc)

Charter

Chair(s):
Jeff Hilgeman, jeffh@apertus.com

Mailing Lists:
General Discussion: snadlcmib@apertus.com
To Subscribe: snadlcmib-request@apertus.com
Archive:

Description of Working Group:

The SNA DLC Working Group is chartered to define a set of managed objects for the SDLC and LLC-2 data link controls for SNA networks. These objects will be the minimum necessary to provide the ability to monitor and control those devices, providing fault, configuration, and performance management, and will be consistent with the SNMP framework and existing SNMP standards.

The Working Group will consider existing enterprise-specific MIB modules that define objects which support management of these devices. The Group may choose to consider any work done by the IEEE in the area of managed object definition for LLC-2. It will also make sure that its work is aligned with the SNA NAU Services MIB Working Group, due to the close relationship between the devices being worked on by the two Groups.

The Working Group recognizes that managed objects for other SNA data link controls and related components (e.g., QLLC, System/370 Channel, DLS (Data Link Switching) and ESCON) may need to be identified in the future. These objects are out of scope for the current Charter; however, once the Group completes its Charter, a new Charter identifying some or all of these components may be considered.

Goals and Milestones:

Apr 1993 Mailing List discussion of vendor proprietary MIBs.

Jul 1993 Post an Internet-Draft of the SNA DLC MIB.

Dec 1993 Submit the SNA DLC MIB to the IESG for consideration as a Proposed Standard.

2.3.14 SNA NAU Services MIB (snanau)

Charter

Chair(s):

Zbigniew Kielczewski, zbig@eicon.qc.ca
Deirdre Kostick, dck2@sabre.bellcore.com

Mailing Lists:

General Discussion: snanaumib@thumper.bellcore.com
To Subscribe: snanaumib-request@thumper.bellcore.com
Archive:

Description of Working Group:

The SNA NAU Services MIB Working Group is chartered to define a set of managed objects for PU type 2.0, and LU type 1, 2, and 3 devices for SNA networks. These objects will be the minimum necessary to provide the ability to monitor and control those devices, providing fault, configuration, and performance management, and will be consistent with the SNMP framework and existing SNMP standards.

The Working Group will consider existing enterprise-specific MIB modules that define objects which support management of these devices. It will also make sure that its work is aligned with the SNA DLC Services MIB Working Group, due to the close relationship between the devices being worked on by the two Groups.

The Working Group recognizes that managed objects for other components (e.g., PU Type 4, PU Type 5, LU Types 1, 3, 4, 6.2 (APPC), APPN EN, APPN NN and APPI) may need to be identified in the future. These objects are out of scope for the current Charter; however, once the Group completes its Charter, a new Charter identifying some or all of these components may be considered.

Goals and Milestones:

Jul 1993 Begin discussion of proprietary MIBS and develop a single proposal.

Oct 1993 Post an Internet-Draft of the SNA NAU Services MIB.

Dec 1993 Submit the SNA NAU Services MIB to the IESG fo consideration as a Proposed Standard.

2.3.15 Token Ring Remote Monitoring (trmon)

<u>Charter</u>

Chair(s):

Michael Erlinger, mike@jarthur.claremont.edu

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 (RFC1271) 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.
Done	Working Group meeting at San Diego IETF.
Mar 1992	Post Internet-Draft of the Token Ring Monitoring MIB.
Done	Working Group meeting at Cambridge IETF.
Nov 1992	Submit the Token Ring MIB to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Steven Waldbusser/CMU

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

The TRMON Working Group met for one session. At that session, the final outstanding technical issues for the draft were identified and resolved. There was consensus that the resulting draft should be submitted to the Network Management Directorate for eventual publication as a Proposed Standard.

The Group then discussed priorities for future work and where a next meeting might take place. There was no clear resolution on these issues. Finally, in the remaining minutes, a few implementation issues for RFC1271 were discussed.

Identify/Resolve Outstanding Issues

Editor's Note (md): A detailed listing of the outstanding issues is available via ftp under trmon-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Call for Consensus

The Chair notes that there were no outstanding issues left for the MIB. He asked if there was consensus that the Group should submit this MIB to the SNMP Directorate for review and publication as a Proposed Standard. No objections were noted.

Discussion of Future Work Priorities

The following items were discussed:

1. Updating RFC1271 from Proposed to Draft Standard (this involves fixing bugs and known interoperability problems).
2. FDDI RMON Extensions.
3. RMON2 - New features for RMON (Network layer information, protocol distributions, etc.).
4. WAN interface Extensions. It was agreed that moving RFC1271 from Proposed to Draft should be the highest priority.

Discussions of Future Meetings for RMON-related Activity

If the Group were to tackle one of these issues, the earliest opportunity to meet would be at July IETF in Amsterdam. The merits of this were discussed, but no resolution was reached.

Discussion of Implementation Issues

1. A vendor noted an interoperability problem when his probe was having an alarm entry created by certain other managers. These managers would create the entry and then immediately issue get requests for the entire row (include alarmValue). His probe would return noSuchName because the first alarm interval had not passed yet and the alarmValue instance had not yet been created. The managers would crash or otherwise refuse to interoperate due to the noSuchName error.

   The consensus was that the probe was compliant to RFC1271 and that management stations should not break if the alarmValue is not present immediately after row creation. An even more robust strategy would be to handle the lack of alarmVariable at any time.

2. A vendor noted an interoperability problem when his probe was being queried by a certain manager. This manager would query the history table, assuming that three history entries had been configured at startup, apparently as that vendor does on its own probe. When the three entries weren't found, the manager crashed or otherwise refused to interoperate.

   The consensus was that the manager was the cause of the interoperability problem and that it shouldn't assume any configuration.

   It was noted however, that while the probe was not obligated to have any particular setup of history entries, that the two entries suggested in the MIB should be configured for the probe to be the "best citizen" possible.

3. A recent mailing list discussion was brought up. The RFC1271 host table specifies that "The host group discovers new hosts on the network by keeping a list of source and destination MAC Addresses seen in good packets." The rational for only using good packets is that bad packets may fill the table up with random MAC addresses. It was noted that short and long packets (with correct CRC) are likely to have correct MAC addresses, and might be appropriate for gleaning new hosts.
   Everyone agreed that the specification is currently very clear on this issue, but that it might be reasonable to modify it in future versions to allow the use of MAC addresses in short and long addresses.

Attendees

David Arneson	arneson@ctron.com
David Battle	battle@cs.utk.edu
Andy Bierman	abierman@synoptics.com
John Chang	changj@ralvm6.vnet.ibm.com
Anthony Chow	chow_a@wwtc.timeplex.com
Manuel Diaz	diaz@davidsys.com
Sandra Durham	sdurham@synoptics.com

David Engel	david@ods.com
Kenneth Giusti	kgiusti.chipcom.com
Daniel Hansen	dan@ngc.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
Jeff Hughes	jeff@col.hp.com
Kenneth Key	key@cs.utk.edu
Carl Madison	carl@startek.com
Evan McGinnis	bem@3com.com
Tom Nisbet	nisbet@tt.com
Jon Penner	jjp@bscs.uucp
David Perkins	dperkins@synoptics.com
Narendra Popat	albin@frontier.com
Venkat Rangan	venkat@.metrix.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Paul Serice	serice@cos.com
Ira Steckler	isteckle@chipcom.com
Richard Sweatt	rsweatt@synoptics.com
Geoffrey Thompson	thompson@synoptics.com
Stephen Tsun	snt@3com.com
Peter Wilson	peter_wilson@3com.com
Kiho Yum	kxy@nsd.3com.com

2.3.16 Uninterruptible Power Supply (upsmib)

Charter

Chair(s):

Jeff Case, case@cs.utk.edu

Mailing Lists:

General Discussion: ups-mib@cs.utk.edu
To Subscribe: ups-mib-request@cs.utk.edu
Archive: ucs.utk.edu:~/pub/ups-mib/mail-archive

Description of Working Group:

This Working Group will produce a document that defines MIB objects for use in monitoring and (possibly) control of both high-end and low-end UPSs and related systems (e.g., power distribution systems or power conditioning systems). Related devices may be addressed in this effort to the extent that the primary focus on UPSs is not compromised.

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

At its discretion, the Working Group may fulfill its Charter by the development of distinct MIB definitions for UPS systems of differing capabilities, but the number of MIB definitions produced by the Working Group will not exceed two.

At its discretion, the Working Group may produce an additional document defining traps that support the management of UPSs.

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 UPS implementations varies widely, the working group shall take special care that its definitions reflect a generic and consistent architectural model of UPS management rather than the structure of particular UPS implementations.

Goals and Milestones:

Done Hold Interim Working Group meeting to review draft.

Done Post initial draft MIB to Internet-Drafts.

Mar 1993 Meet at March IETF meeting to reach closure on MIB document.

Apr 1993 Submit the UPS MIB to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Uninterruptible Power Supply Working Group (UPSMIB)

The UPSMIB Working Group held a meeting at the IETF meeting in Columbus, Ohio, on Thursday, 1 April 1993. Jeff Case and Ken Key presided and shanghaied Bob Stewart to record.

General Discussion

There was a review of the IETF process based on RFC1310, mentioning the terms Internet-Draft, Proposed/Draft Internet Standard. Final review and approval is by the IESG.

The history of the document is: two company submissions at the BOF in Cambridge, synthesis, discussion, strawman proposals, surveys and analysis. The draft is not final yet and still requires consensus on what to include.

Cycle time for this Group is slower due to less direct Internet involvement.

UPSMIB is the first IETF MIB that is not a direct component of communication; a "toaster" MIB. UPSs have been remotely managed for a long time, providing an important experience base. A Group member suggested that the Group should get on with the real work.

Brief introductions around the room revealed about fifteen vendors, two users, and a scattering of telecommunications, NMS, and agent people.

The Group needs to decide about its next meeting based on progress in the near future. About 8-10 of the present attendees could come to Amsterdam. It's a long trip for a single meeting, and most have no other IETF interests.

Review of Strawman

The Strawman had been sent out via email a week prior to the IETF meeting. Several questions were raised.

- Should enterprise-specific OIDs be one per MIB or Group?
- Should input frequence resolution be 0.1 Hz?
- What is the definition of blackout versus brownout? Should the MIB include brownout? The ensuing discussion of the blackout counter took most of the meeting.

Editor's Note (md): A detailed listing of the discussion of the blackout counter is available via ftp under upsmib-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Presentation of a Contact Closure MIB - by Ray Wasson

Ray's presentation included the following points:

- UPSs need a simple MIB group for contact closure devices. A handout of four objects: upsControlOutputOffDelay, upsControlAudioOutput, upsTestBattery, and upsTestBatteryStatus. This is a minimum, and less than the previously-proposed basic MIB.

- It could use some additions from the original basic proposal. Some basic objects, such as location, are in MIB-II.

- This is preferable as a simple, separate document, forwarded separately. It is best if basic information is common across all levels of the MIB. A small MIB is not a direct subset, but picks from groups and loses context.

- Jeff will discuss separate documents with the SNMP Directorate. Compliance groups will cover this without separate documents. There is concern that waiting for the whole thing takes too long.

- There was no consensus on the proposal. The Group was not willing to decide without understanding new compliance specifications. The Group needs to resolve conflicts and move forward.

- For email: Is the contact closure proposal or the original basic MIB a possible separate document. Compliance with this MIB should require naming compliance groups which are implemented. The consensus was for Jeff to see if that is allowable. He expects the SNMP Directorate to allow two but not more. Jeff will email the SNMPv2 compliance document.

MIB Group Spokesmen

Volunteers were found to act as spokesmen to push progress on each MIB group:

Tom Brennan	Configuration
Adam Stolinski	Control
Terry Zumwalt	Alarm
Roger Draper	Bypass Output Input
Phillip Epps	Battery
Steve Held	Identity

Doug Rademacher	Test
Jeff Case	Traps

Attendees

Jeff Case	case@cs.utk.edu
Shane Dawalt	sdawalt@desire.wright.edu
Manuel Diaz	diaz@davidsys.com
Roger Draper	
William Elliot	oneacwre@aol.com
Phillip Epps	76410.327@compuserve.com
Ronald Pitt	
Doug Rademacher	
Rick Royston	rick@lsumvs.sncc.lsu.edu
Timon Sloane	timon@timon.com
Bob Stewart	rlstewart@eng.xyplex.com
Adam Stolinski	stolinsk@cerf.net
Geoffrey Thompson	thompson@synoptics.com
Brian Young	byoung@macc.wisc.edu
Terry Zumwalt	

2.4 Operational Requirements Area

Director(s):

- Phill Gross: pgross@nis.ans.net
- Bernhard Stockman: boss@ebone.net

Area Summary reported by Bernhard Stockman/SUNET

Generic Internet Service Specification BOF (GISS)

Tony Bates gave a brief overview of a project to try to produce a specification for a "Generic Internet Service Specification". The primary emphasis of this work is at the "provider/provider" interface rather than at the "user/provider" interface. The goal is to make it easier for new service providers to understand and interwork with various other service providers. The plan is to have a specification document that will need to be updated and also highlight areas for further study or beyond scope.

A pointer was raised to the FYI16 document which could be augmented slightly to cover the "user/provider" interface in a less "U.S.-centric" approach. However, this is not the primary focus of this current specification.

Within the brainstorming session some concerns were raised as to whether such a document could mandate items that a service provider should provide. The document would raise issues rather than mandating anything. It was clear that the document would have to be revised on a regular basis. A list of "first-pass" items for the specification were worked through. Many of the items could easily fall into more than one category. The first pass list will cover the following items:

- Routing issues
- Addressing
- Information Provision
- Operations
- Connectivity
- Engineering and Maintenance
- Attachment
- Generic Services Coordination
- Other networks
- AUP (more than routing)
- Remote/Local management

Tony Bates will draft the details and circulate to the giss list. Those wishing to join the giss list can do so by sending a message to: giss-wg-request@ripe.net

Dan Long, John Curran and David Conrad will act as reviewers. Daniel Karrenberg will follow up on the revision of FYI16.

It was decided not to ask for an IETF working group at this stage. The draft will be sent out to the ORAD list to see whether or not there is enough interest to create a working group.

MBONE Engineering and Operations BOF (MBONE)

Initially some general issues were discussed.

The Mbone is dangerous, but the applications are very useful. Critical problems are largely due to the lack of tools to manage the mbone. The meeting discussed different approaches to deal with problems from various groups involved in the mbone community such as network operators, network subscribers and end users. The Group unanimously agreed that it is bad practice for network operators to support mbone tunnels into other regions to solve the problem that a subscriber does not get satisfactory service from his network operator.

The meeting continued with a discussion of the new encapsulated tunneling code. The original code is a seriously bastardized use of LSRR seriously violating the IP specifications. The new mrouted supports both, defaulting to encapsulation. There was talk of adding clean source routing options to the encapsulating code such that network operators could prevent tunnels from moving to fall back infrastructure during network failures.

Most of the rest of the meeting was spent drafting a wish list. Some of the items were appropriate for a future meeting of this BOF or Working Group. Other items were appropriate for specific groups or projects involved in multicast research. The items are ordered by priority within each section, but the sections were independent of each other. Throughout the discussion it was understood that resources are tight, and in many cases people were being asked to contribute effort without additional support.

The wish lists were presented to both the AVT Working Group and to the Operational Requirements Area Directorate (ORAD). There was general agreement that the items on the wish list are desirable and appropriate, but nobody agreed to implement anything.

There were some network operators present at the ORAD meeting who were upset that the MBONE BOF did not become a Working Group or take stronger positions regarding operational practices. However most of these people were operators who had chosen not to participate in the mbone, and were therefore not in control of its impact on their own facilities.

BGP Deployment Working Group (BGPDEPL)

BGP-4 deployment:

ANS/NSFnet/GATED	BGP4 is working in a test mode.

CISCO	BGP4 is under still under development.
3-COM	Expects to support BGP4 in early this fall.
Wellfleet	Anticipates rolling out BGP4 support this summer.
Telebit/EuropaNet	No current plans for BGP4 deployment.

CIDR core plans (Alternet, CIX, EBONE, NSFnet/ANS, NSI, PSI, Sprint):

1. Start deploying BGP4 code as soon as possible.
2. NSI (or some alternative) starts announcing one aggregated network.
3. Additional CIDR core members start aggregating networks.
4. Aggregation is officially "turned on" in the Internet.

The members of the CIDR core are progressing as fast as possible, and are well coordinated among themselves.

CIDR configuration issues:

- There is some controversy over how to do global configuration checking. In addition, how can one ensure topology matches policy? Merit presented a preliminary plan for aggregation support in the NSFnet. This support would: 1) accept aggregate routes from a midlevel, or 2) accept site routes from a midlevel, and aggregate on the midlevels behalf. A strawman database format proposal is documented in the "Inter-domain Routing Policy Description and Sharing" Internet-Draft (draft-yu-rpd.00.txt).

- The representatives from RIPE pointed out that the existing U.S. databases, including the current Merit configuration database and the above proposal are not adequate to solve international routing problems. In particular none can be used to determine which backbone (CIDR core member) is the preferred path to a given U.S. network.

- The sense of urgency came primarily from concerns about configuration management and database issues. Although there is still a lot of work to be done to complete the BGP4 roll out, it seems to be a fairly well understood problem except for configuration management. CIDR and BGP4 do impose some new requirements on the databases but the majority of the issues center around topology and AUP enforcement. For these reasons it makes sense to broaden the scope of this Working Group from just BGP deployment to the wider task of fostering sanity in topology, routing policy, and configuration databases.

Network OSI Operations Working Group (NOOP)

Russell Blaesing talked about his Transport MIB. He is going to put his implementation up for anonymous ftp. He is also going to submit a new version of the draft because the old one expired. It was unclear whether the MIB needed variables added or removed, so the Group will put it out there and see how it works.

The Essential Tools for the OSI Internet was discussed. With some minor changes it is going to be submitted as a Proposed Standard. The Echo draft will also be submitted as a Proposed Standard.

Work is going to be put into trying to keep the OSI infrastructure up and operational. Sue Hares is going to have one of her folks put the EON tunnel configuration up for anonymous ftp along with information about reachable CLNS hosts throughout the Internet. She will also try to get someone to start pinging these hosts (via CLNS) and sending mail if there are outages.

The deployment of TUBA was discussed. Several implementations are available and folks are interested in trying them.

The Group concluded that unless there is more work that needs to be done with the deployment of TUBA that NOOP will probably suspend their work for a time until they are needed again.

Operational Requirements Area Directorate (ORAD)

The ORAD mandates were discussed. It was agreed that ORAD was a good forum for discussing operational related items among network service providers. Thus the purpose of the meeting should be to coordinate operation of individual networks, not to change each networks own policy.

It was pointed out that many Standard RFC's have fallen through the cracks towards complete implementation without operational concerns having been addressed. John Curran agreed to make sure that at least a fraction of new RFC's are read for operational impact. Bill Manning agreed to do some reviewing, but cannot do the whole job himself.

There is a need for a working group to deal with a policy routing description language, Many of the routing efforts (BRG, SDRIP, etc.,) are defining a need for a common routing policy language. ORAD needs to form a liaison with the protocol developers to help define such a language.

The Operations Area working groups were discussed and a new scheme was proposed by Dan Long. The intentions and actual planning according to this scheme will continue on email.

The need for a tunnel coordination working group was discussed. There are MBONE tunnels which could be removed. TUBA tunnels may in the future need coordination for the same reasons. ORAD did not see a need for such a working group in the immediate future.

CIDR issues were treated with respect to timeliness. Will CIDR deployment be in time before router hard- and software start to hit the limits or has this already begun to happen.

There is a need to find adequate measurements here. ANS is doing some investigation within this area.

Operational Statistics Working Group (OPSTAT)

Reports on current efforts in deploying the OPSTAT model (RFC1404) revealed some work. RIPE NCC has a tool known as Monster which could be adapted but manpower is lacking. Craig Haney has written a PERL parser. Some other efforts are also known. FARNET has promised funding if a site could be found to take on the implementation work.

Some problems with the RFC1404 storage format were reported and it was decided to make the necessary changes to fix them.

Henry Clark has done some implementation work on the OPSTAT client/server draft specification. This was discussed and Ittai Hershman reported on a similar tool implemented by Merit. The specification of the client/server query language was extensively discussed.

Finally it was agreed that the SNMP/MIB people should be contacted with respect to variables needed in statistical gathering but which as of today are not present in the Internet Standard MIB.

CURRENT MEETING REPORT

Reported by Daniel Karrenberg

Minutes of the Generic Internet Service Specification BOF (GISS)

Presentation

Tony Bates gave a brief presentation of the project. The project is a RARE/RIPE joint project and is funded by SURFnet through RARE. The project is open and would like input from Internet Service (IS) providers wherever possible. It was still at the stage of defining the scope and structure for a specification document which was hoped could be used, not as a mandatory document, but as a document specifying the relevant issues in Internet Service provision. The primary focus is at the service provider level focusing on coordination and information on what new (and possibly old?) service providers need to know. The Internet Service interface is between different service providers. The user/provider interface appears to be covered other documents. FYI16 being a notable example. A plan for the project has been produced plus some very draft text to start the discussion rolling.

Brainstorming

If things are too fixed, it could be difficult for people who try to provide a useful service to do things a bit differently. For instance, CIDR is an issue now but it may not be an issue next year because everyone is doing it. The bottom line is that this document needs to be revised regularly.

Customers should be able to see from the document how innovative the provider is. It should not be a constraining item for providers but more of a helping document. The problem that we want to solve is that it becomes more unclear what an Internet Service is. Performance is just a part of it. There is something like a service level agreement, but no service guarantees. We want to say things that can be agreed upon between a customer and provider, but without figures. Maybe just hints to what kind of things a customer can ask his provider (i.e., last statistics, access to first router).

The document should talk about the issues, but not mandate. Enumerating the issues would not be a controversy. Preferably not "Thou shall." All the bases should be covered. Maybe some advise on the speed an organisation may need (organisations without prior knowledge). All existing attempts at this type of document have died because providers did not want the IETF to tell them what they should do.

What kind of reporting do you want to have from your provider? What kind of backup of general routing arrangements does a provider have?

Some details of items to be in the specification were worked through. *Editor's Note (md): A detailed listing of these items is available via ftp under giss-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Concrete Actions

Tony Bates	Tony Bates would try to get all these items done using the basic structure of the first draft. The document would be reviewed by the following reviewers: Dan Long, David Conrad and John Curran
Daniel Karrenberg	Follow-up on the revision of FYI16. Make sure that it is less U.S. centric.

To Working Group or Not to Working Group?

Service providers need to be kept informed of the document progress. The draft will be sent to ORAD. The Group will see if there is enough critical mass to create a working group at the next IETF in Amsterdam. The word must be spread among providers to make sure they do not feel left out and to ensure they have more input.

Attendees

Tony Bates	tony@ripe.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Henry Clark	henryc@oar.net
Robert Collet	rcollet@icm1.icp.net
David Conrad	davidc@iij.ad.jp
John Curran	jcurran@nic.near.net
Kishan Dudkikar	kishan@icm1.icp.net
Alisa Hata	hata@cac.washington.edu
Daniel Karrenberg	daniel@ripe.net
Daniel Long	long@nic.near.net
James Miner	jjm@fibercom.com
Peder Chr. Noergaard	pcn@tbit.dk
Vilson Sarto	vilson@fapq.fapesp.br
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Kirk Williams	kirk@sbctri.sbc.com

"GENERIC INTERNET SERVICE" SPECFICATION (BOF)

Tony Bates

RARE

RIPE

What is Giss ?

- RARE project as part of the RARE Technical Program being carried out at the RIPE NCC
- Project funded by SURFnet through RARE
- Open project
  - No Hidden agendas
  - Should not be seen in any way as a competitive / political document
- Not just a paper exercise
  - Should be worthwhile and useful for new and old in the internet

What is Giss ?

- Production of a document describing a **'useful'** Internet service in terms of service provision
  - Should cover ALL aspects of Internet Service Provision
- Needs to look at Internet services from both sides of the fence:
  - USERS

    "I've heard about IP and the Global Internet but what can I expect from my service provider"
  - PROVIDERS (PRIMARY EMPHASIS)

    "What do I need to do as a service provider to be part of the Internet"

Ideas

- Original idea for Internet Service Model

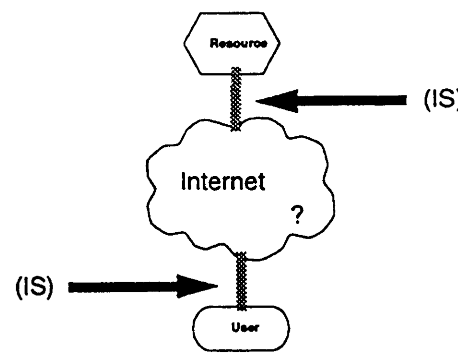

Inside the Cloud

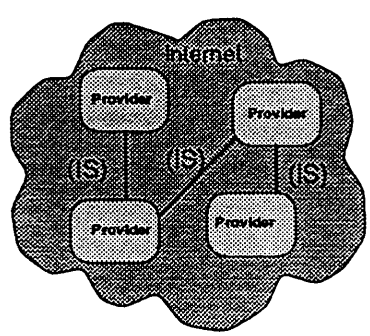

One step Further

- Need to focus on the Internet provision to the Internet rather than to the user so...

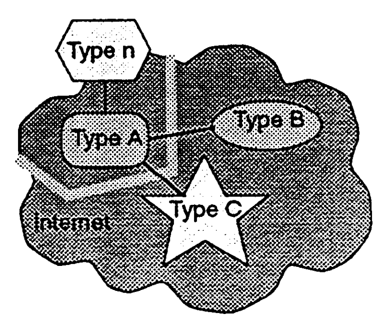

- Various catergories of provider

Ideas

- Need Definitions for clarity
- WHY ?
  - Create consistency and also make specification easier to read
  - Possibly of use to wider audience
- WHAT ?
  - Service provider
  - User
  - Resource
  - Internet
  - others ???

Aspects

- WHY ?
  - Use as a pointer to relevant parts of the specification
  - Not to be used as a competition tool -- this MUST be clear in the specification
- WHAT ?
  - Focus on "IP" based interface between service providers
  - Create a list of "service requirements"
    - -- Access
    - -- Connectivity
    - -- Support

Aspects

- Service support
  - -- Cover
  - -- Helpdesk facility
  - -- UPS
  - -- Information Services
- Management Aspects
  - -- Remote monitoring
  - -- Statistics
  - -- Configuration management
  - -- Change control
  - -- MTBF

Scope

- Aimed at service providers
- Guidance on Routing and connectivity co-ordination (IP addressing issues, Registration, CIDR, etc.)
- Deal with provider/provider interface
- Stress the "**global**" aspects of the Internet
- Need to look at provider relevant aspects
  - In detail where needed. i.e.
    DNS, Routing, etc......
- Check the scope is realistic in timescales

Outcome

- Service specification document
  - requirements (checklist)
- Act as a tutorial for service providers
  - "do's and don'ts"
  - Hints (possible soap boxes)
- Mark areas as for "further study"
- Work towards a proposal for a "service provider" profile

Anything Else ?

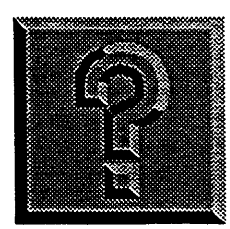

All ideas Welcome !!!!!!

CURRENT MEETING REPORT

Reported by Matt Mathis/PSC

Minutes of the MBONE Engineering and Operations BOF (MBONE)

IETF Organizational Discussion

The Group debated the need for a formal MBONE Working Group. This requires someone to volunteer to be the Chair and to draft a Charter. After some inconclusive discussion it was observed that no one was willing to volunteer. The people present were mostly network operators who are participating in the mbone. Unfortunately a number of network operators who feel victimized by the mbone were not present.

General Issues

- The Mbone is dangerous, but the applications are very useful, and may drive the next generation Internet technology.
- Critical problems are largely due to the lack of tools to manage the mbone.
- There are at least three constituents in the Mbone/multicast community: Network operators, who are providing the mbone core (at least in the NSFnet context), network subscribers, who are typically mbone stubs and finally, mbone end users. There was considerable discussion about these groups, and their identity. These three groups have substantially different cultures, approaches to problems and worries.
  - The network operators have a service oriented perspective and an intimate understanding of the underlying topology. Almost all of the people present were network operators.
  - The subscribers are typically researchers who want mbone connectivity but are not really interested in the details. They usually operate stub tunnels to connect campuses into the mbone.
  - The mbone end users are most likely applications developers or true users, and use local area multicasting to reach a subscriber tunnel. This community is likely to have no knowledge about operational details of the mbone.
  - [I have since realized that there are a number of core mbone hubs which are managed by multicast researchers on the premises of a network operator, with only minor supervision by the operator. This straddles the operator/subscriber distinction above.]

- The Group discussed the mbone mailing list. Some people wanted to split the mailing list such that operators could have candid discussions about debugging without kibitzing from overzealous users. The consensus was not to make any changes and use the mbone list as it stands.

- The hypothetical problem came up about a subscriber who was not getting satisfactory service from his network operator. Would operators support tunnels into other regionals? The Group was unanimous that this is a bad practice, and shouldn't be allowed. Furthermore, subscriber to subscriber tunnels are even worse, and the operators should not provide such poor service that their subscribers resort to such tactics. There was some discussion of the business implications to network operators.

- We discussed the new encapsulated tunneling code. The original code is a seriously bastardized use of LSRR. The new mrouted supports both, defaulting to encapsulation. Use the "srcrt" directive in mrouted.conf to get LSRR. Matt pointed out that those who most endanger the rest of the Internet were also those who didn't require the new encapsulation code for performance. A long discussion of capabilities ensued. One salient point was that the LSRR code seriously violates the IP specifications. There was talk of adding clean source routing options to the encapsulating code such that network operators could prevent tunnels from moving to fall back infrastructure during network failures. This would cause multicasting load to be shed during failures of primary IP connectivity.

Most of the rest of the meeting was spent drafting a wish list. Some of the items are appropriate for a future meeting of this BOF or Working Group. Other items are appropriate for specific groups or projects involved in multicast research. The items are ordered by priority within each section, but the sections are independent of each other. Throughout the discussion it was understood that resources are tight, and in many cases we were asking people to contribute effort without additional support.

Items for Network Operators or Some Future Mbone Working Group

- Put more pressure on router vendors to provide implementations that perform well with LSRR.

- Generate an mbone operational guidelines document. This could be started by splitting the existing FAQ document into separate user and operator documents.

- Explicitly engineer the mbone topology, metrics, and thresholds. This is a daunting task with insufficient tools or poor knowledge of the actual Internet topology.

- There needs to be a global policy on bandwidth budgeting and allocation. The mbone is currently a single global resource, and must be allocated as such. Sometimes there will be two groups with legitimate reasons to do multichannel world-wide multicasts, and there must be some mechanism to arbitrate between their needs.

Infrastructure Tools, to Help Engineer and Operate the Mbone Itself

- Three different mapping tools:
  - A configuration map which shows all configured tunnels, even if they are not "up" - to discover configuration anomalies.
  - A tunnel map (done: Pavel Curtis's tool).
  - A flow map, showing the distribution tree for an arbitrary transmitter.
- Alarms that can be triggered by excessive mbone traffic, such that sites with large pipes can protect the rest of the Internet.
- A "tunnel trap" to detect unauthorized (e.g., client-to-client) tunnels passing through a regional. Several algorithms were discussed.
- SNMP/opstats style tools for logging load (traffic) levels.
- Map tracing - proxy IP traceroute along tunnels to detect when they share the common infrastructure.
- Traceroute (follow the path from the transmitter to a receiver) - deemed to be almost the same as the flow map but not as useful. (A Flow map gives the entire flow topology.)

Transport Tools, to Verify, Monitor and Diagnose Signal Quality

These should use the Audio Video Transport protocol (AVT) directly without specific knowledge of the applications, except to mimic aggregate traffic statistics. All should be supported on all platforms supporting mrouted. (Not just the platform supporting some particular application.)

- Signal quality meters - To display packet loss and delay statistics throughout the distribution tree. It is critical that this tool run on every mrouted platform.
- (Talk) Spurt correlated signal quality - Display packet loss and delay statistics correlated with the start of talk spurts - to detect interactions due to competition between the mbone and other Internet traffic. (TCP congestion avoidance takes one round trip time to back off, so congested links are often late/lossy during the initial second of a talk spurt.)
- Basic test generator - generate traffic streams to mimic the first order statistics of the more popular applications (Correct average packet rate and size.) Again it is important that this tool not require the actual transmitting platform, because the lead time to acquire the transmitters has historically prevented the IETF multicasts from being tested in advance.

- Modulation test generator: Transmit 5 second bursts ("spurts") of the above generator separated by 5 seconds of silence, synchronized with GMT. This can be used to detect many different problems using clock based monitors. It is important that the generators be (roughly) synchronized so this technique can be used with an entire suite of sources. A typical use might be to emulate up to two video and two audio sources for an IETF, and correlating network events, such as device errors and ethernet collision rates, against the clock to determine if the mbone is causing a particular problem.

Applications which do not use AVT should have their own set of the above tools.

It is imperative that the majority of these tools run on all mrouted capable platforms. It is currently very difficult to sectionalize distribution problems on paths through multiple mrouted tunnels that are incapable of running the specific applications.

Mrouted/tunnel Implementation Features

- Encapsulated tunnels (already done).

- Support for "on demand" host tunnels, such that mrouted can be used as a reflector for hosts that don't support multicast, such as Mac's. (This isn't really high priority, but it is relatively easy to implement.)

- Implement pruning. Currently all multicast traffic goes to all tunnel connected nodes within the TTL limit, even if there are no listeners. Note that TTL mechanism is still needed because pruning is not fast enough to protect slow links.

- Implement aggregate packet rate and bandwidth limits, to protect the underlying IP infrastructure from being flooded.

- The routing protocol, DVMRP, should use the same tunnels as the data to prevent the situations we see today where the unicast routing can follow some path, but the data can not. DVMRP sees the tunnel as up but all of the data is discarded.

- Support LSRR on encapsulated tunnels, so tunnels can be anchored to specific IP infrastructure.

- Correct/work around the BSD bug which prevents DVMRP from asserting the source address of a tunnel on a multi-interfaced host. This causes IP routing to break tunnels when they move to other interfaces because the other end does not recognize the source IP address. [A possible solution, which also partially addresses the tunnel anchors, would be to specify the first hop address and have mrouted install static host routes for tunnels.]

- Mrouted support for traceroute and mapping infrastructure tools.

- There needs to be a logging facility/level for monitoring DVMRP routing problems. Such a facility would report tunnel up/down events and routing table changes.

Steve Deering was present and both contributed and noted our comments.

The wish lists were presented to both the AVT Working Group and to the Operational Requirements Area Directorate (ORAD). There was general agreement that the items on the wish list are desirable and appropriate, but nobody agreed to implement anything.

There were some network operators present at the ORAD meeting who were upset that the MBONE BOF did not become a working group or take stronger positions regarding operational practices. However most of these people were operators who had chosen not to participate in the mbone, and were therefore not in control of its impact on their own facilities.

Thanks to Jamshid Mahdavi for taking the Minutes. It should be noted that the attendee list below was reconstructed after the fact and may not contain the names of all participants.

Attendees

Erik-Jan Bos	erik-jan.bos@surfnet.nl
Douglas Carson	carson@utcc.utoronto.ca
Henry Clark	henryc@oar.net
Steve Deering	deering@parc.xerox.com
Dale Finkelson	dmf@westie.mid.net
Eugene Hastings	hastings@psc.edu
Daniel Long	long@nic.near.net
Jamshid Mahdavi	Mahdavi@a.psi.edu
Bill Manning	bmanning@sesqui.net
Matt Mathis	mathis@a.psc.edu
David O'Leary	doleary@cisco.com
Curtis Villamizar	curtis@ans.net
Linda Winkler	lwinkler@anl.gov
Paul Zawada	Zawada@ncsa.uiuc.edu

CURRENT MEETING REPORT

Reported by Bernhard Stockman/EBONE

Minutes for the Operational Area Directorate (ORAD)

Agenda

- Operations Area Mandates
  - Operational Standards
  - Operational Requirements
  - Operation Coordination
  - Review of Standards

- Coordination Issues
  - Routing Policy Language
  - Routing Policy Registration

- Coordination of Tunnels
  - MBONE
  - TBONE
  - *BONE

- Operational Impact if IPng

- Class C Allocation Forecasts

Side notes

Do we need a single Working Group to handle issues regarding "virtual Internets" sitting on top of today's IP internet?

How will the large blocks of Class C Addresses allocated affect current routing platforms?

Operational Area Mandates

Coordinations of network service providers; is it the job of the IETF? "I can see no less biased venue." - Gene Hastings

Historically the IETF has been a venue for providers to discuss the leading edge technology they have been deploying and to give developers feedback.

It seems as if providers are swamped by the rapid growth of the Internet. This forum allows providers to work on problems "at the top of the list." As the Internet continues to grow, more "fire-fighting" resources may become available. At that time some other forum for NSP coordination may be appropriate (i.e., someone else may be willing to host such a forum).

Is it good to meet at the IETF? A great deal of technology transfer takes place. If a separate operational venue is developed, it should collocate with the IETF. Another separate set of meetings would be poorly attended as network operators have too many meetings to attend as it is.

As for reading for RFCs, if the Operational Directorate read RFCs, would that alleviate some of the problems we have now? Many RFCs have fallen through the cracks towards complete implementation without operational issues being addressed (e.g., RFC1400). Informational RFCs need to be addressed as they don't follow a standards track.

What about "executive cooperation?" A lot of informal agreements are made at the IETF meetings, but they can't be backed without network higher-ups agreeing to it. Thus the purpose of the meeting should be to coordinate operation of individual networks, not to change each network's own policy.

A certain amount of consensus is built up in the operational area. The operational area seems to have fuzzy objectives and less concrete standards.

NSP coordination actually also takes place within other organizations, e.g., FARNET.

We'll continue to do things the way we are. If another forum develops, it would be wise to collocate it with the IETF meetings.

One thing that ORAD needs to address is a mechanism to apply an "ORAD seal of approval." There needs to be a mechanism to alert ORAD to [informational] RFCs which have significant operational impact.

John Curran agreed to make sure that the job of reading at least a fraction of new RFCs for operational impact is handled. Bill Manning agreed to do some reviewing, but cannot do the whole job himself.

There is a need for a working group to deal with a policy routing description language. Many of the routing efforts (BRG, SDRIP, etc.,) are defining a need for a common routing policy language. ORAD needs to form a liaison with the protocol developers to help define such a language.

It is possible that the Internet Working Group should be revived to deal with topology configuration. Such a group could form the liaison with the routing protocol developers to give input as to how develop future protocols.

Dan Long's proposed structure (areas to be addressed):

- BGP Deployment - protocol/CIDR

- IWG - NSP Routing Coordination
- NSR - Growth
- Tunnel OPS - Coordinate the deployment of things like MBONE, etc.
- OPSTATS
- NOOP
- ORAD
- UCP (???) - Gas running out - Possible to roll into NJM
- Benchmarking - Bradner
- NJM (new) - coordination of NSP, trading troubleshooting techniques, operational experience

Three different types of working groups:

1. Exchanging Information.
2. Coordinated/establishing operational procedures.
3. Engineering standards (e.g., benchmarking).
4. IWG/Tunnel Ops - Operational planning.
5. NJM/UCP/NOOP - Diagnostic procedures.

Tunnel Coordination

There was no critical mass to form an MBONE engineering group. However, a number of different tunnel types may need to be organized to keep the collection of "BONES" from melting down the Internet.

MBONE seems to be causing operational problems. It is quite possible that this is the first of many other tunnel types which will appear again.

Matt Mathis reported the happenings at the MBONE BOF.

- General issues
- Procedures
- Topology, metric, and threshold engineering
- Bandwidth budget and policy
- Infrastructure oriented tools (wish list)
- Transport and application diagnostic tools (wish list)

MBONE is only an example of very steep growth rates. It is hard to tell end users not to use their connection whether it be MBONE, AFS, Internet Talk Radio, etc.

At this point there is no need to address these issues with a working group, however it would probably be wise to hold some sort of meeting before or at the next IETF. Discussion will be held on the ORAD list to organize such a meeting.

Class C Allocation Forecasts

Will CIDR save us in time to keep our routers from falling apart from too many routes? The number of Class C nets being allocated is quite high. The NSFNET routing table has grown at a rate of 6.5% per month. We'll probably see 10,000 routes by July. Is deploying CIDR going to save us?

Which networks will hit the limit before CIDR is deployed? ANS feels their situation is under control, but other service providers may feel the crunch.

Controlled deaggregation may need to be dealt with for those providers who can't speak BGP4 but can't default. Ground rules need to be laid to define this policy. A mechanism also needs to be defined for negotiating the amount of aggregation which takes place between two networks.

It is not clear when things will fall apart, so the problems may just have to be dealt with when it occurs. All that can be done is to keep trying to get CIDR deployed in time and try to beat the clock.

ANS is also performing tests with cisco and IBM routers to see how many routes can be flapped in and out before they suffer server performance degradation.

Attendees

Tony Bates	tony@ripe.net
Serpil Bayraktar	sbb@noc.ans.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Douglas Carson	carson@utcc.utoronto.ca
Henry Clark	henryc@oar.net
David Conrad	davidc@iij.ad.jp
John Curran	jcurran@nic.near.net
Tom Easterday	tom@cic.net
Dale Finkelson	dmf@westie.mid.net
Francois Fluckiger	fluckiger@vxcern.cern.ch
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Ittai Hershman	ittai@ans.net
Steven Hubert	hubert@cac.washington.edu
Dale Johnson	dsj@merit.edu
Merike Kaeo	merike@alw.nih.gov
Daniel Karrenberg	daniel@ripe.net
Charley Kline	cvk@uiuc.edu
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se

Glenn Mackintosh	glenn@canet.ca
Bill Manning	bmanning@sesqui.net
Matt Mathis	mathis@a.psc.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
Peder Chr. Noergaard	pcn@tbit.dk
David O'Leary	doleary@cisco.com
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@sura.net
Kim Smith	kas@noc.ans.net
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Kaj Tesink	kaj@cc.bellcore.com
Claudio Topolcic	topolcic@cnri.reston.va.us
Curtis Villamizar	curtis@ans.net
Ruediger Volk	rv@informatik.uni-dortmund.de
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net
Paul Zawada	Zawada@ncsa.uiuc.edu

2.4.1 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: `merit.edu:~/pub/bgpd-archive`

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 generated on this topic.
Done	Post as an Internet-Draft, a report of BGP deployment status.
Done	Post an Internet-Draft, defining a mechanism to share policy information between Administrative Domains.

Internet-Drafts:

"Notes of BGP-4/CIDR Coordination Meeting of 11 March 93", 03/24/1993, C. Topolcic <draft-ietf-bgpdepl-minutes-93feb-00.txt>

"Aggregation Support in the NSFNET Policy Routing Database", 03/26/1993, M. Knopper <draft-nsfnet-aggregation-00.txt>

CURRENT MEETING REPORT

Reported by Matt Mathis/PSC

Minutes of the BGP Deployment Working Group (BGPDEPL)

Administrivia

A question of venue was discussed, but not settled. A hand vote indicated the majority of those present were planning to attend the Amsterdam meeting. However, several of the key players would be unable to attend. There was also a question about whether an additional meeting was needed before the next IETF. This question was deferred pending organizational changes.

Later during the meeting it was observed that most of the configuration discussion was not really BGP related, but more apropos of the original "Internet Working Group (IWG)", which was tasked with fostering sanity in topology, routing policy, and configuration databases. It is interesting to note that the original BGP1 arose out of the IWG.

It was suggested that the IWG be reconstituted, and that the BGP Deployment Working Group be folded in as one of its key tasks.

Vendor Reports

- ANS/NSFnet/GATED.

  Dennis Ferguson reported that BGP4 is working in a test mode. He also reported that new IGP code is under development. This new code is needed to interoperate with the existing routed code in the backbone.

- cisco.

  Paul Traina indicated that BGP4 is still under development. The development effort has been hit with some pretty significant delays and is going to be late. BGP4 was not approved for inclusion into 9.21, so there will be a special software release based upon 9.21 with BGP4 support added. This release will be available for testing and limited deployment before 9.21 has completed beta cycle. The BGP4 special release should be ready for general availability near 9.21 FCS (no date available). [This is an updated report to make it more accurate (for the worse).]

- 3COM.

  Nagaraj Arunkumar expects to support BGP4 in release 6.2 due sometime early this fall.

- Wellfleet.

  John Krawczyk anticipates rolling out BGP4 support this summer.

- Telebit Communications.

  Peder Chr. Noergaard reports that EUROPAnet is in the process of deploying BGP3 with no current plans for BGP4 deployment.

CIDR core plans (Alternet, CIX, EBone, NSFnet/ANS, NSI, PSI, Sprint)

As there appeared to be a critical mass of people interested in the BGP4 deployment who were also attending DC INTEROP, Claudio Topolcic convened a meeting to update plans. The Minutes for this meeting are available in the usual IETF directories as draft-ietf.bgpdepl-minutes-feb93-00.txt (even though the meeting took place in March!). That meeting was also summarized for the Group with clarifications and expansions.

There was quite a bit of discussion concerning the deployment plan. It now looks like this:

1. Start deploying BGP4 code as soon as possible. It now appears that this may be delayed to as late as June. The goal here is just to verify that the code works. Exercise no new features of the protocol in the production Internet.

2. NSI (or some alternative) starts announcing one aggregated network. This step has been split into two pieces: 2a) Initially announce an aggregated test network (assigned but non-production). After verification that it is propagating properly to the rest of the core, 2b) aggregate ONE site (several production class C networks), and verify correct interoperability with the rest of the Internet.

3. Additional CIDR core members start aggregating networks, first with test network and then with one production site each. Steps 2 and 3 can be partially overlapped as long as there are no adverse side effects of the announced aggregated nets, and that the selected aggregated sites can make arrangements to reach any portion of the Internet not yet supporting CIDR.

4. Aggregation is officially "turned on" in the Internet. This is a pseudo flag day because all sites requiring full routing tables must be either running BGP4 or must have made alternative arrangements (e.g., default routing). Aggregation should be phased in incrementally (a few sites at a time) and continue to be restricted to the site level (aggregate only multiple class C networks at one site). Aggregation of larger blocks of networks requires better solutions to some configuration management issues. Particularly mixed traffic types, etc., (e.g., AUP/non-AUP, multi-homed sites).

5. Think....

At this point there was a discussion of a number of side issues. Tony Hain of ESnet indicated that he was not sure what would happen in the early phases, he would not be ready to support BGP4 by June. ESnet may have to use default routing to survive.

Paul Traina of cisco mentioned the possibility of adding something to statically manage "controlled de-aggregation" using access lists and last resort approach to CIDR. Dennis Ferguson quiped that he would probably add something to gated in that case since "gated should be able to do anything the cisco can do".

It was generally agreed that another meeting is needed before step 4. It is unclear at this time if there would be contention and a real flag day. Hopefully all seven members of the CIDR core would either be fully BGP4 or have made peace with default routing well in advance of step 4, such that its precise timing is unimportant. If there is a contentious straggler, the community will eventually be forced to choose a flag day over their protests.

The Group decided not to attempt to place precise dates on the schedule. The members of the CIDR core are progressing as fast as possible and are well coordinated among themselves. The schedule has already slipped by about two months from projections made at the DC IETF (in November).

The next Regional-Techs meeting was mentioned as a possibility for another BGP deployment Working Group meeting. It is tentatively scheduled for May or June in Washington, DC. Matt felt that this would be a little too early.

CIDR Configuration Issues

There is some controversy over how to do global configuration checking. In addition, how can we ensure topology matches policy?

Mark Knopper presented a preliminary plan for aggregation support in the NSFnet which would:

1. Accept aggregate routes from a midlevel, or
2. Accept site routes from a midlevel, and aggregate on the midlevels behalf.

A strawman database format had netprefix and length, source (aggregating router) AS, and a destination AS list as components. This proposal is documented in the "Inter-domain Routing Policy Description and Sharing" Internet-Draft written by Yun, Yu, Chen and Rekhter (draft-yu-rpd.00.txt). This presentation resulted in a suggestion to split the single view into two views keyed on source AS. There was quite a bit of discussion on where and how to split this to best support debugging connectivity problems.

Matt Mathis argued fairly strongly for aggregation being controlled by the site owning the networks. The argument is that configuration control is far easier to manage if it is local. Sue Hares felt fairly strongly that central management is better. Matt did concur that

Merit will have the capability of aggregating routes on behalf of regionals who cannot, but would like to.

The representatives from RIPE pointed out that the existing U.S. databases, including the current Merit configuration database and the above proposal are not adequate to solve international routing problems. In particular none can be used to determine which backbone (CIDR core member) is the preferred path to a given U.S. network. Consider for a moment several sites with external connectivity to both NSFnet and PSI. Each site may prefer one or the other for various reasons including differing bandwidth, AUP, etc. This is further complicated because the ANS AS path does not reveal if the connection is "blessed" for non-NSF AUP traffic. Ideally the traffic from Europe could be routed solely on the basis of AS path but essential information is missing. Alternatively there should be a way to glean from our configuration databases which backbone the site prefers, but again there is not.

Global Configuration Issues

Daniel Karrenberg presented the RIPE efforts in the Global configuration database area. He indicated the real focus of this effort was to provide a tool their operators could use. This database also contains enough information to allow someone to compile suitable router network configuration files. It is documented as "Representation of IP Routing Policies in the RIPE Database", and is available from the RIPE repositories: ftp.ripe.net:ripe/docs/ripe-docs/ripe-81.[txt,ps].

Things the RIPE effort cannot do include an inability to process and propagate policies information on transit networks. It cannot use unpublished AS's, and it does nothing to solve the "half baked" AS problem, outside of pointing out inconsistencies.

Closing Remarks

The sense of urgency came from concerns about configuration management and database issues. Although there is still a lot of work to be done to complete the BGP4 roll out, it seems to be a fairly well understood problem except for configuration management. CIDR and BGP4 do impose some new requirements on the databases but the majority of the issues center around topology and AUP enforcement. For these reasons it makes sense to broaden the scope of this Working Group from just BGP deployment to the wider task of fostering sanity in topology, routing policy, and configuration databases.

Thanks to Robert Reschly and Gene Hastings for taking meeting notes.

Attendees

Nagaraj Arunkumar	`nak@3com.com`
William Barns	`barns@gateway.mitre.org`
Tony Bates	`tony@ripe.net`
Jordan Becker	`becker@ans.net`

David Bolen	db3l@ans.net
Rebecca Bostwick	bostwick@es.net
Jeffrey Burgan	jeff@nsipo.nasa.gov
Robert Calderon	calderon@noc.ans.net
Henry Clark	henryc@oar.net
Robert Collet	rcollet@icm1.icp.net
Tom Easterday	tom@cic.net
Stefan Fassbender	stf@easi.net
Mark Fedor	fedor@psi.com
Dennis Ferguson	dennis@ans.net
Vince Fuller	vaf@stanford.edu
Tony Hain	alh@es.net
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Eugene Hastings	hastings@psc.edu
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Ittai Hershman	ittai@ans.net
Jeffrey Honig	Jeffrey_C_Honig@Cornell.edu
Laurent Joncheray	lpj@merit.edu
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
Daniel Karrenberg	daniel@ripe.net
John Krawczyk	jkrawczy@wellfleet.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Frank Liu	fcliu@pacbell.com
Daniel Long	long@nic.near.net
Peter Lothberg	roll@stupi.se
Glenn Mackintosh	glenn@canet.ca
Jamshid Mahdavi	Mahdavi@a.psi.edu
Bill Manning	bmanning@sesqui.net
Jun Matsukata	jm@eng.isas.ac.jp
Daniel McRobb	dwm@noc.ans.net
Dennis Morris	morrisd@imo-uvax.disa.mil
Peder Chr. Noergaard	pcn@tbit.dk
David O'Leary	doleary@cisco.com
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@sura.net
Michael Patton	map@bbn.com
Willi Porten	porten@gmd.de
Selina Priestley	sfp@noc.ans.net
Yakov Rekhter	yakov@watson.ibm.com
Robert Reschly	reschly@brl.mil
Tony Richards	richards@icm1.icp.net
Vilson Sarto	vilson@fapq.fapesp.br
John Scudder	jgs@merit.edu
Paul Serice	serice@cos.com

Kim Smith	kas@noc.ans.net
Bernhard Stockman	boss@ebone.net
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
John Tavs	tavs@vnet.ibm.com
Marten Terpstra	marten@ripe.net
Paul Traina	pst@cisco.com
Curtis Villamizar	curtis@ans.net
Jack Waters	waters@sura.net
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net
Paul Zawada	Zawada@ncsa.uiuc.edu

2.4.2 Benchmarking Methodology (bmwg)

Charter

Chair(s):
Scott Bradner, sob@harvard.edu

Mailing Lists:
General Discussion: bmwg@harvard.edu
To Subscribe: bmwg-request@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:

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

Done The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.

Done Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Request For Comments:

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices"

2.4.3 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

Minutes of the Network Joint Management Working Group (NJM)

Report not submitted. The attendees listed below were participants in the joint NJM/NETSTAT sessions.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Serpil Bayraktar	sbb@noc.ans.net
Jordan Becker	becker@ans.net
B. Sue Blair	bsb@noc.ans.net
David Bolen	db31@ans.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Deborah Boyer	djb@noc.ans.net
Jeffrey Burgan	jeff@nsipo.nasa.gov
James Cassell	jcassell@dsac.dla.mil
Enke Chen	enke@merit.edu
Thomas Christie	tfc@aa.ans.net
David Conklin	conklin@jvnc.net
David Conrad	davidc@iij.ad.jp
John Curran	jcurran@nic.near.net
Chas DiFatta	chas@cmu.edu
Tom Easterday	tom@cic.net
Robert Enger	enger@reston.ans.net
Hans Eriksson	hans@sics.se
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Dale Finkelson	dmf@westie.mid.net
Peter Ford	peter@goshawk.lanl.gov
Christine Fredenburg	cfredenburg@dsac.dla.mil
Vince Fuller	vaf@stanford.edu
Kenneth Goodwin	goodwin@a.psc.edu
Robert Gutierrez	gutierre@nsipo.nasa.gov
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Steven Hubert	hubert@cac.washington.edu
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Mark Knopper	mak@merit.edu

Padma Krishnaswamy	kri@sabre.bellcore.com
John Labbe	jcl@ans.net
Kim Long	klong@sura.net
Glenn Mackintosh	glenn@canet.ca
Bill Manning	bmanning@sesqui.net
Glenn Mansfield	glenn@aic.co.jp
Matt Mathis	mathis@a.psc.edu
Daniel McRobb	dwm@noc.ans.net
Milo Medin	medin@nsipo.nasa.gov
Dennis Morris	morrisd@imo-uvax.disa.mil
David O'Leary	doleary@cisco.com
Maryann Perez	perez@cmf.nrl.navy.mil
Marsha Perrott	mlp+@andrew.cmu.edu
Narendra Popat	albin@frontier.com
Willi Porten	porten@gmd.de
Edward Reed	eer@cinops.xerox.com
Robert Reschly	reschly@brl.mil
Steven Richardson	sjr@merit.edu
Timothy Salo	tjs@msc.edu
Vilson Sarto	vilson@fapq.fapesp.br
Tim Seaver	tas@concert.net
Charlie Smith	crsmith@osvi.edu
Kim Smith	kas@noc.ans.net
Ira Steckler	isteckle@chipcom.com
Bernhard Stockman	boss@ebone.net
Claudio Topolcic	topolcic@cnri.reston.va.us
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Kannan Varadhan	kannan@oar.net
Ruediger Volk	rv@informatik.uni-dortmund.de
Hung Vu	hungv@fonorola.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net
Jane Wojcik	jwojcik@bbn.com
Richard Woundy	rwoundy@vnet.ibm.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov
Paul Zawada	Zawada@ncsa.uiuc.edu

Overview of Merit's Policy-Based Routing Database (PRDB)

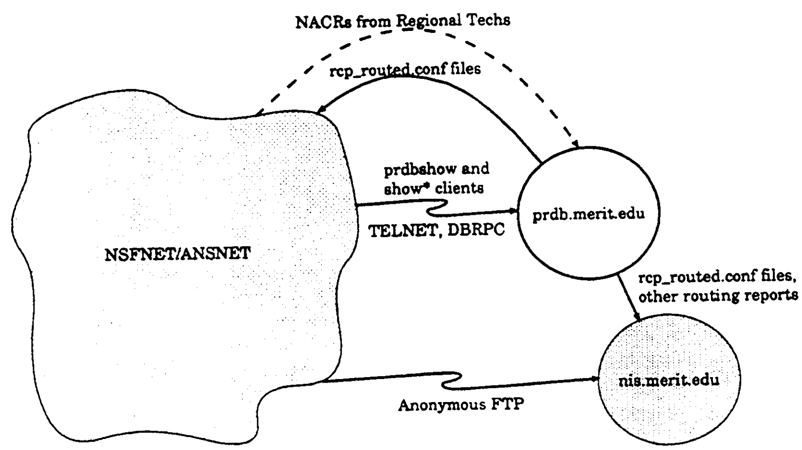

NJM meeting, 26th IETF

Prepared by Steven J. Richardson (sjr@merit.edu) 27 Mar 93 for Merit Network, Inc.

2.4.4 Network OSI Operations (noop)

Charter

Chair(s):

Susan Hares, skh@merit.edu
Cathy Wittbrodt, cjw@barrnet.net

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:

Ongoing	Provide a forum to discuss OSI routing plans by email or in group discussions.
Jan 1992	Post as an Internet-Draft, a tutorial for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.
Apr 1992	Post as an Internet-Draft, a requirements document specifying what OSI network tools are needed on every host and router.
Jul 1992	Post as an Internet-Draft, a collection of regional Routing and Addressing plans.
Done	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.

Done 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.

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.

Internet-Drafts:

"An Echo Function for ISO 8473", 11/10/1992, S. Hares, C. Wittbrodt <draft-ietf-noop-echo-02.txt>

"Essential Tools for the OSI Internet", 11/10/1992, S. Hares, C. Wittbrodt <draft-ietf-noop-tools-03.txt>

CURRENT MEETING REPORT

Reported by Sue Hares/Merit

Minutes of Network OSI Operations Working Group (NOOP)

Agenda

- Transport MIB (Russ Blaesing)
- Tools RFC
- Echo RFC
- TUBA Presentation

The Network OSI Operations Working Group met once during the Columbus IETF. Russ Blaesing led a discussion of the Transport MIB. Suggestions included the following:

- Post the MIB to Venera.ISI.Edu under experimental.
- Post a short form MIB to the NOOP List.

Tools and Echo RFCs

After minor editing the Tools RFC and Echo RFC will go to the Area Director for forwarding to the "requirements" track.

Some questions/comments:

Host requirements specifying how the OSI hosts behave on the Internet should include the slow start from Van Jacobson. Sue Hares expressed her concern that this would slow down the basic format.

TUBA Presentation

Mark Knopper gave a presentation on TUBA which raised the following issues:

1. BSD on 486 or BSDI/BSD4.4. Has TUBA, TCP, telnet, and does new ping
2. NCSA telnet machine. No packet driver. Packet on driver looks at OSI packets and dumps them.
3. CLNPDUMP on SUNOS.
4. Cisco does TUBA for UDP, telnet, TFTP.

A problem with TUBA is that there is no end host vendor buy in. An EON configuration needs to be posted. Sue Hares will make it available via anonymous ftp on merit.edu under /pub/noop/config. She will also try to improve the daily CLNP service in the Internet infrastructure by creating host.db, bis.db and is.db files.

The Group will clnpping all hosts or routers listed in the host.db, bis.db, and is.db.

A concern was expressed that the infrastructure was only up for demos. People will be solicited for weekly pings. There is a catch-22 that you must have a working infrastructure in order to have people try out a TUBA infrastructure.

Attendees

William Barns	barns@gateway.mitre.org
Russell Blaesing	rrb@one.com
Rebecca Bostwick	bostwick@es.net
Ross Callon	rcallon@wellfleet.com
Richard Colella	colella@nist.gov
Tom Easterday	tom@cic.net
Dale Finkelson	dmf@westie.mid.net
Peter Ford	peter@goshawk.lanl.gov
Susan Hares	skh@merit.edu
John Heimaster	jwh@mps.ohio-state.edu
Matthew Jonson	jonson@server.af.mil
Mark Knopper	mak@merit.edu
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Bill Manning	bmanning@sesqui.net
Daniel McRobb	dwm@noc.ans.net
Dennis Morris	morrisd@imo-uvax.disa.mil
Peder Chr. Noergaard	pcn@tbit.dk
Andrew Partan	asp@uunet.uu.net
Laura Pate	pate@gateway.mitre.org
David Piscitello	dave@mail.bellcore.com
Willi Porten	porten@gmd.de
Selina Priestley	sfp@noc.ans.net
April Richstein	amr@tycho.ncsc.mil
Vilson Sarto	vilson@fapq.fapesp.br
John Scudder	jgs@merit.edu
Craig Smith	smith@ohio.gov
Bernhard Stockman	boss@ebone.net
Catherine Summers	cfs@cos.com
Richard Thomas	rjthomas@bnr.ca
Louisa Thomson	louisa@whitney.hac.com
Jack Waters	waters@sura.net
Cathy Wittbrodt	cjw@barrnet.net
Paul Zawada	Zawada@ncsa.uiuc.edu

2.4.5 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, boss@ebone.net
Phillip Gross, pgross@ans.net

Mailing Lists:

General Discussion: oswg-l@wugate.wustl.edu
To Subscribe: oswg-l-request@wugate.wustl.edu
Archive: wuarchive.wustl.edu:~doc/mailing-lists/oswg-l

Description of Working Group:

Today there exists 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.

Done Identify similar efforts being performed by other groups.

Done Define a common minimal set of metrics.

Done 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.

Done Develop outline, and make writing assignments for paper (Opstat1) documenting March 1991 milestones.

Done Complete paper Opstat1.

Done	Possible mid-term meeting to review Opstat1.
Done	Submit Opstat1 as Internet-Draft.
Done	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Done	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Done	Propose old tools to be retrofitted.
Done	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools.
Sep 1992	Submit Internet-Draft of new and retrofitted tools.
Dec 1992	Submit new and old tools Internet-Draft to IESG as an Informational RFC.
Mar 1993	Post an Internet-Draft defining the client/Server Opstat protcool.
Mar 1993	Post the Opstat Statistical MIB as an Internet-Draft.
Jul 1993	Submit the Client/Server Opstat Protocol to the IESG for consideration as a Proposed Standard.
Jul 1993	Submit the Statistical Opstat MIB to the IESG for consideration as a Proposed Standard.

Request For Comments:

RFC 1404 "A Model for Common Operational Statistics"

CURRENT MEETING REPORT

Reported by Bernhard Stockman/EBONE

Minutes of the Operational Statistics Working Group (OPSTAT)

Agenda

- RFC1404 Deployment Experience
- Client/Server Protocol: Presentation by ANS
- Statistical MIB

RFC1404 Deployment Experience

Daniel Karrenberg reported that the efforts to convert MONSTER to use RFC1404 have not progressed. Due to resource constraints further development of MONSTER has been put on indefinite hold. What would be needed to make it use RFC1404 is just a relatively simple reformatter. Further it could need a rewrite of the actual SNMP polling function since this uses the snmp capable awk from ISODE which is very hard to install. Daniel renewed his offer to get started with those people who are interested in continuing work on MONSTER.

Craig Haney who was not present at the Working Group meeting has a PERL parser for the RFC1404 format.

A list of problems were reported with RFC1404:

- In the syntax for device-fields multiple tag-tables are not delimited syntactically. Proposed solutions:
  - Count preceding
  - Different field separator withing tag-table
  - Bracketing tag-table with something BEGIN-TAGT, END-TAGT
  - Bracketing with {}

The Group decided to use {}s around tag-table. This will lead to a revision of the BNF:

```
device-section ::= "BEGIN_DEVICE" <FS>
                   <device-field> <FS>
                   "END_DEVICE"

device-field   ::= <networkname><FS><routername><FS><linkname><FS>
                   <bw-value><FS><bw-sort><FS><proto-type><FS>
                   <proto-addr><FS><time-zone><FS> "{" <tag-table> "}"
                   [ "{" <tag-table> "}" ]
```

Action: Bernhard volunteered to check for similar problems and report via mail.

After the discussion the Group decided to collect more problem reports from implementors before trying to issue another RFC.

There is interest in using tools following the RFC out there. There is not much interest in actually producing the tools. FARNET might produce some funding for developing PD tools. Question: Who would be interested? No immediate interest.

Action: Bernhard will send a message to the list about this and try to find capable implementors. Any interested persons should contact Bernhard.

Client/Server Protocol

The intention is to make (aggregated) statistical information available from a server. A strawman document is being developed. The following reports were given:

- Henry Clark reported on implementation experience. Login is primitive direct string comparison. Separate configuration file defining the view of the users. Server runs use PSI formant but don't use opstat.

- Ittai Herhman reported on a similar system currently being used at ANS which originated at Merit. Refer to slides.

- Vikas Aggarwal also showed some examples of reports. Both Vikas and Ittai remarked that from limited experiments they have done it seems that for line utilisation there is little difference between fifteen minute and one minute sampling. A preliminary document is being done.

The Group went into architectural brainstorm mode and reached the following conclusions:

- There is agreement to generalise the authorisation passed by login to generic authentication strings and leave in the authentication method used as a local matter beyond the scope.

- Select shouldn't return the data, there should be a separate "get" command to initiate that, multiple selections are possible and identified by a selection ID which can be specified by client and will be returned by server.

- Status command to find out status of what is currently selected especially the amount of data.

Action: Bernhard Stockman will change the strawman accordingly.

Statistical MIB

Consensus is that a short list of highly desirable variables should be agreed upon and forwarded to the Working Group for inclusion in the next Internet Standard MIB.

Action: Bernhard will forward a list, of all variables that have been proposed so far, to the mailing list and re-start discussion.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
David Bolen	db31@ans.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Douglas Carson	carson@utcc.utoronto.ca
James Cassell	jcassell@dsac.dla.mil
Henry Clark	henryc@oar.net
David Conrad	davidc@iij.ad.jp
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Dale Finkelson	dmf@westie.mid.net
Eric Fleischman	ericf@act.boeing.com
Paul Franchois	paulf@bldrdoc.gov
Kenneth Goodwin	goodwin@a.psc.edu
Ittai Hershman	ittai@ans.net
Daniel Karrenberg	daniel@ripe.net
Kim Long	klong@sura.net
David Lovering	lovering@bldrdoc.gov
Daniel McRobb	dwm@noc.ans.net
Greg Minshall	minshall@wc.novell.com
Bob Morgan	morgan@networking.stanford.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Steven Richardson	sjr@merit.edu
Vilson Sarto	vilson@fapq.fapesp.br
Richard Woundy	rwoundy@vnet.ibm.com

IETF OpStat WG - March 1993

Advanced Network & Services
100 Clearbrook Road
Elmsford, New York 10523
914-789-5337
ittai@ans.net

Data Collection (Current)

- Centralized Approach
- SNMP:
  - Collected once per day
    - ifDescr, sysUpTime, ipAdEntIfIndex, ifSpeed, is-is-Index
  - Collected every 15 minutes
    - ifInOctets, ifOutOctets, ifInUcastPkts, ifOutUcastPkts, ifInErrors, ifOutErrors, ifOperStatus
- ARTS (ANSNet Router Traffic Statistics):
  - Statistical sampling of packet headers for protocol and src/dest data
- Delay Matrix:
  - ENSS-ENSS latency measurements (RS6000 µsec clocked 100 byte pings)

Data Collection (Current)

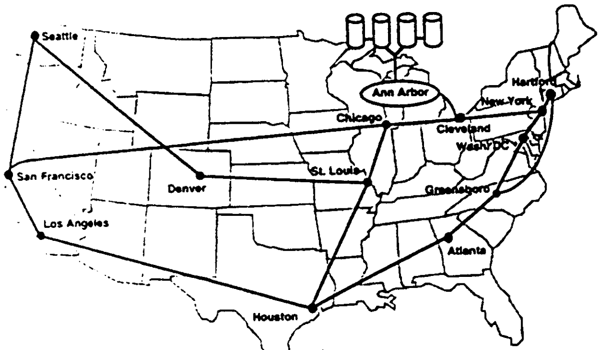

Data Collection (Strategic)

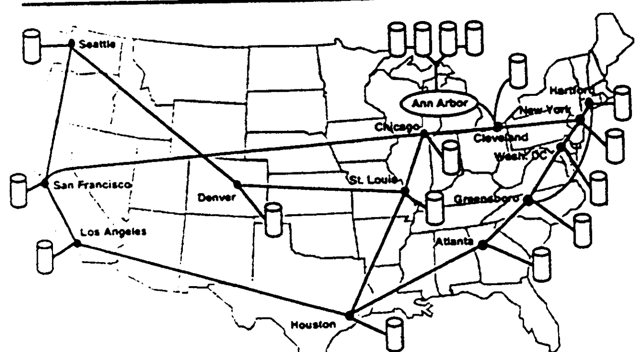

Data Collection (Strategic)

- IBM Data Engine (DRAGONS)
  - Distributed & Fault-Tolerant
  - Object Oriented
  - Persistent Store
- Approach
  - Poll hosts asynchronously using Data Engine thread model
    - Serial/Synchronous approach suffers when hosts become unreachable
    - Threads can be scheduled across CPU's
  - Minimize number of SNMP requests per host
    - Pack lots of variables into 1 SNMP GET-NEXT request
  - Store results efficiently

Data Storage (Strategic)

- Requirements
  - Efficient Storage (disk space)
  - Easy access (for usage reporting)
- Choices
  - Current: ASCII flat files
  - Encoded Data
    - Encode data structures in a machine independent format
  - Store delta's only
    - Save as ASCII, then compress
  - PersistentStore
    - Data Engine core object

Data Engine Overview

- Object Oriented
  - Data Engine "core" objects
  - User implemented objects
  - Multiple inheritance among objects supported
- Multi-Threaded
  - New thread created for each method invocation
- Distributed
  - Threads can be scheduled across CPU's
  - Objects can be replicated across Data Engines
- See paper in 1992 LISA VI Proceedings

Data Engine Objects

- Core Objects
  - SNMPhostinfo
  - PersistentStore
  - XGMON_Subscription
- User Implemented Objects
  - JobController
  - ANS_SNMP_PollNetwork
  - StatisticsCollector

Current Data Storage Architecture

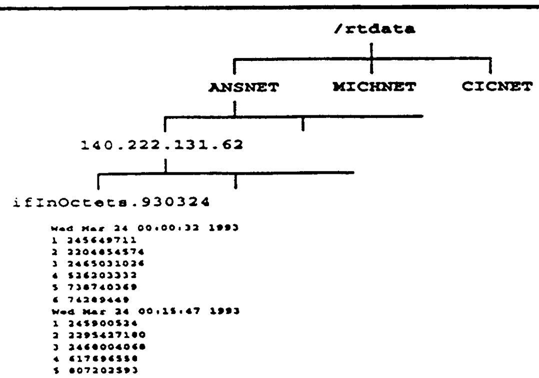

RTDATA Client/Server Interface

- Standard Inetd TCP-Socket server model
- Commands allow traversal of data tree, "object" selection, and whole file retrieval
  - e.g. ls, pwd, cd, head, get, select, streams, go, select <variable name> <begin time> <end time>
- Unix-oriented authentication (trusted hosts, rhosts, password)
- Data access controlled by managing group permissions

Sample Server Interaction

```
 Welcome to real-time dataserver.
 Version Thu Jan 24 21:09:40 CUT 1991

.ls
 ANSNET/
.
cd ANSNET/snmp/140.222.160.62
.
select ifInOctets 930324 9303240030
* begin: 732931200   midnight: 732931200
* stream: 01 new file pointer: 536877112
>01 /ANSNET/snmp/140.222.160.62/ifInOctets 732931200 732933000
.
go -t
T01  732931227 Wed Mar 24 00:00:27 1993
 01 1 12744173
 01 2 539010353
 01 3 639460290
T01  732932135 Wed Mar 24 00:15:35 1993
 01 1 13047323
 01 2 549078533
 01 3 717810163
<01 732933031
.
quit
```

Usage Reporting Pipeline

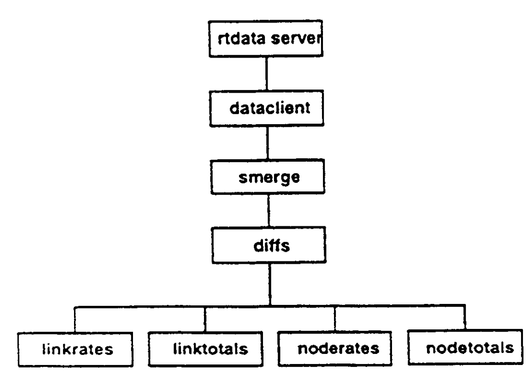

Usage Reporting Chronology

- Standard automated reporting of bandwidth utilization and offered load
- Report-time detection of gaps in data flow with data normalization
- Intelligent error handling for automated reports
- Generalized tool to produce ad-hoc reports: given a customer id, time interval and degree of granularity, produce a series of link utilization reports for that time frame
- Developed a tool that generates postscript graphs of utilization/time and utilization distribution (using S-Plus)

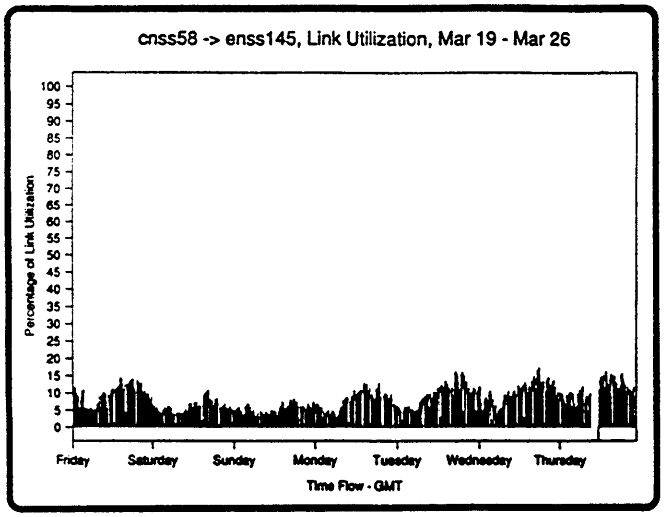
cnss58 -> enss145, Link Utilization, Mar 19 - Mar 26
Percentage of Link Utilization
Friday
Saturday
Sunday
Monday
Tuesday
Wednesday
Thursday
Time Flow - GMT

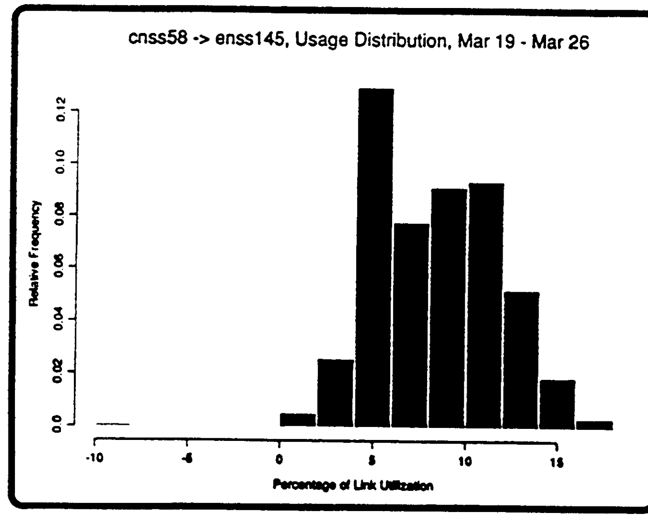
cnss58 -> enss145, Usage Distribution, Mar 19 - Mar 26
Relative Frequency
Percentage of Link Utilization

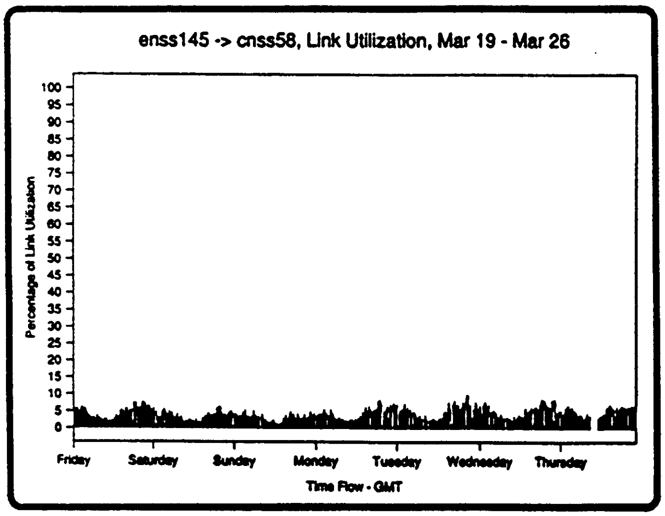
enss145 -> cnss58, Link Utilization, Mar 19 - Mar 26
Percentage of Link Utilization
Friday
Saturday
Sunday
Monday
Tuesday
Wednesday
Thursday
Time Flow - GMT

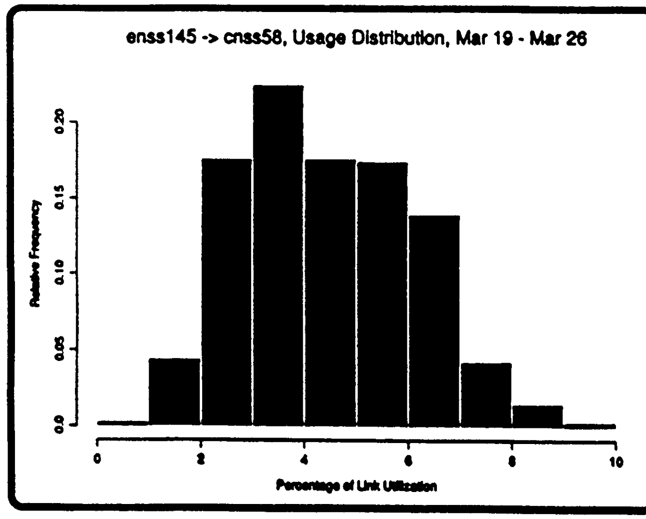
enss145 -> cnss58, Usage Distribution, Mar 19 - Mar 26
Relative Frequency
Percentage of Link Utilization

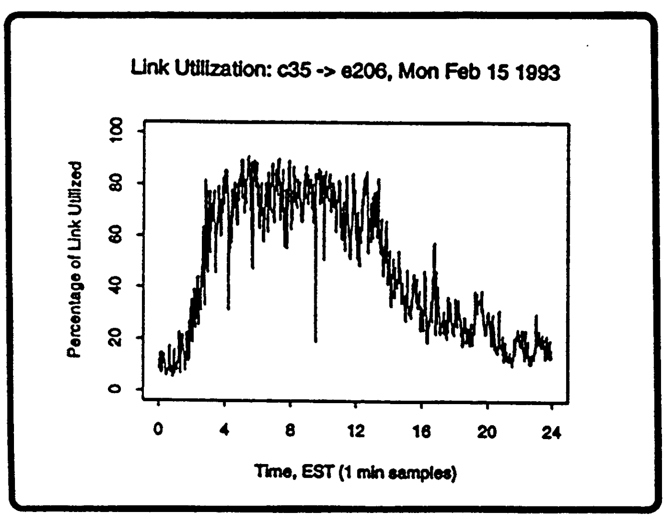
Link Utilization: c35 -> e206, Mon Feb 15 1993
Percentage of Link Utilized
Time, EST (1 min samples)

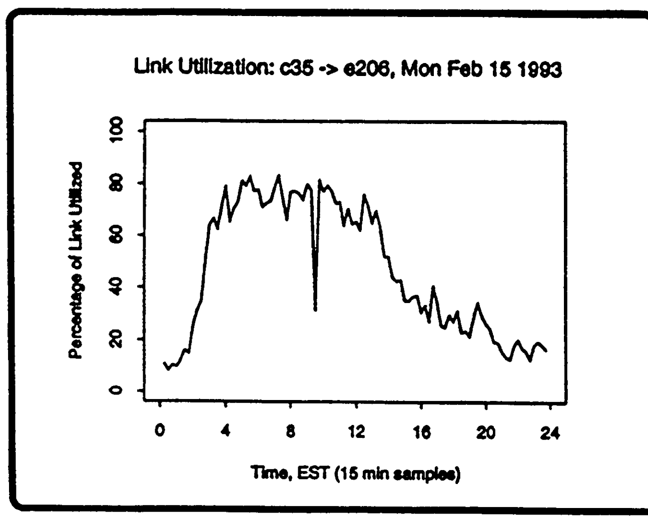
Link Utilization: c35 -> e206, Mon Feb 15 1993
Percentage of Link Utilized
Time, EST (15 min samples)

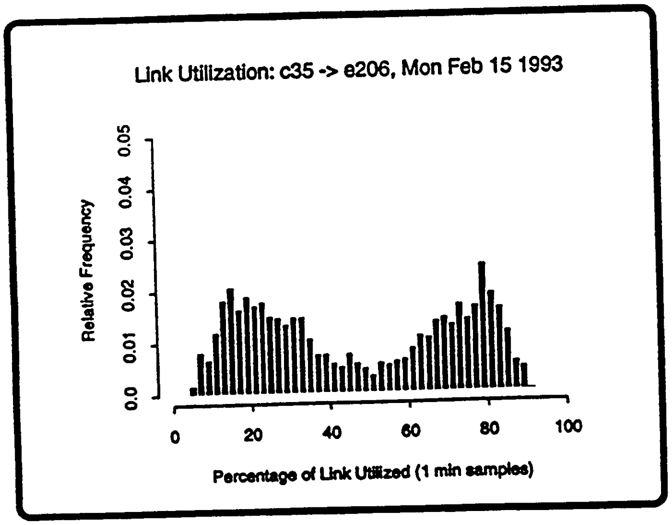
Link Utilization: c35 -> e206, Mon Feb 15 1993
Relative Frequency
Percentage of Link Utilized (1 min samples)

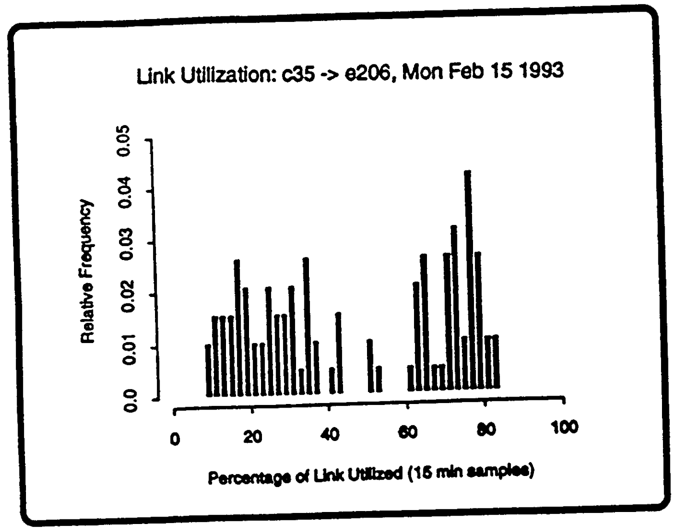
Link Utilization: c35 -> e206, Mon Feb 15 1993
Relative Frequency
Percentage of Link Utilized (15 min samples)

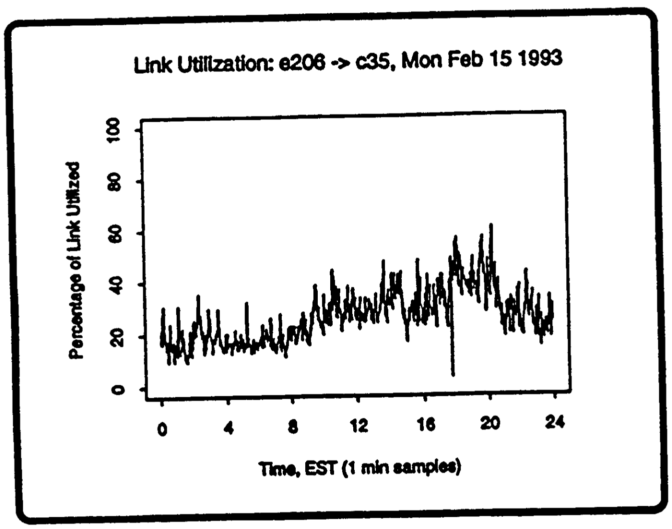
Link Utilization: e206 -> c35, Mon Feb 15 1993
Percentage of Link Utilized
Time, EST (1 min samples)

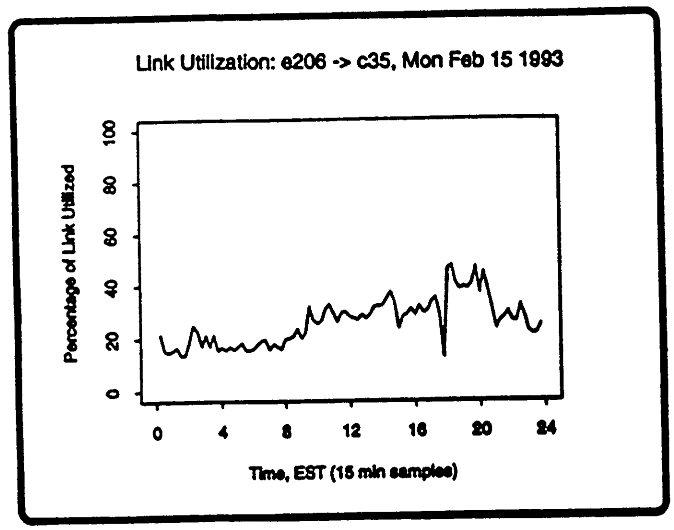
Link Utilization: e206 -> c35, Mon Feb 15 1993
Percentage of Link Utilized
Time, EST (15 min samples)

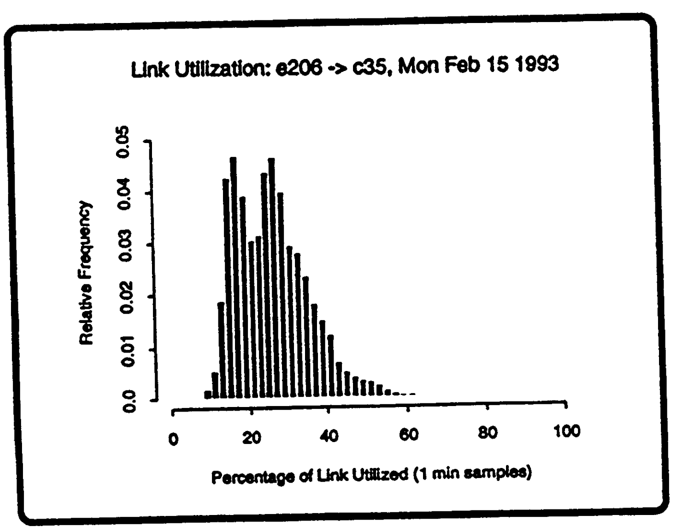
Link Utilization: e206 -> c35, Mon Feb 15 1993
Relative Frequency
Percentage of Link Utilized (1 min samples)

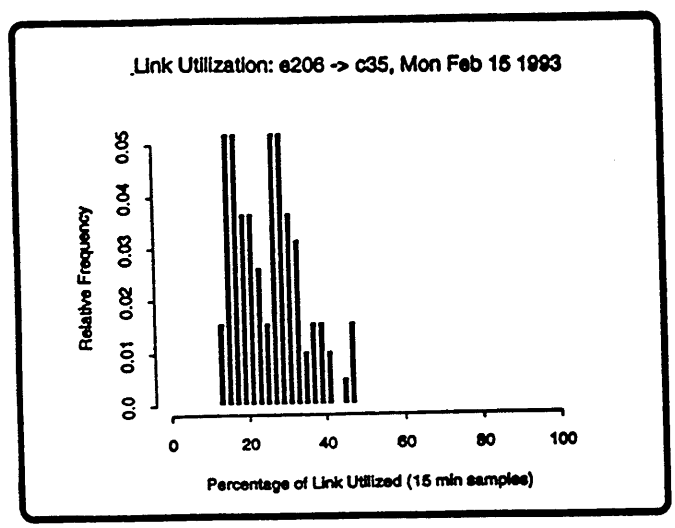
Link Utilization: e206 -> c35, Mon Feb 15 1993
Relative Frequency
Percentage of Link Utilized (15 min samples)

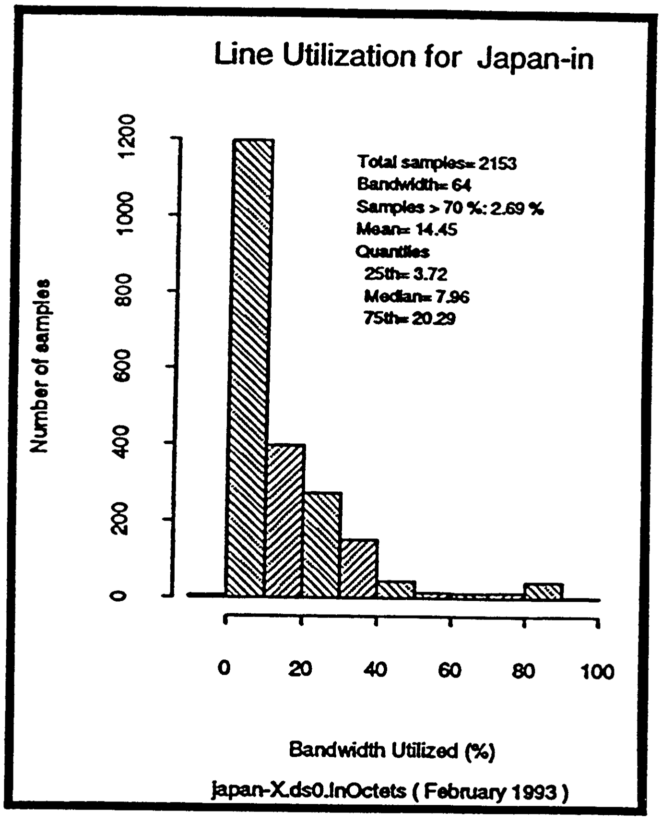
Line Utilization for Japan-in
Number of samples
0
200
400
600
800
1000
1200
Total samples= 2153
Bandwidth= 64
Samples > 70 %: 2.69 %
Mean= 14.45
Quantiles
25th= 3.72
Median= 7.96
75th= 20.29
0
20
40
60
80
100
Bandwidth Utilized (%)
japan-X.ds0.InOctets (February 1993)

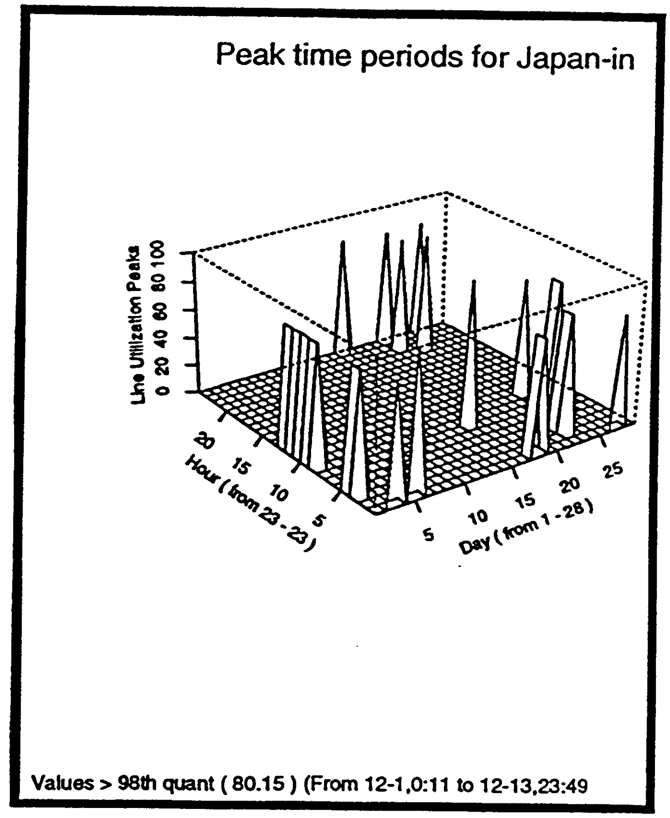
Peak time periods for Japan-in
Line Utilization Peaks
0 20 40 60 80 100
Hour (from 23 - 23)
Day (from 1 - 28)
Values > 98th quant (80.15) (From 12-1,0:11 to 12-13,23:49

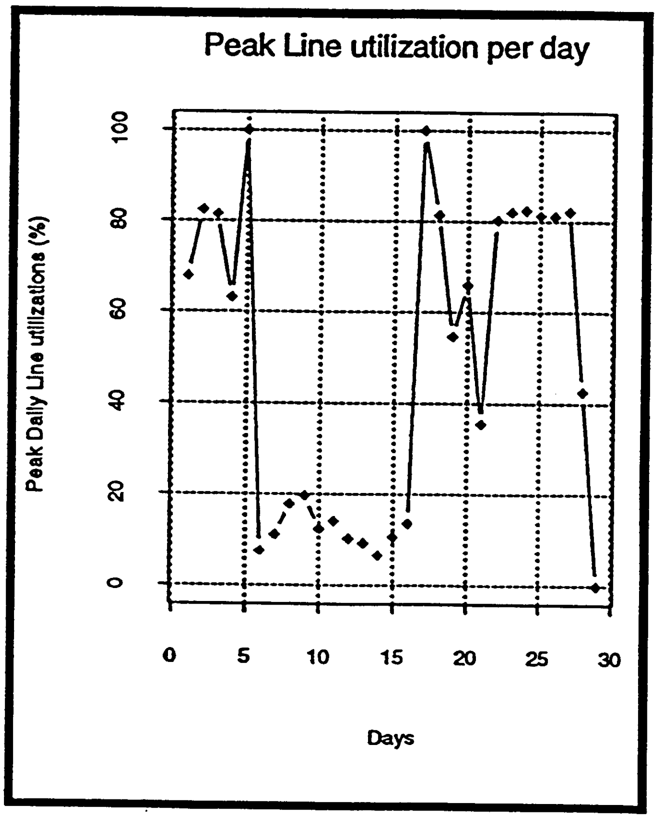
Peak Line utilization per day
Peak Daily Line utilizations (%)
0
20
40
60
80
100
0
5
10
15
20
25
30
Days

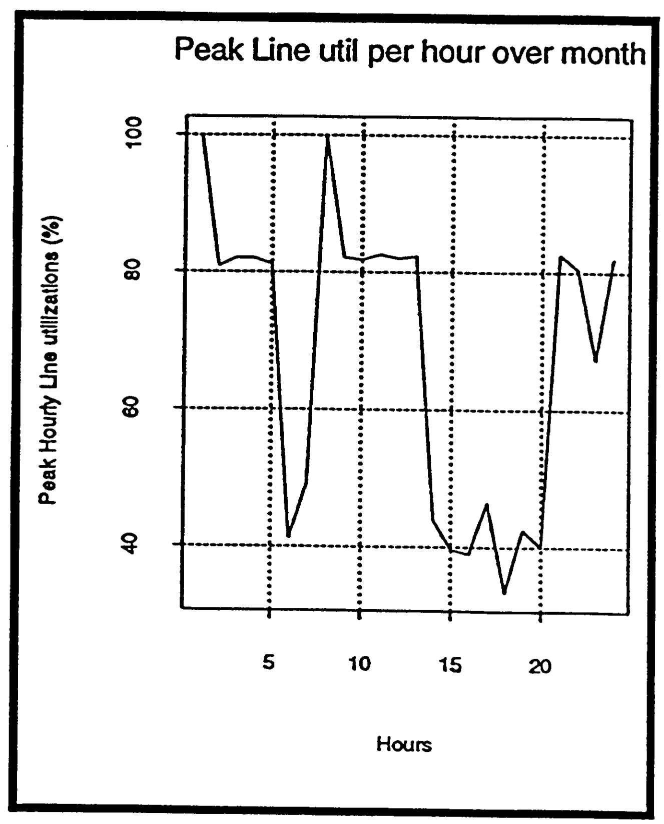
Peak Line util per hour over month
Peak Hourly Line utilizations (%)
40
60
80
100
5
10
15
20
Hours

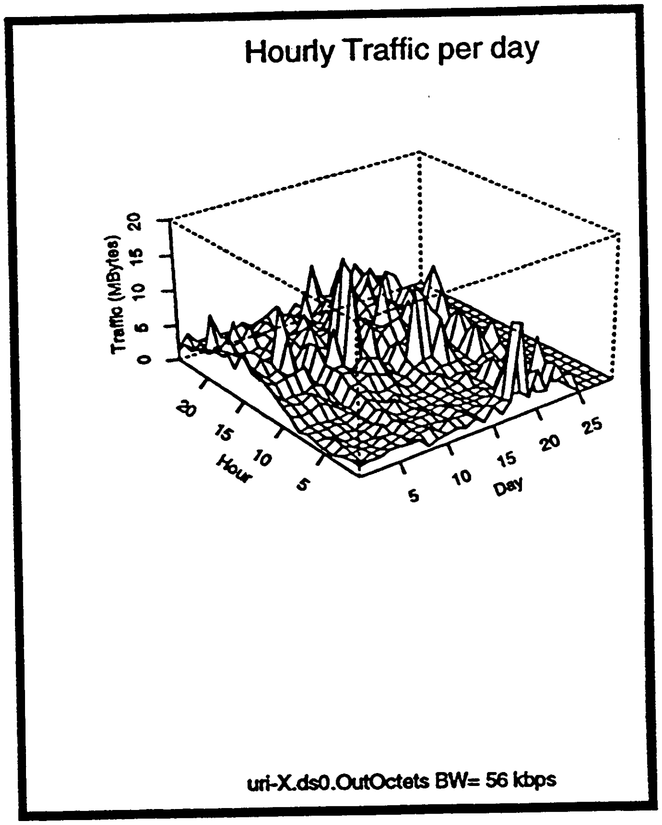
Hourly Traffic per day
Traffic (MBytes)
0 5 10 15 20
Hour
20 15 10 5
Day
5 10 15 20 25
uri-X.ds0.OutOctets BW= 56 kbps

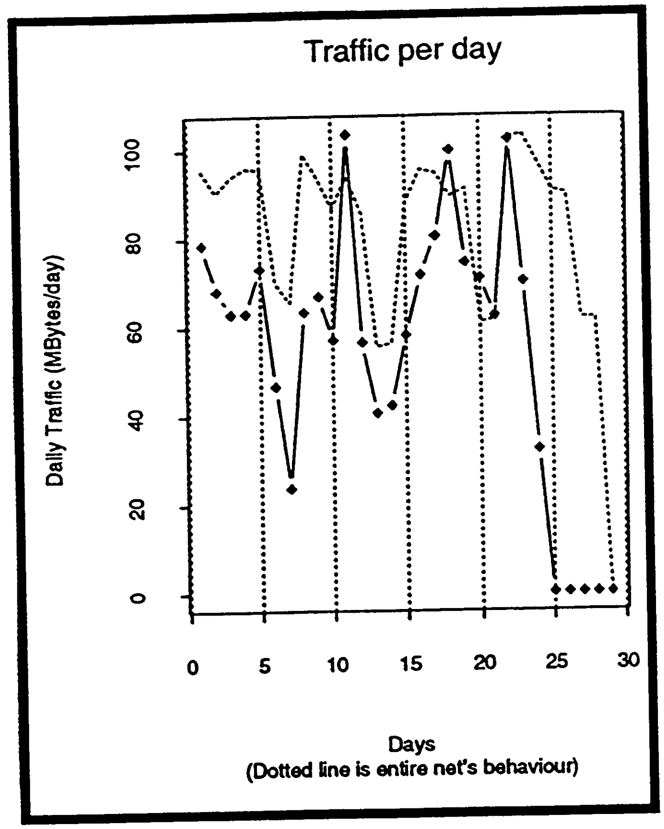
Traffic per day
Daily Traffic (MBytes/day)
0 20 40 60 80 100
0 5 10 15 20 25 30
Days
(Dotted line is entire net's behaviour)

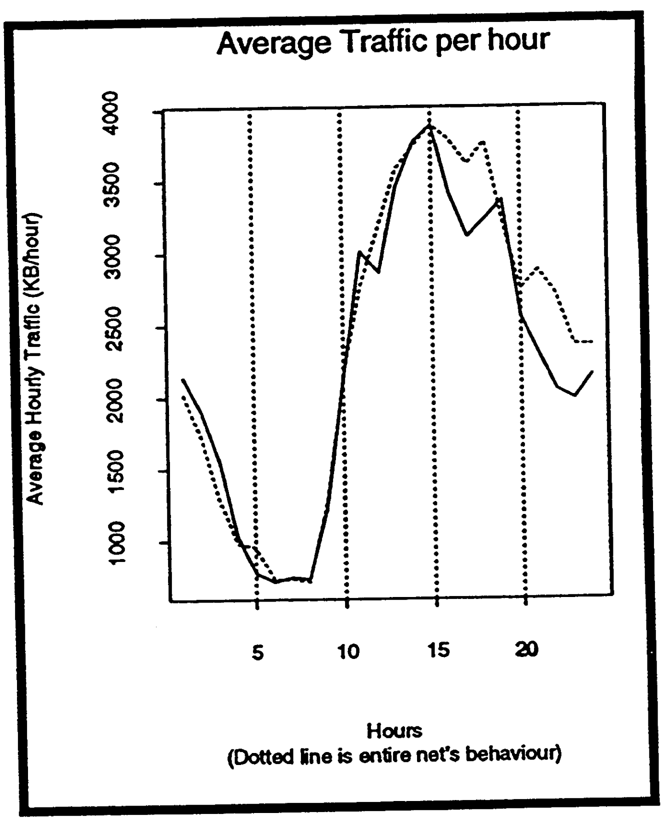
Average Traffic per hour
Average Hourly Traffic (KB/hour)
1000 1500 2000 2500 3000 3500 4000
5 10 15 20
Hours
(Dotted line is entire net's behaviour)

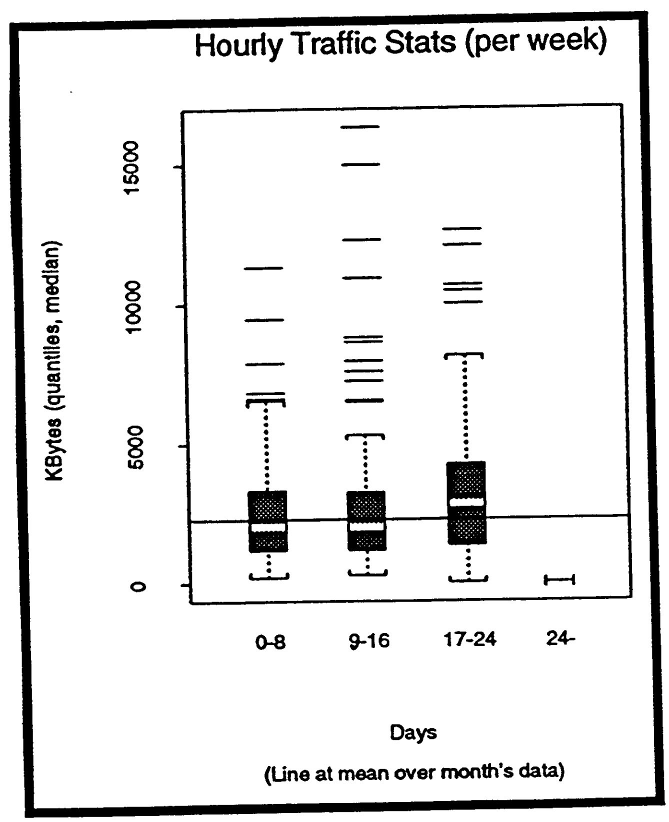
Hourly Traffic Stats (per week)
KBytes (quantiles, median)
0 5000 10000 15000
0-8 9-16 17-24 24-
Days
(Line at mean over month's data)

2.4.6 User Connectivity (ucp)

Charter

Chair(s):

Dan Long, long@nic.near.net

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done — Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.

TBD — Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate.

Request For Comments:

RFC 1297 "NOC Internal Integrated Trouble Ticket System Functional Specification Wishlist ("NOC TT REQUIREMENTS")"

2.5 Routing Area

Director(s):

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun

Inter-Domain Multicast Routing BOF (IDMR)

The entire IDMR session was spent in a discussion of the CBT protocol. Particular attention was given to the changes in the protocol since the November IETF.

- **Multicast Scope Control**

  One possible solution to multicast scope control was presented, based on having a separate group per level of scope required. This resulted in a considerable debate as to how multicast scoping should be defined and the requirements of users it should be able to satisfy. The solution presented was deemed unsuitable, and it was agreed to continue the discussion on the IDMR mailing list. The conclusion was that the Group should work towards a concise definition of multicast scope control.

- **Multicast Data Packets**

  There was a brief discussion on the issue of multicast data packets carrying the group-id as an IP option. The conclusion however, was that there was no more suitable alternative. It was also decided that non-primary cores should be less stringent in accepting join-requests. Further, an additional error detection mechanism is required by routers to distinguish on-tree packets arriving via a child as link-level unicast.

Paul Francis (formerly Tsuchiya) concluded the session with a description of how CBT will run over PIP.

Virtual Circuit Routing BOF (VCROUT)

The VCROUT BOF met twice. On Monday, Rob Coltun led a discussion on routing criteria for a seamless Virtual Circuit internet. Much of the discussion centered on address as well as terminology. Drew Perkins presented Fore Systems' routing strategy.

On Wednesday, Marco Sosa led a discussion on the scope of the Working Group's direction and goals. The Group agreed that they would include both intra-domain and inter-domain routing within their scope, but initially focus on intra-domain routing. The protocol for topology notifications and methods for aggregating topology information to provide possible routes to a call set-up will be addressed. Rob then provided an overview of the proposal drafted by Marco and himself.

Allison Mankin gave a presentation on congestion control implementation, signaling, and its relationship to QOS and routing. Finally, some modifications to the draft proposal were suggested and discussed.

The Group plans to meet in Amsterdam as a working group.

Inter-Domain Policy Routing Working Group (IDPR)

There are two new Internet-Drafts, one by Rob Austein on DNS extensions for IDPR and one by Woody Woodburn which is the latest version of the IDPR MIB. The Group encourages people to read these Drafts and send comments to the IDPR mailing list.

The Internet pilot demonstration of IDPR is scheduled to begin next week and will run for approximately one month. The results of the pilot, as well as a description of the installation, will be published in an Informational RFC at the conclusion of that period.

Discussion topics included the implications of domain hierarchies ("super-domains" in the architecture document terminology) and resource allocation in the context of IDPR.

Super-domain discussion included domain address representation; policies of super-domains; and obtaining more detailed information about the contents of a super-domain through mechanisms such as active distribution by constituent domains and queries from external domains.

Resource allocation discussion included a description of the "fair share" resource allocation mechanism as well as a general discussion of how to integrate resource allocation and policy routing. Topics included route generation heuristics to improve the probability of generating routes that supply the necessary resources, as well as passing flow control information back to the beginning of a path.

IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

The MOBILEIP Working Group met twice during the Columbus IETF. There were six formal presentations on different approaches to mobile IP followed by a discussion of how to actually make progress.

After listening to 6 talks on the status of old proposals and on new proposals, the Working Group decomposed the problem into 5 pieces (with an additional 5 "cross matrix" pieces). Each of the pieces was assigned to a Working Group member for them to edit a document documenting a solution to that part of the puzzle.

ISIS for IP Internets Working Group (ISIS)

The ISIS Working Group met for one session. The Group agreed to try and advance the Proposed Standard (RFC1195) and the MIB up the standard track to Draft and Proposed status respectively at the next IETF meeting.

The Group also discussed a number of new enhancements and extensions to the protocol. These include:

- Support for CLNP Multicast (maybe define something new - not what ANSI/ISO are defining).
- Adding a designated router "feature" to the protocol specification.
- Defining how to support multiple level 1 areas in one router.
- Defining multiple levels of hierarchy.
- Adding Appletalk and IPX integration into the protocol.
- Increasing the LSP number limit to 64K bytes.
- Increasing the metric range to 16 bit internal and 32 bit external.
- Methods to run over non-broadcast multi-access networks (e.g. SMDS, ATM, X.25, etc.).

Multicast Extensions to OSPF Working Group (MOSPF)

The MOSPF Working Group met for one session. John Moy discussed the MOSPF Analysis and Experience Draft that he prepared to accompany the MOSPF protocol specification, as required for all routing protocols submitted to the standards track. Christian Huitema raised a concern about the scaling properties of MOSPF, and suggested the use of Reverse Path Forwarding with on-demand pruning as a backup mechanism for cases of router memory or processor exhaustion. In the discussion that ensued, it was pointed out that, for the size of domains for which MOSPF is intended, the overhead of MOSPF is well within the capabilities of contemporary routers, given certain assumptions of worse-case behavior of multicast group members and senders. However, it was observed that a good model for multicast workload does not yet exist, thus making it difficult to judge the value of Christian's proposed extensions. The Group decided to submit the MOSPF draft, as is, to the IESG for publication as a Proposed Standard.

Open Shortest Path First IGP Working Group (OSPF)

After reviewing the four OSPF documents that were pending, the Group decided to:

- Submit the updated OSPF V2 spec for RFC publication, obsoleting RFC1247 (some urgency exists, since the Group wants CIDR changes communicated to the larger community);

- Submit the OSPF Trap MIB as a Proposed Standard;
- Publish a document describing how to implement OSPF on Frame Relay as an Informational RFC; and
- Delay the OSPF NSSA area document for small modifications.

The Group spent the majority of the remaining time discussing a proposal for carrying BGP path information in OSPF (to eliminate Internal BGP). At the end of the meeting, the Group outlined a document describing RIP to OSPF transition strategies.

OSI IDRP for IP over IP Working Group (IPIDRP)

The IPIDRP and BGP Working Groups met jointly with over 80 people in attendance for the two sessions. Issues discussed during the sessions include the following:

- PIP's requirements for BGP/IDRP.
- Status of BGP-4 documents.
- Size of Local Preference in BGP-4.
- Size of MULTI_EXIT_DISC in BGP-4.
- IDRP for IP documents.
  - IDRP for IP document.
  - IDRP for IP family document.
  - IDRP MIB.
- BGP-4 Transport Session Statistics and Routing Statistics.
- IDRP/BGP-4 to OSPF.
- OSPF Paper.

It was recommended that IDRP for IP be progressed to Proposed Standard as soon as one implementation was completed. The BGP-4 has one implementation and was recommended to be progressed to Proposed Standard.

Source Demand Routing Working Group (SDR)

A brief tutorial on SDRP was given. Changes to the packet format and forwarding specification since the last IETF were reviewed and approved without comment. Prototype development on this portion of the protocol will continue. Preliminary discussions were held on the contents of the usage document and on a proposed "futures" document. A list of other tasks were enumerated and volunteers were drafted.

CURRENT MEETING REPORT

Minutes of the IPv7 Addressing Working Group (BIGADDR)

Report not submitted.

Attendees

Robert Aiken	aiken@es.net
Randall Atkinson	atkinson@itd.nrl.navy.mil
Fred Baker	fbaker@acc.com
Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Tom Benkart	teb@acc.com
Andy Bierman	abierman@synoptics.com
Richard Bjers	rich.bjers@uc.edu
Russell Blaesing	rrb@one.com
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Scott Brim	Scott_Brim@cornell.edu
Sandy Bryant	slb@virginia.edu
Ross Callon	rcallon@wellfleet.com
David Carr	Carr@acsu.buffalo.edu
John Chang	jrc@uswest.com
Yee-Hsiang Chang	yhc@hpl.hp.com
Anthony Chow	chow_a@wwtc.timeplex.com
David Clark	ddc@lcs.mit.edu
Richard Cogger	R.Cogger@cornell.edu
Richard Colella	colella@nist.gov
James Davin	davin@thumper.bellcore.com
Steve Deering	deering@parc.xerox.com
Dino Farinacci	dino@cisco.com
Eric Fleischman	ericf@act.boeing.com
Paul Franchois	paulf@bldrdoc.gov
John Gawf	gawf@compatible.com
Eugene Geer	ewg@cc.bellcore.com
Robert Gilligan	gilligan@Eng.Sun.Com
Fengmin Gong	gong@concert.net
John Hascall	john@iastate.edu
Robert Hinden	hinden@eng.sun.com
John Hopprich	hopprich@davidsys.com
Woody Huang	yuh@merit.edu
Christian Huitema	christian.huitema@sophia.inria.fr
John Ioannidis	ji@cs.columbia.edu
Phil Irey	pirey@relay.nswc.navy.mil
David Jacobson	dnjake@vnet.ibm.com

Ronald Jacoby	rj@sgi.com
Frank Kastenholz	kasten@ftp.com
Stev Knowles	stev@ftp.com
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Tony Li	tli@cisco.com
David Lovering	lovering@bldrdoc.gov
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Paul Lustgraaf	grpjl@iastate.edu
Charles Lynn	clynn@bbn.com
Carl Madison	carl@startek.com
Bill Manning	bmanning@sesqui.net
Michael Marcinkevicz	mdm@csu.net
Donald Merritt	don@arl.army.mil
David Meyer	meyer@ns.uoregon.edu
Greg Minshall	minshall@wc.novell.com
Erik Nordmark	nordmark@eng.sun.com
William Owens	owens@cc.rochester.edu
John Penners	jpenners@advtech.uswest.com
James Philippou	japhilippou@eng.xyplex.com
Jon Postel	postel@isi.edu
Bala Rajagopalan	braja@qsun.att.com
David Reese	dave@csu.net
Ben Robinson	ben_robinson@vnet.ibm.com
Benny Rodrig	4373580@mcimail.com
Steve Romig	romig@cis.ohio-state.edu
Michael Safly	saf@tank1.msfc.nasa.gov
Tim Seaver	tas@concert.net
Allyson Showalter	allyson@nsipo.arc.nasa.gov
W. David Sincoskie	sincos@thumper.bellcore.com
Keith Sklower	sklower@cs.berkeley.edu
Martha Steenstrup	msteenst@bbn.com
Steve Suzuki	suzu@fet.com
Kamlesh Tewani	ktt@arch2.att.com
Richard Thomas	rjthomas@bnr.ca
Susan Thomson	set@bellcore.com
Robert Ullmann	Ariel@Process.com
L. Stuart Vance	vance@tgv.com
Kannan Varadhan	kannan@oar.net
Mario Vecchi	mpv@thumper.bellcore.com
Ruediger Volk	rv@informatik.uni-dortmund.de
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
William Warner	warner@ohio.gov
Kirk Williams	kirk@sbctri.sbc.com
John Wroclawski	jtw@lcs.mit.edu
Liang Wu	ltw99@bellcore.com

CURRENT MEETING REPORT

Reported by Tony Ballardie/UCL

Minutes of the Inter-Domain Multicast Routing BOF (IDMR)

The IDMR BOF was chaired by Tony Ballardie. The Minutes were taken by Benny Rodrig. The CBT Protocol was discussed for the duration of the session, in particular, the changes in the protocol since last November's IDMR BOF session.

Multicast Scope Control

One possible solution to multicast scope control was presented based on having a separate group per level of scope required. This resulted in a considerable debate as to how multicast scoping should be defined and the requirements of users it should be able to satisfy. The solution presented was deemed unsuitable, and it was agreed to continue the the discussion on the IDMR mailing list. The conclusion was that the Group should work towards a concise definition of multicast scope control.

Part of the group initiation procedure includes the group initiator (host) updating DNS (via a system administrator) with the relevant group information. The concern was raised that the procedure is likely to be too slow. This part of group initiation may need to be revised to keep group join latency to a minimum.

Multicast Data Packets

There was a brief discussion on the issue of multicast data packets carrying the group-id as an IP option. The conclusion however, was that there was no more suitable alternative.

When a parent/parent link fails, it was decided that both the flush-tree and re-join mechanisms should be implemented. Which of the two mechanisms eventually becomes redundant can only be decided after a period of performance testing.

It was decided that non-primary cores should be less stringent in accepting join-requests. Also, an additional error detection mechanism is required by routers to distinguish on-tree packets arriving via a child as link-level unicast.

The Draft includes no mention of child-parent "keepalives" and no mention of a handshake protocol between a group initiator and potential cores for that group. The Draft needs updating in these respects.

The control packet format should contain a group-id TYPE as an additional protocol independence feature.

Finally, Paul Francis gave a brief description of how CBT will run over PIP.

Attendees

Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Jim Beers	Jim.Beers@cornell.edu
Robert Braden	braden@isi.edu
Scott Brim	Scott_Brim@cornell.edu
David Carr	Carr@acsu.buffalo.edu
Stephen Casner	casner@isi.edu
Steve Deering	deering@parc.xerox.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Dino Farinacci	dino@cisco.com
Paul Francis	Francis@thumper.bellcore.com
Chris Gunner	gunner@dsmail.enet.dec.com
Joel Halpern	jmh@network.com
Ken Hayward	Ken.Hayward@bnr.ca
Robert Hinden	hinden@eng.sun.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Gerd Holzhauer	holzhauer1@applelink.apple.com
Tony Li	tli@cisco.com
Charles Lynn	clynn@bbn.com
Kent Malave	kent@bach.austin.ibm.com
Jun Matsukata	jm@eng.isas.ac.jp
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
Greg Minshall	minshall@wc.novell.com
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
William Owens	owens@acsu.buffalo.edu
Laura Pate	pate@gateway.mitre.org
Geir Pedersen	Geir.Pedersen@usit.uio.no
Willi Porten	porten@gmd.de
Benny Rodrig	4373580@mcimail.com
Sally Tarquinio	sallyt@gateway.mitre.org
Fumio Teraoka	tera@csl.sony.co.jp
Paul Traina	pst@cisco.com

CURRENT MEETING REPORT

Reported by Rob Coltun/Consultant

Minutes of the Virtual Circuit Routing BOF (VCROUT)

The Group's first session began with a discussion of the VC Routing Architecture Draft and ended somewhere on planet ATM-Forum-Flashback. Several interesting issues were raised here.

Radia Perlman suggested that the routing protocol should be scalable to any level of hierarchy so that a prefix address advertised by a switch may represent directly connected terminals or an aggregation address of a directly connected sub-network.

Juha Heinanen was concerned that a 32-bit address is not sufficient. It is big enough to identify the exchange portion of a North American Numbering Plan address, should be big enough for the routable portion of other E.164 addresses, and is big enough for the routable part of a GOSIP-style NSAP (RD/Area). However, it was pointed out that the GOSIP-style NSAP format may be used in a more efficient manner than suggested in RFC1237 (i.e., the Rsvd field may have significance and the RD/Area separation may not be clear cut). This makes a 32-bit address useless.

Dan Grossman had a strong objection to using the transit carrier ID for call setup in Virtual Circuit networks.

Fore's SPANS

Drew Perkins from Fore Systems presented an overview of Fore's SPANS (Simple Protocol for ATM Network Signaling) protocol. The presentation included Fore's internal addressing, UNI, NNI, QoS, SPF routing algorithm, multicast solution and connectionless server.

Goals of Proposed Working Group

The second session started with a discussion of the goals and milestones of the proposed VCROUT Working Group. The following goals were agreed to:

- The Working Group will initially focus on protocols for use within a single Virtual Circuit network (intra-domain routing). Routing for a Virtual Circuit internet (inter-domain routing) is also within the Charter.

- Initial specifications to focus on what each switch needs to monitor to describe its local topology (including neighbor interactions), and the protocol for distributing/updating topology information should be in Internet-Draft form by the end of June and submitted as a Proposed Standard by the end of 1993.

- Interaction between the CAC and the routing information should be worked out, such that by June, a description of how it would work, with a list of requirements needed from the UNI signaling and NNI signaling detailed for liaison with the ATM Forum. Any modifications should have worked out trade-offs described.

- As part of the interaction between the CAC and routing, ways of describing the performance expected by a switch rather than the algorithm used by the switch should be explored.

VC-Routing Proposal

The Group then had a brief overview of the VC-routing proposal, focusing on routing metrics and how they may be used by signaling.

Marek suggested that for high-speed networking, the minimum propagation delay between switches is related to the geographic distance between the switches.

Tony questioned the stability of a system that would be doing call set-up based on metrics that are in part reflecting available resources.

Allison gave a presentation on congestion control implications in signaling which began with the question "Do we really gain if routing protocol 'helps' call admission control?" and continued with an overview of a number of (as of yet unresolved) signaling issues regarding the terminal (hosts and routers) to switch interaction. This presentation (vcroute_mankin.txt) is available on the machine gated.cornell.edu in the pub/ospf directory.

Distant-Vector Protocol

There was a short discussion on the possible use of a distant-vector protocol to be used for a Virtual Circuit routing protocol. Yakov brought up the point that there is usually no policy routing within the IGP and that congestion avoidance and QoS requirements can be met on a hop-by-hop basis. Others countered that keeping the map at the source (i.e., using a link-state routing protocol) and doing a source route call setup was a superior technical approach.

Tony suggested that a Virtual Circuit domain should be able to inform IDRP of reachability information.

The Group agreed to meet at the Amsterdam IETF.

Attendees

David Arneson	`arneson@ctron.com`
Fred Baker	`fbaker@acc.com`
Jim Beers	`Jim.Beers@cornell.edu`

Nutan Behki	Nutan_Behki@qmail.newbridge.com
Lou Berger	lberger@bbn.com
Caralyn Brown	cbrown@wellfleet.com
Jeff Carman	tcarman@bnr.ca
George Clapp	clapp@ameris.center.il.ameritech.com
Robert Cole	rgc@qsun.att.com
Rob Coltun	rcoltun@ni.umd.edu
Dave Cullerot	cullerot@ctron.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Osmund DeSouza	osmund.desouza@att.com
Chas DiFatta	chas@cmu.edu
M.J. Dixon	mjd@att.com
Kurt Dobbins	kurtdob@ctron.com
Bob Downs	bdowns@combinet.com
Chip Elliott	celliot@bbn.com
Robert Enger	enger@reston.ans.net
Michael Fidler	fidler@mitre.org
Karen Frisa	karen.frisa@andrew.cmu.edu
Mike Goguen	goguen@synoptics.com
Kenneth Goodwin	goodwin@a.psc.edu
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.enet.dec.com
Joel Halpern	jmh@network.com
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Ken Hayward	Ken.Hayward@bnr.ca
Juha Heinanen	juha.heinanen@datanet.tele.fi
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Kathy Huber	khuber@bbn.com
Fong-Ching Liaw	fong@eng.sun.com
Andrew Malis	malis_a@timeplex.com
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@cmf.nrl.navy.mil
Keith McCloghrie	kzm@hls.com
Gerry Meyer	gerry@spider.co.uk
David O'Leary	doleary@cisco.com
Zbigniew Opalka	zopalka@agile.com
Ayal Opher	aopher@synoptics.com
John Penners	jpenners@advtech.uswest.com
Maryann Perez	perez@cmf.nrl.navy.mil
Drew Perkins	ddp@fore.com
Radia Perlman	perlman@dsmail.enet.dec.com
Hal Sandick	sandick@vnet.ibm.com
Shiva Sawant	shiva@synoptics.com
Andrew Schmidt	ags@uius.edu
Kanan Shah	kshag@cmf.nrl.navy.mil

Andrew Smith	asmith@synoptics.com
Marco Sosa	mxs@sabre.bellcore.com
Subbu Subramaniam	subbu@cup.hp.com
Terry Sullivan	terrys@newbridge.com
Sally Tarquinio	sallyt@gateway.mitre.org
Marek Tomaszewski	marek@net.com
Scott Wasson	sgwasson@eng.xyplex.com
James Watt	james@newbridge.com
Rick Wilder	wilder@ans.net
Steven Willis	steve@wellfleet.com

Routing/Addressing

- Provider-Based Routing
- Hierarichal Within Provider Space
  -- Subscriber Address Assigned From Provider Space
- Transit Carrier For Next-Hop Provider
  -- Within Net Dest Is Terminal Or Transit Net
- Need Loop Avoidance (Inter-Provider)

- Support FR and ATM Networks
- Logically Distince Subscribers On Public Net
- Corporate Backbone As Provider
- Mixed Address Space (E.164 and NSAP)

---

1) Corporation's Address
2) Backbone As Provider

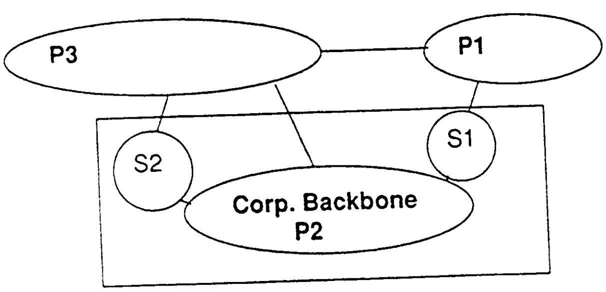

---

1) Subscriber net split across Provider net
2) Subscriber nets logically independent

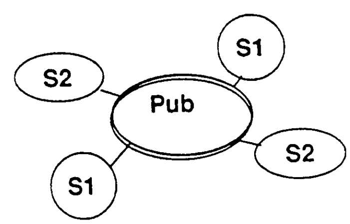

---

1) Subscriber Address
2) Destination Address
3) How Does S1 Reach S2

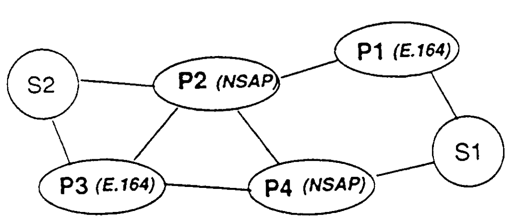

Route Calculation

- Up-Front
  -- All Destinations Determined When There Is A Change In Topology
    > Link Up/Down
    > Metric Change
  -- Alternate Routes Kept
    > Un-equal Cost

> Best Choice Based On Metrics

- On Demand
  - -- Use
    - > Best Guess Is Congested & No Alternate
    - > Best Guess Doesn't Have Sufficient Resources
  - -- Signaling Returns Id Of Failed Node
    - > Failed Node May Be Time Stamped As Congested (ie cached)

-- Criteria Based

> Bypass Congested Switches

> Resource Request

Bandwidth, Load, Closest

Metrics

- Available Bandwidth (Per Link)
  - -- Normalize Capacity/Available Capacity
    - > Normalize = 20*STS-12C
  - -- Dampened
    - > if New > 0.5 * MetricWindow Then Update
  - -- Total For Path Is Min Or Total (if inverse)

- Administrative Distance
  - -- Hop Count
- Load Factor
  - -- Total Available For All Links/ Total Capacity
  - -- Somewhat Related To Delay And Error
  - -- Updated With Available Bandwidth
  - -- Total For Path Is Min Or Total (if inverse)
- Best Path = Most Bandwidth & Closer & Least Loaded

- Source Route
  - -- Uses Less Network Resources
  - -- HBH Requires All Nodes To Agree On Metrics Chosen
    - > No Policy Choices (on demand)
  - -- Loop Avoidance

2.5.1 Border Gateway Protocol (bgp)

Charter

Chair(s):

Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:

General Discussion: bgp@ans.net
To Subscribe: bgp-request@ans.net
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:

Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Complete development of Version 2 of the Border Gateway Protocol (BGP).
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 Version 3.
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.
Done	Post an Internet-Draft specifying multicast extensions to BGP.
Done	Post the specfication of BGP 4 as an Internet-Draft.
Done	Post an Internet-Draft specifying a MIB for BGP Version 4.
Jan 1993	Submit the multicast extensions to BGP to the IESG as a Proposed Standard.
Jan 1993	Submit the specification for BGP Version 4 to the IESG for consideration as a Proposed Standard.
Jan 1993	Submit the BGP Version 4 MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"IP Multicast Communications Using BGP", 08/26/1991, Scott Brim, Yakov Rekhter <draft-ietf-bgp-multicast-02.txt>

"A Border Gateway Protocol 4 (BGP-4)", 05/05/1992, Y. Rekhter, T. Li <draft-ietf-bgp-bgp4-05.txt>

"Definitions of Managed Objects for the Border Gateway Protocol (Version 4)", 09/01/1992, S. Willis, J. Burruss, J. Chu <draft-ietf-bgp-mibv4-01.txt>

"BGP4/IDRP for IP—OSPF Interaction", 09/15/1992, K. Varadhan, S. Hares, Y. Rekhter <draft-ietf-bgp-bgp4ospf-interact-01.txt>

"Application of the Border Gateway Protocol in the Internet", 09/28/1992, Y. Rekhter, P. Gross <draft-ietf-bgp-application-01.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"

RFC 1397 "Default Route Advertisement In BGP2 And BGP3 Versions Of The Border Gateway Protocol"

RFC 1403 "BGP OSPF Interaction"

CURRENT MEETING REPORT

Minutes of the Border Gateway Protocol Working Group (BGP)

Report not submitted. The attendees listed below were participants in the joint BGP/IPIDRP session.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Nagaraj Arunkumar	nak@3com.com
Dennis Baker	dbaker@wellfleet.com
Tony Bates	tony@ripe.net
Jordan Becker	becker@ans.net
Rebecca Bostwick	bostwick@es.net
Ed Brencovich	edb@dss.com
Robert Brenner	Robert.W.Brenner@gte.sprint.com
Robert Calderon	calderon@noc.ans.net
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
Michael Collins	collinsms@es.net
David Conrad	davidc@iij.ad.jp
Osmund DeSouza	osmund.desouza@att.com
David Dubois	dad@pacersoft.com
Kishan Dudkikar	kishan@icm1.icp.net
Tom Easterday	tom@cic.net
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Michael Fidler	fidler@mitre.org
Peter Ford	peter@goshawk.lanl.gov
Vince Fuller	vaf@stanford.edu
Darren Griffiths	dag@ossi.com
Susan Hares	skh@merit.edu
Ittai Hershman	ittai@ans.net
David Jacobson	dnjake@vnet.ibm.com
Ronald Jacoby	rj@sgi.com
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Laurent Joncheray	lpj@merit.edu
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
John Krawczyk	jkrawczy@wellfleet.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Duane Kuang	duanek@kalpana.com
Tony Li	tli@cisco.com

Frank Liu	fcliu@pacbell.com
Peter Lothberg	roll@stupi.se
Glenn Mackintosh	glenn@canet.ca
Jamshid Mahdavi	Mahdavi@a.psi.edu
James Mahon	mahonj@honfsi1.att.com
Bill Manning	bmanning@sesqui.net
Jun Matsukata	jm@eng.isas.ac.jp
Daniel McRobb	dwm@noc.ans.net
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
Dennis Morris	morrisd@imo-uvax.disa.mil
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
Julie Myers	jmyers@network.com
Erik Nordmark	nordmark@eng.sun.com
Peder Norgaard	pcn@tbit.dk
Brad Parker	brad@fcr.com
Andrew Partan	asp@uunet.uu.net
Laura Pate	pate@gateway.mitre.org
Michael Patton	map@bbn.com
Willi Porten	porten@gmd.de
Selina Priestley	sfp@noc.ans.net
Bala Rajagopalan	braja@qsun.att.com
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	amr@tycho.ncsc.mil
Ben Robinson	ben_robinson@vnet.ibm.com
Manoel Rodrigues	manoel_rodrigues@att.com
Michael Safly	saf@tank1.msfc.nasa.gov
John Scudder	jgs@merit.edu
Paul Serice	serice@cos.com
Kim Smith	kas@noc.ans.net
Steve Suzuki	suzu@fet.com
Steven Sweeney	steves@farallon.com
John Tavs	tavs@vnet.ibm.com
Marten Terpstra	marten@ripe.net
Kamlesh Tewani	ktt@arch2.att.com
Richard Thomas	rjthomas@bnr.ca
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Curtis Villamizar	curtis@ans.net
Hung Vu	hungv@fonorola.com
Scott Wasson	sgwasson@eng.xyplex.com
Jack Waters	waters@sura.net
Steven Willis	steve@wellfleet.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net

Jane Wojcik jwojcik@bbn.com
Richard Woundy rwoundy@vnet.ibm.com
Paul Zawada Zawada@ncsa.uiuc.edu

2.5.2 IP Routing for Wireless/Mobile Hosts (mobileip)

Charter

Chair(s):

Steve Deering, `deering@parc.xerox.com`
Greg Minshall, `minshall@wc.novell.com`

Mailing Lists:

General Discussion: `mobile-ip@parc.xerox.com`
To Subscribe: `mobile-ip-request@parc.xerox.com`
Archive: `parcftp.xerox.com:~/pub/mobile-ip/mail-archive`

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	Review the Charter of the Mobile IP Working Group for additional work required to facilitate non-host mobility.
Mar 1993	Submit the Mobile Host Protocol to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Minutes of the IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

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

Attendees

Kannan Alagappan	kannan@dsmail.lkg.dec.com
Steve Alexander	stevea@lachman.com
Jim Barnes	barnes@xylogics.com
William Barns	barns@gateway.mitre.org
Jim Beers	Jim.Beers@cornell.edu
Tom Benkart	teb@acc.com
Richard Bjers	rich.bjers@uc.edu
Monroe Bridges	monroe@cup.hp.com
David Bridgham	dab@epilogue.com
Gregory Bruell	gob@wellfleet.com
Ramon Caceres	ramon@mitl.com
David Carr	Carr@acsu.buffalo.edu
Stephen Casner	casner@isi.edu
Cyrus Chow	cchow@ames.arc.nasa.gov
Simon Coppins	coppins@arch.adelaide.edu.au
Shane Dawalt	sdawalt@desire.wright.edu
Steve Deering	deering@parc.xerox.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Tony DeSimone	tds@hoserve.att.com
Pierre Dupont	dupont@mdd.comm.mot.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Eric Fleischman	ericf@act.boeing.com
Francois Fluckiger	fluckiger@vxcern.cern.ch
Craig Fox	foxcj@network.com
Paul Franchois	paulf@bldrdoc.gov
Paul Francis	Francis@thumper.bellcore.com
Ron Frederick	frederick@parc.xerox.com
Robert Gilligan	gilligan@Eng.Sun.Com
Danny Hanson	hanson.tic@commlan.safb.af.mil
John Hascall	john@iastate.edu
Frank Heath	heath@cmc.com
Robert Hinden	hinden@eng.sun.com
John Hopprich	hopprich@davidsys.com
John Ioannidis	ji@cs.columbia.edu

Ronald Jacoby	rj@sgi.com
David Johnson	dbj@cs.cmu.edu
Scott Kaplan	scott@wco.ftp.com
Phil Karn	karn@qualcomm.com
Charles Kaufman	kaufman@zk3.dec.com
Charley Kline	cvk@uiuc.edu
Giri Kuthethoor	giri@ms.uky.edu
Paul Lambert	paul_lambert@email.mot.com
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Mark Lewis	Mark.S.Lewis@telebit.com
Tony Li	tli@cisco.com
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Paul Lustgraaf	grpjl@iastate.edu
Charles Lynn	clynn@bbn.com
Bruce Mackey	brucem@cinops.xerox.com
James Mahon	mahonj@honfsi1.att.com
Michael Marcinkevicz	mdm@csu.net
Brian Marsh	marsh@mitl.com
Donald Merritt	Don@brl.mil
Greg Minshall	minshall@wc.novell.com
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Erik Nordmark	nordmark@eng.sun.com
Masataka Ohta	mohta@cc.titech.ac.jp
William Owens	owens@cc.rochester.edu
Brad Parker	brad@fcr.com
John Penners	jpenners@advtech.uswest.com
Charles Perkins	perk@watson.ibm.com
Drew Perkins	ddp@fore.com
Marcus Ranum	mjr@tis.com
Jim Rees	jim.rees@umich.edu
David Reese	dave@csu.net
Tony Richards	richards@icm1.icp.net
Ben Robinson	ben_robinson@vnet.ibm.com
Manoel Rodrigues	manoel_rodrigues@att.com
Hal Sandick	sandick@vnet.ibm.com
Vilson Sarto	vilson@fapq.fapesp.br
William Simpson	Bill.Simpson@um.cc.umich.edu
Michael St. Johns	stjohns@darpa.mil
Martha Steenstrup	msteenst@bbn.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Terry Sullivan	terrys@newbridge.com
Steve Suzuki	suzu@fet.com
Steven Sweeney	steves@farallon.com
John Tavs	tavs@vnet.ibm.com

Fumio Teraoka	tera@csl.sony.co.jp
Richard Thomas	rjthomas@bnr.ca
Susan Thomson	set@bellcore.com
Paul Traina	pst@cisco.com
John Veizades	veizades@apple.com
Warren Vik	wmv@lachman.com
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Von Welch	vwelch@ncsa.uiuc.edu
Kirk Williams	kirk@sbctri.sbc.com

Extensions to Protocols for Mobile Internetworking

Kannan Alagappan
Digital

MH Adaptability

o Success of mobile ip dependent on mobile ip client integrated into as many portable devices and as many TCP/IP stacks.

o Using MicP protocol between MH-MSR may require kernel changes for some MHs.

o Suggest that all MH-MSR communication occur over UDP packets to a well known port.

MH Adaptability (cont'd.)

o The following packet types would go over UDP.

- o ping
- o beacon
- o great
- o grack
- o grnack
- o popup

MH over Point-to-Point Links

o Mobile ip needs to work over wireless, directly connected, and point-to-point links.

o Problem when an MH connects over a point-to-point link through a network access server.

o NAS may do proxy-ARPing for attached IP hosts and not full bridging.

o In this case, MSR beacon packets sent via broadcast never reach the MH.

MH over Point-to-Point Links (cont'd.)

o Two Approaches
- o Fix NAS
- o Extend MicP Protocol

o MH uses IP multicast to solicit an MSR to send it to a unicast beacon.

o Suggest a new packet type called MicP_SOLICIT

o packet type	(1 octet)
o subcode	(1 octet)
o checksum	(2 octet)
o IP addr of MH	(4 octet)
o Vendor extensions	(? octect)

MH over Point-to-Point Links (cont'd.)

o When a MH is booted up, it has the option of either listening for beacons or multicasting a solicit packet.

o When a MSR receives a solicit packet, it sends a unicast beacon to the specified MH.

o MH registers itself with the MSR and turns off beacon tracking if connected via a point-to-point link.

o Other uses of solicit protocol.
- o timely registration in a cell with slow beacon rate.
- o assist in inter-campus popup security.

Mobile Host Internetworking Using IP Loose Source Routing

David B. Johnson

School of Computer Science
Carnegie Mellon University

`dbj@cs.cmu.edu`

Mobile Hosts

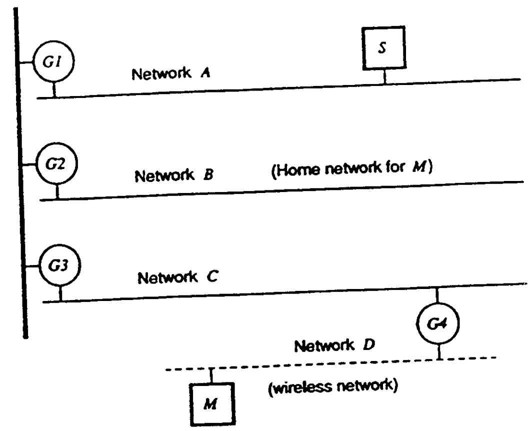

Examples:

- Carry somewhere and use while away from home
- Handheld, continuously moving hosts

Keep same IP address when move to new network

Goals

- Transparent above IP layer
- Use existing standard IP features where possible
- Minimize changes to existing software
- No changes required in non-mobile hosts
- Scalable
- Low overhead

IP Loose Source Routing

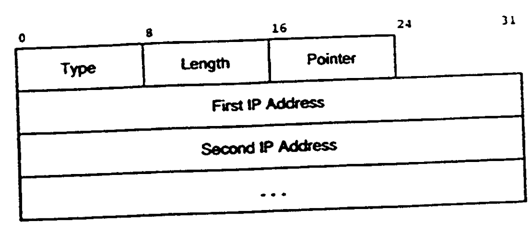

Deliver to destination address in IP header

While address list in option is not finished:

- Copy next address to destination in IP header
- Increment pointer in option
- Deliver to new destination address in IP header

Basically, a list of gateways to be sure the datagram goes through

Use normal IP routing, any number of hops to get to each gateway in turn

IP Address Assignment

IP network number assigned to organization by IANA

Host number assigned by that organization

Mobile host address assignment is *not* a special case

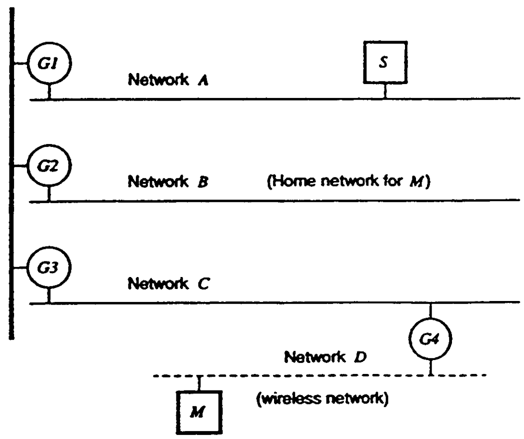

Mobile host always uses *only* its home IP address

The Foreign Gateway

Keeps list of visiting mobile hosts currently on subnet

Transmits arriving datagrams locally to mobile host

May instead be a separate support host on subnet

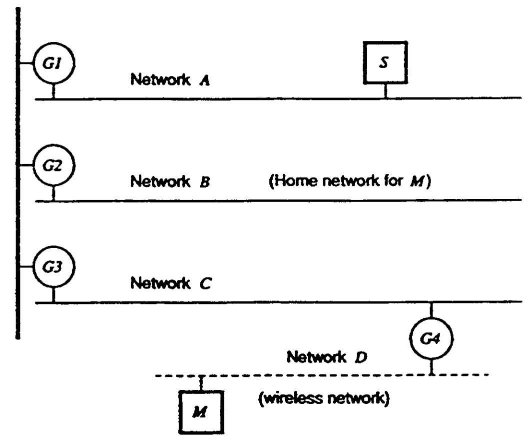

The Home Gateway

Maintains database of location of subnet's mobile hosts (foreign gateway IP address)

May have more than one gateway

May instead be a separate support host on subnet

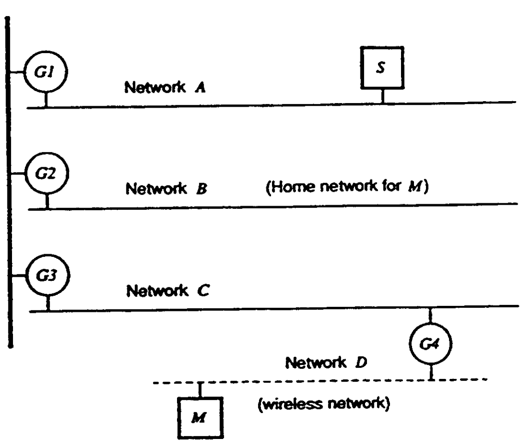

IP Delivery to Mobile Hosts

Loose source route datagrams through foreign gateway

Uses normal IP routing to get there

The foreign gateway:

- Optionally removes source route option from datagram
- Transmits datagram over local subnet to mobile host

Source route option is put into datagram by:

- The sending host, or
- Any intermediate gateway, or
- The home gateway

Anybody can cache location of a mobile host

Nobody is required to cache it

Nothing special to do for sending *from* a mobile host

Moving a Mobile Host

Disconnecting from the network:

- Mobile host notifies old foreign gateway and home gateway

Reconnecting to the network:

- Mobile host notifies new foreign gateway and home gateway
- Can disconnect from old foreign gateway at the same time

Old foreign gateway removes mobile host from list

New foreign gateway adds mobile host to list

Home gateway:

- Maintains record of host's current location
- Answers ARP requests with proxy ARP while mobile host is not connected to home network

Mobile Host Movement Notification

Sent to home gateway and new foreign gateway

0 8 16 31

Type	Code	Checksum
IP Address of Foreign Gateway		
Sequence Number		

Home gateway notifies old foreign gateway if needed

Retransmitted until acknowledged

Sequence number allows duplicate detection

Locating a Mobile Host

Sender never knows (does not need to know) that destination host is mobile

A mobile host's IP address will route datagrams for it to its home network

The home gateway:

- Adds loose source route option to datagram and resends it
- Returns ICMP "mobile host redirect" to sender
- Sender or any gateway can cache this information

Old cached location information in sender or other gateways is automatically corrected when/if needed

No overhead when mobile host is "at home"

Any host running the right software can become mobile without special arrangement

ICMP Mobile Host Redirect

Define new code number within standard ICMP redirect message type

0 8 16 31

Type	Code	Checksum
Gateway IP Address		
IP Header and First 64 Bits of Original Datagram		
...		

Sent to original sender of datagram

Handled similar to existing host-specific ICMP redirect:

- Store in same table (new type field on records)
- Lookup already required for host-specific redirect
- Can find foreign gateway address for mobile host in table with little or no extra cost

Error Recovery

If foreign gateway forgets mobile host, IP datagrams will end up at the home gateway

Will contain "spent" loose source route option

Home gateway sends "mobile host query" to foreign gateway

0	8	16	31
Type	Code	Checksum	
IP Address of Foreign Gateway			
IP Address of Mobile Host			

Foreign gateway transmits onto indicated local subnet

Mobile host performs reconnect

Use of IP Options

The loose source route option:

- Hosts need not see loose source route option
- Foreign gateway must interpret it
- Home gateway must be able to add it to datagrams
- Other gateways only must pass option through unmodified
- Sending host and/or intermediate gateways may add option to datagrams
- Does not rely on option's recorded route

No other IP options used

Overhead

- Adds only 8 bytes to IP datagram instead of 20 or 24
- Increases minimum IP datagram size with TCP only 20%, not 50%
- No overhead on sending from mobile host, only on sending to mobile host
- No overhead when mobile host is "at home"
- Allows optimal routing when sending host or sender's gateway implements mobile IP protocol
- Near-optimal routing when some intermediate gateway understands mobile IP
- Movement does not result in timeout or "destination unreachable" on next datagram from sender

Scalability

- No global database
- No broadcast or multicast
- No temporary IP address assignment
- No need to maintain forwarding pointers at old locations
- Each organization manages its own IP address assignment for its own mobile hosts
- Each organization serves as database of its own mobile hosts' current locations
- No special configuration of mobile host IP address
- Any host running the right software can become mobile at any time

Summary

- Transparent above IP layer
- Uses existing standard IP features where possible
- Few changes to existing software
- No changes required in non-mobile hosts
- Scalable
- Low overhead
- Simple protocol

For More Information

The tech report:

David B. Johnson.
Mobile Host Internetworking Using IP Loose Source Routing.
Technical Report CMU-CS-93-128, School of Computer Science, Carnegie Mellon University, February 1993.

Available by anonymous FTP:

Host: `reports.adm.cs.cmu.edu`
File: `1993/CMU-CS-93-128.ps`

E-mail address:

`dbj@cs.cmu.edu`

Mobile IP Working Group

Charlie Perkins

- o Usenix Symposium on Mobile ... Computing
- o Areas of interest in Mobile Computing
- o WG focus areas
- o A future, a view of the present, and a road map
- o Our status at IBM

Parts of the Mobile Networking Problem

- o Link layer, MAC layer
  - o Beaconing vs. power-saving H/W?
  - o What is the H/W like??
    - - Infrared, RF
  - o Bridging for single-net installations
- o "Ad-hoc" networking?
- o Disseminate host routes to all routers?

Parts of the Mobile Networking Problem (cont'd.)

- o Other layer-3 solutions (e.g., US)
- o Disconnected filesystem operations (e.g., CODA)
- o Resource Location
  - o DHCP
  - o SNMP
- o Dynamic DNS, with "Mobile Server" RR ("MR", "RD")

WG Focus Areas

- o Framework, terminology
  - o Basic tools for good design (need sharpening)
- o Clear, clean problem statement
- o Location update protocol
  - o ICMP Host_Moved
  - o Port {UDP}
  - o SSR ping ☺
- o Vision of the future, road map

WG Focus Areas (cont'd.)

- o Connections to other WG's
  - o DNS o DHCP, SNMP o IPSEC
- o Monitor MAC-layer development {IEEE}
  - ===> Express our needs
- o DOS? OS/2? NT? Pen-Windows?
- o Other connectionless layer-3 protocols?
- o Pick an ecapsulation protocol
  - o IPIP? o SIPIP

A View of the Future

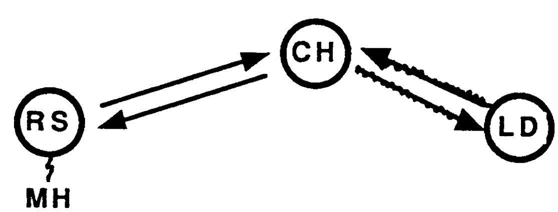

- o Each CH is its own RD/RS (i.e., gets location & re-addresses)

A View of the Future (cont'd.)

RS = Re-addresing Server (≃ "base station")

LD = Location Directory (HomeAddr <-> Forw.Addr)

CH = Correspondent host

RD = Redirector (specialized router) ("Mobile Server")

A View of the Present

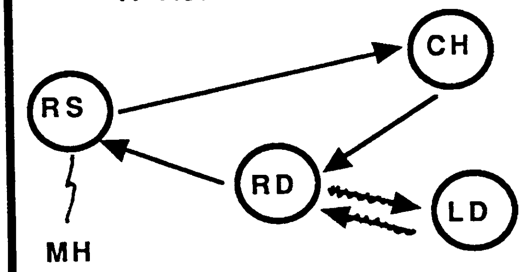

o If there are many Δs, Δ routing is ok

A Road Map

- o Gracefully phase in RDs/LD
- o Equip CH with encapsulation
- o Equip CH with LD interface
  - o e.g., ICMP
- o Create modified "route" command syntax {RT_ENC} --->Encapsulate outgoing packets plus rtioctl ().
- o Fix transport layer route caching
- o Fix DHCP for popups
- o Create new "Mobile Server" RR

Our Status

- o "Productized" code, ready for encapsulation drop-in.
- o Kernel works "vanilla" unless mobile ioctl's are isued, for all three agents MH, MAS, MR.
- o Product group interest
  - o We cannot offer total solution from our project

Our Status (cont'd.)

- o Created "Framework" Draft RFC
- o Papers almost done
- o Getting some real MHs
- o OS/2
- o Source Code giveaway

Internet Packet Transmission Protocol (IPTP) Packet Forwarding for Mobile Hosts

Hiromi Wada
Tatsuya Ohnishi
Matsushita Electric Industrial Co., Ltd.

Brian Marsh
Matsushita Information Technology Laboratory

April 1993

Overview

Brief IPTP recap

- Architecture
- Benefits/Goals

Recent Work

- Streamlining
- Robustness
- Implementation

Future Plans

Why IPTP?

IPTP Recap

Functions provided to IP client protocols

- Host location tracking
- Packet forwarding

3 entities

- Mobile Hosts (MH)
- Stationary Hosts (SH)
- Packet Forwarding Server (PFS)

2 communication modes

- *Forwarding mode* for backward compatibility
- *Autonomous mode* for optimal routing

MH has two IP addresses

- Home address that does not change
- Temporary address that changes when roaming
- Applications use only home address

Benefits/Goals

Backward compatibility

- User code uses only home IP address
- Forwarding mode allows mobile communication with unmodified stationary hosts
- Encapsulation within standard IP packets requires no router modifications

Efficient packet routing

- Autonomous mode allows mobile traffic to find optimal route

Low traffic load

- Primary copy location database reduces host lookup time
- Lazy propagation of updates reduces unnecessary handoff traffic

Recent Work

Streamline protocol

- Removed one redundant message type: "MH Visiting"

Fix problem

- Failure of home PFS disables roaming

Prototype implementation

- Under SunOS for Sun3s and SparcStations

Single-Point Failure: The Home PFS

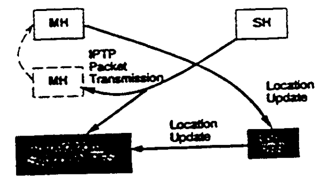

Home PFS is a single point of vulnerability

- Location updates cannot be propagated
- Roaming hosts will vanish

Replicate Updates

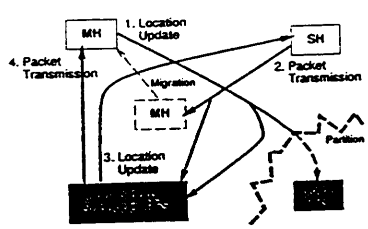

On MH migration

- Location update goes to both Autonomous Supporter and Home PFS

On SH transmission

- Autonomous Supporter (AS) PFS intercepts packet
- AS sends location update to SH
- AS forwards packet to MH

Implementation

SunOS prototype

- Sun3/60s and SparcStations

PFS

- User-level server
  - Location Management
  - Encapsulation
- Utilizes NIT
- 560 lines of C

MH

- User-level
  - Address allocation/reconfiguration
  - 760 lines of C
- Kernel-level
  - Location management (address mapping table)
  - Encapsulation/Decapsulation
  - 450 lines of C

SH

- Forwarding Mode – No modifications
- Autonomous Mode – 260 lines of kernel code
  - Location management
  - Encapsulation

Transmission Throughput

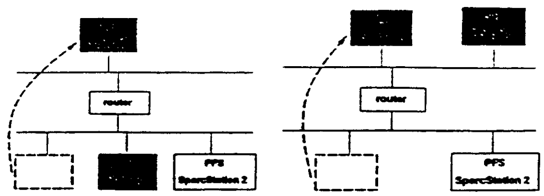

Case A   Case B

Transmission Rates (Mb/s)

- TCP data stream between SH and MH

Configuration Mode	Case A	Case B
Normal IP	1.6	2.6
Forwarding	1.2 (78%)	1.1 (40%)
Autonomous	1.4 (88%)	2.3 (86%)

A Real Application

Normal IP

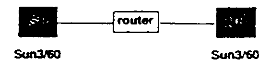

Autonomous Mode

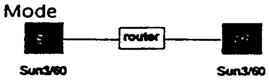

Forwarding Mode

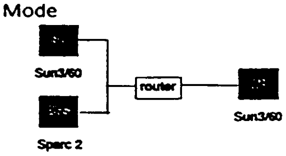


```
% ls -l /vmunix
-rwxr-xr-x  1 root     1240596 Jul 29  1992 /vmunix*
% time cp /vmunix /net/MH/tmp
```

Transmission Time

Configuration Mode	Time (seconds)
Normal IP	17.4
Autonomous	18.17
Forwarding	19.24

Future Work

Multi-PFS protocol

- Load-sharing to allow the protocol to scale in the number of MHs supported
- Replication to improve fault tolerance by increasing the number of PFSs that can provide "home" service

Conclusions

- IPTP provides transparent mobile internetworking
- Working implementation
- Current work increases tolerance to single-point failures
- Future work will add PFSs for scalability and fault tolerance

Why IPTP?

- One mechanism for intra- and inter-campus
- Connections persist across home network failures
- Modifications to SH not necessary but can improve fault tolerance and performance
- Simple – Approximately 2000 lines in prototype
- Works with existing infrastructure

Mobility in Pip

(such as it its)
Paul Francis

- o Current solution not "complete"
  - o Missing element of injecting mobility information into local routing.
  - o Current solution doesn't negate such technique.

Basic Technique

- o Pip header has End System ID and Address.
  - o ID just used to identify hosts.
  - o Address <u>not</u> used to identify hosts.
- o Mobile host has one or more "Mobile Address Servers".

Basic Technique (cont'd)

- o Mobile Address Server holds current address of mobile hosts.
  - o Mobile Address Server is stationary.
- o Address of Mobile Address Servers learned from DNS

Basic Technique (cont'd)

- o Two modes of operation "secure" and "insecure".
- o With insecure, when mobile host gets new address:
  - o Inform Mobile Address Servers.
  - o Inform active destinations, possibly by just sending data packet.
    - o ID binds new address to existing connection.

Basic Technique (cont'd)

- o If stationary host can't reach mobile host:
  - o Send query to Mobile Address Server.
  - o If address Assignment technique has "I'm leaving" message, unreachable message can indicate movement.

Basic Technique (cont'd)

- o If stationary host doesn't trust new address message from mobile host, can query Mobile Address Server.

The Virtual Network Protocol for Host Mobility

Fumio Teraoka (tera@csl.sony.co.jp)

March 31, 1993

Sony Computer Science Laboratory Inc.
3-14-13 Higashigotanda, Shinagawa-ku, Tokyo 141, Japan

Overview

- Two network addresses of a host:
  - VN-addresses: immutable logical addresses.
  - PN-addresses: indicate host location.
  - Both have the same format.
  - The transport layer uses VN-addresses.
- Two sublayers in the network layer:
  - VN-sublayer: resolves VN-addresses into PN-addresses.
  - PN-sublayer: the conventional network layer.

⇓

The Virtual Network Protocol

Terminology

- Home network: the subnetwork indicated by the VN-address of a host.
- Address resolution: processing in which a VN-address is resolved into a corresponding PN-address.
- Address resolver: a router which executes address resolution.
  - primary address resolvers: address resolvers in the home network.
  - secondary address resolvers: address resolvers in subnetworks other than the home network.
  - temporary proxy: address resolvers in the subnetworks to which a host was previously connected.

The Virtual Network Protocol

source PN-address
destination PN-address
flags
source VN-address
source address version number
destination VN-address
destination address version number
holding time
resolver PN-address
(option) authentication data

(a) data packet header

source PN-address
destination PN-address
flags
source VN-address
source address version number
target system VN-address
target system PN-address
target system address version number
target system holding time
(option) authentication data

(b) control packet header

Figure 1: Packet formats for VN-protocol

The Virtual Network Protocol

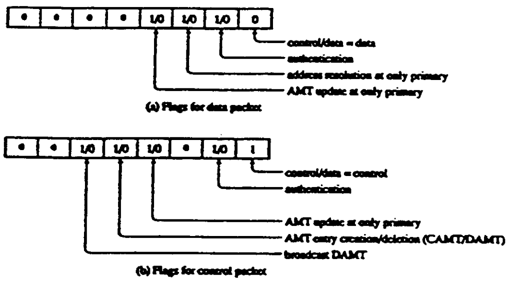

Figure 2: Flags field of the packet

The Virtual Network Protocol

VN-address
PN-address
address version number
holding time
flags

Figure 3: AMT entry for the VN-protocol

The Virtual Network Protocol

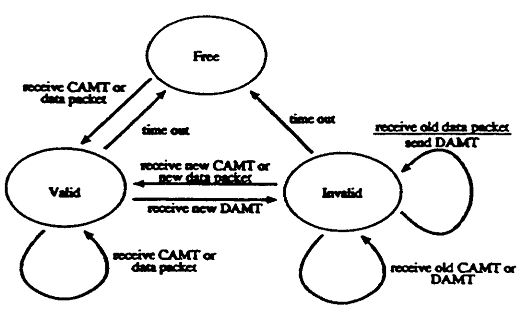

Figure 4: AMT entry state transition diagram

The Virtual Network Protocol

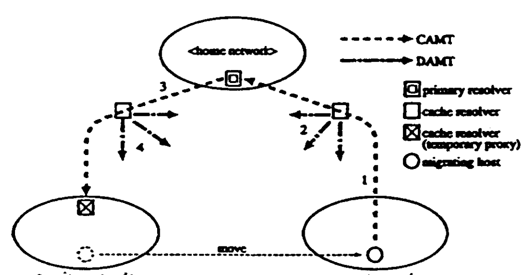

Figure 5: Host connection

The Virtual Network Protocol

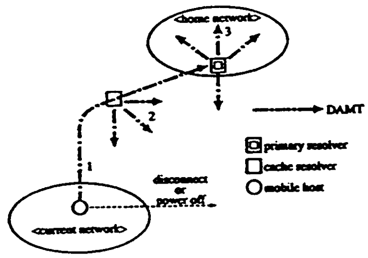

Figure 6: Host disconnection or power off

The Virtual Network Protocol

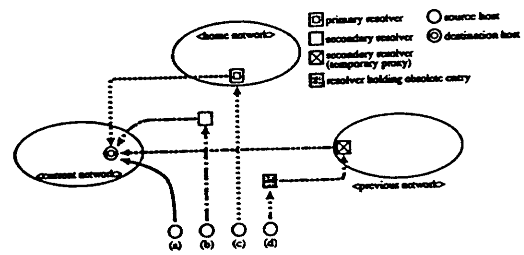

Figure 7: Data communication

VIP: Virtual Internet Protocol

§ VIP = IP + VN-protocol

- VN-addresses
  - ⇒ VIP addresses
- PN-addresses
  - ⇒ IP addresses
- home network
  - ⇒ the subnetwork indicated by the VN-address
- primary resolver
  - ⇒ routers in the home network

VIP: Virtual Internet Protocol

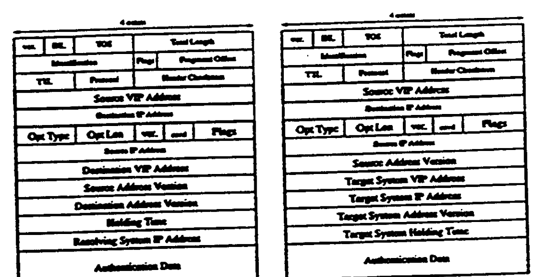

(a) Data packet (b) Control packet

Figure 8: VIP header formats

VIP: Virtual Internet Protocol

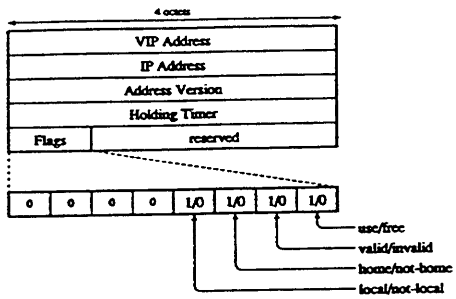

Figure 9: AMT entry for VIP

VIP: Virtual Internet Protocol

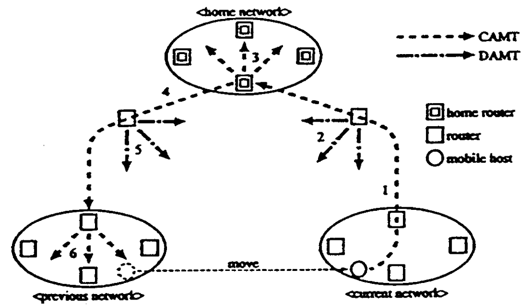

Figure 10: Packet flow upon host migration in VIP

ISO-VIP: VN-Protocol for CLNP

§ ISO-VIP = CLNP + VN-protocol

- VN-addresses
  ⇒ V-NSAP addresses (Virtual NSAP addresses) or V-NET (Virtual Network Entity Title)
- PN-addresses
  ⇒ NSAP addresses or NET
- primary resolver
  ⇒ the level1-IS that advertises connectivity.
- home network
  ⇒ the subnetwork to which the primary resolver is connected.

ISO-VIP: VN-Protocol for CLNP

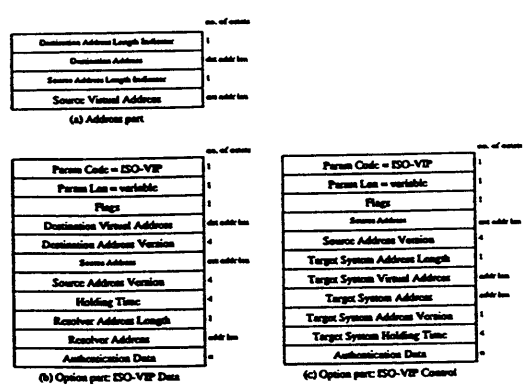

Figure 11: ISO-VIP packet header formats

ISO-VIP: VN-Protocol for CLNP

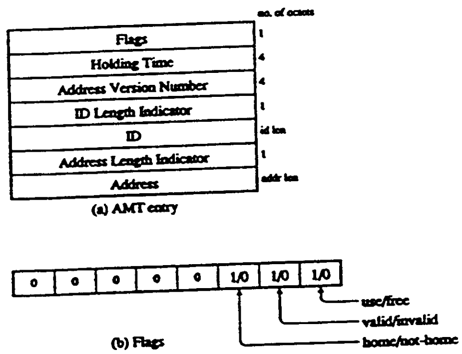

Figure 12: AMT entry for ISO-VIP

2.5.3 ISIS for IP Internets (isis)

Charter

Chair(s):

Ross Callon, rcallon@wellfleet.com
Chris Gunner, gunner@dsmail.enet.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	Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
Done	Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.
Done	Post a revision of the IS-IS as an Internet-Draft.
Mar 1993	Submit the revised IS-IS to the IESG as a Draft Standard.
Mar 1993	Submit the IS-IS MIB to the IESG as a Proposed Standard.

Internet-Drafts:

"Integrated IS-IS Management Information Base", 11/05/1991, Chris Gunner <draft-ietf-isis-mib-02.txt>

"Use of OSI IS-IS for Routing in TCP/IP and Multi-Protocol Environments", 01/11/1993, R. Callon <draft-ietf-isis-tcpip-00.txt, .ps>

Request For Comments:

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

CURRENT MEETING REPORT

Reported by Dino Farinacci/cisco

Minutes of the ISIS for IP Internets Working Group (ISIS)

Standards Progression (Integrated IS-IS)

- Proposed standard for over 2 years - what should we do?
- Write an interworking proof.
- Testing all features.

Radia Presented IS-IS Multicasting

- The scheme she describes was designed in DECNET Phase V but was taken out.
  - Uses multicast bit in the system-id field of an NSAP. Packet is routed to area and then distributed across a multicast tree in the area. (Multicast to whole area).

Radia Presented NLSP Differences From IS-IS

- Latch on to DR so changes are minimized.
- Radia proposes a configuration knob to have a priority increase after a router becomes DR.

Radia Presented the Multi-Area Router Concept

- A router has connection to two LANs that are in different areas. The router internally is two virtual routers.

Radia Presents How IS-IS Can Implement Multiple Levels of Hierarchy

- Nested circles can be configured to multiple levels of abstraction. This allows IS-IS to scale to larger routing domain sizes.

Discussed the Issue About a Maximum of 256 LSP Fragments

- Discussed increasing the size of the LSP fragment number from the LSPID.
- Discussion shifted to scaling issues related to injecting externals into an IGP.

Range of Metrics in IS-IS

- Increase metric to 16-bits for internal and 32-bits for external

IS-IS Over NBMA WANs

- Radia describes NSAP addresses that have E.164 IDPs where you can "get routing for free." This can be used for CLNP.
- For non-CLNP protocols, Radia wants to avoid SMDS multicast because of the coordination required with the phone company.
- Each node knows about a subset of the routers on the WANs. With transitive closure a DR is elected and responsible for knowing all nodes on the WAN. The non-DR routers flood LSPs to the DR only and the DR floods to the rest of the nodes.

Attendees

Michael Anello	mike@xlnt.com
Ross Callon	rcallon@wellfleet.com
Jeff Carman	tcarman@bnr.ca
John Chang	jrc@uswest.com
Richard Colella	colella@nist.gov
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Eric Fleischman	ericf@act.boeing.com
Eugene Geer	ewg@cc.bellcore.com
Chris Gunner	gunner@dsmail.enet.dec.com
Robert Hinden	hinden@eng.sun.com
Kathy Huber	khuber@bbn.com
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@server.af.mil
Tracy Mallory	tracym@3com.com
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
Radia Perlman	perlman@dsmail.enet.dec.com
Benny Rodrig	4373580@mcimail.com
Michael Safly	saf@tank1.msfc.nasa.gov
John Scudder	jgs@merit.edu
Kamlesh Tewani	ktt@arch2.att.com
Curtis Villamizar	curtis@ans.net

2.5.4 Inter-Domain Policy Routing (idpr)

<u>Charter</u>

Chair(s):

Martha Steenstrup, `msteenst@bbn.com`

Mailing Lists:

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

Description of Working Group:

The Inter-Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Done	Submit the IDPR Specification to the IESG as a Proposed Standard.

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, .ps>

"Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)", 07/22/1991, R.A. Woodburn <draft-ietf-idpr-mib-02.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 for two consecutive sessions during the March 1993 IETF meeting. We discussed the status of IDPR as a standard as well as experimental work that the Group is currently pursuing.

Two new IDPR Internet-Drafts were issued prior to the IETF meeting:

1. An updated version of the IDPR MIB, written by Woody Woodburn of Sparta.
2. A discussion of the DNS modifications to support IP address to administrative domain mappings, written by Rob Austein of Epilogue.

If you have not yet obtained a copy of these new Internet-Drafts, you are encouraged to do so. Please send all comments on the drafts to idpr-wg@bbn.com.

Internet Pilot Demonstration

The Internet pilot demonstration of IDPR is proceeding. We will demonstrate the functioning of IDPR in the presence of policies ("acceptable use policies") supplied by the transit networks NSFnet, NSInet, and TWBnet. No routers in any of these networks will actually run IDPR. Rather, special IDPR routers ("policy gateways" in the IDPR terminology) will act on behalf of NSFnet, NSInet, and TWBnet to supply policy routing. These policy gateways will only handle traffic designated as IDPR traffic. IDPR traffic will be generated by a small set of hosts at BBN, SRI, UCL, and ISI; no other Internet hosts will generate IDPR traffic.

Policy gateways will be located at BBN, SRI, UCL, and ISI and will act on behalf of these sites to handle IDPR traffic. Policy gateways will also be located at the FIXes and will act on behalf of NSFnet, NSInet, and TWBnet to handle IDPR traffic. There will be two SPARCstations installed at the FIXes, one SPARCstation per FIX Ethernet. Each SPARCstation will act as a set of three policy gateways, one policy gateway per participating transit network per FIX.

We will show that:

- IDPR selects routes that respect the access restrictions imposed by the transit networks and the service requirements (for example, low delay) of the sources.
- One may reconfigure transit network policies to suit current needs, and IDPR routes will reflect these changes.
- IDPR is robust in the presence of connectivity failures and quickly learns new routes.

- If a source attempts to setup a route that violates transit network access restrictions, the transit network refuses the route.

The policy gateways handle IDPR traffic only and will not touch other traffic. Hence, this pilot demonstration will not affect regular (non-IDPR) traffic traveling over the FIXes or over any of the transit networks attached to them.

Two years ago, the Group demonstrated this functionality in networks such as DARTnet; this will be the first demonstration of IDPR functionality in the general Internet.

Two of the more "experimental" topics the Group discussed included adding the super-domain functionality, described in the IDPR architecture, to the IDPR protocols and integrating resource reservation with IDPR.

To handle hierarchies of domains, IDPR must have a way to represent hierarchical domain addresses and to define the distribution scope of routing information in the hierarchy. Routing information from a given domain may be visible at multiple levels within the hierarchy if the distribution scope for that domain includes multiple levels.

"Routing context" is defined as the level in the domain hierarchy at which the routing will occur. For example, suppose domain A contains domain B which in turn contains a set of domains C1,...,Cn. Furthermore, suppose we want to route among the C domains. Then the routing context is represented as A/B. Thus, when we refer to Ci after defining the routing context, it is clear that we mean Ci contained in B contained in A, rather than Ci contained in some other domain Y. Routing context must be distributed in routing information messages as well as in path setup messages, to identify the context of the information contained within the message.

When integrating resource reservation and IDPR, there must exist mechanisms to generate routes that are likely to have the requested resources, to reserve resources, and to control traffic such that reservations are honored. IDPR relies on intra-domain mechanisms to support resource reservation and traffic control across a domain, and hence there must exist an interface for communicating reservation information to the intra-domain resource reservation mechanism.

IDPR domains supporting resource reservations should distribute the mean and standard deviation of the bandwidth available for reservation between relevant virtual gateway pairs, in their routing information messages. Route servers can select routes based on the mean values of available bandwidth or even on more pessimistic views such as available bandwidth equal to n standard deviations lower than the mean. Generating routes with such bandwidth metrics should increase the chances of producing routes that can in fact provide the requested bandwidth.

Using statistical measures of available bandwidth, a domain need not generate a routing information message for each resource reservation or release. Thus, the amount of routing information messages can be obtained. However, when bandwidth available for reservation is almost exhausted, a domain should immediately generate a routing information message indicating this state. This will prevent generation of new paths through this domain. Moreover, when the available bandwidth increases to a reasonable amount again, the domain should generate a routing information message to inform other domains that it has sufficient resources to accept more traffic. We have simulated such algorithms, but have not yet implemented them in an actual network.

Attendees

Robert Austein	sra@epilogue.com
David Bridgham	dab@epilogue.com
Al Costanzo	al@akc.com
Shane Dawalt	sdawalt@desire.wright.edu
Chip Elliott	celliot@bbn.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Frank Kastenholz	kasten@ftp.com
Tony Li	tli@cisco.com
Charles Lynn	clynn@bbn.com
Glenn Mackintosh	glenn@canet.ca
Brad Passwaters	bjp@sura.net
Manoel Rodrigues	manoel_rodrigues@att.com
Shawn Routhier	sar@epilogue.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Martha Steenstrup	msteenst@bbn.com
Stuart Stubblebine	stubblebine@isi.edu

2.5.5 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):
John Moy, jmoy@proteon.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 Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.

Done Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.

Done The Group 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 of the University of Maryland.

Done Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations.

Done Submit the MOSPF Specification to the IESG as a Proposed Standard.

Internet-Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-03.txt, .ps>

"IP Multicast over Token-Ring Local Area Networks", 11/09/1992, T. Pusateri <draft-pusateri-tokenring-lan-02.txt>

"MOSPF: Analysis and Experience", 04/16/1993, J. Moy <draft-ietf-mospf-analysis-01.txt>

CURRENT MEETING REPORT

Minutes of the Multicast Extensions to OSPF Working Group (MOSPF)

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

Attendees

Jim Beers	Jim.Beers@cornell.edu
Ed Brencovich	edb@dss.com
Jeff Carman	tcarman@bnr.ca
James Davin	davin@thumper.bellcore.com
Steve Deering	deering@parc.xerox.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Dennis Ferguson	dennis@ans.net
Ron Frederick	frederick@parc.xerox.com
Kenneth Giusti	kgiusti.chipcom.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Robert Hinden	hinden@eng.sun.com
Jeffrey Honig	Jeffrey_C_Honig@Cornell.edu
Steven Hubert	hubert@cac.washington.edu
Christian Huitema	christian.huitema@sophia.inria.fr
Ronald Jacoby	rj@sgi.com
Charley Kline	cvk@uiuc.edu
Robin Littlefield	rlittlef@wellfleet.com
Yu-Lin Lu	yulin@hpinddu.cup.hp.com
Charles Lynn	clynn@bbn.com
Bruce Mackey	brucem@cinops.xerox.com
Glenn Mackintosh	glenn@canet.ca
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Julie Myers	jmyers@network.com
Erik Nordmark	nordmark@eng.sun.com
Laura Pate	pate@gateway.mitre.org
Geir Pedersen	Geir.Pedersen@usit.uio.no
Maryann Perez	perez@cmf.nrl.navy.mil
Drew Perkins	ddp@fore.com
Narendra Popat	albin@frontier.com
Ben Robinson	ben_robinson@vnet.ibm.com
Benny Rodrig	4373580@mcimail.com
Kanan Shah	kshag@cmf.nrl.navy.mil

William Simpson	Bill.Simpson@um.cc.umich.edu
Terry Sullivan	terrys@newbridge.com
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net

2.5.6 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 interactions document.

4) IDRP for IP Usage document (standards track)

Most of the IDRP for IP Usage document 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.

Internet-Drafts:

"IDRP for SIP", 03/22/1993, S. Hares <draft-ietf-ipidrp-sip-00.txt>

CURRENT MEETING REPORT

Minutes of the OSI IDRP for IP over IP Working Group (IPIDRP)

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

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Nagaraj Arunkumar	nak@3com.com
Dennis Baker	dbaker@wellfleet.com
Tony Bates	tony@ripe.net
Jordan Becker	becker@ans.net
Rebecca Bostwick	bostwick@es.net
Ed Brencovich	edb@dss.com
Robert Brenner	Robert.W.Brenner@gte.sprint.com
Robert Calderon	calderon@noc.ans.net
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
Michael Collins	collinsms@es.net
David Conrad	davidc@iij.ad.jp
Osmund DeSouza	osmund.desouza@att.com
David Dubois	dad@pacersoft.com
Kishan Dudkikar	kishan@icm1.icp.net
Tom Easterday	tom@cic.net
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Michael Fidler	fidler@mitre.org
Peter Ford	peter@goshawk.lanl.gov
Vince Fuller	vaf@stanford.edu
Darren Griffiths	dag@ossi.com
Susan Hares	skh@merit.edu
Ittai Hershman	ittai@ans.net
David Jacobson	dnjake@vnet.ibm.com
Ronald Jacoby	rj@sgi.com
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Laurent Joncheray	lpj@merit.edu
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
John Krawczyk	jkrawczy@wellfleet.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Duane Kuang	duanek@kalpana.com
Tony Li	tli@cisco.com
Frank Liu	fcliu@pacbell.com

Peter Lothberg	roll@stupi.se
Glenn Mackintosh	glenn@canet.ca
Jamshid Mahdavi	Mahdavi@a.psi.edu
James Mahon	mahonj@honfsi1.att.com
Bill Manning	bmanning@sesqui.net
Jun Matsukata	jm@eng.isas.ac.jp
Daniel McRobb	dwm@noc.ans.net
David Meyer	meyer@ns.uoregon.edu
Gerry Meyer	gerry@spider.co.uk
Dennis Morris	morrisd@imo-uvax.disa.mil
Matthew Morrisey	morrisey@wpsp01.hq.aflc.af.mil
Julie Myers	jmyers@network.com
Peder Chr. Noergaard	pcn@tbit.dk
Erik Nordmark	nordmark@eng.sun.com
Brad Parker	brad@fcr.com
Andrew Partan	asp@uunet.uu.net
Laura Pate	pate@gateway.mitre.org
Michael Patton	map@bbn.com
Willi Porten	porten@gmd.de
Selina Priestley	sfp@noc.ans.net
Bala Rajagopalan	braja@qsun.att.com
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	amr@tycho.ncsc.mil
Ben Robinson	ben_robinson@vnet.ibm.com
Manoel Rodrigues	manoel_rodrigues@att.com
Michael Safly	saf@tank1.msfc.nasa.gov
John Scudder	jgs@merit.edu
Paul Serice	serice@cos.com
Kim Smith	kas@noc.ans.net
Steve Suzuki	suzu@fet.com
Steven Sweeney	steves@farallon.com
John Tavs	tavs@vnet.ibm.com
Marten Terpstra	marten@ripe.net
Kamlesh Tewani	ktt@arch2.att.com
Richard Thomas	rjthomas@bnr.ca
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Curtis Villamizar	curtis@ans.net
Hung Vu	hungv@fonorola.com
Scott Wasson	sgwasson@eng.xyplex.com
Jack Waters	waters@sura.net
Steven Willis	steve@wellfleet.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@barrnet.net
Jane Wojcik	jwojcik@bbn.com

Richard Woundy rwoundy@vnet.ibm.com
Paul Zawada Zawada@ncsa.uiuc.edu

2.5.7 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-02.txt>

"The OSPF NSSA Option", 10/13/1992, R. Coltun, V. Fuller <draft-ietf-ospf-nssa-option-00.txt>

"OSPF Version 2 Management Information Base", 11/03/1992, F. Baker, R. Coltun <draft-ietf-ospf-mib-00.txt>

"OSPF Version 2", 11/11/1992, J. Moy <draft-ietf-ospf-version2-02.txt, .ps>

"The OSPF External Attributes LSA", 03/23/1993, D. Ferguson <draft-ietf-ospf-extattr-00.txt>

"Guidelines for Running OSPF Over Frame Relay Networks", 05/03/1993, O. deSouza, M. Rodrigues <draft-ietf-ospf-guidelines-frn-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 on Tuesday, March 30th at the Columbus, IETF. The meeting began with a discussion of the six documents that the Working Group had in progress:

1. The latest draft of the OSPF V2 specification was reviewed. The only change that had been made since the last meeting was to handle the following case: A router becomes Designated Router, originates a network-LSA, then later restarts with a different Router ID and becomes Designated Router again. At this point, the router originates a network-LSA having the same Link State ID, but different Advertising Router, than the previous network-LSA. The fact that these two network-LSAs can both exist in the OSPF domain concurrently can confuse the Dijkstra calculation. Text has been added to the specification ensuring that the network-LSA originated before the last router restart will be flushed.

   After some discussion, it was decided to submit the latest draft (which had been published as an Internet-Draft) to replace RFC1247 at the current standards status (Draft Standard).

2. The OSPF Trap MIB had remained unchanged, except for editorial comments, since the last meeting. It was decided to submit this to be published as a Proposed Standard RFC.

3. The OSPF NSSA area specification was reviewed. A problem was found in the aggregation of multiple type-7 LSAs into a single type-5 LSA, involving the choice of metric. Two solutions were discussed:

   - Have the type-7 LSAs always take precedence, and
   - Set the metric to be the largest of any of the component metrics. Rob Coltun will investigate these options further.

4. Osmund deSouza presented a document describing how to run OSPF over Frame Relay. The document describes how to split the Frame relay PVCs into collections of OSPF point-to-point networks and NBMA networks. A comment was made that treating PVCs as unnumbered links was problematic, due to the inability to assign an ifIndex to individual PVCs. It was decided that after adding comments, the document will be submitted for publication as an Informational RFC.

5. No progress had been made on the OSPF MIB, which needs some additions before it can be republished.

6. No progress had been made on the "OSPF Database Over-flow" document.

The "OSPF for SIP" Internet-Draft written by Christian Huitema was summarized as: regular OSPF, running over IPv4, with two additional LSAs to import SIP information and an additional bit in the router-LSA to indicate SIP capability. This was intended to allow a more or less seamless migration from IPv4 to SIP, after which a native OSPF for SIP would be defined. Detailed discussion of the Draft was carried on in the SIP Working Group. In fact, it was decided that all detailed discussions of OSPF in IPv7 would be carried on in the appropriate IPv7 working groups.

Dennis Ferguson then presented an overview of his "OSPF external attributes" proposal, which is an addition to the OSPF<->BGP routing interchange and can be used as a substitute for Internal BGP. Dennis also presented operational statistics from the NSFnet that indicated his proposal would be quite efficient. Discussion indicated that efficiency, measured in terms of the percentage of the database dedicated to this scheme, would decrease when CIDR was deployed. Tony Li mentioned that it will also be necessary to know whether all routers participating in the "OSPF external attributes" are BGP-4 speakers, or whether some are BGP-3 speakers, in order to decide whether BGP aggregation should be done. Most of the discussion then centered on the problem that, since the external attributes (type-8 LSAs) can only be flooded through supporting OSPF routers, it is possible that the database of external attributes could get out of synch with the type-5 LSAs (which in turn could lead to problems in BGP routing). Dennis suggested three ways of dealing with this:

1. Choose the Link State IDs for type-8 LSAs in a random fashion, so that lack of synchronization would be obvious.

2. Potentially run Dijkstra a second time to ensure that there is type-8 flooding connectivity between BGP speakers, or

3. Change the document so that "most" routers must be capable of flooding type-8 LSAs. No decision was made on these options. Finally, it was noted that a combination of Internal BGP and the new "OSPF external attributes" cannot be run since the tag field in the OSPF type-5 LSAs would then have two conflicting requirements.

Tom Pusateri presented an outline of a proposed RIP to OSPF transition document, based on a talk he gave at INTEROP. He solicited suggestions for additional items to cover (send to pusateri@cs.duke.edu). Suggestions given at the meeting were:

1. Warning against running OSPF and RIP in parallel.
2. Give an example of a real, non-trivial network and how to transition it.
3. How to do address assignment.
4. How to decide what is in an area.

Lastly, a conflict between OSPF and Router Requirements was mentioned. Router Requirements states that you cannot follow the default route to get to subnets. This rule does not work for OSPF stub areas, and several people mentioned that the rule, while consistent with RIP, should not really apply to other protocols. A different rule, along the lines of: "when aggregating, create a discard route for the aggregate", was suggested. Philip Almquist, the editor of the Router Requirements documents, was present and participated in the discussion.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Dennis Baker	dbaker@wellfleet.com
Fred Baker	fbaker@acc.com
Jim Beers	Jim.Beers@cornell.edu
Nutan Behki	Nutan_Behki@qmail.newbridge.com
Richard Bjers	rich.bjers@uc.edu
John Boatright	bryan_boatright@ksc.nasa.gov
Robert Calderon	calderon@noc.ans.net
Douglas Carson	carson@utcc.utoronto.ca
James Cassell	jcassell@dsac.dla.mil
Rob Coltun	rcoltun@ni.umd.edu
David Conrad	davidc@iij.ad.jp
Wayne Cullen	wnc@netlink.com
Kurt Dobbins	kurtdob@ctron.com
Kishan Dudkikar	kishan@icm1.icp.net
Dennis Ferguson	dennis@ans.net
Paul Franchois	paulf@bldrdoc.gov
Christine Fredenburg	cfredenburg@dsac.dla.mil
Vince Fuller	vaf@stanford.edu
Darren Griffiths	dag@ossi.com
Patrick Hanel	hanel@yoyodyne.trs.ntc.nokia.com
Jeffrey Honig	Jeffrey_C_Honig@Cornell.edu
David Jacobson	dnjake@vnet.ibm.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
John Krawczyk	jkrawczy@wellfleet.com
Duane Kuang	duanek@kalpana.com
Tony Li	tli@cisco.com
Robin Littlefield	rlittlef@wellfleet.com
Glenn Mackintosh	glenn@canet.ca
Jamshid Mahdavi	Mahdavi@a.psi.edu
Glenn Mansfield	glenn@aic.co.jp
Jun Matsukata	jm@eng.isas.ac.jp
James Miner	jjm@fibercom.com
John Moy	jmoy@proteon.com
Julie Myers	jmyers@network.com
Shannon Nix	sdn@netlink.com
Zbigniew Opalka	zopalka@agile.com

Ayal Opher	aopher@synoptics.com
Joe Pagan	jrp@afterlife.ncsc.mil
Thomas Pusateri	pusateri@cs.duke.edu
Edward Reed	eer@cinops.xerox.com
Ben Robinson	ben_robinson@vnet.ibm.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Hal Sandick	sandick@vnet.ibm.com
Shiva Sawant	shiva@synoptics.com
Kanan Shah	kshag@cmf.nrl.navy.mil
Andrew Smith	asmith@synoptics.com
Martha Steenstrup	msteenst@bbn.com
Steve Suzuki	suzu@fet.com
John Tavs	tavs@vnet.ibm.com
Marek Tomaszewski	marek@net.com
Kannan Varadhan	kannan@oar.net
Linda Winkler	lwinkler@anl.gov
Jane Wojcik	jwojcik@bbn.com

2.5.8 RIP Version II (ripv2)

<u>Charter</u>

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:

RIP Version 2 and the Version 2 MIB was approved as a Proposed Standard in January 1993. They were published as RFC1388 and RFC1389. Since the mimimum required period has elapsed for a protocol to remain as a Proposed Standard, RIP V2 can now be considered for advancement to Draft Standard.

The RIP Version 2 Working Group will prepare a recommendation to the IESG evalating the standards track status of RIP Version 2 and the RIP Version 2 MIB. The recommendation will document implementation, interoperability and deployment experience as required by RFC1264 "Routing Protocol Criteria".

This Group is chartered to prepare revisions of RFC1388, RIP Version 2, RFC1389, the RIP Version 2 MIB, and RFC1387, analysis of the protocol if necessary.

The RIP Version 2 Working Group is further chartered to evaluate the proposal for "Routing over Demand Circuits using RIP" for standards track consideration.

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.
Jul 1993	Hold Working Group meetings to review RIP Version 2 implementations and make any changes needed to the specifications.
Jul 1993	Review the RIP over Demand Circuits Internet Draft.
Aug 1993	Submit the RIP over Demand Circuits to the IESG for consideration as a Proposed Standard.

Nov 1993 Post as an Internet-Draft a report describing the implementation and operational experience of the RIP v2 protocol in accordance with the RFC1264 "Routing Protocol Criteria."

Mar 1994 Submit the RIP Version 2 protocol to the IESG for consideration as a Draft Standard.

Request For Comments:

RFC 1387 "RIP Version 2 Protocol Analysis"

RFC 1388 "RIP Version 2 Carrying Additional Information"

RFC 1389 "RIP Version 2 MIB Extension"

2.5.9 Source Demand Routing (sdr)

Charter

Chair(s):

Deborah Estrin, estrin@isi.edu
Tony Li, tli@cisco.com

Mailing Lists:

General Discussion: sdrp@caldera.usc.edu
To Subscribe: sdrp-request@caldera.usc.edu
Archive: jerico.usc.edu:~/pub/sdrp

Description of Working Group:

The SDR Working Group is chartered to specify and promote the use of SDR (Source Demand Routing Protocol) as an interdomain routing protocol capability in conjunction with IDRP and BGP interdomain routing protocols. The purpose of SDR is to support source-initiated selection of interdomain routes, to complement the intermediate node selection provided by BGP/IDRP.

The goal of the SDR Working Group is to release the components of SDR as IETF prototypes and to obtain operational experience with SDR in the Internet. Once there is enough experience with SDR the Working Group will submit the SDR components to the IESG for standardization.

SDR has four components: Packet formats for protocol control messages and encapsulation of user datagrams, Processing and forwarding of user data and control messages, Routing information distribution/collection and route computation, Configuration and usage.

The Group's strategy is to:

1. Define the format, processing and forwarding of user datagram and control messages so that SDR can be used very early on as an efficient means of supporting "configured" inter-domain routes. User packets are encapsulated along with the source route and forwarded along the "configured" route. Routes are static at the inter-domain level, but are not static in terms of the intra-domain paths that packets will take between specified points in the SDR route. The impact of encapsulation on MTU, ICMP, performance, etc., are among the issues that must be evaluated before deployment.

2. Develop simple schemes for a) collecting dynamic domain-level connectivity information, and b) route construction based on this information, so that those domains that want to can make use of a richer, and dynamic set of SDR routes.

3. In parallel with 1 and 2, develop usage and configuration documents and prototypes that demonstrate the utility of static-SDR and simple-dynamic-SDR.

4. After gaining some experience with the simple schemes for distribution, develop a second generation of information distribution and route construction schemes. The Group hopes to benefit from discussions with IDPR and NIMROD developers at this future stage because the issues faced are similar.

5. The Group will also investigate the addition of security options into the SDRP forwarding and packet format specifications.

Goals and Milestones:

Mar 1993 Post an Internet-Draft of packet forwarding and control message format and protocol for IP.

Jun 1993 Post as an Internet-Draft the SDR MIB.

Jun 1993 Post as an Internet-Draft the SDR Usage and Configuration document. This is the highest priority after the draft specification in order to demonstrate how even static-SDR can be used to achieve concrete objectives.

Sep 1993 Post as an Internet-Draft the BGP/IDRP Extensions Specification. As mentioned in the Internet Draft there are a few extensions to BGP/IDRP needed to support SDR. These must be detailed and documented.

Sep 1993 Submit as an Internet-Draft a specification for Route Setup.

Nov 1993 Post as an Internet-Draft a SDR Deployment Plan.

Dec 1993 Post as an Internet-Draft a document describing the distribution/acquisition of Information to construct richer SDR routes. The initial versions of SDR will use only configured information (some of which may be derived from BGP/IDRP) as the basis for constructing source routes.

Dec 1993 Post as an Internet-Draft a specification for SDR Multicast.

Mar 1994 Submit the set of SDR specifiations to the IESG for consideration as a Proposed Standard.

Mar 1994 Submit the set of SDR specifications to the IESG for consideration as a Prototype protocol.

CURRENT MEETING REPORT

Reported by Deborah Estrin/USC and Tony Li/cisco

Minutes of the Source Demand Routing Working Group (SDR)

Changes to the Specification

Changes to the specification were presented and discussed. Major modifications were made to support interior SDRP. The new packet header format was presented. All packets now carry a hop count, which formerly was only in data packets. All packets now carry target router and notification fields, even though only control packets use them. Notification uses a byte which would be necessary for alignment anyhow, so this causes four bytes of overhead on data packets. The source route length is now the number of IP addresses, not the number of bytes. The next hop pointer also now is in terms of addresses. The source route now supports interior routing due to the need expressed at the previous SDRP BOF for source demand routing within domains.

Source routes now contain three types of entries, all of which are syntactically IP addresses. An entry may be a normal IP address, or an AS number, or a change in source route attributes. An AS number is encoded in the low order two octets of network 128.0.0.0. Changed source route attributes are encoded in the low order three octets of 127.0.0.0. Currently, the only change possible is to change the strict/loose source route bit. This accommodates source routes which need a mix of strict and loose source routing.

There are changes to forwarding to match the new source route format. If the address in the source route that is currently being processed is a normal IP address, then forwarding checks to see if it matches the local address and if so, looks at the next address in the source route. Otherwise the packet is forwarded to the indicated address using normal IP forwarding. If the address in the source route encodes an AS number that matches the local AS#, then forwarding looks at the next entry in the source route; otherwise the packet is forwarded to the indicated AS looking at D-FIB. If the address in the source route encodes a change in attribute type, then the SDRP speaker reaches in and sets the attribute bit accordingly and looks at the next source route entry for processing.

SDRP Overview

A brief SDRP overview was presented for new folks; see the BOF Minutes from the previous IETF or the Unified Architecture document for background.

SDRP Usage Document

The Group discussed a draft of the SDRP usage document distributed before the IETF.

SDRP can be used in the near-term to provide special routes that compliment existing IGP and BGP/IDRP routing. SDRP can be phased into the operational Internet without wholesale replacement of routing.

At the same time as the Group is proceeding with protocol specifications for nearer-term experimentation, longer-term issues are already under consideration. To provide a sense of "where we are headed" with this protocol and the Unified Architecture in general, a companion document on SDRP futures has also been drafted.

In the packet format and forwarding protocol specification it is not specified how an SDRP router that originates an SDRP packet acquires an SDRP route. An SDRP route is defined as a sequence of domain identifiers and/or IP addresses, or a combination; the route may be strict or loose.

The usage document should discuss mechanisms for acquiring SDRP routes using EXISTING routing information distribution mechanisms (BGP/IDRP). In particular, it will cover the following three sources of routes:

1. BGP/IDRP routes
2. Manually configured routes
3. Route fragments

Any legal BGP or IDRP route is, by definition, a legal SDRP route, so long as there are SDRP speakers at appropriate points along the path.

Every BGP/IDRP speaker may maintain information about multiple feasible routes to a destination (routes advertised by different neighbors). But a BGP speaker chooses at most one route to be active (selected), and an IDRP speaker may choose more than one route to be active (selected) only if all selected routes have different "distinguishing attributes". As a result, the currently active (selected) IDRP/BGP route may not be appropriate for the packet.

One of the simplest forms of SDRP route acquisition is to select among the alternative routes advertised by the node's neighbors. This requires NO modifications to BGP/IDRP. It does require development of appropriate route selection rules, both manual and semi-automated, for selecting particular BGP/IDRP routes to be used as SDRP routes.

Network administrators can also create SDRP routes by examination of network topology BGP/IDRP databases, or manually collecting network information through active probing (traceroute).

The operational status of routes can be determined dynamically using the passive and active mechanisms defined in SDRP packet forwarding, allowing the scheme to adapt to topological changes.

For the usage document, examples of useful manual configurations need to be given. It must be emphasized that PROBE needs to be used to detect black hole routes and the utility of having several SDRP routes as fallback routes to somewhat make up for the fact that these will be "static" due to manual configuration.

Route fragments from different BGP/IDRP routes can be used, in part or whole, to create desired SDRP routes that do not appear in the node's neighbors' BGP/IDRP tables. This allows the administrator to "cut and paste" to create new routes.

If SDRP is used within a domain, an IGP route can be used as an SDRP routes.

Additional information derived from IGP can also be used to construct routes, e.g., the OSPF link state database for reachability within the OSPF system.

Interior SDRP is an area that in particular needs further discussion and development of a usage model. For example, there is a need to:

- Clarify how you get information about exit points into the interior.
- Investigate the use of information that OSPF and ISIS carry already.
- Consider adding the ability to query BGP speakers internally.

Another mechanism not given in the specification is how a source host's SDRP-speaking border router maps a particular packet to a particular SDRP route. This is not part of the protocol specification because it can be left to local control; we need not be coordinated across the Internet, or even across the set of routers on a single path. However, to use SDRP, the network administrator must be able to configure the information used to map host-generated payload packets to appropriate routes, therefore it must be addressed in the usage document. The mapping indicates whether a packet can be sent out using the BGP/IDRP route; and if not, which available SDRP route can be used (if any).

A domain may choose any mapping function that is unambiguous and whose input information can be found in the payload packet or locally to the router (e.g., based upon incoming interface); but may "pay for" more sophisticated mappings.

Good examples need to be developed for the usage documents as well as clarification on where the mapping/classification is done and note the tradeoffs between doing it closer to the host and at the border router.

BGP/IDRP and SDRP routes have transit policy qualifications associated with them. The syntax and semantics of SDRP policies should be consistent with transit IDRP/BGP policies. The Group should probably proceed by initially using the existing BGP/IDRP policy semantics and syntax and evaluating the need for extensions after gaining some experience.

For the next IETF the Group will review the IDRP policy language and identify if there are unmet needs for SDRP.

In the current specification, the Working Group disclaims any attempt to provide secure/verifiable enforcement of transit policies. The essential tools needed for this security service are more a function of the authentication and integrity mechanisms available in the protocol providing delivery service for SDRP, than of SDRP itself. However, transit policy conformance can be audited by sampling data to identify violators. Spot checks can be effective and are used in many other kinds of systems (computerized and manual). Auditing procedures and sampling rules are a subject for local control and may vary across different SDRP routes. It would be useful to develop some examples for the usage document.

The only planned modification of BGP/IDRP is an optional attribute indicating that a particular domain supports SDRP and optionally specifies address(es) of SDRP speaker(s) in the domain. This is important for route selection and forwarding decisions. There are two proposals for this function so far and the arguments for and against will be discussed shortly on the SDRP mailing list.

SDRP supports interior policy routing by allowing SDRP routes to carry IP addresses. This can be used to direct traffic via configured paths in the source domain. It can also be used to direct routing of packets within other domains; for example, by specifying a particular exit router for a transit domain. Particular routes within the destination domain can also be specified; but this requires detailed knowledge of the topology and addressing of other domains which requires mutual agreements for information update between domain administrators.

The possible use of OSPF and ISIS information and the implications of attempting to use this (or not) with other interior routing protocols such as RIP, or IGRP should be discussed in the usage document. The use of IBGP for this purpose should also be documented.

The Unified Architecture is designed to allow evolution. SDR was also designed to allow innovation without global coordination. The Group is working to specify parts of the protocol that could be implemented and used in the short-term such that they will interwork with other parts of the architecture still under development. In particular, the packet format and forwarding protocol have been specified while details of SDR route computation are still under development.

Mechanisms for route computation and even information distribution/collection can be changed more readily than packet forwarding mechanisms because route computation is a local matter. Information distribution concerns some subset of routers or domains whereas packet forwarding procedure must be agreed upon by all routers that implement SDRP.

Important but evolving aspects of the architecture include:

- Route construction.
- Policy language.
- Route setup.
- Multicast routing.
- Alternate path routing for reservation-oriented (virtual circuit) traffic.

The Group wants to extend route construction mechanisms to obtain routes that conform to source-specific policies where a route's use is restricted to certain sources, or QOS requirements where a route supports a particular performance or policy related QOS (color), and/or path-constraint policies where a route must pass through or avoid particular transit domain(s).

Routes available via IDRP are the result of path selection processes in all the intermediate IDRP speakers between the source and destination. Mechanisms are needed for the source to obtain information about other routes that it is allowed to use but that intermediate domains filtered out as a result of their path preferences.

Different approaches can be characterized to route construction according to whether construction is based on distributed or centralized processing. For example: using an IDRP route is a form of distributed processing since the route is constructed hop-by-hop by nodes on a path. Collecting inter-domain topology/policy information from around the network and computing a route at the source is a form of centralized processing. Route fragments represent intermediate points where the source centrally controls the acquisition and concatenation of fragments, but the fragments themselves represent the result of a distributed computation.

Query is one example mechanism where a source domain SDRP speaker queries its immediate neighbor IDRP domains to get all available routes to a particular destination (possibly with QOS specified as well).

The SDRP speaker could also query non-neighbor IDRP speakers; but this raises the question of heuristics for deciding whom to query, which is still a subject for further research.

Query is an example of centralized processing and can also be used to obtain route fragments.

The Extract mechanism is a second proposed mechanism for on-demand SDRP route acquisition. For example, the source could send an extract request to the destination indicating desired QOS and possibly exclusionary transit information (e.g., what transit it does NOT want to use). The destination would then cause IDRP to propagate back routes that fit the characteristics specified by the source. The routes would NOT be stored in the RIBs en route back to the source; rather the information would be passed along on an FYI basis.

Extract is an example of distributed processing and could also be extended to send extract requests to a preferred transit domain for it to initiate the extract. Extract could also be used to obtain route fragments. The big question is how to constrain the propagation of the return information; hop-count limits, limits on the number of routes propagated by each domain are possibilities, each of which trades off overhead for some loss of information.

Other schemes for collecting information and computing routes are the subject of ongoing research. However, the combination of extract, query, and route fragment mechanisms may be adequate to meet most needs; this needs further study.

A common language is needed for specifying policy constraints on all routes. This would allow other domains to do policy computations to determine feasible routes. The language must be extensible. For example, in response to a policy query, a domain may respond with its policy configuration. The policy language would look like a boolean expression; and policy computation would consist of evaluating this expression. Syntactically, the expression appears as a series of terms; satisfying any term satisfies the expression. Possible variables include:

- QOS of the packet.
- Source domain.
- Source address.
- Destination domain.
- Destination address.
- Transport protocol.
- Application protocol.
- Time of day.
- Inter-domain path in use.

Terms that contain unrecognized variables would be ignored.

The initial specification for packet format and forwarding includes a full SDRP route in every packet sent.

When the duration of a packet stream is significantly longer than the end-to-end delay, and if the payload in the packets is small, it is worth establishing state information in SDRP speakers along route, instead of carrying a full SDRP route in every packet, i.e., "setup". Once state is established, the source can rely on a route identifier in each packet and thereby reduce SDRP packet header size and processing time. However in designing a setup protocol it is important to not IMPOSE setup on all SDRP speakers (might be short on state space or might not otherwise wish to support setup).

Strawman Proposal

A strawman proposal for setup operations was presented.

SDRP multicast would coexist and interoperate with IDRP/BGP multicast routing mechanisms. The Group anticipates more than the single IP multicast routing model currently used in the Internet. IDRP may be used for setting up the multicast distribution trees (or branches thereof) when the generic routes satisfy the requirements of the application and group (i.e., QOS). In particular there will be complementary mechanisms that are more efficient than DVMPR or MOSPF style multicast for supporting sparse multicast groups. Both IDRP and SDRP will be used to support these mechanisms.

SDRP would set up multicast distribution trees (or branches thereof) when the generic routes do not meet the needs of the application and group.

SDRP can be used to support alternate path routing for reservation (or more generally virtual circuit) traffic. Source routing is good for achieving alternate path routing because it has inherent loop avoidance and it avoids placing burden on intermediate switches to compute and retain multiple routes to each destination.

Alternate path routing is particularly important for reservation traffic where a call setup request may be rejected due to insufficient resources at some intermediate switch/link as a result of heavy utilization. In this case, the source would like to attempt alternate routes that do not go through the bottleneck link. SDRP can provide a source with alternate, loop free, routes; particularly appropriate when SDRP setup is used. A recent Internet-Draft by Rob Coltun and Marco Sosa also concluded that source routing is the best means of achieving alternate path routing for virtual circuit routing.

Given that a route must have sufficient resources to accommodate a reservation flow (i.e., stream, call), it might be useful for the source to maintain recently measured load levels on those links in the network that it uses frequently; for example from those links used by active flows. There are open research issues to resolve in the inter-domain case where detailed information of remote domains is not available.

Because SDRP can be used to support interior routing, SDRP could be used for alternate path routing within areas of a domain and within domains.

Initially, it may be simplest to have the source try to use an alternate domain level route when a reject is received from a remote domain; this may be justified if one assumes that the hop-by-hop routing choice used in that domain to traverse the domain does reflect long-term utilization in that domain.

There is much more to be said on all of these subjects.

Projects and Milestones

Projects and milestones were discussed. The following is a list of topics to be discussed and people interested in working on them.

Usage Document	(Draft before July IETF.) Deborah Estrin, Yakov Rekhter and Peter Ford.
BGP/IDRP Attributes Draft	(Draft by May.) Tony Li, Yakov Rekhter and Deborah Estrin (referee).
Prototype	(Working prototype for others to see by June.) Daniel Zappala and Tony Li will look it over.
Setup Specification	(Draft before July IETF.) Deborah Estrin, Tony Li and Osmund deSouza.

Info. Distribution/Route Selection	(Draft description of Extract Mechanism (not specification) and more detailed plan for how to proceed in short and mid-term by July IETF.) Tony Li, Steve Hotz, David Bridgam dab@epilogue.com, Yakov Rekhter and Brijesh Kumar.
Policy Language	(Presentation and discussion at July IETF; draft document for November.) Tony Li, David Karrenberg, Peter Lothberg, Steve Hotz, Sue Hares and Steve Willis.
Multicasting	(Possible draft for November.) Deborah Estrin and Osmund deSouza.
Use of SDRP for Adaptive Routing	(Discuss at July or November IETF; In the mean time discuss with VCROUTE BOF.) Deborah Estrin and Daniel Zappala.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Michael Anello	mike@xlnt.com
Tony Bates	tony@ripe.net
David Bolen	db31@ans.net
Ed Brencovich	edb@dss.com
David Bridgham	dab@epilogue.com
Jeff Carman	tcarman@bnr.ca
James Cassell	jcassell@dsac.dla.mil
David Conklin	conklin@jvnc.net
Dave Cullerot	cullerot@ctron.com
Tony DeSimone	tds@hoserve.att.com
Osmund DeSouza	osmund.desouza@att.com
Kurt Dobbins	kurtdob@ctron.com
Kishan Dudkikar	kishan@icm1.icp.net
Deborah Estrin	estrin@isi.edu
Peter Ford	peter@goshawk.lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Robert Hinden	hinden@eng.sun.com
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@server.af.mil
Merike Kaeo	merike@alw.nih.gov
Daniel Karrenberg	daniel@ripe.net
Padma Krishnaswamy	kri@sabre.bellcore.com
Giri Kuthethoor	giri@ms.uky.edu
Mark Laubach	laubach@hpl.hp.com

Tony Li	tli@cisco.com
Kim Long	klong@sura.net
Charles Lynn	clynn@bbn.com
Glenn Mansfield	glenn@aic.co.jp
Jun Matsukata	jm@eng.isas.ac.jp
David Meyer	meyer@ns.uoregon.edu
Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.disa.mil
Peder Chr. Noergaard	pcn@tbit.dk
Erik Nordmark	nordmark@eng.sun.com
Zbigniew Opalka	zopalka@agile.com
Laura Pate	pate@gateway.mitre.org
John Penners	jpenners@advtech.uswest.com
Maryann Perez	perez@cmf.nrl.navy.mil
Yakov Rekhter	yakov@watson.ibm.com
Robert Reschly	reschly@brl.mil
Ben Robinson	ben_robinson@vnet.ibm.com
Manoel Rodrigues	manoel_rodrigues@att.com
Shawn Routhier	sar@epilogue.com
Hal Sandick	sandick@vnet.ibm.com
Vilson Sarto	vilson@fapq.fapesp.br
John Scudder	jgs@merit.edu
Kanan Shah	kshag@cmf.nrl.navy.mil
Andrew Smith	asmith@synoptics.com
Martha Steenstrup	msteenst@bbn.com
Bernhard Stockman	boss@ebone.net
Terry Sullivan	terrys@newbridge.com
John Tavs	tavs@vnet.ibm.com
Marten Terpstra	marten@ripe.net
Kannan Varadhan	kannan@oar.net
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Steven Willis	steve@wellfleet.com
Richard Woundy	rwoundy@vnet.ibm.com

2.6 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker/TIS and Jim Galvin/TIS

The Security Area within the IETF is responsible for development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet.

Much of the work of the Security Area is performed in coordination with working groups in other areas. The Security Area Advisory Group (SAAG) is a group of security experts which provides both consulting help to other areas and direct management of working groups within the security area.

The main bulk of the work for the SAAG consists of a set of formal work items. These work items correspond to working groups within the IETF Security Area, security relevant developments within working groups in areas other than security, and internal SAAG work items which do not merit the creation of formal working groups but which do need some level of attention.

Below is the status of each of the Working Groups and/or BOFs officially chartered or initiated within the Security Area. Immediately following those reports is an update on other security issues as well as security related work in other IETF areas.

Authorization and Access Control BOF (AAC)

A Charter has been submitted to the IESG. Its official ratification is waiting for a statement indicating its relationship to other security related activities in the IETF.

The Authorization and Access Control BOF met on Wednesday afternoon. Common characteristics of several distributed authorization mechanisms were discussed. The Group will compile a common list of restrictions and/or privilege attributes sufficient to support DCE, ECMA/Sesame, and restricted proxies, as well as the needs of applications. The specification for an authorization API was refined with the form of several arguments defined, and others sketched. Work items were assigned to further refine these definitions and to specify the form of access control list entries themselves.

Internet Protocol Security Protocol BOF (IPSEC)

A Charter has been submitted to the IESG. Its official ratification is waiting for a statement indicating its relationship to other security related activities in the IETF.

A review of initial experimental implementations was conducted. A preliminary list of IPSEC protocol features/requirements was discussed and will be posted to the mailing list. There was a brief discussion of key management issues but it was deferred to be conducted on the mailing list.

Common Internet Protocol Security Option Working Group (CIPSO)

The CIPSO Working Group meets principally under the auspices of the Trusted Systems Interoperability Group. A revised Internet-Draft was posted for discussion at the Columbus IETF meeting. A few changes were discussed, that were primarily structural with some additions to provide more detail.

The majority of the Working Group believes its work is done. Steve Crocker will coordinate a team of experts to review the current specification prior to its submission to the IESG for publication as a Proposed Standard.

Common Authentication Technology Working Group (CAT)

The GSS-API base specification, GSS-API C Language Bindings, and Kerberos Version 5 documents are to be submitted for consideration as Proposed Standards.

The DASS document is to be submitted for consideration as an Experimental Protocol.

The CAT Working Group met for two sessions at the Columbus IETF. The primary agenda item was integration of security features into FTP, a topic for which Sam Sjogren is acting as task leader and on which Steve Lunt has generated a working document shortly to be released as an Internet-Draft. The FTP security discussions were quite fruitful, both in terms of providing feedback for improving the draft proposal for FTP as well as fine tuning the GSS-API requirements and specifications.

Privacy Enhanced Mail Working Group (PEM)

The PEM specifications have been published as RFCs 1421, 1422, 1423, and 1424. This work item was officially closed at the Columbus IETF meeting.

SNMP Security Working Group (SNMPSEC)

In conjunction with the SNMPv2 Working Group, twelve documents have been completed and adopted by the IESG as Proposed Standards. They are currently in the hands of the RFC editor for processing for publication.

By agreement with the new Network Management Area Director, Marshall Rose, further work on SNMP security will be carried within the existing SNMP Working Group with assistance provided by the Security Area.

TCP Client Identity Protocol Working Group (IDENT)

The protocol specification has been published in RFC 1413 as a Proposed Standard. A network management MIB document was published in parallel as RFC1414. Using this MIB, a SNMP client can ascertain the same information that an Indent client can, thereby giving clients two options for implementing this service. This work item was officially closed at the Columbus IETF meeting.

OSI Directory Services Working Group (OSIDS) - Applications

There is no security activity in this area at this time. This work item was officially closed at the Columbus IETF meeting.

TELNET Working Group (TELNET) - Applications

A document specifying a combination authentication-encryption option was discussed, including replacing the individual option documents with this one document. A revised Internet-Draft will be posted.

A Kerberos version 5 sub-option document was also discussed. A revised Internet-Draft will be posted.

Router Requirements Working Group (RREQ) - Internet

The previous single document has been split into four documents and a number of auxiliary documents. Philip Almquist has responsibility for finishing the documents and submitting them to the IESG for publication.

Mobile IP Security Working Group (MOBILEIP) - Routing

If there existed an IP security option Mobile IP would not have to create its own. This raises the question of what the relationship between this security work item and the IP security work item is. This will be addressed in a document to be posted to Internet-Drafts.

Audio/Video Transport Working Group (AVT) - Transport

This activity will be reviewed to identify the security issues for the Amsterdam meeting.

Domain Name System Working Group (DNS) - Transport

A subcommittee will be created to deal with security issues. A mailing list will be created for use by the subcommittee.

Trusted Network File System Working Group (TNFS) - Transport

The TNFS Working Group meets principally under the auspices of the Trusted Systems Interoperability Group.

No progress to report.

Integrated Directory Services Working Group (IDS) - User Services

This activity will be reviewed to identify the security issues for the Amsterdam meeting.

Export Control Issues

Vint Cerf and Steve Crocker need to press forward on drafting a document.

IP: The Next Generation

A plan for processing a security review of the competing next generation proposals will be drafted for the Amsterdam meeting.

ITAR Publication

An on-line version of the U.S. International Traffic in Arms Regulations (ITAR) will be created. In addition, it was noted that the ISSA published a summary of U.S. export law that would be useful to include.

Key Management Strategies

A review of key management strategies and activities will be drafted for the Amsterdam meeting.

Network Database Privacy

There is no activity in this area. This work item was officially closed at the Columbus IETF meeting.

PEM and MIME Integration

The meeting began with discussions of implementation status and deployment strategies. There will soon be PEM implementations available in the UK and Germany as a result of work under the EC PASSPORT program. Interoperability testing is in progress. In support of the Internet certification hierarchy RSADSI and TIS announced the availability of PCAs.

In addition to the PEM and MIME integration, the use of email addresses in distinguished names and the relaxation of the trust model for the current hierarchy were discussed, but no consensus was reached. The PEM and MIME integration was also not settled since there was a fair amount of disagreement about the issues. A revised Internet-Draft will be posted.

Random Number Generation Issues

A document has been posted as an Internet-Draft that identifies the issues to be concerned about when generating random numbers. However, the document does not have a conclusion on how to generate random numbers given a set of requirements. A revision will be prepared.

Routing Security Plan

Radia Perlman will submit a brief white paper identifying the issues.

Security Area Architecture

A short description of the relationship between the IETF security activities will be drafted for the Amsterdam meeting.

Working Group Liaison Checklist

A checklist for use by security liaisons to working groups that will assist in tracking progress will be drafted for the Amsterdam IETF.

CURRENT MEETING REPORT

Reported by Clifford Neuman/Information Sciences Institute

Minutes of the Authorization and Access Control BOF (AAC)

The third meeting of the Authorization and Access Control BOF was held at the March IETF. This should be the last BOF meeting as the Charter has already been submitted to the IESG and is currently under review.

Agenda

- Review the AAC Charter as submitted to the IESG for review.
- Identify common characteristics of evolving distributed authorization mechanisms and begin discussion on how best to encourage interoperability across mechanisms.
- Review the authorization API from the last meeting, in light of the DCE and Sesame documents made available by Piers McMahon, and consider characteristics as identified in the preceding item.
- Assign work items to generate an access control API Internet-Draft and assign other items to be accomplished before the Amsterdam IETF.

The draft Charter, past Minutes, mailing list discussions, and other documents mentioned in these Minutes are available by anonymous FTP from prospero.isi.edu in the directory /pub/aac.

Review of Charter

The Working Group Charter, discussed at the previous meeting, was submitted to the IESG for approval. Based on several comments from IESG members, the Charter was revised to make the relationship with other working groups clearer. A new version was provided to the IESG and distributed to the AAC mailing list. Everyone at the meeting seemed comfortable with the revised Charter and no additional changes were suggested. The Charter is once again in the hands of the IESG, which seems reluctant to form new groups without understanding the "big picture" of how all the groups in an area fit together. Steve Crocker, the Security Area Director, will provide the IESG with such a statement.

Evolving Distributed Authorization Mechanisms

The goals of distributed authorization mechanisms include avoiding the need to duplicate common information about principals' access rights on each end-system and facilitating ease of management, often by enabling centralized maintenance of access control and revocation information. Evolving mechanisms for distributed authorization were discussed, including DCE, ECMA/Sesame, and restricted proxies (Kerberos V5). In each of these mechanisms, a certificate grants privileges to a principal, either by explicitly enumerating privileges, or by restricting privileges. Positive privileges are represented in the latter case as the set of rights available to the issuer of the certificate which are not explicitly restricted.

Clifford Neuman argued that the restricted forms of credentials are more general since they allow the addition of restrictions during delegation, and also because the rights of the issuer of a certificate naturally limit what can be granted, rather than requiring a separately maintained list of which principals are authorized to grant particular privileges. Piers McMahon and John Linn suggested that the term restriction is a confusing term for permission. It also wasn't clear to them how one could compose privileges from different sources. Cliff responded that the permissions are granted by certificates which provide all rights possessed by the signer and that the restrictions apply to what is granted by a particular certificate, limiting the permissions. By providing multiple certificates one gets a sum of products; i.e., the sum of the remaining rights from each of the restricted certificates. There was some skepticism about this role of restrictions in the representation of what appear to be positive rights; the issue will be revisited later.

It would improve interoperability across mechanisms if there were a common set of restrictions or privileges. This would make it easier to translate from one mechanism to another (i.e., the translation would be syntactic, rather than semantic). Clifford Neuman will compile a list of restrictions sufficient to represent privileges supported by DCE, Sesame, restricted proxies, and any other mechanisms brought to the Group's attention. He will also accept input from application developers on the restrictions and privileges they require. Submissions can be in either form, restrictions or permissions. They will be represented at least initially as restrictions in the compiled list to asses the benefits of such a representation.

Example permissions/restrictions:

DCE	A way of specifying global userID and groups the user is in.
DCE	"Optional" restrictions.
NAS	A way of restricting the subnets to which a user may connect.
Sesame	A way to represent each of its privilege attributes.
Others	Roles.

Authorization API (and more)

The Working Group began refining the strawman ACL API presented at the Washington D.C. IETF. There is more than just an API involved. In particular, it will be necessary to define the authorization information that can be stored. The API will be designed to allow easy integration with distributed authorization methods, including support for the specification of restrictions on ACL entries. Ted Ts'o suggested that since the API is not restricted to what people normally think of as ACLs, that it should be called an Authorization API, rather than an ACL API. This change in terminology was adopted.

The Group started to define the input and output arguments to the API. To set the stage for this discussion, Piers McMahon and Clifford Neuman described the anticipated flow of control for the end-server. When a request is received, calls are made to the GSSAPI or other authentication routines. The output of the GSSAPI is a Security Context which contains information about the client. This might then be passed into an additional function to verify authorization credentials. The distributed authorization mechanisms might instead be handled as part of the GSSAPI itself.

The output of the combined GSSAPI and distributed authorization functions will be a Security Context which is fed into the authorization API. For simplicity at this point in the discussion, the authorization API will return a yes/no answer indicating whether a particular operation is allowed. We need to define what needs to be part of the Security Context, an input to the authorization API. It might then be possible to extend the GSSAPI Security Context definition so that it can be used directly. The diagram below shows the flow of control.

```
                          This is the Security Context to be defined
                                 /
-----+  +---------+  +-------X---------+ +-------> yes/no+restrictions
     |  |         |  |                 | |
     V  |         V  |                 V |
   GSSAPI        Dist             Authorization
            Authorization              API
+----possibly combined----+
```

The authorization API will provide mechanisms to query a local authorization database to check authorization.

answer = check_authorization(sc, target, operation)

SC is the Security Context, containing information about the identity of the client and about authorization credentials that have been received (and possibly verified). Piers McMahon agreed to work on defining the form of the Security Context for our present needs, trying to define it in such a way that it will later be able to accommodate information from additional distributed authorization mechanisms. Defining the Security Context will require interaction with those defining other parts of the API, and should include consultation with John Linn and the CAT Working Group to make sure that it is consistent with, and perhaps ultimately included in, the GSSAPI.

TARGET specifies the target of the attempted access. If the authorization database were represented as a set of access control lists, the target would identify which list is to be consulted. TARGET should have a two part name, one part identifying an application or set of applications that share the same set of objects, and the name of the individual object. This is useful when the authorization database is maintained on a system wide basis (for example by an ACL manager) rather than by individual applications. For now we will define TARGET to be two strings, an identifier for an application and an application specific object name. This can be represented as a single string using a ":" as a separator (this means : can not appear in the identifier of the application).

OPERATION identifies the operation to be performed on the target. The definition of this field will correspond to the rights field in the ACLs in the authorization database. The Group decided that this field would be an extensible bit vector of application specific rights, together with a tag identifying the application (the tag for the operation bit vector should be compared with the tag for the bit vector in the ACL to make sure the bits are interpreted correctly).

At this point it was felt that the Group needed to specify the form of the authorization database. The access control list method was proposed, with rights specified as described in the preceding paragraph, but with the ACL entry extended so entries would include an optional list of restrictions that further specify the conditions under which the operation would be allowed. Among these restrictions might be time of day restrictions, as well as application specific restrictions such as a netmask of network access servers, or perhaps a restriction for a mailing list server that allowed modification of a mailing list, but only to add or delete oneself. These restrictions would exist in the same space as those defined for distributed authorization mechanisms. In fact, it should be possible to apply the restrictions returned from an authorization database directly to distributed authorization credentials issued by security servers (authorization servers, privilege attribute servers, etc).

If restrictions are supported, the value returned by check_authorization would no longer be yes/no, but instead a list of restrictions. An answer of "no" would be represented by the restriction "not authorized". An answer of "yes" would be represented as an empty list of restrictions. A conditional yes would be represented by the list of conditions/restrictions. Some restrictions, such as time of day, could be checked directly by the authorization API, while others would be returned by the API as unresolved. These unresolved restrictions might be application specific, and they would be handled by a function wrapped around the API, avoiding the need for the authorization API to understand the needs of all potential applications. For example, consider the pseudo code to check authorization to use a network access server:

```
check_nas_acl(sc, aclid) {
   resp = check_authorization(sc, "nas:network", 0x1);
   for restrict in resp {
       case: net_mask;
             ...
             break;
       case: toll_lines_only;
             ...
             break;
       case: spare_capacity_only;
             ...
             break;
       default:
             reject;
    }
    accept;
 }
```

Technically, the principal identifier in an entry can be thought of as a restriction on who may use it, but we'll leave it as a separate field since that is what most people expect in an ACL entry. Multiple principal identifiers will be supported for each ACL entry, allowing the specification of compound principals. Principal identifiers will be typed. We considered using the GSSAPI representation for the principal. John Linn pointed out that the opaque nature of GSSAPI names makes them difficult to store in a persistent and shared access control list. Naming issues have been a stumbling block in other groups, and wanting to avoid that issue for now, we decided that we will support typed names. When the GSSAPI naming issues are resolved it might be one of the types. For now, the type would indicate the namespace (which at the present time is closely tied to the authentication method) from which the name is drawn.

The definition of the get_acl call as an alternate interface to the authorization API was deferred until the authorization database itself is better defined.

Work Items

Clifford Neuman	Compile a list of restrictions or privilege attributes addressing the needs of DCE, Sesame, restricted proxies, and various applications.
	Contact Russ Hobby for input on other application requirements.
Piers McMahon	Define the information that needs to be part of the Security Context.
Piers/Clifford	Work to further refine the form of the authorization database.

Allan Rubens/John Vollbrecht	Provide a description of the authorization needs for the network access server.
Sam Sjogren	Provide the same for FTP.

Thanks to Richard Graveman for his notes which were helpful in the preparation of these Minutes.

Attendees

Jim Barnes	barnes@xylogics.com
John Campbell	jrcamp@nosc.mil
Stephen Crocker	crocker@tis.com
Dale Dougherty	dale@ora.com
Jonathan Fellows	jonf@gdstech.grumman.com
Antonio Fernandez	afa@thumper.bellcore.com
Barbara Fraser	byf@cert.org
James Galvin	galvin@tis.com
Jisoo Geiter	geiter@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Richard Harris	rharris@atc.boeing.com
Laurent Joncheray	lpj@merit.edu
David Katinsky	dmk@pilot.njin.net
John Linn	linn@gza.com
Steven Lunt	lunt@bellcore.com
Cynthia Martin	martin@spica.disa.mil
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Clifford Neuman	bcn@isi.edu
Radia Perlman	perlman@dsmail.enet.dec.com
Christopher Provenzano	proven@csi.compuserve.com
Robert Raisch	raisch@ora.com
April Richstein	amr@tycho.ncsc.mil
Allan Rubens	acr@merit.edu
Paul Sangster	sangster@reston.ans.net
Wolfgang Schneider	schneider@gmd.de
Sam Sjogren	sjogren@tgv.com
Stuart Stubblebine	stubblebine@isi.edu
Louisa Thomson	louisa@whitney.hac.com
Klaus Truoel	truoel@gmd.de
Theodore Ts'o	tytso@mit.edu
John Vollbrecht	jrv@merit.edu
James Weatherford	weatherf@nosc.mil

CURRENT MEETING REPORT

Minutes of the Internet Protocol Security Protocol Working Group (IPSEC)

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

Attendees

Randall Atkinson	atkinson@itd.nrl.navy.mil
Jim Barnes	barnes@xylogics.com
Tom Benkart	teb@acc.com
Richard Bjers	rich.bjers@uc.edu
John Campbell	jrcamp@nosc.mil
Richard Colella	colella@nist.gov
David Conrad	davidc@iij.ad.jp
Stephen Crocker	crocker@tis.com
Chas DiFatta	chas@cmu.edu
Kishan Dudkikar	kishan@icm1.icp.net
Stefan Fassbender	stf@easi.net
Jonathan Fellows	jonf@gdstech.grumman.com
Antonio Fernandez	afa@thumper.bellcore.com
Barbara Fraser	byf@cert.org
Dan Frommer	dan@jeremy.enet.dec.com
Jisoo Geiter	geiter@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Robert Hagens	hagens@ans.net
Danny Hanson	hanson.tic@commlan.safb.af.mil
Richard Harris	rharris@atc.boeing.com
Alton Hoover	hoover@ans.net
John Ioannidis	ji@cs.columbia.edu
Dale Johnson	dsj@merit.edu
Phil Karn	karn@qualcomm.com
Charles Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
Paul Lambert	paul_lambert@email.mot.com
John Linn	linn@gza.com
Philip Lozowick	lozowick@isv.dec.com
Steven Lunt	lunt@bellcore.com
Cynthia Martin	martin@spica.disa.mil
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Joe Pagan	jrp@afterlife.ncsc.mil
Andrew Partan	asp@uunet.uu.net
Charles Perkins	perk@watson.ibm.com

Radia Perlman	perlman@dsmail.enet.dec.com
Marcus Ranum	mjr@tis.com
Paul Sangster	sangster@reston.ans.net
Jeffrey Schiller	jis@mit.edu
Wolfgang Schneider	schneider@gmd.de
Ron Sharp	rls@neptune.att.com
Robert Shirey	shirey@mitre.org
Allyson Showalter	allyson@nsipo.arc.nasa.gov
Michael St. Johns	stjohns@darpa.mil
Richard Thomas	rjthomas@bnr.ca
Paul Traina	pst@cisco.com
Klaus Truoel	truoel@gmd.de
Theodore Ts'o	tytso@mit.edu
Huyen Vu	vu@polaris.disa.mil
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
James Weatherford	weatherf@nosc.mil
James Zmuda	zmuda@mls.hac.com

2.6.1 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. Department 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:

Ongoing	Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard.
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.

Internet-Drafts:

"COMMON IP SECURITY OPTION", 03/10/1993, R. Sharp <draft-ietf-cipso-ipsec-option-00.txt>

CURRENT MEETING REPORT

Reported by Ron Sharp/AT&T

Minutes of the Commercial Internet Protocol Security Option Working Group (CIPSO)

The CIPSPO meeting started with Ron Sharp providing a brief introduction of the Working Group's purpose and history. This was followed by a concise description of CIPSO. The current status of the Internet-Draft was given with emphasis on the fact that the specification has still not reached Proposed Standard after more than a year of trying.

Attendees were warned that there were two versions of CIPSO in the Internet-Draft remote directories. Since then the situation has been corrected. The latest version of CIPSO is dated March 9, 1993. There was some concern about the file name of the Internet-Draft in the directories and Steve Crocker agreed to look into it and change it.

Ron described the changes made in the most recent version of CIPSO. Most of the changes were organizational with some addition of information to improve clarity. There was a new release tag added in this version. A couple of members were concerned as to the proper processing of the release tag and Ron agreed to follow up on this.

Noel Nazario of NIST provided a description of the work being done on the new NIST GOSIP security label. He is soliciting comments for this new label, particularly on how it relates or should relate to CIPSO.

The remainder of the time was left for discussion. Three of the people present were concerned that the specification needed more information such as pseudo-code to assist implementors. It was noted that over ten companies have successfully implemented CIPSO with an earlier specification and that it has gone through interoperability testing successfully.

The discussion ended with Steve Crocker agreeing to send detailed comments to Ron Sharp as to what is needed in order for the specification to go to Proposed Standard.

Attendees

John Campbell	jrcamp@nosc.mil
Stephen Crocker	crocker@tis.com
Kishan Dudkikar	kishan@icm1.icp.net
Jonathan Fellows	jonf@gdstech.grumman.com
Antonio Fernandez	afa@thumper.bellcore.com
Jisoo Geiter	geiter@mitre.org
Richard Harris	rharris@atc.boeing.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
Alton Hoover	hoover@ans.net
Charles Kaufman	kaufman@zk3.dec.com

Stephen Kent	kent@bbn.com
Tony Li	tli@cisco.com
James Mahon	mahonj@honfsi1.att.com
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Noel Nazario	nazario@csmes.ncsl.nist.gov
Christopher Provenzano	proven@csi.compuserve.com
April Richstein	amr@tycho.ncsc.mil
Paul Sangster	sangster@reston.ans.net
Vilson Sarto	vilson@fapq.fapesp.br
Ron Sharp	rls@neptune.att.com
Robert Shirey	shirey@mitre.org
Michael St. Johns	stjohns@darpa.mil
Klaus Truoel	truoel@gmd.de
James Weatherford	weatherf@nosc.mil
Charles Young	Charles.E.Young@att.com
Kiho Yum	kxy@nsd.3com.com
James Zmuda	zmuda@mls.hac.com

2.6.2 Common Authentication Technology (cat)

Charter

Chair(s):

John Linn, linn@gza.com

Mailing Lists:

General Discussion: cat-ietf@mit.edu
To Subscribe: cat-ietf-request@mit.edu
Archive: bitsy.mit.edu:~/cat-ietf/archive

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	Progress Internet-Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms.

Done Preliminary BOF session at IETF meeting, discussions with Telnet and Network Printing Working Groups.

Done Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet-Draft process.

Done First IETF meeting as full Working Group: review Charter distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.

Oct 1991 Update mechanism-independent Internet-Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.

Nov 1991 Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.

Dec 1991 Submit service interface specification to to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Generic Security Service Application Program Interface", 06/12/1991, John Linn <draft-ietf-cat-genericsec-04.txt>

"The Kerberos Network Authentication Service (V5)", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-02.txt, .ps>

"Generic Security Service API : C-bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-02.txt>

"Distributed Authentication Security Service", 11/04/1991, Charles Kaufman <draft-ietf-cat-dass-02.txt, .ps>

"FTP Security Extensions", 04/06/1993, S. Lunt <draft-ietf-cat-ftpsec-01.txt>

CURRENT MEETING REPORT

Reported by John Linn/Geer Zolot Associates and Sam Sjogren/TGV

Minutes of the Common Authentication Technology Working Group (CAT)

The Common Authentication Technology Working Group (CAT) met for two sessions at the Columbus IETF. The primary Agenda item was integration of security features into FTP, a topic for which Sam Sjogren is acting as task leader and on which Steve Lunt has generated a working document which will shortly be released as an Internet-Draft. Additional discussion topics were the advancement status of currently active CAT Internet-Drafts (GSS-API, GSS-API C bindings, and Kerberos V5), and a working proposal by Ted Ts'o for a CATS stream-oriented protocol overlay to be used in conjunction with GSS-API.

Status Of Specifications

The CAT Internet-Drafts have been pending administrative action for some time, and an action plan was evolved for their advancement recommendation. An updated Kerberos V5 specification was produced by Cliff Neuman; comments received from its review list by April 10th will be reflected in an Internet-Draft to be issued by April 17th. Primary late-breaking changes include added detail on an encrypted timestamp preauthentication type, and a new message type for credential exchange when proxies are being provided to an end server.

Concurrently with the Kerberos specification review, an advancement memo will be prepared for the set of Internet-Drafts (GSS-API, GSS-API C bindings, Kerberos V5). Intended follow-on documents will include a CAT mechanism definition (with token format) based on Kerberos V5. It was determined that use of the mechanism tagging recommended in GSS-API Appendix B should be adopted as mandatory, and message stream facilities were strongly desired within GSS-API mechanisms in order to satisfy FTP functional and integration requirements; a new appendix to the GSS-API specification has been drafted and distributed to the CAT mailing list codifying the results of this discussion.

CATS Stream Protocol

The CATS proposal, presented by Ted Ts'o, was engendered by a desire to provide prospective protocol integrators with a subprotocol specification for security, along with a stream-oriented API. An explicit CATS goal is that it be implementable quickly and without kernel modifications (unlike, e.g., IP- level security). CATS presents an alternative, generic form for protocol integration, contrasting with the protocol-specific integration being employed in Telnet options and FTP commands. Submission of a revised CATS specification as an Internet-Draft is planned.

In evolving CATS, Ted observed that the status of the tagging scheme in GSS-API Appendix B as a recommendation to mechanism designers led to undesirable ambiguity, and suggested making it mandatory; this suggestion was adopted and would be integrated into the Kerberos V5 implementation. Among other discussion, it was suggested that the CATS protocol be extended in order to exchange preamble tokens in advance of context establishment for mechanism type negotiation, and this prospect will be examined.

FTP Security

Sam Sjogren began the FTP part of the meeting by describing the goals for the FTP security work, as initiated at a BOF during the previous IETF and to be continued under the CAT Working Group framework. Support for authentication, as well as integrity verification and privacy, via any of a number of different mechanisms will be provided by this work.

Steve Lunt presented his proposal for extensions to the FTP protocol (additional commands) to implement the security goals; he intends to make an FTP security implementation publicly available. An overhead presentation supplemented the draft that had previously been posted to the Group's mailing list. Lively discussion occurred during and after this presentation, to the point that an additional evening meeting was scheduled to continue the discussions. The resulting FTP security discussions were quite fruitful, both in terms of providing feedback for improving the draft proposal for FTP, as well as fine tuning the GSS-API requirements and specifications. It was decided that the case of a three-party FTP interaction is sufficiently complex and rarely enough used that specification of how to do it securely will be deferred, probably to a separate document to be produced later.

Once peer entity authentication has been completed (and a session key is available), the proposed FTP security approach protects all subsequent FTP commands for at least integrity (encapsulation of base-64 encoding of gss_seal() output within MIC command), and optionally for confidentiality as well as integrity (encapsulation within ENC command, gss_seal() conf_flag TRUE). It was observed that underlying GSS-API mechanisms must represent and protect the conf_avail flag value and other service availability indicators within their tokens to prevent active attacks; the to-be-written "Kerberos Vn for GSS-API" and other mechanism design documents will describe how this protection is provided. (These indicator protection requirements apply independently of whether GSS-API is employed.) To protect against active attackers corrupting a data or control stream by changing the order of data or commands in the stream, protection via sequence numbers or some other such technique must be provided by the FTP security standard or the underlying mechanisms.

Given that the FTP control connection is a Telnet stream, questions arose about the rationale for not using the Telnet authentication option as an approach for FTP. Steve cited two reasons:

1. Because the Host Requirements RFC precludes use of Telnet options on the FTP control connection.

2. Because of a desire to provide data integrity, which Telnet security does not offer.

The Group discussed the fact that no facility was exposed at the level of the FTP security commands for negotiating the type and strength of cryptography to be used for data protection although, it was recognized that such features can exist within underlying mechanisms yet remain transparent at the FTP level. The ordering of USER and AUTH commands will be made such that an FTP server may make a decision on what authentication and encryption mechanisms are acceptable based on the identity of the user.

Kerberos ticket granularity and naming issues were discussed. For Kerberos V4, the form of a target FTP server's name should be "ftp.<simple-host>", and for Kerberos V5, "ftp/<domain-qualified-host>". It was noted that use of different target names for different server processes within a host is not primarily an access control measure, but does permit the servers to run in different protection domains (as might be desirable for an FTP server available to casual remote users).

Areas to be revised in the FTP draft based on discussion include:

- Additional discussion of requirements.
- Approach for transfer of arbitrarily long tokens from client to server.
- Alignment of base-64 encoding technique with RFC1421.
- Statements regarding mechanism negotiation (including a statement that an FTP server may refuse to accept anything less than suitably strong authentication).
- Further work on the data stream protection approach.

Attendees

Steve Alexander	stevea@lachman.com
John Campbell	jrcamp@nosc.mil
William Chung	whchung@watson.ibm.com
Stephen Crocker	crocker@tis.com
Dave Cullerot	cullerot@ctron.com
Antonio Fernandez	afa@thumper.bellcore.com
James Galvin	galvin@tis.com
Jisoo Geiter	geiter@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Frank Kastenholz	kasten@ftp.com
David Katinsky	dmk@pilot.njin.net
Charles Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Andrew Knutsen	andrewk@sco.com
Paul Lambert	paul_lambert@email.mot.com
John Linn	linn@gza.com
James Mahon	mahonj@honfsi1.att.com
Cynthia Martin	martin@spica.disa.mil
Evan McGinnis	bem@3com.com
P.V. McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Bob Morgan	morgan@networking.stanford.edu

Clifford Neuman	bcn@isi.edu
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Christopher Provenzano	proven@csi.compuserve.com
Steven Richardson	sjr@merit.edu
April Richstein	amr@tycho.ncsc.mil
Shawn Routhier	sar@epilogue.com
Jeffrey Schiller	jis@mit.edu
Wolfgang Schneider	schneider@gmd.de
Sam Sjogren	sjogren@tgv.com
Stuart Stanley	stuarts@apertus.com
Bob Stewart	rlstewart@eng.xyplex.com
Stuart Stubblebine	stubblebine@isi.edu
Louisa Thomson	louisa@whitney.hac.com
Klaus Truoel	truoel@gmd.de
Theodore Ts'o	tytso@mit.edu
James Weatherford	weatherf@nosc.mil
Peter Wilson	peter_wilson@3com.com

2.6.3 Network Access Server Requirements (nasreq)

Charter

Chair(s):

Allan Rubens, acr@merit.edu
John Vollbrecht, jrv@merit.edu

Mailing Lists:

General Discussion: nas-req@merit.edu
To Subscribe: nas-req-request@merit.edu
Archive:

Description of Working Group:

The Network Access Server Requirements Working Group has as its primary goal, to identify functions and services that should be present in IP Network Access Servers (NAS's) and to specify the standards that provide for these functions and services. The term "Network Access Server" is used instead of the more conventional term "Terminal Server" as it more accurately describes the functions of interest to this Group. A "Network Access Server" is a device that provides for the attachment of both traditional "dumb terminals" and terminal emulators as well as workstations, PC's or routers utilizing a serial line framing protocol such as PPP or SLIP. A NAS is viewed as a device that sits on the boundary of an IP network, providing serial line points of attachment to the network. A NAS is not necessarily a separate physical entity; for example, a host system supporting serial line attachments is viewed as providing NAS functionality and should abide by NAS requirements.

This Group will adopt (or define, if need be) a set of standard protocols to meet the needs of organizations providing network access. The immediate needs to be addressed by the Group are in the areas of authentication, authorization, and accounting (AAA). In general, this Group will select a set of existing standards as requirements for a NAS. If necessary, the Group will identify areas of need where Internet standards don't already exist and new standardization efforts may be required.

Initially the Group will independently investigate the two cases of character and frame oriented access to the NAS. This investigation will be aimed at determining what work is being done, or needs to be done, in this and other working groups in order to be able to define the set of NAS requirements. While the ultimate goal of this Group is to produce a NAS Requirements document, it may be necessary to define standards as well. This initial investigation will help determine what the goals of this Group need to be. The Group will also work with appropriate Working Groups to define required NAS standards that fall into the areas of these other groups.

Goals and Milestones:

Done	NAS Requirements Document posted as an Internet-Draft.
Nov 1992	Post an Internet-Draft on Character oriented Authentication, Authorization, and Accounting(AAA).
Nov 1992	Post an Internet-Draft on frame oriented AAA requirements.
Nov 1993	Submit the NAS Requirements document to the IESG as a Proposed Standard.

Internet-Drafts:

"Network Access Server Proposed Requirements Document", 10/01/1992, J. Vollbrecht, A. Rubens, G. McGregor, L. Blunk, Richard Conto <draft-ietf-nasreq-nasrequirements-00.txt>

CURRENT MEETING REPORT

Reported by Larry Blunk/Merit, Allan Rubens/Merit and John Vollbrecht/Merit

Minutes of the Network Access Server Requirements Working Group (NASREQ)

Agenda

- Quick Introduction and Review.
- User/NAS Authentication Issues. Review of Write Ups.
- Server Requirements Discussion.
- NAS/AAA Architecture Options Vendor Presentations/discussions.
- Plan for Future Work.

Review

The meeting started out with a review of what this Group is trying to accomplish. Its primary purpose is to establish a set of requirements for Authentication, Authorization, and Accounting (AAA) capabilities in a Network Access Server (NAS). This has been difficult because standards to address the required interfaces are not well defined. The focus has shifted to outlining requirements for standards that need to be developed.

This meeting attempted to focus on the NAC/NAS interface and the NAS/AAA server interface separately. It was pointed out that all parts of the authentication process need to be considered when evaluating a particular scheme, so such a distinction needed to be made carefully.

The term Network Access Client (NAC) was suggested quite early in the meeting to replace the term "user" in order to prevent confusion about the problem being addressed. A NAC is as a device such as a workstation or router that wishes to attach to the NAS.

NAC/NAS Authentication

John Vollbrecht presented three models for NAC/NAS authentication which he has described in postings to the NASREQ mailing list. The models presented were:

1. NAC id and password in the clear.
2. One way authentication of NAC to NAS, password not in clear.
3. Mutual authentication of NAC and NAS, again no password in the clear.

There was some discussion of whether the first approach was desirable since it puts the user password in the clear. It was pointed out that this is the predominant way of doing authentication at this time. It was also pointed out that the connection between NAC and

NAS would be coming under more serious attacks as time goes on and technology becomes more dispersed and better understood.

Jeff Schiller pointed out, and the Group endorsed, that any authentication methods or protocols adopted by this Group need to preserve the same degree of security provided by the underlying authentication mechanism. That is, inserting a NAS should not make the process more vulnerable than it was without a NAS.

There was some concern that the NAS needs to be an identifiable entity to prevent spoofing. There was a question as to whether mutual NAS/NAC authentication is necessary. One might be willing to trust that the connection is being made to the phone number being called.

Although many people at the meeting thought that it wasn't really necessary at this time, there still may be a need in the larger community. It is not the role of the Group to decide what level of security is needed but to provide for mechanisms to achieve various levels of security that can be selected by the NAS providers. Jeff Shiller noted that it is not very difficult to provide mutual authentication, so it should be provided for.

Another question that arose was "Even after you mutually authenticate with the NAS, how can you be assured that the NAS really is one provided by the network you think you're calling into?" One could envision someone advertising a phone number to use that is really connected to a machine that picks up user ids and passwords or logs all data. To avoid this, you may want to know that the particular NAS you've called into is really a NAS provided by the administrator of the network the NAS resides on. Using the concept of Group membership for this problem makes it addressable as an authorization (as opposed to authentication) issue.

It was also pointed out during this discussion that only NAC/NAS authentication is being considered here. While the same mechanisms might be used between a host system or a server and the authentication server, this would be done separately from the NAC/NAS authentication. The NAS is not meant to act as a security front end for a set of servers.

Jeff Schiller agreed to review the proposed authentication mechanisms and provide a set of specifications. Bill Simpson said that he would insure that those authentication requirements that impact PPP authentication protocols are addressed by the PPP Working Group.

NAS/AAA Interface Requirements

John Vollbrecht then opened discussion of the NAS/AAA interface. The environment that a NAS works in may be quite complex. The MichNet environment is one where several institutions provide dialup NAS service and allow the others to share usage of the NAS. In this environment it is necessary for a NAS to be able to authenticate a user at the user's home authentication server. The home authentication servers may be quite diverse, including Kerberos, Unix pw/id, MTS, etc. The NAS would also need to interface to multiple authorization and accounting servers. In this environment it seems that architecturally a common AAA server that interfaces to the NAS and then to all the servers would be appropriate - a "helper".

The interface to the common AAA server could be some existing standard, if one existed, that would serve the purpose. The Group had some discussion of whether such a standard did exist. No standard appeared to be appropriate.

John next presented in greater detail, the idea of the helper and the potential benefits of having a standard protocol between the NAS and the helper. The helper is a process that interfaces with other servers for the NAS. A standard NAS/helper interface is then required.

This approach could benefit NAS purchasers, vendors, and users. The NAS/helper concept allows NAS vendors to implement a standard protocol for interfacing with AAA rather than implementing a separate interface for each type of server to be supported. Users, or third party implementors, could then provide the special interfaces to AAA servers in stand alone packages that could augment or replace vendor implementations.

The hope is that the NAS/helper protocol could be defined in a manner that provides enough functionality and expandability to allow adaptation to evolving AAA standards. Network providers could base their choice of NAS on factors other than the AAA server types supported.

NAS/helper Proposal

Steve Willens then presented Livingston's RADIUS implementation. RADIUS is a system recently developed by Livingston that parallels the NAS/helper model. Livingston is willing to make code supporting this available as part of a NAS/helper standard if the Group thinks that would be a good idea.

Steve described the protocol used and the functions provided by this system. He also described how MD5 is used to pass unencrypted passwords securely from the NAS to the helper. This allows authentication systems that need to receive an unencrypted password to be used for the NAC/NAS authentication.

It was suggested that the entire data stream be encrypted, not just the packet containing the id and password. There are problems with this due to PPP and also due to the extra processing load or cost that would be imposed on the typically inexpensive NAS. It would be unreasonable to require sending all PPP data encrypted. It also would require DES or something which has much tighter export controls.

Representatives from PSI and JVNCnet saw a real need for a mechanism like RADIUS. Steve offered to make both the RADIUS client and server software available to others, including other vendors.

Plan for Future Work

Allan Rubens reminded the Group that there are still a few other issues that need to be addressed. One issue deals with how the per port packet filters are specified and updated. One proposal was that there should be a standard MIB for packet filters and that the filters should be updated using SNMP.

There is also the issue of how a NAC specifies what authentication domain an id belongs to. It was pointed out that Kerberos provides for this capability. There are still major accounting issues that need to be investigated and discussed. Steve Kent thought that the issue of distributed names could be significant.

Allan Rubens proposed that the Group attempt to finish up the "Requirements Document" by November. Some thought that it was too early. Others thought that we needed an "Architecture Document" first to clarify what a NAS is and the environment in which it is expected to operate. There was some concern that protocols were being defined by a requirements Group.

Action Items

Steve Willens	Agreed to head up an effort to define a NAS-"helper" protocol as an RFC. Representatives from at least two other vendors volunteered to help with this effort. The NASREQ Charter will be updated to reflect the current goals of this Group and discussions will take place to determine if this protocol should be defined as a product of this Working Group. It was pointed out that the NAS-helper model is just one way of addressing NAS/AAA issues and that the Group needs to still be open to other approaches. Steve will publish the names of people participating in the protocol definition and invites participation from anyone else interested. A goal is to have some working documents available for the Amsterdam IETF and a Draft RFC for November. Steve will advise if these goals are realistic.
Jeff Schiller	Volunteered to describe a set of NAC/NAS authentication protocols. Which will hopefully be available by the Amsterdam IETF to be passed to the PPPEXT Working Group.
Bill Simpson	Agreed to take those descriptions to the PPPEXT Group.
Rubens or Vollbrecht	Will rework the Charter to include the possibility of coming up with an RFC for a NAS/helper interface protocol.
Jim Barnes	Volunteered to Chair a meeting at the IETF in Amsterdam. A meeting had not been scheduled earlier because neither of the Chairs was able to commit to being at the meeting. It seems that

it would be good to have a meeting there, if only to allow more people to hear about what we are doing and to get more input from a European view. Barring strong objections, a meeting will be held in Amsterdam.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
Jim Barnes	barnes@xylogics.com
Ed Brencovich	edb@dss.com
Sandy Bryant	slb@virginia.edu
John Campbell	jrcamp@nosc.mil
John Chang	changj@ralvm6.vnet.ibm.com
David Conrad	davidc@iij.ad.jp
Avi Elenko	avi@dss.com
Jonathan Fellows	jonf@gdstech.grumman.com
Antonio Fernandez	afa@thumper.bellcore.com
Jisoo Geiter	geiter@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Terry Gray	gray@cac.washington.edu
Richard Harris	rharris@atc.boeing.com
Susan Harris	srh@umich.edu
Frank Heath	heath@cmc.com
David Katinsky	dmk@pilot.njin.net
Charles Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
Hock-Koon Lim	lim@po.cwru.edu
John Linn	linn@gza.com
Steven Lunt	lunt@bellcore.com
Cynthia Mills	cmills@bbn.com
Bob Morgan	morgan@networking.stanford.edu
Clifford Neuman	bcn@isi.edu
David O'Leary	doleary@cisco.com
Brad Parker	brad@fcr.com
Brad Passwaters	bjp@sura.net
April Richstein	amr@tycho.ncsc.mil
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
William Simpson	Bill.Simpson@um.cc.umich.edu
Bob Stewart	rlstewart@eng.xyplex.com
Theodore Ts'o	tytso@mit.edu
John Vollbrecht	jrv@merit.edu
Richard Warwick	richard@dss.com
James Weatherford	weatherf@nosc.mil

Steve Willens	steve@livingston.com
Cathy Wittbrodt	cjw@barrnet.net

2.6.4 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.
Ongoing	Revise Proposed Standards and submit to IESG for consideration as a Draft Standard, and repeat for consideration as an Internet Standard.
Done	Submit second document as an 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.

Done Submit Internet-Drafts to IESG for consideration as Proposed Standards.

Done Post an Internet-Draft of the MIME/PEM Interaction specification.

Apr 1993 Submit the PEM/MIME Specification to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"MIME-PEM Interaction", 11/23/1992, S. Crocker, N. Freed, M. Rose <draft-ietf-pem-mime-02.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"

RFC 1421 "Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures"

RFC 1422 "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management"

RFC 1423 "Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers"

RFC 1424 "Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services"

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

Minutes of the Privacy Enhanced Mail Working Group (PEM)

The PEM Working Group meeting began with a status update. It was noted that the PEM RFCs were published in February as Proposed Standards. The TISPEM implementation has been available for a while and a revised version, with improved documentation and CRL support, will become available in April. Plans are underway to transition TISPEM from use of BSAFE to RSAREF for cryptographic algorithms and to make installation easier, especially for single- user systems. Plans for easier distribution of TISPEM, perhaps via anonymous FTP with suitable export control warnings, are also under review.

PEM implementations will soon become available in the United Kingdom and Germany as a result of work under the EC PASSPORT program. Testing of these implementations is taking place with the TISPEM and other PEM implementations. An authentication-only (MIC-ONLY and MIC-CLEAR) version will be available from University College London (UCL) to facilitate international distribution. A full PEM (including encryption) implementation is expected to be available via anonymous FTP in Germany.

RSADSI noted that an updated PEM toolkit will be available for testing this summer. RSADSI announced that a "commercial assurance PCA," was suitable for use in a variety of high assurance environments. It was noted that users of Apple's OCE will be registered under this PCA. RSADSI also announced the creation of a "low assurance" PCA with two classes of CAs. One class is for CAs suitable for support of testing, and these CAs can be registered at no cost now and at a minimal cost in the future. A persona CA, operated as an automated responder, soon will be instantiated under this PCA, providing free certificates with unauthenticated names.

Trusted Information Systems (TIS) announced that it was planning to operate a "medium assurance" PCA and had developed a draft policy statement. UCL observed that it was planning to become a PCA and wished to examine the PCA statements from TIS and RSADSI in order to get a better understanding of the form of such statements.

In addition to providing its TISPEM status update, TIS introduced two discussion topics: use of mailbox names, (rather than distinguished names), in certificates and relaxing the certification hierarchy constraints. It was suggested that Domain Name System (DNS) names could be encoded in X.509 certificate format by creating a new attribute, the value of which would be a DNS mailbox name. An alternative suggestion was put forth to represent a mailbox name as a sequence of AVAs, each corresponding to one component of a DNS name (plus the user name). Several motivations for this proposal were provided:

- It allows a user to employ his existing mailbox name for authentication.
- It permits a user to become a PEM user more easily as it does not require the user to determine his distinguished name.

It was observed that this naming proposal is not likely to scale as well as the use of true distinguished names nor does it provide users with authenticated, descriptive identifiers. It was suggested that the use of DNS names in certificates might facilitate storage of certificates in the DNS, but this claim was not substantiated.

TIS also suggested that it would be appropriate to consider relaxing the current certification system. The motivations cited here included a desire to facilitate the deployment of PEM to users not connected to the Internet and to incorporate trust models not accommodated by the current certification system. No specific, detailed proposals were put forth, but TIS announced plans to make material available later. It was observed that if the current certification path validation rules are substantially relaxed, the result may be a system with functionality no better than systems such as PGP.

PEM and MIME Integration

The rest of the meeting was devoted to a discussion of the integration of PEM and MIME. The proposal put forth by the MIME Working Group Chair was to formally supersede RFC1421 with a specification for PEM which requires a minimal MIME implementation. The requirements for a minimally compliant MIME mailer, from both submission and delivery perspectives was presented. Changes required relative to existing PEM processing includes recognition of static MIME header fields, adding the quoted-printable transformation, generation and parsing of multipart and text content types, support of ISO 8859-n character sets, and support for parameterized boundary markers (not compatible with the current PEM boundary markers).

There were some strong objections to this proposal, on several grounds. For example, the MIME-PEM specification is not yet complete and there are a number of details in the currently published MIME-PEM specification which are contentious (e.g., the use of inband markers for local control of PEM submission processing and delivery indication of applied PEM processing), so a gap of six or more months would result if deployment of 822-PEM implementations were discouraged or halted until MIME-PEM is available. Also, this proposal would disenfranchise non-MIME users, which was viewed as an unacceptable result. Other attendees argued that they would not consider implementing 822-PEM and would wait for MIME-PEM because of a desire to support only MIME users.

There was general agreement that it would be advantageous for users to migrate to MIME-PEM. There was little or no support for the concept of requiring gateways to translate between 822-PEM and MIME-PEM users. Thus an outstanding question is whether there is any way to provide interoperability between 822-PEM and MIME-PEM users with minimal changes to 822-PEM and/or minimal extensions to MIME-PEM. A related question is the extent to which MIME-PEM can be profiled for 822 users, to minimize the burden for such users.

The Working Group agreed that revision of the MIME-PEM specification should be aggressively pursued with the intent that the revised specification should be thoroughly reviewed prior to the next IETF meeting in Amsterdam, when the PEM Working Group will meet.

In parallel, an examination of options for 822-PEM and MIME-PEM interoperability, e.g., minor 822-PEM revisions and suitable profiling of MIME-PEM for use with text-only messages, will be undertaken.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jules Aronson	aronson@nlm.nih.gov
Richard Bjers	rich.bjers@uc.edu
Nathaniel Borenstein	nsb@bellcore.com
Sandy Bryant	slb@virginia.edu
John Campbell	jrcamp@nosc.mil
William Chung	whchung@watson.ibm.com
Steve Dusse	spock@rsa.com
Antonio Fernandez	afa@thumper.bellcore.com
Michael Fidler	fidler@mitre.org
Barbara Fraser	byf@cert.org
Ned Freed	ned@innosoft.com
Raphael Freiwirth	5242391@mcimail.com
James Galvin	galvin@tis.com
Christine Garland	garland@ihspa.att.com
Jisoo Geiter	geiter@mitre.org
Terry Gray	gray@cac.washington.edu
Richard Harris	rharris@atc.boeing.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
Alton Hoover	hoover@ans.net
Christian Huitema	christian.huitema@sophia.inria.fr
David Katinsky	dmk@pilot.njin.net
Charles Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@binky.arc.nasa.gov
John Linn	linn@gza.com
Kent Malave	kent@bach.austin.ibm.com
Bob Morgan	morgan@networking.stanford.edu
Noel Nazario	nazario@sdns.ncsl.nist.gov
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Bradley Rhoades	bdrhoades@mail.mmmg.com
April Richstein	amr@tycho.ncsc.mil
Gary Rowe	gjrowe@attmail.com
Jeffrey Schiller	jis@mit.edu
Wolfgang Schneider	schneider@gmd.de
Einar Stefferud	stef@nma.com
Stuart Stubblebine	stubblebine@isi.edu
Louisa Thomson	louisa@whitney.hac.com

Klaus Truoel	truoel@gmd.de
Theodore Ts'o	tytso@mit.edu
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
James Weatherford	weatherf@nosc.mil

2.7 Transport and Services Area

Director(s):

- David Borman: dab@cray.com

Area Summary reported by David A. Borman/Cray Research

A major item of interest in the Transport and Services Area is that, as of the end of the Columbus IETF meeting, the Area no longer exists. In its place are now two new Areas, the Transport Area and the Service Applications Area. David Borman has stepped down from the IESG. Allison Mankin is the new Area Director for Transport, and David Crocker is the new Area Director for Service Applications.

For the Working Groups that were in the old Transport and Services Area, they will now be assigned as:

- Transport:
  - Audio/Video Transport Working Group
  - TCP Large Windows Working Group
- Service Applications:
  - Distributed File Systems Working Group
  - Domain Name System Working Group
  - Service Location Protocol Working Group
  - Trusted Network File Systems Working Group

Of these Working Groups, three of them met at the Columbus IETF meeting. Summaries of their meetings follow. Please refer to the Minutes of each Working Group for more details.

Audio/Video Transport Working Group (AVT)

In the AVT sessions, some open issues were discussed that had come up since the presentation of the draft Realtime Transport Protocol specification at the previous meeting. The primary items discussed at this meeting were:

- Elimination of IPv4 addresses carried within RTP.
- Separation of RTCP control functions not related to transport.
- The addition of security services and mechanisms to the protocol based on a proposal presented by Stuart Stubblebine.

No roadblocks were identified, but resolution of some of the questions will be left to email discussion. The Group intends to update the Internet-Drafts in May to be ready for RFC approval in June.

The third Working Group session was an "implementors agreement" discussion to promote convergence and interoperation among the audio and video programs using RTP. Ron Frederick presented a proposed API that would allow the various video programs to share decoding subroutines so they could all decode each other's data. Christian Huitema described the complex state involved in H.261 decoding and explained how it would be difficult to implement a decoder to process each packet separately as proposed in the API. The solution may be to treat the list of packets in an RTP synchronization unit as an application layer frame to be processed as a unit for decoding. The software video programs are now being converted from monochrome to color, and Paul Milazzo proposed a standard color representation based on CCIR 601 YCrCb.

Domain Name System Working Group (DNS)

The DNS Working Group spent some time discussing the future of the Group. Items that had been completed or were no longer applicable were removed from the Charter. Three sub-groups were formed to address specific issues:

- Scaling problems of "big zones" like the .COM zone, to be led by Bill Manning.
- DNS Security, to be led by James Galvin.
- Load balancing, to be led by Tom Brisco and Stuart Vance.

The DNS MIB was discussed. There still exist some non-trivial problems due to differences in the "SNMP way" and "DNS way" of doing things. Frank Kastenholz and the DNS MIB authors will be addressing this.

Susan Thomson gave a presentation on DNS support for PIP. Three problems were discussed by the Group. Recommendations were made on how to address two of the problems, but the third was left open due to time constraints.

The final discussion item was the proposed X.400 "temporary" routing mechanism, which uses the DNS to replace X.400 routing tables. Serious flaws in the proposal were discussed. The results of the discussion were presented to the authors of the document two days later. A better solution was agreed upon, and will be documented by Claudio Allocchio.

Service Location Protocol Working Group (SVRLOC)

The SVRLOC Working Group met twice. The sessions were spent conducting a technical review of just over half of the current Internet-Draft. Some of the items reviewed during the first session include:

- Network interactions between the user agent and the servers.
- Server timeout of cached information.
- Resending of broadcast queries.

The second session reviewed the packet structure, and had a presentation and discussion of internationalization issues.

2.7.1 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:~/ietf/rem-conf/av-transport-archive

Description of Working Group:

The Audio/Video Transport Working Group 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 Working Group 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 Working Group 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 Working Group 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. The first step will be to solicit contributions of potential protocols from projects that have already developed packet audio and video. From these contributions the Group will distill the appropriate protocol features.

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.

Done Post an Internet-Draft of the lightweight audio/video transport protocol.

May 1993 Post a revision of the AVT protocol addressing new work and security options as an Internet-Draft.

Jun 1993 Submit the AVT protocol to the IESG for consideration as an Experimental Protocol.

Internet-Drafts:

"Issues in Designing a Transport Protocol for Audio and Video Conferences and other Multiparticipant Real-Time Applications", 10/27/1992, H. Schulzrinne <draft-ietf-avt-issues-00.txt, .ps>

"A Transport Protocol for Real-Time Applications", 12/16/1992, H. Schulzrinne, S. Casner <draft-ietf-avt-rtp-01.txt>

"Media Encodings", 12/16/1992, H. Schulzrinne <draft-ietf-avt-encodings-00.txt>

"Sample Profile for the Use of RTP for Audio and Video Conferences with Minimal Control", 12/16/1992, H. Schulzrinne <draft-ietf-avt-profile-00.txt>

"Packetization of H.261 video streams", 03/11/1993, T. Turletti, C. Huitema <draft-ietf-avt-video-packet-01.txt>

CURRENT MEETING REPORT

Reported by Steve Casner/Information Sciences Institute

Minutes of Audio/Video Transport Working Group (AVT)

The AVT Working Group met for three sessions. The first two sessions were used to discuss some open issues on the Draft specification for the Real-Time Transport Protocol (RTP); the third session was an "implementors agreement" session focusing on software video encoding.

Status of the Working Group and Review of the Draft RTP Specification

The goal of the AVT Working Group is to specify a set of experimental protocols for real-time transmission of audio and video. The emphasis is short-term to promote experimentation now so that standards-track protocols can be developed based on what is learned. The first Working Group meeting was a year ago, so one might expect a conclusion soon. In fact, many issues were resolved during the last two IETF meetings and the specification is nearing readiness for submission as an RFC.

A set of four Internet-Drafts specifying the RTP protocol were issued in December 1992, and an Internet-Draft on packetization of H.261 coded video was issued in March, 1993. In addition, RTP has been implemented to varying degrees in the following programs:

IVS	Thierry Turletti and Christian Huitema
Maven	Charley Kline
NEVOT	Henning Schulzrinne
"nv"	Ron Frederick
PictureWindow	Paul Milazzo and Bob Clements (an older draft)
PVP	Steve Casner

This session began with a brief review of RTP. It consists primarily of protocol headers for real-time data packets, which is just eight octets long in the typical case. RTP supports the following functions:

- Transfer of media data.
- Demultiplexing of multiple flows.
- Content identification (e.g., the media encoding used).
- Synchronization and sequencing.
- Options for simple control functions such as identification of participants.

Full details on the protocol are available in the Internet-Drafts, which are listed later in this document.

Discussion of Open Issues

A few open issues had developed since the last meeting and were discussed. The primary items were:

- Elimination of IPv4 addresses carried within RTP.
- Separation of RTCP control functions not related to transport.
- Security services and mechanisms.

These issues are expanded in the following paragraphs. No roadblocks were identified, but resolution of some of the questions was left to email discussion.

The CSRC and SSRC options in the December RTP specification carried globally unique identifiers for the "content source" and "synchronization source" respectively, in data packets that have been processed through transport/application-level gateways such as audio mixers. These globally unique identifiers were based on IPv4 addresses, but this is not acceptable considering the impending transition to a next-generation IP. Therefore, the Group considered two revision choices:

1. Use a (type, length, value) triple to allow any form of address to be carried.
2. Carry only a 16-bit or 32-bit identifier that is locally unique to the gateway that inserts the option, and then define the mapping to a globally unique address, or other information, through an extended SDESC (source description) option or a higher-layer protocol.

Since the CSRC and SSRC options may be carried in every data packet for some flows, the long addresses in the first choice might impose an uncomfortably large overhead. Furthermore, the Group recognized that the locally unique identifier is sufficient for an RTP receiver to distinguish the source and process the packet; it is only necessary to map to a globally unique address or user name for purposes of monitoring, control, and user interface. Therefore, the second choice was selected, and it was generally agreed that 16 bits would be sufficient because the identifiers are unique only with respect to the full source address of the gateway (network address plus UDP port, for example).

Mapping of the identifiers to other information, such as the name of a conference participant, should be accomplished through a higher-layer control protocol. Note that the gateway entities must be involved in that protocol, but this would be true even if globally unique addresses were used since the addresses include port numbers to allow more than one entity per host. One way to accomplish the mapping is through the Source Description (SDESC) option in the RTP control "sub-protocol" RTCP. The SDESC option includes the 16-bit identifier plus one or more items to which it is mapped. Examples include the user name, as before, and various forms of globally unique addresses.

A full address must also be specified for the return of RTP reverse control packets containing options such as the Quality of Service Measurement (QOS) option. In the December Draft, a return port number was specified in the Content Description (CDESC) option, but it was agreed that a full address is required because the return information may go to the content or synchronization source or to third party monitoring host. Furthermore, it was suggested that the return address be separated into its own option because it may be desirable to establish a return address without defining a new content type or redefining an old one. One problem is that the sender and receiver of the reverse control information may be using different forms of addressing, for example IPv4 and SIP or PIP. This problem extends beyond RTP, and its solution will depend on IPng transition plans. Two solutions the Group has considered are to use a DNS name or to allow multiple forms of binary address to be specified.

Some members of the Working Group objected to including within the RTP specification those RTCP options unnecessary for transport functions. However, it seems impractical to split off a few pages of RTCP option definitions into another RFC. It was agreed to keep the RTCP options within the RTP specification as a separate section with appropriate disclaimers about these functions being replaced by higher-layer control protocols when they are available.

The RTP specification Draft includes a brief Security Considerations section, but the protocol will be inadequate for teleconferencing applications, at least for confidentiality of voice communication, without access to some security mechanisms. In the future, it may be possible to depend on security mechanisms at the IP layer, but for near-term use, possibly including experimentation with new security mechanisms specifically for real-time applications, the Group believes it is appropriate to define security mechanisms at the RTP layer.

Stuart Stubblebine made a presentation reviewing the security services and mechanisms that might be needed for applications that use RTP. The most needed services are confidentiality and integrity of the payload, and authentication of the source. Integrity is considered separately for individual packets and for the stream of packets. A set of three new RTP options was proposed to implement these security services:

- ENC, to indicate the start point for encryption and carry the initialization parameters.

- MIC, which may be used with a variety of security schemes, for example to carry a Message Integrity Check.

- KDEF, an RTCP option to define key identifiers.

Details of these security options will be found in a new Draft of the RTP specification to be released in mid-May.

In addition to the major topics, the Group also discussed a request from Frank Hoffmann for two additions to the protocol to allow higher reliability levels of service. The first was that a checksum option be provided for use when RTP is carried over ST-II, which does not provide a checksum. However, there is a problem with carrying a checksum in an option: an error in the header may make it appear that there is no checksum option, so the checksum would not be checked and the error would not be detected. Also, it is inconvenient that an RTP-level reflector from ST-II to IP/UDP would have to check and remove the checksum option to avoid presenting the UDP destination with two potentially conflicting checksums. As an alternative, it is suggested that there be a separate specification for an encapsulation of RTP in ST-II that would include a checksum. That encapsulation might define a service like that of UDP over IP.

The second request was for an option to request retransmission to implement a "reliable" class of service. It is expected that most real-time applications will not want to incur the delay imposed by retransmission. A generic retransmission function probably does not make sense in RTP, but reverse control options can be used to request retransmission in an application-specific manner when appropriate. For example, negative acknowledgements are defined in the INRIA H.261 packetization protocol Draft. For those applications that want to use the services of RTP but do require reliable delivery, RTP can be transported in a TCP stream.

One final topic that the Group did not have time to discuss adequately was how RTP profiles should be defined and used. Would it be reasonable to include as part of a profile definition that the RTP framing field would always be included in order to allow multiple RTP PDUs to be assembled into, for example, one UDP packet? The Group has also not defined how the selection of a particular profile will be identified for applications that can operate with more than one profile. More work is needed on this topic.

"Implementors Agreement" Session

The third Working Group session was devoted to an "implementors agreement" discussion to promote convergence and interoperation among the software video compression programs being developed.

At the previous meeting, it was observed that the tight coupling between the video frame grabbing procedure and the encoding process might mean that it would be infeasible to define an API between these two steps. However, Ron Frederick has observed that the coupling between software decoding and the display process can be much more flexible. Therefore, it may be feasible to achieve interoperation by allowing each hardware and software platform to encode and transmit data in its native format, but to incorporate multiple decode routines using a common API so that each program can decode many or all of the other programs' native formats.

Ron Frederick gave a presentation on the implementation of a decoder API in version 3.0 of the "nv" program. Decoding and rendering of the image data and decoupled: "nv" does all the network I/O, RTP processing, and X window system interaction; the image decode routines just convert each packet of compressed bits into uncompressed pixels for a portion of the image. When the "S" bit (end of synchronization unit) is set in the RTP header, "nv" will display the new image on the screen. This scheme allows support for multiple display depths, brightness/contrast mapping, and image scaling with simultaneous display of multiple sizes. Three decoders have been incorporated into "nv" to process video from "nv", from the hardware Bolter codec, and from the Cornell CU-SeeMe program.

Christian Huitema identified a problem with the "nv" API when he described the H.261 encoding in the IVS program and the complexity that results from a combination of image structure and Huffman encoding of the image coefficients. There is a lot of state information implicit in the decoding process in the middle of processing "Group of Blocks" (GOB). Since one GOB may occupy more than one packet, it might be infeasible to try to save and restore the state at the boundary between packets so that control could be returned across the API from a decode routine. Therefore, IVS uses the "S" bit to indicate the GOB boundary so that all of the packets of a GOB can be handed together to the decode routine.

This conflict in the interpretation of the "S" bit will have to be resolved to allow interoperation of "nv" and IVS. It was suggested that IVS is using the "S" bit for a transport function, whereas "nv" is using it for a presentation function, and therefore the former is correct. Ron said "nv" could be modified to get the display-image indication as a return valued from the decode routine, so this may be the solution.

Paul Milazzo has implemented color video transmission in the BBN PictureWindow program. Paul proposed that the YCrCb color encoding from the CCIR 601 digital video standard be used as the color representation for software encoding. This is similar to the YUV analog encoding. The proposed scheme would keep the luminance (Y) pixels in one array, and chrominance (CrCb or UV) pixel pairs, subsampled by 2 in the X dimension, in another array. Because of the subsampling, the two arrays would be the same size. This scheme allows easy rendering into monochrome or color images.

Further Working Group Activities

A set of Internet-Drafts on RTP was issued in December 1992:

- draft-ietf-avt-rtp-00.txt
- draft-ietf-avt-encoding-00.txt
- draft-ietf-avt-profile-00.txt
- draft-ietf-avt-issues-00.ps, .txt

The first Draft is the specification of the Real-Time Transport Protocol itself. The second and third Drafts define a set of media encodings and a sample profile for use of those encodings to implement audio and video multiparticipant conferences with minimal control. The last Draft is an updated discussion of the issues and decisions involved in the design of the protocol.

Revised Drafts incorporating changes discussed at this meeting will be issued in May. After review and possible further revision, it is expected that the Internet-Drafts will be submitted for approval as RFCs in June, completing the Working Group's Charter.

Attendees

Lou Berger	lberger@bbn.com
Larry Blunk	ljb@merit.edu
John Boatright	bryan_boatright@ksc.nasa.gov
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Monroe Bridges	monroe@cup.hp.com
Sandy Bryant	slb@virginia.edu
Stephen Casner	casner@isi.edu
Yee-Hsiang Chang	yhc@hpl.hp.com
William Chimiak	chim@relito.medeng.wfu.edu
Richard Cogger	R.Cogger@cornell.edu
David Conklin	conklin@jvnc.net
Simon Coppins	coppins@arch.adelaide.edu.au
Mark Davis-Craig	mad@merit.edu
Shane Dawalt	sdawalt@desire.wright.edu
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Tony DeSimone	tds@hoserve.att.com
Ed Ellesson	ellesson@vnet.ibm.com
Chip Elliott	celliot@bbn.com
Hans Eriksson	hans@sics.se
Francois Fluckiger	fluckiger@vxcern.cern.ch
Paul Franchois	paulf@bldrdoc.gov
Ron Frederick	frederick@parc.xerox.com
Jerry Friesen	jafries@sandia.llnl.gov
Marcello Frutig	frutig@rnp.impa.br
Gwen Funchess	funchess@magnus.acs.ohio-state.edu
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Fengmin Gong	gong@concert.net
Kenneth Goodwin	goodwin@a.psc.edu
Mark Green	markg@apple.com
Robert Gutierrez	gutierre@nsipo.nasa.gov
Don Hoffman	hoffman@eng.sun.com
Frank Hoffmann	hoffmann@dhdibm1.bitnet
Christian Huitema	christian.huitema@sophia.inria.fr
Phil Irey	pirey@relay.nswc.navy.mil
Ronald Jacoby	rj@sgi.com
Peter Kirstein	P.Kirstein@cs.ucl.ac.uk
Charley Kline	cvk@uiuc.edu
Lakshman Krishnamurthy	lakashman@ms.uky.edu
Giri Kuthethoor	giri@ms.uky.edu
Paul Lambert	paul_lambert@email.mot.com
Ruth Lang	rlang@nisc.sri.com
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Yu-Lin Lu	yulin@hpinddu.cup.hp.com

Marjo Mercado	marjo@cup.hp.com
Donald Merritt	Don@brl.mil
Paul Milazzo	milazzo@bbn.com
Joseph Pang	pang@bodega.stanford.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
Jim Rees	jim.rees@umich.edu
Michael Safly	saf@tank1.msfc.nasa.gov
Carl Schoeneberger	70410.3563@Compuserve.com
Eve Schooler	schooler@isi.edu
Stuart Stubblebine	stubblebine@isi.edu
Sally Tarquinio	sallyt@gateway.mitre.org
Craig Todd	ctodd@desire.wright.edu
Claudio Topolcic	topolcic@cnri.reston.va.us
Hung Vu	hungv@fonorola.com
Huyen Vu	vu@polaris.disa.mil
Abel Weinrib	abel@bellcore.com
John Wroclawski	jtw@lcs.mit.edu
Yow-Wei Yao	yao@chang.austin.ibm.com

2.7.2 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.

2.7.3 Domain Name System (dns)

Charter

Chair(s):
Rob Austein, sra@epilogue.com

Mailing Lists:
General Discussion: namedroppers@nic.ddn.mil
To Subscribe: namedroppers-request@nic.ddn.mil
Archive: nicfs.nic.ddn.mil:~/namedroppers/*.Z

Description of Working Group:

The DNS Working Group is concerned with the design, operation, and evolution of the Domain Name System within the Internet. As the Internet continues to grow, we expect to serve as a focal point for work on scaling problems within the current framework, work on protocol evolution as new mechanisms become necessary, and documentation of current practice for DNS implementors and administrators. We are also responsible for oversight of DNS activities by other groups within the IETF to the extent that we review the impact such work will have on the DNS and make recomendations to the working groups and IESG as necessary. Since some of these are ongoing tasks, we do not expect the working group to disband anytime soon.

Several issues are of particular concern at this time:

Scaling. The DNS is the victim of its own success. The global DNS namespace has grown to the point where administering the top levels of the tree is nearly as much work as the old NIC host table used to be. We need to work on ways to distribute the load. Some of the solutions are likely to be technical, some political or economic; we still treat the top-level DNS service the way we did when DARPA was footing the bill, and the funding for that service is in the process of going away.

Security. The DNS is a zero-security system; it is not even as strong as the IP layer above which it operates. As a result, accidental spoofing (cache pollution) is an all-too-frequent occurance. We need to make the DNS more robust against accidental corruption, and must provide at least an optional authentication mechanism for that portion of the community that wants one. At the same time, we must not cripple the existing system by drasticly increasing its bandwidth consumption or by mandating use of cryptographic techniques that would preclude worldwide distribution of DNS software. The global DNS database is exactly that, an existing world-wide database representing hosts on six continents and (at least) forty-five countries. A solution that does not take this into account is not acceptable.

Management. The Group has a draft document describing MIB extensions to manage the DNS. It also needs to specify a standard way to dynamically create and destroy DNS records; SNMP may be an appropriate tool for this task, but we haven't yet specified enough of the details to know for certain. The impact that a dynamic update mechanism will have on the DNS needs to be examined, with particular attention given to security and scaling issues.

IPng/Routing. As the fur starts flying in the battle between the IPng proponants and the new-routing-architecture proponants, it is expected that groups on both sides will need some amount of support from the DNS. Such support is likely to be minimal and straightforward, but these proposals are likely to need "rush service" for whatever support they require. So the Working Group needs to monitor these activities, stay involved, and generally do what it can to make sure that DNS support is not a bottleneck.

The DNS Group also needs to examine the impact that any proposed IPng system would have on the DNS, since the DNS database and protocols have special provision for IP addresses.

Goals and Milestones:

Done	Post as an Internet-Draft a description of the Responsible Person Record.
Done	Post an Internet-Draft specifying the addition of network naming capability to the DNS.
TBD	Submit to the IESG the document for load balancing in the DNS as an Informational document.
Done	Submit the Responsible Person Record to the IESG for consideration as a Proposed Standard.
Ongoing	Monitor and offer technical support to the various groups working on the next version of IP.
TBD	Post an Internet-Draft of the "Big Zone" policy recommendations for root and first-level zone adminstraton.
TBD	Submit the "Big Zone" policy document to the IESG for consideraton as a policy statement.
Done	Submit the specification for network naming to the IESG for consideration as a Proposed Standard.
Done	Post the DNS MIB as an Internet-Draft.
Feb 1993	Submit the DNS MIB to the IESG for consideration as a Proposed Standard.
Mar 1993	Post an Internet-Draft specifying the dynamic resource record creation and deletion.

Mar 1993 Submit to the IESG the incremental zone transfer mechanism as a Proposed Standard.

Mar 1993 List and prioritize the Working Group's goals, and pick a subset that is appropriate to pursue at the present time.

Jun 1993 Post an Internet-Draft for adding load balancing capability to the DNS.

Nov 1993 Submit the proposal for dynamic resource record creation/deletion to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"DNS MIB Extensions", 03/05/1992, R. Austein, J. Saperia <draft-ietf-dns-mibext-05.txt, .ps>

"DNS Support for IDPR", 03/22/1993, R. Austein <draft-ietf-dns-idpr-00.txt>

CURRENT MEETING REPORT

Reported by Rob Austein/Epilogue Technology

Minutes of the Domain Name System Working Group (DNS)

The meeting opened with a few comments by the Chair regarding the future of the Working Group. In brief, the Working Group has been in existence for a long time, with ill-defined goals and an odd mixture of tasks. Some of the Working Group's tasks are of an ongoing nature, some are very specific with definite goals and milestones. As currently organized, the Working Group does not fit well into the IETF administrative structure. At some point in the relatively near future it may be necessary to restructure the Group to address this problem. The Chair will be working with the Director of the new IETF Service Applications Area to address this issue.

Given the state of flux that the IESG and the IETF itself were in at the time of the Working Group meeting, the Working Group decided to spin off a few well-defined tasks to "sub-groups" but otherwise leave the current DNS Working Group structure unchanged. Whether the "sub-groups" would be full-fledged IETF working groups or just subcommittees of the DNS Working Group was left unspecified; the organizers of the three sub-groups spun off at this meeting all subsequently expressed a distinct preference for being subcommittees rather than independent working groups.

The next item of business was to dispose of some old tasks that were still listed on the Group's Charter. The following tasks were removed from the task list:

1. Discussion of adding a Responsible Person RR.
2. Discussion of adding network naming capability to the DNS.
3. Implementation catalogue for DNS software.
4. Writing a "DNS operators guide".

Tasks (1) and (2) were done years ago and published in RFCs 1183 and 1101, respectively. Task (3) was killed for lack of volunteers who cared enough to write it; if such a volunteer exists, that person should just write the document, a Group effort is not needed.

The Group also briefly discussed writing a "DNS Operators guide", Task (4). Since both RFCs 1032-1033 and an O'Reilly Associates book exist on the subject, the Group felt that writing such a document would be a waste of time. Stuart Vance agreed to write a one-page document referring readers to the O'Reilly Associates book; if there are other published documents on this subject, whether commercial or free, please announce them on the namedroppers list for possible mention in Stuart's reference.

The Group felt the following three tasks were of sufficient interest to form a subgroup for each. Each task was discussed long enough to air some of the issues that needed to be addressed, the Chair then asked if someone was willing to organize a subgroup, attack the problem and report back.

1. Scaling problems of "big zones", particularly the .COM zone.

   The Working Group has been kicking this problem around for a year without much progress. Discussion was limited to a summary of the problem by Mike St. Johns and to new suggestions by the Working Group. One-line summary: some of the first-level DNS zones are approaching the size of the old NIC HOSTS.TXT file, which leads to various technical and political problems. Please see the Minutes of previous DNS Working Group meetings for background on this discussion.

   John Romkey pointed out that, as the Internet grows and more hosts are registered in the DNS, it will be increasingly hard to maintain any kind of obvious correlation between DNS names and "real" names. All the proposed "solutions" to the .COM zone problem will exacerbate this situation; the only exception is the "XYZ-CORP.X.COM" partitioning scheme, which, while ugly, is at least obvious. In the long run we need a solution to the problem of unintuitive names in the DNS. Whatever solution is picked for the .COM zone should be chosen with this in mind, so that the problem is not made worse.

   Bill Manning agreed to lead a subgroup to investigate the problem and report back. The subgroup's mailing list is "bigz@rice.edu", send mail to "bigz-request@rice.edu" to join the list.

2. DNS security.

   The Working Group heard a report from Steve Crocker, Director of the Security Area. There are really two tasks under discussion.

   (a) The first security task is bulletproofing the DNS so that it cannot be spoofed without ability to spoof IP addresses; this is, essentially, Paul Mockapetris's "Just As Good As IP" security mechanism, as described by Paul at the "DNS-II" BOF held at the 25th IETF in Washington DC. This level of security could also be described as "robustness," that is, implementation of Paul's techniques would help defend the global DNS database against some of the accidental spoofing that has happened over the years. This is primarily an implementation issue, and wheels are already rolling to get this mechanism into a near-future release of BIND. Watch the namedroppers list for announcements.

   (b) The second security task is specifying a mechanism by which RRs could be accompanied by digital signature information for authentication. Both for backwards compatibility with existing DNS software and in order to avoid running afoul of U.S. export controls on cryptographic software, this signature must be an optional mechanism, added in such a way that non-players can ignore it and still comply with the DNS protocols. The general idea is to define a new RR type, the RDATA portion of which would be used to contain the digital signature. James Galvin agreed to lead a subgroup to work on this project and report back to the Working Group. The subgroup's mailing list is "dns-security@tis.com", send mail to "dns-security-request@tis.com" to join the list.

3. Load balancing.

   This task has been with the Working Group for a long time. At the time of this meeting there were at least four mechanisms proposed to solve some aspect of the "load balancing" problem, some already in widespread use, one requiring no protocol changes at all. Tom Brisco and Stuart Vance agreed to lead a subgroup to investigate this problem and document their conclusions. The subgroup's mailing list is "dns-wg-lb@ns1.rutgers.edu", send mail to "dns-wg-lb-request@ns1.rutgers.edu" to join the list.

Next, the Working Group heard a report on the current status of the DNS MIB from Frank Kastenholz, representing the IETF Network Management Directorate. In brief, the DNS MIB has been stalled since the 25th IETF for two reasons:

1. The Network Management Area Directorate has been busy with the emerging SNMPv2 specification and suffered some disorganization due to the abrupt resignation of its Area Director.

2. The DNS MIB presents some non-trivial problems, because there are some cases where the "SNMP way" and the "DNS way" of doing things are diametrically opposed.

Problem (1) is already being repaired: a new Area Director for Network Management was elected two days after the DNS Working Group meeting, and the SNMPv2 specification is on its way out the door.

Problem (2) will be addressed by Frank (representing the Network Management Area Directorate) and the authors of the DNS MIB.

Frank suggested that the DNS Working Group consider whether or not SNMP was really the right tool for all of the management tasks addressed by the DNS MIB, and for the proposed dynamic update mechanism (dynamic creation and deletion of RRs via a network protocol); it's possible that extensions to the DNS protocol might be more appropriate for some of these tasks. James Galvin pointed out a counter-argument: if DNS management is not done via SNMP, the extended DNS protocol would need to duplicate much of the mechanism of SNMP, including the security features. The Working Group decided to press ahead with cleaning up the DNS MIB, given that we've already spent so much time on it.

Next, Susan Thomson gave a presentation on DNS support for PIP. The presentation was basically the same material covered in the PIP-DNS Internet-Draft ("draft-ietf-pip-dns-00.txt"). The Group discussed three problems with the PIP-DNS proposal:

1. The PIP-DNS proposal creates two new semi-wildcard QTYPEs (the PIP document calls these "special-purpose query types"). Semi-wildcard QTYPEs have known prob-

lems when combined with DNS caching algorithms, as documented in RFC-973, page 4, "MD and MF replaced by MX". The Working Group recommended that the PIP Group redesign their proposal to get rid of the semi-wildcard QTYPEs.

2. The proposed BBD (BackBone Descriptor) RR format defined a one-octet field to select the character set used in the variable length text portions of the RR. The Working Group recommended that this be eliminated and that the variable length text portions of the RR be limited to the NETASCII character set.

3. PIP includes a mechanism by which a PIP implementation can know that it needs to get a fresh copy of an RR. The PIP implementors would like to have a way to speed propagation of fresh information when this happens. Discussion of this subject on the namedroppers list suggests that one way to accomplish this would be the addition of an RR timestamping mechanism to the DNS. To date we do not know how to add such timestamps without an incompatible change to the DNS protocols, so we may not be able to help the PIP implementors here. This issue was left open, due to time pressure at the Working Group meeting.

There was a brief discussion of the proposed X.400 "temporary" routing mechanism (using the DNS to replace X.400 routing tables). Those members of the Working Group who had read the proposal felt that it was seriously flawed and could not be implemented as specified. Rob Austein and Jon Postel volunteered to meet with the authors of the proposal in order to find the right answer to this problem. Said meeting took place two days later, and resulted in a better solution, to be documented by Claudio Allocchio, one of the authors of the original proposal.

At this point, having run out of time and having covered all the major items on the Agenda, the meeting was adjourned.

Attendees

N. Akiko Aizawa	akiko@nacsis.ac.jp
Philip Almquist	almquist@jessica.stanford.edu
Randall Atkinson	atkinson@itd.nrl.navy.mil
Robert Austein	sra@epilogue.com
David Battle	battle@cs.utk.edu
David Bolen	db31@ans.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
David Bridgham	dab@epilogue.com
Thomas Brisco	brisco@pilot.njin.net
Sandy Bryant	slb@virginia.edu
Henry Clark	henryc@oar.net
Wayne Clark	wclark@cisco.com
David Conklin	conklin@jvnc.net
Stephen Crocker	crocker@tis.com

John Curran	jcurran@nic.near.net
Chas DiFatta	chas@cmu.edu
James Galvin	galvin@tis.com
Richard Graveman	rfg@ctt.bellcore.com
Steven Hubert	hubert@cac.washington.edu
Daniel Karrenberg	daniel@ripe.net
Frank Kastenholz	kasten@ftp.com
David Katinsky	dmk@pilot.njin.net
Kenneth Key	key@cs.utk.edu
John Klensin	klensin@infoods.unu.edu
John Linn	linn@gza.com
Daniel Long	long@nic.near.net
Steven Lunt	lunt@bellcore.com
Paul Lustgraaf	grpjl@iastate.edu
Bruce Mackey	brucem@cinops.xerox.com
Bill Manning	bmanning@sesqui.net
Clifford Neuman	bcn@isi.edu
William Nowicki	nowicki@legato.com
Masataka Ohta	mohta@cc.titech.ac.jp
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@sura.net
Michael Patton	map@bbn.com
John Penners	jpenners@advtech.uswest.com
Jon Postel	postel@isi.edu
Robert Reschly	reschly@brl.mil
John Romkey	romkey@asylum.sf.ca.us
Jon Saperia	saperia@lkg.dec.com
Carl Schoeneberger	70410.3563@Compuserve.com
Tim Seaver	tas@concert.net
Allyson Showalter	allyson@nsipo.arc.nasa.gov
Michael St. Johns	stjohns@darpa.mil
Martha Steenstrup	msteenst@bbn.com
Bernhard Stockman	boss@ebone.net
Susan Thomson	set@bellcore.com
L. Stuart Vance	vance@tgv.com
Ruediger Volk	rv@informatik.uni-dortmund.de
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Jane Wojcik	jwojcik@bbn.com

2.7.4 Service Location Protocol (svrloc)

Charter

Chair(s):

John Veizades, veizades@apple.com
Scott Kaplan, scott@wco.ftp.com

Mailing Lists:

General Discussion: srv-location@apple.com
To Subscribe: srv-location-request@apple.com
Archive: apple.com:~/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 service 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 service 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 service location and then determine the correct action in their view whether it be to use current protocols to suggest a service 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?

Internet-Drafts:

"Resource Location Protocol", 03/12/1993, S. Kaplan <draft-ietf-svrloc-resloc-00.txt, .ps>

CURRENT MEETING REPORT

Reported by Scott Kaplan/FTP

Minutes of the Service Location Protocol Working Group (SVRLOC)

The Service Location Working Group met twice to conduct a technical review of the current Internet-Draft and got a little over half way through the document.

The first session focused on the basic network interactions between the user agent and the servers (resource agent/binding broker). The main issues were how servers should timeout information that they cache and backoff strategies for resending broadcast queries. There was little disagreement. The Group consensus will appear in the next version of the Draft.

There was also a discussion about whether or not a client could ever converge on receiving database updates if a resource agent rapidly changed its database. Scott Kaplan agreed to review this issue and offer a solution in the next Draft.

The second session focused on the packet structure. There were minor editorial changes. The session concluded with a presentation from John Veizades about internationalization. A discussion regarding internationalization issues followed.

Attendees

Gregory Bruell	gob@wellfleet.com
Ross Callon	rcallon@wellfleet.com
David Crocker	dcrocker@mordor.stanford.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Kenneth Goodwin	goodwin@a.psc.edu
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Don Hoffman	hoffman@eng.sun.com
Scott Kaplan	scott@wco.ftp.com
Frank Kastenholz	kasten@ftp.com
Andrew Knutsen	andrewk@sco.com
Janet L. Marcisak	jlm@ftp.com
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
William Nowicki	nowicki@legato.com
Charles Perkins	perk@watson.ibm.com
Edward Reed	eer@cinops.xerox.com
Bradley Rhoades	bdrhoades@mail.mmmg.com
Steven Richardson	sjr@merit.edu
John Veizades	veizades@apple.com
Les Wibberley	lhw24@cas.org

2.7.5 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 e2e-interest@isi.edu mailing list.

The TCP Large Windows Working Group is being resurrected for a one time meeting, to review and if appropriate, approve this new document.

Goals and Milestones:

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"

2.7.6 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:

Mar 1991 Verify the interoperability of TNFS implementations at the 1992 NFS Connectathon.

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.

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.

Nov 1991 — Publish a Proposed Standard following the July meeting.

Mar 1992 — Request IESG to make the revised document a Draft Standard.

Internet-Drafts:

"A Specification of Trusted NFS (TNFS) Protocol Extensions", 07/23/1991,
Fred Glover <draft-ietf-tnfs-spec-03.txt>

2.8 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/Information Sciences Institute

Three BOFs (Birds of a Feather) and ten Working Groups met in Columbus.

Gopher BOF (GOPHER)

This BOF focused on the Gopher protocol and insuring its existence in a well-defined, standard fashion. Current implementations/licensing issues were not discussed at this meeting as it is outside the scope of involvement of the Chairs and the BOF.

Comments on the current Gopher protocol (RFC1436) were fielded and discussed. Discussion and consensus of moving this BOF to the Integration of Internet Information Resources Working Group (IIIR) for formal actions was approved. The GOPHER Chairs and the IIIR Chair, Chris Weider, will report the results of this BOF to the people working on Gopher at the University of Minnesota.

Low Cost IP Hardware Wish List BOF (LOIP)

The purpose of this BOF was to draw together interested IETFers and multi-vendors to define the next generation of its router/modem products. The issues were mostly operational (functions, features, price, performance). The FARNET community is interested in these issues because it provides services across a spectrum of users (from cost-conscious schools and colleges to supercomputer centers).

World-Wide Web BOF (WWW)

The World-Wide Web has an established user base, many compatible implementations, and many new ideas. This BOF discussed how the current WWW standards should be put into the RFC process, and what are the future directions for the Web.

Tim Berners-Lee mentioned a few recent developments, including NCSA's "Mosaic" WWW client for X11, a windows client for MS-DOS machines, and the growth rate of access to the CERN server.

The three "standards" important for WWW at this stage include the Universal Resource Locators specification which defines the addressing syntax used by WWW, the Hypertext Markup Language (HTML) specification that defines the transport format for menus, hypertext and general on-line documentation, and the HTTP access protocol which is used by WWW servers.

Integrated Directory Services Working Group (IDS)

The IDS Working Group is chartered to facilitate the integration and interoperability of current and future directory services into a unified directory service. This work will unite directory services based on a heterogeneous set of directory services protocols (X.500, WHOIS++, etc.). In addition to specifying technical requirements for the integration, the IDS Group will also contribute to the administrative and maintenance issues of directory service offerings by publishing guidelines on directory data integrity, maintenance, security, and privacy and legal issues for users and administrators of directories.

The final draft of the X.500 Advanced Usages Survey was released and will be moved to RFC publication status. The final draft of FYI 11 revision "A Catalog of Available X.500 Implementations" will be out in three weeks for movement to RFC publication. The Pilot Project Catalog was not released in time for review and discussion for this IETF.

The Whois++ - X.500 Interoperability document was assigned to Chris Weider and Mark Prior. The Whois++ Implementation Catalog will be assigned at the next IETF in Amsterdam. The "Directory Administrator's Guide" has been assigned to a team of seven people and will be called, "The Directory Policy Handbook." A "Directory Users Rights" document will be discussed at the next IETF.

Integration of Internet Information Resources Working Group (IIIR)

IIIR is chartered to facilitate interoperability between Internet Information Services, and to develop, specify, and align protocols designed to integrate the plethora of Internet information services (WAIS, archie, Prospero, etc.) into a single "virtually unified information service".

The Agenda included: the introduction and correction of Minutes from last meeting, discussions on a Resource Transponder document, and the draft, "Vision of Integrated Information Services document." A taxonomy document was discussed, but not assigned at this session. The Gopher protocol work will be brought into IIIR.

Internet School Networking Working Group (ISN)

The Internet School Networking Working Group is chartered to facilitate the connection of the United States' K-12 (Kindergarten through 12th Grade) schools, public and private, to the Internet, and school networking in general.

Gene Hastings led a discussion on his draft document of connectivity models. Jennifer Sellers discussed ISN's FAQ draft document. Jon Postel described the history and concept of domain names, and the United States domain. K12 placement in the naming system was discussed. Art St. George and Connie Stout went over new/updated ISN goals and milestones.

The Network Information Services Infrastructure Working Group (NISI)

NISI 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.

The NISI Working Group discussed the current relationships between NICs of different "levels" especially as they relate to the roles of the new InterNIC and to NISI. From this discussion, new goals were defined, including writing an informational paper describing these NIC roles and relationships and starting work toward developing guidelines for coordinating cooperative user assistance when more than one NIC is involved.

Network Training Materials Working Group (TRAINMAT)

The Network Training Materials Working Group is chartered to enable the research community to make better use of the networked services. Towards this end, the Working Group will work to provide a comprehensive package of "mix and match" training materials for the broad academic community which will: 1) enable user support staff to train users to use the networked services and 2) provide users with self-paced learning material. In the first instance, it will not deal with operational training. This Working Group is the IETF component of a joint RARE/IETF group working on Network Training Materials.

Thirty people attended the Working Group session to discuss training. This is the first official meeting of the Group following a BOF at the last IETF in Washington, D.C. The Group discussed a current, on-going project at the University of Newcastle, to catalog training materials and agreed to use this work as the basis for an informational RFC on training materials. Jill Foster lead the discussion on a template to be used to catalog the materials. The template has been developed based on input from several other projects, including NIR, IAFA, and the Coalition for Networked Information's TopNode. The Group had open discussions on what is needed in training materials and what projects are being worked on. The Group agreed to work on an informal posting system to the USWG mailing list to share information about training aimed at the trainers. (This Working Group is a joint project with RARE ISUS.)

Networked Information Retrieval Working Group (NIR)

NIR is chartered to increase the useful base of information about networked information retrieval tools, their developers, interested organizations, and other activities that relate to the production, dissemination, and support of NIR tools. NIR is a cooperative effort of the IETF, RARE, and CNI.

Jane Smith and Jill Foster provided updates on the Clearinghouse for Networked Information Discovery and Retrieval and the RARE ISUS Working Group NIR activities respectively. The Group reviewed the draft Status Report on NIR Tools and Groups (version 3.0) and reviewed the templates in light of experience in completing them. Dates were scheduled for the next revision and the next step of evaluation of the various tools.

Uniform Resource Identifiers Working Group (URI)

URI is chartered to define a set of standards for the encoding of system independent resource location and identification information for the use of Internet information services.

URI held three sessions at this IETF. Sessions one and two focused on Uniform Resource Locators (URLs). In the first session, Tim Berners-Lee led a discussion of URLs and review of the draft URL document. Incorporation of final changes to the draft were discussed, and it is hoped that the document will be an Internet-Draft before the Amsterdam meeting in July. The second session included presentation and discussion of proposals for URIs (open admission) and discussion of transitioning to UR*s. It was agreed that the URI document will be ready for the Amsterdam IETF. Session three focused on future plans beyond URLs and URIs.

After much spirited discussion, URI was able to finalize both the URL documents (which will now be submitted as Internet-Drafts) and a conceptual and syntactical framework for Uniform Resource Names (URNs). Documents will be written for URNs and should be submitted to the list in the next couple of weeks. Work continues to define elements needed for these objects to be useful to the end-user.

User Documents Revisions (USERDOC2)

The USERDOC2 Working Group is preparing a revised bibliography of on-line and hard copy documents, reference materials, and training tools addressing general networking information and how to use the Internet. The target audience includes those individuals who provide services to end users and end users themselves.

The short bibliography is now complete and has been submitted to the User Services Area Director as a first step in becoming an official FYI RFC. The "Introducing the Internet" document archive which contains the files documented in the short bibliography is on-line at four sites and four other sites have expressed interest in mirroring the files for anonymous FTP. Merit has the files in Gopher and WAIS. The group's next project is completing a comprehensive bibliography to supplement RFC1175. A first draft was distributed and discussed, with a final version to be completed by the next IETF in Amsterdam. Future goals were considered, including a new FYI aimed at those who are not connected to the Internet, a documents location system aimed at making materials in the long bibliography available, and consideration of developing materials aimed specifically at making information available to help librarians. The first of these projects will be started by looking at materials already produced by several Working Group members.

User Services Working Group (USWG)

The USWG 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.

Joyce Reynolds reported on the IETF User Services Area activities including: working groups coming to closure, new working groups starting up, new publications, and current User Services related Internet-Drafts postings. Discussion also focused on the "new NICs" (i.e., the InterNIC and the proposed/experimental APNIC - Asia/Pacific NIC). Jill Foster presented an update on RARE activities, including a report on the RARE Information Services/User Services (ISUS) activities. Susan Calcari presented a talk on the new InterNIC.

Gary Malkin led a session on two FYI RFC updates; FYI 7 "FYI on Questions and Answers: Answers to Commonly Asked "Experienced Internet User" Questions" (Also RFC1207), February 1991 and FYI 4 "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions" (Also RFC1325), May 1992.

Whois and Network Information Lookup Service (WNILS)

The purpose of WNILS is to expand and define the standard for WHOIS services, to resolve issues associated with the variations in access, and to promote a consistent and predictable service across the network.

Peter Deutsch led a review of Whois++ - Architecture. Dave Crocker briefly presented his two drafts to the WNILS session. Chris Weider led a review of Distributed Whois++ Model - Centroids, and Jim Fullton lead a review of Front End to Database Integration.

Discussion of Projects: A simple server, centroid and client will be available by April 30. Sources will be available on ftp.cnidr.org. The Working Group also discussed and revised their goals and milestones.

CURRENT MEETING REPORT

Reported by Joyce Reynolds/Information Sciences Institute

Minutes of the Gopher BOF (GOPHER)

This BOF focused on the Gopher protocol and insuring its existence in a well-defined, standard fashion. Current implementations/licensing issues were not discussed at this meeting as it was outside the scope of involvement of the Chairs and the BOF.

Comments on the current Gopher protocol (RFC1436) were fielded and discussed. A game plan was discussed on where this Group would like to go. The Gopher protocol is one of various documents which have not come into the "standardization" process. The intent is to start with this informational document, make some minor additions to bring it into agreement with the protocol as it is currently deployed, and place it into the protocol standardization track of the IETF. Discussion also focused on a liaison "reporter" that could interact between the IETF participants and the Gopher folks at the University of Minnesota, as they were unable to attend this IETF.

Discussion and consensus of moving this BOF to the Integration of Internet Information Resources Working Group (IIIR) for formal actions was approved. IIIR will assist by providing integration guidance of the Gopher protocol with other services, progressing RFC 1436 into the protocol standardization track, and looking into the next generation of Gopher protocols.

The GOPHER and IIIR Chairs will report the results of this BOF to the Gopher developers at the University of Minnesota.

Formal Declaration to be sent to the Gopher attendees at the Gopher Conference to be held April 13th:

> The intended collaboration between the IETF and the Gopher group is welcomed and necessary. The Gopher Protocol document (RFC1436) is nearly complete. The consensus of the attendees at the Gopher BOF last week is that minor corrections need to be made to the document that will bring the paper into compliance with existing working software and thus put it on the standardization track in the IETF. Discussion and consensus of moving this BOF to the IIIR (Integration of Internet Information Resources) Working Group for formal actions was approved. We look forward to working with you.

Attendees

N. Akiko Aizawa	akiko@nacsis.ac.jp
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jules Aronson	aronson@nlm.nih.gov
Robert Beer	r-beer@onu.edu
Tim Berners-Lee	timbl@info.cern.ch
Fred Bohle	fab@interlink.com
Deborah Boyer	djb@noc.ans.net
Sandy Bryant	slb@virginia.edu
Susan Calcari	calcaris@cerf.net
Corinne Carroll	ccarroll@bbn.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
David Conklin	conklin@jvnc.net
Naomi Courter	naomi@concert.net
Mark Davis-Craig	mad@merit.edu
Peter Deutsch	peterd@bunyip.com
Kishan Dudkikar	kishan@icm1.icp.net
Jill Foster	Jill.Foster@newcastle.ac.uk
Paul Franchois	paulf@bldrdoc.gov
Marcello Frutig	frutig@rnp.impa.br
Jim Fullton	Jim.Fullton@cnidr.org
Kevin Gamiel	kevin.gamiel@concert.net
Greg Gicale	gicale@ohio.gov
Terry Gray	gray@cac.washington.edu
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Erik Jul	jul@oclc.org
Bill Kelly	kellywh@mail.auburn.edu
Michael Khalandovsky	mlk@ftp.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Mark Kosters	markk@internic.net
John Kunze	jak@violet.berkeley.edu
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Hock-Koon Lim	lim@po.cwru.edu
Samir Malak	malaks@alice.uncg.edu
April Marine	april@atlas.arc.nasa.gov
Michael Mealling	michael@fantasy.gatech.edu
David Meyer	meyer@ns.uoregon.edu
Mitra	mitra@pandora.sf.ca.us
Joe Pagan	jrp@afterlife.ncsc.mil

Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Joyce K. Reynolds	jkrey@isi.edu
Francois Robitaille	francois.robitaille@crim.ca
Richard Rodgers	rodgers@nlm.nih.gov
Steve Romig	romig@cis.ohio-state.edu
Marshall Rose	mrose.iesg@dbc.mtview.ca.us
Larry Snodgrass	snodgrass@bitnic.educom.edu
Michael St. Johns	stjohns@darpa.mil
Craig Todd	ctodd@desire.wright.edu
Matt Tuttle	snodgras@cren.net
Robert Ullmann	Ariel@Process.com
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Richard Wiggins	wiggins@msu.edu

CURRENT MEETING REPORT

Reported by Mark Alexander Davis-Craig/Merit

Minutes of the Low-Cost IP Hardware Wish List BOF (LOIP)

Mark Davis-Craig chaired the BOF in place of Laura Breeden. The charge of the Group was to come up with a wish list of what it would like to see in a low-cost piece of IP hardware, and to give that list to FARNET.

The Group had the critical mass necessary for a 2.5 hour discussion with adequate representation from network providers, vendors, and people trying to connect K-12. The following recommendations were made:

- FARNET should submit an RFI to vendors with requirements stipulated by FARNET members (network providers) and not with requirements of those who may ultimately use said box for their Internet connectivity.
- An inventory of today's existing solutions should be taken.
- A matrix of costs and responsibilities should be created to help buyers of network access make good choices with respect to the balance between low-cost and low-maintenance.

With that data, FARNET should be able to bundle some "startup kits" of software and hardware for various environments to make it easier for small institutions to get connected.

The unedited comments were written on transparencies and the Chair will rewrite them later in a more coherent fashion. *Editor's Note (md): A complete listing of the "unedited" comments is available via ftp under loip-minutes-93mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Attendees

Mark Davis-Craig	`mad@merit.edu`
Sallie Fellows	`sallie@ed.unh.edu`
Terry Gray	`gray@cac.washington.edu`
John Hascall	`john@iastate.edu`
John Klensin	`klensin@infoods.unu.edu`
Daniel Long	`long@nic.near.net`
Paul Lustgraaf	`grpjl@iastate.edu`
Bruce Nelson	`bnelson@novell.com`
David O'Leary	`doleary@cisco.com`
Roy Perry	`rperry@advtech.uswest.com`

CURRENT MEETING REPORT

Reported by Tim Berners-Lee/CERN

Minutes of the World Wide Web BOF (WWW)

Status

Tim Berners-Lee mentioned a few recent developments:

- NCSA's "Mosaic" WWW client for X11 was agreed by the Group to be the best information access tool to date. This is at an advanced beta state (version 0.12). FTP pickup rate for this product now exceeds that for NCSA telnet.
- A Windows client for MS-DOS machines exists, the equivalent of XMosaic but for PCs. It is user configurable and looks good. A problem with release is that a runtime license is needed for the TCP stack used (Distinct).
- The growth rate of access to the CERN server has continued to double every four months for the last two years with no sign of tapering off.

Manual Pages on the Web

Steve Romig of Ohio State was congratulated on his server for Unix manual pages. This is a good example of an automatic server (a Perl script in this case) which provides a powerful data access. Steve agreed to make his tools available on the Web.

Web Agenda

There was a discussion of things which should be done to make it easier for the Web to spread. These include:

- Better packaging of CERN software, in particular the server: simpler installation, irrelevant file hiding, etc.
- Licensing of CERN software: There was deep concern expressed that investment by others in WWW related projects would be jeopardized later by CERN's licensing conditions, in the way that the University of Minnesota's have pulled the rug out from under the Gopher project. The General Public (Gnu-style) license was not so acceptable as pure public domain code. The Group was quite intent on solving this very quickly.
- A better vt100 full-screen browser is needed for a large dial-up user base.

NCSA is likely to bring out a public domain very simple to install server, as a partial solution to the first two issues issues.

There was a proposal to create a small group to implement the vt100 client, which formed and aims to implement the product for public release.

vt100 client: [Note: Shortly after the session ended, the Group received an announcement of the "Lynx" 2.0 beta release by Lou Montulli. This may solve the problem or at least provide a base. Lynx is a vt100 full screen hypertext WWW client. Check it out by telnetting to ukanaix.cc.ukans.edu and logging in as www. Source and binaries are available by anon FTP from acs220.cc.ukans.edu]

Standards

The three standards important for WWW at this stage are:

1. The Universal Resource Locators specification defines the addressing syntax used by WWW. This is to be released immediately as an Internet-Draft for the standards track.

2. The Hypertext Markup Language (HTML) specification defines the transport format for menus, hypertext and general on-line documentation. This must be specified at RFC standards track level in order to be submitted as a MIME content-type. This in turn is necessary as HTTP uses MIME format for returned multimedia messages, and we want it to stay MIME-compliant. HTML will therefore be released as an Internet-Draft as soon as possible.

3. HTTP is the access protocol used by WWW servers. An extended (backward-compatible) version is in use in the 2.0 and above CERN software. This is quite open to extension and so any comment on the protocol should be discussed on the www-talk list.

It was agreed to put these documents through the IIIR Working Group whose Chair has agreed to oversee their passage. The BOF will not, therefore, meet again or propose a WWW working group.

Mailing lists and Newsgroup

To join the mailing list, send a mail message to: www-talk-request@info.cern.ch. There is a newsgroup currently being formed. Its name is comp.infosystems.www. To vote on its creation, please send mail to either: www-yes@msen.com, or www-no@msen.com depending whether you wish the Group to be formed or not.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
N. Akiko Aizawa	akiko@nacsis.ac.jp

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Robert Beer	r-beer@onu.edu
Richard Bjers	rich.bjers@uc.edu
Thomas Brisco	brisco@pilot.njin.net
Sandy Bryant	slb@virginia.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
David Conklin	conklin@jvnc.net
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Mark Davis-Craig	mad@merit.edu
Alan Emtage	bajan@bunyip.com
Roger Fajman	raf@cu.nih.gov
Jill Foster	Jill.Foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
Marcello Frutig	frutig@rnp.impa.br
Jim Fullton	Jim.Fullton@cnidr.org
Maria Gallagher	maria@nsipo.nasa.gov
Kevin Gamiel	kevin.gamiel@concert.net
Joan Gargano	jcgargano@ucdavis.edu
Greg Gicale	gicale@ohio.gov
Terry Gray	gray@cac.washington.edu
Susan Harris	srh@umich.edu
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Inna Il'yasova	ilyasova@meg.uncg.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Neil Katin	katin@eng.sun.com
John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Edward Krol	e-krol@uiuc.edu
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
David Lovering	lovering@bldrdoc.gov
Bruce Mackey	brucem@cinops.xerox.com
Samir Malak	malaks@alice.uncg.edu
Kent Malave	kent@bach.austin.ibm.com
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Keith Moore	moore@cs.utk.edu
Masataka Ohta	mohta@cc.titech.ac.jp

Pete Percival	percival@indiana.edu
David Perkins	dperkins@synoptics.com
Marsha Perrott	mlp+@andrew.cmu.edu
Mark Prior	mrp@itd.adelaide.edu.au
Robert Raisch	raisch@ora.com
Joyce K. Reynolds	jkrey@isi.edu
Francois Robitaille	francois.robitaille@crim.ca
Carl Schoeneberger	70410.3563@Compuserve.com
Jennifer Sellers	sellers@nsipo.nasa.gov
Jane Smith	Jane.Smith@cnidr.org
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgrass@bitnic.educom.edu
Craig Todd	ctodd@desire.wright.edu
Matt Tuttle	snodgras@cren.net
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
Ruediger Volk	rv@informatik.uni-dortmund.de
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Les Wibberley	lhw24@cas.org

2.8.1 Integrated Directory Services (ids)

Charter

Chair(s):

Chris Weider, clw@merit.edu
Tim Howes, tim@umich.edu

Mailing Lists:

General Discussion: ids@merit.edu
To Subscribe: ids-request@merit.edu
Archive: merit.edu:~/pub/ids-archive

Description of Working Group:

The Integrated Directory Services Working Group is chartered to facilitate the integration and interoperability of current and future directory services into a unified directory service. This work will unite directory services based on a heterogeneous set of directory services protocols (X.500, WHOIS++, etc.). In addition to specifying technical requirements for the integration, the IDS Group will also contribute to the administrative and maintenance issues of directory service offerings by publishing guidelines on directory data integrity, maintenance, security, and privacy and legal issues for users and administrators of directories.

IDS will also assume responsibility for the completion of the outstanding Directory Information Services Infrastructure (DISI) Internet-Drafts, which are all specific to the X.500 protocol, and for the maintenance of FYI 11, "A catalog of available X.500 implementations".

IDS will need to liase with the groups working on development and deployment of the various directory service protocols.

The IDS Working Group is a combined effort of the Applications Area and the User Services Area of the IETF.

Goals and Milestones:

Ongoing	Track emerging directory service protocols to specify standards for interoperation with existing protocols.
Ongoing	Liase with groups working on deployment and development of directory services to locate and fix interoperability problems.
Ongoing	Identify unfilled needs of directory service offerers, administrators, and users.
Done	Submit to the IESG the DISI "Advanced Usages of X.500" paper as an informational document.

Jun 1993	Submit to the IESG the 1993 revision of FYI 11, "A catalog of available X.500 implementations" as an informational document.
Jul 1993	Submit as an Internet-Draft a "Specifications for interoperability between WHOIS++ and X.500".
Jul 1993	Submit as an Internet-Draft a "Guide to administering a directory service", which covers data integrity, maintenance, privacy and legal issues, and security.
Jul 1993	Submit as an Internet-Draft a "Catalog of available WHOIS++ implementations".
Jul 1993	Post the "X.500 Pilot Project Catalog" paper as an internet Draft.
Nov 1993	Submit to the IESG the DISI "X.500 Pilot Project Catalog" paper as an informational document.
Nov 1993	Submit to the IESG the "Specifications for interoperability between WHOIS++ and X.500" as a standards document.
Nov 1993	Submit as an Internet-Draft a "User's guide to directory services on the Internet".
Mar 1994	Submit to the IESG the "Guide to administering a directory service" as an informational document.
Mar 1994	Submit to the IESG the 1994 revision of FYI 11.
Mar 1994	Submit to the IESG the "Catalog of available WHOIS++ implementations" as an informational document.

Internet-Drafts:

"A Survey of Advanced Usages of X.500", 10/07/1992, Chris Weider, Russ Wright <draft-ietf-ids-x500-survey-02.txt>

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Integrated Directory Services Working Group (IDS)

Review of Charter and Scope

The Charter and scope of the IDS Working Group were reviewed and generally accepted. IDS will deal with issues common to all types of directory service (e.g., security/privacy/data maintenance issues), as well as issues of integration between directory services, and user-oriented directory service issues (e.g., the implementation surveys).

Progression of Documents

"Advanced Usages of X.500" - This document has been published as an Internet-Draft. The authors got about 20 responses.

"X.500 Implementation Survey" - Progress is being made, but the authors feel that they have gotten as many responses as they are going to get from their initial call. The next step is to contact people who submitted surveys to the original document and others known to have X.500 implementations and to try to "encourage" them directly to submit.

"X.500 Pilot Project Catalog" - Still waiting for responses from NADF and NIST on whether they would like to be included in this document. April Marine, the document's original author, has agreed to take over this document again.

Status of WHOIS++

Chris Weider and Peter Deutsch gave a quick status report on WHOIS++. Peter will fill in some holes in the specification in the next couple of weeks, after which the protocol documents will be submitted as Internet-Drafts. An initial implementation will be available by April 30, 1993, and will include a server and clients for several platforms.

New Documents

"WHOIS++ and X.500 Interoperability Guide" - There was general agreement that although this was a useful document, it was premature to begin writing it since WHOIS++ has not been officially released. Mark Prior, who has been writing a WHOIS++ server that accesses X.500, and Chris Weider agreed to revisit this topic off-line once Mark's server and the WHOIS++ implementation are further along.

"Catalog of WHOIS++ Implementations" - Ditto here, though Peter volunteered to be the contact point for collecting this information. When there is enough to warrant an RFC, he will put one out.

"Directory Security and Privacy Policy Administrators Guide" - There was much discussion on this topic, and many volunteers to help write it. Among the issues people want to see addressed:

- Security
- Notification of inclusion
- Source of data available to users
- Character sets
- Support of expected load
- Schema management
- Authentication
- Application gateways
- Update policy
- Access control
- Bulk access control
- Good neighbor policy
- Data management techniques
- Organizational ownership of data
- Privacy issues
- Sources of documentation

Since there were so many issues people wanted covered, the title of the document was changed to "Directory Policy Handbook: A guide to good directories." The authors include but are not limited to:

- Tim Howes
- Barbara Jennings
- Marco Hernandez
- Chris Weider
- Mark Prior
- and perhaps Virginia Razmersky

Anyone wishing to contribute to this document as an author is encouraged to sign on as soon as possible.

Date of Next Meeting

The next meeting of the IDS Working Group will be at the July IETF in Amsterdam, The Netherlands.

Attendees

N. Akiko Aizawa	akiko@nacsis.ac.jp
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Paul Barker	p.barker@cs.ucl.ac.uk
Cyrus Chow	cchow@ames.arc.nasa.gov
James Conklin	jbc@bitnic.educom.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Peter Deutsch	peterd@bunyip.com
Taso Devetzis	devetzis@bellcore.com
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Urs Eppenberger	eppenberger@switch.ch
Jonathan Fellows	jonf@gdstech.grumman.com
Kevin Gamiel	kevin.gamiel@concert.net
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Marco Hernandez	marco@mh-slip.cren.edu
Jeroen Houttuin	houttuin@rare.nl
Tim Howes	tim@umich.edu
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Kevin Jordan	Kevin.E.Jordan@cdc.com
David Katinsky	dmk@pilot.njin.net
Steve Kille	S.Kille@isode.com
Andrew Knutsen	andrewk@sco.com
Mark Kosters	markk@internic.net
Sylvain Langlois	Sylvain.Langlois@exp.edf.fr
Ronald Lanning	lanning@netltm.cats.ohiou.edu
David Lovering	lovering@bldrdoc.gov
Ignacio Martinez	martinez@rediris.es
Mitra	mitra@pandora.sf.ca.us
Clifford Neuman	bcn@isi.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
Mark Prior	mrp@itd.adelaide.edu.au
Sheri Repucci	smr@merit.edu
Robert Reschly	reschly@brl.mil
Francois Robitaille	francois.robitaille@crim.ca
Richard Rodgers	rodgers@nlm.nih.gov
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Srinivas Sataluri	sri@qsun.att.com
Chris Shaw	cshaw@banyan.com
Karen Shay	kshay@cs.cmu.edu
Gregory Sheehan	gregory.c.sheehan@att.com
Charlie Smith	crsmith@osvi.edu
Mark Smith	mcs@umich.edu

Sue Smith	smiths.es.net
Larry Snodgrass	snodgrass@bitnic.educom.edu
Stuart Stanley	stuarts@apertus.com
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Les Wibberley	lhw24@cas.org
Russ Wright	wright@lbl.gov

2.8.2 Integration of Internet Information Resources (iiir)

Charter

Chair(s):
Chris Weider, clw@merit.edu

Mailing Lists:
General Discussion: iiir@merit.edu
To Subscribe: iiir-request@merit.edu
Archive: merit.edu:~/pub/iiir-archive

Description of Working Group:

The Integration of Internet Information Resources Working Group (IIIR) is chartered to facilitate interoperability between Internet Information Services, and to develop, specify, and align protocols designed to integrate the plethora of Internet information services (WAIS, ARCHIE, Prospero, etc.) into a single "virtually unified information service" (VUIS). Such protocols would include (but are not limited to) update protocols for distributed servers, a 'query routing protocol' to pass queries between existing services, protocols for gateways between existing and future services, and standard exchange formats (perhaps based on Z39.50) for cross-listing specific information.

Also, where necessary, IIIR will create technical documentation for protocols used for information services in the Internet.

Goals and Milestones:

Ongoing	Track emerging Internet information services in order to specify technical requirements for their integration into the VUIS.
Ongoing	Liaise with other groups working on deployment and integration of Internet information services: e.g., The Coalition for Networked Information, RARE Working Group 3, etc.
Ongoing	Create specifications for interoperability between Internet information systems.
Done	Post an Internet-Draft on 'A vision of integrated information resources.'
Jul 1993	Post an Internet-Draft on 'Taxonomy of Internet Information Services.'
Jul 1993	Submit final version of 'A vision of integrated information resources' to the IESG as an Informational RFC.
Jul 1993	Submit final version of 'Taxonomy of Internet Information Services' to the IESG as an Informational RFC.
Nov 1993	Post an Internet-Draft defining common exchange formats.

Nov 1993 Post an Internet-Draft defining a Query Routing Protocol.

Mar 1994 Submit final version of common exchange format to the IESG as a Proposed Standard.

Jul 1994 Submit final version of Query Routing Protocol to the IESG as a Proposed Standard.

Internet-Drafts:

"Resource Transponders", 03/22/1993, C. Weider <draft-ietf-iiir-transponders-00.txt>

"A Vision of an Integrated Internet Information Service", 03/26/1993, C. Weider, P. Deutsch <draft-ietf-iiir-vision-00.txt>

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Integration of Internet Information Resources Working Group (IIIR)

Agenda

- Introduction and Correction of Minutes from Last Meeting.
- Resource Transponders.
- Vision of Integrated Information Services Document.
- Start Work on Taxonomy Document.
- Internet Information Framework.
- AOB.

There were no corrections offered to the Minutes from Washington.

Resource Transponders

Chris Weider discussed his Resource Transponder Internet-Draft. Scott Kaplan of FTP mentioned that a similar concept was already under development in the Service Location Protocol Working Group for modems and other hardware; Chris agreed to look at the SVRLOC stuff to see if there was any synergy.

Vision of Integrated Information Services Document

Peter Deutsch presented the Vision of Integrated Information Services document that he and Chris Weider had developed. It was mentioned that this architecture provides a trivial level of interoperability between currently deployed services, and can provide a foundation for integrating new services. Also, in the discussions of the 'Axioms of Information Services', an axiom was added by the Group: "Users want to have the ability to look under the hood".

Start Work on Taxonomy Document

The taxonomy document's contents were discussed; it was finally settled that it would categorize the various available tools by functionality (along several different schemes) and would point out gaps in the current tool mix in light of the architecture provided in the Vision document. Although there was not a universal consensus that such a document was needed, there was a rough consensus that such a document could teach us a lot about what tools we have and what tools we need. Peter Deutsch and Chris Weider volunteered to take a first cut at the paper for Amsterdam.

Internet Information Framework

Chris Weider discussed the Internet Information Framework developed by the IESG and several working group Chairs. This framework lays out the administrative structure for working groups working in the information services area. In particular, a new working group will not be formed for every information services protocol developed; instead IIIR will take the responsibility of documenting new and existing protocols and shepherding their standardization. Examples of the protocols IIIR will cover include GOPHER (RFC1436), World Wide Web, and Z39.50. As there are active communities for these protocols, IIIR will simply help the standardization process and may also serve as a mediator if many incompatible versions of the same base protocol appear. For those new protocols which are needed by the integration effort or are identified in the taxonomy paper, IIIR will work with the IESG to determine if a new working group should be formed or if the work can be done in IIIR.

Attendees

N. Akiko Aizawa	akiko@nacsis.ac.jp
Jules Aronson	aronson@nlm.nih.gov
Robert Beer	r-beer@onu.edu
Tim Berners-Lee	timbl@info.cern.ch
Cyrus Chow	cchow@ames.arc.nasa.gov
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Mark Davis-Craig	mad@merit.edu
Peter Deutsch	peterd@bunyip.com
Alan Emtage	bajan@bunyip.com
Jill Foster	Jill.Foster@newcastle.ac.uk
Steven Foster	foster@cs.unr.edu
Marcello Frutig	frutig@rnp.impa.br
Jim Fullton	Jim.Fullton@cnidr.org
Kevin Gamiel	kevin.gamiel@concert.net
Joan Gargano	jcgargano@ucdavis.edu
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Alisa Hata	hata@cac.washington.edu
Susan Horvath	shorvath@merit.edu
Scott Kaplan	scott@wco.ftp.com
Michael Khalandovsky	mlk@ftp.com
Edward Krol	e-krol@uiuc.edu
John Kunze	jak@violet.berkeley.edu
Hock-Koon Lim	lim@po.cwru.edu
Samir Malak	malaks@alice.uncg.edu
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov

Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Clifford Neuman	bcn@isi.edu
Geir Pedersen	Geir.Pedersen@usit.uio.no
Pete Percival	percival@indiana.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Steven Richardson	sjr@merit.edu
Richard Rodgers	rodgers@nlm.nih.gov
Mark Smith	mcs@umich.edu
Sue Smith	smiths.es.net
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu

2.8.3 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: archive.cc.mcgill.ca:~/pub/iafa-archive

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.

(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:

Done	First IETF Meeting: review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended procedures and impact on site administrators. Assign writing assignments for the first draft of the documents.
Mar 1992	Review first draft and determine necessary revisions. Follow up discussion will occur on mailing list.

Jun 1992 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.

2.8.4 Internet School Networking (isn)

<u>Charter</u>

Chair(s):

John Clement, clement@educom.edu
Arthur St. George, stgeorge@bootes.unm.edu
Connie Stout, cstout@tenet.edu

Mailing Lists:

General Discussion: isn-wg@unmvma.unm.edu
To Subscribe: listserv@unmvma.unm.edu
In Body: subscribe isn-wg <first name> <last name>
Archive:

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 post-secondary 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:

Done — 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 Art St. George/UNMexico

Minutes of the Internet School Networking Working Group (ISN)

The Internet School Networking session was chaired by Art St. George and Connie Stout.

Jennifer Sellers, representing the NASA NREN Group, said that the K-12 networking Frequently Asked Questions document would be made into a standard before the Amsterdam IETF. She solicited comments for changes and said she will post the document to Kidsnet and other mailing lists for review.

A quick poll revealed that there were no K-12 educators in the room, with Connie Stout having the most direct involvement. Many in the room were involved in helping to get K-12 schools on the Internet, either as consultants, or as part of their jobs as network support people for universities or regional networks, or out of personal interest. The observation was made that people supporting K-12 networking needed to continue their dialogue with the K-12 community, since their physical attendance at IETF was unlikely. Connie observed that some technical staff working for K-12 schools might attend future IETFs.

There was some discussion of how K-12 schools should pursue funding for networking. Several observed that in some cases funding can be easy to come by with the right demonstration of need. Some attendees pointed out that inter-district funding disparities affect ability to network. Bill Manning, Rob Raisch, and a couple of other attendees said they would help devise a list of benefits of networking. Ray Perry of US West said his organization has prepared a video that shows advantages of network connectivity.

Connie stated that she has met with twenty-two institutions involved with K-12 networking initiatives. Someone expressed a need for a list of all K-12 initiatives under way. Another attendee asked if private and parochial schools should be part of these initiatives; Connie's answer was "yes."

A distinction was raised between promotional versus "how to" documents; there is a need for both.

It was observed that there is a wide range of skill levels among those interested in starting K-12 networks, and that there is a bootstrap process required to bring new schools and support staff on-line. There was some discussion on how to spread the word more broadly. One suggestion was giving talks at teacher conferences. Brochures from the Consortium for School Networking were handed out. (Connie Stout is the Chair; St. George is the Secretary- Treasurer. Their Listserv is COSNDISC@bitnic.bitnet; subscribe via a mail message with "SUBSCRIBE COSNDISC First_name Last_name". General email goes to cosn@bitnic.bitnet.)

Gene Hastings handed out a document that lists a variety of connection options for K-12 schools. The document was a menu of options from dialup to Frame Relay and all points in between. There was some concern this document would overwhelm new sites. Consensus was eventually reached that the document fills a definite need and it should be refined and enhanced with another, more general document. Gene mentioned documents available for anonymous ftp from ftp.cc.berkeley.edu, under /k12.

Connie observed that the average school has two phone lines serving twenty-two instruments. Connectivity can be a challenge.

John Postel, one of the architects of the Domain Name System, spoke on the evolution of DNS and the challenges presented by growing use of the name space. He proposed a model for delegating management of K-12 names to the states, where state authorities would manage names of the form:

computer.school.k12.state.us

This announcement led to some lively discussions. Connie observed that the "edu" domain was being used for universities exclusively in that model, and that K-12 educators felt that, as educators, the name belonged to them as well. She suggested that universities ought to use a new "uni" domain, and leave "edu" for the schools. This led to some discussion as to how this is handled abroad. In the U.K. it's ".ac" for "academic community." A gentleman from Germany told their practice (which sounded like the word "school" is used). Connie observed that statewide education networks also need to be given .edu names; e.g., tenet.edu. Bill Manning pointed out that the question of whether domain names should be geographical or organizational is an old one. With 10 million Internet nodes on-line now and another 10 million coming something must be done to manage the namespace in a distributed fashion.

Bruce Nelson spoke on ISOC/K12 committee issues. He said there is a need for K-12 focus within the Internet Society – someone who serves as spokesperson. He said the Group needs to be sure to include an international perspective. He does not see the need for an advocacy group per se, but he does see a need to infuse ISOC with a K-12 perspective.

There was more discussion of the various initiatives for K-12 networking. Art St. George asked if there was a need for a registry of consultants/advisers who can help new efforts. Discussion was inconclusive on this point.

Attendees

Robert Beer	r-beer@onu.edu
John Boatright	bryan_boatright@ksc.nasa.gov
Thomas Brisco	brisco@pilot.njin.net
Al Broscius	broscius@bellcore.com

Sandy Bryant	slb@virginia.edu
Susan Calcari	calcaris@cerf.net
David Carr	Carr@acsu.buffalo.edu
Corinne Carroll	ccarroll@bbn.com
Douglas Carson	carson@utcc.utoronto.ca
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Mark Davis-Craig	mad@merit.edu
Dale Dougherty	dale@ora.com
Hans Eriksson	hans@sics.se
Sallie Fellows	sallie@ed.unh.edu
Dale Finkelson	dmf@westie.mid.net
Marcello Frutig	frutig@rnp.impa.br
Maria Gallagher	maria@nsipo.nasa.gov
Greg Gicale	gicale@ohio.gov
Terry Gray	gray@cac.washington.edu
Robert Gutierrez	gutierre@nsipo.nasa.gov
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Raymond Harder	rharder@eis.calstate.edu
Susan Harris	srh@umich.edu
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Barbara Jennings	bjjenni@sandia.gov
John Klensin	klensin@infoods.unu.edu
Edward Krol	e-krol@uiuc.edu
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
Paul Lustgraaf	grpjl@iastate.edu
Gary Malkin	gmalkin@xylogics.com
Bill Manning	bmanning@sesqui.net
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Cynthia Mills	cmills@bbn.com
Bruce Nelson	bnelson@novell.com
William Owens	owens@acsu.buffalo.edu
Brad Parker	brad@fcr.com
Marsha Perrott	mlp+@andrew.cmu.edu
Roy Perry	rperry@advtech.uswest.com
Jon Postel	postel@isi.edu
Robert Raisch	raisch@ora.com
Joyce K. Reynolds	jkrey@isi.edu

Jennifer Sellers	sellers@nsipo.nasa.gov
Charlie Smith	crsmith@osvi.edu
Robert Ullmann	Ariel@Process.com
Mario Vecchi	mpv@thumper.bellcore.com
Ruediger Volk	rv@informatik.uni-dortmund.de
Richard Wiggins	wiggins@msu.edu
Liang Wu	ltw99@bellcore.com

2.8.5 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@reston.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.
Done	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.
Done	Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI RFC for publication.

Internet-Drafts:

"FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices", 11/11/1992, R. Enger, J. Reynolds <draft-ietf-noctool2-debug-tcpip-00.txt>

2.8.6 Network Information Services Infrastructure (nisi)

Charter

Chair(s):

April Marine, april@atlas.arc.nasa.gov
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. A primary goal of the Group is to facilitate the development of relationships between NICs that will result in the presentation of a seamless user support service. NISI will work with all NICs, including the InterNIC, to achieve the goal of a fully-functioning, cooperative mesh of worldwide NICs. In addition to creating policies for interaction, NISI will 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	Complete draft for phase 2 suggesting cooperative agreements for NICs.
Done	Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of 'cooperative agreements' which will allow NICs to work together more effectively to serve users.
Done	Revised draft document ready for Working Group review. Document defines NIC functions and suggests some standardizations for NIC services, as well as offers new mechanisms for exchanging information between NICs.
Done	Document submitted as Internet-Draft for comment from a wider Internet audience.
Done	Working Group discussed current Internet-Draft and suggested minor revisions. Decision made to continue Working Group activity beyond this document.
Done	First document released as informational RFC. Outline and discuss new NISI tasks at IETF meeting.

Done	Write a document explaining the security issues of privacy and accuracy in Internet databases. Publish as an informational RFC.
Jun 1993	Post an Internet-Draft describing NIC interelationships.
Jun 1993	Post an Internet-Draft of a NIC user handoff procedure based on the UCP work.
Jun 1993	Post an Internet-Draft describing accessing the nic-profiles data in the X.500 database at Merit.
Dec 1993	Submit the NIC Interelationship document to the IESG for consideration as an FYI RFC.
Dec 1993	Submit the User Handoff procedures to the IESG for consideration as an FYI RFC.
Dec 1993	Submit the Nic-Profiles paper to the IESG for consideration as an FYI 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 Marsha Perott/PREPNet

Minutes Network Information Services Infrastructure Working Group (NISI)

Agenda

- Introductions and Agenda Modifications
- NIC-profiles Status
- Net-help Status
- Coordination with New InterNIC Information Services
- New Projects
- New Goals and Milestones
- Summary and End

NIC-profiles

Since the last meeting, three volunteers (Marsha Perrott, Bill Yurcik, Charlotte Mooers) have been working to populate this database of information about NIC contacts. They have sent email to many NIC sites asking contacts to fill out an information template. There are now about forty-five entries in the database.

The database is currently maintained at Merit. To access the information there, send a message to x500test@merit.edu with "help" as the subject of the message and you will be returned access directions.

The template will be put on-line at merit.edu in /pub/nisi/nic-template.

The future home of the database is not clear. Some discussion with Susan Calcari of the new InterNIC prompted ideas such as eventually having each NIC maintain a template of information locally which an archie-like service would update and make available for distributed access. Or possibly the whole database would just move to the new InterNIC. Certainly, it was agreed that the most important goal at this time is simply to collect the information.

Jerry Martin of Ohio State mentioned that he had been doing something similar in trying to poll university computer centers for NIC contacts, and he has already collected information on abut sixty-two sites. He kindly agreed to forward that information to merit.

Net-help

After having this project stagnate for a couple of meetings now, ye olde co-Chairs were going to vote to boot it off of the ol' goals and milestones. The suggestions received before the meeting for implementing the idea (which, by the way, is for a simple command to allow a user to see a referral to local help for network questions), seemed to involve quite a bit of change to network software, etc., or to demand quite a bit of a knowledge base from users before they could even issue a query.

In mid-boot however, a couple of good suggestions were given as to easy ways of getting and providing the information. Susan Calcari suggested that since the registration folks were just now making changes to their net registration templates, that it would be a relatively simple thing to ask them to add a "NIC Services Contact" to the template. That way, everyone requesting a new number would need to provide that information and any user doing a whois on his network can see it. Since whois is a common application, an extension to that seemed easier than inventing a whole new command. Plus there was a suggestion to implement a common address that anyone could telnet to, i.e., telnet help.merit.edu. Several other ideas were kicked around.

The Group recognizes that "any" effort in this area does assume that the user knows "something" to type first to get that NIC contact information, which is rather a leap of faith. But it does have to start somewhere. It is clear that this project is still alive and will probably require that NISI communicate with the Applications Area and figure out what to do next.

Coordination with New InterNIC Information Services

The Group took some time drawing pictures to make sure there was a clear understanding of the relationships between the InterNIC, existing NICs, and the NISI Working Group. The Group was especially trying to make sure that overlaps between NISI and the InterNIC were avoided. No problem. Basically the InterNIC is on a par with other NICs, e.g., the DDN NIC, the JNIC in Japan, the RIPE NCC, etc., while NISI is a forum outside of the NIC structure for people working in information services and user support to develop policies, procedures, and tools that would be generally useful for them all.

One area of effort that would seem to be really useful is to develop more formal procedures for coordinating user response between NICs, rather in the sense of a trouble ticket system that could be used among NICs when a user problem or question requires a coordination of services.

New Projects and New Goals and Milestones

Given the discussion during the meeting, new goals and milestones were created for the Group based on existing tasks and on new ones suggested by the Group.

Possible tasks were identified as:

- Implementing net-help.
- Developing guidelines for NIC hand offs and feedback regarding user questions.
- Developing a guide to NICs.

- Writing an informational paper describing the current structure and interrelationships of the various NICs (because, after all, if *we* were confused!...).

This last was seen to possibly overlap with IRG/Directory efforts and the nic-profiles, so it has been tabled for now until the Group has time to clarify it better and act on it. Of the others:

- Net-help.

  - Talk to NSI regarding adding a hook for NIC POC to their registration template. Done.
  - Have NIC point-of-contact information show in WHOIS. May 30th.
  - NISI Group to campaign for information to be added to this field retroactively for those sites that already have their net number. Start by May 30th.

  Susan Calcari and the Group will be working on this.

- NIC Structure Paper.

  - Define scope of paper, write outline. April 30th.
  - Write draft. June 15th.
  - Get feedback on paper by or at the July IETF.

  April Marine, Susan Calcari and Maria Gallagher have been assigned to this task.

- NIC Hand-off Procedures.

  - Review UCP work. July IETF.
  - Modify the UCP documents to fit NIC rather than NOC interaction. July IETF.
  - Come up with an outline. July IETF.

  April Marine, Pat Smith, Joyce Reynolds, Martyne Hallgren, Susan Calcari, and Jerry Smith will be working on this.

The above action items are the new goals and milestones for the Group. Well, of course, the milestones will be extended beyond the "draft" stages of these papers. The Group also approved the minor change to the Charter that was sent to the mailing list before the meeting.

Attendees

Corinne Carroll	ccarroll@bbn.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu

Cynthia Clark	cclark@cnri.reston.va.us
John Curran	jcurran@nic.near.net
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Marcello Frutig	frutig@rnp.impa.br
Maria Gallagher	maria@nsipo.nasa.gov
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Susan Horvath	shorvath@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Erik Jul	jul@oclc.org
Mary La Roche	maryl@cos.com
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Daniel Long	long@nic.near.net
Gary Malkin	gmalkin@xylogics.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Keith Moore	moore@cs.utk.edu
Jim Naro	jnaro@nic.near.net
Judy Nasar	jdnasar@magnus.acs.ohio-state.edu
Bruce Nelson	bnelson@novell.com
Clifford Neuman	bcn@isi.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Mark Prior	mrp@itd.adelaide.edu.au
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Francois Robitaille	francois.robitaille@crim.ca
Jennifer Sellers	sellers@nsipo.nasa.gov
Patricia Smith	psmith@merit.edu
Craig Todd	ctodd@desire.wright.edu
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.8.7 Network Training Materials (trainmat)

Charter

Chair(s):

Ellen Hoffman, `ellen@merit.edu`
Jill Foster, `Jill.Foster@newcastle.ac.uk`

Mailing Lists:

General Discussion: `us-wg@nnsc.nsf.net`
To Subscribe: `us-wg-request@nnsc.nsf.net`
Archive: `nnsc.nsf.net:~/nsfnet/us-wg*`

Description of Working Group:

Widespread familiarity with global network services and competence in using them brings benefit to individual users, enriches the information skills and resources of the community and optimises the return in investment in networked services.

The Network Training Materials Working Group is chartered to enable the research community to make better use of the networked services. Towards this end, the Working Group will work to provide a comprehensive package of "mix and match" training materials for the broad academic community which will: 1) enable user support staff to train users to use the networked services and 2) provide users with self-paced learning material. In the first instance, it will not deal with operational training.

This Working Group is the IETF component of a joint RARE/IETF group working on Network Training Materials.

The Working Group will create a catalogue of existing network training materials (using the TopNode cataloguing fields where appropriate), identify the gaps in Network Training Materials and work to identify the problems associated with hands on training workshops using networked services providing a real service.

Goals and Milestones:

Done	First Working Group meeting. Review and approve the Charter with a review of documents and materials to be written.
Jul 1993	Post the catalogue of training materials as an Internet-Draft.
Dec 1993	Submit the catalogue of training materials to the IESG for review and publication as an Informational RFC.

CURRENT MEETING REPORT

Reported by Ellen Hoffman/Merit

Minutes of the Network Training Materials Working Group (TRAINMAT)

Agenda

- Review of Updated Charter, Introduction of Chairs.
- Review and Discuss Catalog of Training Materials.
- Top Node-style Template for Collecting Training Information.
- Identify Documents/Materials that Need to be Produced.
- Liaisons with Other Groups.
- Training Activities by Working Group Members.

Thirty people attended the Working Group session to discuss training. This is the first official meeting of the Group following a BOF at the last IETF.

Review of Updated Charter, Introduction of Chairs

Jill Foster described changes to the Charter based on the discussions at the BOF held during the previous IETF meeting. Dates for production of a guide to training materials were adjusted to represent current projections. The Group agreed to continue to use the us-wg@nnsc.nsf.net mailing list rather than creating a new one.

Jill described her role heading ISUS in Europe and its relationship to this Working Group, which is a joint RARE/IETF initiative. She also described the ITTI Network Training Materials Project which she supervises, and distributed copies of their latest efforts. These are available via Anonymous FTP from mailbase.ac.uk in the /pub/itti-networks directory. Jill also described a two-year project to train subject specialists and tailor training materials for that Group. Ellen Hoffman described projects discussed at the Coalition for Networked Information meeting in the working group on Teaching and Learning, noting that end user training had recently been added to this Group's responsibilities within CNI.

Review and Discuss Catalog of Training Materials

The ITTI Training Project cataloging effort is over, but Jill recommended that it become the basis for this Group's document on training materials. The Group will focus on developing an FYI RFC on training materials, and agreed to work with the User Documents Revisions Working Group so as not to duplicate efforts in documenting materials. Jodi Chu and Jerry Martin agreed to assist in this cataloging effort. CNIDR agreed to house a database of training materials to make it easier to maintain and update the listings.

Top Node-style Template for Collecting Training Information

Jill noted she had been working with the CNI Top Node project as well as others (NIR, IAFA) to further develop a template for describing training materials. She showed a recent version, and comments were made to improve the template. Jill will make these changes. The template will become an appendix to the FYI RFC so others can use it. One area the Group agreed to work on was a controlled list of "Subject Headings." Lenore Jackson and Jennifer Sellers volunteered to help finalize these by looking at other projects, such as ERIC, that catalog training materials.

Identify Documents/Materials that Need to be Produced

Jill described the latest developments of the ITTI training project, which include a survey that was distributed, a basic training program done in PowerPoint that was available through the mailbase.ac.uk server along with speaker notes, and evaluation sheets. These can be used and modified by others doing training.

Michael Mealling described several examples where MUDs are being used for training. He gave an example of one for biochemists which includes a link to Gopher. Jennifer Sellers described MUSE based training projects. Jane Smith is collecting examples of MUSE based teaching projects.

Jill led a discussion of current needs for network training, which focused on the need for on-line training. Also discussed was the problem of overloading network resources when a large class tries to access a server en masse. Suggestions were made for client/server based training programs as one solution to this. This area was identified as one for potential future focus of the TRAINMAT Working Group. Jill also discussed the potential of Internet Talk Radio as a possible means of distributing training.

Liaisons with Other Groups

The Group discussed methods for distributing training information, noting that often mail groups become "noisy" so it becomes hard to sort out information. Susan Harris and Jane Smith volunteered to work out ways to summarize discussions and use the mailing list to send it out to the Working Group members.

Training Activities by Working Group Members

The Group discussed a project at the University of Newcastle to catalog training materials and agreed to use this work as the basis for an Informational RFC on training materials. Jill Foster lead the discussion on a template to be used to catalog the materials. The template has been developed based on input from several other projects, including the Networked Information Retrieval and Internet Anonymous FTP Archives Working Groups, and the Coalition for Networked Information's Top Node. The Group had open discussions on what is needed in training materials and what projects are being worked on. The Group agreed to work on an informal posting system to the User Services Working Group mailing list to share information about training aimed at trainers. (This Working Group is a joint project with RARE ISUS).

Attendees

Robert Beer	r-beer@onu.edu
Al Broscius	broscius@bellcore.com
Susan Calcari	calcaris@cerf.net
Corinne Carroll	ccarroll@bbn.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Mark Davis-Craig	mad@merit.edu
Dale Dougherty	dale@ora.com
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Marcello Frutig	frutig@rnp.impa.br
Jim Fullton	Jim.Fullton@cnidr.org
Maria Gallagher	maria@nsipo.nasa.gov
Greg Gicale	gicale@ohio.gov
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Susan Harris	srh@umich.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Hock-Koon Lim	lim@po.cwru.edu
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Michael Mealling	michael@fantasy.gatech.edu
Cynthia Mills	cmills@bbn.com
Mark Prior	mrp@itd.adelaide.edu.au
Joyce K. Reynolds	jkrey@isi.edu
Francois Robitaille	francois.robitaille@crim.ca
Jennifer Sellers	sellers@nsipo.nasa.gov
Jane Smith	Jane.Smith@cnidr.org
Craig Todd	ctodd@desire.wright.edu
Chris Weider	clw@merit.edu
Richard Wiggins	wiggins@msu.edu

2.8.8 Networked Information Retrieval (nir)

Charter

Chair(s):

Jill Foster, Jill.Foster@newcastle.ac.uk
George Brett, George.Brett@cnidr.org

Mailing Lists:

General Discussion: nir@mailbase.ac.uk
To Subscribe: mailbase@mailbase.ac.uk
In Body: subscribe nir <first name> <last name>
Archive: mailbase.ac.uk:~/pub/nir

Description of Working Group:

As the network has grown, along with it there has been an increase in the number of software tools and applications to navigate the network and make use of the many, varied resources which are part of the network. Within the past year and a half we have seen a wide spread adoption of tools such as the Archie servers, the Wide Area Information Servers (WAIS), the Internet Gopher, and the WorldWide Web (WWW). In addition to the acceptance of these tools there are also diverse efforts to enhance and customize these tools to meet the needs of particular network communities.

There are many organizations and associations that have recently begun to focus on the proliferating resources and tools for Networked Information Retrieval (NIR). The Networked Information Retrieval Group will be a cooperative effort of three major players in the field of NIR: IETF, RARE, and the Coalition for Networked Information (CNI) specifically tasked to collect and disseminate information about the tools and to discuss and encourage cooperative development of current and future tools.

The NIR Working Group intends to increase the useful base of information about NIR tools, their developers, interested organizations, and other activities that relate to the production, dissemination, and support of NIR tools, to produce documentation that will enable user services organizations to provide better support for NIR tools, to develop materials that will assist the support and training of end users and to evolve in the future as necessary to meet and anticipate changes in the field (i.e., NIR tools, protocols, network topology, etc.).

Goals and Milestones:

Done	Review and comment on proposed Charter. Discuss Applications Template and Organizational Template.

Sep 1992 Post an Internet-Draft containing the Applications and Organizational Templates.

Oct 1992 Post an Internet-Draft of the "Consumer Report" with introductory material and completed templates.

Dec 1992 Submit "Consumer Report" to the IESG for publication as an Informational RFC.

Internet-Drafts:

"A Status Report on Networked Information Retrieval: Tools and Groups", 03/24/1993, J. Foster, G. Brett, P. Deutsch <draft-ietf-nir-status-report-00.txt>

CURRENT MEETING REPORT

Reported by Jill Foster/University of Newcastle upon Tyne and George Brett/MCNC-CNIDR

Minutes of the Networked Information Retrieval Working Group (NIR)

The Networked Information Retrieval Working Group was convened by Jill Foster and Jane Smith (substituting for George Brett). The Agenda was discussed and approved as presented.

Updates of Other Projects

Updates on other related projects were presented by Jane and Jill. Jane presented material on the Clearinghouse for Networked Information Discovery and Retrieval (CNIDR), funded by the National Science Foundation and directed by George Brett. Jill presented an update of the RARE ISUS NIR activities.

Review of the Draft Status Report on NIR Tools and Groups (Version 3.0)

The Draft was reviewed and discussed. The Draft has been submitted as an Internet-Draft, to eventually become an FYI document. Some formatting work remains to be done to qualify as an official draft, but essentially the document is now in draft status.

Additional tools and groups were suggested and discussed, and some tools will be moved from the "coming attractions" appendix to the regular listings section. In particular, there was discussion about individually listing implementations of X.500 and Z39.50. Tim Berners-Lee will submit listings for X.500, and Jill suggested that someone from the ZIG submit listings for Z39.50. Additional groups suggested were: OCLC, USMARC, MARBI, LC, NLM and Paradise.

The Group voted to omit listings in Appendix F, "Other Interested Parties," but to provide a pointer to where such information might be obtained.

Review of the Templates in Light of Experience in Completing Them

There was little discussion about this. It was suggested that perhaps, if and when more people become involved in completing the templates, more feedback will be forthcoming.

Schedule Next Revision

The document as a whole will continue to be updated three times a year just prior to IETF meetings. CNIDR will split the document into logical pieces and make it available on-line both for retrieval and for more frequent updating.

Next Stage: "Evaluation" of the Various Tools

There was much discussion about the definition and depth of the "evaluation" portion of the report. Suggestions ranged from a simple checklist to a formal study of user interactions with the various programs. Volunteers were solicited to begin work on some of the basic elements of the evaluation portion.

- Checklist of features and capability.
- Glossary/definition of terms.
  - April Marine.
  - Jim Fullton.
  - Jill will solicit UK and European volunteers.

It was suggested that after this Group makes a first pass a separate group should review the work.

- Evaluation criteria

There was a lot of discussion about how to establish evaluation criteria. Some of the consensus is shown below in the breakdown of tasks.

- Document current, advanced, and innovative uses of the tools, giving 5-10 examples of each. Get information from tool developers and from mailing lists, newsgroups, etc.
  - April Marine (tentative).
  - Jane Smith.
  - Lenore Jackson.
  - Jill will solicit UK and European volunteers.
- Evaluations, both from the information consumer and the information provider points of view, will be approached first on an anecdotal basis, then perhaps on formal studies should they become available.
  - Erik Jul
  - Jill Foster
  - Chris Weider

All tasks listed above may be discussed on the mailing list and preliminary results will be compiled and issued as a *separate* report.

Identifying Other Tasks (need for documentation, training materials, etc.)

It was noted that this task overlaps the Network Training Materials Working Group somewhat, but that compiling a bibliography of documentation and training materials specific to NIR would be valuable. CNIDR will be collecting such materials and will assist both in distribution and in compiling a bibliography. Such a document will also be considered a separate part of the report.

Update Goals and Milestones in the Light of Current Discussion

Jill will be making suggested changes to this part of the NIR Charter, pursuant to discussions with the IESG. It was noted that the current list is more like an Agenda, when what is basically required is a schedule for publications.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Paul Barker	p.barker@cs.ucl.ac.uk
Robert Beer	r-beer@onu.edu
Tim Berners-Lee	timbl@info.cern.ch
John Boatright	bryan_boatright@ksc.nasa.gov
Fred Bohle	fab@interlink.com
Corinne Carroll	ccarroll@bbn.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Henry Clark	henryc@oar.net
David Conklin	conklin@jvnc.net
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Jill Foster	Jill.Foster@newcastle.ac.uk
Steven Foster	foster@cs.unr.edu
Marcello Frutig	frutig@rnp.impa.br
Gwen Funchess	funchess@magnus.acs.ohio-state.edu
Greg Gicale	gicale@ohio.gov
Robert Gutierrez	gutierre@nsipo.nasa.gov
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Susan Horvath	shorvath@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Thomas Johannsen	Thomas.Johannsen@ebzaw1.et.tu-dresden.de
Dale Johnson	dsj@merit.edu
Erik Jul	jul@oclc.org
Andrew Knutsen	andrewk@sco.com
John Kunze	jak@violet.berkeley.edu
Hock-Koon Lim	lim@po.cwru.edu
Samir Malak	malaks@alice.uncg.edu
Glenn Mansfield	glenn@aic.co.jp
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Michael Mealling	michael@fantasy.gatech.edu
Keith Moore	moore@cs.utk.edu
Jim Naro	jnaro@nic.near.net
Mark Needleman	mhn@stubbs.ucop.edu

Bruce Nelson	bnelson@novell.com
Clifford Neuman	bcn@isi.edu
Bill Norton	wbn@merit.edu
William Nowicki	nowicki@legato.com
Marsha Perrott	mlp+@andrew.cmu.edu
Mark Prior	mrp@itd.adelaide.edu.au
Steven Richardson	sjr@merit.edu
Francois Robitaille	francois.robitaille@crim.ca
Richard Rodgers	rodgers@nlm.nih.gov
Jennifer Sellers	sellers@nsipo.nasa.gov
Chris Shaw	cshaw@banyan.com
Sue Smith	smiths.es.net
Larry Snodgrass	snodgrass@bitnic.educom.edu
Craig Todd	ctodd@desire.wright.edu
Matt Tuttle	snodgras@cren.net
Hung Vu	hungv@fonorola.com
Chris Weider	clw@merit.edu
Les Wibberley	lhw24@cas.org

2.8.9 Uniform Resource Identifiers (uri)

Charter

Chair(s):

Jim Fullton, Jim.Fullton@cnidr.org
Alan Emtage, bajan@bunyip.com

Mailing Lists:

General Discussion: uri@bunyip.com
To Subscribe: uri-request@bunyip.com
Archive: archives.cc.mcgill.ca:~/pub/uri-archive

Description of Working Group:

The Uniform Resource Identifiers Working Group is chartered to define a set of standards for the encoding of system independent Resource Location and Identification information for the use of Internet information services.

This Working Group is expected to produce a set of documents that will specify standardized representations of Uniform Resource Locators (URLs) which specify a standardized method for encoding location and access information across multiple information systems. Such standards are expected to build upon the document discussed at the UDI BOF session held during the 24th IETF meeting in Boston, Unique Resource Serial Numbers (URSNs) which specify a standardized method for encoding unique resource identification information for Internet resources and Uniform Resource Identifiers (URIs), which specify a standardized method for encoding combined resource identification and location information systems to be used for resource discovery and access systems in an Internet environment.

Such a set of standards will provide a framework that: allows the Internet user to specify the location and access information for files and other resources on the Internet, allows users and network-based tools to uniquely identify specific resources on the Internet, and allows the creation and operation of resource discovery and access systems for the Internet. The security of such resource discovery services will also be considered to be an integral part of the work of this Group.

Goals and Milestones:

Done — Review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended documents. Review the first draft of a proposal for Uniform Resource Locators already available.

Done — Submit URL document as an Internet-Draft. Review additional draft documents and determine necessary revisions. Follow up discussion will occur on mailing list.

Nov 1993 Submit the URL document to the IESG for publication as a Proposed Standard RFC.

Internet-Drafts:

"Uniform Resource Locators", 04/26/1993, T. Berners-Lee <draft-ietf-uri-url-00.txt, .ps>

"Uniform Resource Names", 05/17/1993, C. Weider, P. Deutsch <draft-ietf-uri-resource-names-00.txt>

CURRENT MEETING REPORT

Reported by Alan Emtage/BUNYIP

Minutes of the Uniform Resource Identifiers Working Group (URI)

The Uniform Resource Identifiers Working Group held three sessions in Columbus. These Minutes are separated on a per-session basis.

Agenda

- Uniform Resource Locators (URLs) - Session 1
- Uniform Resource Names (URNs) - Session 2
- Discussion of Other Necessary Objects - Session 3

Uniform Resource Locators (URLS) - Session 1

In order to try to prevent further confusion it was agreed that the following terminology would be used:

For all the various UR* objects being discuss, the "U" would stand for "Uniform" and the "R" would stand for "Resource".

URL	Uniform Resource Locator
URN	Uniform Resource Name
URI	Unform Resource Identifier (the collective name for UR*)

A discussion of Tim Berners-Lee's current draft on Uniform Resource Locators (URLs) followed. The following points were made:

- It is expected that in the near future character sets other than ASCII will need to be addressed. However in the short run it was decided that ASCII would be adequate for the task. The point was made by several European members of the Working Group that while other character sets would be necessary, it is important to get the current draft out and implement the protocols discussed. Wording to the effect that this matter has been addressed should be incorporated into the current text. The mechanisms defined need to be extensible to allow for expansion in this area.

- The issue of "fragments" was raised. While the current draft addresses "large scale" objects such as entire files and services, it makes no attempt at defining sub-objects (such as a paragraph, word or individual letter in text file). For example, how does one define a "paragraph" in a PostScript file, given that this is effectively an interpreted language? The general consensus was that we still do not have an adequate understanding of the underlying principles involved and that this discussion should be pursued on the mailing list.

- The issue of OSI distinguished names in URLs was discussed. While further discussion is probably warranted, consensus held that this would probably be too "heavy" for the current proposals.

- MIME encodings are also one possible avenue for describing network objects. It was agreed that the Working Group should work closely within the framework of existing RFCs for such descriptions.

- It was agreed that the current URL draft should include an example URL specification for each access method defined in order to guide implementors.

- Again the issue of partial URLs in the current draft was raised. It was agreed that while systems may choose to use such constructs internally, at no time would they be valid at the inter-system interface. Consensus was reached that stronger warnings need be placed in the current draft to that effect. It was also agreed that in the interests of time further discussions of the issue should be taken to the mailing list. The definition of partial URLs should also be moved to an appendix of the current document since they are not part of the official specification. Any algorithm for determining partial URLs should also be moved to the appendix.

- The mechanism for registering new access methods with the Internet Assigned Numbers Authority (IANA) should be more prominently placed in the current draft. Also, some mechanism for defining experimental access methods should be included.

- Several issues were raised which, it was decided, were better suited for the upcoming URN discussion since they fell into that domain.

  - Some form of integrity test was suggested (check digits) for URLs. It was decided that since URLs are inherently transitory in nature, that such tests would not be necessary.
  - Versioning
  - Security issues
  - Time to Live (TTLs)

- It is expected that with this input the current draft can be submitted for Internet-Draft status within a few weeks following the meeting.

Uniform Resource Names (URNs) - Session 2

Before the meeting Cliff Lynch had posted an overview document to the mailing list titled "A Framework for Identifying, Locating, and Describing Networked Information Resources" and at the beginning of the session, he described the major points that the paper contained.

John Kunze had also posted a document entitled "Resource Citations for Electronic Discovery and Retrieval" and made a short presentation about the paper.

Peter Deutsch made a presentation as to a possible architecture for URNs.

A "spirited" discussion followed as a result.

There was much discussion as to what properties a URN should and should not have and the resulting fracas was in the best tradition of IETF "consensus building".

It was agreed that while some of the underlying data of particular network objects changed (for example, a video feed), that the URN associated with such an object would remain essentially the same. However, the URNs for the underlying data would have to change as the data changed.

Several suggestions for the type of information to be included in URNs were discussed and it was decided that a final decision would be made at the final session.

Discussion of Other Necessary Objects - Session 3

After canvassing several members of the Working Group, the session started with a short presentation by one of the co-Chairs, Alan Emtage. It was proposed that the URN have a very simple structure. In order to be able to completely distinguish URNs from URLs the following structure was proposed:

URN:<ID Authority>:<URS>

The string "URN" is part of the structure. <ID Authority> is the unique identifier for the issuing authority. <URS> is the Uniform Resource String which is unique (as determined by the ID Authority) for that ID Authority. No assumptions may be made about the substructure of the URS which is effectively opaque to any entity other than the ID Authority. The ID Authority would be registered with the IANA to ensure uniqueness.

This proposal was endorsed and the corresponding document would be written by Alan Emtage, Jim Fullton and Chris Weider and submitted to the mailing list as soon as possible. It is hoped that the document can become an Internet-Draft at or before the Amsterdam meeting.

A presentation was made by Rob Raisch which suggested some revisions to the above scheme.

Questions about architecture were raised and it was suggested that the current draft architecture document from the Integration of Internet Information Resources (IIIR) Working Group be consulted.

Further discussion illustrated the fact that the combination of URL and URN would not be sufficient for an effective infrastructure since much of the data needed by the user to determine desirability of an object located through a search was not present in these structures. These include such things as:

- Versioning
- Language
- Character Sets
- Representation (e.g., PostScript, bitmaps, ASCII etc.)
- A whole array of non-static/non-text attributes

Tim Berners-Lee, John Kunze and Michael Mealling made presentations as to how to handle this "meta data" or "factoids".

It was decided that defining the semantics and syntax of these attributes would take careful work and should be the focus of upcoming meetings.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jules Aronson	aronson@nlm.nih.gov
Karl Auerbach	karl@empirical.com
Paul Barker	p.barker@cs.ucl.ac.uk
Robert Beer	r-beer@onu.edu
Tim Berners-Lee	timbl@info.cern.ch
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
William Chung	whchung@watson.ibm.com
David Conklin	conklin@jvnc.net
James Conklin	jbc@bitnic.educom.edu
Naomi Courter	naomi@concert.net
John Curran	jcurran@nic.near.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Mark Davis-Craig	mad@merit.edu
Peter Deutsch	peterd@bunyip.com
Thomas DeWitt	tdewitt@osi.ncsl.nist.gov
Dale Dougherty	dale@ora.com
Alan Emtage	bajan@bunyip.com
Urs Eppenberger	eppenberger@switch.ch
Roger Fajman	raf@cu.nih.gov
Jill Foster	Jill.Foster@newcastle.ac.uk
Steven Foster	foster@cs.unr.edu

Jim Fullton	Jim.Fullton@cnidr.org
Kevin Gamiel	kevin.gamiel@concert.net
Joan Gargano	jcgargano@ucdavis.edu
Greg Gicale	gicale@ohio.gov
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Inna Il'yasova	ilyasova@meg.uncg.edu
Erik Jul	jul@oclc.org
Scott Kaplan	scott@wco.ftp.com
Michael Khalandovsky	mlk@ftp.com
John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Andrew Knutsen	andrewk@sco.com
Edward Krol	e-krol@uiuc.edu
John Kunze	jak@violet.berkeley.edu
Ronald Lanning	lanning@netltm.cats.ohiou.edu
Hock-Koon Lim	lim@po.cwru.edu
Clifford Lynch	calur@uccmvsa.ucop.edu
Bruce Mackey	brucem@cinops.xerox.com
Samir Malak	malaks@alice.uncg.edu
Gary Malkin	gmalkin@xylogics.com
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Ignacio Martinez	martinez@rediris.es
Michael Mealling	michael@fantasy.gatech.edu
Keith Moore	moore@cs.utk.edu
Jim Naro	jnaro@nic.near.net
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Bill Norton	wbn@merit.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Geir Pedersen	Geir.Pedersen@usit.uio.no
Pete Percival	percival@indiana.edu
Charles Perkins	perk@watson.ibm.com
Marsha Perrott	mlp+@andrew.cmu.edu
Cecilia Preston	cpreston@info.berkeley.edu
Robert Raisch	raisch@ora.com
Joyce K. Reynolds	jkrey@isi.edu
Francois Robitaille	francois.robitaille@crim.ca
Richard Rodgers	rodgers@nlm.nih.gov
Charlie Smith	crsmith@osvi.edu

Jane Smith	Jane.Smith@cnidr.org
Patricia Smith	psmith@merit.edu
Sue Smith	smiths.es.net
Karen Sollins	sollins@lcs.mit.edu
Wayne Tackabury	wayne@cayman.com
Matt Tuttle	snodgras@cren.net
Ruediger Volk	rv@informatik.uni-dortmund.de
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Les Wibberley	lhw24@cas.org
Richard Wiggins	wiggins@msu.edu

2.8.10 User Documents Revisions (userdoc2)

Charter

Chair(s):

Ellen Hoffman, ellen@merit.edu
Lenore Jackson, jackson@nsipo.arc.nasa.gov

Mailing Lists:

General Discussion: user-doc@merit.edu
To Subscribe: user-doc-request@merit.edu
Archive:

Description of Working Group:

The focus of the USER-DOC2 Working Group is on identifying and locating documentation about the Internet. A major activity is the revision of an existing bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "How to use the Internet" (RFC1175, FYI 3). This effort will also be used to help locate documentation produced by other organizations and examine the means by which such documents are made available on the Internet. The target audience is those individuals who provide services to end users and end users themselves. The Group is also developing a new FYI RFC document designed as a very short bibliography targeted at novice users.

The USER-DOC2 Working Group will:

(1) Identify and categorize useful documents, reference materials, training tools, and other publications about the Internet, particularly those available on-line.

(2) Publish on-line and hard copies of the bibliography(s) produced and other reference material on documentation as needs are identified.

(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 and identify additional needs which are required for locating documentation and other publications.

(5) Review procedures for periodic review of the bibliography by the User Services Working Group.

(6) Examine methods for delivering documentation and work with providers to improve the availability of basic Internet documentation.

Goals and Milestones:

Done	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 IETF.

Done Publish an Internet-Draft of the short bibliography for novice users.

Done Submit the revised FYI document to the IESG for publication as an RFC.

Done Post a revised version of FYI3, "A bibliography of Internetworking Information" as an Internet-Draft.

Apr 1993 Submit the revised FYI3 to the IESG for publication as an Informational RFC.

Request For Comments:

RFC 1463 "FYI on Introducing the Internet–A Short Bibliography of Introductory Internetworking Readings for the Network Novice"

CURRENT MEETING REPORT

Reported by Ellen Hoffman/Merit

Minutes of the User Documents Revisions Working Group (USERDOC2)

Introduction

The Group's Charter was described and mailing addresses were reviewed. Subscriptions to the mailing list should now be sent to user-doc-request@merit.edu.

Bibliography for New Users Completed

The final version of the short bibliography was distributed. This has been submitted to the Area Director as the first step in finalizing it as an FYI RFC.

The on-line documents listed in the soon-to-be FYI are all collected in an archive available from four servers on the Internet for anonymous FTP called "Introducing the Internet." Two sites have email servers (nic.merit.edu and nisc.sri.com). Gopher and WAIS servers are available at Merit. Four additional sites expressed interest in mirroring the archive. The documents can also be reached via dialup services at Merit for those with a modem and communications software. For more information, send an email message to nis-info@merit.edu with the text: send access.guide.

To provide an ongoing source of information about new documents for the FYI, a file will be kept in the "Introducing the Internet" archive listing new materials aimed at beginning Internet users. The Working Group recommends updating the FYI itself in a year.

First Draft of Comprehensive Bibliography Distributed

A nine-page draft bibliography was distributed as a first step in updating RFC1175. The Group discussed goals and directions for the document. The Group agreed that the new bibliography would supplement, rather than obsolete, the existing RFC. It will cover hardcopy and on-line documents about the Internet from 1990-93. With nine pages in a very incomplete listing, it was also agreed that abstracts for all listings would make the bibliography too long, so some form of keywords would be developed to help provide content information.

The Group agreed to work on developing the list on-line through the mailing list, and to try to have a final version for approval at the July meeting. A suggestion was made to develop an on-line version that would allow navigation to each of the electronic documents, and it was agreed this would be discussed at the next meeting. There was also interest in having a document that could have "icons" that would make it easy to identify different subject areas. Creating a tool useful to librarians was another area of interest. A previous bibliography done by OCLC was presented as a possible source tool.

Input was solicited on corrections and additions to the document. A revised version will be submitted as an Internet-Draft and Working Group members were asked to send suggestions to the mailing list or authors.

Revised Goals and Directions

A brief discussion was held of new and future projects for the Working Group. A possible FYI aimed at those who are not yet connected, which had been mentioned in the earlier User Services Working Group meeting, was suggested as a project. With the Group having taken on introductory materials as part of the short bibliography, it was agreed that this was a reasonable project. A couple of authors said they had already written materials on this and offered to send them to the mailing list for consideration. There was continued interest in creating a "living document" to keep the bibliographies updated, and a locator resource for documentation on the Internet. More work will be done on these topics on the mailing list.

Attendees

Robert Beer	r-beer@onu.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Erik Jul	jul@oclc.org
Edward Krol	e-krol@uiuc.edu
Hock-Koon Lim	lim@po.cwru.edu
Gary Malkin	gmalkin@xylogics.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Bruce Nelson	bnelson@novell.com
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Jennifer Sellers	sellers@nsipo.nasa.gov
Charlie Smith	crsmith@osvi.edu
Patricia Smith	psmith@merit.edu
Craig Todd	ctodd@desire.wright.edu
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.8.11 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: nnsc.nsf.net:~/nsfnet/us-wg*

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"

RFC 1462 "FYI on "What is the Internet?""

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

Minutes of the User Services Working Group (USWG)

Joyce Reynolds reported on the IETF User Services Area activities including: working groups coming to closure and new working groups starting up, new publications, and current user services related Internet-Drafts postings.

Working Groups coming to closure include:

- Directory Information Services Infrastructure (DISI).
- NOC-Tool Catalogue Revisions (NOCTool2).
- Internet User Glossary (USERGLOS).

New working groups initiating their first sessions at this IETF include:

- Training Materials (TRAINMAT).
- Integrated Directory Services (IDS).

New FYI RFC publications since the last IETF are:

- FYI 18 "Internet Users' Glossary", (Also RFC 1392), January 1993.
- FYI 17 "The Tao of IETF - A Guide for New Attendees of the Internet Engineering Task Force", (Also RFC 1391), January 1993.

Discussion also focused on the "new NICs" (i.e., the InterNIC and the proposed/experimental APNIC - Asia/Pacific NIC).

Jill Foster presented an update on RARE activities, including a report on the RARE Information Services/User Support (ISUS) activities. The ISUS Working Group is broken down into several sub-areas: Network User Support, Asynchronous Group Communication, and Networked Information Retrieval and Services Liaison. WG-ISUS Task Forces are being set up to cover the main tasks outlined in the workplan. The following task forces are proposed (and many have started work):

- User Documentation.
- Document Delivery.
- Networked Information Retrieval (joint RARE/IETF/CNI endeavor).
- Coordination of Networked Information Retrieval Services.
- Maintenance and gathering of information on networks and networked resources.
- UNITE "Total Solution" User Interface.
- RARE Technical Report 1.
- Training, Publicity and Awareness (joint RARE/IETF Working Group).
- Support for Special Interest Communities.

- Automatic Mailing List Servers.
- Multi Media Information Services.

Susan Calcari presented a talk on the new InterNIC. The InterNIC information services include: Network Solutions provides the registration services, AT&T provides the Directory and Database Services, and General Atomics provides the information services. NIS fests and NIS manifests are planned. The InterNIC reference desk will be open from 5am to 7pm Pacific time. For more information on information services send email to: info@internic.net, Database and Directory Services: admin@ds.internic.net, and for Internet Registry services, hostmaster@internic.net.

Gary Malkin led a session on two FYI RFC updates:

- FYI 7 "FYI on Questions and Answers: Answers to Commonly Asked "Experienced Internet User" Questions", (Also RFC1207), February 1991.

- FYI 4 "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions", (Also RFC1325), May 1992.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Robert Beer	r-beer@onu.edu
Susan Calcari	calcaris@cerf.net
Corinne Carroll	ccarroll@bbn.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Al Costanzo	al@akc.com
Naomi Courter	naomi@concert.net
Mark Davis-Craig	mad@merit.edu
Peter Deutsch	peterd@bunyip.com
Robert Enger	enger@reston.ans.net
Jill Foster	Jill.Foster@newcastle.ac.uk
Marcello Frutig	frutig@rnp.impa.br
Gwen Funchess	funchess@magnus.acs.ohio-state.edu
Maria Gallagher	maria@nsipo.nasa.gov
Greg Gicale	gicale@ohio.gov
Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Deborah Hamilton	debbie@qsun.att.com
Susan Harris	srh@umich.edu
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen@merit.edu
Alton Hoover	hoover@ans.net
Susan Horvath	shorvath@merit.edu
Edward Krol	e-krol@uiuc.edu

Hock-Koon Lim	lim@po.cwru.edu
Gary Malkin	gmalkin@xylogics.com
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Cynthia Mills	cmills@bbn.com
Jim Naro	jnaro@nic.near.net
Bill Norton	wbn@merit.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Richard Rodgers	rodgers@nlm.nih.gov
Jennifer Sellers	sellers@nsipo.nasa.gov
Karen Shay	kshay@cs.cmu.edu
Charlie Smith	crsmith@osvi.edu
Patricia Smith	psmith@merit.edu
Ruediger Volk	rv@informatik.uni-dortmund.de
Janet Vratny	janet@apple.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.8.12 Whois and Network Information Lookup Service (wnils)

Charter

Chair(s):

Joan Gargano, jcgargano@ucdavis.edu

Mailing Lists:

General Discussion: ietf-wnils@ucdavis.edu
To Subscribe: ietf-wnils-request@ucdavis.edu
Archive: ucdavis.edu:~/archive/wnils

Description of Working Group:

The Network Information Center (NIC) maintains the central NICNAME database and server, defined in RFC954, providing online look-up of individuals, network organizations, key nodes, and other information of interest to those who use the Internet. Other distributed directory information servers and information retrieval tools have been developed and it is anticipated more will be created. Many sites now maintain local directory servers with information about individuals, departments and services at that specific site. Typically these directory servers are network accessible. Because these servers are local, there are now wide variations in the type of data stored, access methods, search schemes, and user interfaces. The purpose of the Whois and Network Information Lookup Service (WNILS) Working Group is to expand and define the standard for WHOIS services, to resolve issues associated with the variations in access and to promote a consistent and predictable service across the network.

Goals and Milestones:

Done	Review and approve the Charter making any changes deemed necessary. Examine the particular functional needs for expanded whois directory service. Begin work on a framework for recommendations. Assign writing assignments for first draft of document.
Apr 1993	Post the Whois and Network Information Lookup Service Recommendations document as an Internet-Draft.
Apr 1993	Post the revised WHOIS protocol and index service document to the IESG as an Internet-Draft.
Done	Post the "Architecture of the Whois++ Index Service" as an Internet-Draft.
Jun 1993	Submit the Whois and Network Information Lookup Service Recommendations document to the IESG as an Informational RFC.
Jun 1993	Submit the "Architecture of the WHOIS++ Index Service" to the IESG for consideration as an Informational RFC.

Sep 1993 Submit a revised WHOIS protocol specification and index service document to the IESG for consideration as a Draft Standard.

Internet-Drafts:

"Architecture of the Whois++ Index Service", 11/23/1992, C. Weider, J. Fullton, S. Spero <draft-ietf-wnils-whois-01.txt>

CURRENT MEETING REPORT

Reported by Joan Gargano/UCDavis

Minutes of the Whois and Network Information Lookup Service Working Group (WNILS)

Joan Gargano reviewed the Agenda and the ways in which WNILS information is available over the network. Each author of the papers was asked to provide a review of work completed to date.

Review of Whois++ Architecture - Peter Deutsch

Peter provided a brief overview of the Whois++ model and architecture for new members of the Working Group. He then discussed the following extensions to the document which incorporated suggestions from the November 1992 IETF meeting and other discussions. The Working Group was asked for consensus on the inclusion of these features into the architecture.

1. Options and Extensions

   To maintain the simplicity of the original model and provide the features that may be needed in the future, a facility for options and extensions has been described and is recommended for inclusion in the architecture.

2. Multilingual Support

   Multilingual support is very important to the European community and needs to be supported. Whois++ should support the specification of multilingual queries and responses and the handling of multiple character sets. This option will be included and should draw upon the MIME standard.

3. Security and Authentication

   There is a need to restrict the output of information based upon the identity of the requester. An option should be implemented which provides for authentication.

4. Multiple Views

   The Group suggested this feature was already implemented in other query languages such as SQL and Z39.50. It was recommended that this feature be discussed further.

5. MIME Support

   The ability to support the delivery of non-ASCII text has been defined in the MIME standard. It is recommended that Whois++ incorporate the appropriate portion of the MIME standard to handle these documents.

6. Error Message Specification

   Error handling and error messages need to be defined.

7. Extending Centroids

   This will be covered by Chris Weider.

8. Changing Ports and Developing a More Complex Protocol

   This Group felt this should not be pursued.

9. Ability to Hold Connections Open

   Multiple queries would benefit from the ability to leave a connection open. This should be included as an option.

10. Template Clearinghouse

    CNIDR has volunteered to perform this function.

11. Synonyms

    This feature would allow the ability to perform implicit "or" queries. It was recommended that this feature not be included. This can be accomplished by the client issuing multiple queries.

12. Database Write Option

    This feature would provide a distributed data management feature and should be included as an option.

Review of the Distributed Whois++ model - Centroids - Chris Weider

Chris gave a brief overview of the Centroid model of data propagation and information retrieval. The following extensions are recommended for the centroid model.

1. A free text abstract of the contents of each centroid server would provide a mechanism for finding servers which contain topical information of interest. This would provide an efficient way of providing a "Yellow Pages" directory server.

2. Polled by and polled for features would provide a mechanism for determining server location.

3. Bookmarks allow users to return to servers they have found useful in the past.

Review of Front End to Database Integration - Jim Fullton

A Whois++ client can be implemented with simple shell scripts. Centroid implementation on a UNIX platform currently require gdbs. Clients have been implemented on RS6000, Sun and Linex platforms.

Review of Recommended Modifications to the Whois Protocol - Joan Gargano

Joan discussed the collation of data elements underway by the IAFA, NIR, TopNode and WNILS members. A revised list of data elements will be incorporated into the document prior to submission as an Internet-Draft.

Discussion of Projects

A simple server, centroid and client will be available by April 30th. Sources will be available on ftp.cnidr.org.

Revise Goals and Milestones

April 30, 1993	Submit the "Whois and Network Information Lookup Service Recommendations" and "Architecture of the Whois++ Index Service" documents as Internet-Drafts.
April 30, 1993	A working implementation of the Whois++ protocol will be available for demonstration. Sources will be available from ftp.cnidr.org.
June 11, 1993	Submit the "Whois and Network Information Lookup Service Recommendations", "Architecture of the Whois++ Index Service" and "Distributed Whois++ Model - Centroids" documents as Proposed Standards.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
N. Akiko Aizawa	akiko@nacsis.ac.jp
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Robert Beer	r-beer@onu.edu
Richard Bjers	rich.bjers@uc.edu

Thomas Brisco	brisco@pilot.njin.net
Sandy Bryant	slb@virginia.edu
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
David Conklin	conklin@jvnc.net
Naomi Courter	naomi@concert.net
Brent Curtiss	bcurtiss@magnuss.ocs.ohio-state.edu
Mark Davis-Craig	mad@merit.edu
Alan Emtage	bajan@bunyip.com
Roger Fajman	raf@cu.nih.gov
Jill Foster	Jill.Foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
Marcello Frutig	frutig@rnp.impa.br
Jim Fullton	Jim.Fullton@cnidr.org
Maria Gallagher	maria@nsipo.nasa.gov
Kevin Gamiel	kevin.gamiel@concert.net
Joan Gargano	jcgargano@ucdavis.edu
Greg Gicale	gicale@ohio.gov
Terry Gray	gray@cac.washington.edu
Susan Harris	srh@umich.edu
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Ellen Hoffman	ellen@merit.edu
Susan Horvath	shorvath@merit.edu
Inna Il'yasova	ilyasova@meg.uncg.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Neil Katin	katin@eng.sun.com
John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Edward Krol	e-krol@uiuc.edu
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
David Lovering	lovering@bldrdoc.gov
Bruce Mackey	brucem@cinops.xerox.com
Samir Malak	malaks@alice.uncg.edu
Kent Malave	kent@bach.austin.ibm.com
Janet L. Marcisak	jlm@ftp.com
April Marine	april@atlas.arc.nasa.gov
Jerry Martin	jerry+@osu.edu
Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Keith Moore	moore@cs.utk.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Pete Percival	percival@indiana.edu
David Perkins	dperkins@synoptics.com
Marsha Perrott	mlp+@andrew.cmu.edu

Mark Prior	mrp@itd.adelaide.edu.au
Robert Raisch	raisch@ora.com
Joyce K. Reynolds	jkrey@isi.edu
Francois Robitaille	francois.robitaille@crim.ca
Carl Schoeneberger	70410.3563@Compuserve.com
Jennifer Sellers	sellers@nsipo.nasa.gov
Jane Smith	Jane.Smith@cnidr.org
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgrass@bitnic.educom.edu
Craig Todd	ctodd@desire.wright.edu
Matt Tuttle	snodgras@cren.net
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
Ruediger Volk	rv@informatik.uni-dortmund.de
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Les Wibberley	lhw24@cas.org

Chapter 3

Network Status Briefings

3.1 ANS

Reported by Jordan Becker/ANS

ANSnet Network Status Report
April 1, 1993

AIX 3.2 Installation on RS6000 Routers

T3 Network Status

ANSnet Usage Reports & Performance Analysis

T3 Router Software Migration Plan

Summer '93 Planned Network Enhancements

ans

AIX 3.2 Release 1 Deployment

New BSD 4.4 Based Kernel
- Radix Trie Forwarding Table
- Use of Route Sockets

Already Deployed on 4 Nodes
- CNSS51, ENSS205, ENSS160, ENSS131

Full Scale Production Network Deployment Begins on 4/2
- Separated into 4 Weekend Maintenance Windows
- Involves Hard Drive Swaps for CNSS & Adjacent ENSS Nodes
- 4/2 : Washington D.C.
- 4/16 : Seattle/Denver, San Francisco/Los Angeles
- 4/23 : Greensboro/Atlanta, Houston/St. Louis
- 4/30 : Hartford/New York City, Cleveland/Chicago

Fall Back Procedures Established If Necessary

ans

AIX 3.2 DEPLOYMENT PLAN

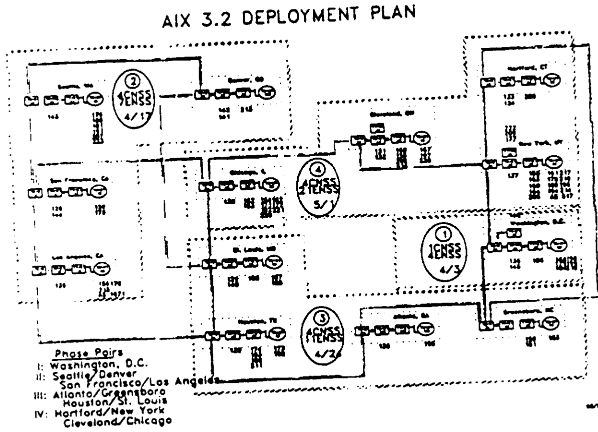

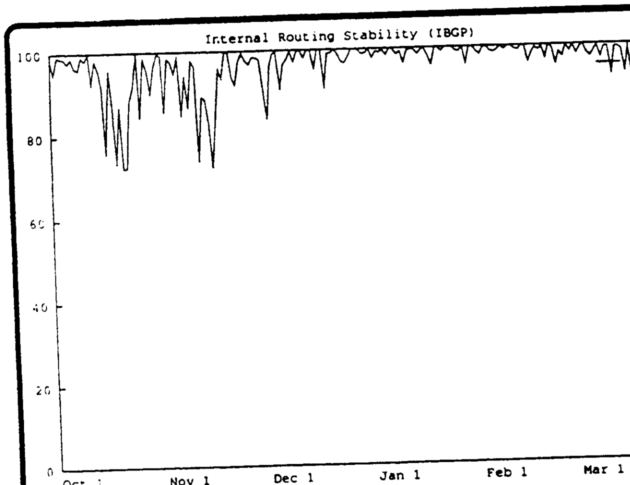

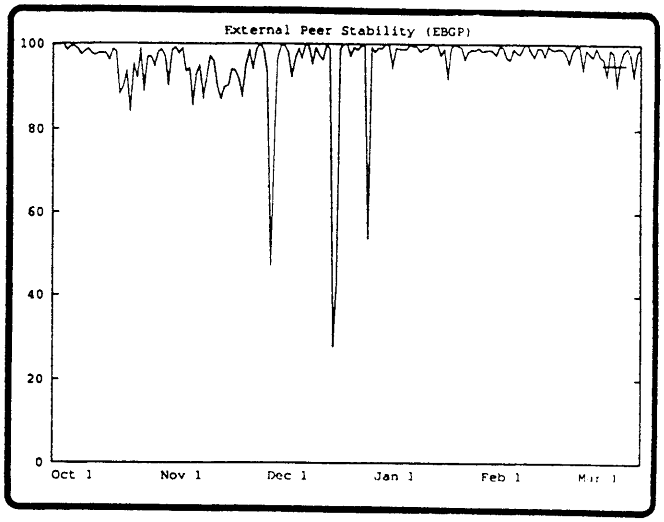

T3 Network Status

- RS960 FDDI now Installed at 14 ENSS Sites, Including...
  - Boston (E134), Argonne (E130), FIX-E (E145)
  - Houston (E139), FIX-W (E144), Salt Lake City (E142)
  - Seattle (E143), Ann Arbor (E131), Denver (E141)
- **T1 NSFNET Backbone Dismantled in February '92**
- **EON OSI (RFC1070) Encapsulator Support**
  - **Installed EON Router at OARnet for TUBA Demo**
- **PSC-SDSC T3 Network Performance Testing with UDP-Ping**
  - **17.8Mbps (550PPS) to 20.8Mbps (650 PPS)**
- **ANSnet Router Traffic Statistics Reports**

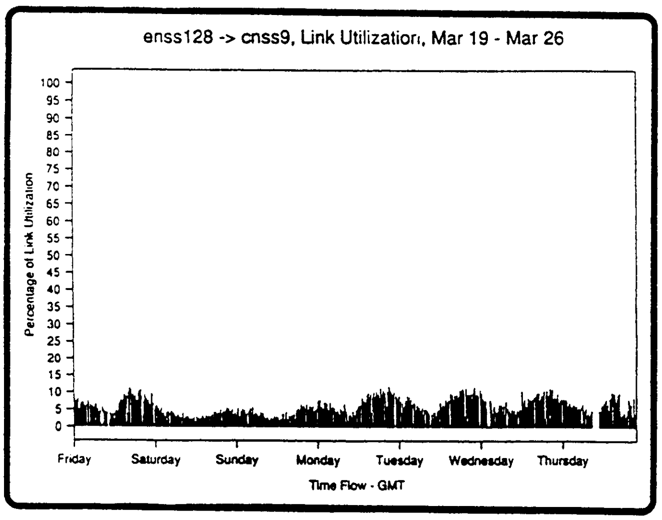

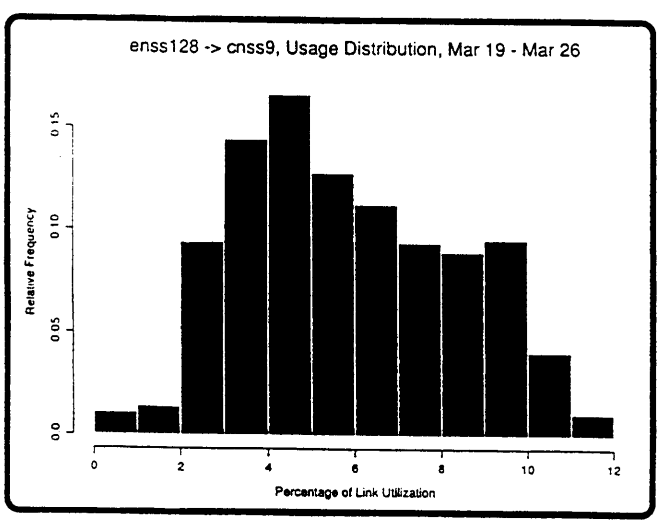

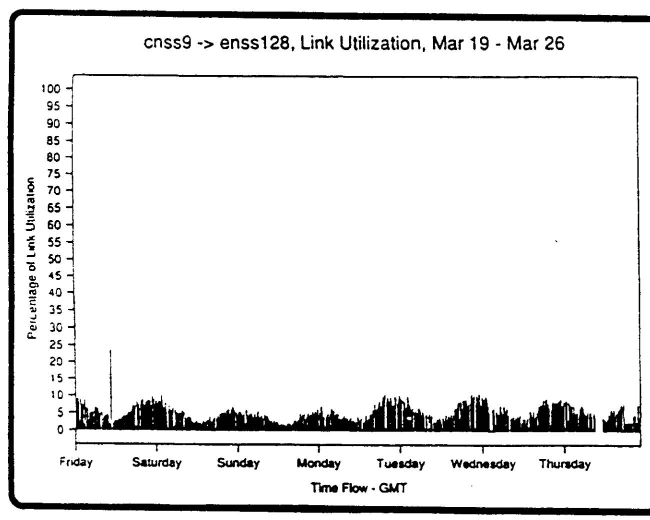

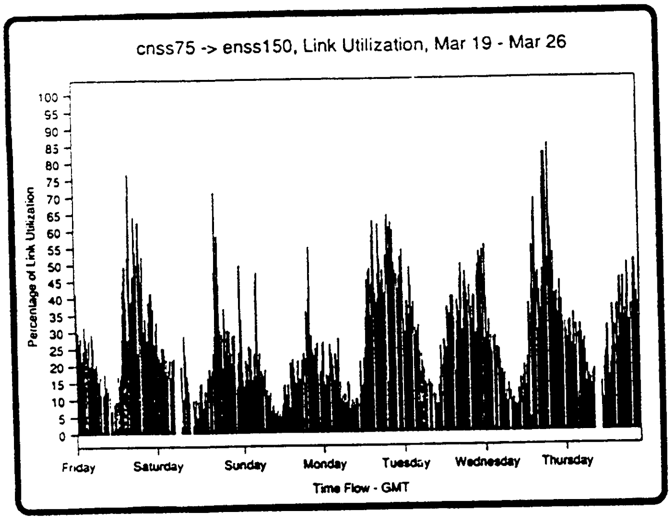
cnss75 -> enss150, Link Utilization, Mar 19 - Mar 26
Percentage of Link Utilization
Friday
Saturday
Sunday
Monday
Tuesday
Wednesday
Thursday
Time Flow - GMT

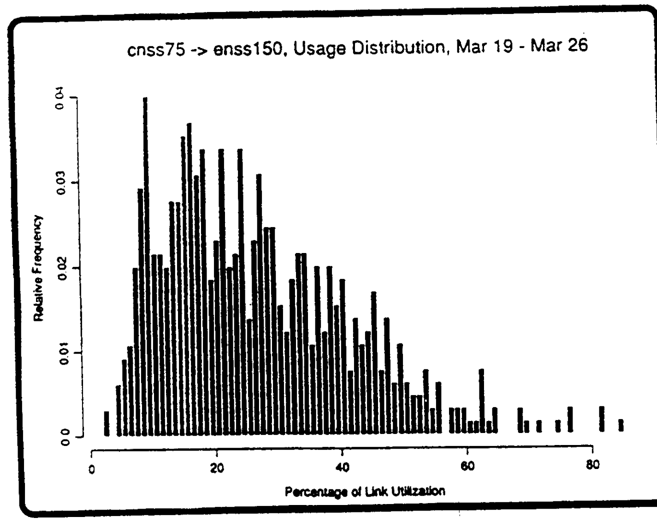
cnss75 -> enss150, Usage Distribution, Mar 19 - Mar 26
Relative Frequency
Percentage of Link Utilization

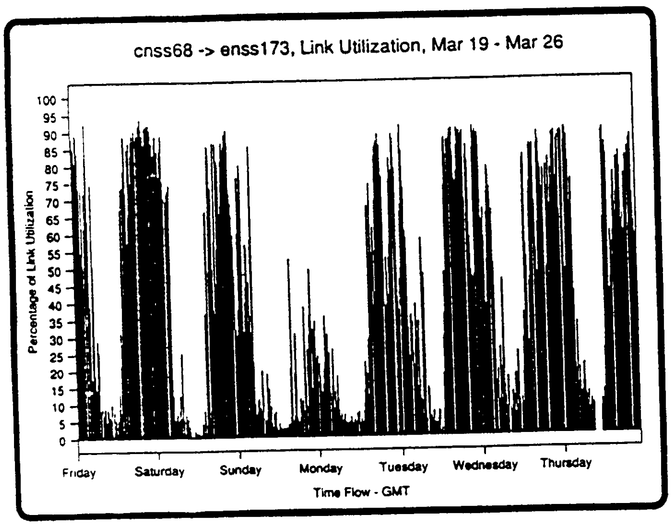
cnss68 -> enss173, Link Utilization, Mar 19 - Mar 26
Percentage of Link Utilization
Friday
Saturday
Sunday
Monday
Tuesday
Wednesday
Thursday
Time Flow - GMT

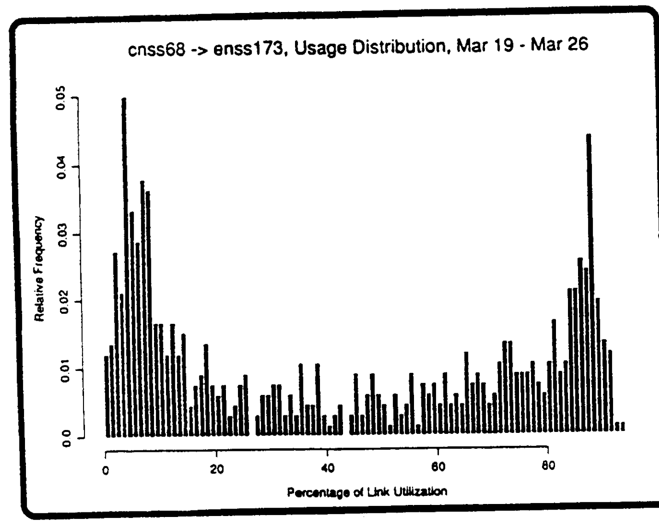
cnss68 -> enss173, Usage Distribution, Mar 19 - Mar 26
Relative Frequency
Percentage of Link Utilization

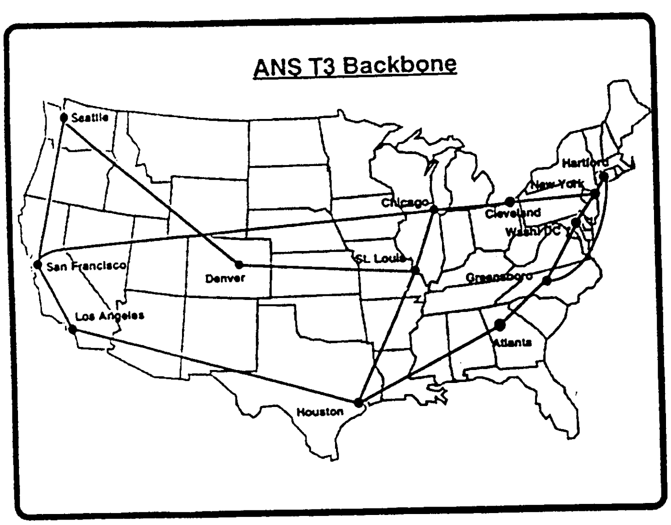
ANS T3 Backbone
Seattle
San Francisco
Los Angeles
Denver
St. Louis
Chicago
Cleveland
Hartford
New York
Wash DC
Greensboro
Atlanta
Houston

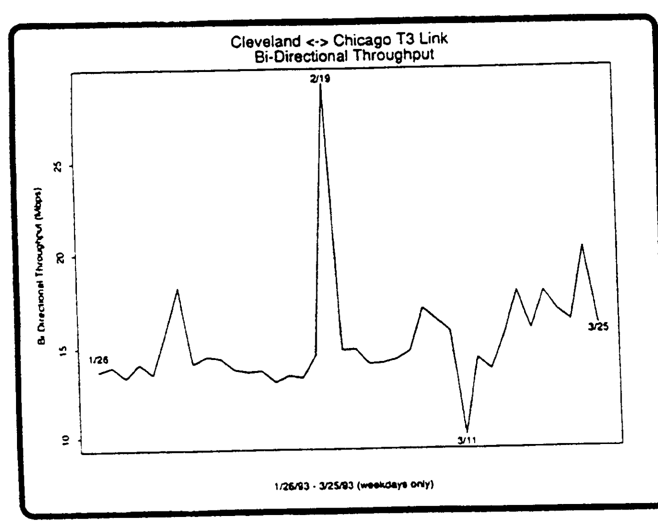
Cleveland <-> Chicago T3 Link
Bi-Directional Throughput
Bi-Directional Throughput (Mbps)
2/19
1/26
3/11
3/25
1/26/93 - 3/25/93 (weekdays only)

T3 Router Software Migration Plan

April '93

- AIX 3.2 Rel. 1
  - 12,000 Destination Networks Supported
  - Radix Trie Forwarding Table, Route Sockets
- BGP4 CIDR Open Interoperability Testing Invitation
  - Vendors Welcome to Tunnel BGP4 Through ANSnet & Testnet
  - May Bring Equipment to Testnet

May '93

- AIX 3.2 Rel. 2
  - 25,000 Destination Networks Supported
  - Limited Support for CIDR/Supernetting
- GateD Rel. 1 Deployment
  - SLSP IGP Co-Existing with Rcp_routed
  - New BGP4, BGP3, BGP2, I-BGP

ans

T3 Router Software Migration Plan

June '93

- AIX 3.2 Rel. 3
  - Full On-Card CIDR Support for Class-less Address/Masks
  - Native OSI Support (ES-IS, CLNP)
  - PPP on T960/RS960
- GateD Rel. 2 Deployment
  - BGP4 with CIDR Card/Kernel Support
    - Turn on CIDR Aggregation To/From CIDR Peers?
    - Controlled De-Aggregation Supported
    - Non-BGP4 Neighbors will Default
  - Dual IS-IS IGP Routing
    - Injection of Prefix Path Attributes from egps
      - Eliminate I-BGP
    - Multipath Forwarding
    - IS-IS MIB

ans

Summer '93 Planned T3 Network Enhancements

- T3 & FDDI RS960 Interface Upgrades
  - New HSSI Daughter Card with Full OC-1 Bandwidth
  - Base Card Upgrade for Lower Power, Higher Performance
    - 40KPPS per Card @ 250 Byte Packet Size
  - Fixes Memory Parity Problem
- CNSS Router Reconfiguration
  - Common FDDI Ring for CNSS Routers in each POP
  - Replace T3-DSUs with T3Plus Bandwidth Managers
    - Additional Fault Tolerance

ans

Network Status Report Summary

Continue to Focus on Improved System Stability

- Support for 12K -> 25K Destination Networks
- GateD Migration and BGP4/CIDR Support

Performance & Function Enhancements

- Native CLNP Switching Services
- New RS960 HSSI and Base Cards
- CNSS Reconfiguration with FDDI Rings

ans

3.2 JvNCnet

Reported by Vikas Aggarwal/JvNCnet

Global Enterprise Services, Inc.

Vikas Aggarwal
JvNCnet

GES Background

- 1986: Established as part of the John von Neumann National Supercomputer Center
- Built the first T1 (1.544 million bits per second) research network worldwide
- 1990: Moved to Princeton University
- 1992: Spun off of Princeton as a private company

GES Overview

- Headquartered in Princeton, NJ
- Access points in nine states: NJ, PA, NY, CT, RI, MA, MD, DE and Washington DC
- International connections to Japan (64kbps), Singapore (128kbps), Taiwain (256kbps), and Venezuela (19.2kbps)
- 235 Members

GES Organizational Chart

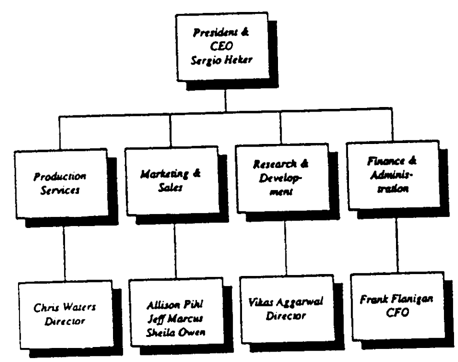

JvNCnet Connectivity

- T1 backbone–multiple links between POPs
- T3 ENSS to the NSFNET
- T1 to ESnet
- CIX connection 2Q 1993

JvNCnet Backbone Network

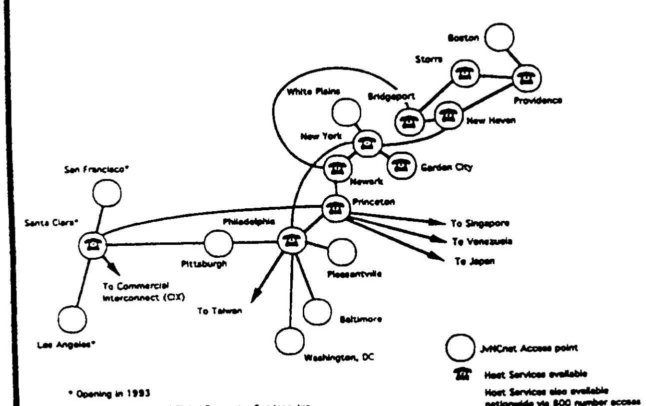

Connectivity Options

- Dedicated leased line connections
  - Bandwidths from 19.2kbps to T1
  - Cisco CSC/3, CSC/4, IGS
- Dialin'Gateway connections
  - Datability Mux/Router
  - Morning Star Technologies PPP
  - Telebit NetBlazer
- Host connections
  - SLIP access
  - V.32bis

Routing

- IGRP on backbone
  - No metric holddown
  - Fast IGRP (timings halved)
- RIP to endsites
  - Usually advertise default only
- BGP3 to ENSS

Operations

- Monitoring
  - NOCOL (freely available)
  - SNMP for statistics
- Trouble Ticketing System
  - Netlog (freely available)
- Monthly reports
  - Sent to members
  - Downtime and traffic stats
  - Line utilization
- 24 hour/7 day NOC operation

Trouble Ticketing- Netlog

- Developed within JvNCnet
- OPEN, UPDATE, CLOSE, INFO ticket types
- Runs on Unix platforms
- Requires no database
- Used to generate monthly summaries
- Available from ftp.jvnc.net : pub/jvncnet-packages

```
                    ********************
                    * NETLOG MAIN MENU *
                    ********************

      * Create an entry
        List open tickets
        Read log
        Search for an entry
        Process ticket(s)
        Edit log
        Index Utilities
        Quit

Enter option:
```

Network Monitoring- NOCOL

- Developed at JvNCnet
- Multiple, independent monitoring agents
- Curses display module
- Monitors for:
  - ICMP, SNMP traps, nameserver, thruput, terminal server tty's
- New release:
  - Logging + notification, novell monitors
- Test- nocol.jvnc.net as user nocol
- Availability- ftp.jvnc.net : pub/jvncnet-packages

```
04                                    JvNCnet NOCOL        Fri Mar 26 19:34:14 1993

           Site          Address       Time   <-Variable->  <-Value->  Condition

           mitl     130.94.2.10        12:45  Reachability      0      Critical
         mitl-c     150.169.1.1        12:45  Reachability      0      Critical
      rmorris-c     192.203.233.10     12:45  Reachability      0      Critical
       st-joe-c     192.245.90.9       19:33  Reachability      6      Warning
  sacredheart-c     148.166.1.11       19:33  Reachability      6      Warning
       JPMorgan     wtmp.jpmorgan      19:34   tslineusage      1      Warning



                                                           Display Level: WARNING
                                  MESSAGES


Enter option, or any other key for next screen:
```

Traffic Statistics

- 15 minute samples using SNMP
- Munged using snmppoll-dc (data cruncher)
- Processed using Splus
- Obtain line utilization, number of packets/bytes
- Helps plan upgrades, redesign backbone,

Additional GES Services

- Network News Feeds (45 sites, INN)
- Domain Name Service (350 domains)
- RFC Repository
- Symposium Series (monthly)
- Quarterly Members' Meetings
- Consulting Services
- GES Newsletter

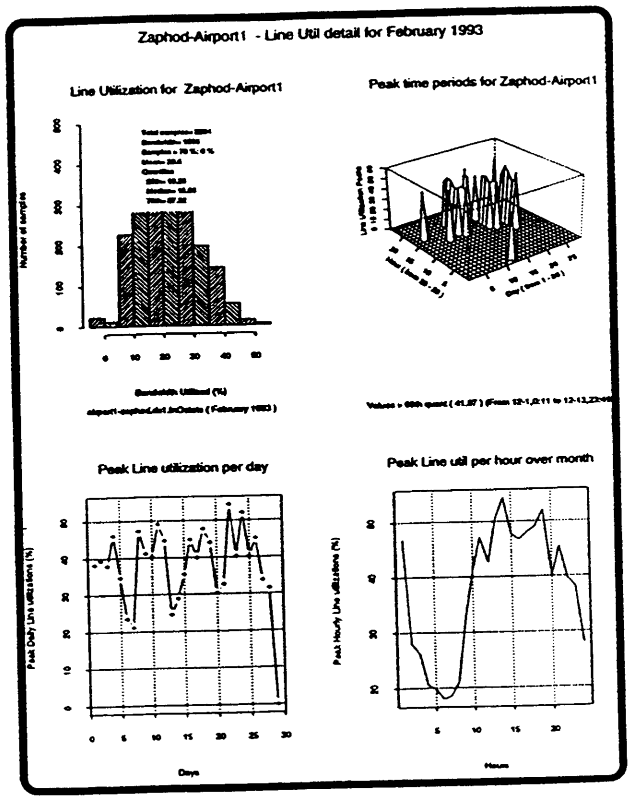

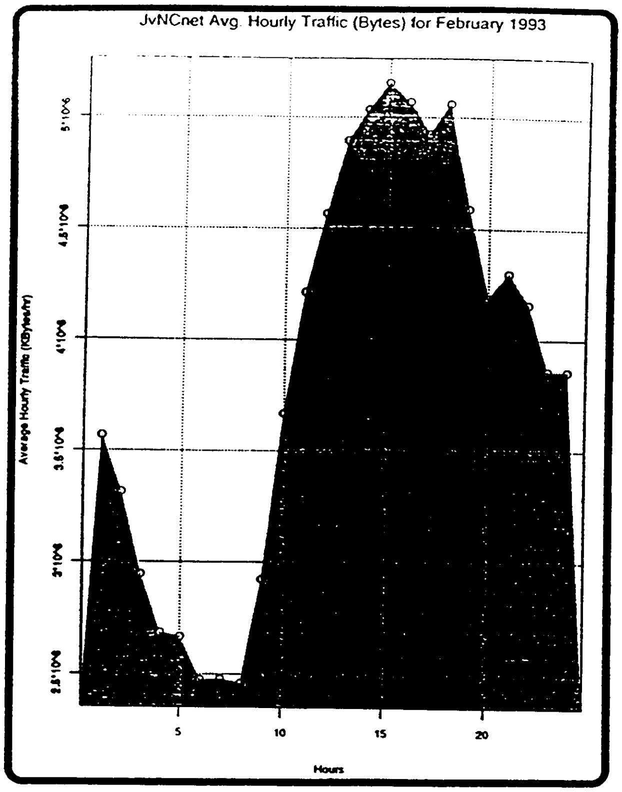

New Projects

- Multicast MBONE node
- SMDS with Bellcore
- FieldWAIS
  - Card Catalog for RFB
  - InterNIC
  - Freely available
- Gopher Services

3.3 NSFNET

Reported by Mark Knopper/Merit

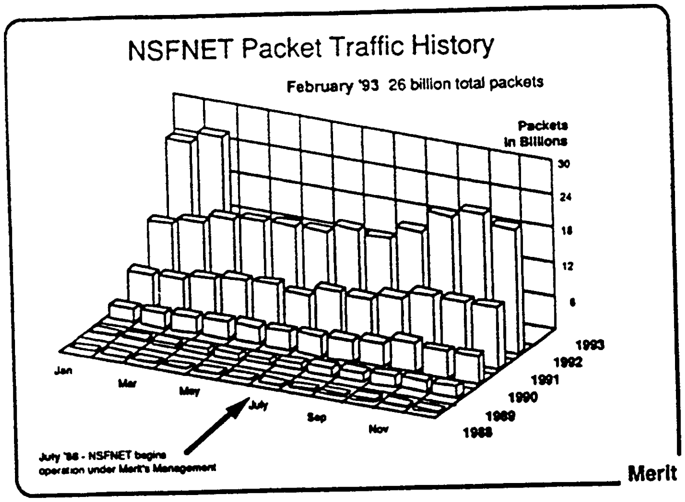
NSFNET Packet Traffic History
February '93 26 billion total packets
Packets in Billions
30
24
18
12
6
1993
1992
1991
1990
1989
1988
Jan
Mar
May
July
Sep
Nov
July '88 - NSFNET begins operation under Merit's Management
Merit

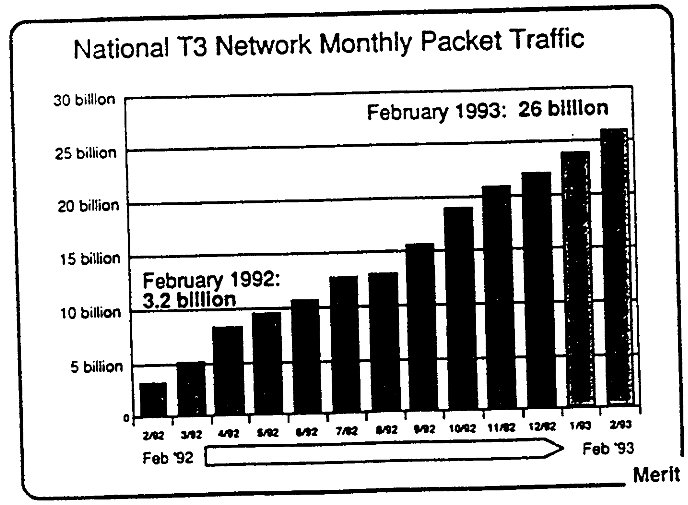
National T3 Network Monthly Packet Traffic
30 billion
25 billion
20 billion
15 billion
10 billion
5 billion
0
February 1993: 26 billion
February 1992: 3.2 billion
2/92
3/92
4/92
5/92
6/92
7/92
8/92
9/92
10/92
11/92
12/92
1/93
2/93
Feb '92
Feb '93
Merit

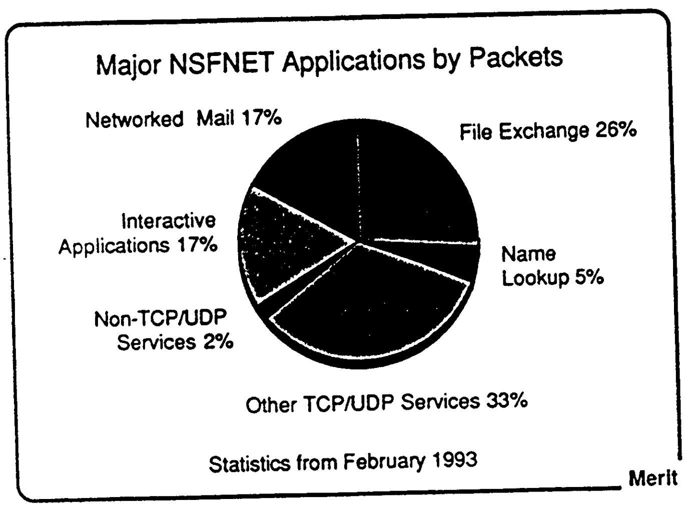
Major NSFNET Applications by Packets
Networked Mail 17%
File Exchange 26%
Interactive Applications 17%
Name Lookup 5%
Non-TCP/UDP Services 2%
Other TCP/UDP Services 33%
Statistics from February 1993
Merit

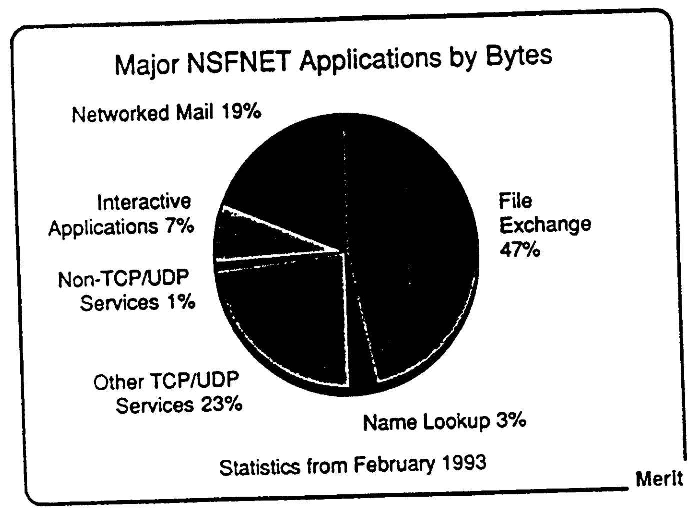
Major NSFNET Applications by Bytes
Networked Mail 19%
Interactive Applications 7%
File Exchange 47%
Non-TCP/UDP Services 1%
Other TCP/UDP Services 23%
Name Lookup 3%
Statistics from February 1993
Merit

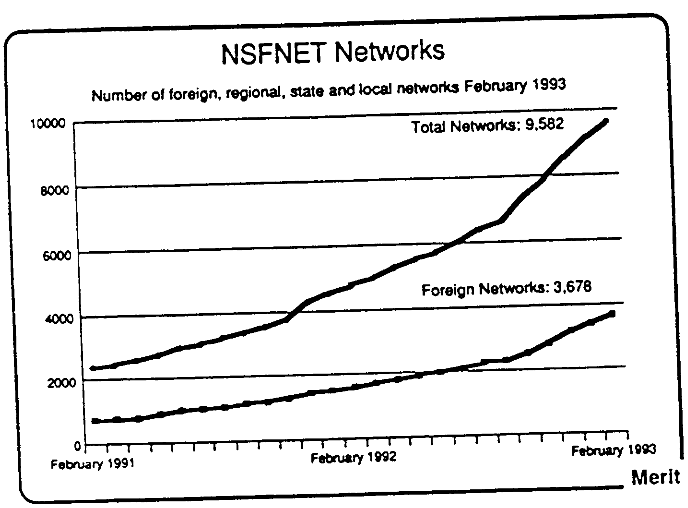
NSFNET Networks
Number of foreign, regional, state and local networks February 1993
10000
8000
6000
4000
2000
0
Total Networks: 9,582
Foreign Networks: 3,678
February 1991
February 1992
February 1993
Merit

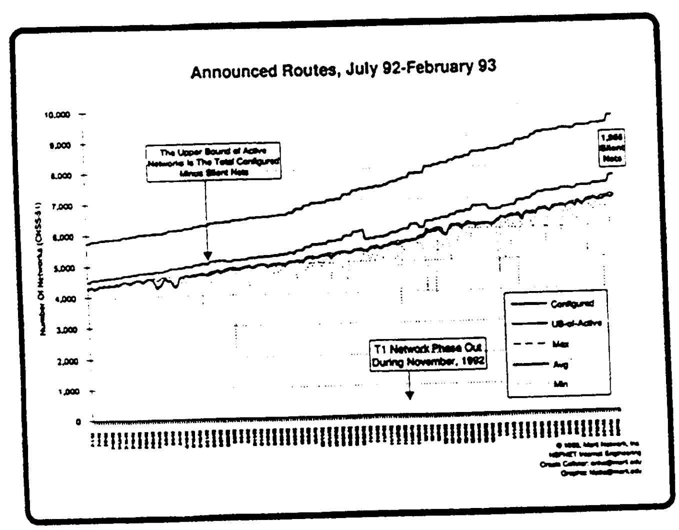
Announced Routes, July 92-February 93
T1 Network Phase Out During November, 1992
Configured
Max
Avg
Min

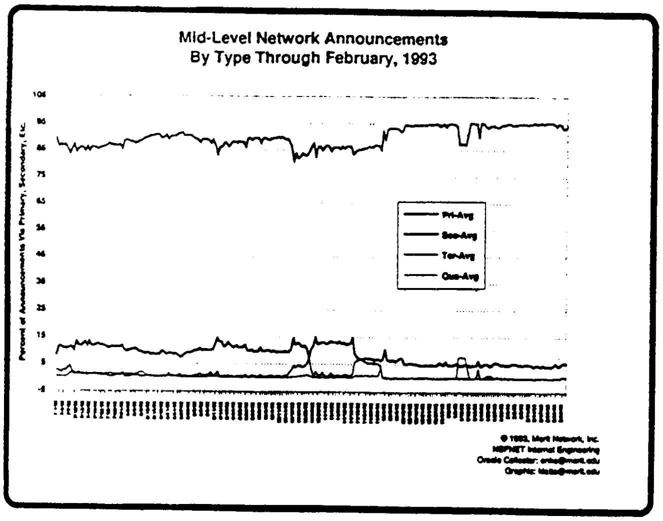
Mid-Level Network Announcements
By Type Through February, 1993

Chapter 4

IETF Protocol Presentations

4.1 Next Generation of IP

Presented by Robert Ullmann/Process Software Corporation

Internet Version 7

Robert Ullmann
Process Software Corporation

26th Internet Engineering Task Force
Columbus, Ohio
March 1993

The Trouble With Tribbles

- Address space
- Route scaling
- TCP window and sequence space
- Interoperation with V4

IPv7 Addresses

- Small number of top level administrative domains
- Ability to subnet every network assignment

Admin Domain	Network	Host

- Mapping of IPv4 numbers
- NSF Administrative Domain

Admin Domain	Network	Host
[fixed at 7E 00 00]	[1st 24 bits of V4 IP]	[1] [last 8]

IPv4 Forwarding Model

- Hop by Hop decision making
- Information being discarded

Hop by Hop Routing

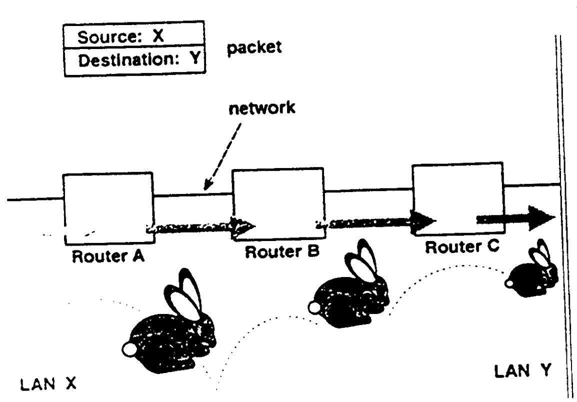

Circuit Routing

- Fast forwarding
- Setup restrictions

Circuit Model

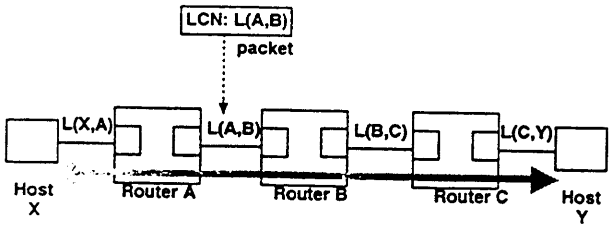

IPv7 Forwarding

- Forward route identifier
- Information from the routing protocol
- Fast switching
- Datagrams can be re-railed

IPv7 Forwarding

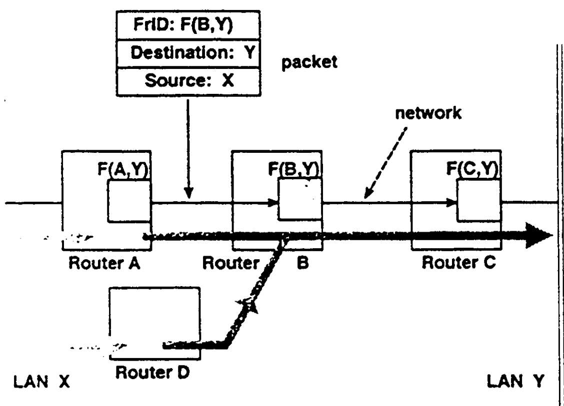

Upstream Decision Making

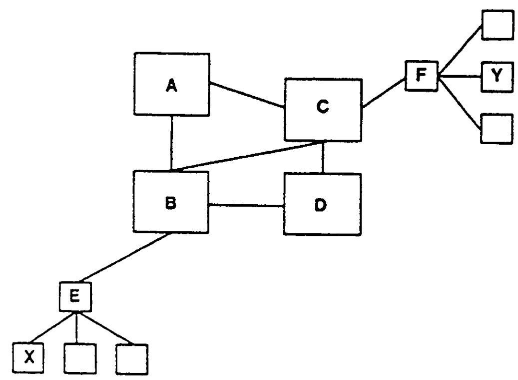

IPv7 Datagram Header

Version	Header Length	Time to Live
Total Datagram Length		
Forward Route Identifier		
Destination Address		
Source Address		
Protocol		Checksum

How Fast Can We Forward?

Steps in IPv7 datagram forwarding:

- Check that version is 7
- Decrement TTL
- Look up forward route identifier and replace
- Adjust checksum

Fields read and written in order

Datagrams may be cut-through routed

Digital Alpha AXP Numbers

- 160 nanoseconds
- 6.24 million times/second

TCP Window and Sequence Space

- TCPv4 limits
- TCPv7 uses 64 bit sequence/acknowledgement
- 32 bit window
- 32 bit port numbers

Interoperation With Version 4

- No dual stacking
- Incremental deployment
- Leaves no legacy systems

Engineering Design, not Research

Protocol status now:

- Design complete
- Start of production prototyping

Status allows vendors to deploy products rapidly:

- RAP in mid 1993
- IPv7 by late 1993 or early 1994

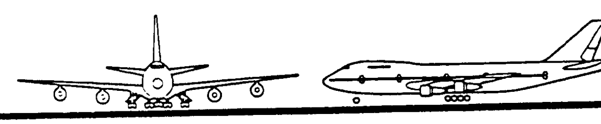

Internet Version 7:

- Addresses the addressing problem
- Extends the transport layer performance envelope
- Advances the network layer technology

"...where they'll be no tribble at all."

4.2 TUBA

Presented by Mark Knopper/Merit

Bio: *Mark Knopper is the manager of Merit's Internet Engineering group which is responsible for NSFNET engineering. Before moving over to the IE group, Mark was the manager of Merit's Michnet engineering group. Mark is also active in the IETF OSI-DS Working Group and was co-investigator on the DARPA/NSF FOX (Field Operational X.500) project.*

The Next Generation of IP

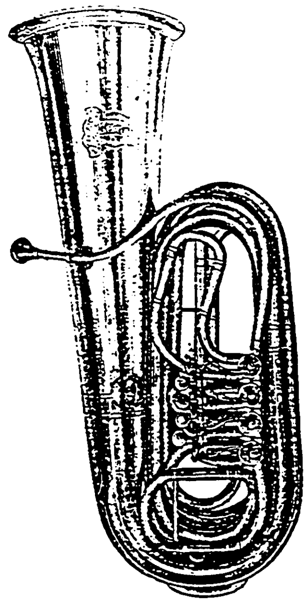

Reproduced by permission of the Trustees of the British Library

Fig. 13. Cerveny Kaiserbass, c.1850.

TUBA:The Next Generation of IP

- Status after last IETF
- Documents
- TUBA implementors agreement
- TUBA demo at this IETF
- Work left to do

TUBA: Status Since Last IETF

- November meeting:
  - Early implementation report (Keith Sklower)
  - Working implementation agreement
  - Transition architecture discussed
  - Document drafts presented
- Mailing list discussions have mainly focused on implementation and interoperability

TUBA: Documents

- RFC 1347:
  TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing, June 1992 (Ross Callon)
- RFC 1348:
  DNS NSAP RRs, July 1992 (Bill Manning)
- draft-ietf-tuba-clnp-02.txt
  Use of ISO CLNP in TUBA Environments, January 1993 (Dave Piscitello)
- draft-ietf-tuba-address-00.ps
  Addressing and End Point Identification For Use with TUBA, October 1992 (Ross Callon)
- Survey of OSI in the Internet
  Work of NOOP group (Sue Hares)
- draft-ietf-noop-echo-01.txt (RFC 1139)
  An Echo Function for ISO 8473, February 1993 (Cathy Wittbrodt, Sue Hares)
- draft-ietf-noop-tools-00.txt
  Essential Tools for the OSI Internet, March 1993 (Sue Hares, Cathy Wittbrodt)

TUBA Implementors Agreement

- TUBA spec "frozen" after November to allow implementation
- Connection identifier and pseudoheader contain full NSAP address
- NSEL = IP protocol type
- Standard CLNP used for network layer

TUBA Implementations

- cisco
- 3com
- BSDI/BSD 4.4
- SunOS (based on BSD 4.4 NET II)
- NCSA Telnet (for MS-DOS)
- RS/6000 AIX 3.2

TUBA Infrastructure

- In the current demo, TUBA hosts are located at Merit, Los Alamos, NIST, 3com, and cisco.
- Connectivity is provided by:
  - NSFNET (using EON encapsulation)
  - ESnet to LANL
  - SURAnet to NIST
  - OSINET X.25 to 3com
  - BARRnet to cisco
  - RARE/COSINE CLNS Pilot to European sites
- Gated implementation of IDRP by Sue Hares and John Scudder is progressing.

IETF Tuba Demonstration: March, 1993

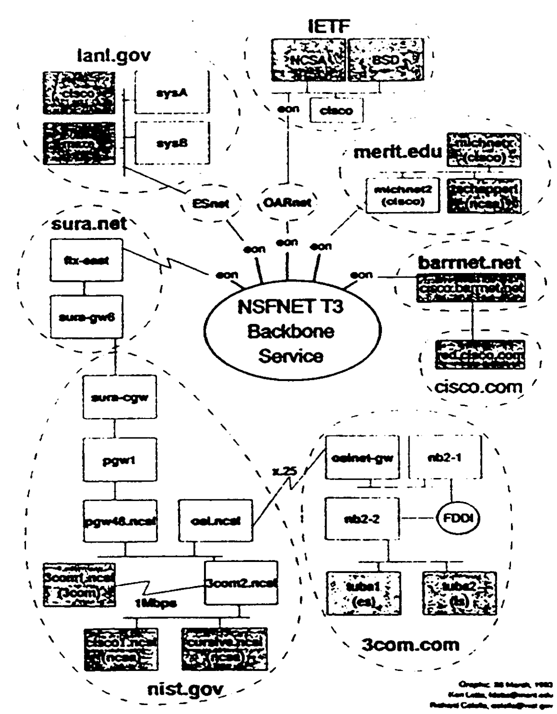

Work Left to Do
(or Agenda for TUBA WG mtg)

- Work with other groups (SIP/IPAE, PIP, etc.) on higher layer integration
- Continue implementation. Identify new TUBA platforms and applications.
- NOOP group continues to expand infrastructure (eg. ANS backbone to support CLNP native on AIX 3.2).
- Work with X3S3.3 on enhancements to CLNP for new functionality.
- Develop long term IPv7 addressing plan for TUBA.

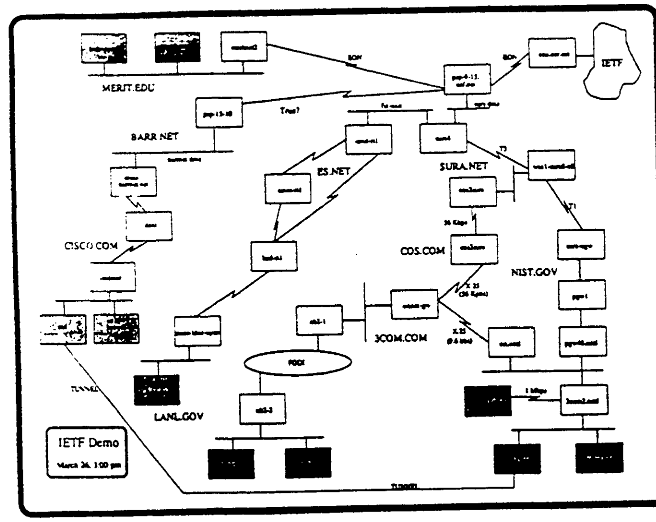

4.3 Next Generation of IP - SIP

Presented by Dave Sincoskie/Bellcore

Bio:*Dave Sincoskie is Executive Director of Computer Networking Research at Bellcore. He manages a group of about 30 researchers working on gigabit networking, high speed packet switching, internetworking, transport protocols, and computer network management. He is responsible for coordinating all of Bellcore's research work on the national Gigabit Testbeds, and has been heavily involved with the AURORA project. In 1989, Dr. Sincoskie was appointed Adjunct Associate Professor of Computer and Information Sciences at the University of Pennsylvania.*

Pip

Columbus IETF

Dave Sincoskie

Bellcore

Partridge/Kastenholz IPv7 Essential Criteria

- Scale to 10^12 hosts
- Robust service
- Transition
- Media
- Unicast and multicast
- Configuration, administration, operation
- Accounting
- Extensibility

Partridge/Kastenholz IPv7 Strongly Desired Items

- Guaranteed flows
- Mobility
- Secure operation

Current Pip Effort

- Paul Tsuchiya (Bellcore), Pip architect
- Chuck Davin (Bellcore), Pipe (Real-time Pip)
- Ramesh Govindan (Bellcore), Pip design and implementation
- Sue Thomson (Bellcore), Pip design and implementation
- Bala Rajagopalan (Bell Labs), Pip routing algorithm
- Tony Ballardie (UCL), CBT multicast design and implementation
- Zheng Wang (UCL), CBT implementation

Pip Features

- Separate naming from addressing
  - Host ID is name
- Pip "address" merges addressing and routing information
  - Host can have many addresses
  - Flexible ROAD without compromising performance
- Handling Directive encodes QoS/flow information
- Evolution mechanisms

Pip Overview

- New address encoding in Pip header (FTIF Chain)
  - Solves scaling, address depletion problems
  - Greatly simplifies address administration
  - Allows for efficient encoding of policy routes (provider selection)
  - Fast routing table lookup, even with policy routes
- Multiple addressing types, easily extendable
  - Unicast, "Class D" multicast, CBT multicast
- Evolvable
  - Especially evolution to real-time flows

Pip Implementation

- Host
- Router
- DNS
- Static configuration plus support software
- First implementation produced at Bellcore
  - code available on thumper.bellcore.com:pub/pip

Host Implementation

- Transmit and receive Pip packets
- Query DNS, receive Pip Addresses, ID, related information
- Choose among multiple Pip Addresses (source and destination) to format header
- Backwards compatible with existing IP application binaries
- Fragmentation/reassembly
- Pip ARP

Router Implementation

- Receive and forward Pip packets
  - Complete forwarding engine, including multicast, and evolution mechanisms
- Translate between Pip and IP
- Tunnel over IP

DNS Implementation

- Carries:
  - Host Pip ID
  - Pip Addresses of host or Pip/IP translation gateway
  - Provider information
  - Mobile Host Server name
  - Exit public data network addressing information
- Inverse queries on Pip ID or Pip Address

Pip Forwarding Engine Performance

Relatively un-tuned implementation of Pip versus SunOS 4.1.2

- Sparc 1
- Measured:
  - Time from packet in to packet out on Ethernet
  - Just forwarding part of header processing
- SunOS has a "last packet" only cache
  - Compared both cache hits and cache misses

Pip Forwarding Engine Performance

- Numbers are in µsec

	Pip	IP Cache Hit	IP Cache Miss
Ethernet	512	620	693
Forwarding Engine Only	25	11	61

Comments on Pip Performance

- Pip performance not dependent on cache hit
  - Performs uniformly well at low and high ends of hierarchy
  - Performs uniformly well for large and small routing tables
- Pip performance better than IP if IP routing table accessed
  - Even more so if routing table is bit-wise (CIDR), not Class A/B/C
- Pip performance comparable to IP when last packet cache hit

Pip Demo

- In terminal room
- Demonstrating:
  - All TCP/UDP applications running over Pip
  - Provider selection
  - Transition from IP
  - DNS

Demo (Logical) Topology

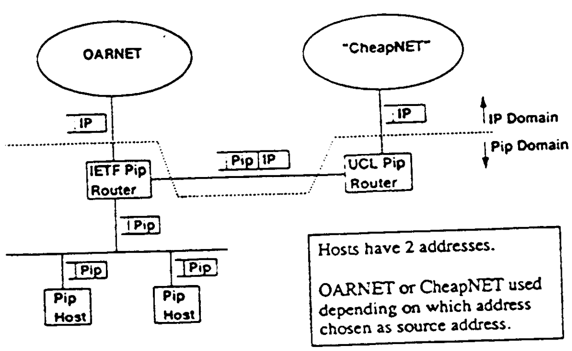

Demo User Interface

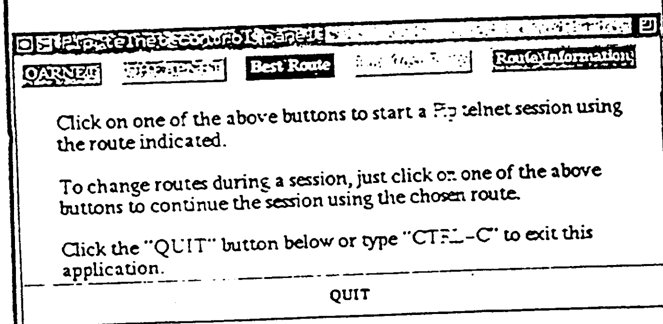

Transition Features in the Demo

- Outgoing packets translated from Pip to IP by terminal room and UCL Pip routers
- Incoming packets translated to Pip by terminal room Pip router
  - Incoming packets through OARNET only
- Pip packets transmitted from terminal room Pip router to UCL Pip router by tunneling over IP

> Pip systems can be introduced anywhere and be made part of Pip infrastructure
>
> Even a few Pip systems bring added functionality to user

Next Steps

- Start internet-wide testing of statically configured Pip in May
- Demo Pip with dynamic algorithms in Amsterdam
  - Routing
  - Router discovery/host address administration
  - Automatic domain-wide prefix administration
  - PMTU discovery
  - CBT multicast algorithm
- Internet-wide testing of dynamic Pip after July

Pipe (Pip Enhanced): Pip with Real-time Flow Support

- Pipe is protocol for real-time flow handling
  - Reservations
  - High-speed
- Leverages Pip ROAD technology
- Efficient interworking of public/private nets
- Demo at winter IETF
- Internet-wide testing afterwards

Conclusions

- Pip solves all existing internet protocol problems
- IPv7 will have to last at least 10 years
- Pip is flexible enough to solve future problems
- Pip can be deployed in time
- No performance penalty
- Interested in collaborators and deployment opportunities

4.4 IPAE

Presented by Bob Hinden/Sun

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).*

SIP / IPAE STATUS REPORT

March 29, 1993

Bob Hinden

OVERVIEW

- Introduction
- Status Summary
- Who is Involved
- Protocol Development
- Implementation and Testing
- Working Group Information

Sun Microsystems March 29, 1993

SIP / IPAE OVERVIEW

- SIP: Simple Internet Protocol
  - Evolution of IPv4
  - 64-bit Addresses
  - Header Simplification
  - Options moved to Separate Headers
- IPAE: IP Address Encapsulation
  - Transition Scheme for SIP
  - Uses Encapsulation and Translation
  - Flexible Deployment Scheme
  - Self Configuring IPv4 Compatibility

Sun Microsystems March 29, 1993

SUMMARY OF PROGRESS

- Lots of Meetings
  - Used MCI Phone Conferences and VAT
- Seven Implementations
- First SIP/IPAE Interoperability Event
  - March 24, 1993
- SIP Routing Protocols
  - SIP RIP
  - SIP OSPF
  - SIP IDRP
- DNS Specification
- BSD API Specification

Sun Microsystems March 29, 1993

ACTIVE PARTICIPANTS

• Beame & Whiteside	Implementation (PC)
• INRIA	Implementation (BSD, BIND), DNS & OSPF Specifications
• Intercon	Implementation (MAC)
• MCI	Phone Conferences
• Merit	IDRP for SIP Specification
• Network General	Implementation (Sniffer)
• SGI	Implementation (IRIX, NetVisulizer)
• Sun	Implementation (Solaris 2.x, Snoop, KA9Q)
• TGV	Implementation (VMS)
• Xerox PARC	*Steve Deering*
• Bill Simpson	Address Resolution / Router Discovery

Sun Microsystems March 29, 1993

PROTOCOL CLARIFICATIONS

- Pseudo Header Checksums
- Format of Addresses
  xxxx:xxxx:ddd.ddd.ddd.ddd
  xxxx:xxxx:xxxx:xxxx
- Options
  - Hop by Hop Options
  - Order of Option Processing

Sun Microsystems March 29, 1993

NEW CAPABILITIES

- Auto Configuration
  - *Local Use Addresses*

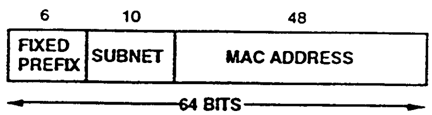

  - Plug & Play inside of a Site
  - Lookup *Global Use Address* from DNS inside of Site

Sun Microsystems — March 29, 1993

NEW CAPABILITIES (CONTINUED)

- Cluster Addresses
  - Identifier for Topological Region

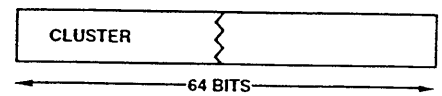

  - Very Loose Source Routes Through Cluster
  - Allows Source Selection of Providers
  - Mixture of Metro and Provider Based Addressing supports Changing Provider without Changing Host Addresses

Sun Microsystems — March 29, 1993

NEW DOCUMENTS

- *SIP: A Simple Internet Protocol*, S. Deering, to be published in May '93 IEEE Network.
- *SIP-RIP*, G. Malkin, C. Huitema, Internet Draft, draft-ietf-sip-rip-00.txt
- *IDRP for SIP*, S. Hares, Internet Draft, draft-ietf-ipidrp-sip-00.txt
- *OSPF for SIP*, C. Huitema, to be published as an Internet Draft.
- *SIP Addresses in the Domain Name Service Specifications*, C. Huitema, to be published as an Internet Draft.
- *SIP Program Interfaces for BSD Systems*, R. Gilligan, to be published as an Internet Draft.

Sun Microsystems — March 29, 1993

IMPLEMENTATIONS

OS	Organization	Status
BSD 4.3	INRIA	In progress.
DOS &Windows	Beame & Whiteside	Completed (telnet, ftp, tftp, ping)
IRIX	Silicon Graphics	In progress (ping)
KA9Q	Sun	In progress (ping)
Mac OS	Intercon	Completed (telnet, ftp, finger, ping)
Solaris	Sun	Completed (telnet, ftp, tftp, ping)
VMS	TGV	In Progress (ping)

TOOLS	Organization	Status
NetVisualizer	Silicon Graphics	Completed (SIP & IPAE)
Sniffer	Network General	Completed (SIP & IPAE)
Snoop	Sun	Completed (SIP & IPAE)

MISC	Organization	Status
Bind	INRIA	Code done

Sun Microsystems — March 29, 1993

TESTING CONFIGURATION

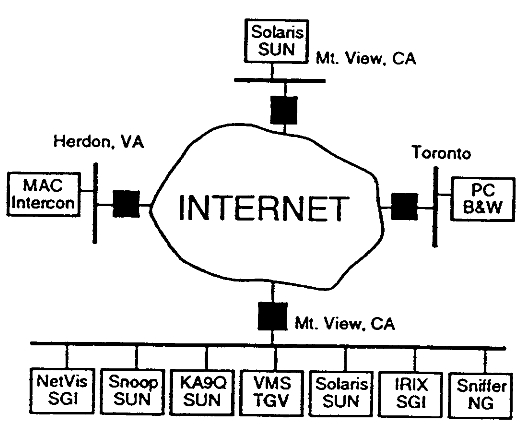

Sun Microsystems — March 29, 1993

WORKING GROUP INFORMATION

- Archive

  parcftp.xerox.com /pub/sip/

- SIP Working Group

  sip-request@caldera.usc.edu

- IPAE Working Group

  ip-encaps-request@sunroof.eng.sun.com

Sun Microsystems — March 29, 1993

IETF ACTIVITIES

- Demonstration of Implementations
  - Tuesday 6:00pm in Terminal Room
- SIP Meetings
  - Tuesday 9:30am - 12:00 Noon
  - Wednesday 1:30pm - 3:30pm
- IPAE Meeting
  - Monday 4:00pm - 6:00pm

Sun Microsystems — March 29, 1993

Chapter 5

Technical Presentations

5.1 An Architecture for Resource Management Networks

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.*

SUPPORTING REAL-TIME APPLICATIONS

An Architecture for Resource Management in Networks

David D. Clark
M.I.T. Laboratory for Computer Science

Joint work with lots of others....

SUPPORTING REAL-TIME APPLICATIONS

Who...

Scott Shenker, Lixia Zhang and Steve Deering
- **Xerox PARC**

Deborah Estrin, Sugih Jamin
- **USC**

Craig Partridge
- **BBN**

Chuck Davin, Andrew Heybey and others
- **MIT**

Members of the End-2-End Research Group:
Bob Braden, Van Jacobson, Sally Floyd

Abhay Parekh
- **MIT and IBM**

SUPPORTING REAL-TIME APPLICATIONS

What are the key questions?

What set of services can we implement?
- **Technical question: theory and practice.**

What set of services are needed?
- **What do applications actually require?**
- **What can be marketed?**

How do we write standards?
- **Specify mechanism.**
- **Specify objective.**

SUPPORTING REAL-TIME APPLICATIONS

What are the service requirements?

We identify two kinds of service requirements:

Performance requirements of individual applications.
- **Relate to delay objectives.**
- **Real time, interactive, bulk data.**
- **Ergonomic requirements.**

Controlled sharing among traffic aggregations.
- **Relates to bandwidth.**
- **Institutions, protocol families, applications.**
- **Economic requirements.**

SUPPORTING REAL-TIME APPLICATIONS

What is real-time?

Real-time traffic is characterized by a "playback" point.

- **Packets "on-time" or early are useful.**
- **Late packets are not useful.**

Note: not the same as interactive traffic.

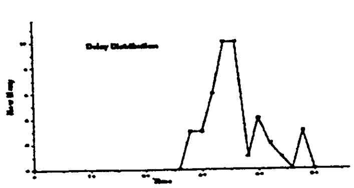

Playback time?

Loss % ?

SUPPORTING REAL-TIME APPLICATIONS

The playback point

Where does it come from?

Simple model (guaranteed service):
- **Negotiated at service setup time.**
- **Network assures it.**

Alternate model (predictive service):
- **Receiver observes arriving traffic.**
- **Receiver adapts to observed arrival time.**
- **Network provided stable service.**

A receiver that adapts must tolerate some late packets.
We can build guaranteed service.
Predictive service will give lower bounds.

Admission control

If we are to meet specified service requirements, then:

- We must have a characterization of the offered load.
  - Token bucket, etc.
- We must have an algorithm to decide if the request can be met.
  - Admission control.
- We must be able to refuse the request.

Other application requirements

Real-time traffic has playback point.

Elastic traffic can tolerate variation in delay.

- Continuous curve of user satisfaction.
- Wide range of delay objectives.
  - Can be as demanding as real-time.
- Much harder to characterize offered load.
  - Most systems today do not try.
  - No admission control.
- No attempt to characterize the service.
  - Best effort delivery.
  - Manage aggregates of sources.
- Management of QOS in background.
  - Purchase larger share, install more fiber.
  - Like highway system.

Link Sharing

Controlled sharing:

- Among customers
- Among protocols

Want overload assurance and statistical sharing.

Worst-case objective: nested sharing.

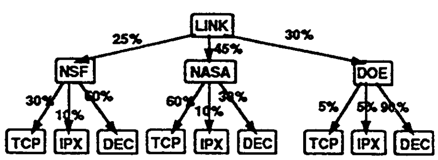

Summary: Service Requirements

Ergonomic:

- Providing real-time service
  - Guaranteed and predictive
- Providing lower delays for more demanding elastic services.

Economic:

- Provide isolation among protocol families.
  - Not just economic, but technical
- Provide isolation among customers.
- Achieve effective link utilization.
  - Aggregation of traffic.
  - Statistical sharing or real low loading.

Per-packet scheduling

We need one scheme that accomplishes all the goals:

- Ergonomic :guaranteed real time, predictive real time, and a range of elastic objectives.
- Economic: nested link sharing.

How can these be combined?

- Need an ordering on the requirements.

Real time "comes first".

- Delay of each individual packet matters.
- Admission control can achieve link sharing.

Guaranteed service takes precedence over predictive.

- Need not do better than guarantee.

Link sharing takes precedence over elastic objectives.

- No admission control.

The other parts of the solution

To this point I have discussed per packet scheduling.

Here are the other parts of the solution:

- Resource setup protocol -- propagate a request across the network.
- Route establishment -- selecting suitable links.
- The "flow spec" -- the description of the needed service.
- Admission control -- the algorithm to run at each switch to determine if the request can be met.
- Packet classifier -- map packets to resource class.

SUPPORTING REAL-TIME APPLICATIONS

The FLOW SPEC

Flow Spec – a specification of the needed service.
- RFC by Craig Partridge

Bandwidth:
- Average rate.
- Token bucket size.
- Peak rate.

Delay:
- Target max delay. (Not a hard limit.)
- Max jitter. (Not useful, I think.)

Sort of service:
- Guaranteed or Best Effort.

Errors:
- Approx. maximum tolerable.

R/T 13

SUPPORTING REAL-TIME APPLICATIONS

Admission Control

An algorithm run at each node to see if the requested service can be accepted.

Key issue: how to characterize the existing traffic.
- From the flow spec.
- From actual behavior.

We propose to measure actual usage.
- MUCH better link utilization.
- Risk of being misled.
- Presume "well-behaved" sources.
- Also sell the guaranteed service.

Any model like token bucket will (must be) a very loose bound. A tight model is much to complex.
- Consider video.

R/T 14

SUPPORTING REAL-TIME APPLICATIONS

Resource Setup Protocol

A protocol to propagate a resource request across a network.

One proposal: ST-II.
- "Hard setup".
- Uses special header.

"Our" proposal: RSVP
"Our" = Zhang, Shenker, Deering, Estrin, et al
- Supports "hard" and "soft" multicast.
- Works with IP header.
- One resource flow ≠ one end-point connection.

R/T 15

SUPPORTING REAL-TIME APPLICATIONS

Resources and connections

There is not a one-to-one relation between a data source and a resource reservation.

- Audio: Only a few speak at once, so share the reservation.
- Video: Might want to watch one source, or switch attention among sources.

The same is true of data destinations.

- If several destinations are jointly listening, only one reservation is required over shared links.

R/T 16

SUPPORTING REAL-TIME APPLICATIONS

Summary – project status

Per-packet scheduling: second generation implementation exists, and has been simulated. Must demonstrate efficiency.

Flow spec: Internet RFC published.

Admission control: preliminary paper published, more testing and evolution needed.

Resource setup: Algorithm has been implemented and tested in simulation. Currently being integrated into scheduling code.

Packet classifier: Hoping for code from Van Jacobson.

Overall IP test: DARTnet demo in spring/summer 1993.

R/T 17

SUPPORTING REAL-TIME APPLICATIONS

QOS and ATM

Could we do all this for ATM?

Remember efficiency and functionality?
- 53 byte cells = 2.7 us. @155 mb/s.

VCs provide instant classification.
Fixed cell size eliminates a multiply.

Big question:
Are we sharing bandwidth statistically among VCs?

Our research objective in the next year.
Show we can do this sort of control for ATM.
Decide if we need to.

R/T 18

What service is needed?

An application example: guaranteed vs. predicted real-time bounds.

- Are we selling a characterized service or a good one?
  - Predicted RT, delay targets, aggregated classes.

A marketing example: commercial services.

- This discussion of real-time suggests fine grained reservations.
- Customers want predictable rates.
  -> Aggregation for charging?

Two key questions?
- Can one create a cost-effective switched service?
- How does one add QOS to this service?

How will this be standardized?

Traditional model: standardize the formats and algorithms.
- Our queue scheme.

But...
There may be several valid ways to achieve the same results.
- Low link loading.
- Alternative realizations in different switch designs.

Alternative model: functional or performance spec.
- Building standards have both.
- Networks: routing standards?

Performance specs are HARD TO WRITE.

5.2 Internet Talk Radio

Presented by Carl Malamud

Bio: *Carl Malamud is the author of several professional reference books and the founder of Internet Talk Radio.*

INTERNET
TALK
RADIO

"Flame of the Internet"

INTERNET TALK RADIO

- Internet Service Provider
- Public Radio Analogy
- ~Trade Press, ~Journal
- News and Information for Internet

INTERNET
TALK
RADIO

WHY "RADIO"?

- Too Much ASCII in World
- Higher Semantic Content in Audio
- Radio Familiar Metaphor
- Video Too Hard, Too Big

INTERNET
TALK
RADIO

SOUND FORMATS

- Sun .au as Native Format
- PCM, 8-bit, 8 khz, mu-law
- 64,000 bps = 30 Mbytes/hour
- Conversion: Sound Exchange (SoX)
- [Your Format Here]

INTERNET
TALK
RADIO

DISTRIBUTION

- UUNET, EUnet, IIJ as Primaries
- FTP to Regionals
- MIME, NNTP, FTP to Locals
- Multicasting
- [Your Protocol Here]

INTERNET
TALK
RADIO

LEGAL MODEL

- Copyrighted Material
- Unlimited Noncommercial Copies
- No Derivative Works
- Heads & Tails Allowed

INTERNET
TALK
RADIO

GEEK OF THE WEEK!

INTERNET TALK RADIO

GEEK OF THE WEEK

- Flagship Show
- Interview Format
- Prominent Techno-Types
- Technology, Politics, Personalities

INTERNET TALK RADIO

SOME PROTO-GEEKS

- Marshall T. Rose
- Eric Huizer
- Daniel Karrenberg
- Daniel Lynch
- Milo Medin

INTERNET TALK RADIO

FEATURETTES

- Dear Ada ("Miss Manners")
- Internet Hall of Flame
- The Incidental Tourist
- Book Byte

INTERNET TALK RADIO

FINANCIAL MODEL

- Sponsors
- Sun Microsystems
- O'Reilly & Associates
- Short, Tasteful ACKs

INTERNET TALK RADIO

AUP

INTERNET TALK RADIO

CONFERENCE COVERAGE

- IETF and RIPE
- NET93 and INET93
- INTEROP
- [Your Name Here]

INTERNET
TALK
RADIO

LEVELS OF COVERAGE

- Rebroadcast
- Daily Summaries
- Analysis
- Live Coverage
- Sportscasters

INTERNET
TALK
RADIO

NEXT STEPS

VAPORWARE

- Daily News
- Mailing List Summaries
- Internet Traffic Reports
- Other Techie Topics (UNIX)
- Other Research Topics (Physics)

INTERNET
TALK
RADIO

BEYOND RADIO

VAPORWARE++

- Multicasting Key
- Interactive: Locally and Globally
- Multimedia: Video, Text, ...
- Multiprotocol: WWW, Gopher, ...
- Desktop Broadcasting?

INTERNET
TALK
RADIO

INTERNET TALK RADIO

carl@radio.com

5.3 Introducing the InterNIC - General Atomics/CERFnet

Presented by Susan Calcari/CERFnet

Bio:*Susan Calcari is the Information Scout for the InterNIC, tracking developments in network services, especially new resources, and acting as a liaison between the InterNIC and portions of the networking community. In her previous position with Merit Network, Inc., she spoke to national audiences about NSFNET and developed and produced the Merit Networking Seminars. Susan has over a decade of experience in voice and data networking with such companies as Sprint and Bank of America, and has a degree in Biological Sciences from Michigan Technological University.*

As a result of the NSF solicitation released last year for a "Network Information Services Manager for NSFNET and the NREN", three organizations have each been awarded a five year cooperative agreement to provide networking services to the research and education community. General Atomics/CERFnet will provide Information Services; AT&T will provide Directory and Database Services, and Network Solutions, Inc., will provide Registration Services. This presentation will discuss the services to be provided by each organization.

The Information Services provided by General Atomics/CERFnet will encompass three major areas - Reference Desk, Coordination, and Education. Information Services will provide services to the midlevel and campus NIC organizations so that they can then serve their customers more efficiently. In addition, InterNIC Information Services will act as the "NIC of first resort" and the "NIC of last resort" for end users. These and other aspects of the InterNIC project will be discussed, as well as coordination between InterNIC team organizations.

InterNIC Information Services

Susan Calcari
InterNIC Info Scout
General Atomics/CERFnet

GENERAL ATOMICS

The InterNIC

- Registration Services
  - Network Solutions Inc. (NSI)
- Directory and Database Services
  - AT&T
- Information Services
  - General Atomics/CERFnet

GENERAL ATOMICS

InterNIC Information Services

- NIC of NICs
  - service to midlevel and campus NICS
- Feedback sought and appreciated
- Responsive and dynamic operation

GENERAL ATOMICS

InterNIC Information Services

- Reference Desk
- Coordination
- Education
- InterNIC Team Coordination

GENERAL ATOMICS

Reference Desk

- "NIC of first and last resort"
- Midlevel, campus, international NICs
- InterNIC Directory and Database Services
- InterNIC Registration Service

GENERAL ATOMICS

Info Source Contents

- Starter materials
- Selected resource lists
- Discipline Specific Packets
- Pointers to other archives

GENERAL ATOMICS

Info Source Online Access

- FTP, telnet, mail
- Archie, WAIS, gopher
- NICLink

NICLink

- FrameViewer and FrameReader
- PC, Macintosh, NeXT, and UNIX
- Diskette, CD-ROM

Discipline Specific Info Packets

- Field Specialists
  - Biology
  - Chemistry
  - Networking
  - Libraries
  - K-12

Info Scout

- A five year mission to to explore new worlds, seek out new tools and resources, to boldly go where no Internaut has gone before...

Info Scout

- Scout out new resources
- Scout out new tools
- Stay current with tools under development

InterNIC Mailing List

- Intended for end users, NICs
- Announcements only
- Collaborative project
  - nis@ InterNIC.net
  - listserv@InterNIC.net
  - subscribe nis your name

Quality Evaluation

- Tracking
- Trouble Tickets
- Reports
- IQ Scores

GENERAL ATOMICS

Coordination Services

- InterNIC Liaison Council
- International Cooperation
- Representation to the Community

GENERAL ATOMICS

InterNIC Liaison Council

- InterNIC Liaison Council Membership
- Encourage information sharing
- Assure end user's needs are recognized

GENERAL ATOMICS

Community Outreach

- IETF, FARNET, CNI, CIX, ISOC, EDUCOM, SIGGUCCS
- NIS Fest
- NIS Minifests

GENERAL ATOMICS

Training

- Use the experts
- Work with the midlevels and campuses
- Offer established courses at a discount

GENERAL ATOMICS

InterNIC Team Coordination

- Unified InterNIC Interface
- Common trouble ticket system
- Joint community activities

GENERAL ATOMICS

Unified InterNIC Interface

- Common telephone identity
  - 800.444.4345
- Common electronic identity
  - is@InterNIC.net

Individual Contact Info

- Information Services
  - info@InterNIC.net
- Directory and Database Services
  - admin@ds.InterNIC.net
- Registration Services
  - hostmaster@InterNIC.net

5.4 Introducing the InterNIC - Network Solutions, Inc.

Presented by Scott Williamson/Network Solutions, Inc.

InterNIC Registration Services
IETF

Presented by

Scott Williamson

scottw@internic.net

April 1, 1993

NETWORK SOLUTIONS

InterNIC Registration Services
IETF

April 14, 1993

NETWORK SOLUTIONS

Registration Services

- Internet protocol numbers assignment
- Autonomous system number assignment
- Domain registration
- Inverse address registration
- Registration help service

NETWORK SOLUTIONS

Internet Address and ASN Assignment

- Assign IP network numbers

	NIC	IANA
Class A	Pre-evaluate	Approve
Class B	Evaluate and assign	Provide criteria
Class C	Evaluate/Assign/Delegate	Provide guidance

- Assign blocks of numbers ot delegated registries
- Assist with the management of the Class B address space
- Assign and delegate autonomous system numbers

NETWORK SOLUTIONS

Enhancements to Internet Address Assignments

- Automation
- Increased error checking
- Further delegation (Pacific)

NETWORK SOLUTIONS

Domain Registration

- Operate a root domain server
- Generate/Manage the root zone files
- Register top and second level domains
- Assist development of policy and naming conventions (K–12)
- Administer the .US domain using USC ISI

NETWORK SOLUTIONS

Domain Registration Enhancements

- Fully automated registration with oversight
- Zone releases Monday, Wednesday, and Friday
  (Data cutoff Sunday/Tuesday/Thursday at 11:59 p.m.)
- Development and integration of .US domain by USC ISI

NETWORK SOLUTIONS

Delegated Registration Processing

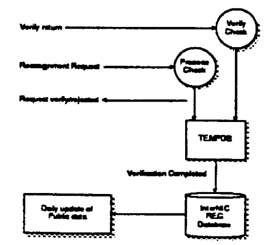

NETWORK SOLUTIONS

Registration Help Desk

- Operating from 7:00 a.m. to 7:00 p.m. Monday – Friday
- Assist users in answering policy or registration status questions
- General questions will be answered by the InterNIC Information Services Provider

NETWORK SOLUTIONS

Information Via Public FTP

- Domain
  - Domain zone files
- Templates
  - Current IP and domain templates
- Policy
  - Documents related to IP and domain policy

NETWORK SOLUTIONS

Transition Issues

- Splitting DDN data from non-DDN data
- Moving non-DDN to InterNIC registration facility
- Informing the community of schedule and change (RFC)

NETWORK SOLUTIONS

Zone File Release

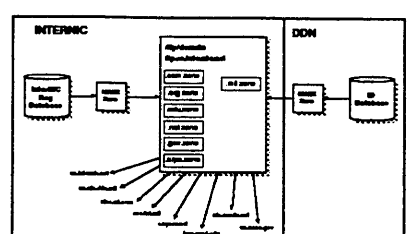

NETWORK SOLUTIONS

Current E-Mail Processing

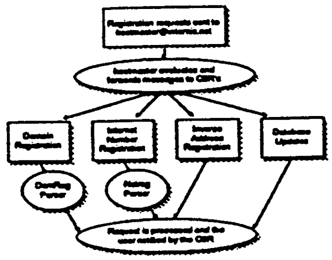

NETWORK SOLUTIONS

Automatic Registration Processing

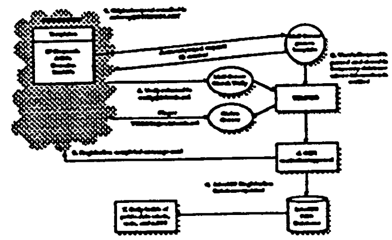

NETWORK SOLUTIONS

The Future

- Client tools
- Internet addressing schemes
- Additional delegated registries

NETWORK SOLUTIONS

Delegated Registries

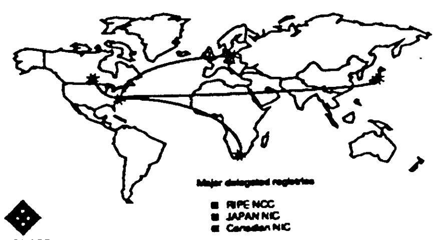

NETWORK SOLUTIONS

5.5 Introducing the InterNIC - AT&T

Presented by Rick Huber/AT&T

Bio:*Rick Huber received a BS in Electrical Engineering from MIT in 1970, and an MS and PhD in Computer Science from SUNY Stony Brook in 1972 and 1975. From 1975 to 1977, he was Assistant Professor of Computer Science at Texas A&M University. In 1977 he joined the technical staff at AT&T Bell Laboratories where he has worked on a variety of data communications and computer projects. He is currently Technical Manager of the Network Servers Planning Group at Bell Laboratories in Holmdel, NJ.*

InterNIC
Directory and Database Services

Rick Huber

AT&T Bell Laboratories

AT&T

Service Philosophy

Provide the user community with timely, accurate, and easy access to a wide range of information resources on the Internet

AT&T

Directory and Database Services
Service Summary

- Directory of directories
  - Help users locate information resources in the Internet
- Directory Services
  - Users accessible by name, discipline, and organization
  - Institutions on the network and characteristics of their connections
  - Organizations accessible by function
  - Resources available on the network
- Database Services
  - Databases of contributed materials
  - Databases of communications documents (RFCs, FYIs, IENs, etc.)
  - Databases maintained for other groups

AT&T

Directory of Directories

- Contents
  - Resource description files
    - Name, keywords, brief description, access information
  - Validation
- Tools
  - WAIS
  - Archie
  - FTP
  - Gopher (July 1993)
  - X.500 (future)
- Fees
  - No fees for access
  - Basic listings (limited space for description) are free
  - Extended listings available for a fee

AT&T

Directory Services

- "White pages"
  - Search by name
- "Yellow pages"
  - Search by category
- Contents
  - Registration information
  - Individual listings
  - Links to other directory servers
  - Encourage and assist organizations setting up their own systems
- Tools
  - X.500 - primary tool
  - WHOIS via X.500
  - Netfind
- Fees
  - No fees

AT&T

Database Services

- Contents
  - Files requested by NSF
  - Communications documents
  - Data stored for other organizations
  - Database design, maintenance, and management by special arrangement
- Tools
  - FTP
  - Gopher (July 1993)
- Fees
  - No fees for access
  - Storage fees for data stored for other organizations
  - Fees for design, maintenance, and management of databases

AT&T

InterNIC

Access to Services

- Access from remote WAIS and X.500 clients
- Telnet access to WAIS, X.500, Archie, and Netfind on our server
- Mail access to WAIS, X.500, Archie, and file transfer
- FTP access to files on our server, including resource description files

AT&T

InterNIC

Directory of Directories

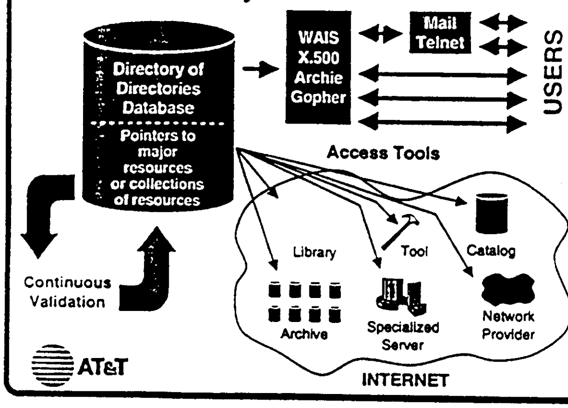

AT&T

InterNIC

Quality

- Customer support objectives
  - Service available 24 hours a day, seven days a week
  - 24 hour emergency support for equipment
  - Primary and two backup servers at different sites
  - Staffing during regular business hours (8AM-8PM Eastern time) to address customer questions and concerns
  - Staff accessible by email, fax, and phone

AT&T

InterNIC

Quality

- Response time and content
  - Measure and report response times, concurrent sessions, etc.
  - Use data to tune system
  - Collect and analyze data on failed requests to find areas where more data is needed
  - User satisfaction surveys

AT&T

InterNIC

Quality

- Privacy and accuracy
  - No resource will be listed without permission from owner
  - Periodically check with owner of resource to see if listing is still current

AT&T

InterNIC

Service Evolution

- User outreach
  - Trade shows
  - Affinity groups
- Make the data useful
  - Surveys
  - Analysis of queries
- Train users and administrators
  - Introduction to directory and database services
  - How to set up X.500
  - X.500 "starter kit"
- Encourage discussion and growth
  - Organize seminars for discussion and to track trends

AT&T

Technology Evolution

- Track enhancements to our tools
- Provide feedback to tool developers based on our users' experience
- Consider new tools for inclusion in our services as they become available

AT&T

Directory and Database Services

Contact Information

Mail Server:	mailserv@ds.internic.net
Anonymous FTP:	ds.internic.net, introductory info is in directory pub/InterNIC-info
Telnet:	ds.internic.net, login as guest for general info
Administrator:	admin@ds.internic.net

AT&T

5.6 Qualcomm Wireless

Presented by Phil Karn/Qualcomm

Wireless Data Communications

Phil Karn, KA9Q

Qualcomm, Inc
San Diego, CA
karn@qualcomm.com

with credit to:
Dewayne Hendricks, WA8DZP
Tetherless Access, Ltd
Fremont, CA

Wireless Data Overview

- Amateur (Part 97)
- Unlicensed (Part 15)
- Mobile Satellite
- Private Microwave (Part 94)
- Land Mobile (Part 90)

  Specialized Mobile Radio (SMR)
- Cellular (Part 22)
- PCS/PCN

Amateur (Part 97)

- Wide range of frequency bands from MF thru microwave
- Power levels up to 2KW
- Relatively flexible technical rules
- Individual operators must pass a license test
- The mother of all "acceptable use" policies: personal, noncommercial use only, no encryption, third party traffic must be monitored
- Early pioneering efforts, slow to evolve since

Unlicensed (Part 15)

- Rules also cover kiddie walkie-talkies, baby monitors, cordless phones, garage door openers, etc.
- "Steerage class" (lowest priority access to spectrum)
- Rules added to permit up to 1W of spread spectrum on Industrial, Scientific & Medical (ISM, aka "garbage") bands, including

  902 MHz (shared with amateur, automatic vehicle monitoring)
  2.4 GHz (shared with microwave ovens)
  5.7 GHz
- Many products: NCR Wavelan, Proxim, Western Multiplex LYNX, etc; data rates up to several megabits/s
- IEEE 802.11

Omnitracs

- Qualcomm satellite mobile messaging service
- Uses existing Ku-band satellite transponders (Gstar-1)
- Spread spectrum
- 5-15 kb/s total forward link capacity, 50 b/s/user reverse link (100's simultaneous users)
- Position reporting
- Primary market: long-haul trucking of high-value cargo
- Coverage in CONUS, Europe; expanding to Australia, South America, Japan

SMR

- "Two way" radios on 850 MHz band, next to cellular
- Used mainly for voice, with dynamic channel allocation (trunking)
- Two large dedicated packet data networks: ARDIS and RAM
- Narrowband RF modems, relatively low capacity and high cost
- RAM being resold into Internet community: RadioMail, PSIlink

Cellular

- Analog FM cellular a real phenomenon - 5 million US subscribers in 1990
- Analog systems hitting growth limits in major markets
- Capacity needs main motivation for digital
- Original digital standard, TDMA, promises 3x FM capacity
- Newer digital technology, CDMA, demonstrates 10-15x FM capacity
- Digital cellular either bane or boost to data services

Cellular Digital Packet Data

- Aka "Celluplan II"
- Data overlay for existing FM cellular system
- Packet radio on idle FM cellular channels, secondary to voice
- Slow frequency hopping to even channel loading
- Will support TCP/IP and OSI
- Players: 6 RBOCs (all but Bell South), McCaw, Contel, GTE, IBM, PCSI
- Stopgap until digital cellular

CDMA Digital Cellular

- Begun by Qualcomm in 1989
- Underdog technology (TDMA already standard)
  Now in ballot as TIA PN-3118/3119
- Greater system capacity and better quality:
  1:1 frequency reuse, multipath discrimination,
  soft handoff, strong FEC, tight power control
- 1.25 MHz spreading bandwidth
- 50 Hz (20 ms) frame rate for variable rate vocoder
- Four frame sizes: 16, 40, 80, 171 data bits
- Complete technical specs available by anonymous FTP:
  ftp.qualcomm.com in /pub/cdma/*ps.z

PCS

- New allocations proposed by FCC at 1.8 GHz
- Becoming all things to all people:

  Alternative to cellular duopoly
  Wireless local loop
  Intra-building telephony ("cordless phones")
  Local area computer networks
- Qualcomm running existing 1.25 MHz CDMA on 1.8 GHz
  "PCS 2000" in initial design stages (higher data rates and/or wider spreading bandwidths)

Data PCS

- Aspect of PCS - FCC proposing 1.91-1.93 GHz for "user provided" local data networks
- Similar in principle to Part 15, except with dedicated spectrum
- Unlicensed operation, no carriers, no charges - user stations cooperate to minimize interference, maximize capacity
- Closest in principle to DARPA packet radio (e.g, SURAN)
- WINforum vendor group

CDMA Data

- Prototype philosphy - use existing CDMA system and air interface with minimal changes to mobile to prove concept, develop protocols
- Support three main data services:
  - Asynchronous dialup modem emulation
  - Group III Fax
  - Direct access to packet switched network (e.g., Internet)
- Approach: first build general purpose packet data facility, use as foundation for other services (fax, modem, etc)
- Packet data well suited to variable rate CDMA system

CDMA Mobile Phone

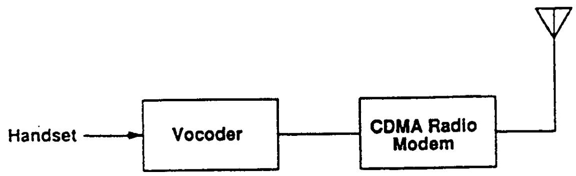

- Since CDMA is digital, we want to use the CDMA Radio Modem directly - there's no need for an external modem
- Vocoders can't pass modem signals, so we must bypass them on both ends, mobile and base

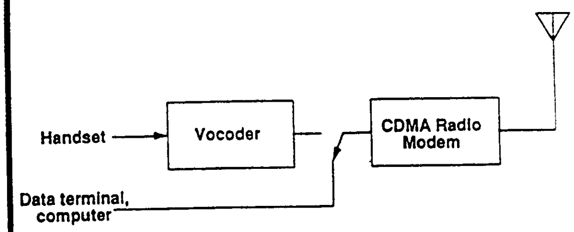

- Add mobile software to allow external PC to place calls, send and receive raw CDMA data through mobile diagnostic port
- Temporarily use CDMA mobile-to-mobile capability until data support is added to QTSO

Prototype Data Mobile Station

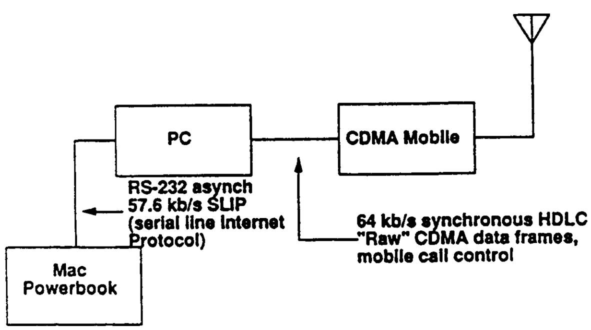

Prototype Data Base Station

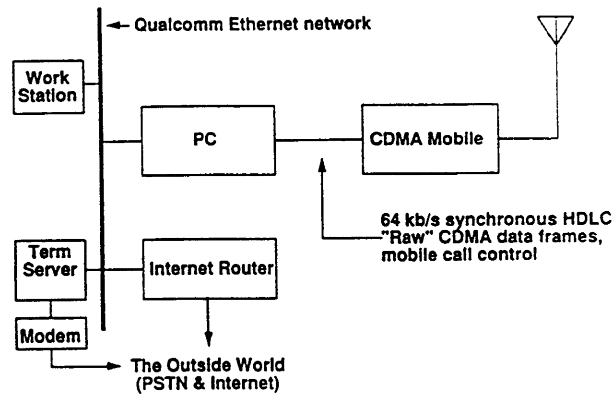

PC Protocol Functions

- Runs KA9Q NOS networking package (TCP/IP) with Internet client, server and IP packet switching capabilities
- Modules added: CDMA-specific radio link protocol, call control
- Convenient platform for protocol development and experimentation
- Lower level functions can move into mobile firmware

CDMA Radio Link Protocol

- Custom protocol designed to carry simple, general purpose data packets through a standard CDMA system
- Packets divided into CDMA frames for transmission, reassembled into packets at receiver
- Brings up CDMA call when needed, drops after an idle timeout
- Selectively retransmits erased CDMA frames
- "Lightweight" protocol intended to improve performance of upper layer protocols (e.g., TCP/IP), not replace them.
- Performs very well in both lab and mobile field tests
- Typical user data throughput: 850-950 bytes/sec (6800-7600 bits/sec), excluding all protocol overhead

RLP Structure

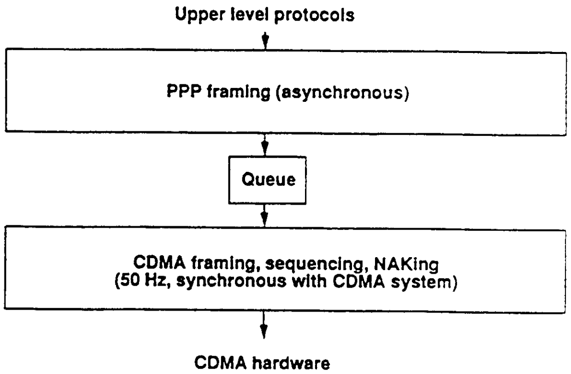

CDMA RLP Data Packet

(as seen by upper layer protocols)

PID	User Data	CRC-CCITT	Flag 0x7e
8	0-16384 bits	16	8

0x21 = IP
0x2d = compressed TCP/IP
0x2f = uncompressed TCP/IP
etc (standard PPP values)

RLP Selective Retransmission

- RLP data frames have 8-bit sequence numbers, incremented for each data frame
- Receiver resequences frames and NAKs missing sequence numbers; sender selectively retransmits requested frames
- No positive acknowledgements in RLP, but upper layer protocol (e.g., TCP) may provide them
- NAKs retransmitted (with increasing urgency) up to three times
- Packets are still lost if/when:

  The call drops
  NAKs aren't answered after three round trip times
  The CDMA 12-bit (full rate) and 8-bit (half rate) CRCs fail to detect a data error

CDMA Data Demo

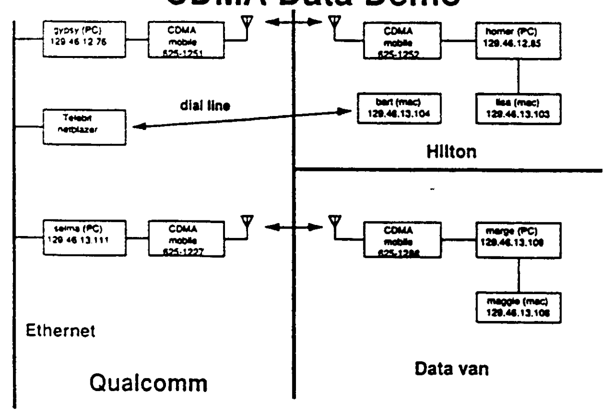

Future Directions

- Add data support to base station (QTSO vocoder selector)
- Move radio link protocol into mobile
- Add modem emulation (AT command set) to mobile
- Add support for G3 fax (local RJ-11 jack & modem) to mobile or external fax adaptor
- Simultaneous data and voice? (Requires PN-3118 extensions)
- Connectionless channel access? (e.g., credit card verification)
- TR 45.5.1.5 CDMA data service task group

Chapter 6

Workshop for Working Group Chairs

6.1 Workshop for Working Group Chairs

Presented by David Crocker/Silicon Graphics

Chairing a working group is not all fame and glory. Most experienced Chairs can be distinguished by the hollow look in their eyes from trying to walk the fine line between being fair to working group members, while still making forward progress. They have had the pleasure of doing this while receiving only informal guidance from other, experienced members of the IETF. This session attempted to provide a cohesive view of the responsibilities and authorities of a working group Chair.

In terms of content, the training session reviewed the formal structure of the IETF, the formal process for creating working groups and standards-track specifications and then ventured into the dangerous territory of building rough concensus within a reasonable amount of time.

The workshop was intended for potential working group Chairs, and those of recent vintage. Experienced Chairs and all others with an interest in the training were eagerly welcomed. The more voices of experience, the better.

Workshop for Working Group Chairs

D. Crocker

Silicon Graphics
dcrocker@sgi.com / +1 415 390 1804

Hello

- INTRODUCTION
  - What are we doing here, at this hour?
  - There is only rough consensus about the process
- AGENDA
  - I. IETF structure
  - II. Formal process
  - III. The inner working group
  - IV. Conflict resolution

The need for working group chair training

- IETF LARGE, DIVERSE
  - Process increasingly formal
  - No voting means (very) rough consensus
- DIFFICULTY MAKING PROGRESS AND BEING FAIR
  - Listen to all points of view
  - Keep working group focus
- CHAIRS OFTEN UNCLEAR ABOUT LIMITATIONS AND AUTHORITIES

Documents

- THE INTERNET STANDARDS PROCESS (RFC 1310)
- WORKING GROUP GUIDELINES (IN PROGRESS)
- ["EVOLVING THE SYSTEM" IN INTERNET SYSTEM HANDBOOK, LYNCH & ROSE, EDS.]

I. IETF Structure

ISOC	Internet Society legal cover
IAB	Internet Architecture Board Design cohesion, process appeals, IETF liaisons
IETF secretariat	Staff support
IESG	Internet Engineering Steering Group IETF oversight
AD	Area director Oversight for specific working groups
WG chair	Manage a working group to a productive end
Working group	The people who do the work

Working group roles

- CHAIR OVERSEES ENTIRE PROCESS, BUT:

Facilitator	Process management, things fair, focused, on time
Judge	Evaluation of technical options and driver towards "right" choice
Scribe	Record-keeper and editor of documents

- WORKING GROUP IS JURY, PROVIDING IDEAS, REVIEW, CONSENSUS
- DESIGN TEAM IS PRIMARY ADVOCATE AS SELF-SELECTING GROUP WITH COMMON VISION, PROVIDING CORE EFFORT

Formal Process

- Formal labels for a specification
- Developmental steps
- Acceptance criteria

Formal labels

Internet draft:	no official standing, fluid working document
Proposed Standard:	stable spec, no known errors, might have implementation
Draft Standard:	multiple, interoperable implementations testing all functionality
Internet Standard:	field experience and clear community acceptance (and use)

IETF acceptance criteria

Competence:	technically sound
Constituency:	providers & users
Coherence:	clear writing
Consensus:	rough but clear

II. Developmental steps

0. Birds of a feather (BOF)
   - "Market research" to determine interest and ability to pursue topic
   - Optional, one-shot meeting
1. Charter

Role:	Public announcement & project management plan
Scope:	What is to be pursued
Approach:	How it will be pursued
Product:	What will be delivered
Checkpoints:	Milestones and dates

II. Developmental steps

2. Document specification
   - Clarity of purpose
   - Clarity of writing
   - Clarity of solution
3. WG consensus
   - Clearly dominant agreement
   - Diversity of opinion about solution may be resolved by agreement to make *some* decision
   - Agreement about parts may permit eventual agreement about whole

II. Developmental steps

4. Area director approval
   - Technical review
   - Process review
   - Independent review when results of wg in question
5. Submission to IESG
   - Via secretariat & AD

II. Developmental steps

6. LAST CALL
   - Request for final feedback from IETF
   - Intended to detect major errors in process or content that might have slipped through cracks
   - Not intended as formal, full review
7. IESG REVIEW (& APPROVAL)
   - May conduct independent review

II. Developmental steps

(7.5) IAB CONFLICT RESOLUTION
   - If formal challenge not resolved by IESG

8. RFC PUBLICATION
   - RFC editor has publication criteria

III. The inner working group

- THE LIVES OF A CHAIR
- GROUP STYLE
- GROUP ROLES
- DEVELOPMENTAL PHASES (PROBLEM SOLVING 101)
- VENUES
- DEBATE
- CONFLICT MANAGEMENT

The lives of a chair

- HOW TO KEEP FROM BEING SAT ON
  - Agenda & schedule
  - Adequate debate, but not more than that
  - Maintain clear focus
  - Rehash *only* if constructive *and* working group desires
- PROACTIVE MANAGEMENT
  - Maintain pressure for forward progress
  - Escalate to IETF management when progress stalled

Working group style

- FREE-FLOWING
  - Cohesive group
  - Clear purpose
- TIGHTLY-MANAGED
  - Complex topic
  - Group diversity
  - Major differences in philosophy

WG management roles

Facilitator :	ensuring fairness and a thorough airing of views and alternatives
Judge :	evaluation of choices and movement towards choice
Scribe:	keeping track of things
Design team:	Primary advocates for the core effort, when wg diverse & topic complex; <u>must</u> work to keep wg consensus
Working group:	Jury & other contributors

Problem solving 101

- Problem statement
- Solution exploration
- Solution adoption
- Specification refinement

Discussion & decision venues

- Email
  - International participation
  - Inefficient, but extensive
  - The real place for consensus
  - Can be run as "meetings"
- Face-to-face meetings
  - Well-advertised ahead of time
  - Inherently restricted attendance
  - Limited time
  - Need for clear agenda and crisp management

Email vs. Meetings (One person's perspective)

- WG results must show approval based on *entire* working group
- Meetings have limited attendance
- Treat meetings as "strong indicator" prima facie basis for decisions
- Ensure verification through email

Debate

- Can clarify purpose, implications, alternatives
- Can tear the group apart
- Must be tolerated and even encouraged, until resolution or impasse

IV. Conflict Resolution

- Preferable to solve within working group
  1. Conflict types
  2. Timing of objections
- Often can't
  3. Chain of appeal

1. Conflict types

- Technical

Specific detail:	minor vs. show-stopper
Basic philosophies:	rarely resolved

- Process

Unfair practice:	usually claim against wg chair
Topic missed:	oops. (showstopper?)

2. Timing of objections

- TECHNICAL SHOWSTOPPERS WELCOME ANYTIME
- SMALL DETAILS WELCOME *ONLY* AT TIME WG COVERS THE SUBJECT
- PHILOSOPHICAL DEBATE WELCOME *ONLY* AT TIME WG MAKING DECISIONS ABOUT APPROACH
- UNFAIR PRACTICE COMPLAINTS ALLOWED *WHENEVER* INFRACTION FELT
- WG MAY ALLOW TOPIC TO BE RE-OPENED IF WG FEELS ISSUE COMPELLING OR NEW ALTERNATIVE INTRIGUING.

3. Chain of appeal

- WG CHAIR
- AREA DIRECTOR
- AREA DIRECTOR FOR STANDARDS MANAGEMENT
- IETF CHAIR
- IESG
- IAB

If you can keep your head when those around you...

- Most IETF members are remarkably well-intentioned
- Differences happen
  - Tempers often flare, but then settle down
  - Not all differences can be settled
  - When minority view clearly will not sway working group, respect the opinion, but move on
- Ask questions
- Make it happen!

Appendix A

Attendees

Vikas Aggarwal
JvNCnet
Phone: 609-258-2403
Fax: 609-258-2424
aggarwal@jvnc.net

Masuma Ahmed
Bell Communications Research
Phone: 908-758-2515
Fax: 908-758-4192
mxa@sabre.bellcore.com

Robert Aiken
Department of Energy
Phone: 301-903-5800
Fax: 301-913-7774
aiken@es.net

N. Akiko Aizawa
NACSIS National Center for
Science Information Systems
Phone: +81 3 3942 6994
Fax: +81 3 5395 7064
akiko@nacsis.ac.jp

Kannan Alagappan
Digital Equipment Corporation
Phone: 508-486-5953
Fax: 508-486-5279
kannan@dsmail.lkg.dec.com

Steve Alexander
Lachman Technology, Inc.
Phone: 708-505-9555
Fax: 708-505-9574
stevea@lachman.com

Michael Allen
IBM Corporation
Phone: 919-254-4442
Fax: 919-254-5483
moallen@ralvmg.vnet.ibm.com

Claudio Allocchio
National Institute for
Nuclear Physics - Italy
Phone: +39 40 3758523
Fax: +39 40 226338
Claudio.Allocchio@elettra.trieste.it

Guy Almes
Advanced Network and
Services, Inc.
Phone: 914-789-5333
Fax: 914-789-5310
almes@ans.net

Philip Almquist
Consultant
Phone: 415-752-2427
almquist@jessica.stanford.edu

Harald Alvestrand
SINTEF DELAB
Phone: +47 7 59 70 94
Fax: +47 7 53 25 86
Harald.Alvestrand@delab.sintef.no

Michael Anello
XLNT Designs, Inc.
Phone: 619-487-9320
Fax: 619-487-9768
mike@xlnt.com

David Arneson
Cabletron Systems, Inc.
Phone: 603-337-5064
Fax: 603-424-8121
arneson@ctron.com

Jules Aronson
National Library of Medicine
Phone: 301-496-9300
Fax: 301-496-0673
aronson@nlm.nih.gov

Nagaraj Arunkumar
3Com Corporation
Phone: 408-764-5104
Fax: 408-764-5002
nak@3com.com

Randall Atkinson
Naval Research Laboratory
atkinson@itd.nrl.navy.mil

Karl Auerbach
Empirical Tools and Technologies
Phone: 408-427-5280
Fax: 408-427-5281
karl@empirical.com

Robert Austein
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sra@epilogue.com

Dennis Baker
Wellfleet Communications, Inc.
Phone: 617-280-2349
Fax: 617-275-5977
dbaker@wellfleet.com

Fred Baker
Advanced Computer Communications
Phone: 805-685-4455
Fax: 805-685-4465
fbaker@acc.com

Anthony Ballardie
University College London
Phone: +44 71 387 7050
Fax: +44 71 387 1397
A.Ballardie@cs.ucl.ac.uk

Paul Barker
University College London
Phone: +44 71 380 7366
Fax: +44 71 387 1397
p.barker@cs.ucl.ac.uk

Jim Barnes
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-3159
barnes@xylogics.com

William Barns
The MITRE Corporation
Phone: 703-883-6832
Fax: 703-883-7142
barns@gateway.mitre.org

Tony Bates
RIPE Network Coordination Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
tony@ripe.net

David Battle
SNMP Research, Inc.
Phone: 615-573-1434
Fax: 615-573-9197
battle@cs.utk.edu

Serpil Bayraktar
Advanced Network and Services, Inc.
Phone: 313-764-3574
sbb@noc.ans.net

Jordan Becker
Advanced Network and Services, Inc.
Phone: 914-789-5308
Fax: 914-789-5310
becker@ans.net

Robert Beer
Ohio Northern University
Phone: 419-772-2362
Fax: 419-772-1932
r-beer@onu.edu

Jim Beers
Cornell University
Phone: 607-255-2627
Fax: 607-255-9086
Jim.Beers@cornell.edu

Gabe Beged-Dov
Hewlett-Packard
Phone: 503-750-4415
Fax: 503-750-3788
gabe@cv.hp.com

Nutan Behki
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
Nutan_Behki@qmail.newbridge.com

Tom Benkart
ACC Systems
Phone: 410-290-8100
Fax: 410-290-8106
teb@acc.com

Lou Berger
Louis Berger Consulting
Phone: 301-530-9790
lberger@bbn.com

Jeffrey Berk
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-332-1739
berk@ctron.com

Tim Berners-Lee
CERN
Phone: +41 22 767 3755
Fax: +41 22 767 7155
timbl@info.cern.ch

Mahesh Bhatia
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-424-8121
bhatia@ctron.com

Andy Bierman
SynOptics Communications, Inc.
Phone: 408-764-1673
Fax: 408-988-5525
abierman@synoptics.com

Richard Bjers
University of Cincinnati
Phone: 513-556-9012
Fax: 513-556-2042
rich.bjers@uc.edu

Russell Blaesing
Open Networks
Engineering, Inc.
Phone: 313-996-9900
Fax: 313-996-9908
rrb@one.com

B. Sue Blair
Advanced Network and
Services, Inc.
Phone: 313-763-3448
Fax: 313-764-6429
bsb@noc.ans.net

Larry Blunk
Merit Network, Inc.
Phone: 313-763-6056
Fax: 313-747-3745
ljb@merit.edu

John Boatright
NASA-KSC
Phone: 407-867-7278
Fax: 407-867-4079
bryan_boatright@ksc.nasa.gov

Kevin Bohan
Proginet Corporation
Phone: 516-228-6613
Fax: 516-228-6605
KBohan@mcimail.com

Fred Bohle
Interlink Computer Sciences
Phone: 301-317-6600
Fax: 301-317-6608
fab@interlink.com

David Bolen
Advanced Network and Services, Inc.
Phone: 914-789-5327
Fax: 914-789-5310
db31@ans.net

Nathaniel Borenstein
Bell Communications Research
Phone: 201-829-4270
Fax: 201-829-5963
nsb@bellcore.com

Erik-Jan Bos
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
erik-jan.bos@surfnet.nl

Rebecca Bostwick
Lawrence Livermore National Laboratory
Phone: 510-422-4017
Fax: 510-422-1482
bostwick@es.net

Rich Bowen
IBM Corporation
Phone: 919-543-9851
rkb@ralvm11.vnet.ibm.com

Michael Bowman
Netlink, Inc.
Phone: 919-878-8612
Fax: 919-872-2132
meb@netlink.com

Deborah Boyer
Advanced Network and Services, Inc.
Phone: 313-763-3448
Fax: 313-764-6429
djb@noc.ans.net

Robert Braden
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
braden@isi.edu

Ed Brencovich
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
edb@dss.com

Robert Brenner
GTE Mobile Communications
Phone: 404-391-8450
Fax: 404-395-8537
Robert.W.Brenner@gte.sprint.com

Monroe Bridges
Hewlett-Packard
Phone: 408-447-3957
Fax: 408-447-3660
monroe@cup.hp.com

David Bridgham
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
dab@epilogue.com

Scott Brim
Cornell University
Phone: 607-255-9392
Fax: 607-255-9086
Scott_Brim@cornell.edu

Thomas Brisco
Rutgers University
Phone: 908-932-5020
Fax: 908-932-2968
brisco@pilot.njin.net

Al Broscius
Bell Communications Research
Phone: 201-829-4781
Fax: 201-829-2504
broscius@bellcore.com

Caralyn Brown
Wellfleet Communications, Inc.
Phone: 617-280-2335
Fax: 617-280-2300
cbrown@wellfleet.com

Jack Brown
United States Army
Computer Engineer Center
Phone: 602-533-2834
Fax: 602-533-5676
jbrown@huachuca-emh8.army.mil

Gregory Bruell
Wellfleet Communications, Inc.
Phone: 617-250-2339
Fax: 617-275-5001
gob@wellfleet.com

Theodore Brunner
Bell Communications Research
Phone: 201-829-4578
Fax: 201-829-2504
tob@thumper.bellcore.com

Sandy Bryant
University of Virginia
Phone: 804-982-4708
Fax: 804-982-4715
slb@virginia.edu

Jeffrey Burgan
NASA Ames Research Center
Phone: 415-604-5705
Fax: 415-604-0063
jeff@nsipo.nasa.gov

Randy Butler
National Center for
Supercomputing Applications
Phone: 217-333-1189
Fax: 217-244-1987
rbutler@ncsa.uiuc.edu

Ramon Caceres
Matsushita Information
Technology Laboratory
Phone: 609-497-4600
Fax: 609-497-4013
ramon@mitl.com

Susan Calcari
CERFnet
Phone: 619-455-3900
Fax: 619-455-3990
calcaris@cerf.net

Robert Calderon
Advanced Network and
Services, Inc.
Phone: 313-764-3575
Fax: 313-764-6429
calderon@noc.ans.net

Ross Callon
Wellfleet Communications, Inc.
Phone: 617-280-2436
Fax: 617-280-2300
rcallon@wellfleet.com

John Campbell
Naval Command, Control and
Ocean Surveillance Center
Phone: 619-553-4103
Fax: 619-553-6288
jrcamp@nosc.mil

Larry Campbell
The Ohio State University
Phone: 614-851-0223
Fax: 614-292-4512
larry@psycho.psy.ohio-state.edu

C. Allan Cargille
University of Wisconsin
Phone: 608-262-5084
Fax: 608-262-9777
allan.cargille@cs.wisc.edu

Jeff Carman
Bell Northern Research
Phone: 613-763-3921
Fax: 613-763-9496
tcarman@bnr.ca

Kevin Carosso
Innosoft International, Inc.
Phone: 909-624-7907
Fax: 909-621-5319
kvc@innosoft.com

David Carr
University of New York at
Buffalo
Phone: 716-645-6265
Fax: 716-645-3081
Carr@acsu.buffalo.edu

Corinne Carroll
Bolt Beranek and Newman
Phone: 617-873-3087
Fax: 617-873-5620
ccarroll@bbn.com

Douglas Carson
University of Toronto
Phone: 416-978-3855
Fax: 416-978-6620
carson@utcc.utoronto.ca

Jeff Case
University of Tennessee
Phone: 615-573-1434
Fax: 615-573-9197
case@cs.utk.edu

Stephen Casner
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
casner@isi.edu

James Cassell
DLA Systems Automation Center
Phone: 614-692-9971
Fax: 614-692-9731
jcassell@dsac.dla.mil

Vinton Cerf
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
vcerf@cnri.reston.va.us

George Chang
Bell Communications Research
Phone: 908-699-3879
Fax: 908-336-2969
gkc@ctt.bellcore.com

John Chang
US West Advanced Technologies
Phone: 303-541-6255
Fax: 303-541-6250
jrc@uswest.com

John Chang
IBM Corporation
Phone: 919-254-4393
Fax: 919-254-5483
changj@ralvm6.vnet.ibm.com

Yee-Hsiang Chang
Hewlett-Packard
Phone: 415-857-7398
yhc@hpl.hp.com

A. Lyman Chapin
Bolt Beranek and Newman
Phone: 617-873-3133
Fax: 617-873-4086
lyman@bbn.com

Enke Chen
Merit Network, Inc.
Phone: 313-936-3335
enke@merit.edu

Jia-bing Cheng
IBM Corporation
Phone: 919-254-4434
Fax: 919-254-5483
cheng@ralvm6.vnet.ibm.com

William Chimiak
The Bowman Gray School of Medicine
Phone: 919-716-2815
chim@relito.medeng.wfu.edu

Anthony Chow
Ascom Timeplex, Inc.
Phone: 301-443-4032
Fax: 301-443-7978
chow_a@wwtc.timeplex.com

Cyrus Chow
NASA Ames Research Center
Phone: 415-604-6843
Fax: 415-604-7300
cchow@ames.arc.nasa.gov

Thomas Christie
Advanced Network and Services, Inc.
Phone: 914-789-5300
tfc@aa.ans.net

Jodi-Ann Chu
University of Hawaii
Phone: 808-956-2400
Fax: 808-956-2412
jodi@uhunix.uhcc.hawaii.edu

William Chung
IBM Corporation
Phone: 914-945-3934
Fax: 914-945-2141
whchung@watson.ibm.com

George Clapp
Ameritech Advanced Data Services
Phone: 708-248-6507
Fax: 708-248-6038
clapp@ameris.center.il.ameritech.com

Cynthia Clark
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
cclark@cnri.reston.va.us

David Clark
Massachusetts Institute of Technology
Phone: 617-253-6002
Fax: 617-253-2673
ddc@lcs.mit.edu

Henry Clark
OARnet
Phone: 614-292-6483
Fax: 614-292-7168
henryc@oar.net

Wayne Clark
cisco Systems
Phone: 415-688-4627
Fax: 415-688-8282
wclark@cisco.com

Brian Coan
Bell Communications Research
Phone: 201-829-4489
Fax: 201-829-5889
coan@faline.bellcore.com

Richard Cogger
Cornell University
Phone: 607-255-7566
Fax: 607-255-5771
R.Cogger@cornell.edu

Robert Cole
AT&T Bell Laboratories
Phone: 908-949-1950
Fax: 908-949-1726
rgc@qsun.att.com

Richard Colella
National Institute of Standards and Technology
Phone: 301-975-3627
Fax: 301-590-0932
colella@nist.gov

Robert Collet
Sprint Government Systems Division
Phone: 703-904-2230
Fax: 703-904-2251
rcollet@icm1.icp.net

Michael Collins
Lawrence Livermore National Laboratory
Phone: 510-422-4018
Fax: 510-422-1482
collinsms@es.net

Rob Coltun
Consultant
Phone: 301-340-9416
rcoltun@ni.umd.edu

David Conklin
JvNCnet
Phone: 609-258-2414
Fax: 609-258-2424
conklin@jvnc.net

James Conklin
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
jbc@bitnic.educom.edu

David Conrad
Internet Initiative Japan
Phone: +81 3 3580 3781
Fax: +81 3 3580 3782
davidc@iij.ad.jp

Robert Cooney
Naval Computer and Telecommunications Station
Phone: 202-685-1050
Fax: 202-433-3188
cooney@wnyose.nctsw.navy.mil

Simon Coppins
University of Adelaide
Phone: +61 8 228 5836
Fax: +61 8 223 7239
coppins@arch.adelaide.edu.au

Al Costanzo
AKC Consulting
Phone: 908-298-9000
al@akc.com

Naomi Courter
MCNC - CONCERT Network
Phone: 919-248-1807
Fax: 919-248-1405
naomi@concert.net

Tracy Cox
Bell Communications Research
Phone: 908-758-2107
Fax: 908-758-4177
tacox@sabre.bellcore.com

Stephen Coya
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
scoya@cnri.reston.va.us

David Crocker
Silicon Graphics, Inc.
Phone: 415-390-1804
Fax: 415-962-8404
dcrocker@mordor.stanford.edu

Stephen Crocker
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
crocker@tis.com

Wayne Cullen
Netlink, Inc.
Phone: 919-878-8612
Fax: 919-872-2132
wnc@netlink.com

Dave Cullerot
Cabletron Systems, Inc.
Phone: 603-337-1409
Fax: 603-332-1739
cullerot@ctron.com

John Curran
NEARnet
Phone: 617-873-4398
Fax: 617-873-5620
jcurran@nic.near.net

Brent Curtiss
The Ohio State University
Phone: 614-292-4843
Fax: 614-292-7081
bcurtiss@magnuss.ocs.ohio-state.edu

Megan Davies
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
mdavies@cnri.reston.va.us

James Davin
Bell Communications Research
Phone: 201-829-4798
Fax: 201-829-2504
davin@thumper.bellcore.com

Mark Davis-Craig
Merit Network, Inc.
Phone: 313-936-3000
mad@merit.edu

Shane Dawalt
Wright State University
Phone: 513-873-4037
sdawalt@desire.wright.edu

Steve Deering
Xerox Corporation
Phone: 415-812-4839
Fax: 415-812-4471
deering@parc.xerox.com

Steve DeJarnett
IBM Corporation
Phone: 415-694-3896
Fax: 415-694-3073
steve@ibmpa.awdpa.ibm.com

Tony DeSimone
AT&T Bell Laboratories
Phone: 908-949-4521
Fax: 908-834-5906
tds@hoserve.att.com

Osmund DeSouza
AT&T Bell Laboratories
Phone: 908-949-1393
Fax: 908-949-4673
osmund.desouza@att.com

Peter Deutsch
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8611
peterd@bunyip.com

Taso Devetzis
Bell Communications Research
Phone: 201-829-4764
Fax: 201-829-2504
devetzis@bellcore.com

Thomas DeWitt
General Services Administration
Phone: 202-501-0139
Fax: 202-208-5555
tdewitt@osi.ncsl.nist.gov

Manuel Diaz
DAVID Systems, Inc.
Phone: 408-541-6865
Fax: 408-541-6980
diaz@davidsys.com

Chas DiFatta
Software Engineering Institute
Phone: 412-268-7788
Fax: 412-268-5758
chas@cmu.edu

M.J. Dixon
AT&T Bell Laboratories
Phone: 908-582-5460
Fax: 908-582-5857
mjd@att.com

Kurt Dobbins
Cabletron Systems, Inc.
Phone: 603-337-5002
Fax: 603-424-8121
kurtdob@ctron.com

Dale Dougherty
O'Reilly and Associates, Inc.
Phone: 707-829-0515
Fax: 707-829-0104
dale@ora.com

Bob Downs
Combinet
Phone: 408-522-9020
Fax: 408-732-5479
bdowns@combinet.com

Roger Draper
Liebert Corporation
Phone: 714-457-3638
Fax: 714-457-3633

Ralph Droms
Bucknell University
Phone: 717-524-1145
Fax: 717-524-1760
droms@bucknell.edu

David Dubois
Pacer Software, Inc.
Phone: 508-898-3300
Fax: 508-366-1356
dad@pacersoft.com

Kishan Dudkikar
SKD Systems
Phone: 301-984-7657
Fax: 301-816-2982
kishan@icm1.icp.net

Pierre Dupont
Motorola, Inc.
Phone: 604-241-6243
Fax: 604-241-6042
dupont@mdd.comm.mot.com

Sandra Durham
SynOptics Communications, Inc.
Phone: 408-764-1285
Fax: 408-988-5525
sdurham@synoptics.com

Steve Dusse
RSA Data Security, Inc.
Phone: 415-595-8782
Fax: 415-595-1873
spock@rsa.com

Tom Easterday
CICNet, Inc.
Phone: 313-998-6285
Fax: 313-998-6105
tom@cic.net

Avi Elenko
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
avi@dss.com

Ed Ellesson
IBM Corporation
Phone: 919-254-4115
Fax: 919-254-5410
ellesson@vnet.ibm.com

William Elliot
ONEAC Corporation
Phone: 708-816-6000
Fax: 708-680-5124
oneacwre@aol.com

Chip Elliott
Bolt Beranek and Newman
Phone: 617-873-2615
celliot@bbn.com

Alan Emtage
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8134
bajan@bunyip.com

David Engel
Optical Data Systems, Inc.
Phone: 214-234-6400
Fax: 214-234-4059
david@ods.com

Robert Enger
Advanced Network and Services, Inc.
Phone: 703-758-7722
Fax: 703-758-7717
enger@reston.ans.net

Urs Eppenberger
SWITCH
Phone: +41 1 261 8112
Fax: +41 1 261 8133
eppenberger@switch.ch

Phillip Epps
Technipower
Phone: 203-748-7001
Fax: 203-797-9285
76410.327@compuserve.com

Hans Eriksson
Swedish Institute of Computer Science
Phone: +46 8 752 1527
Fax: +46 8 751 7230
hans@sics.se

Deborah Estrin
University of Southern California
Phone: 310-822-1511
Fax: 310-823-6714
estrin@isi.edu

Erik Fair
Apple Computer, Inc.
Phone: 408-996-1010
Fax: 408-974-3103
fair@apple.com

Roger Fajman
National Institutes of Health
Phone: 301-402-4768]
Fax: 301-480-6245
raf@cu.nih.gov

Dino Farinacci
cisco Systems
Phone: 415-688-4696
Fax: 415-688-4575
dino@cisco.com

Stefan Fassbender
IBM Corporation
Phone: +49 6221 594 446
Fax: +49 6221 593 300
stf@easi.net

Mark Fedor
Performance Systems International, Inc.
Phone: 518-283-8860
Fax: 518-283-8904
fedor@psi.com

Jonathan Fellows
Grumman Data Systems
Phone: 619-224-3500
Fax: 619-224-4937
jonf@gdstech.grumman.com

Sallie Fellows
National Forum on Education Statistics
Phone: 603-271-2778
Fax: 603-271-1953
sallie@ed.unh.edu

Dennis Ferguson
Advanced Network and Services, Inc.
Phone: 914-789-5367
Fax: 914-789-5310
dennis@ans.net

Antonio Fernandez
Bell Communications Research
Phone: 201-829-4585
Fax: 201-829-5963
afa@thumper.bellcore.com

Michael Fidler
The MITRE Corporation
Phone: 703-883-6859
Fax: 703-883-6436
fidler@mitre.org

Dale Finkelson
MIDnet
Phone: 402-472-5032
Fax: 402-472-8486
dmf@westie.mid.net

Eric Fleischman
Boeing Computer Services
Phone: 206-957-5334
Fax: 206-957-5050
ericf@act.boeing.com

Francois Fluckiger
CERN
Phone: +41 22 767 4991
Fax: +41 22 767 7155
fluckiger@vxcern.cern.ch

Wayne Foco
IBM Corporation
Phone: 919-254-4422
Fax: 919-254-5483
foco@ralvm6.vnet.ibm.com

Peter Ford
Los Alamos National Laboratory
Phone: 505-665-0058
Fax: 505-665-4939
peter@goshawk.lanl.gov

Jill Foster
University of Newcastle upon Tyne
Phone: +44 91 222 8250
Fax: +44 91 222 8765
Jill.Foster@newcastle.ac.uk

Steven Foster
UCCSN
University of Nevada
Phone: 702-784-6557
Fax: 702-784-1108
foster@cs.unr.edu

Craig Fox
Network Systems Corporation
Phone: 612-424-4888
foxcj@network.com

Karl Fox
Morning Star Technologies
Phone: 614-451-1883
Fax: 614-459-5054
karl@morningstar.com

Paul Franchois
National Institute of Standards and Technology
Phone: 303-497-5674
Fax: 303-497-3012
paulf@bldrdoc.gov

Paul Francis
Bell Communications Research
Phone: 201-829-4484
Fax: 201-829-2504
Francis@thumper.bellcore.com

Barbara Fraser
CERT Coordination Center
Phone: 412-268-5010
Fax: 412-268-6989
byf@cert.org

Christine Fredenburg
DLA Systems Automation Center
Phone: 614-692-9970
Fax: 614-692-9731
cfredenburg@dsac.dla.mil

Ron Frederick
Xerox Corporation
Phone: 415-812-4459
Fax: 415-812-4471
frederick@parc.xerox.com

Ned Freed
Innosoft International, Inc.
Phone: 714-624-7907
Fax: 714-621-5319
ned@innosoft.com

Anita Freeman
Pacific Bell
Phone: 510-901-9294
Fax: 510-277-0176

Raphael Freiwirth
RCI, Inc.
Phone: 508-453-7857
Fax: 508-934-9066
5242391@mcimail.com

Jerry Friesen
Sandia National Laboratories
Phone: 510-294-3144
Fax: 510-294-1225
jafries@sandia.llnl.gov

Karen Frisa
Carnegie Mellon University
Phone: 412-268-8576
Fax: 412-268-4987
karen.frisa@andrew.cmu.edu

Dan Frommer
Digital Equipment Corporation
Phone: +972 2 782 551
Fax: +972 2 782 373
dan@jeremy.enet.dec.com

Marcello Frutig
Rede Nacional de Pesquisa
RNP/IMPA
Phone: +55 21 259 77 96
Fax: +55 21 259 77 96
frutig@rnp.impa.br

Vince Fuller
BARRNet
Phone: 415-723-6860
Fax: 415-723-0010
vaf@stanford.edu

Jim Fullton
MCNC - CNIDR
Phone: 919-248-9247
Fax: 919-248-1405
Jim.Fullton@cnidr.org

Gwen Funchess
The Ohio State University
Phone: 614-292-4843
Fax: 614-292-7081
funchess@magnus.acs.ohio-state.edu

Peter Furniss
Peter Furniss Consultants
Phone: +44 81 313 1833
Fax: +44 81 313 1833
p.furniss@ulcc.ac.uk

Maria Gallagher
NASA Science Internet
Phone: 415-604-3601
Fax: 415-604-0063
maria@nsipo.nasa.gov

James Galvin
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
galvin@tis.com

Kevin Gamiel
MCNC Center for Communications
Phone: 919-248-7679
Fax: 919-248-1405
kevin.gamiel@concert.net

Joan Gargano
University of California, Davis
Phone: 916-752-2591
Fax: 916-752-9017
jcgargano@ucdavis.edu

Christine Garland
AT&T Bell Laboratories
Phone: 708-979-5430
Fax: 708-979-8822
garland@ihspa.att.com

John Gawf
Compatible Systems Corporation
Phone: 303-444-9532
Fax: 303-444-9595
gawf@compatible.com

Eugene Geer
Bell Communications Research
Phone: 908-758-2149
Fax: 908-758-4184
ewg@cc.bellcore.com

Jisoo Geiter
The MITRE Corporation
Phone: 703-883-6571
Fax: 703-883-5200
geiter@mitre.org

Greg Gicale
DAS/Division of Computer Services, State of Ohio
Phone: 614-466-6516
Fax: 614-752-6108
gicale@ohio.gov

Robert Gilligan
Sun Microsystems, Inc.
Phone: 415-336-1012
Fax: 415-336-6015
gilligan@Eng.Sun.Com

Kenneth Giusti
Chipcom Corporation
Phone: 508-460-8922
Fax: 508-460-8952
kgiusti.chipcom.com

Joseph Godsil
National Center for Supercomputing Applications
Phone: 217-333-1989
Fax: 217-244-1987
jgodsil@ncsa.uiuc.edu

Mike Goguen
SynOptics Communications, Inc.
Phone: 408-764-1289
Fax: 408-764-1299
goguen@synoptics.com

Fengmin Gong
MCNC
Phone: 919-248-9214
Fax: 919-248-1405
gong@concert.net

Kenneth Goodwin
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
goodwin@a.psc.edu

Ramesh Govindan
Bell Communications Research
Phone: 201-829-4406
rxg@thumper.bellcore.com

Richard Graveman
Bell Communications Research
Phone: 908-699-4611
Fax: 908-336-2943
rfg@ctt.bellcore.com

Cleve Graves
OpenConnect Systems, Inc.
Phone: 214-484-5200
Fax: 214-888-0688
cvgpc@oc.com

Terry Gray
University of Washington
Phone: 206-543-3880
Fax: 206-543-3909
gray@cac.washington.edu

Mark Green
Apple Computer, Inc.
Phone: 510-204-0934
Fax: 510-649-4008
markg@apple.com

Darren Griffiths
Fujitsu Open Systems Solutions
Phone: 510-652-6200
Fax: 510-652-5532
dag@ossi.com

Phillip Gross
Advanced Network and Services, Inc.
Phone: 914-789-5335
Fax: 914-789-5310
pgross@ans.net

Daniel Grossman
Motorola Codex
Phone: 617-821-7333
Fax: 617-821-4218
dan@merlin.dev.cdx.mot.com

Chris Gunner
Digital Equipment Corporation
Phone: 508-486-7792
Fax: 508-486-5279
gunner@dsmail.enet.dec.com

Robert Gutierrez
NASA Science Internet
Phone: 800-424-9920
Fax: 415-604-0063
gutierre@nsipo.nasa.gov

Robert Hagens
Advanced Network and Services, Inc.
Phone: 703-758-7700
Fax: 703-758-7717
hagens@ans.net

Tony Hain
Lawrence Livermore National Laboratory
Phone: 510-422-4200
Fax: 510-422-1482
alh@es.net

Martyne Hallgren
Cornell University
Phone: 607-254-8767
Fax: 607-255-5771
martyne@nr-tech.cit.cornell.edu

Joel Halpern
Network Systems Corporation
Phone: 612-424-1606
Fax: 612-424-1736
jmh@network.com

Deborah Hamilton
AT&T Bell Laboratories
Phone: 908-949-9459
Fax: 908-949-1726
debbie@qsun.att.com

Patrick Hanel
Nokia Telecommunications
Phone: +358 0 51044927
Fax: +358 0 51044952
hanel@yoyodyne.trs.ntc.nokia.com

Daniel Hansen
Network General Corporation
Phone: 201-616-0786
Fax: 201-616-8311
dan@ngc.com

Danny Hanson
AFCC/TNACT
Phone: 618-256-3720
Fax: 618-256-8952
hanson.tic@commlan.safb.af.mil

Raymond Harder
Consultant
Phone: 909-983-4713
rharder@eis.calstate.edu

Susan Hares
Merit Network, Inc.
Phone: 313-936-2095
Fax: 313-747-3745
skh@merit.edu

Richard Harris
Boeing Computer Services
Phone: 206-865-4922
Fax: 206-865-6903
rharris@atc.boeing.com

Susan Harris
Merit Network, Inc.
Phone: 313-936-3000
srh@umich.edu

John Hascall
Iowa State University
Phone: 515-294-9551
Fax: 515-294-1717
john@iastate.edu

Eugene Hastings
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
hastings@psc.edu

Alisa Hata
University of Washington
Phone: 206-543-7732
Fax: 206-543-3909
hata@cac.washington.edu

Ken Hayward
Bell Northern Research
Phone: 613-763-4042
Fax: 613-763-2626
Ken.Hayward@bnr.ca

Frank Heath
Rockwell International Corporation
Phone: 714-833-6864
Fax: 714-833-6863
heath@cmc.com

Roland Hedberg
Technische Universiteit Delft
Phone: +33 15 785210
Fax: +33 15 783787
Roland.Hedberg@rc.tudelft.nl

John Heimaster
The Ohio State University
Phone: 614-292-1435
Fax: 614-292-8261
jwh@mps.ohio-state.edu

Juha Heinanen
Telecom Finland
Phone: +358 49 500958
Fax: +358 31 2432211
juha.heinanen@datanet.tele.fi

Marco Hernandez
CREN
Phone: 202-872-4200
Fax: 202-872-4318
marco@mh-slip.cren.edu

Ittai Hershman
Advanced Network and Services, Inc.
Phone: 914-789-5337
Fax: 914-789-5310
ittai@ans.net

Scott Hertzog
DCA, Inc.
Phone: 512-244-3871
Fax: 512-388-7090

Jeff Hilgeman
Apertus Technologies, Inc.
Phone: 612-828-0668
Fax: 612-828-0773
jeffh@apertus.com

Robert Hinden
Sun Microsystems, Inc.
Phone: 415-336-2082
Fax: 415-336-6015
hinden@eng.sun.com

Russ Hobby
University of California, Davis
Phone: 916-752-0236
Fax: 916-752-9154
rdhobby@ucdavis.edu

Don Hofacker
Network Solutions, Inc.
Phone: 513-427-3737
Fax: 513-427-8482
hofacker@dtedi.hq.aelc.af.mil

Don Hoffman
Sun Microsystems, Inc.
Phone: 415-336-4339
Fax: 415-965-4903
hoffman@eng.sun.com

Ellen Hoffman
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
ellen@merit.edu

Frank Hoffmann
IBM Deutschland Information System GmbH
Phone: +49 6221 594427
Fax: +49 6221 593300
hoffmann@dhdibm1.bitnet

Gerd Holzhauer
Apple Computer, Inc.
Phone: +33 1 4901 4704
Fax: +33 1 4090 0951
holzhauer1@applelink.apple.com

Jeffrey Honig
Cornell University
Phone: 607-255-6460
Fax: 607-255-9086
Jeffrey_C_Honig@Cornell.edu

Alton Hoover
Advanced Network and Services, Inc.
Phone: 914-789-5341
Fax: 914-789-5310
hoover@ans.net

John Hopprich
DAVID Systems, Inc.
Phone: 408-541-7459
Fax: 408-541-6980
hopprich@davidsys.com

Susan Horvath
Merit Network, Inc.
Phone: 313-936-3000
shorvath@merit.edu

Jack Houldsworth
ICL
Phone: +44 438 738 941
Fax: +44 438 738 904
J.Houldsworth@ste0906.wins.icl.co.uk

Jeroen Houttuin
RARE
Phone: +31 20 6391131
Fax: +31 20 6393289
houttuin@rare.nl

Tim Howes
University of Michigan
Phone: 313-747-4454
Fax: 313-764-5140
tim@umich.edu

Woody Huang
Merit Network, Inc.
Phone: 313-936-3335
yuh@merit.edu

Kathy Huber
Bolt Beranek and Newman
Phone: 617-873-2520
Fax: 617-873-4086
khuber@bbn.com

Richard Huber
AT&T Bell Laboratories
Phone: 908-949-4524
Fax: 908-949-1726
rvh@qsun.att.com

Steven Hubert
University of Washington
Phone: 206-543-6384
Fax: 206-543-3909
hubert@cac.washington.edu

Jeff Hughes
Hewlett-Packard
Phone: 719-531-4777
Fax: 719-531-4505
jeff@col.hp.com

Christian Huitema
INRIA
Phone: +33 93 65 77 15
Fax: +33 93 65 77 66
christian.huitema@sophia.inria.fr

Erik Huizer
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
Erik.Huizer@SURFnet.nl

Inna Il'yasova
CNIDR
Phone: 919-299-8508
ilyasova@meg.uncg.edu

John Ioannidis
Columbia University
Phone: 212-939-7038
Fax: 212-666-0140
ji@cs.columbia.edu

Phil Irey
NSWC
Phone: 703-663-1582
Fax: 703-663-8673
pirey@relay.nswc.navy.mil

Keisuke Ito
Network Express
Phone: 313-761-5051
Fax: 313-995-1114
itoh@msen.com

Britt Jackman
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
bjackman@cnri.reston.va.us

Lenore Jackson
NASA Science Internet
Phone: 415-604-0455
Fax: 415-604-0063
jackson@nsipo.arc.nasa.gov

David Jacobson
IBM Corporation
Phone: 919-254-5590
Fax: 919-254-9515
dnjake@vnet.ibm.com

Ronald Jacoby
Silicon Graphics, Inc.
Phone: 415-390-1371
Fax: 415-969-2314
rj@sgi.com

Barbara Jennings
Sandia National Laboratories
Phone: 505-845-8554
Fax: 505-844-2067
bjjenni@sandia.gov

Thomas Johannsen
Dresden University of Technology
Phone: +49 351 463 4621
Fax: +49 351 463 7163
thomas@ebzaw1.et.tu-dresden.de

Dale Johnson
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
dsj@merit.edu

David Johnson
Carnegie Mellon University
Phone: 412-268-7399
Fax: 412-681-5739
dbj@cs.cmu.edu

Laurent Joncheray
Merit Network, Inc.
Phone: 313-936-3335
lpj@merit.edu

Matthew Jonson
United States Air Force
Phone: 205-416-4075
Fax: 205-416-4738
jonson@server.af.mil

Claudia Jordan
Ohio State University
Phone: 614-292-4843
Fax: 614-292-7028
jordan.76@osu.edu

Kevin Jordan
Control Data Systems, Inc.
Phone: 612-482-6835
Fax: 612-482-4455
Kevin.E.Jordan@cdc.com

Dan Jordt
NorthWestNet
Phone: 206-562-3000
Fax: 206-562-4822
danj@nwnet.net

Erik Jul
OCLC Online Computer Library Center
Phone: 614-764-4364
Fax: 614-764-2344
jul@oclc.org

Merike Kaeo
National Institutes of Health
Phone: 301-402-1548
Fax: 301-402-2867
merike@alw.nih.gov

Marko Kaittola
FUNET
Phone: +358 31 162 933
Fax: +358 31 162 913
Marko.Kaittola@funet.fi

Scott Kaplan
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
scott@wco.ftp.com

Doug Karl
The Ohio State University
Phone: 614-292-4843
Fax: 614-292-7081
dkarl@osu.edu

Phil Karn
Qualcomm, Inc.
Phone: 619-597-5501
Fax: 619-455-0571
karn@qualcomm.com

Daniel Karrenberg
RIPE Network Coordination Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
daniel@ripe.net

Frank Kastenholz
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
kasten@ftp.com

Neil Katin
SunSoft
Phone: 415-336-3618
Fax: 415-336-6776
katin@eng.sun.com

David Katinsky
Rutgers University
Phone: 908-932-3380
Fax: 908-932-2968
dmk@pilot.njin.net

Charles Kaufman
Digital Equipment Corporation
Phone: 603-881-1495
kaufman@zk3.dec.com

Doug Kay
SPRINT
Phone: 703-689-6514
Fax: 703-689-5177
doub.kay@sprintintl.sprint.com

Bill Kelly
Auburn University
Phone: 205-844-4512
Fax: 205-844-9680
kellywh@mail.auburn.edu

Stephen Kent
Bolt Beranek and Newman
Phone: 617-873-3988
Fax: 617-873-4086
kent@bbn.com

Mark Kepke
Hewlett-Packard
Phone: 303-229-4201
Fax: 303-229-3526
mak@fc.hp.com

Kenneth Key
SNMP Research, Inc.
Phone: 615-573-1434
Fax: 615-573-9197
key@cs.utk.edu

Michael Khalandovsky
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
mlk@ftp.com

Zbigniew Kielczewski
Eicon Technology Corporation
Phone: 514-631-2592
Fax: 514-631-3092
zbig@eicon.qc.ca

Steve Kille
ISODE Consortium
Phone: +44 71 223 4062
Fax: +44 71 223 3846
S.Kille@isode.com

Glenn Kime
Network Express
Phone: 313-761-5051
Fax: 313-995-1114
itoh@msen.com

Peter Kirstein
University College London
Phone: +44 71 380 7286
Fax: +44 71 387 1397
P.Kirstein@cs.ucl.ac.uk

John Klensin
INFOODS Secretariat
Phone: 617-864-3247
Fax: 617-491-6266
klensin@infoods.unu.edu

Charley Kline
University of Illinois
Phone: 217-333-3339
Fax: 217-244-7089
cvk@uiuc.edu

Mark Knopper
Merit Network, Inc.
Phone: 313-763-6061
Fax: 313-747-3745
mak@merit.edu

Jim Knowles
Sterling Software
Federal Systems
Phone: 415-604-3056
Fax: 415-604-7300
jknowles@binky.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
stev@ftp.com

Andrew Knutsen
The Santa Cruz Operation, Inc.
Phone: 408-427-7538
Fax: 408-458-0811
andrewk@sco.com

Moshe Kochinski
Fibronics, Ltd.
Phone: +972 4 313619
Fax: +972 4 313342
moshek@FibHaifa.com

Mark Kosters
Network Solutions, Inc.
Phone: 703-742-4795
Fax: 703-742-4811
markk@internic.net

Deirdre Kostick
Bell Communications Research
Phone: 908-758-2642
Fax: 908-758-4177
dck2@sabre.bellcore.com

John Krawczyk
Wellfleet Communications, Inc.
Phone: 617-280-2311
Fax: 617-280-2300
jkrawczy@wellfleet.com

Lakshman Krishnamurthy
University of Kentucky
Phone: 606-257-8737
Fax: 606-358-1029
lakashman@ms.uky.edu

Padma Krishnaswamy
Bell Communications Research
Phone: 908-758-2362
Fax: 908-758-4177
kri@sabre.bellcore.com

Edward Krol
University of Illinois
Urbana
Phone: 217-333-7886
Fax: 217-244-7089
e-krol@uiuc.edu

Duane Kuang
Kalpana, Inc.
Phone: 408-749-1900
Fax: 408-988-1361
duanek@kalpana.com

John Kunze
University of California,
Berkeley
Phone: 510-642-1530
Fax: 510-643-5385
jak@violet.berkeley.edu

Giri Kuthethoor
University of Kentucky
Phone: 606-257-8781
Fax: 606-258-1928
giri@ms.uky.edu

William Kwan
Rabbit Software Corporation
Phone: 215-993-1135
Fax: 215-640-1379
kwan@rabbit.com

Mary La Roche
Corporation for Open Systems International
Phone: 703-205-2741
Fax: 703-846-8590
maryl@cos.com

John Labbe
Advanced Network and Services, Inc.
Phone: 313-936-1540
Fax: 313-764-6429
jcl@ans.net

Paul Lambert
Motorola, Inc.
Phone: 602-441-3646
Fax: 602-441-8377
paul_lambert@email.mot.com

Ruth Lang
SRI International
Phone: 415-859-5608
Fax: 415-859-6028
rlang@nisc.sri.com

Sylvain Langlois
Electricite de France
Phone: +33 1 47 65 44 02
Fax: +33 1 47 65 35 23
Sylvain.Langlois@exp.edf.fr

Ronald Lanning
Ohio University Communication Network Services
Phone: 614-593-0451
lanning@netltm.cats.ohiou.edu

Mark Laubach
Hewlett-Packard
Phone: 415-857-3513
Fax: 415-857-8526
laubach@hpl.hp.com

Debra Legare
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
dlegare@cnri.reston.va.us

Barry Leiner
Universities Space Research Association
Phone: 415-390-0317
Fax: 415-390-0318
leiner@nsipo.nasa.gov

Patrick Leung
Eicon Technology Corporation
Phone: 514-631-2592
Fax: 514-631-3092
patrickl@eicon.qc.ca

Mark Lewis
Telebit Corporation
Phone: 408-745-3232
Fax: 408-745-3810
Mark.S.Lewis@telebit.com

Tony Li
cisco Systems
Phone: 415-688-8186
Fax: 415-688-4575
tli@cisco.com

Fong-Ching Liaw
Sun Microsystems, Inc.
Phone: 415-336-5415
Fax: 415-336-3156
fong@eng.sun.com

Hock-Koon Lim
Case Western Reserve University
Phone: 216-368-2982
Fax: 216-368-4693
lim@po.cwru.edu

John Linn
Geer Zolot Associates
Phone: 617-374-3700
linn@gza.com

Robin Littlefield
Wellfleet Communications, Inc.
Phone: 617-280-2312
Fax: 617-280-2300
rlittlef@wellfleet.com

Frank Liu
Pacific Bell
Phone: 510-867-6167
Fax: 510-867-1405
fcliu@pacbell.com

Daniel Long
Bolt Beranek and Newman
Phone: 617-873-2766
Fax: 617-873-5620
long@nic.near.net

Kim Long
SURAnet
Phone: 301-982-3214
Fax: 301-482-4605
klong@sura.net

Peter Lothberg
EBONE
Phone: +46 8 6699720
Fax: +46 8 102510
roll@stupi.se

David Lovering
National Institute of Standards and Technology
Phone: 303-497-5662
Fax: 303-497-3012
lovering@bldrdoc.gov

Philip Lozowick
Digital Equipment Corporation
Phone: +972 2782 551
Fax: +972 2782 373
lozowick@isv.dec.com

Yu-Lin Lu
Hewlett-Packard
Phone: 408-447-2470
Fax: 408-447-3660
yulin@hpinddu.cup.hp.com

Steven Lunt
Bell Communications Research
Phone: 908-699-4244
Fax: 908-336-2943
lunt@bellcore.com

Paul Lustgraaf
Iowa State University
Phone: 515-294-0324
grpjl@iastate.edu

Clifford Lynch
University of California, Berkeley
Phone: 510-987-0522
Fax: 510-839-3573
calur@uccmvsa.ucop.edu

Charles Lynn
Bolt Beranek and Newman
Phone: 617-873-3367
clynn@bbn.com

Bruce Mackey
Xerox Corporation
Phone: 716-422-9165
brucem@cinops.xerox.com

Glenn Mackintosh
University of Toronto
Phone: 416-978-1792
Fax: 416-978-6620
glenn@canet.ca

Carl Madison
Star-Tek, Inc.
Phone: 508-393-9393
Fax: 508-365-6934
carl@startek.com

Jamshid Mahdavi
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
Mahdavi@a.psi.edu

James Mahon
Honeywell Federal Systems
Phone: 703-827-3955
Fax: 703-827-3729
mahonj@honfsi1.att.com

Samir Malak
MCNC Center for Communications
Phone: 919-248-1499
Fax: 919-248-1405
malaks@alice.uncg.edu

Carl Malamud
Internet Talk Radio
Phone: 703-548-1126
carl@malamud.com

Kent Malave
IBM Corporation
Phone: 512-838-3486
Fax: 512-838-3509
kent@bach.austin.ibm.com

William Malicoate
The Ohio State University
Phone: 614-292-1835
Fax: 614-292-7081
wjm+@osu.edu

Andrew Malis
Ascom Timeplex, Inc.
Phone: 508-470-2565
Fax: 508-264-4999
malis_a@timeplex.com

Gary Malkin
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-3159
gmalkin@xylogics.com

Tracy Mallory
3Com Corporation
Phone: 408-764-6518
Fax: 408-764-5002
tracym@3com.com

Allison Mankin
Locus, Inc. at Naval Research Laboratory
Phone: 202-404-7030
Fax: 202-767-1683
mankin@cmf.nrl.navy.mil

Bill Manning
SESQUINET
Phone: 713-285-5415
Fax: 713-527-6099
bmanning@sesqui.net

Glenn Mansfield
AIC Systems Laboratory
Phone: +81 22 279 3310
Fax: +81 22 279 3460
glenn@aic.co.jp

Michael Marcinkevicz
California State University
Phone: 714-773-3009
Fax: 310-985-9400
mdm@csu.net

Janet L. Marcisak
FTP Software, Inc.
Phone: 508-659-6362
Fax: 508-794-4488
jlm@ftp.com

April Marine
NASA NAIC
april@atlas.arc.nasa.gov

Brian Marsh
Matsushita Information Technology Laboratory
Phone: 609-497-4600
Fax: 609-497-4013
marsh@mitl.com

Cynthia Martin
Defense Information Systems Agency
Phone: 703-487-3038
Fax: 703-487-3019
martin@spica.disa.mil

Jerry Martin
The Ohio State University
Phone: 614-292-4843
Fax: 614-242-7081
jerry+@osu.edu

Ignacio Martinez
RedIRIS
Phone: +34 1 4351214
Fax: +34 1 5781773
martinez@rediris.es

Matt Mathis
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
mathis@a.psc.edu

Jun Matsukata
Institute of Space and Astronautical Science
Phone: +81 427 51 3911
Fax: +81 427 59 4252
jm@eng.isas.ac.jp

Keith McCloghrie
Hughes LAN Systems
Phone: 415-966-7934
Fax: 415-966-3738
kzm@hls.com

Evan McGinnis
3Com Corporation
Phone: 408-764-6064
Fax: 408-764-6134
bem@3com.com

William McKenzie
IBM Corporation
Phone: 919-254-5705
mckenzie@ralvma.vnet.ibm.com

P.V. McMahon
ICL
Phone: +44 734 586211
Fax: +44 734 855106
p.v.mcmahon@rea0803.wins.icl.co.uk

Donna McMaster
SynOptics Communications, Inc.
Phone: 408-764-1206
Fax: 408-764-1299
mcmaster@synoptics.com

Daniel McRobb
Advanced Network and Services, Inc.
Phone: 313-763-3448
Fax: 313-764-6429
dwm@noc.ans.net

Michael Mealling
Georgia Institute of Technology
Phone: 404-894-4669
Fax: 404-894-6157
michael@fantasy.gatech.edu

Milo Medin
NASA Ames Research Center
Phone: 415-604-6440
Fax: 415-604-0063
medin@nsipo.nasa.gov

Marjo Mercado
Hewlett-Packard
Phone: 408-447-2826
Fax: 408-447-3660
marjo@cup.hp.com

Donald Merritt
Army Research Laboratory
Phone: 410-278-6808
Fax: 410-278-5075
don@arl.army.mil

David Meyer
University of Oregon
Computing Center
Phone: 503-346-1747
Fax: 503-346-4397
meyer@ns.uoregon.edu

Gerry Meyer
Spider Systems Limited
Phone: +44 31 554 9424
Fax: +44 31 554 0649
gerry@spider.co.uk

Paul Milazzo
Bolt Beranek and Newman
Phone: 617-873-4325
Fax: 617-873-3776
milazzo@bbn.com

Cynthia Mills
Bolt Beranek and Newman
Phone: 617-873-4143
Fax: 617-873-5620
cmills@bbn.com

James Miner
FiberCom, Inc.
Phone: 703-342-6700
Fax: 703-342-5961
jjm@fibercom.com

Robert Mines
Sandia National Laboratories
Phone: 510-294-2548
rfm@sandia.llnl.gov

Greg Minshall
Novell, Inc.
Phone: 510-975-4507
Fax: 510-947-1238
minshall@wc.novell.com

Mitra
Pandora Systems
Phone: 415-488-0944
Fax: 415-931-7931
mitra@pandora.sf.ca.us

Ray Mlejnek
AT&T Paradyne
Phone: 813-530-8346
Fax: 813-530-8224
r.mlejnek@pdn.paradyne.com

Keith Moore
University of Tennessee
Phone: 615-974-3735
Fax: 615-974-8296
moore@cs.utk.edu

Bob Morgan
Stanford University
Phone: 415-723-9711
Fax: 415-723-0010
morgan@networking.stanford.edu

Dennis Morris
Defense Information Systems
Agency
Phone: 703-487-3329
Fax: 703-487-3190
morrisd@imo-uvax.disa.mil

Matthew Morrisey
C.E.T.A. Corporation
Phone: 513-257-4479
Fax: 513-427-2896
morrisey@wpsp01.hq.aflc.af.mil

Robert Moskowitz
Chrysler Corporation
Phone: 313-758-8212
Fax: 313-758-8173
3858921@mcimail.com

George Mouradian
AT&T Bell Laboratories
Phone: 908-949-7671
Fax: 908-949-8569
gvm@arch3.att.com

John Moy
Proteon, Inc.
Phone: 508-898-2800
Fax: 508-898-3176
jmoy@proteon.com

Satinder Mundra
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-424-8121
mundra@ctron.com

Baktha Muralidharan
Digital Equipment Corporation
Phone: 508-486-7724
Fax: 508-486-7417
murali@smaug.enet.dec.com

Julie Myers
Network Systems Corporation
Phone: 612-424-4888
Fax: 612-424-2853
jmyers@network.com

Jim Naro
NEARnet
Phone: 617-873-2899
Fax: 617-873-5620
jnaro@nic.near.net

Judy Nasar
O.S.U. Hospitals
Phone: 614-488-6965
Fax: 614-293-3700
jdnasar@magnus.acs.ohio-state.edu

Rina Nathaniel
RAD Network Devices, Ltd.
Phone: +972 3 6458555
Fax: +972 3 5447368
rina!rnd!rndi@uunet.uu.net

Noel Nazario
National Institute of Standards and Technology
Phone: 301-975-2837
Fax: 301-948-0279
nazario@sdns.ncsl.nist.gov

Mark Needleman
University of California
Phone: 510-987-0530
Fax: 510-763-9673
mhn@stubbs.ucop.edu

Bruce Nelson
Novell, Inc.
Phone: 801-429-3875
Fax: 801-429-3086
bnelson@novell.com

Steve Neruda
Morning Star Technologies
Phone: 614-451-1883
Fax: 614-459-5054
neruda@morningstar.com

Clifford Neuman
Information Sciences Institute
Phone: 310-822-1511
bcn@isi.edu

Tom Nisbet
Telecommunications Techniques Corporation
Phone: 301-353-1550
Fax: 301-353-1536
nisbet@tt.com

Shannon Nix
Netlink, Inc.
Phone: 919-878-8612
Fax: 919-872-2132
sdn@netlink.com

Erik Nordmark
Sun Microsystems, Inc.
Phone: 415-336-2788
Fax: 415-336-6015
nordmark@eng.sun.com

Peder Chr. Noergaard
Telebit Communications ApS
Phone: +45 86 28 81 76
Fax: +45 86 28 81 86
pcn@tbit.dk

Bill Norton
Merit Network, Inc.
Phone: 313-936-2656
Fax: 313-747-3745
wbn@merit.edu

William Nowicki
Legato Systems, Inc.
Phone: 415-329-7880
Fax: 415-329-8898
nowicki@legato.com

David O'Leary
cisco Systems
Phone: 714-261-5681
Fax: 714-261-5684
doleary@cisco.com

Masataka Ohta
Tokyo Institute of Technology
Phone: +81 3 5499 7084
Fax: +81 3 3729 1940
mohta@cc.titech.ac.jp

Eric Olinger
Peregrine Systems
Phone: 619-431-2400
Fax: 619-431-0696
eric@peregrine.com

Zbigniew Opalka
Agile Networks
Phone: 508-287-0700
Fax: 508-287-0606
zopalka@agile.com

Ayal Opher
SynOptics Communications, Inc.
Phone: 408-764-1559
Fax: 408-988-5525
aopher@synoptics.com

William Owens
University of Rochester
Phone: 716-275-9120
owens@cc.rochester.edu

Joe Pagan
Department of Defense
Phone: 301-688-8283
jrp@afterlife.ncsc.mil

Joseph Pang
Starlight Networks
Phone: 415-967-0574
Fax: 415-967-0686
pang@bodega.stanford.edu

Brad Parker
FCR Software, Inc.
Phone: 617-424-6552
Fax: 617-424-6552
brad@fcr.com

Andrew Partan
UUNET Technologies, Inc.
Phone: 703-204-8000
Fax: 703-204-8001
asp@uunet.uu.net

Emmanuel Pasetes
Premenos Corporation
Phone: 510-602-2000
Fax: 510-602-2133
ekp@enlil.premenos.sf.ca.us

Brad Passwaters
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
bjp@sura.net

Laura Pate
The MITRE Corporation
Phone: 703-883-5400
Fax: 703-883-7142
pate@gateway.mitre.org

Michael Patton
Bolt Beranek and Newman
Phone: 617-873-2737
map@bbn.com

Gaige Paulsen
InterCon Systems Corporation
Phone: 703-709-5509
Fax: 703-709-5555
gaige@intercon.com

Chris Payson
Octocom Systems, Inc.
Phone: 508-441-2181
Fax: 508-441-9060
!sharps!bar!c_payson@telebit.com

Geir Pedersen
University of Oslo UNINETT
Phone: +47 22 85 24 78
Fax: +47 22 85 27 30
Geir.Pedersen@usit.uio.no

Jon Penner
DCA, Inc.
Phone: 512-244-3871
Fax: 512-388-7090
jjp@bscs.uucp

John Penners
U.S. West Adv. Technologies
Phone: 303-541-6106
Fax: 303-541-8239
jpenners@advtech.uswest.com

Pete Percival
Indiana University
Phone: 812-855-9146
Fax: 812-855-8299
percival@indiana.edu

Maryann Perez
Locus, Inc. at
Naval Research Laboratory
Phone: 202-404-7026
perez@cmf.nrl.navy.mil

Charles Perkins
IBM Corporation
Phone: 914-784-7350
Fax: 914-784-7007
perk@watson.ibm.com

David Perkins
SynOptics Communications, Inc.
Phone: 408-764-1516
Fax: 408-988-5525
dperkins@synoptics.com

Drew Perkins
Fore Systems, Inc.
Phone: 412-967-4040
Fax: 412-967-4044
ddp@fore.com

Radia Perlman
Digital Equipment Corporation
Phone: 508-486-7648
Fax: 508-486-5279
perlman@dsmail.enet.dec.com

Marsha Perrott
PREPnet
Phone: 412-268-7870
Fax: 412-268-7875
mlp+@andrew.cmu.edu

Roy Perry
US West Advanced Technologies
Phone: 303-541-6245
Fax: 303-541-6276
rperry@advtech.uswest.com

James Philippou
Xyplex, Inc.
Phone: 508-264-9901
Fax: 508-264-9930
japhilippou@eng.xyplex.com

David Piscitello
Bellcore
Phone: 908-758-2286
Fax: 908-758-4177
dave@mail.bellcore.com

Ronald Pitt
Network Security Systems, Inc.
Phone: 619-622-8712
Fax: 619-552-9162

Narendra Popat
Frontier Software Development
Phone: 508-851-8872
Fax: 508-851-6956
albin@frontier.com

Willi Porten
GMD - Gessellschaft fuer Datenverarbeitung mbH
Phone: +49 2241 14 3190
Fax: +49 2241 14 3002
porten@gmd.de

Jon Postel
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
postel@isi.edu

Cecilia Preston
SCR
Phone: 415-644-4216
cpreston@info.berkeley.edu

Selina Priestley
Advanced Network and Services, Inc.
Phone: 313-769-3448
Fax: 313-764-6429
sfp@noc.ans.net

Mark Prior
The University of Adelaide
Phone: +61 8 228 5680
Fax: +61 8 223 6245
mrp@itd.adelaide.edu.au

Christopher Provenzano
CompuServe, Inc.
Phone: 614-457-8600
Fax: 614-457-0348
proven@csi.compuserve.com

Thomas Pusateri
Thomas & Willow Consulting
Phone: 919-676-9941
pusateri@cs.duke.edu

Kelly Quaal
Network Systems Corporation
Phone: 612-242-4888
Fax: 612-424-2853
quaalk@network.com

Doug Rademacher
American Power Conversion
Phone: 401-789-5735
Fax: 401-789-3710
apc!dougr@uunet.uu.net

Robert Raisch
O'Reilly and Associates, Inc.
Phone: 617-354-5800
Fax: 617-661-1116
raisch@ora.com

Bala Rajagopalan
AT&T Bell Laboratories
Phone: 908-949-8017
Fax: 908-949-1726
braja@qsun.att.com

Dave Rand
Novell, Inc.
Phone: 408-321-1259
Fax: 408-946-9761
dave_rand@novell.com

Venkat Rangan
Metrix Network Systems, Inc.
Phone: 603-888-7000
Fax: 603-891-2796
venkat@.metrix.com

Marcus Ranum
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
mjr@tis.com

Owen Reddecliffe
Walker Richer & Quinn Inc.
Phone: 206-726-7345
Fax: 206-860-9079
owen%wrq@mcimail.com

Edward Reed
Xerox Corporation
Phone: 716-422-4221
eer@cinops.xerox.com

Jim Rees
University of Michigan
Phone: 313-763-4174
Fax: 313-763-4434
jim.rees@umich.edu

David Reese
California State University
Phone: 714-773-2159
Fax: 310-985-9400
dave@csu.net

Louise Reingold
Sprint
Phone: 703-689-5940
Fax: 703-689-6998
l.reingold@sprint.sprint.com

Yakov Rekhter
IBM Corporation
Phone: 914-945-3896
Fax: 914-945-2141
yakov@watson.ibm.com

Sheri Repucci
Merit Network, Inc.
Phone: 313-936-3335
smr@merit.edu

Robert Reschly
US Army Research Laboratory
Phone: 410-278-6808
Fax: 410-278-5075
reschly@brl.mil

Joyce K. Reynolds
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
jkrey@isi.edu

Bradley Rhoades
Metal Matrix Composites
Phone: 612-736-2874
Fax: 612-736-0431
bdrhoades@mail.mmmg.com

Tony Richards
Sprint Government Systems Division
Phone: 703-904-2460
Fax: 703-904-2463
richards@icm1.icp.net

Steven Richardson
Merit Network, Inc.
Phone: 313-747-4813
Fax: 313-747-3745
sjr@merit.edu

April Richstein
Department of Defense
Phone: 301-688-0849
Fax: 301-688-0289
amr@tycho.ncsc.mil

Michael Roberts
EDUCOM
Phone: 202-331-5347
Fax: 202-872-4318
mmr@darwin.ptvy.ca.us

Sam Roberts
Farallon Computing, Inc.
Phone: 510-814-5215
Fax: 510-814-5024
sroberts@farallon.com

Ben Robinson
IBM Corporation
Phone: 919-254-5734
ben_robinson@vnet.ibm.com

Francois Robitaille
Centre De Recherche Informatique De Montreal
Phone: 514-398-1234
Fax: 514-398-1244
francois.robitaille@crim.ca

Kenneth Rodemann
AT&T Bell Laboratories
Phone: 908-949-6229
Fax: 908-949-1726
krr@qsun.att.com

Richard Rodgers
National Library of Medicine
Phone: 301-496-9300
Fax: 301-496-0673
rodgers@nlm.nih.gov

Benny Rodrig
RAD Network Devices, Ltd.
Phone: +972 3 6458552
Fax: +972 3 5447368
4373580@mcimail.com

Manoel Rodrigues
AT&T Bell Laboratories
Phone: 908-949-4655
Fax: 908-949-1726
manoel_rodrigues@att.com

Jim Romaguera
NetConsult AG
Phone: +41 31 235750
Fax: +41 31 235792
romaguera@cosine-mhs.switch.ch

Dan Romascanu
Lannet Data Communications
Phone: +972 3 6458414
Fax: +972 3 5447146
dan@lannet.com

Steve Romig
The Ohio State University
Phone: 614-292-8018
Fax: 614-292-9021
romig@cis.ohio-state.edu

John Romkey
ELF Communications
Phone: 617-942-7846
romkey@asylum.sf.ca.us

Yzhak Ronen
AT&T Bell Laboratories
Phone: 908-949-1013
Fax: 908-949-8569
y.ronen@homxa.att.com

Marshall Rose
Dover Beach Consulting, Inc.
Phone: 415-968-1052
Fax: 415-968-2510
mrose@dbc.mtview.ca.us

Shawn Routhier
Epilogue Technology
Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sar@epilogue.com

Gary Rowe
Rapport Communication
Phone: 513-779-8750
Fax: 513-779-8758
gjrowe@attmail.com

Rick Royston
U.S. Robotics
Phone: 708-933-5430
Fax: 708-982-1348
rick@lsumvs.sncc.lsu.edu

Allan Rubens
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
acr@merit.edu

Joseph Rumolo
CNR, Inc.
Phone: 908-291-7934
Fax: 908-872-2063
jrumolo@attmail.com

Michael Safly
Boeing Computer Services
Phone: 205-544-6445
Fax: 205-544-4618
saf@tank1.msfc.nasa.gov

Timothy Salo
Minnesota Supercomputer Center
Phone: 612-626-0347
Fax: 612-624-6550
tjs@msc.edu

Hal Sandick
IBM Corporation
Phone: 919-254-4614
Fax: 919-254-5483
sandick@vnet.ibm.com

Paul Sangster
Advanced Network and Services, Inc.
Phone: 703-758-7706
Fax: 703-758-7717
sangster@reston.ans.net

Jon Saperia
Digital Equipment Corporation
Phone: 508-952-3171
saperia@tay.dec.com

Vilson Sarto
Rede Nacional de Pesquisa
Phone: +55 11 837 03 11
Fax: +55 11 261 41 67
vilson@fapq.fapesp.br

Srinivas Sataluri
AT&T Bell Laboratories
Phone: 908-949-7782
Fax: 908-949-1726
sri@qsun.att.com

Shiva Sawant
SynOptics Communications, Inc.
Phone: 408-764-1612
Fax: 408-764-1699
shiva@synoptics.com

Jeffrey Schiller
Massachusetts Institute of Technology
Phone: 617-253-0161
Fax: 617-258-8736
jis@mit.edu

Andrew Schmidt
University of Illinois
Phone: 217-244-0582
Fax: 217-244-7089
ags@uius.edu

Wolfgang Schneider
RARE/GMD
Phone: +49 6151 869 262
Fax: +49 6151 869 785
schneider@gmd.de

Carl Schoeneberger
Tigon Corporation
Phone: 214-733-2704
Fax: 214-733-8666
70410.3563@Compuserve.com

Martin Schoffstall
Performance Systems International, Inc.
Phone: 703-620-6651
Fax: 703-620-4586
schoff@psi.com

Eve Schooler
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
schooler@isi.edu

John Scudder
Merit Network, Inc.
Phone: 313-936-3335
jgs@merit.edu

Tim Seaver
MCNC - CONCERT Network
Phone: 919-248-1973
Fax: 919-248-1405
tas@concert.net

Jennifer Sellers
NASA NREN
Phone: 202-434-8954
Fax: 202-434-4599
sellers@nsipo.nasa.gov

Paul Serice
Corporation for Open Systems International
Phone: 703-205-2713
Fax: 703-846-8590
serice@cos.com

Kanan Shah
Locus, Inc.
Phone: 703-960-1000
Fax: 703-960-0643
kshag@cmf.nrl.navy.mil

Ron Sharp
AT&T Bell Laboratories
Phone: 201-386-4479
Fax: 201-386-5549
rls@neptune.att.com

Chris Shaw
Banyan Systems, Inc.
Phone: 508-836-2898
Fax: 508-898-3604
cshaw@banyan.com

Karen Shay
CMU
Phone: 412-268-3617
kshay@cs.cmu.edu

Gregory Sheehan
AT&T Bell Laboratories
Phone: 908-949-9442
Fax: 908-949-8569
gregory.c.sheehan@att.com

Kitty Shih
Novell, Inc.
Phone: 408-747-4305
Fax: 408-747-4900
kmshih@novell.com

Robert Shirey
The MITRE Corporation
Phone: 703-883-7210
Fax: 703-883-1397
shirey@mitre.org

Allyson Showalter
NASA Science Internet
Phone: 202-434-8954
Fax: 202-434-4599
allyson@nsipo.arc.nasa.gov

Mark Silverman
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
mark@dss.com

William Simpson
Computer Systems Consulting Services
Phone: 313-623-7979
Bill.Simpson@um.cc.umich.edu

W. David Sincoskie
Bell Communications Research
Phone: 201-829-4426
Fax: 201-829-2504
sincos@thumper.bellcore.com

Sam Sjogren
TGV, Inc.
Phone: 408-427-4366
Fax: 408-427-4365
sjogren@tgv.com

Keith Sklower
University of California, Berkeley
Phone: 510-642-9587
Fax: 510-642-5775
sklower@cs.berkeley.edu

Timon Sloane
Timon Sloane Consulting Services
Phone: 415-515-0741
Fax: 415-941-4679
timon@timon.com

Andrew Smith
SynOptics Communications, Inc.
Phone: 408-764-1574
Fax: 408-988-5525
asmith@synoptics.com

Charlie Smith
The Ohio State University
Phone: 614-292-2899
Fax: 614-292-7081
crsmith@osvi.edu

Craig Smith
DAS/Division of Computer Services, State of Ohio
Phone: 614-466-9602
Fax: 614-752-6108
smith@ohio.gov

Jane Smith
MCNC - CNIDR
Phone: 919-248-9213
Fax: 919-248-1405
Jane.Smith@cnidr.org

Kim Smith
Advanced Network and Services, Inc.
Phone: 313-936-2090
Fax: 313-763-4050
kas@noc.ans.net

Mark Smith
University of Michigan
Phone: 313-764-2277
Fax: 313-764-5140
mcs@umich.edu

Patricia Smith
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
psmith@merit.edu

Suzanne Smith
Lawrence Livermore National Laboratory
Phone: 510-422-4017
Fax: 510-422-1482
smith@es.net

Larry Snodgrass
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
snodgrass@bitnic.educom.edu

Karen Sollins
Massachusetts Institute of Technology
Phone: 617-253-6002
Fax: 617-253-2673
sollins@lcs.mit.edu

Marco Sosa
Bell Communications Research
Phone: 908-758-2651
Fax: 908-758-4177
mxs@sabre.bellcore.com

Raj Srinivasan
Sun Microsystems, Inc.
Phone: 415-336-2478
Fax: 415-336-3156
raj.srinivasan@eng.sun.com

Michael St. Johns
Advanced Research Projects Agency
Phone: 703-696-2271
Fax: 703-696-2202
stjohns@darpa.mil

Stuart Stanley
Apertus Technologies, Inc.
Phone: 612-828-0391
Fax: 612-828-0773
stuarts@apertus.com

Ira Steckler
Chipcom Corporation
Phone: 508-460-8922
Fax: 508-460-8952
isteckle@chipcom.com

Martha Steenstrup
Bolt Beranek and Newman
Phone: 617-873-3192
Fax: 617-873-3776
msteenst@bbn.com

Einar Stefferud
Network Management
Associates, Inc.
Phone: 714-842-3711
Fax: 714-848-2091
stef@nma.com

Scott Stein
Apple Computer, Inc.
Phone: 408-974-3124
Fax: 408-974-9793
scotts@apple.com

Barbara Sterling
McDATA
Phone: 303-460-4211
Fax: 303-465-4996
bjs@mcdata.com

Bob Stewart
Xyplex, Inc.
Phone: 508-264-9900
Fax: 508-264-9930
rlstewart@eng.xyplex.com

Bernhard Stockman
EBONE
Phone: +46 8 790 6519
Fax: +46 8 241 179
boss@ebone.net

Adam Stolinski
Liebert Corporation
Phone: 714-457-3600
Fax: 714-457-3787
stolinsk@cerf.net

Roxanne Streeter
NASA NREN
Phone: 415-604-4845
Fax: 415-604-0063
streeter@nsipo.arc.nasa.gov

Richard Stuart
PENRIL
Phone: 301-921-8600
Fax: 301-921-8376
dickstuart@attmail.com

Stuart Stubblebine
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
stubblebine@isi.edu

Subbu Subramaniam
Hewlett-Packard
Phone: 408-447-2693
Fax: 408-447-3660
subbu@cup.hp.com

Subu Subramanian
AT&T Bell Laboratories
Phone: 908-234-3247
Fax: 908-719-8058
ssubramanian@attmail.com

Terry Sullivan
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
terrys@newbridge.com

Catherine Summers
Corporation for Open Systems
Phone: 703-205-2801
Fax: 703-846-8590
cfs@cos.com

Bob Sutterfield
Morning Star Technologies
Phone: 614-451-1883
Fax: 614-459-5054
bob@morningstar.com

Steve Suzuki
Furukawa Electric Tech., Inc.
Phone: 408-248-4884
Fax: 408-248-8815
suzu@fet.com

Richard Sweatt
SynOptics Communications, Inc.
Phone: 408-764-1260
Fax: 408-764-1299
rsweatt@synoptics.com

Steven Sweeney
Farallon Computing, Inc.
Phone: 510-814-5203
Fax: 510-814-5024
steves@farallon.com

Wayne Tackabury
Cayman Systems
Phone: 617-494-1999
Fax: 617-494-5167
wayne@cayman.com

Sally Tarquinio
The MITRE Corporation
Phone: 703-883-7987
Fax: 703-883-7142
sallyt@gateway.mitre.org

John Tavs
IBM Corporation
Phone: 919-254-7610
Fax: 919-254-0151
tavs@vnet.ibm.com

Fumio Teraoka
Sony Computer Science Laboratory, Inc.
Phone: +81 3 3448 4380
Fax: +81 3 3448 4273
tera@csl.sony.co.jp

Marten Terpstra
RIPE Network Coordination Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
marten@ripe.net

Kaj Tesink
Bell Communications Research
Phone: 908-758-5254
Fax: 908-758-4196
kaj@cc.bellcore.com

Kamlesh Tewani
AT&T Bell Laboratories
Phone: 908-949-5369
Fax: 908-949-1196
ktt@arch2.att.com

Richard Thomas
Bell Northern Research
Phone: 613-763-7025
Fax: 613-765-4920
rjthomas@bnr.ca

Geoffrey Thompson
SynOptics Communications, Inc.
Phone: 408-764-1339
Fax: 408-988-5525
thompson@synoptics.com

Steven Thompson
Soft-Switch, Inc.
Phone: 215-640-9600
Fax: 215-651-5358
sjt@gateway.ssw.com

Louisa Thomson
Hughes Aircraft Company
Phone: 310-364-6369
Fax: 310-364-6368
louisa@whitney.hac.com

Susan Thomson
Bell Communications Research
Phone: 201-829-4514
Fax: 201-829-2504
set@bellcore.com

Craig Todd
Wright State University
Phone: 513-873-4827
ctodd@desire.wright.edu

Marek Tomaszewski
Network Equipment Technologies
Phone: 415-780-4235
Fax: 415-780-4272
marek@net.com

Kim Toms
Morning Star Technologies
Phone: 614-451-1883
Fax: 614-459-5054
kim@morningstar.com

Claudio Topolcic
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
topolcic@cnri.reston.va.us

Paul Traina
cisco Systems
Phone: 415-688-8189
Fax: 415-688-1945
pst@cisco.com

Klaus Truoel
RARE/GMD
Phone: +49 6151 869 700
Fax: +49 6151 869 785
truoel@gmd.de

Theodore Ts'o
Massachusetts Institute of Technology
Phone: 617-253-8091
Fax: 617-258-8736
tytso@mit.edu

Panos-Gavriil Tsigaridas
Research Center for Open Communication Systems FOKUS
Phone: +49 30 25499232
Fax: +49 30 25499202
Tsigaridas@fokus.berlin.gmd.dbp.de

Stephen Tsun
3Com Corporation
Phone: 408-764-6978
Fax: 408-764-5002
snt@3com.com

James Tufano
The Ohio State University
Phone: 614-292-4843
Fax: 614-292-7081
jtufano@magnus.ohio-state.edu

Matt Tuttle
MCNC Center for Communications
Phone: 919-248-1499
Fax: 919-248-1405
snodgras@cren.net

Robert Ullmann
Process Software Corporation
Phone: 508-879-6994
Fax: 508-879-0042
Ariel@Process.com

L. Stuart Vance
TGV, Inc.
Phone: 408-427-4366
Fax: 408-427-4365
vance@tgv.com

Kannan Varadhan
OARnet
Phone: 614-292-4137
Fax: 614-292-7168
kannan@oar.net

Gregory Vaudreuil
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
gvaudre@cnri.reston.va.us

Mario Vecchi
Bell Communications Research
Phone: 201-829-4374
Fax: 201-829-2504
mpv@thumper.bellcore.com

John Veizades
Apple Computer, Inc.
Phone: 408-974-2672
Fax: 408-862-1159
veizades@apple.com

Dinesh Verma
IBM Corporation
Phone: 914-784-6403
Fax: 914-784-7455
verma@watson.ibm.com

Warren Vik
Lachman Technology, Inc.
Phone: 708-505-9555
Fax: 708-505-9574
wmv@lachman.com

Curtis Villamizar
Advanced Network and
Services, Inc.
Phone: 914-789-5300
Fax: 914-789-5310
curtis@ans.net

Ioannis Viniotis
North Carolina State
University
Phone: 919-515-5148
Fax: 919-515-5523
candice@ececho.ncsu.edu

Ruediger Volk
University of Dortmund
Phone: +49 231 755 4760
Fax: +49 231 755 2386
rv@informatik.uni-dortmund.de

John Vollbrecht
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
jrv@merit.edu

Janet Vratny
Apple Library
Phone: 408-974-2048
Fax: 408-725-8502
janet@apple.com

Hung Vu
Fonorola
Phone: 613-235-3666
Fax: 613-232-4329
hungv@fonorola.com

Huyen Vu
Defense Information Systems
Agency
Phone: 703-487-3038
Fax: 703-487-3190
vu@polaris.disa.mil

Steven Waldbusser
Carnegie Mellon University
Phone: 412-268-6628
Fax: 412-268-4987
waldbusser@andrew.cmu.edu

Chuck Warlick
Boeing Computer Services
Phone: 205-544-6270
Fax: 205-544-4618
warlick@theophilis.nsfc.nasa.gov

William Warner
State of Ohio,
Telecommunications
Phone: 614-466-6683
Fax: 614-644-3349
warner@ohio.gov

Richard Warwick
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
richard@dss.com

Scott Wasson
Xyplex, Inc.
Phone: 508-264-9901
Fax: 508-264-9930
sgwasson@eng.xyplex.com

Jack Waters
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
waters@sura.net

James Watt
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
james@newbridge.com

James Weatherford
Naval Command, Control and
Ocean Surveillance Center
Phone: 619-553-4090
Fax: 619-553-6288
weatherf@nosc.mil

Chris Weider
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3185
clw@merit.edu

Terry Weigler
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
tweigler@cnri.reston.va.us

Abel Weinrib
Bell Communications Research
Phone: 201-829-4454
Fax: 201-829-5889
abel@bellcore.com

Von Welch
National Center for
Supercomputing Applications
Phone: 217-333-1989
Fax: 217-244-1987
vwelch@ncsa.uiuc.edu

Jonathan Wenocur
Shiva Corporation
Phone: 617-270-8347
jhw@shiva.com

Moira West
CERT Coordination Center
Phone: 412-268-7090
Fax: 412-268-6989
mjw@cert.org

Fred Whiteside
Beame and Whiteside Software
Phone: 416-765-0822
Fax: 416-765-0815
fred@bws.com

Les Wibberley
Chemical Abstracts Service
Phone: 614-447-3600
Fax: 614-447-3713
lhw24@cas.org

John Wieronski
OARnet
Phone: 614-292-9290
Fax: 614-292-7168
john@oar.net

Richard Wiggins
Michigan State University
Phone: 517-353-4955
Fax: 517-353-9847
wiggins@msu.edu

Rick Wilder
Advanced Network and Services, Inc.
Phone: 703-758-7703
Fax: 703-758-7717
wilder@ans.net

Steve Willens
Livingston Enterprises, Inc.
Phone: 510-426-0770
Fax: 510-426-8951
steve@livingston.com

Douglas Williams
IBM Corporation
Phone: 919-254-4801
dougw@ralvmg.vnet.ibm.com

Jim Williams
Merit Network, Inc.
Phone: 313-936-3600
jimw@merit.edu

Kirk Williams
Southwestern Bell Technology Resources
Phone: 314-529-7586
Fax: 314-529-7674
kirk@sbctri.sbc.com

Alan Williamson
Network Solutions, Inc.
Phone: 703-742-4820
Fax: 703-742-4811
scottw@nic.ddn.mil

Steven Willis
Wellfleet Communications, Inc.
Phone: 617-280-2427
Fax: 617-280-2300
steve@wellfleet.com

Peter Wilson
3Com Limited
Phone: +44 442 231000
Fax: +44 442 254701
peter_wilson@3com.com

Steve Wilson
Morning Star Technologies
Phone: 614-451-1883
Fax: 614-459-5054
steve@morningstar.com

Walter Wimer
Carnegie Mellon University
Phone: 412-268-6252
Fax: 412-268-4987
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
Phone: 708-252-7236
Fax: 708-252-5983
lwinkler@anl.gov

Dan Wintringham
OARnet
Phone: 614-292-0901
Fax: 614-292-7168
danw@oar.net

Cathy Wittbrodt
BARRNet
Phone: 415-725-5481
Fax: 415-723-0100
cjw@barrnet.net

Jane Wojcik
Bolt Beranek and Newman
Phone: 617-873-2764
Fax: 617-873-4086
jwojcik@bbn.com

Richard Woundy
IBM Corporation
Phone: 203-783-7996
Fax: 203-783-7026
rwoundy@vnet.ibm.com

Russ Wright
Lawrence Berkeley Laboratory
Phone: 510-486-6965
Fax: 510-486-6363
wright@lbl.gov

John Wroclawski
Massachusetts Institute of Technology
Phone: 617-253-6002
Fax: 617-253-2673
jtw@lcs.mit.edu

Liang Wu
Bell Communications Research
Phone: 201-829-4365
Fax: 201-829-2504
ltw99@bellcore.com

Yow-Wei Yao
IBM Corporation
Phone: 512-818-3548
yao@chang.austin.ibm.com

Brian Young
Best Power Technology, Inc.
Phone: 800-356-5794
Fax: 608-565-2221
byoung@macc.wisc.edu

Charles Young
AT&T Bell Laboratories
Phone: 908-949-6234
Fax: 908-949-1196
Charles.E.Young@att.com

Chin Yuan
Pacific Bell
Phone: 510-867-5016
Fax: 510-867-1405
cxyuan@pacbell.com

Kiho Yum
3Com Corporation
Phone: 408-764-6400
Fax: 408-754-5003
kxy@nsd.3com.com

William Yurcik
NASA
Phone: 301-286-1376
Fax: 301-286-5152
yurcik@dftnic.gsfc.nasa.gov

Paul Zawada
National Center for Supercomputing Applications
Phone: 217-333-1189
Fax: 217-244-1987
Zawada@ncsa.uiuc.edu

James Zmuda
Hughes Aircraft Company
Phone: 714-732-7240
Fax: 714-732-2000
zmuda@mls.hac.com

Terry Zumwalt
International Power Machines Corporation
Phone: 214-205-1540
Fax: 214-494-2690
71151.2550@compuserve.com