ACKNOWLEDGEMENTS

November 1-5, 1993 marked the twenty-eighth meeting of the Internet Engineering Task Force. The meeting was held in Houston, Texas and was hosted by Rice University and SESQUINET. Both organizations provided funding and SESQUINET provided routing and oversight as well. Bill Manning of Rice University was our primary contact and did a wonderful job of overseeing the set up of the on-site terminal room. As Bill acknowledged during Monday morning's Open Plenary his goal was to make it very difficult to top the setup in Houston. Judging by the compliments which flowed from attendees, Bill and those who worked with him may well have done just that. Among those who assisted with the setup and operation of the terminal room were: Stan Barber, Blair Copland, Cathy Foulston, Farrell Gerbode, William soren Deigaard, Jain Li, Kevin Mullet and Steve Blair. A special thanks goes to David Boyes for his supervision of the multicasting process.

The following organizations deserve recognition and thanks for ensuring that the terminal room was exceptionally furnished.

Cabletron	3COM HUB
University North Texas	Cabletron HUB
Rice University	Gatorbox and Localtalk Hardware
	Laserwriters
Sun Microsystems	Sun Workstations
Network Computing Devices	NCD Terminals
Baylor College of Medicine	NCD Terminals
Apple Computer	Macintosh Computers
Dell Computer	10baseT Hardware
	Localtalk Hardware
	10base2 Cable
TRI/Southwestern Bell	FDDI MAN and cisco Routers

The Houston IETF marked the third time that presentations were given relating to the Next Generation IP. Monday's opening session included scheduled presentations by Allison Mankin and Scott Bradner (IP: Next Generation Overview), Frank Solensky (ALE), Steve Deering (SIPP), Rob Ullman (TP/IX) and Peter Ford (TUBA). Presentations later in the week were made by Tony Li (CIDR Status), Craig Partridge (ST2), Harald Alvestrand (Using E-Mail in Europe) and Vint Cerf (Intellectual Property Rights and the ISO/MOU Draft). Dick DesJardins gave a brief presentation on the Federal Internetworking Requirements Panel.

Closing acknowledgements and thanks go to Debra Legare, Cynthia Clark, Lois Keiper and Terry Weigler for their tireless efforts to ensure a successful and smooth running meeting.

Megan Walnut
IETF Meeting Coordinator

Contents

Director's Message **1**

Agenda of the Twenty-Eighth IETF **7**

1 IETF Overview **17**
- 1.1 Future IETF Meeting Sites 21
- 1.2 On-Line IETF Information 23
- 1.3 Guidelines to Authors of Internet-Drafts 27

2 Area and Working Group Reports **29**
- 2.1 Applications Area 31
  - 2.1.1 Internet Message Access Protocol (IMAP) 51
  - 2.1.2 OSI Directory Services (OSIDS) 57
  - 2.1.3 TELNET (TELNET) 65
  - 2.1.4 Telnet TN3270 Enhancements (TN3270E) 69
  - 2.1.5 X.400 Operations (X400OPS) 77
- 2.2 Internet Area 85
  - 2.2.1 Dynamic Host Configuration (DHC) 91
  - 2.2.2 Internet Stream Protocol V2 (ST2) 103
  - 2.2.3 IP Over Large Public Data Networks (IPLPDN) 111
  - 2.2.4 IP Over AppleTalk (APPLEIP) 117
  - 2.2.5 IP Over Asynchronous Transfer Mode (ATM) 119
  - 2.2.6 Point-to-Point Protocol Extensions (PPPEXT) 127
  - 2.2.7 Router Requirements (RREQ) 141
- 2.3 IP: Next Generation Area 143
  - 2.3.1 P. Internet Protocol (PIP) 157
  - 2.3.2 Simple Internet Protocol (SIP) 159
  - 2.3.3 TP/IX (TPIX) 187
  - 2.3.4 TCP/UDP Over CLNP-Addressed Networks (TUBA) 193
- 2.4 Network Management Area 201
  - 2.4.1 ATM MIB (ATOMMIB) 207
  - 2.4.2 Character MIB (CHARMIB) 213
  - 2.4.3 Frame Relay Service MIB (FRNETMIB) 215

2.4.4 Interfaces MIB (IFMIB) 219
2.4.5 Modem Management (MODEMMGT) 225
2.4.6 Remote LAN Monitoring (RMONMIB) 229
2.4.7 SNA DLC Services MIB (SNADLC) 235
2.4.8 SNA NAU Services MIB (SNANAU) 243
2.4.9 Uninterruptible Power Supply (UPSMIB) 249
2.5 Operational Requirements Area 253
2.5.1 BGP Deployment and Application (BGPDEPL) 267
2.5.2 Benchmarking Methodology (BMWG) 277
2.5.3 Generic Internet Service Description (GISD) 281
2.5.4 Network Joint Management (NJM) 287
2.5.5 Network OSI Operations (NOOP) 293
2.5.6 Operational Statistics (OPSTAT) 295
2.5.7 User Connectivity (UCP) 297
2.6 Routing Area 299
2.6.1 Border Gateway Protocol (BGP) 311
2.6.2 IP Routing for Wireless/Mobile Hosts (MOBILEIP) 329
2.6.3 IS-IS for IP Internets (ISIS) 345
2.6.4 Inter-Domain Multicast Routing (IDMR) 349
2.6.5 Inter-Domain Policy Routing (IDPR) 357
2.6.6 Multicast Extensions to OSPF (MOSPF) 361
2.6.7 OSI IDRP for IP Over IP (IPIDRP) 363
2.6.8 Open Shortest Path First IGP (OSPF) 367
2.6.9 RIP Version II (RIPV2) 373
2.6.10 Routing over Large Clouds (ROLC) 377
2.6.11 Source Demand Routing (SDR) 391
2.7 Security Area 395
2.7.1 Authorization and Access Control (AAC) 401
2.7.2 Commercial Internet Protocol Security Option (CIPSO) . . . 413
2.7.3 Common Authentication Technology (CAT) 415
2.7.4 Internet Protocol Security Protocol (IPSEC) 425
2.7.5 Network Access Server Requirements (NASREQ) 445
2.7.6 Privacy-Enhanced Electronic Mail (PEM) 451
2.8 Service Applications Area 459
2.8.1 Domain Name System (DNS) 461
2.8.2 MHS-DS (MHSDS) 469
2.8.3 Minimal OSI Upper-Layers (THINOSI) 471
2.8.4 Network Database (NETDATA) 477
2.8.5 Network Printing Protocol (NPP) 479
2.8.6 Service Location Protocol (SVRLOC) 481
2.8.7 Trusted Network File Systems (TNFS) 485
2.9 Transport Area 487

2.9.1 Audio/Video Transport (AVT) 513
2.9.2 Multiparty Multimedia Session Control (MMUSIC) 521
2.9.3 TCP Large Windows (TCPLW) 543
2.10 User Services Area 545
2.10.1 Integrated Directory Services (IDS) 551
2.10.2 Integration of Internet Information Resources (IIIR) 557
2.10.3 Internet Anonymous FTP Archives (IAFA) 563
2.10.4 Internet School Networking (ISN) 565
2.10.5 Network Information Services Infrastructure (NISI) 569
2.10.6 Network Training Materials (TRAINMAT) 575
2.10.7 Networked Information Retrieval (NIR) 579
2.10.8 Uniform Resource Identifiers (URI) 585
2.10.9 User Documents Revisions (USERDOC2) 593
2.10.10 User Services (USWG) 597
2.10.11 Whois and Network Information Lookup Service (WNILS) . . 613

3 Network Status Briefings **619**
3.1 NASA Science Internet 621
3.2 ANSNet 623

4 IETF Protocol Presentations **629**
4.1 A Report on the Temporary IPng Area 631
4.2 IPng: ALE 637
4.3 IPng: SIPP 639
4.4 IPng: TUBA 641
4.5 CIDRD Status 645
4.6 ST-II 651

5 Technical Presentations **655**
5.1 Using E-mail in Europe: Opportunities and Challenges 657
5.2 Federal Internetworking Requirements Panel (FIRP) 661
5.3 Intellectual Property Policy 663
5.4 Status of the ISOC/ISO Liaison MOU 667

6 Workshop for Working Group Chairs **671**
6.1 Workshop for Working Group Chairs 673

A Attendees **679**

Director's Message

The 28th meeting of the Internet Engineering Task Force was held in Houston, Texas, November 1-5, 1993. The meeting was co-hosted by SESQUINET and Rice University, and our thanks and appreciation go out to Bill Manning and to all the others who helped with the terminal room and the social event.

I'm not sure how much longer the growth of the IETF meetings can continue. This was one of the top IETF meetings based on attendance with just over 635 registered attendees in Houston, making this the third largest meeting to date (slightly exceeding the Washington D.C. meeting which had 633 attendees).

The number of first time attendees remained close to the 200 mark. There were 173 first time attendees at the Houston meeting (approximately 27%). About 100 people showed up for the Newcomers' Orientation on Sunday afternoon.

Multicast IETF Meetings

From the multicast perspective, things keep growing and growing, and it is impossible to guess how many people are "listening in" during the IETF meetings. There were more than 600 hosts on the receiving end of the Houston IETF multicast effort, up from approximately 400 hosts on-line during the Amsterdam meeting. The number of countries listening in stayed the same at 16. This technology is increasing the number of virtual attendees at these meetings, and coverage is not limited to the technical presentations and plenaries. The broadcast system is on wheels, and it is not uncommon to see the multicast volunteers (from the host group) wheeling a cart from working group to working group meeting!

The IETF is Going "Green"

Recognizing the concern for the environment, and the focus on electronic distribution of information for which the Internet is famous, the IETF consensus during the Thursday afternoon Open Plenary was to begin moving towards an electronic version of the meeting proceedings. This effort will reduce the paper requirements (saving trees) and the overall cost of printing and distributing the proceedings.

The IETF Secretariat has always made the minutes available in the IETF shadow directories, but this is not an electronic version of the proceedings. They do not include the overheads from the technical presentations or those used during working group meetings. Obviously these must be included in the electronic proceedings.

It is also understood that merely having files available for copying via FTP is not sufficient, and there is much more that can, and will, be done.

The Secretariat is already examining what must be done in order to provide electronic proceedings. We are looking into how we might be able to store and provide the overhead materials, and we are looking into the hardware and software needed to scan the overheads and create some transportable file (i.e. Postscript). Many of the overheads used are created on workstations and computers that have the capability of producing Postscript files (and other formats as well), and this will be factored into our process.

Obviously, there will be no shortage of suggestions and techniques (indeed, these are already coming in). It should be noted that this will be conducted as an experiment for the Seattle IETF meeting and there will be many changes suggested in the future as we gather experience and take advantage of new technologies and capabilities.

The Secretariat will also be looking into how this information can be provided to the Internet community. Initially, we are planning to make everything available via gopher. Future plans include investigating the options of other mechanisms such as hyper-text and possibly distributing the proceedings on CD-ROM. In fact, depending on storage requirements and capabilities (not to mention demand), it might be possible to create a single CD with the proceedings for an entire year. I can just see it now, the IETF's Greatest Hits of 1994!

There will be a change in the registration form for the Seattle IETF meeting in 1994. People will be asked to indicate if they want to receive a printed copy of the proceedings. Remember that this is an experiment, and make your choice appropriately.

What a Year!

This has been quite a year for the IETF. We began the year with IPNG candidate demonstrations in the terminal room at Columbus, status updates in Amsterdam, and the consolidation of SIP and PIP into a single effort by the November meeting.

Another first-time event at the Columbus IETF meeting was the announcement of new members to the IESG and IAB. This was the first implementation of the selection process defined by the POISED Working Group.

The IESG established a special ad-hoc Area for all the IPNG related working groups, and by November the IPNG Directorate had been announced, along with a six month plan of action.

We held the first IETF meeting outside of North America, and future non-North American meetings are being planned as I write this message.

Multi-casting is no longer a "special" component but an integral part of the meetings themselves. We've seen the number of receiving sites grow to more than 600 in over 15 countries.

And the world has discovered the Internet. A significant number of books have appeared in bookstores, many articles are printed in the press, a cartoon appeared in the New Yorker Magazine (see the Amsterdam proceedings), and even Doonesbury has gotten into the act. More and more "mainstream" publications are carrying information on the Internet. More and more services are being offered and discussed.

There are a number of new products (user interfaces) that are available to all Interneters; new tools and features are anticipated all the time, and are being worked on today. Capabilities we are only now beginning to conceptualize will probably be designed, implemented, distributed, and re-implemented (good ol' Version 2, eh?) by this time next year. Traditional concepts are being challenged and rethought as the general public moves into cyberspace.

Consider electronic publication... this is/will be much more than merely having the articles and pictures, along with the cover and title pages, available on-line for electronic distribution or browsing. The entire concept of books will be re-examined as one considers the capabilities available today (and conceptualize what *could* be available tomorrow)... additional references, use of new technologies such as hypertext, knowbots, links to reference material and even more... two-way communication! Just imagine an application where a "reader" can ask the author to elaborate on a concept, clarify with additional examples, or even to submit additional queries.

"May you live through interesting times" is an ancient Chinese curse. However, I am looking forward to more interesting times as new capabilities are provided and we improve our ability to perceive what cyberspace has to offer.

Future Meetings

The next IETF meeting will be in Seattle, Washington the last week of March (March 28 - April 1, 1993). This meeting is being hosted by NorthWestNet. Following Seattle, we will be travelling to Canada for the summer meeting which is scheduled to be in Toronto, Ontario from July 25-29, 1994. The final meeting of 1994 will be in the San Francisco Bay area; presently, the Secretariat staff are working with the host group to identify the meeting time and place. Once this information is known it will be broadcasted to the IETF Announcement mailing list.

Note that information on future IETF meetings can always be found in the file `/ietf/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 Amsterdam IETF Meeting last July; over 125 Internet-Drafts, 26 Protocol Actions, and over 65 RFCs.

Between the IETF meetings in Amsterdam, The Netherlands and Houston, Texas, there were four new working groups created:

1. Remote Lan Monitoring (RMBONMIB)
2. Generic Internet Service Description (GISD)
3. Internet Stream Protocol V2 (ST2)
4. Routing over Large Clouds (ROLC)

and ten working groups were concluded:

1. FDDI MIB (FDDIMIB)
2. Bridge MIB (BRIDGE)
3. IP Over Large Public Data Networks (IPLPDN)
4. Internet Message Extensions (822EXT)
5. Network News Transport Protocol (NNTP)
6. IEEE 802.3 HUB MIB (HUBMIB)
7. Token Ring Remote Monitoring (TRMON)
8. Chassis MIB (CHASSIS)
9. Host Resources MIB (HOSTMIB)
10. MIME-MHS Internetworking (MIMEMHS)

Additionally, 66 RFCs have been published since the Amsterdam IETF meeting in July, 1993:

RFC	Status	Title
RFC1440	E	SIFT/UFT: Sender-Initiated/Unsolicited File Transfer
RFC1467	I	Status of CIDR Deployment in the Internet
RFC1477	I	IDPR as a Proposed Standard
RFC1478	PS	An Architecture for Inter-Domain Policy Routing
RFC1479	PS	Inter-Domain Policy Routing Protocol Specification: Version 1
RFC1482	I	Aggregation Support in the NSFNET Policy Routing Database
RFC1483	PS	Multiprotocol Encapsulation over ATM Adaptation Layer 5
RFC1484	E	Using the OSI Directory to achieve User Friendly Naming (OSI-DS 24 (v1.2))
RFC1485	PS	A String Representation of Distinguished Names (OSI-DS 23 (v5))
RFC1486	E	An Experiment in Remote Printing

RFC1487	PS	X.500 Lightweight Directory Access Protocol
RFC1488	PS	The X.500 String Representation of Standard Attribute Syntaxes
RFC1489	I	Registration of a Cyrillic Character Set
RFC1490	DS	Multiprotocol Interconnect over Frame Relay
RFC1491	I	A Survey of Advanced Usages of X.500
RFC1492	I	An Access Control Protocol, Sometimes Called TACACS
RFC1493	DS	Definitions of Managed Objects for Bridges
RFC1494	PS	Equivalences between 1988 X.400 and RFC-822 Message Bodies
RFC1495	PS	Mapping between X.400 and RFC-822 Message Bodies
RFC1496	PS	Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages
RFC1497	DS	BOOTP Vendor Information Extensions
RFC1498	I	On the Naming and Binding of Network Destinations
RFC1500	S	INTERNET OFFICIAL PROTOCOL STANDARDS
RFC1501	I	OS/2 User Group
RFC1502	PS	X.400 Use of Extended Character Sets
RFC1503	I	Algorithms for Automating Administration in SNMPv2 Managers
RFC1504	I	Appletalk Update-Based Routing Protocol: Enhanced Appletalk Routing
RFC1505	E	Encoding Header Field for Internet Messages
RFC1506	I	A tutorial on gatewaying between X.400 and Internet mail
RFC1507	E	DASS - Distributed Authentication Security Service
RFC1508	PS	Generic Security Service Application Program Interface
RFC1509	PS	Generic Security Service API : C-bindings
RFC1510	PS	The Kerberos Network Authentication Service (V5)
RFC1511	I	Common Authentication Technology Overview
RFC1512	PS	FDDI Management Information Base
RFC1513	PS	Token Ring Extensions to the Remote Network Monitoring MIB
RFC1514	PS	Host Resources MIB
RFC1515	PS	Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs)
RFC1516	DS	Definitions of Managed Objects for IEEE 802.3 Repeater Devices
RFC1517	PS	Applicability Statement for the Implementation of Classless Inter-Domain Routing (CIDR)
RFC1518	PS	An Architecture for IP Address Allocation with CIDR

RFC1519	PS	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
RFC1520	I	Exchanging Routing Information Across Provider Boundaries in the CIDR Environment
RFC1521	DS	MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies
RFC1522	DS	MIME (Multipurpose Internet Mail Extensions) Part Two: Message Header Extensions for Non-ASCII Text
RFC1523	I	The text/enriched MIME Content-type
RFC1524	I	A User Agent Configuration Mechanism For Multimedia Mail Format Information
RFC1525	PS	Definitions of Managed Objects for Source Routing Bridges
RFC1526	I	Assignment of System Identifiers for TUBA/CLNP Hosts
RFC1527	I	What Should We Plan Given the Dilemma of the Network?
RFC1528	E	Principles of Operation for the TPC.INT Subdomain: Remote Printing – Technical Procedures
RFC1529	I	Principles of Operation for the TPC.INT Subdomain: Remote Printing – Administrative Policies
RFC1530	I	Principles of Operation for the TPC.INT Subdomain: General Principles and Policy
RFC1531	PS	Dynamic Host Configuration Protocol
RFC1532	PS	Clarifications and Extensions for the Bootstrap Protocol
RFC1533	PS	DHCP Options and BOOTP Vendor Extensions
RFC1534	PS	Interoperation Between DHCP and BOOTP
RFC1535	I	A Security Problem and Proposed Correction With Widely Deployed DNS Software
RFC1536	I	Common DNS Implementation Errors and Suggested Fixes
RFC1537	I	Common DNS Data File Configuration Error
RFC1538	I	Advanced SNA/IP : A Simple SNA Transport Protocol
RFC1539	I	The Tao of IETF - A Guide for New Attendees of the Internet Engineering Task Force
RFC1540	S	INTERNET OFFICIAL PROTOCOL STANDARDS
RFC1541	PS	Dynamic Host Configuration Protocol
RFC1542	PS	Clarifications and Extensions for the Bootstrap Protocol
RFC1543	I	Instructions to RFC Authors

Final Agenda of the Twenty-Eighth IETF

(November 1-5, 1993)

MONDAY, November 1, 1993

0800-0900	IETF Registration and Continental Breakfast	
	Working Group Chairs Workshop (Dave Crocker/Silicon Graphics)	
0900-0930	Introductions	
0930-1200	Technical Presentations	
	• IP: Next Generation	
Break		
1330-1530	Afternoon Sessions I	
	APP	Internet Message Access Protocol WG **(imap)** (Terry Gray/UWash)
	INT	IP Over Large Public Data Networks WG **(iplpdn)** (Caralyn Brown/Wellfleet)
	IPNG	P. Internet Protocol WG **(pip)** (Paul Francis/Bellcore)
	IPNG	TCP/UDP over CLNP-addressed Networks WG **(tuba)** (Peter Ford/LANL and Mark Knopper/Merit)
	MGT	ATM MIB WG **(atommib)** (Kaj Tesink/Bellcore)
	OPS	Generic Internet Service Description WG **(gisd)** (Tony Bates/RIPE and Danniel Karrenberg/RIPE)
	SEC	Security Area Advisory Group **(saag)** (Steve Crocker/TIS)
	TSV	Multiparty Multimedia Session Control WG **(mmusic)** (Eve Schooler/ISI and Abel Weinrib/Bellcore)
	USV	Whois and Network Information Lookup Service WG **(wnils)** (Joan Gargano/UCDavis)
1530-1600	Break (Refreshments provided)	
1600-1800	Afternoon Sessions II	
	INT	Whither ATM - an Update BOF **(atminfo)** (Mark Laubach/Hewlett-Packard)
	IPNG	Simple Internet Protocol WG **(sip)** (Steve Deering/Xerox PARC and Bob Hinden/Sun)

1600-1800 Monday, November 1, 1993 - Afternoon Sessions II (cont'd.)

MGT	Remote LAN Monitoring WG **(rmonmib)** (Mike Erlinger/Harvey Mudd College)
OPS	Network Joint Management WG **(njm)** (Gene Hastings/PSC) [1]
OPS	Network Status Reports WG **(netstat)** (Gene Hastings/PSC)[1]
SEC	Authorization and Access Control WG **(aac)** (Cliff Neuman/ISI)
USV	Uniform Resource Identifiers WG **(uri)** (Alan Emtage/Bunyip and Jim Fullton/UNC)

[1] NJM and NETSTAT will be meeting in joint session.

TUESDAY, November 2, 1993

0830-0900	Continental Breakfast	
0900-0930	IETF Technical Presentations	
	• "CIDR Status" (Tony Li/cisco)	
0930-1200	Morning Sessions	
	APP	Mime Content BOF **(mimecont)** (John Klensin/UNU)
	APP	TELNET WG **(telnet)** (Steve Alexander/Lachman Technology)
	INT	IP Over Asynchronous Transfer Mode WG **(atm)** (Mark Laubach/Hewlett-Packard)
	MGT	Network Management Area: Open Meeting **(nmarea)** (Marshall T. Rose/DBC)
	RTG	Inter-Domain Multicast Routing WG **(idmr)** (Tony Ballardie/UCL and Paul Francis/Bellcore)
	RTG	ISIS for IP Internets WG **(isis)** (Ross Callon/Wellfleet and Chris Gunner/DEC)
	SAP	Service Location Protocol WG **(svrloc)** (Scott Kaplan/FTP Software and John Veizades/FTP Software)
	SEC	Internet Protocol Security Protocol WG **(ipsec)** (Al Hoover/ANS and Paul Lambert/Motorola)
	USV	Internet School Networking WG **(isn)** (Art St. George/UNMexico and Jennifer Sellers/NASA)
	USV	Uniform Resource Identifiers WG **(uri)** (Alan Emtage/Bunyip and Jim Fullton/UNC)
Break		
1330-1530	Afternoon Sessions I	
	IPNG	Address Lifetime Expections WG **(ale)** (Frank Solensky/FTP Software)
	MGT	Interfaces MIB WG **(ifmib)** (Ted Brunner/Bellcore)
	MGT	SNA DLC Services MIB WG **(snadlc)** (Jeff Hilgeman/Aperatus Technologies)
	RTG	IP Routing for Wireless/Mobile Hosts WG **(mobileip)** (Steve Deering/Xerox PARC and Greg Minshall/Novell)

1330-1530 Tuesday, November 2, 1993 - Afternoon Sessions I (cont'd.)

TSV Multiparty Multimedia Session Control WG **(mmusic)** (Eve Schooler/ISI and Abel Weinrib/Bellcore)

SAP Minimal OSI Upper-Layers WG **(thinosi)** (Peter Furniss/Consultant)

SEC Common Authentication Technology WG **(cat)** (John Linn/OpenVision Technologies)

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

1530-1600 Break (Refreshments provided)

1600-1800 Afternoon Sessions II

APP Internet Message Access Protocol WG **(imap)** (Terry Gray/UWash)

MGT Frame Relay Service MIB WG **(frnetmib)** (James Watt/Newbridge Networks)

OPS Operational Statistics WG **(opstat)** (Phill Gross/ANS and Bernhard Stockman/SUNET)

RTG Routing over Large Clouds BOF **(rolc)** (Joel Halpern/Network Systems)

RTG Source Demand Routing Protocol WG **(sdr)** (Deborah Estrin/USC and Tony Li/cisco)

SAP Service Location Protocol WG **(svrloc)** (Scott Kaplan/FTP Software and John Veizades/FTP Software)

SEC Network Access Server Requirements WG **(nasreq)** (Allan Rubens/Merit and John Vollbrecht/Merit)

TSV Audio/Video Transport WG **(avt)** (Steve Casner/ISI)

USV Integration of Internet Information Resources WG **(iiir)** (Kevin Gamiel/CNIDR and Chris Weider/Merit)

1930-2200 Evening Sessions

GEN Open IAB Meeting

INT Dynmaic Host Configuration WG **(dhc)** (Ralph Droms/Bucknell)

WEDNESDAY, November 3, 1993

0800-0900	Working Group Chairs Workshop (Dave Crocker/Silicon Graphics)
0830-0900	Continental Breakfast
0900-0930	Technical Presentations • ST-II (Craig Partridge/BBN)
0930-1200	Morning Sessions

APP	X.400 Operations WG **(x400ops)** (Alf Hansen/UNINETT and Tony Genovese/LLNL)
INT	Point-to-Point Protocol Extensions WG **(pppext)** (Fred Baker/ACC)
INT	Internet Stream Protocol V2 WG **(st2)** (Steve DeJarnett/IBM and Luca Delgrossi/IBM)
MGT	Modem Management WG **(modemmgt)** (Mark Lewis/Telebit)
RTG	Inter-Domain Multicast Routing WG **(idmr)** (Tony Ballardie/UCL and Paul Francis/Bellcore)
RTG	Inter-Domain Policy Routing WG **(idpr)** (Martha Steenstrup/BBN)
SEC	Internet Protocol Security Protocol WG **(ipsec)** (Al Hoover/ANS and Paul Lambert/Motorola)
TSV	TCP Multiplexing BOF **(tmux)** (Jim Barnes/Xylogics)
USV	User Services WG **(uswg)** (Joyce K. Reynolds/ISI)

Break

1330-1530	Afternoon Sessions I

APP	Telnet TN3270 Enhancements WG **(tn3270e)** (Robert Moskowitz/Chrysler Corporation)
APP	X.400 Operations WG **(x400ops)** (Alf Hansen/UNINETT and Tony Genovese/LLNL)
IPNG	TCP/UDP over CLNP-addressed Networks WG **(tuba)** (Peter Ford/LANL and Mark Knopper/Merit)[2]

[2]TUBA and TP/IX will be meeting in joint session

1330-1530 W...

THURSDAY, November 4, 1993

0830-0900 Continental Breakfast

0900-0930 Technical Presentations

- "Using E-mail in Europe: Opportunities and Challenges" (Harald Alvestrand/UNINETT)

0930-1200 Morning Sessions

APP	Applications Area and Area Directorate **(apples)**: Open (Erik Huizer/SURFnet and John Klensin/UNU)[4]
APP	Integrated Information Architecture **(iia)** (Erik Huizer/SURFnet, John Klensin/UNU and Joyce K. Reynolds/ISI)[4]
INT	IP Over Asynchronous Transfer Mode WG **(atm)** (Mark Laubach/Hewlett-Packard)
INT	Point-to-Point Protocol Extensions WG **(pppext)** (Fred Baker/ACC)
MGT	Remote LAN Monitoring WG **(rmonmib)** (Mike Erlinger/Harvey Mudd College)
OPS	Global Printing on Facsimile Devices BOF **(tpcint)** (Carl Malamud/Internet Multicasting and Marshall T. Rose/DBC)
RTG	Border Gateway Protocol WG **(bgp)** (Yakov Rekhter/I...
RTG	New Internet Routing and Addressing Architecture BOF **(nimrod)** (Noel Chiappa and Isidro Castineyra/BBN)
RTG	OSI IDRP for IP over IP WG **(ipidrp)** (Sue Hares
SAP	Domain Name System WG **(dns)** (Rob Austein/
USV	User Documents WG **(userdoc2)** (Ellen Hoffm... and Lenore Jackson/NASA)

[4] Apples and IIA will meet consecutively, from 0930-1045 and 1045-1200 re...
...DRP will be meeting in joint session.

Break

1330-1530 Thursday, November 4, 1993 - Afternoon Sessions I

APP	OSI Directory Services WG **(osids)** (Steve Kille/ISODE)
INT	Internet Stream Protocol V2 WG **(st2)** (Steve DeJarnett/IBM and Luca Delgrossi/IBM)
IPNG	Simple Internet Protocol WG **(sip)** (Steve Deering/Xerox PARC and Bob Hinden/Sun)
MGT	ATM MIB WG **(atommib)** (Kaj Tesink/Bellcore)
OPS	Benchmarking Methodology WG **(bmwg)** (Scott Bradner/Harvard)
RTG	New Internet Routing and Addressing Architecture BOF **(nimrod)** (Noel Chiappa and Isidro Castineyra/BBN)
RTG	RIP Version II WG **(ripv2)** (Gary Malkin/Xylogics)
SAP	Mail-based File Distribution BOF **(mailftp)** (Marko Kaittola/FUNET and Urs Eppenberer/SWITCH)
SEC	Security Area Advisory Group **(saag)** (Steve Crocker/TIS)
USV	Network Training Materials WG **(trainmat)** (Jill Foster/UNewcastle-Upon-Tyne)
USV	Uniform Resource Identifiers WG **(uri)** (Alan Emtage/Bunyip and Jim Fullton/UNC)

1530-1600 Break (Refreshments provided)

1600-1700 Technical Presentations

- "Federal Internetworking Requirements Panel (FIRP) (Richard desJardins/NASA)
- "Intellectual Property Rights" (Vint Cerf/CNRI)
- "Draft ISO/MOU" (Vint Cerf/CNRI)

1700-1930 Open Plenary and IESG

FRIDAY, November 5, 1993

0830-0900	Continental Breakfast	
0900-1200	Morning Sessions	
	APP	OSI Directory Services WG **(osids)** (Steve Kille/ISODE)
	MGT	Uninterruptible Power Supply WG **(upsmib)** (Jeff Case/UTenn)
	OPS	Network Joint Management WG **(njm)** (Gene Hastings/PSC)[6]
	OPS	Network Status Reports WG **(netstat)** (Gene Hastings/PSC)[6]
	SAP	Domain Name System WG **(dns)** (Rob Austein/Epilogue)

Key to Abbreviations

APP	Applications	Erik Huizer/SURFnet and John Klensin/UNU
GEN	General Interest	
INT	Internet	Stev Knowles/FTP Software and Dave Piscitello/Bellcore
IPNG	IP: Next Generation	Scott Bradner/Harvard and Allison Mankin/NRL
MGT	Network Management	Marshall T. Rose/DBC
OPS	Operational Requirements	Scott Bradner/Harvard
RTG	Routing	Bob Hinden/Sun
SAP	Service Applications	Dave Crocker/SGI
SEC	Security	Steve Crocker/TIS
TSV	Transport	Allison Mankin/NRL
USV	User Services	Joyce K. Reynolds/ISI

[6]NJM and NETSTAT will be meeting in joint session.

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 then US Defense Advanced Projects Agency (DARPA), working on the ARPANET, US 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. Facilitating technology transfer from the Internet Research Task Force (IRTF) to the wider Internet community; and

4. Providing a forum for the exchange of relevant 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 ten 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, the 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
	John Klensin/UNU
Internet	Stev Knowles/FTP Software
	Dave Piscitello/Bellcore
IP: Next Generation	Scott Bradner/Harvard
	Allison Mankin/NRL
Network Management	Marshall Rose/DBC
Operational Requirements	Scott Bradner/Harvard
Routing	Robert Hinden/Sun
Security	Steve Crocker/TIS
Service Applications	Dave Crocker/SGI
Transport	Allison Mankin/NRL
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	John Stewart
IETF Meeting Coordinator	Megan Davies Walnut
IETF Meeting Registrar	Debra Legare
IETF Internet-Drafts Administrator	Cynthia Clark
IETF Administrative Support	Lois Keiper

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 meetings 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 the IETF announcement list, send a request to:

`ietf-announce-request@cnri.reston.va.us`

To join the IETF general discussion list, send a request to:

`ietf-request@cnri.reston.va.us`

To join other mailing lists, 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

Spring 1994

Seattle, Washington
NorthWestNet and
The University of Washington
Host(s): Dan Jordt and Terry Gray
March 28 - April 1, 1994
Status: CONFIRMED

Summer 1994

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

Fall 1994

San Francisco Bay Area
Sun Microsystems, Inc.
Host: Bob Hinden
Possible Dates: Nov. 14-18, 1994
Dec. 5-9, 1994
Status: TENTATIVE

1.2 On-Line IETF Information

The Internet Engineering Task Force maintains up-to-date, on-line information on all of its activities. This information is available via FTP and e-mail. 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", RFC 1539.
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 e-mail 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 (e.g., 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 may eventually be submitted to the IESG and/or the RFC Editor to be considered for publication as RFCs. 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 e-mail 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

e.g.,

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.

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

US West Coast Address: venera.isi.edu (128.9.0.32)

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 e-mail 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
```

E-mail Access

Internet-Drafts are available by mail server from ds.internic.net. To retrieve a file, mail a request:

To: mailserv@ds.internic.net
Subject: Anything you want

In the body, put a command of the form:

FILE /internet-drafts/1id-abstracts.txt
FILE /ietf/1wg-summary.txt
FILE /ietf/822ext/822ext-minutes-91jul.txt
PATH jdoe@somedomain.edu

where PATH lists the e-mail address where the response should be sent.

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 may 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 RFCs and are NOT a numbered document series. The string "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, venera.isi.edu, 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 1id-abstracts.txt index and in the announcement of the Internet-Draft. The abstract should follow the "Status of this Memo" section.

All Internet-Drafts should contain the full filename (beginning with draft- and including the version number) in the text of the document. The filename information should, at a minimum, appear on the first page (possibly with the title).

For those authors submitting updates to existing Internet-Drafts, the choice of the file name is easily determined (increase the version by 1). For new documents, send a message to `internet-drafts@cnri.reston.va.us` with the document title, if it is a product of a working group (and the name of the group), and an abstract. The filename to be assigned will be included in a response. Simply add the filename text to the document (ASCII AND PostScript versions) and submit the Internet-Draft.

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

Directors:

- Erik Huizer: erik.huizer@surfnet.nl
- John Klensin: klensin@infoods.unu.edu

Area Summary reported by Erik Huizer/SURFnet

This is a short report on the Applications Area, with respect to the Houston IETF meeting November 1993.

The Applications Area currently contains the following working groups:

- Internet Message Access Protocol (IMAP)
- OSI Directory Services (OSIDS)
- TELNET (TELNET)
- TELNET TN3270 Enhancements (TN3270E)
- X.400 Operations (X400OPS)

In addition, the Applications Area and the User Services Area jointly oversee the following working groups:

- Integrated Directory Services (IDS)
- Integration of Internet Information Resources (IIIR)
- Internet Anonymous FTP Archives (IAFA)
- Networked Information Retrieval (NIR)
- Uniform Resource Identifiers (URI)
- Whois and Network Information Lookup Service (WNILS)

The status of these groups is described in the User Services Area report.

The Internet Message Extensions (822EXT), MIME-MHS Interworking (MIMEMHS), and Network News Transport Protocol (NNTP) Working Groups have disbanded since the last meeting.

An open meeting of the Applications Area Directorate (APPLES) as well as a MIME Content BOF (MIMECONT) were held in Houston.

Applications Area Directorate (APPLES)

The goal of the meeting was to present an overview of the applications work that is going on in the Applications and User Services area, possibly identify areas of common interest/overlap and discuss possible coordination.

MIME Content BOF (MIMECONT)

The reviews covered:

- Handling of SGML files over MIME

  There was a lengthy discussion and a new proposal will be prepared that reflects the comments.

- Structuring, beyond the "mixed," "alternative," "digest," and "parallel" constructions of RFC 1521, of Multipart MIME messages

  A new proposal is being prepared.

- Models for attribute-value (or name-value) pairs over MIME, such as for personal contact information

  Discussions will continue using the 822ext mailing list. Formation of a working group in this area is likely.

- Mail delivery reports

  There have been several proposals for specific formats for automatically-generated reports about mail delivery or non-delivery. The review concluded that a working group was needed in this area.

- Language directionality

  The group reviewed a proposal for specifying the relationship between presentation order (e.g., on a screen) and characters in the data stream for languages whose characters were written other than left-to-right. The conclusion was that this capability should not be added to text/plain, but should either use a different "text" subtype or that the information needed should be identified by multiple special character set names.

- Macintosh files over MIME

  The group reviewed the new proposals. Some tuning is still needed. A new draft will be produced and reviewed via an extended Last Call.

Internet Message Access Protocol Working Group (IMAP)

A total of nineteen agenda items were considered. Considerable progress was made on all fronts. One notable result: on Monday the group agreed that the acronym "IMAP" should be remapped to the words "Internet Message Access Protocol" to better reflect what the protocol has evolved into.

OSI Directory Services Working Group (OSIDS)

The OSIDS Working Group will disband and their work items will be partitioned among a number of new groups:

- Schema and naming (Sri Sataluri)
- Lightweight protocols (Tim Howes)
- Indexing DSAs and centroids (Simon Spero)
- IP representation in X.500 (Glenn Mannsfield)

Detailed proposed charters will be submitted by the proposed chairs and discussed in a wider directory forum (e.g. with WHOIS++ included). The OSIDS mailing list will remain for discussion of umbrella X.500 issues.

Document status:

- The CLDAP (Connectionless LDAP) document will be submitted as a Proposed Standard.

- The RFC 1384 update will be submitted as a Proposed Standard.

- Two Internet-Drafts on representing IP information in the DIT will be submitted for approval as Experimental RFCs.

- The "Schema" subgroup is established, and will submit various documents to replace RFC 1274.

TELNET Working Group (TELNET)

Unless the charter is revised, the group will conclude after the environment and authentication documents are final. Much of the meeting focused on discussion of the environment Internet-Draft, "Telnet Environment Option" and the authentication Internet-Draft, "Telnet Authentication and Encryption Option."

Sam Sjogren raised the issue of interoperability testing. The group was receptive, and may try to schedule an event prior to the Seattle meeting.

TELNET TN3270 Enhancements Working Group (TN3270E)

The current-practices Internet-Draft, "TN3270 Current Practices," was agreed on, but minor typographical errors were found. It will be reposted, and at that time should be forwarded for consideration as an Informational RFC.

The LUnames-printer Internet-Draft, "TN3270 Extensions for LUname and Printer Selection," was discussed. The two outstanding issues were resolved and an new draft will be created in a couple of weeks. That should go through a quick internal review. At that point, it should be reviewed by key members of the TELNET Working Group and then forwarded for consideration as an RFC. The question is: should this document go in as an Informational RFC or as a Proposed Standard which gets changed to Informational when the "TN3270 Enhancements" Internet-Draft gets published?

The "TN3270 Enhancements" Internet-Draft went through extensive discussion. All of the known issues were covered and an approach for each was devised. A new Internet-Draft will be created for review by the working group members. It is expected that a consensus can be reached on this document by the end of the year to be published as a Proposed Standard.

X.400 Operations Working Group (X400OPS)

- Allan Cargille's Internet-Draft, "Postmaster Convention for X.400 Operations," will undergo minor editorial changes scheduled for November 8.

- Alf Hansen's Internet-Draft, "Operational Requirements for X.400 Management Domains in the GO-MHS Community," will undergo one more editorial pass with the final version scheduled for November 15.

- Claudio Allocchio gave a status report on the DNS mapping table experiment. When the working group concludes, this work will be transferred to the RARE Working Group MSG.

- The group reviewed the draft CXII Charter and developed a workplan for addressing this as a new IETF working group.

- ADMD=IMX was presented by Allan Cargille. Erik Huizer reiterated that this is a User Services issue that is inappropriate for this IETF.

- Erik also reported that the establishment of an "IOTF" was agreed in principle by the IESG and the IAB.

- With the completion of the above documents, the working group has completed all its goals and will conclude. The mailing list for the group will be kept active to work new items as they may come up.

CURRENT MEETING REPORT

Reported by Erik Huizer/SURFnet bv

Minutes of the Applications Area Directorate (APPLES)

Thanks to Keith Moore for his contribution to the minutes.

The meeting was chaired by John Klensin, Erik Huizer and Joyce Reynolds. The goal of the meeting was to present an overview of the applications work that is going on in the Applications Area and User Services Area, possibly identify areas of common interest/overlap and discuss possible coordination.

The various working group chairs gave a very brief description of their applications-related working group and the current issues.

- Integrated Directory Services (IDS) - User Services Area
  - non-technical documentation for directory services users
  - X.500 catalogs
  - WHOIS++ catalogs
  - gateways between X.500/WHOIS++
  - legal issues with running directory services
- Integration of Internet Information Resources (IIIR) - User Services Area
  - vision of the future
  - quality of service issues
  - standardizing types for applications/data
  - documenting existing protocols like WWW, Gopher and WAIS
- Internet Message Access Protocol (IMAP) - Applications Area
  - manipulate remote mailbox
  - access sequences of messages (news, FTP archives, etc.)
  - selectively access MIME body parts
- Internet Anonymous FTP Archives (IAFA) - User Services Area
  - help anonymous FTP administrators organize FTP sites
  - FTP site administrators guide (FYI)
  - FTP guide for users (FYI)
  - set of data templates to catalog FTPable information
- Networked Information Retrieval (NIR) - User Services Area
  - status report on NIR tools/groups

- OSI Directory Services (OSIDS) - Applications Area
  - deployment of X.500-based directory services on the Internet
  - LDAP - lightweight directory access protocol
  - CLDAP - connectionless lightweight directory access protocol
  - representing bibliographic/text information

- TELNET (TELNET) - Applications Area
  - environment option
  - authentication

- TELNET TN3270 Enhancements (TN3270E) - Applications Area
  - document TN3270 (Informational RFC)
  - extensions: LUnames/printer data streams
  - enhanced TN3270: full function 3270 over TCP

- Uniform Resource Identifiers (URI) - User Services Area
  - URLs
  - URNs
  - citations/characteristics (whatever URCs are)
  - naming/address resolution for objects (possibly DNS)

- WHOIS and Network Information Lookup Service (WNILS) - User Services Area
  - WHOIS network information and lookup service
  - development of WHOIS++

- X.400 Operations (X400OPS) - Applications Area
  - closing down
  - proposes new working group: Commercial X.400 Internet Interconnection (CXII)
    - how do commercial providers of X.400 connect to Internet mail?
      - go-mhs
      - RFC 1327 gateway operations
    - must get collaboration from X.400 service providers

After these presentations the discussion went on to define sub-areas of related working groups. A discussion ensued that concentrated on the topic of white pages and directory services. Erik Huizer proposed a strawman model of how efforts can be coordinated in this sub-area. The strawman proposes two things:

1. Coordination of WHOIS++ and X.500 working groups working on specific well defined issues, that cover the "common ground" between these two otherwise competing proposals.

2. The establishment of a directorate for directory service issues that would act as an advisory board to the area directors, and would help to coordinate the related working groups.

Discussion

- It was perceived that coordination of WHOIS++ and X.500 efforts in the white pages area is useful. However, the ultimate goal of this is *not* to come up with the *one* protocol to solve the directory problem. Rather, the goal should be to pool expertise to solve common problems, thus increasing the level of interworking and the perceived functionality for the end-users.

  It was pointed out that this effort should not preclude other directory services efforts from working within the Applications Area and/or User Services Area.

  This will be discussed further in the OSIDS and WNILS Working Groups.

- The establishment of a separate directorate for directory services in the Applications Area was generally not perceived as a good idea. Fear was expressed that this tended to solve the problems for the directory service-related working groups, but would not improve communication with those working groups that work on protocols that make use of the directory. Suggestions were made for enlargement of the Applications Area Directorate, at least with the working group chairs, and possibly with support of some pools of experts for some fields of interest.

  The Area Directors will write another strawman proposal for this specific coordination issue, and it will be discussed on the list: `apples@surfnet.nl`. To subscribe to the list send a request to: `apples-request@surfnet.nl`.

  After consensus is reached, the proposal will be submitted to the IESG and IAB for approval.

Attendees

Steve Alexander	stevea@lachman.com
Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Glen Cairns	cairns@mprgate.mpr.ca
James Conklin	jbc@bitnic.educom.edu
Ann Cooper	cooper@isi.edu
Alan Emtage	bajan@bunyip.com
Urs Eppenberger	eppenberger@switch.ch
Erik Fair	fair@apple.com
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com

Kevin Gamiel	kgamiel@cnidr.org
Tony Genovese	genovese@es.net
Arlene Getchell	getchell@es.net
Shawn Gillam	shawn@timonware.com
Mei-Jean Goh	goh@mpr.ca
Judith Grass	grass@cnri.reston.va.us
Terry Gray	gray@cac.washington.edu
Phillip Gross	pgross@ans.net
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Deborah Hamilton	debbieh@internic.net
Alf Hansen	Alf.Hansen@uninett.no
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Marco Hernandez	marco@cren.net
Russ Hobby	rdhobby@ucdavis.edu
Tim Howes	tim@umich.edu
Richard Huber	rvh@ds.internic.net
Erik Huizer	Erik.Huizer@SURFnet.nl
Marko Kaittola	Marko.Kaittola@funet.fi
Neil Katin	katin@eng.sun.com
Steve Kille	S.Kille@isode.com
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
John Klensin	Klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Mark Kosters	markk@internic.net
John Kunze	jak@violet.berkeley.edu
Barry Leiner	leiner@nsipo.nasa.gov
Ben Levy	seven@ftp.com
Clifford Lynch	calur@uccmvsa.ucop.edu
Glenn Mansfield	glenn@aic.co.jp
April Marine	april@atlas.arc.nasa.gov
Larry Masinter	masinter@parc.xerox.com
Mitra	mitra@pandora.sf.ca.us
Keith Moore	moore@cs.utk.edu
Robert Moskowitz	3858921@mcimail.com
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Rakesh Patel	rapatel@pilot.njin.net
Marsha Perrott	perrott@prep.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Cecilia Preston	cpreston@info.berkeley.edu
Kenneth Rossen	kenr@shl.com
Srinivas Sataluri	sri@internic.net

Rickard Schoultz	schoultz@sunet.se
Henry Sinnreich	hsinnreich@mcimail.com
Mark Smith	mcs@umich.edu
Karen Sollins	sollins@lcs.mit.edu
Simon Spero	ses@unc.edu
Anthony Valle	valle@huntsville.sparta.com
John Veizades	veizades@ftp.com
Chris Weider	clw@merit.edu
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov
Russ Wright	wright@lbl.gov

CURRENT MEETING REPORT

Reported by Luc Boulianne/McGill University

Minutes of the Integrated Information Architecture BOF (IIA)

Introduction - Phill Gross

Phill Gross took a few moments to introduce the reasons behind the creation of this group:

> A year or so ago, the IESG was hit all at once with the creation of a large set of working groups in the general areas of network information discovery, retrieval, and user information handling. It became apparent that many of these working groups were related, or should be. There appears to be two ways in which the IETF operates: top-down and bottom-up.
>
> 1. Top-down: (or pro-active) such as the IPNG.
> 2. Bottom-up: the usual way things are done. Usually the 'right' things come out of this approach. And yet, it would appear that sometimes, the area directors are still needed for pro-active planning.
>
> When the working groups were chartered, they were made jointly part of an "Internet Information Architecture" (IIA) activity. The expectation was that these groups would work together, as well as on their own primary foci, and would do so under the joint supervision of the User Services and Applications Areas.

Phill suggested that the area directors now write a new overview of the IIA, providing a framework only. Because of the importance of this issue, Phill suggested that the IESG request a working group be charged to create an IIA architecture framework definition citing as an example: IPNG (Allison Mankin and Scott Bradner).

Summary of the Issues - John Klensin

Working groups were formed, work was done and documents began to appear. Some concluded that there was a lack of coordination among the working groups, but that the current meeting is an effort to reconcile this lack.

IIA is comprised of several working groups, overlapping work with an overlapping cast of characters. The working groups should be coordinated technically, but it often appears that they are not.

Characteristics of the IIA working groups and their membership:

- Very expert in several types of work.
- There is, however, some evidence that, in protocol design areas, they may be moving out of their depth or succumbing to the Not Invented Here syndrome.
- There are interactions with the "real world" that one must consider, e.g., librarians and other information specialists, external standards.
- Most of the groups seem to have nearly the same membership, with topics and issues flowing back and forth between them.

Finally, there were these questions to ask the group:

- Is the current model as efficient as possible? If it is not, what can be done to improve things?
- Is there a structural way of going about this?
- What about working group functional boundaries?
- What is the definition of a functional boundary?
- What can be done to not break anything that is now working, while we try to increase efficiency and productivity?

A suggestion was made by the group that multiple solutions to this problem (i.e. working groups) which have trivial differences, should be merged into a homogeneous solution. This would help to avoid diluting the merged efforts.

User Services Area - Joyce K. Reynolds

Joyce believes that it is important to make sure there is communication between areas. A meeting of the User Services Area Council (USAC) was held on Tuesday evening. USAC observed that developers and users are well represented in these gatherings, but operators (information providers) are not. The following items are lacking:

- Tools for maintaining information
- Support tools
- How does one share information
- An adequate level of cooperation
- An adequate level of operational effectiveness

General Discussion

Many suggestions for better communication and exchange of information among the working groups were presented and discussed.

Editor's Note: An itemized list of suggestions and discussion topics is available via FTP or mail server from the remote directories as /ietf/93nov/iia-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Conclusion

Using the APPLES mailing list for further discussion, it should be determined 1) how to improve communications, and 2) what structure might work to propagate this newly acquired information.

Attendees

Steve Alexander	stevea@lachman.com
Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Jules Aronson	aronson@nlm.nih.gov
Luc Boulianne	lucb@cs.mcgill.ca
James Conklin	jbc@bitnic.educom.edu
John Curran	jcurran@nic.near.net
Alan Emtage	bajan@bunyip.com
Urs Eppenberger	eppenberger@switch.ch
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Paul Francis	Francis@thumper.bellcore.com
Ned Freed	ned@innosoft.com
Kevin Gamiel	kgamiel@cnidr.org
Tony Genovese	genovese@es.net
Judith Grass	grass@cnri.reston.va.us
Terry Gray	gray@cac.washington.edu
Phillip Gross	pgross@ans.net
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Deborah Hamilton	debbieh@internic.net
Alf Hansen	Alf.Hansen@uninett.no
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Marco Hernandez	marco@cren.net
Russ Hobby	rdhobby@ucdavis.edu
Jeroen Houttuin	houttuin@rare.nl
Tim Howes	tim@umich.edu

Richard Huber	rvh@ds.internic.net
Christian Huitema	Christian.Huitema@sophia.inria.fr
Erik Huizer	Erik.Huizer@SURFnet.nl
Steve Kille	S.Kille@isode.com
John Klensin	Klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Barry Leiner	leiner@nsipo.nasa.gov
Ben Levy	seven@ftp.com
Clifford Lynch	calur@uccmvsa.ucop.edu
Glenn Mansfield	glenn@aic.co.jp
April Marine	april@atlas.arc.nasa.gov
Larry Masinter	masinter@parc.xerox.com
Mitra	mitra@pandora.sf.ca.us
Keith Moore	moore@cs.utk.edu
Robert Moskowitz	3858921@mcimail.com
Chris Newman	chrisn+@cmu.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Rakesh Patel	rapatel@pilot.njin.net
Pete Percival	percival@indiana.edu
Marsha Perrott	perrott@prep.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Cecilia Preston	cpreston@info.berkeley.edu
Joyce K. Reynolds	jkrey@isi.edu
Srinivas Sataluri	sri@internic.net
Richard Schmalgemeier	rgs@merit.edu
Rickard Schoultz	schoultz@sunet.se
Henry Sinnreich	hsinnreich@mcimail.com
Mark Smith	mcs@umich.edu
Karen Sollins	sollins@lcs.mit.edu
Chris Weider	clw@bunyip.com
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov

CURRENT MEETING REPORT

Reported by John Klensin/United Nations University

Minutes of the Mime Content BOF (MIMECONT)

Because the following reviews are informal, there are not, in general, topic-specific mailing lists. However, the "822ext" list, available for MIME implementation issues, is the generic location for discussions of these types until topic-specific lists are spun off.

General discussion:	ietf-822@dimacs.rutgers.edu
To subscribe:	ietf-822-request@dimacs.rutgers.edu

Background

The MIME RFC (RFC 1521) specifies that anyone can register a content subtype under one of the major types simply by supplying a name and specification to IANA. However, there are cases where something is important enough to justify special review or when there appears to be an opportunity to draw competing proposals together and avoid the interoperability problems that would otherwise arise from differently profiled MIME applications. Rather than charter a working group for each topic, or create a standing working group that would review all such proposals (but probably have expertise specific to few of them), the Applications Area Directors have established an ad hoc review mechanism by which interested people can discuss the proposals and recommend to the area directors what, if any, further action should be taken.

Six of these reviews occurred during this IETF. In several cases, the fact of scheduling the reviews and asking people to be prepared to present and discuss their proposals produced significant convergence, and the reviews were devoted to overviews and discussion of the new proposals that were emerging from that process.

While the reviews appeared together in the agenda, and are consolidated in these minutes, it is important to understand that they are independent events and activities.

Full SGML Over MIME and SGML Introduction

Current document:

- draft-levinson-sgml-00.txt

Supplemental tutorial on SGML:

- SGML Tutorial that appears as Part 1 of 'Guidelines for Electronic Text Encoding and Interchange,' draft version 2 of document TEI P2 from the Text Encoding Initiative (TEI), edited by C. M. Sperberg-McQueen and Lou Burnard

This review dealt with moving general SGML document files over MIME as distinguished from specially-profiled files that use SGML syntax and concepts (see the review on "Structured Information and Personal Contact Information"). SGML was defined and its important characteristics identified. The group discussed the relationship between SGML external references and various existing, and proposed, Internet mechanisms including MIME external bodies and Content-IDs and the various URIs. Another issue was the SGML, like Postscript permits embedding executable structures (normally used to interpret graphics) and raw device commands that could create significant security risks. Since these are implementation- and site-dependent, the group concluded that it would be sensible to significantly restrict their use.

There was also some consensus that the present document should be modified to utilize more general mechanisms for aggregating files within a MIME message, rather than inventing its own. This would leverage existing mechanisms for cross-references, external documents, and so on.

Conclusion: No new working group is needed, at least at present. Discussion will continue on the 822 list. Ed Levinson will reissue the document with changes suggested in this session and additional discussions.

MIME Multipart Structuring: Header-Sets and References

Current documents:

- draft-crocker-headerset-00.txt, .ps
- draft-moore-mime-reference-00.txt
  Major consolidation and revision pending, see below.

Many people have observed that many different multipart types – beyond the "mixed," "alternative," "digest," and "parallel" constructions specified in RFC 1521 – are being specified for MIME. Many of these have most of their features in common, and a generic strategy would ease implementations, simplify design of future ones, and possibly reduce burdens on the registration process. Preparation for this review resulted in the combination of several alternatives into a new proposal, which has not yet been written up.

A new proposal for multipart "families" is being prepared to define multipart as a general facility and provide guidance for handling aggregate objects. One important aspect of this work will be to specify how gateways to non-MIME environments should behave when these types are encountered.

Conclusion: The headerset proposal is to be dropped. The multipart/alternative one will be dropped until and unless applications appear that actually require that level of generality and complexity. The new "families" proposal will be written up and discussed, then we will review what to do with it, since it is likely to be rather more a set of guidelines than an actual protocol specification.

Structured Information and Personal Contact Information

Current documents:

- draft-crocker-stif-00.txt, .ps
- draft-crocker-pci-00.txt, .ps
- draft-adie-shave-00.txt, .ps
- draft-adie-spci-00.txt, .ps
- draft-vaudreuil-mime-sig-00.txt

Many situations need structured attribute (or name)/value pairs within MIME messages and in other applications. Personal contact information, such as one might find on business cards or in a Rolodex are among the often-cited examples.

Several people have independently tried to develop standard ways to represent this type of information. The two major proposals to do this are based on an extension of the RFC 822 header field model to accommodate nested structures (STIF) and a profile and set of definitions for using SGML to accomplish the same purpose (SHAVE). The 822-like format (at least without the extensions) is very familiar in the Internet community and feels quite natural. The SGML one feels natural to communities that have been using SGML and provides solutions to problems that still must be worked out with STIF, but SGML is not familiar to most of the IETF community and looks foreign and complex to a major subset of it. STIF is certainly easier to read for simple cases; but SHAVE might be easier in very complex ones.

The familiarity with STIF-like arrangements, the installed base of data embedded in SGML formats, and the availability of a formal, executable definition against which validity can be determined, are all important considerations. It is important to note that SHAVE, unlike the general SGML model of the "Full SGML Over MIME and SGML Introduction" review, does not contemplate sending SGML Document Type Definitions (DTDs) around: at most one DTD would be defined per application, and processing the application would just imply applying it. This is similar to having a definition of a set of fields for use with STIF.

STIF will not be registered as a content subtype. It is really a framework for constructing such subtypes. SHAVE could go either way, either as such a framework, or in a model that might use, e.g., content-type: text/SHAVE; dtd=SPCI

Conclusion: Discussions will continue using the 822ext mailing list. It is not clear whether a separate signature subtype is needed or desirable, or whether signatures should just be handled as a special case of personal contact information under either the STIF or SHAVE models. Formation of a working group in this area is likely.

Mail Delivery Reports and Notifications

Current documents:

- draft-moore-mime-delivery-00.txt
- draft-moore-smtp-drpt-00.txt
- draft-vaudreuil-mime-delivery-00.txt

There have been several proposals for specific formats for automatically-generated reports about mail delivery or non-delivery. Getting such notices require a model for requesting them that probably must be handled as an SMTP extension, but a standardized format for sending them would greatly facilitate automated processing and building of intelligent user agents.

The two report format proposals differ in level of generality and the problems addressed. All of the problems appear to be important, and a new proposal is needed that would address them.

Conclusion: These reports, and the request mechanism, must be on the standards track to be useful. A working group is needed that will focus on both the report formats and the needed SMTP extensions, probably in that order. Keith Moore and Greg Vaudreuil will start a mailing list, announce it to the 822ext list, and begin to develop a working group charter.

Specification of Presentation Direction for Text/Plain and Languages Whose Natural Order is Not Left-to-Right

Current document:

- draft-nussbacher-mime-direction-01.txt

The group reviewed a proposal for specifying the relationship between presentation order (e.g., on a screen) and characters in the data stream for languages whose characters were written other than left-to-right. The proposal was to handle this by adding an extra parameter to Content-type: text/plain; charset=xxx that would specify the "directionality" of the characters with keywords drawn from applicable ECMA and ISO standards.

While there was some sense that this would have been the right thing to do had it be thought of earlier in MIME's development, the consensus of those present was that it was not possible to add a parameter to text/plain at this time: some implementations might ignore it, others might actually get into trouble.

Two alternate suggestions were made:

1. Extend the character set names to include the directionality, e.g., Content-type: text/plain; charset=iso-8859-8-visual

2. Use a completely different text subtype, e.g., either:

   Content-type: text/directional; charset=iso-8859-9; direction=implicit

   or

   Content-type: text/plain-explicit; charset=iso-8859-9

Macintosh File Transmission With MIME

Current documents:

- draft-faltstrom-macmime1-00.txt
- draft-faltstrom-macmime2-00.txt

There have been discussions in various forums for a year or more about how to best send Macintosh files over MIME. The Internet-Drafts listed above represent consensus among most of the contenders.

The group reviewed them and concluded that, while some tuning is still needed, the concepts are basically sound. The importance of these formats is such that they should be placed on the standards track. A new draft will be produced and reviewed via an extended Last Call.

Attendees

George Chang	`gkc@ctt.bellcore.com`
James Conklin	`jbc@bitnic.educom.edu`
David Crocker	`dcrocker@mordor.stanford.edu`
James Davin	`davin@thumper.bellcore.com`
Donald Eastlake	`dee@skidrow.lkg.dec.com`
Ed Ellesson	`ellesson@vnet.ibm.com`

Urs Eppenberger	eppenberger@switch.ch
Patrik Faltstrom	paf@nada.kth.se
Qin Fang	qin_fang@unc.edu
James Fielding	jamesf@arl.army.mil
Ned Freed	ned@innosoft.com
Tony Genovese	genovese@es.net
Shawn Gillam	shawn@timonware.com
Erik Huizer	Erik.Huizer@SURFnet.nl
Neil Katin	katin@eng.sun.com
John Klensin	Klensin@infoods.unu.edu
Edward Levinson	elevinson@accurate.com
Lars-Johan Liman	liman@ebone.net
Wayne McDilda	wayne@dir.texas.gov
Keith Moore	moore@cs.utk.edu
Robert Moskowitz	3858921@mcimail.com
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Masataka Ohta	mohta@cc.titech.ac.jp
Henry Sinnreich	hsinnreich@mcimail.com
Simon Spero	ses@unc.edu
Rick Troth	troth@rice.edu
Gregory Vaudreuil	gvaudre@cnri.reston.va.us

2.1.1 Internet Message Access Protocol (IMAP)

Charter

Chair(s)

Terry Gray: gray@cac.washington.edu

Mailing Lists

General Discussion: imap@cac.washington.edu
To Subscribe: imap-request@cac.washington.edu
Archive: ftp.cac.washington.edu:~/imap/imap_archive

Description of Working Group

The Interactive Mail Access Protocol (IMAP) Working Group is chartered to refine and extend the current IMAP2 protocol as a candidate standard for a client-server Internet email protocol to manipulate remote mailboxes as if they were local. An explicit objective is to retain compatibility with the growing installed base of IMAP2-compliant software. It is expected that the resulting specification will replace both RFC 1176 and the more recent (as yet unplublished) IMAP2bis extensions document.

The IMAP Working Group will also investigate how to provide for "disconnected operation" capabilities similar to the DMSP protocol (RFC 1056, with Informational Status) with a goal of making it possible for IMAP to replace DMSP.

An email access protocol provides a uniform, operating system-independent way of manipulating message data (email or bulletin board) on a remote message store (repository). Mail user agents implementing such a protocol can provide individuals with a consistent view of the message store, regardless of what type of computer they are using, and regardless of where they are connected in the network. Multiple concurrent sessions accessing a single remote mailbox, and single sessions accessing multiple remote mailboxes are both possible with this approach.

This differs from POP3 (RFC 1225) in that POP is a store-and-forward transport protocol that allows an MUA to retrieve pending mail from a mail drop (where it is then usually deleted automatically), whereas IMAP is focused on remote mailbox manipulation rather than transport. IMAP differs from various vendor-specific remote access approaches in that IMAP is an open protocol designed to scale well and accommodate diverse types of client operating systems.

Security-related tasks include how to incorporate secure authentication mechanisms when establishing a session, and possible interactions with Privacy Enhanced Mail.

It is expected that most of the work of this group will be conducted via email. A goal is to integrate and update RFC 1176 and the existing IMAP2bis draft, then submit the result as an Internet-Draft well before the November IETF meeting, which would then focus on detailed review of the text in preparation for submission as a Proposed Standard before the end of 1993.

Goals and Milestones

Done	Post an Internet Draft of the revised IMAP 2 protocol.
Aug 1993	Hold an Interim Working Meeting at UW or CMU.
Done	Hold a Working Group meeting to review the IMAP document.
Done	Hold a Working Group meeting at the November IETF meeting.
Dec 1993	Submit the IMAP protocol to the IESG for consideration as a Proposed Standard.

Internet-Drafts

"INTERACTIVE MAIL ACCESS PROTOCOL - VERSION 2bis", 10/29/1993, M. Crispin <draft-ietf-imap-imap2bis-02.txt>

INTERIM MEETING REPORT

Reported by Terry Gray/University of Washington

Minutes of the Internet Message Access Protocol Working Group (IMAP)

Summary

An interim IMAP Working Group meeting was held at the University of Washington on August 30 and 31, 1993. Eight people attended. Twenty-three issues were discussed. A consensus position was reached on twenty of those issues. No consensuswas reached, but some progresswas made on the issues related to namespace semantics and hierarchy support.

A new Internet-Draft, incorporating the results of this meeting and several suggestions made via e-mail, will be forthcoming.

Editor's Note: The complete set of minutes for this meeting is available via FTP or mail server from the remote directories as /ietf/imap/imap-minutes-93aug.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

CURRENT MEETING REPORT

Reported by Terry Gray/University of Washington

Minutes of the Internet Message Access Protocol Working Group (IMAP)

Summary

On Monday, twenty-seven people convened for the first of two scheduled IMAP Working Group sessions (although not everyone signed the attendance sheet). On Tuesday, a proper subset of around a dozen stalwarts carried on. For the continuation session after dinner, we were down to six, and two different ad hoc "midnight subcommittee" meetings were held.

The fifteen original agenda items and four new ones were considered. Considerable progress was made on all fronts. One notable result: on Monday the group agreed that the acronym "IMAP" should be remapped to the words "Internet Message Access Protocol" to better reflect what the protocol has evolved into.

Editor's Note: An itemized list of agenda items and their resolutions are available via FTP or mail server from the remote directories as /ietf/imap/imap-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

There are three remaining work items:

1. John Myers to propose an additional set of protocol-specified special information tokens.
2. Chris Newman to propose a revised hierarchy support solution based on conceptual agreement reached at the meeting.
3. Chris Newman to propose a syntax for an IMAP "meta" namespace, to allow unambiguous identification of multiple namespaces.

A new draft, incorporating at least some of the above three pending items and all of the other agreed upon items is expected around 12 November 1993.

All in all, it was a very productive meeting!

Attendees

Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Robert Austein	sra@epilogue.com
Luc Boulianne	lucb@cs.mcgill.ca
James Conklin	jbc@bitnic.educom.edu
Mark Crispin	mrc@cac.washington.edu

Walt Drummond	drummond@noc.rutgers.edu
Urs Eppenberger	eppenberger@switch.ch
Erik Fair	fair@apple.com
Patrik Faltstrom	paf@nada.kth.se
Ned Freed	ned@innosoft.com
Terry Gray	gray@cac.washington.edu
Jeroen Houttuin	houttuin@rare.nl
Ben Levy	seven@ftp.com
Lars-Johan Liman	liman@ebone.net
Chip Matthes	chip@delphi.com
Bob Morgan	morgan@networking.stanford.edu
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Kenneth Rossen	kenr@shl.com
Mark Smith	mcs@umich.edu
Bradley Wilson	wilson@ftp.com
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov
Peter Yee	yee@atlas.arc.nasa.gov

2.1.2 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

"DSA Metrics", 04/30/1993, P. Barker, R. Hedberg <draft-ietf-osids-dsa-metrics-01.txt>

"Representing IP Information in the X.500 Directory", 09/02/1993, T. Johannsen, G. Mansfield, M. Kosters <draft-ietf-osids-ipinfo-x500-dir-00.txt, .ps>

"Charting Networks in the X.500 Directory", 09/02/1993, G. Mansfield, T. Johannsen, M. Knopper <draft-ietf-osids-chart-network-dir-00.txt, .ps>

"Connection-less Lightweight Directory Access Protocol", 10/27/1993, A. Young <draft-ietf-osids-cldap-00.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"
RFC 1484	"Using the OSI Directory to achieve User Friendly Naming (OSI-DS 24 (v1.2))"
RFC 1485	"A String Representation of Distinguished Names (OSI-DS 23 (v5))"
RFC 1487	"X.500 Lightweight Directory Access Protocol"
RFC 1488	"The X.500 String Representation of Standard Attribute Syntaxes"

CURRENT MEETING REPORT

Reported by Paul Barker/University College London

Minutes of the OSI Directory Services Working Group (OSIDS)

Special thanks to Sri Sataluri, Mark Prior and Ken Rossen for their contributions to these minutes.

DSA Performance Study (Roland Hedberg)

DSA performance statistics are being circulated by Leggenhager regularly. But this study is based on study of the logs.

Reachability

Editor's Note: A sketchy account of this issue is available via FTP or mail server from the remote directories as /ietf/osids/osids-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

CLDAP (Steve Kille)

This is a connectionless protocol for retrieving names (something more similar to DNS) from the directory. It is an important element for deploying the directory and it is generally agreed that it should be moved speedily to a Proposed Standard.

Erik Huizer stated that a similar proposal has been discussed by Christian, similar in functionality, but it has not been put on paper.

Steve Kille indicated the group will proceed as if there is no other document. If Christian's document appears, and if it becomes necessary, the group will review the present CLDAP document in that light. Both CLDAP and Christian's proposal are LDAP-compatible. If one needs authentication, Steve said LDAP will be used.

A period of two to three weeks will be allowed for electronic discussion. After that, if there are no comments/changes and if there is no review requirement in the light of the document which Christian may issue, then the document will be submitted for consideration as a Proposed Standard.

The above resolution was approved by a show of hands.

Networks in the Directory (Glenn Mansfield)

Editor's Note: More details of each of the items below are available via FTP or mail server from the remote directories as /ietf/osids/osids-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

- OSI-DS 37/38 present status

  Two Internet-Draft have been in circulation since July 9, 1993:
  - "Charting Networks in the X.500 Directory" (draft-ietf-osids-chart-network-dir-00.txt) explains the necessity of network maps and its possible uses.
  - "Representing IP Information in the X.500 Directory" (draft-ietf-osids-ipinfo-x500-dir-00.txt) contains the schemas for representing IP-networks in the directory.

  Steve said that the following two Internet-Drafts will be moved to Experimental RFC's.
  So far, no negative responses or comments on mailing list or via personal mail (and few positive ;-) ones) have been received. Experiments and/or implementations are being carried out at several sites.

- Deployment strategy for Directory in the Internet

  The deployment document was circulated in Amsterdam and only minor changes were made to that. Steve said that it needs to be made into an OSI-DS document.

- Network Information

  Applications based on this include:
  - Network maps for Configuration management.
  - Connection trees
  - Softpages

- JPNIC WHOIS DB is in Progress

- DNS in the directory

  There are problems with the present schema. Improvements and changes are being made and the group hopes to circulate a draft by the end of November, and to commence deployment by the end of December.

- Application Support

  Glenn indicated that the group was in the process of preparing a document. There are some operational issues to be discussed. Steve stressed the importance of having an operations guide. Glenn said it is being worked on and presently it is an image of the DNS Administrator's Guide. A first draft will be posted to the list for discussion.

- Operational issues

  Real life applications are starting. The reliability of DSAs has to be improved. In case of problems due to other domains:
  - Complain privately to the responsible person for the domain.
  - Complain publicly to the responsible person for the domain.
  - Complain to the parent domain authorities.
  - Ask the parent authorities to excommunicate the domain. [Quote from RFC1033]

Liaisons

Liaison reports were given for the following groups:

- ISO/IEC/ITU-T (Ken Rossen)
- OIW DS SIG (Ken Rossen)
- NADF (Tim Howes)
- AARNet (Mark Prior)
- PARADISE (Roland Hedberg)
- NREN-NIS (Sri Sataluri)

Editor's Note: Reports for each of these liaisons is available via FTP or mail server from the remote directories as /ietf/osids/osids-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Schema Working Group

Members include Sri Sataluri, Tim Howes, Ken Rossen, and Russ Wright. The goals of this group are to:

- Identify a repository and appropriate useful formats for publicizing and distributing schema elements (object classes and attributes) to the Internet community.

- Facilitate broad-based experimentation with new applications of X.500 by publicizing experimental schema elements.

- Maintain a stable production schema for the Internet, including definitions both for common core of elements and application-specific subschemas.

A draft document was sent out to the osi-ds mailing list on November 4, 1993. To receive a copy please send mail to `sri@internic.net`. A revised form of this document (Procedures and Guidelines) will be released as an Internet-Draft in the first week of December and should eventually become an Informational RFC. Please send comments ASAP to `schema@ds.internic.net` and/or `osi-ds@cs.ucl.ac.uk`.

In addition to the procedures document, the schema group will publish a standards-track RFC that will document the "core Internet Schema" (successor to RFC 1274) and an Informational RFC documenting the current Internet Schema will be issued on a six-month update cycle.

The schema group will announce the availability of the "Internet Schema" and will start accepting updates on December 1, 1993.

Comments, suggestions, and submissions should be sent to `schema@ds.internic.net`. As and when the "Internet Schema" gets updated, an announcement will be sent out using a mailing-list `schema-announce@ds.internic.net`. To join this list, please send a message to `schema-announce-request@ds.internic.net`.

RFC 1384 Update - X.500 Naming Guidelines (Steve Kille)

The members agreed to progress this document as an Informational RFC. Comments should be directed to the authors and/or to the list ASAP.

Charter Discussion

Erik Huizer, an Applications Area co-Director, presented a short summary of the previous day's meeting. Here are the salient points:

- The OSI-DS Working Group in its present form cannot continue and should be disbanded and a set of small focussed working groups be set up.

- WHOIS++ and X.500 address similar issues and hence common problems need to be identified and worked on. Resources are too scarce for duplication of effort.

A strategy (proposed by Steve Kille and accepted by everyone) is to identify a set of working groups that will work on the open issues of the OSIDS Working Group and prepare charters for these new groups. The charters, and the need for more or fewer groups, should be discussed in the osi-ds mailing lists. By the Seattle IETF, the new working groups should be constituted and the OSIDS Working Group should be disbanded.

There was strong support for continuing to maintain the osi-ds mailing-list.

A list of possible working groups and chairs was composed at the meeting:

- Lightweight Protocols for Access and Synchronization (LDAP, CLDAP, SOS, SOLO, etc.) – Tim Howes
- Data Structure and Schema Management (will use a design team approach for schema issues; deal with naming issues) – Sri Sataluri
- Index Services and Distributed Search (Index DSAs, Centroids, etc.) – Simon Spero
- Use of Directory for Network Management – Glenn Mansfield
- URN → URL Resolution

The following suggestions were made:

- There is a need for the Data Structure and Schema group. Similar efforts in other groups should also be folded in.
- Access and synchronization issues should be split.
- Do not fragment the work into too many things.
- There is widespread desire for creating a group to discuss operational issues. This motion was seconded later by Linda Millington, Mark Prior, and Arlene Getchell.
- Operations issues must be dealt with in the operations area. There are plans to create an IOTF (Internet Operations Task Force) since, in several projects, the technical work is more or less finished and operations issues are becoming vital. Until the IOTF is formed the IDS Working Group will be used to get the operations work done.
- In the IIIR Working Group meeting the issue of forming a working group to discuss Quality Assurance Issues for X.500, Gopher, WAIS, WHOIS++, etc. was discussed. There will be a BOF at the Seattle IETF. A mailing-list (`quality@sunsite.unc.edu`) is being formed to discuss the quality issues. To join, send a note to (`listserv@sunsite.unc.edu`) with the following body:

  subscribe quality
  <your email address>

Attendees

Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Glen Cairns	cairns@mprgate.mpr.ca
Richard desJardins	desjardi@eos.nasa.gov
Urs Eppenberger	eppenberger@switch.ch
Qin Fang	qin_fang@unc.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Vincent Gebes	vgebes@sys.attjens.co.jp
Arlene Getchell	getchell@es.net
Mei-Jean Goh	goh@mpr.ca
Chris Gorsuch	chrisg@lobby.ti.com
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Tim Howes	tim@umich.edu
Richard Huber	rvh@ds.internic.net
Erik Huizer	Erik.Huizer@SURFnet.nl
Barbara Jennings	bjjenni@sandia.gov
Steve Kille	S.Kille@isode.com
Kanchei Loa	loa@sps.mot.com
Glenn Mansfield	glenn@aic.co.jp
Wayne McDilda	wayne@dir.texas.gov
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Rakesh Patel	rapatel@pilot.njin.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Marshall T. Rose	mrose@dbc.mtview.ca.us
Kenneth Rossen	kenr@shl.com
Srinivas Sataluri	sri@internic.net
Rickard Schoultz	schoultz@sunet.se
Vincent Shekher	vin@sps.mot.com
Mark Smith	mcs@umich.edu
David Staudt	dstaudt@nsf.gov
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov
Russ Wright	wright@lbl.gov
Peter Yee	yee@atlas.arc.nasa.gov
Weiping Zhao	zhao@nacsis.ac.jp

2.1.3 TELNET (TELNET)

Charter

Chair(s)

Steve Alexander: `stevea@lachman.com`

Mailing Lists

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

Description of Working Group

The TELNET Working Group will examine RFC 854, "Telnet Protocol Specification," in light of the last six years of technical advancements, and will determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the TELNET options, and decide which are still germane to current day implementations of the TELNET protocol.

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

(2) Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically: Environment variable passing, Authentication, Encryption, and 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 RFC 854.
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: Kerberos Version 5", 11/18/1993, S. Alexander <draft-ietf-telnet-authker-v5-01.txt>

"Telnet Environment Option", 10/10/1993, S. Alexander <draft-ietf-telnet-envmnt-option-03.txt>

"TELNET Transfer Control Option", 06/22/1993, S. Denton <draft-ietf-telnet-transfer-option-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)

Agenda

- Any feedback on the "Telnet Environment Option" Internet-Draft Last Call.
- Discussion of merged authentication/encryption options, "Telnet Authentication and Encryption Option" Internet-Draft, with emphasis on a plan to get the document finished.
- Any other business.

Steve Alexander presented the agenda and asked if there were other items that needed to be discussed. Marjo Mercado asked about the charter, so a brief discussion was held. Steve stated that the charter was no longer open-ended, and that the group would conclude when the environment and authentication documents were done. If other issues arise the charter will have to be amended. There was general agreement on this point.

Since a Last Call has been issued for the "Telnet Environment Option," Steve asked for any feedback. Marjo pointed out a minor grammatical error which will need to be corrected during the RFC editing process. Steve urged everyone to review the document if they hadn't already.

The bulk of the meeting was devoted to authentication. Dave Borman is currently implementing the merged authentication/encryption options. The group discussed whether it is okay to abandon the output mode DES—this seemed acceptable to all present. Ted Ts'o raised the concern about active attackers forcing the use of a weaker encryption mechanism. There was brief discussion on this point and Ted agreed to write up his view of how this could be avoided.

John Linn expressed concern about getting a Kerberos V authentication document out ahead of the merged mechanism. The group agreed that the current V5 Draft, "Telnet Authentication: Kerberos Version 5" could be issued as an Experimental RFC. Steve will send the current draft to Ted for review.

Dave Borman mentioned that he would like to release his current telnet reference sources in the near-term, but is concerned about the encryption code. Ted suggested that perhaps MIT could be a distribution point, since they have a similar problem with the Kerberos distribution.

Sam Sjogren raised the issue of interoperability testing. The group was receptive, and might try to schedule an event prior to the Seattle meeting. This would most likely be a virtual event held between cooperating parties via the Internet. There was some discussion of whether this would be appropriate to have at an IETF meeting, but no conclusion was reached.

Action Items

Dave Borman	Will finish implementation of the merged authentication/encryption options.
Ted Ts'o	Will write up a discussion of how he would like to see the encryption type negotiation covered by a checksum to prevent active attackers from forcing a weak encryption method to be negotiated.
Steve Alexander	Will fine-tune the Kerberos V draft and send it to Ted for review with the goal of issuing it as an Experimental RFC.
Steve/Dave	Will drive the document editing process so that work on merging the encryption text into the Kerberos documents will be complete by Seattle.

Attendees

Steve Alexander	stevea@lachman.com
David Borman	dab@cray.com
Peter Cameron	cameron@xylint.co.uk
Jonathan Didner	jonb@bangate.compaq.com
John Linn	linn@security.ov.com
Marjo Mercado	marjo@cup.hp.com
Clifford Neuman	bcn@isi.edu
Jon Penner	jjp@bscs.uucp
Vladimir Sukonnik	sukonnik@process.com
Theodore Ts'o	tytso@mit.edu
Raymond Vega	rvega@cicese.mx

2.1.4 Telnet TN3270 Enhancements (TN3270E)

Charter

Chair(s)

Robert Moskowitz: 3858921@mcimail.com

Mailing Lists

General Discussion: tn3270e@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 describing 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 RFC 1041 and the Experimental RFC describing the interim approach to printing and LU name selection.

Goals and Milestones

Done	Submit an Internet-Draft documenting current TN3270 terminal emulation practices.
Done	Post an Internet-Draft describing the Interim approach to printing and LU name selection.
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.

Done Submit 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.

Internet-Drafts

"TN3270 Enhancements", 10/05/1993, B. Kelly <draft-ietf-tn3270e-enhancements-02.txt>

"TN3270 Extensions for LUname and Printer Selection", 07/28/1993, C. Graves <draft-ietf-tn3270e-luname-print-00.txt>

"TN3270 Current Practices", 11/12/1993, J. Penner <draft-ietf-tn3270e-current-pract-03.txt>

INTERIM MEETING REPORT

Reported by Robert Moskowitz/Chrysler Corporation

Minutes of the Telnet TN3270 Enhancements Working Group (TN3270E)

The TN3270E Working Group met Thursday, August 26, 1993, at the Fall '93 INTEROP in San Francisco. Approximately thirty-four people attended, most were new faces. Bob Moskowitz, TN3270E Chair, presented a little background on the work started at the 26th IETF meeting in Columbus, Ohio, and what had been accomplished to date.

The working group is well on the way of meeting its charter to produce three RFCs.

1. "TN3270 Current Practices"

   This one on current TN3270 practices is authored by Jon Penner; it is available as an Internet-Draft and is undergoing what is hoped to be final changes.

2. "TN3270 Extensions for LUname and Printer Selection"

   The Open Connect System's proposal for an interim printing solution is authored by Cleve Graves. It is also available as an Internet-Draft and is now undergoing what is hoped to be final changes.

3. "TN3270 Enchancements"

   This is a full redesign of the TN3270 protocol and is being authored by Bill Kelly. It is in its second version and the next is being worked on. The group does not feel that this document is finished; more customer and vendor input is needed.

Cleve Graves followed with a presentation of his proposal. It quickly became apparent that a name was needed for this approach, and 'Simple IBM Printing' or SIMP was proposed. The only addition suggested was support for IBM's Intelligent Printer Data Stream, or IPDS, and examples of how the query response would work. Cleve agreed that this was needed and committed to adding it. Cleve will also remove descriptive text that has found its way into the current practices document. He will work at getting the changes done quickly so this document can be moved on to full RFC status.

Cleve's presentation was accompanied by a number of customer representatives' requests that this functionality be made available as soon as possible, if not yesterday, to meet their user needs. Based on this, everything will be done to get the document through the IETF/IESG system.

Bill Kelly then presented the 'Enhanced TN3270 protocol'. This protocol will be put forth as an IETF Standard and Bob explained the logistic steps involved in setting standards within the IETF.

There was some concern about 'breaking' the TELNET practice of negotiating Terminal Type first. With Enhanced TN3270, the new TN3270E option is negotiated before Terminal Type. Most of the attendees that understand TELNET did not see this as an issue. There was a side discussion on whether SNA positive and negative acknowledgments were needed and, if so, would SNA sequencing be needed in the protocol as well. No resolution of this point was reached. A question was raised about authentication and encryption. General agreement was reached that the TELNET authentication option will be used when possible, but that encryption will use the IBM SNA encryption methodology, thus making it not a TN3270 protocol issue.

Finally there was discussion on expanding the protocol to include 5250 support. This was viewed by many of the vendors as an unwise burdening of the TN3270 protocol and that it would not bring any value either to it or to 5250 support, as the protocol issue is a very small part of either 5250 or 3270 emulation. However, it was agreed to spend some effort to investigate this.

Bill wrapped up the presentation by stating that there are a number of changes that he will be making to the current document and that he will be calling on those attending to review his document and comment about it on the list.

Attendees

Chris Bickel	bcb9467@scmi.boeing.com
Anders Borg	anders@axis.se
Olaf Burge	unetokb@uicvmsa.aiss.uic.edu
Mike Coleman	mikec@dophin.tdh.texas.gov
David Dedic	Dave@cc.mcgill.ca
H. Dychoff	hdyc@chevron.com
Roger Erickson	
Eric Fleischman	ericf@atc.boeing.com
Nadine Goes	goes@cnmtl.cn.ca
Cleve Graves	cvg@oc.com
Martin Gren	martin@axis.se
Scott Hertzog	
Rich Hogrogian	rhhogrog@vnet.ibm.com
Tim Huntley	timh@vnet.ibm.com
Bill Kelly	kellywh@mail.auburn.edu
William Kwan	kwan@rabbit.com
Marjo Mercado	marjo@cup.hp.com
Jeanine Modic-Carmons	modic@r3vm.dsd.trw.com
Tome Molnar	molnar@utcc.utoronto.ca
Robert Moskowitz	3858921@mcimail.com
Rod Nay	
Yoram Nussenboim	yoram@netmanage.co.il

William Palter	palter@tgv.com
Daniel Pellet	pellet@ibm.unige.ch
Jon Penner	jjp@bscs.uucp
Mike Perelc	xopr40@email.mot.com
Reynold Reimer	
Bill Rubin	rubin@watson.ibm.com
Willie Ryals	unetwar@uicvmsa.aiss.uic.edu
Jim Smith	adp07@msu.edu
John David Smith	jdsmith@ads.tdh.texas.gov
Jim Spink	spink@ralvmm.vnet.ibm.com
Nick Thille	thille@cisco.com
Harrison Warrener	

CURRENT MEETING REPORT

Reported by Robert Moskowitz/Chrysler Corporation

Minutes of the TELNET TN3270 Enhancements Working Group (TN3270E)

"TN3270 Current Practices"

There was some discussion of the "current practices" Internet-Draft, and it was agreed that with a few editorial changes it is ready to be submitted to the IESG with the request that it be published as an Informational RFC. Bob Moskowitz mentioned the possibility that all three of the working group's documents might wind up on the standards track, and become various "flavors" of Standards RFCs. Several members objected to this approach and only wanted one standards document.

"TN3270 Extensions for LUname and Printer Selection"

Discussion centered around two areas: IPDS printer support and handling of errors during term-type negotiation. IPDS problems arise due to the difference in LU1 and LU3 support (function management headers versus structured fields). It was agreed that LU3 IPDS support can be attained by adding a term-type of IBM-4224; this will be a "queryable" device type. LU1 IPDS support will not be included in TN3287 at this time.

The TN3270 Extensions Internet-Draft will document a list of error codes (and their meanings) that represent problems that can occur while negotiating the term-type. If an error occurs, the server will send the error number to the client instead of negotiating the EOR and Binary options; it will then close the connection. The client will be able to take whatever action it deems appropriate, which could include such things as sending a message to the user or attempting to reconnect.

The goal is one more content change on the TN3270 Extension Internet-Draft followed by a quick editorial cleanup and submission to our Area Director for review and forwarding to the RFC Editor.

"TN3270E Enhancements"

First up was the subject of sequence numbers; some members questioned the need for them. It was agreed that sequence numbers will be needed when exception response processing occurs. It was also decided that the sequence number field in the TN3270E header need only be maintained when the RESPONSES function has been agreed to; otherwise, this field will contain binary zeroes.

There was a lively discussion of the initial negotiation of the TN3270E option (i.e., the WILL/DO and WON'T/DON'T TN3270E negotiation). It was pointed out that since TN3270E will be a TELNET option governed by the TELNET RFC, it must be treated like other TELNET options: both parties must be free to send WILL, DO, WON'T and DON'T. Some in the group would like to have the server be the only party allowed to actually initiate the TN3270E negotiation—that if a client sends a DO TN3270E, the server should respond with a DON'T TN3270E, and subsequently send a DO TN3270E when it is ready. It was agreed that input from people such as Steve Alexander, TElNET Working Group Chair, would be helpful in resolving this issue.

TN3270E Term-type Negotiation

Next came a discussion of the TN3270E term-type negotiation. Two of the "gateway-based server" vendors present expressed serious concerns with the recently proposed method of simply negotiating TERMINAL or PRINTER and having the server send out a Read Partition Query. These objections had to do with the notion of sending out 3270 data before a session has actually been established. It was suggested that the best approach would be to leave the Document as it reads now (which includes 3278 models 2, 3, 4 and 5, both with and w/o the "-E" suffix) and to add a "DYNAMIC" term-type, which would allow for the "non-standard" screen sizes. There was also a suggestion that what is really being negotiated are screen sizes and whether or not a device is queryable; therefore, model designations should be done away with and these items should be negotiated directly. More discussion on the list will be required to resolve this issue.

John Klensin, Applications Area co-Director, briefly discussed the question of WILL/DO and WON'T/DON'T TN3270E. He also stated that the current practices Internet-Draft will be published as an Informational RFC, not a Standards one. John also reported that all of our Internet-Drafts will be reviewed by TELNET experts before being submitted to the RFC process to attempt to avoid open discussions during the Last Call process. Further, there will be some further thought on what RFC designation will be used for the TN3270 Extensions Internet-Draft.

RFC 1538

With time running out, a brief discussion of RFC 1538 (SNA/IP) ensued. Two of the vendors present are implementing a form of SNA over IP (although they are not compatible). It was pointed out that IBM would prefer to address the issue through the APPN Implementor's Workshop, rather than in the IETF. Discussion will take place between higher levels of the IETF and IBM as to where best to work on this.

Attendees

Rich Bowen	rkb@ralvm11.vnet.ibm.com
Dan Brendes	brendes@raleng.mtc.com
Thomas Butts	tom@oc.com
Roger Fajman	raf@cu.nih.gov
Cleve Graves	cvg@oc.com
Jeff Hilgeman	jeffh@apertus.com
Bill Kelly	kellywh@mail.auburn.edu
John Klensin	Klensin@infoods.unu.edu
William Kwan	kwan@rabbit.com
David Lapp	lapp@waterloo.hp.com
Marjo Mercado	marjo@cup.hp.com
Robert Moskowitz	3858921@mcimail.com
William Palter	palter@tgv.com
Jon Penner	jjp@bscs.uucp
Eddie Renoux	elr0262@newsit2.mcdata.com
Barbara Sterling	bjs@mcdata.com
John Tavs	tavs@vnet.ibm.com

2.1.5 X.400 Operations (X400OPS)

Charter

Chair(s)

Alf Hansen: Alf.Hansen@uninett.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", 10/15/1993, Robert Hagens, Alf Hansen <draft-ietf-x400ops-mgtdomains-ops-06.txt>

"Postmaster Convention for X.400 Operations", 10/28/1993, C. A. Cargille <draft-ietf-x400ops-postmaster-03.txt>

"C=US; A=IMX", 10/28/1993, E. Stefferud <draft-ietf-x400ops-admd-03.txt>

"Using the Internet DNS to distribute RFC1327 Address Mapping Tables", 12/27/1993, C. Allocchio, A. Bonito, B. Cole <draft-ietf-x400ops-dnsx400maps-04.txt>

"Mail based file distribution Part 1: Dialog between two nodes", 07/06/1993, M. Kaittola <draft-ietf-x400ops-tbl-dist-part1-01.txt>

"Mail based file distribution Part 2: Over-all structure", 07/06/1993, M. Kaittola <draft-ietf-x400ops-tbl-dist-part2-01.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"

RFC 1502 "X.400 Use of Extended Character Sets"

CURRENT MEETING REPORT

Reported by Alan Cargille/University of Wisconsin

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

Executive Summary

- The Amsterdam minutes were not approved. They will be revised.
- The postmaster document will receive final editorial comments and be submitted for consideration as a standards-track RFC.
- The management domains requirements document will receive final editorial comments and be submitted for consideration as an Informational RFC.
- A revised document on storing RFC 1327 mapping rules in the DNS will be released within a few weeks. A new companion document about how this should be administratively implemented and deployed will be written by the next IETF or the meeting of the RARE Working Group on messaging.
- The proposed CXII group will continue to be discussed on the cxii list. If it cannot be finalized by the Seattle IETF, the group will probably not be created.
- The work on ADMD IMX is viewed as a United States national issue and should be developed in some US forum, not the IETF. The work should be fed back into the IETF for comments and publication.
- A sister group to the IETF on operations may be created (the IOTF).
- The X400OPS Working Group will be terminated following this IETF. Outstanding work items can be brought up on the X400OPS mailing list. If worthwhile, a small focused working group will be created to work on the new topic.

Thanks to Tony Genovese and Alf Hansen for chairing this group.

Goodbye, and thanks for all the fish!

Review of Action Items

This was difficult to do because action items were not summarized in previous minutes. This section will be updated as the Amsterdam minutes are revised.

Jim Romaguera conducted the review of the minutes from the last meeting. They are were not approved. They had been submitted in a rush. Alf and Tony apologized for incomplete minutes being published. Marko and Urs had sent messages requesting changes to the minutes which were not made. Marko's name was misspelled in Section 6. He was unhappy with the proposed chairs in Section 10.

The Amsterdam minutes will be reviewed again on the list and revised. Allan Cargille foolishly volunteered to edit the revised minutes. Action items need to be identified, both those from the previous meeting (Columbus) at the beginning of the document and those from Amsterdam in the body. We can also check the X400OPS list archive for comments on the minutes.

Postmaster Convention for X.400 Operations

Allan Cargille reviewed the key idea of the document. He removed the section about supporting an easy way to reach the managers of an X.400 management domain (ADMD or PRMD) out of the document and plans to write that up in a separate document (edit out the part about 84 and 88, just say both are running and reference both standards). There was consensus that the group will forward the document for consideration as a standards-track RFC. Allan will revise the document and clean up the references. He will then publish it as an Internet-Draft and ask for comments for one week on the ops list. This final review is for editorial comments only. Allan will make any necessary corrections and forward the document to be published. Allan will have the revised Internet-Draft out by November 8.

Operational Requirements for X.400 Management Domains

Alf made editorial changes to the document and cleaned up the references. Few people have read the final version. The document is available and key people are asked to review the document for editorial changes: Tony, Urs, Jeroen, Harald, and Allan. We will close discussion by November 15. People who read the document should let Alf and Tony know that they have read it. Tony will buy a beverage for the person who finds the most typos!

DNS support for RFC 1327 Mapping

Claudio has been working on a mechanism to store and look up RFC 1327 mappings using the DNS. The first proposal received some strong requests for changes from the namedroppers mailing list (DNS experts) at the March 1993 IETF. Claudio had also done work on storing X.400 routing and MTA connection information in the DNS. This work has been suspended in favor of using X.500 (the IETF MHSDS Working Group work).

Claudio has developed a second version of his proposal. The document will be published as an Internet-Draft this coming week. He presented the major changes of the new approach.

The new approach defines a new DNS resource record which allows a single DNS query for a lookup. Some extensions are also included for eventual future use. The new approach stores Table 2 (822 to X.400) and "Gateway Table" mappings in the normal DNS domain tree. Table 1 (X.400 to 822) mappings will be stored in a separate tree, rooted at the *national* level. This approach forces coordination between the X.400 and DNS naming authorities. This will require considerable work in explaining concepts and coordinating things. Claudio said there is a need for an API specification.

Mapping coordination at the national level will be achieved in different steps, according to the draft document on mapping authorities. It fits into the regionalization process currently ongoing in the internet. It allows a full authority delegation as a final result of the process. An orderly transition is supported from centralized storage of the mapping rules in Italy to using the new national mapping tree, because software will support checking the national tree first and looking in the Italian tree if nothing was found.

Mapping rule storage and control will proceed in three different steps:

1. The information is maintained centrally in Italy and servers fallback to that location for lookups.

2. The national trees are implemented but things are centralized at the national level.

3. The information is truly distributed in the national trees.

The document also makes it possible to define a DNS/x.500 interface to make LONGBUD and DNS a unique schema for mapping distribution, with no duplication and global accessibility.

There was general concern about an update problem with two distributed mapping storage technologies (DNS and X.500). Urs said that the technical work is done and is solid, and that we need to think about the administrative work that is necessary to use this technology. The group notes that this work has implications on the MHSDS Working Group.

Claudio will write a separate document about information and deployment of this technology by the next RARE MSG or IETF meeting. Further discussion of both documents will proceed in the RARE Working Group on messaging.

Commercial X.400 Interconnection with Internet (CXII) Working Group

A proposed charter was included as input to this meeting. Tony Genovese led this discussion. Tony's slides are on the ESnet file server (FTP to `ftp.es.net`, the directory is `pub/mhs/x400ops/houston`).

Since Amsterdam:

- There are two points of contention: chairs and technical contributors.
- The chairs will be determined solely by the Area Directors.
- Technical leads are needed for document sets.
- There is already one volunteer co-chair, but another is still needed.
- Technical leads for documents are needed.
- The working group will be in the Operations Area but Erik Huizer (without his co-director) will serve as Area Director for the group.

Editor's Note: Questions posed as well as detailed statements by some of the attendees is available via FTP or mail server from the remote directories as /ietf/x400ops/x400ops-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

US-ops

Allan Cargille led this discussion. He outlined two different issues which fall under this agenda item:

- The work on ADMD IMX under C=US.
- The need for a forum for Internet-related issues which are specific to the United States or North America.

There are questions about whether ADMD IMX should be viewed as a United States issue or an Internet-wide issue. It can be viewed as an Internet-wide solution which happens to be stuck under C=US due to the X.400 country-centric addressing structure. For example, if C=WW (worldwide) existed, we would prefer to register ADMD IMX under C=WW and it would not be bound to the Unites States. Alternatively, it can be viewed as the Unites States national solution to X.400 naming in the US Internet, which is US-centric and should be developed in a United States forum.

The second issue is that the IETF developed in the context of the US. Therefore work on an issue which was Internet-related could be conducted in the IETF, even if the work was US-centric. Now that the IETF has developed its identity as an international organization, Internet-related topics which are United States or North American in scope do not have a valid forum. The problem was recognized by the group, but addressing this problem is outside the scope of the X400OPS Working Group.

Editor's Note: Summaries of attendee's comments are available via FTP or mail server from the remote directories as /ietf/x400ops/x400ops-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

What Next

Erik commented that the ops list should be kept open because of the documents being progressed. He also noted that the RARE Working Group on messaging could be used for some topics. If discussions raise technical issues that merit IETF work, he would welcome a proposal for that group (not just an extension of X400OPS but a focused group for 1/2 year or so to work on a specific issue.)

Urs pointed out that there is a specific list for RFC 1465 issues: `rfc1465@chx400.switch.ch`.

Jeroen has copies of tutorial papers, and RARE can send more copies if needed.

Allan sees the following as outstanding work items:

- A document on the long-range plan for X.400 in the Internet.
- Possible work on dynamic X.400 routing using the DNS. X.500 work (mhsds/LONGBUD) is not materializing fast enough.
- X.400(88) in the GO-MHS community.
- A standard way to address the managers of an X.400 management domain (PRMD or MD).
- A document on internal operations of ADMD IMX.
- A document on connections between ADMD IMX and ADMDs.

It appears that the IMX work will not be approved to be done in the context of the IETF.

Steve was also concerned about mhsds delays. It is a very high priority for specifications and for implementation. A global solution is needed for scalable routing, he sees X.500 as the only viable solution.

Erik wants focused groups in future. The problem is that groups can have beautiful ideas about what needs to be done, but there must be volunteers to do the work. People are needed to chair the groups, write the documents, and lead discussions.

Erik thanked Alf and Tony for chairing the group. The working group will be terminated after this IETF.

Attendees

Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Vadim Antonov	avg@icm1.icp.net
C. Allan Cargille	allan.cargille@cs.wisc.edu
Urs Eppenberger	eppenberger@switch.ch
Qin Fang	qin_fang@unc.edu
Tony Genovese	genovese@es.net
Alf Hansen	Alf.Hansen@uninett.no
Jeroen Houttuin	houttuin@rare.nl
Erik Huizer	Erik.Huizer@SURFnet.nl
Barbara Jennings	bjjenni@sandia.gov
Marko Kaittola	Marko.Kaittola@funet.fi
Steve Kille	S.Kille@isode.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Jian Li	jian@rice.edu
Lars-Johan Liman	liman@ebone.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Kenneth Rossen	kenr@shl.com
Srinivas Sataluri	sri@internic.net
Richard Schmalgemeier	rgs@merit.edu
David Staudt	dstaudt@nsf.gov
Panos Tsigaridas	Tsigaridas@fokus.gmd.de
Brien Wheeler	blw@mitre.org
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov
Russ Wright	wright@lbl.gov

2.2 Internet Area

Directors:

- Stev Knowles: `stev@ftp.com`
- Dave Piscitello: `wk04464@worldlink.com`

Area Summary reported by Stev Knowles/FTP Software

The following BOF and working groups met during the November IETF meeting in Houston:

- Whither ATM - An Update BOF (ATMINFO)
- Dynamic Host Configuration Working Group (DHC)
- Internet Stream Protocol V2 Working Group (ST2)
- IP Over Asynchronous Transfer Mode Working Group (ATM)
- IP Over Large Public Data Networks Working Group (IPLPDN)
- Point-to-Point Protocol Extensions Working Group (PPPEXT)

Whither ATM - An Update BOF (ATMINFO)

An information only BOF was held on ATM technology and on developments since the Amsterdam IETF meeting.

- Bryan Lyles presented a summary of ANSI T1S1 activities towards defining a new ATM Bearer Service called "Class Y" for variable bit-rate traffic, i.e. "best effort."

- Eric Hoffman gave a presentation, on behalf of Allison Mankin and Maryann Perez, on the problems of implementing ATM standards.

- Mike Goguen gave a review of the work of the ATM Forum's Private NNI Working Group.

- The ATM Forum UNI 3.0 Specification is out and available.

Internet Stream Protocol V2 Working Group (ST2)

ST2 met for the first time at Houston in two sessions. The first session involved a brief review of the charter, presentations on ST-II experience from BBN and ARPA, and an overview of the IBM ST-II implementation. After the presentations, discussion was started on which portions of the protocol need revision, removal, or clarification. The second session was a more detailed working session in which details of State Transition diagrams, groups of streams, and the join/leave mechanism were discussed. The group also selected Lou Berger, Steve DeJarnett, and Luca Delgrossi as the document editors for the Internet-Drafts and RFC, and decided to issue a further call for writers over the mailing list. An interim meeting may be held in January to write the initial Internet-Draft.

IP Over Asynchronous Transfer Mode Working Group (ATM)

The Classical Internet-Draft Last Call closed on Monday, November 1. All issues raised during the Last Call process were dealt with and closed. One serious technical issue was raised by Dave Sincoskie regarding the arp table entry timeout and n*n InARP transmission characteristics. A paragraph change was presented and adopted by consensus at the Thursday meeting. Dave Piscitello approved the change process; another Last Call is not needed. The changed paragraph will appear in the meeting minutes. The document is awaiting IESG ballot.

Joel Halpern gave a presentation of the Routing Over the Large Clouds Working Group (ROLC) proposed charter. Ran Atkinson presented his MTU draft. Bob Cole led a discussion of the framework document. The chair hopes that this document can be turned into a planning guide for the working group.

The working group hosted other discussions on security, source address, the non-optimal behavior of InARP, selectors and multiple LISs, application binding, and Q.93B parameters. One discussion item, the issue of IP over the ATM Forum's LAN Emulation specification, was not completed.

IP Over Large Public Data Networks Working Group (IPLPDN)

The purpose of re-opening the IP over Large Public Data Networks (IPLPDN) working group was to clean up some unresolved items, and attend to those items which have come up after the group became inactive.

- Keith Sklower presented summaries of the "Determination of Encapsulation of Multiprotocol Datagrams in Circuit-switched Environments" and "Parameter Negotiation over Frame Relay" Internet-Drafts.

- The draft for the updates to RFC 1315 has expired. Caralyn Brown has agreed to repost it and set the wheels in motion to get it forwarded.

- Joel Halpern and Fred Baker volunteered to write an Informational document covering experience in partial mesh networks. The document will be posted on the mailing list and discussed there.

- The group decided that the definition of InARP for IPX might better be handled by Novell.

- During the IP over ATM discussions, it was felt that InARP was not robust enough. The group decided that because this problem was related to ATM's ARP server, the ATM group should pursue this work.

- Those who were most interested in IEEE 802.5 Source Routing over Frame Relay were not present at the meeting. It was decided that this should be taken to the mailing list for further discussion.

Point-to-Point Protocol Extensions Working Group (PPPEXT)

Two documents were referred, without discussion, to the IESG for consideration as Proposed Standards.

- "PPP over ISDN" (draft-ietf-pppext-isdn-03.txt)
- "PPP over SONETE/SDH" (draft-ietf-pppext-sonet-01.txt)

There was discussion about "PPP over X.25" (draft-ietf-pppext-x25-02.txt) and it was decided that the document should be recommended to the IESG for consideration as a Proposed Standard.

"PPP in Frame Relay" (draft-ietf-pppext-frame-relay-02.txt) was also discussed. It was recommended that a single sentence be added, and that the resulting Internet-Draft, draft-ietf-pppext-frame-relay-03.txt, be considered by the IESG as a Proposed Standard.

It was recommended that "PPP LCP Extensions" (draft-ietf-pppext-lcpext-04.txt) be considered by the IESG as a Proposed Standard. Another document will be drawn up describing the LCP option for negotiation of encapsulations.

Dave Rand presented the "PPP Reliable Transmission" document, (draft-ietf-pppext-reliable-00.txt). After some discussion, the document was recommended for consideration by the IESG as a Proposed Standard.

Several other documents were presented and discussed.

CURRENT MEETING REPORT

Reported by Mark Laubach/Hewlett-Packard

Minutes of the Whither ATM - An Update BOF (ATMINFO)

An information-only BOF was held on ATM technology and developments since the Amsterdam IETF meeting. Many people were appreciative that the BOF was held.

Presentations

Bryan Lyles presented a summary of ANSI T1S1 activities towards defining a new ATM Bearer Service called "Class Y" for variable bit-rate traffic, i.e. "best effort." Various public ATM providers are viewing this a boon as it is a different service than constant bit-rate (CBR) traffic (voice) and can be tariffed differently. T1S1 will be working further to create an implementation independent definition. When adopted in T1S1, it should make it into ITU-TSS and then be adopted by the ATM Forum. Bryan's closing said to "tune in again in about six months."

Eric Hoffman gave a talk on behalf of Allison Mankin and Maryann Perez on the problems of implementing ATM standards. Naval Research Laboratory (NRL) has recently released their Q.93B implementation called VINCE into the public domain. In addition, NRL will be implementing the Classic IP and ARP over ATM model. NRL feels the primary "snag" with Q.93B lies in the complicated parsing, forced bitfield parsing for rarely used elements, and spurious ordering rules. They do have a working version that is tracking ATM Forum's UNI 3.0 specification. Ports are being done to support several switch architectures.

Mike Goguen gave a review of the work of the ATM Forum's Private NNI Working Group. The groups main work is NNI signaling, VC routing, and Q.93B as the basis for the NNI signaling. This working group just started in July. The November ATM Forum meeting will focus on finalizing requirements and beginning a draft reference configuration. Some key requirement areas to be agreed on:

- Policy
- To what extent in Phase 1
- Hybrid private/public networks
- Source routing versus hop-by-hop

The ATM Forum UNI 3.0 Specification is out and available.

Attendees

Masuma Ahmed	mxa@mail.bellcore.com
Bart Berger	bart_berger@3com.com
Caralyn Brown	cbrown@wellfleet.com
Tracy Brown	tacox@mail.bellcore.com
Theodore Brunner	tob@thumper.bellcore.com
Ted Chang	tchang@atg.wiltel.com
Chris Chiotasso	chris@andr.ub.com
Blair Copland	copland@unt.edu
Kurt Dobbins	dobbins@ctron.com
Christopher Dorsey	dorsey@es.net
David Engel	david@ods.com
Dario Ercole	Dario.Ercole@cselt.stet.it
Michael Erlinger	mike@jarthur.claremont.edu
William Fardy	billf@frontier.com
Steve Feldman	feldman@mfsdatanet.com
Robert Fenoglio	fenoglio@vnet.ibm.com
Mark Garrett	mwg@faline.bellcore.com
Steve Garritano	steveg@kalpana.com
Robert Grow	bob@xlnt.com
Stuart Hale	stu_hale@vnet.ibm.com
Augusto Ham	ahamm@atg.wiltel.com
Daniel Hansen	danh@ngc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Mike Holloway	mikeh@newbridge.com
Refael Horev	horev@lannet.com
Frank Kastenholz	kasten@ftp.com
Mark Kepke	mak@fc.hp.com
Cheryl Krupczak	cheryl@empiretech.com
Ted Kuo	tik@ececho.ncsu.edu
Andrew Malis	malis@maelstrom.timeplex.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Steven Onishi	sonishi@wellfleet.com
Zbigniew Opalka	zopalka@agile.com
David Perkins	dperkins@synoptics.com
Kenneth Rodemann	krr@qsun.att.com
Hal Sandick	sandick@vnet.ibm.com
Michael Scanlon	scanlon@ftp.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Kanan Shah	kshah@cmf.nrl.navy.mil
Satya Sharma	ssharma@chang.austin.ibm.com
Uttam Shikarpur	uttam@zk3.dec.com
Timon Sloane	timon@timonware.com

Kaj Tesink	kaj@cc.bellcore.com
Panos Tsigaridas	Tsigaridas@fokus.gmd.de
Thomas Walsh	
Alice Wang	alice.wang@eng.sun.com
Honda Wu	honda@nat.com

2.2.1 Dynamic Host Configuration (DHC)

<u>Charter</u>

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 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.

Request For Comments

RFC 1531	"Dynamic Host Configuration Protocol"
RFC 1532	"Clarifications and Extensions for the Bootstrap Protocol"
RFC 1533	"DHCP Options and BOOTP Vendor Extensions"

RFC 1534 "Interoperation Between DHCP and BOOTP"

RFC 1541 "Dynamic Host Configuration Protocol"

RFC 1542 "Clarifications and Extensions for the Bootstrap Protocol"

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell University

Minutes of the Dynamic Host Configuration Working Group (DHC)

Since the last meeting of the DHC Working Group, DHCP was accepted as a Proposed Standard and the protocol specification was published as RFC 1541 (specification), RFC 1533 (options) and RFC 1534 (DHCP-BOOTP interoperation). J. Allard and Fred Lien organized two rounds of interoperability testing. At the second round of testing, 7 servers and 12 clients were tested:

- Microsoft: NT server, NT client, DOS client
- Sun: server and client
- HP: client
- Boeing: server and client
- DEC: client
- WIDE project (Japan): client, server and relay agent
- SGI: server and client
- Competitive Automation: server and client
- FTP Software: Windows and OS/2 servers, Windows and DOS clients

At present, there are no freely-distributable implementations. The WIDE project's implementation, described in a short presentation to the group, may be made available, but needs additional work first. The WIDE project, from Keio University, has implemented a DHCP server, client and relay agent, all based on UNIX and BPF (Berkeley Packet Filter). The server manages three databases: an available address pool, the set of client bindings and the known relay agents. The server uses ICMP echo to test for an address already in use before allocation. The server does not yet support the class identifier and vendor-specific data options, and the use of 'sname' and 'file' fields to hold options. The client is also built on BPF, as a library of functions for the various DHCP state transitions. Thus, the client software can be integrated into a variety of DHCP implementations. The relay agent uses BPF to communicate with the client and a socket to communicate with the server.

The interoperability testing identified a set of "minor" problems. The group discussed these problems and devised solutions as follows:

- Packet size: As BOOTP specifies smaller packets (300 octets) than DHCP (576 octets), the DHCP specification should be changed to explicitly allow smaller BOOTP packets.

- Minimal protocol requirements: DHCP requires some minimal functions from the TCP/IP protocol software on a client(e.g., ability to accept unicast replies before the IP address has been configured); these requirements must be added to the protocol specification.

- Use of 'ciaddr' field: As RFC 1542 requires a client to be able to respond to ARP requests if it puts an address in 'ciaddr', a client must use the 'requested IP address' option in DHCPREQUEST packets.

- Use of 'server ID' in DHCPACK and DHCPNAK packets: Make the use of 'server ID' a MUST requirement.

- Change number of retries of DHCPREQUESTs to 4 (to match other retry specifications).

- Use of 'BROADCAST flag' in DHCPNAKs: Possibly; still under consideration.

- Use of 'XID':

  - Client MUST use unique, random XID (NOT a well-known constant!) for each client DHCP packet to avoid associating reply for client B with request from client A.

  - Changing XID for each retransmission seems to be an implementation detail (client can choose to change XID with each retransmission of a specific DHCP packet).

- The group rejected the idea of a protocol version number.

- Timeouts: The group concluded that the timeout back off mechanism is "over-specified". The specification will be changed to read that the mechanism SHOULD be employed, and the reasoning behind choosing a specific mechanism.

- T2 not explicitly specified to be less than the lease time; specification to be fixed to reflect that requirement.

- Size limit on a single option (255 octets) may be too small: Allow multiple copies of the same option.

Specification and the use of 'client ID,' and 'client class' was discussed. The 'client ID' field is supposed to address the problem of separating client identification by the server from the delivery of DHCP packets from the server to the client. That is, the server always needs a MAC address (supplied in 'chaddr') for the client, through which messages can be delivered to the client, but the server may want to use some other identifier to track the binding of an IP address to that client. BOOTP overloads the MAC address with delivery and identification functions. It was decided to specify that DHCP servers should use 'client ID' if supplied by the client and 'chaddr' otherwise, for binding an IP address to a client. For the purposes of address binding, 'client ID' is to be interpreted as an opaque string of octets. Text will be added to the protocol specification explaining the reasons for using 'client ID' and possible effects of using 'chaddr' (e.g., when Ethernet cards are moved between hosts).

There was some discussion prior to the DHC meeting as to whether the 'client class' option was under-specified. The concern was that, without further specification, interoperability among DHCP participants would be compromised as a result of different interpretations of the the 'client class'. (See the DHC Working Group mailing list archive for more details.) The group felt that the primary use of 'client class' will be in aggregation of clients; i.e., the description of a collection of identical clients by a single entry in the DHCP server database. The attendees concluded that this use can be met as follows:

- Treat the 'client class' option as an character string.
- Recommend that vendors supply an initial value:
  - Should be "descriptive of the product".
  - Must be well-documented.
  - Must be useful in a DHCP database.
  - Must be configurable by the system administrator.
- Allow system administrators to choose local values for 'client class'.
- Add text to the protocol specification suggesting how system administrators can use vendor-supplied or locally-configured 'client identifier's.

The attendees also discussed two issues related to other IETF working groups. First, the Domain Name System Working Group (DNS) is aware of the requirement for a network interface to DNS updates. DHC is not the only group making such a request. DNS is working on the problem. Second, the attendees decided to hold off on any changes to the DHCP specification to accommodate new versions of IP and IP addressing such as SIP or TUBA.

There will be another round of interoperability testing in December after the latest changes to the protocol specification are integrated into the documentation. A copy of the text source used by the RFC Editor to generate the DHCP RFCs has been obtained, so revised documents can be generated that are consistent with the published RFCs.

Attendees

Kannan Alagappan	kannan@dsmail.lkg.dec.com
Steve Alexander	stevea@lachman.com
James Allard	jallard@microsoft.com
Jim Barnes	barnes@xylogics.com
Monroe Bridges	monroe@cup.hp.com
Michael Carney	mwc@sun.com
Jonathan Didner	jonb@bangate.compaq.com
Thomas Dimitri	tommyd@microsoft.com

Ralph Droms	droms@bucknell.edu
Roger Fajman	raf@cu.nih.gov
Shawn Gillam	shawn@timonware.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Marc Hasson	marc@mentat.com
Ronald Jacoby	rj@sgi.com
Rick Jones	raj@cup.hp.com
Akira Kato	kato@wide.ad.jp
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Andrew Knutsen	andrewk@sco.com
David Lapp	lapp@waterloo.hp.com
Fred Lien	fred@speedster.ca.boeing.com
Glenn Mansfield	glenn@aic.co.jp
Jun Matsukata	jm@eng.isas.ac.jp
Sath Nelakonda	sath@lachman.com
Steve Parker	sparker@ossi.com
Isil Sebuktekin	isil@nevin.bellcore.com
Satya Sharma	ssharma@chang.austin.ibm.com
Robert Stevens	robs@join.com
John Tavs	tavs@vnet.ibm.com
Fumio Teraoka	tera@csl.sony.co.jp
Susan Thomson	set@bellcore.com
Akihiro Tominaga	tomy@sfc.wide.ad.jp
Keisuke Uehara	kei@cs.uec.ac.jp
Jean Yao	yao@cup.hp.com
Weiping Zhao	zhao@nacsis.ac.jp

Agenda

- Report on interoperability testing
- Review of outstanding problems
  - 'ciaddr'
  - Other "small" problems
  - Client ID/class ID
  - Interface to DNS
  - SIP or TUBA (hold for now)
- Strategy for advancing protocol specification
- Discussion of "small" problems - solved
- Review of client ID issues
- Discussion of client ID problem

Problems From Testing

- Use of 'ciaddr'
- Inclusion of server ID in acks and naks - must
- Change retries of requests to four - yes
- Use of broadcast flag in naks - yes
- More detail in class ID
- Client ID option "over specified"
- Protocol version number - no
- Timeouts and backoff "over specified" - should - reasoning, especially total time to wait, should be based on local network and site characteristics
- T2 not explicitly < lease - ok
- 255 vendor options insufficient => multiple copies of vendor options
- Size limitation - minimum, maximum DHCP message size

Client ID Issues

- What values can be used as client IDs?

L	T	n-1

2 MAC (not unique)
6 MAC

- Overloading 'client ID' with client-server behavior

Class ID Issues

- Under specified, overspecified - What is it?

- Hosts that cannot accept broadcasts
- Hosts that can accept broadcasts
- Hosts that can go into promiscuous mode
- Hosts that cannot go into promiscuous mode
- Hosts which must receive IP unicasts/MAC broadcasts
- Hosts which cannot receive IP unicasts/MAC broadcasts

Class ID

- Aggregate clients
- Distinguish vendor options
- Optional on server:
  - If implemented, should be given a value and must be documented
  - Must be configured by system administrator
  - Advisory words...descriptive of their products useful in DHCP database

Tastes Great

- Client ID is client "disambiguatory"
  - Wire <- MAC
  - DNS name

Less Filling

- Client ID is always wire <- MAC Uniquely defines and interface
- Disambiguation is elsewhere
- "Indexing"
- Warning - why not use MAC address...
- Allow client ID to override MAC address
- Careful about chaddr vs. client ID in RFC

DHCP implementations

WIDE version

Akihiro Tominaga

WIDE Project

Keio University

DHCP implementations WIDE version

- Software Structure
  - Server
  - Relay agent
  - Client
- Implementation Overview
  - User process implementation
  - Uses BPF and socket
  - DHCP client library
- Implementation Varieties
  - BSD/386 (i486)
  - NEWS-OS(4.3BSD-based) (MC68030)
  - NEWS-OS (R3000)

Why using BPF ?

- BPF is Berkeley Packet Filter.
- BPF can read/write the packets from Interfaces directly.
- BPF sends Limited Broadcast address (255.255.255.255).
- BPF recognizes the incoming interface.
- Most BSD UNIX have BPF.
- Structure.

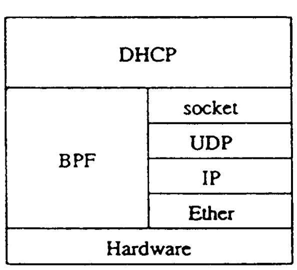

Server Implementation (1)

- Manages three databases
  - Address pool database
  - Binding database
  - Relay agent database
- Checks address duplication by ICMP echo request
- Supports All Address Allocation Models
  - Dynamic Allocation
  - Automatic Allocation
  - Manual Allocation
- Supports BOOTP client

Server Implementation (2)

- Format of Address pool database : modified version of CMU BOOTP Server

```
#dummy entry
global.dummy:snmk=255.255.255.0:\
                :dht1=500:dht2=850:
191.dummy:tblc=global.dummy:dfl1=3600:\
            :maxl=7200:rout=133.138.191.1:\
            :brda=133.138.191.255

#for dhcp client
19107:      :ipad=133.138.191.7:tblc=191.dummy:
19108:      :ipad=133.138.191.8:tblc=191.dummy:

# for bootp client
19106:      :ipad=133.138.191.6:alty=bootp:\
            :tblc=191.dummy:
```


Server Implementation (3)

- Unsupported features
  - IP fragmentation / reassembly
  - "Class identifier" option handling
  - "Vendor specific" option handling
  - "sname" field and "file" field (from client) handling

Client Implementation (1)

- Uses BPF.
- No IP fragment handling.
- libdhcp.a - dhcp client functions

dhcp_init(struct if_info *)	Client initialization
dhcp_discover(struct dhcp_reqinfo *, struct if_info *, struct dhcp_param **)	sends DISCOVER and collects OFFERs
dhcp_request(struct dhcp_reqinfo *, struct if_info *, struct dhcp_param *)	sends REQUEST and waits ACK
dhcp_verify(struct dhcp_reqinfo *, struct if_info *, struct dhcp_param *)	sends REQUEST to verify the address and waits ACK
dhcp_extlease(struct dhcp_reqinfo *, struct if_info *, struct dhcp_param *)	sends REQUEST to extend lease and waits ACK
dhcp_decline(struct dhcp_declinfo *, struct dhcp_ifinfo *)	send DECLINE

dhcp_release(struct dhcp_declinfo *, struct dhcp_ifinfo *)	send RELEASE
clean_param(struct dhcp_param *)	clear the structure which is used in dhcp_request and dhcp_verify
arp_reply(struct in_addr *, struct dhcp_ifinfo *)	send ARP reply
reset_if(struct if_info *)	reset the specified network interface
config_if(struct if_info *, struct in_addr *, struct in_addr *, struct in_addr *)	setup the specified network interface

- struct if_info includes information about interface

```
struct if_info {
  int fd;             /* bpf file descriptor */
  int arpfd;          /* bpf file descriptor for ARP */
  char name[sizeof "enxx"];
  int bufsize;
  int arpbufsize;
  char *buf;
  char *arpbuf;
  struct chaddr haddr;
};
```

- struct dhcp_param includes offered information from server

```
struct dhcp_param {
  struct dhcp_param *next;
  char got_option[GOTOPTSIZ];
  char *sname,          /* server host name */
```

```
        *file,          /* boot file name */
        *hostname,
        *merit_dump,
        *dns_domain,
        *root_path,
        *extensions_path,
        *nis_domain,
        *nb_scope,      /* NetBIOS scope */
        *errmsg;
  unsigned ip_forward          :1,   /* boolean */
           nonlocal_srcroute   :1,   /* boolean */
           all_subnet_local    :1,   /* boolean */
           mask_discover       :1,   /* boolean */
           mask_supplier       :1,   /* boolean */
           router_discover     :1,   /* boolean */
           trailer             :1,   /* boolean */
           ether_encap         :1,   /* boolean */
           keepalive_garba     :1,   /* boolean */
           nb_nodetype,        :4;  /* NetBIOS node type */
#define BNODE 0x1
#define PNODE 0x2
#define MNODE 0x4
#define HNODE 0x8
  u_char default_ip_ttl,  /* time to live (1 - 255) */
           default_tcp_ttl;
  u_short secs,              /* secs from DHCP message */
          bootsize,          /* bootfile size in 16 bits */
          max_dgram_size,    /* max dgram reassembly */
          if_mtu;            /* interface mtu */
  long xid,                  /* xid from DHCP message */
       time_offset;          /* 32 bits integer(network
                                byte order) offset from UTC */
  u_long lease_origin,       /* time that begin the lease */
         lease_duration,      /* lease duration */
         dhcp_t1,
```

```
       dhcp_t2,
       mtu_aging_timeout,   /* unsigned 32 bits integer */
       arp_cache_timeout,
       keepalive_inter;
struct shorts *mtu_plateau_table;/* array of 16 bits int */
struct in_addr server_id,        /* DHCP server identifier */
               ciaddr,
               yiaddr,           /* IP address */
               siaddr,           /* (tftp) server */
               giaddr,
               *subnet_mask,
               *swap_server,
               *brdcast_addr,
               *router_solicit;
struct in_addrs *router,
                *time_server,
                *name_server,
                *dns_server,      /* DNS server */
                *log_server,
                *cookie_server,
                *lpr_server,
                *impress_server,
                *rls_server,
                *policy_filter,
                *static_route,
                *nis_server,
                *ntp_server,
                *nbn_server,      /* NetBIOS name server */
                *nbdd_server,
                *xfont_server,     /* X font erver */
                *xdisplay_manager; /* X display manager */
};
```

- struct dhcp_reqinfo includes requeting parameter such as "Requested Lease".

```
/* parameter for DHCPDISCOVER and DHCPREQUEST */
struct dhcp_reqinfo {
  long            waitsecs;
  int             retry; /* number of retransmission */
  u_long          lease; /* suggestion of lease duration */
  struct in_addr  ipaddr; /* suggestion of IP */
  struct in_addr  ciaddr; /* used when verify */
  struct request_list {
    unsigned char  len;
    char           *list;
  } *reqlist;
  struct client_id *cid;
  struct class_id {
    unsigned char  len;
    char           *id;
  } *class;
};
```

- struct dhcp_declinfo is used in dhcp_release or dhcp_decline.

```
/* parameter for DHCPDECLINE and DHCPRELEASE */
struct dhcp_declinfo {
  struct in_addr ipaddr, /* DHCP client's IP addr */
                 srv_id; /* DHCP server's IP addr */
  struct client_id *cid;
  char             *errmsg;
};
```


Client Implementation (2)

- Client Program Example.

```
int main(int argc, char **argv)
{
  initialize();
  dhcp_init();

  /* main loop */
  while (1) {
    /* get OFFER from server */
    construct_request(req, &ipaddr);
    dhcp_discover(req, &netif, &param);

    /* get ACK from server */
    choose_offer(&param);
    dhcp_request(req, &netif, param);

    /* got ACK and configure */
    config_if(&netif, &param->yiaddr,
              param->subnet_mask,
              param->brdcast_addr);
```

```
    set_route(param);

    while (1) { /* extend lease loop */
      construct_extlease(req, param);
      if (dhcp_extlease(req, &netif, param) < 0) {
        reset_if(&netif);
        break;
      }
    }
  }
}
```


Relay agent Implementation

- Uses BPF and socket
- Relay agent model;

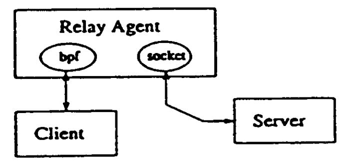

- Example of configuration file

```
2               # Number of servers to forward
133.138.191.1   # server IP
133.138.192.1   # server IP
133.138.193.1   # server IP
133.138.194.1   # server IP
```

in this case, relay agent choose 2 servers from 133.138.19X.1 and relay messages to them. Relay agent choose these servers by hashing messages.

2.2.2 Internet Stream Protocol V2 (ST2)

<u>Charter</u>

Chair(s)

Luca Delgrossi: `luca@ibmpa.awdpa.ibm.com`
Steve DeJarnett: `steve@ibmpa.awdpa.ibm.com`

Mailing Lists

General Discussion: `st@ibminet.awdpa.ibm.com`
To Subscribe: `st-request@ibminet.awdpa.ibm.com`
Archive: `ibminet.awdpa.ibm.com:~/pub/st/st-archive`

Description of Working Group

The Stream Protocol Working Group was formed to clarify and refine the existing specification of the Stream Protocol, Version 2 (ST-II) contained in RFC 1190. Since ST is a protocol that is already used in audio-visual and reserved-resource applications and services, the focus of this group is near-term and its primary purpose is to provide a specification that corrects errors in the existing ST specification and makes it easier to implement ST in a manner that is likely to be interoperable with other ST implementations.

The ST Working Group intends to address several areas of the ST specification including:

a) the formal definition of states and state transitions;

b) the removal of mechanisms which are too complicated as currently designed and which have not shown any use in practice;

c) address the ambiguities caused by the current implementation subsets;

d) definition of a clear IP encapsulation mechanism;

e) minor revisions suggested by experience with ST.

These modifications are expected to reduce implementation time and to improve the utility and interoperability of existing and future ST implementations. The working group may also provide guidance on the use of standard routing protocols to support ST and on the format and use of flow specifications. Finally, particular attention will be given to the specification of groups of streams as required for the efficient sharing of resources. Input from current ST developers and application developers will be solicited to help clarify issues that the working group should address.

It is the goal of the ST Working Group to produce a refined ST specification that can be used to rapidly satisfy operational requirements. The result of this group is expected to be an Experimental RFC. It is not the intention of this Working Group to define a new communication or resource reservation protocol. ST is part of the ongoing IETF efforts to develop protocols that address resource reservation issues. It is possible that future IETF Working Groups will produce other operational protocol options in this area. Related work by other IETF Working Groups shall be carefully monitored to see if the actions of this Working Group should be revised. In particular it is expected that there will be interaction with the AVT Working Group relating to issues of running RTP over ST.

Goals and Milestones

Done	Meet at IETF meeting to identify possible contributors. Review current ST-II specification to identify areas that need revising.
Mar 1994	Address the previously identified areas and complete an Internet-Draftwith a revised protocol specification
Apr 1994	Meet at the IETF meeting to review the completed Internet-Draft
Jul 1994	Submit revised Internet-Draft
Aug 1994	Submit the new ST specification for publication as an RFC

CURRENT MEETING REPORT

Reported by Luca Delgrossi/IBM

Minutes of the Internet Stream Protocol V2 Working Group (ST2)

The ST2 Working Group met in two sessions. Prior to the first session, Craig Partridge of BBN gave a talk on his experiences with ST-II. Craig did one of the first ST-II implementations with Steven Pink at the Swedish Institute of Computer Science in 1992.

Craig pointed out some of the weaker points of ST-II, including the fact that the specification (RFC 1190) forced the implementor to guess at what was intended in certain situations. Protocol complexity was also raised as an issue. The points raised all seemed consistent with the experience of the other implementors and certainly helped to motivate everyone prior to the start of the first working group meeting which was held immediately after the talk.

The working group meeting began with a brief review of the goals of the working group, the milestones specified in the charter, and a review of the agenda for the two working group meetings planned for this week.

IBM Heidelberg Transport System (HeiTS)

Luca Delgrossi of the IBM European Networking Center presented an overview of the IBM HeiTS (Heidelberg Transport System) stack, which includes an ST-II implementation. HeiTS is strongly focused on providing guaranteed quality of service to applications, particularly multimedia applications. HeiTS uses its own FlowSpec which is significantly simpler than the RFC 1190 FlowSpec. In addition to implementing ST-II, HeiTS also provides a Resource Management Subsystem which handles resource reservation for CPU, [memory] buffers, and network and communication adapter resources. In addition, HeiTS will interface with a "Central Resource Allocator" to coordinate network resource reservations in a complex network environment. Luca ended his presentation by discussing some possible protocol extensions or modifications that could make ST-II a more scalable and useful protocol, including the ability for targets to initiate a connection by joining an existing stream at a router instead of communicating directly with the origin to join a stream.

"ST-II Testing and Evaluation"

Doris Roland from Houston Associates, Inc. (HAI) gave a presentation on "ST-II Testing and Evaluation" which discusses some testing that HAI is doing for ARPA and the Defense Simulation Internet (DSInet). The DSInet is an evolution of the Terrestrial Wideband Network (TWBnet) which runs ST-II at about half of the sites on the network. Houston Associates, Inc. provides support for users of the DSInet and is performing their testing independently of other testing being done by BBN, which is the contractor responsible for building and operating the DSInet. The DSInet runs simulation exercises and video conferencing using ST-II to carry the realtime traffic. The HAI test plan consists of multiple stages, each of increasing complexity. They are explicitly testing stream setup, bandwidth reservation, routing, data transfer, stream modification, multicasting, and stream teardown. Their ultimate goal is to run multiple simulation exercises over portions of the backbone to see how well the overall system functions.

HAI Test Plan

After Doris's presentation, the group discussed some of the details of the HAI test plan, which included the measurement of delay variance in the network. Since a relatively low upper delay bound was specified, group members wondered why delay variance needed to be measured. The final answer was that the buffer space on the end systems is limited and excessive delay variance can cause buffers to overflow. An additional discussion item was brought up when it was mentioned that Wellfleet had developed an ST-II router for ARPA and was going to be deploying it on the DSInet. The group wanted to know whether this would be made generally available in Wellfleet's routers, but the Wellfleet representative was not certain at this point, as they had only recently been informed that their routers would be used on the DSInet.

"Preliminary ST-II Evaluation"

The final formal presentation was made by Michael Patton of BBN on "Preliminary ST-II Evaluation." This talk centered around work done by the DSI Network Engineering group at BBN under contract from ARPA. A brief overview of the DSInet was given, including a map showing most of the sites connected to the DSInet. The DSInet has nodes located throughout the US and as far away as Germany and Korea. It is an "around the world network" with over fifty sites connected presently. The DSInet architecture is built on a foundation of "Wideband Packet Switches" (BBN Butterfly's) connected to local BBN T/20V routers which handle routing of IP and ST-II. Local systems are connected to networks attached to the T/20V router. The testing done by BBN is being conducted in phases. The first phase was a simple connection of two Sun workstations on separate Ethernet's connected via a T/20V router. A traffic generator from SRI was used to provide the traffic and the bandwidth utilization was monitored to ensure that ST-II and IP were each running within their allocation limits. The traffic characteristics, network design, and end systems were changed in each phase to increase the stress on the network. Further testing is continuing to stress the network further.

After a minor digression about IP multicast, the group moved on to a list of possible discussion topics. That list included:

- Lack of State Transitions (14)
- FlowSpec issues (1)
- Heterogeneous FlowSpecs (1)
- Timestamps and negotiation (0)
- TargetList parameter (0)
- Header changes (2)
- Reason Codes (1)
- Hello/Status/Notify/Stream Data Flow (1)
- HID negotiation incompatibilities (4)
- Groups of Streams use (8)
- IP Encapsulation (1)
- Transport Protocol Interaction (e.g. RTP) (4)
- Stream naming simplify (0)
- Routing (2)
- Use of Class D addresses (4)
- ST-II MIB (2)
- ReverseCharge option (0)
- Point-to-Point option (0)
- Full Duplex (1)
- MTU discovery (0)
- Source routing (0)
- ErroredPDU pointer (0)
- Use over Ethernet/subnets (0)
- Join/Leave Streams (6)
- Subset implementation (2)

From here the group started to discuss various issues. It was decided, that in spite of IETF tradition, the group would vote on which topics people felt were most important to address, and the preferences are listed in parentheses in the list above. It should be noted that many topics that did not receive votes above were later discussed and it seems clear that many, if not all, will require the attention of the working group.

Editor's Note: A list of discussion topics which followed the vote is available via FTP or mail server from the remote directories as /ietf/st2/st2-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The meeting ended with a discussion of what other people were using ST-II for. IBM will start shipping a multimedia server (Ultimedia Server/6000) that uses ST-II to provide realtime data delivery to clients. Other users had been mentioned previously (BBN and the ARPA DSInet).

On Thursday the discussion turned to finding people willing to work on various issues, defining the scope of various problems, identifying people willing to work on writing the Internet-Drafts and the RFC, classifying protocol issues, and identifying work that needs to be completed prior to the Seattle IETF meeting.

State Transition and State Definition Problem

We started by discussing the State Transition and State Definition problem. Luca Delgrossi presented the state transition diagrams developed by IBM during implementation of the HeiTS stack. Luca agreed to make PostScript and ASCII versions of the state transition diagrams available via anonymous FTP so that others could review them. The PostScript versions should be available by November 19, while the ASCII versions might take a bit longer to create. People agreed that they needed time to study the state diagrams before volunteering to work on updating them, so a call for participation will be done over the mailing list.

Groups of Streams

Lou Berger discussed his ideas for the use of Groups of Streams. This could be used for associating independent streams (to allow "channel switching" while only allocating bandwidth for a small number of channels), bandwidth aggregation/sharing (for teleconferences), subnet multicast address sharing, identifying interdependence of streams, or sending hierarchically encoded data in multiple grouped streams. Lou, Skip Harboth, and Sybille Schaller from IBM in Heidelberg will look at defining Groups of Streams more fully and will then present a proposal to the mailing list.

Join/Leave Stream

Luca presented the Join/Leave stream idea as a way to allow targets to join a stream without having the source send a CONNECT message. This would save 1/2 RTT in the stream setup phase and would be accomplished by having the would-be recipient send a Join message toward the origin. As soon as the Join hit a router that was carrying the stream, that router would send a CONNECT back to the receiver and negotiation would continue "normally," with the exception that the router would be the origin for that receiver instead of the original data sender. A second proposal was that a backward path would be created from the would-be receiver toward the origin. This caused a lot of concern about requiring duplicate state machines in systems to handle a reverse-connection and also because this flows backward from the way routes are traditionally built. There was no consensus on this idea. The group asked IBM to write this up more fully and present it to the mailing list for discussion. After the list determines that this is (or is not) something that should be pursued, volunteers will (or will not) be solicited.

Future Plans

The discussion moved on to who would edit and write the Internet-Drafts and the RFC. Luca and Steve DeJarnett agreed to work on this, and Lou Berger said he would be willing to help out. The editors plan to base the new drafts and RFC on RFC 1190, but expect that a substantial rewrite and reorganization will be required. The editors intend to make PostScript and ASCII text versions available for both the drafts and (hopefully) the RFC.

Mark Pullen suggested that an interim meeting should take place in late January or early February to work on the Internet-Draft. Mark offered to host the meeting. Most people seemed to think this was a good idea and it will be suggested to the mailing list.

Subjects that are likely to be discussed in the near future include:

- HIDs with the possibility of removing the negotiation and just using globally-unique identifiers at each hop instead.

- Groups of Streams, and how you might use them to aggregate streams for bandwidth sharing and multicast address allocation.

- State Transition diagrams. Define them for the current protocol and then update them based on changes made by the working group.

- Join/Leave streams. Further specify how this might work for receiver-initiated communication.

[These minutes, while the product of discussions of the entire group, are quite possibly biased by the thoughts and interests of the author. I've attempted to eliminate some of that bias by asking others to review these notes but in the end they represent what I understood to have happened at the meetings.]

Attendees

Susie Armstrong	susie@mentat.com
William Barns	barns@gateway.mitre.org
Lou Berger	lberger@bbn.com
Stephen Casner	casner@isi.edu
Richard Colella	colella@nist.gov
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Luca Delgrossi	luca@ibmpa.awdpa.ibm.com
Ed Ellesson	ellesson@vnet.ibm.com
Eugene Geer	ewg@cc.bellcore.com
Fengmin Gong	gong@concert.net
John Hanratty	jhanratty@agile.com
W.S. Harborth	wharbort@ghost.darpa.mil
Ken Hayward	Ken.Hayward@bnr.ca
Kathy Huber	khuber@wellfleet.com
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@ddn.af.mil
Lee Kilpatrick	leekil@bbn.com
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Charles Kunzinger	kunzinger@vnet.ibm.com

Ted Kuo	tik@vnet.ibm.com
Thomas Maslen	maslen@eng.sun.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Zbigniew Opalka	zopalka@agile.com
Laura Pate	pate@gateway.mitre.org
Michael Patton	map@bbn.com
J. Mark Pullen	mpullen@cs.gmu.edu
Bala Rajagopalan	braja@qsun.att.com
Doris Roland	droland@ghost.darpa.mil
Hal Sandick	sandick@vnet.ibm.com
Henning Schulzrinne	hgs@research.att.com
Dallas Scott	scott@fluky.mitre.org
Michael See	mikesee@vnet.ibm.com
Vincent Shekher	vin@sps.mot.com
Ming Sheu	msheu@vnet.ibm.com
Michael Speer	michael.speer@sun.com
Michael St. Johns	stjohns@arpa.mil
George Swallow	gswallow@bbn.com
John Tavs	tavs@vnet.ibm.com
Matsuaki Terada	tera@sdl.hitachi.co.jp
Akihiro Tominaga	tomy@sfc.wide.ad.jp
Claudio Topolcic	topolcic@cnri.reston.va.us
Keisuke Uehara	kei@cs.uec.ac.jp
Jean Yao	yao@cup.hp.com
Shinichi Yoshida	yoshida@sumitomo.com

2.2.3 IP Over Large Public Data Networks (IPLPDN)

Charter

Chair(s)

George Clapp: clapp@ameris.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

"Determination of Encapsulation of Multi-protocol Datagrams in Circuit-switched Environments", 09/02/1993, K. Sklower <draft-ietf-iplpdn-multi-isdn-02.txt>

"Parameter Negotiation for the Multiprotocol Interconnect", 09/02/1993, K. Sklower, C. Frost <draft-ietf-iplpdn-para-negotiation-02.txt>

"Management Information Base for Frame Relay DTEs", 11/12/1993, C. Brown, F. Baker, C. Carvalho <draft-ietf-iplpdn-frmib-dte-01.txt>

"A Multilink Protocol for Synchronizing the Transmission of Multi-protocol Datagrams.", 09/02/1993, K. Sklower <draft-ietf-iplpdn-simple-multi-01.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"
RFC 1490	"Multiprotocol Interconnect over Frame Relay"

CURRENT MEETING REPORT

Reported by Caralyn Brown/Wellfleet Communications

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

The purpose of re-opening the IP over Large Public Data Networks Working Group (IPLPDN) was to clean up some unresolved items, and attend to those items which have come up after the group became inactive.

Encapsulation Determination

Keith Sklower presented a summary of the Internet-Draft he has written entitled "Determination of Encapsulation of Multi-protocol Datagrams in Circuit-switched Environments." The objective of this work is to define a way in which a receiving station might determine which type of encapsulation (X.25, Frame Relay or PPP) is used on a ISDN call. This is an issue because ISO prefers X.25, PPP is out there, and the ITU has recently included access to a Frame Relay switch as an access feature. The document is not specific to ISDN, but to circuit switched networks where prior configuration is not easily done.

Keith Sklower agreed to update the document to remove part of section 8, "Out of Band Signaling," change bit inversion parameters ("callee's algorithm") and remove section 10. Keith also agreed to clean up the sections referring to Internet-Drafts.

It was agreed that this document should be published as an Informational RFC as a statement of applicability for various standards. This will be done after Keith has updated the document and circulated it to the working group (via the mailing list) for further comments.

Parameter Negotiation

Keith Sklower presented a summary of the Internet-Draft entitled "Parameter Negotiation over Frame Relay." The fundamental issue is to enable the negotiation of a few options in the context of the existing RFC 1490 encapsulation and philosophy. There is a similar document being worked in the Point-to-Point Protocol Extensions Working Group (PPPEXT) called "PPP over Frame Relay." This document preposes that once an NCP is negotiated, the encapsulation changes to the PPP encapsulation with the CF NLPID identifier. Each document presupposes different goals. Parameter negotiation defines how to add certain negotiations to a 1490 environment, while PPP on frame relay attempts to define how to run the entire PPP suite over frame relay.

The forwarding both documents is the fact that two implementations, one using the parameter negotiations document, and one using PPP over frame relay, might successfully complete negotiation and then be unable to pass data due to differing data encapsulations.

The decision was reached within the group that the parameter negotiations document would be modified to clarify that the final data encapsulation would be as specified in RFC 1490 even after negotiations. It would also be clarified to specify that, should an implementation decide to negotiate a protocol for which a PPP encapsulation is defined, but none is defined within RFC 1490 (VJ compression for example), the PPP encoding would be allowed. Protocols which can be defined within the context of RFC 1490 will continue to be encapsulated in that manner.

Status of Updates to RFC 1315

The draft for the updates to RFC 1315 has expired. Caralyn Brown has agreed to repost it and set the wheels in motion to get it forwarded.

Routing Over Frame Relay

Since the disbanding of the original IPLPDN group, there has been much discussion about how to run various protocols over a frame relay network; in particular DECnet over frame relay. The group decided that there are many ways in which to run a protocol over the frame relay network depending upon what the configuration is. Joel Halpern and Fred Baker volunteered to write an Informational document covering experience in partial mesh networks. The document will be posted on the mailing list and discussed there.

Inverse ARP for IPX

The group decided that the definition of InARP for IPX might better be handled by Novell. There were several companies which already have an implementation of InARP for IPX, but the attendees could not remember details. It was decided that the discussion would take place off-line among those who had already implemented InARP for IPX. Caralyn Brown agreed to be editor for a document describing a common method for IPX InARP.

Inverse ARP Extensions

During the IP over ATM discussions, it was felt that InARP was not robust enough. Specifically, a requesting station could not determine whether an InARP request was lost, or the responding station did not have an appropriate answer. It was suggested that InARP be expanded to contain a NAK. The group did not disagree with the suggestion, but decided that, because this problem was related to ATM's ARP server, the IP over Asynchronous Transfer Mode Working Group (ATM) should pursue this work.

IEEE 802.5 Source Routing Over Frame Relay

Those who were most interested in this topic were not present at the meeting. It was decided that this should be taken to the mailing list for further discussion.

Attendees

Anthony Alles	aalles@cisco.com
Fred Baker	fbaker@acc.com
Caralyn Brown	cbrown@wellfleet.com
Steve Buchko	stevebu@newbridge.com
David Carr	dcarr@gandalf.ca
Cheng Chen	chen@accessworks.com
Frank Ciotti	frankc@telxon.com
George Clapp	clapp@ameris.ameritech.com
Thomas Coradetti	tomc@digibd.com
Robert Downs	bdowns@combinet.com
Craig Fox	craig@ftp.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.lkg.dec.com
Joel Halpern	jmh@network.com
B.V. Jagadeesh	bvj@novell.com
Jan-Olof Jemnemo	jan-olof.jemnemo@farsta.trab.se
David Kaufman	dek@magna.telco.com
Stev Knowles	stev@ftp.com
Sundar Kuttalingam	sundark@wiltel.com
William Kwan	kwan@rabbit.com
Thang Lu	tlu@mcimail.com
Andrew Malis	malis@maelstrom.timeplex.com
William Miskovetz	misko@cisco.com
Orly Nicklass	orly@radmail.rad.co.il
Drew Perkins	ddp@fore.com
David Rand	dave_rand@novell.com
Allen Rochkind	Allen_Rochkind@3com.com
Robert Roden	roden@roden.enet.dec.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Chi Shue	chi@casc.com
Keith Sklower	sklower@cs.berkeley.edu
Douglas Williams	dougw@vnet.ibm.com

2.2.4 IP Over AppleTalk (APPLEIP)

Charter

Chair(s)

John Veizades: veizades@ftp.com

Mailing Lists

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

Description of Working Group

The IP Over AppleTalk 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

"AppleTalk Management Information Base II", 04/30/1993, S. Waldbusser, K. Frisa <draft-ietf-appleip-mib2-01.txt>

"KIP AppleTalk/IP Gateway Functionality", 07/06/1993, P. Budne <draft-ietf-appleip-kip-gateway-00.txt, .ps>

Request For Comments

RFC 1243 "AppleTalk Management Information Base"

2.2.5 IP Over Asynchronous Transfer Mode (ATM)

Charter

Chair(s)

Mark Laubach: laubach@hpl.hp.com

Mailing Lists

General Discussion: atm@hpl.hp.com
To Subscribe: atm-request@hpl.hp.com
Archive: Send message to atm-request@hpl.hp.com

Description of Working Group

The IP Over Asynchronous Transfer Mode 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

"Default IP MTU for use over ATM AAL5", 11/16/1993, R. Atkinson <draft-ietf-atm-mtu-05.txt>

"Classical IP and ARP over ATM", 12/22/1993, M. Laubach <draft-ietf-atm-classic-ip-06.txt>

Request For Comments

RFC 1483 "Multiprotocol Encapsulation over ATM Adaptation Layer 5"

CURRENT MEETING REPORT

Reported by Mark Laubach/Hewlett-Packard

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

The Classical Internet-Draft

The "Classical IP and ARP over ATM" (henceforth called "Classical") Internet-Draft Last Call closed on Monday, November 1. All issues raised during the Last Call process were dealt with and closed. One serious technical issue was raised by Dave Sincoskie regarding the ARP table entry timeout and n*n InARP transmission characteristics. A paragraph change was presented and adopted by consensus at the Thursday meeting. The change is as follows:

Under section 8.5 "ATMARP Table Aging," replace paragraph:

> Prior to aging (removing) an ATMARP table entry, all members MUST generate an InARP_REQUEST on any open virtual circuit (VC) associated with that entry. If an InARP_REPLY is received, that table entry is updated and not deleted. If there is no open VC associated with the table entry, the entry is deleted.

With the following two paragraphs:

> Prior to aging an ATMARP table entry, an ATMARP server MUST generate an InARP_REQUEST on any open VC associated with that entry. If an InARP_REPLY is received, that table entry is updated and not deleted. If there is no open VC associated with the table entry, the entry is deleted.
>
> When an ATMARP table entry ages, an ATMARP client MUST invalidate the table entry. If there is no open VC associated with the invalidated entry, that entry is deleted. In the case of an invalidated entry and an open VC, the ATMARP client must revalidate the entry prior to transmitting any non address resolution traffic on that VC. In the case of a PVC, the client validates the entry by transmitting an InARP_REQUEST and updating the entry on receipt of an InARP_REPLY. In the case of an SVC, the client validates the entry by transmitting an ARP_REQUEST to the ATMARP Server and updating the entry on receipt of an ARP_REPLY. If a VC with an associated invalidated ATMARP table entry is closed, that table entry is removed.

Dave Piscitello approved the change process; another Last Call is not needed. The Classical Internet-Draft is awaiting IESG ballot.

Routing over Large Clouds Working Group Introduction

Joel Halpern gave a presentation of the proposed charter of the Routing Over Large Clouds Working Group (ROLC). Juha's NBMA ARP has been moved into that working group. Issues involved with ARPing beyond the LIS and shortcut routing, et al. for IP over ATM are now in ROLC.

The MTU Internet-Draft

Ran Atkinson presented his Internet-Draft, "Default IP MTU for use over ATM AAL5" (henceforth called "MTU"). There was much discussion over the use of SDU negotiation. Dan Grossman suggested that advantage should be taken of whatever signaling support is available and make it mandatory for SVC negotiations. The working group needs to specify the parameters of UNI 3.0 so that interoperable implementations exist. The issue was raised that a very clearly defined default case exists (classical model) and it is necessary to have a clear plan of how signaling is used, for what, and what the defaults are.

A discussion of the MTU path discovery requirement took place. The question of whether system requirements (IP systems) can be driven by requiring it in IP over ATM was raised. Ran feels that an on/off switch is a implementation optimization; i.e., up to the implementor. Others feel that it is not the ATM Working Group's place to require it. The group reached the following recommendation: use the default MTU size of 9180. IP stations must implement MTU path discovery but are not required to use it. If they do use it, the MTU size may be adjusted, etc.

Ran will be updating the document soon. The MTU path discovery issue is still being debated.

Framework Document

Bob Cole led a discussion of the framework document. Joel Halpern led a short presentation on TUNIC and TULIP. Discussion was plentiful on all issues. Bob will be seeking volunteers for help with a new version. The working group chair hopes that this document can be turned into a planning guide for the working group. Discussions will continue on the mailing list.

Security and Reliability

Bryan Lyles presented a brief introduction of security issues with regards to IP over ATM, in that a firewall-level mechanism is needed that allows certain streams to go through a firewall. Also, as trends will want to multiplex a VC higher in the protocol stack (e.g., TCP ports, or higher) reliability of the VC must be understood. A reliable peer protocol cannot be replaced with an unreliable VC. These issues were presented to the working group as a consideration of areas that might be worked on in the future.

Wrap Up

The group hosted other discussions on source address, the non-optimal behavior of InARP, selectors and multiple LIS's, application binding, and Q.93B parameters.

There was not enough time to complete discussion on the issue of IP over the ATM Forum's LAN Emulation specification.

Action items for the group are:

- Ran and Bob each incorporate comments from the meeting into their respective documents.
- Dan Grossman, Mike Goguen, and George Swallow are forming a small design team to generate an Informational document on how to use the UNI 3.0 for IP over ATM. Sufficient information will be presented to enable consistent implementations but not to duplicate ATM Forum specifications.
- Bryan Lyles and Drew Perkins will collaborate on a draft statement for the framework document on possible methods of supporting IP multicast.
- Andy Malis will follow through on the multiple VC thrashing issue and will generate consensus.

Attendees

Masuma Ahmed	mxa@mail.bellcore.com
Nick Alfano	alfano@mpr.ca
Anthony Alles	aalles@cisco.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Nutan Behki	nebhki@newbridge.com
Ram Bhide	
Richard Binder	rbinder@cnri.reston.va.us
Jon Boone	boone@psc.edu
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Al Broscius	broscius@bellcore.com
Caralyn Brown	cbrown@wellfleet.com
Theodore Brunner	tob@thumper.bellcore.com
Steve Buchko	stevebu@newbridge.com

Glen Cairns	cairns@mprgate.mpr.ca
Enke Chen	enke@merit.edu
Ping Chen	ping@ping2.aux.apple.com
George Clapp	clapp@ameris.ameritech.com
Robert Cole	rgc@qsun.att.com
Rob Coltun	rcoltun@ni.umd.edu
Michael Conn	4387451@mcimail.com
Matt Crawford	crawdad@fncent.fnal.gov
James Davin	davin@thumper.bellcore.com
Thomas Dimitri	tommyd@microsoft.com
Kurt Dobbins	dobbins@ctron.com
Waychi Doo	wcd@berlioz.nsc.com
David Dubois	dad@pacersoft.com
Pierre Dupont	dupont@mdd.comm.mot.com
Dario Ercole	Dario.Ercole@cselt.stet.it
Julio Escobar	jescobar@bbn.com
Stefan Fassbender	stf@easi.net
Steve Feldman	feldman@mfsdatanet.com
William Fenner	fenner@cmf.nrl.navy.mil
Robert Fenoglio	fenoglio@vnet.ibm.com
Carlos Fernandez	carlos@plk.af.mil
Robert Fink	rlfink@lbl.gov
James Forster	forster@cisco.com
Dan Frommer	dan@isv.dec.com
Mark Garrett	mwg@faline.bellcore.com
John Gawf	gawf@compatible.com
Eugene Geer	ewg@cc.bellcore.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Mike Goguen	goguen@synoptics.com
Fengmin Gong	gong@concert.net
Ramesh Govindan	rxg@thumper.bellcore.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Robert Grow	bob@xlnt.com
Stuart Hale	stu_hale@vnet.ibm.com
Joel Halpern	jmh@network.com
John Hanratty	jhanratty@agile.com
W.S. Harborth	wharbort@ghost.darpa.mil
Dimitry Haskin	dhaskin@wellfleet.com
Marc Hasson	marc@mentat.com
Ken Hayward	Ken.Hayward@bnr.ca
Juha Heinanen	juha.heinanen@datanet.tele.fi
Kathryn Hill	khill@newbridge.com
Eric Hoffman	hoffman@cmf.nrl.navy.mil
Mike Holloway	mikeh@newbridge.com
Refael Horev	horev@lannet.com

Kathy Huber	khuber@wellfleet.com
Melanie Humphrey	msh@uiuc.edu
Ronald Jacoby	rj@sgi.com
B.V. Jagadeesh	bvj@novell.com
Merike Kaeo	mkaeo@cisco.com
Yasuhiro Katsube	katsube@mail.bellcore.com
David Kaufman	dek@magna.telco.com
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Charles Kunzinger	kunzinger@vnet.ibm.com
Ted Kuo	tik@ececho.ncsu.edu
Sundar Kuttalingam	sundark@wiltel.com
Olav Kvittem	Olav.Kvittem@uninett.no
Mark Laubach	laubach@hpl.hp.com
Joseph Liu	jliu@atg.wiltel.com
Kim Long	klong@sura.net
Thang Lu	tlu@mcimail.com
Bryan Lyles	lyles@parc.xerox.com
Andrew Malis	malis@maelstrom.timeplex.com
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@cmf.nrl.navy.mil
Matt Mathis	mathis@psc.edu
Jun Matsukata	jm@eng.isas.ac.jp
Keith McCloghrie	kzm@hls.com
Donald Merritt	don@arl.army.mil
Orly Nicklass	orly@radmail.rad.co.il
Dan Nordell	
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Zbigniew Opalka	zopalka@agile.com
Steve Parker	sparker@ossi.com
Craig Partridge	craig@bbn.com
Laura Pate	pate@gateway.mitre.org
Maryann Perez	perez@cmf.nrl.navy.mil
Charles Perkins	perk@watson.ibm.com
Drew Perkins	ddp@fore.com
Radia Perlman	perlman@novell.com
David Piscitello	wk04464@worldlink.com
Robert Roden	roden@roden.enet.dec.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Doris Roland	droland@ghost.darpa.mil
Allan Rubens	acr@merit.edu
Timothy Salo	tjs@msc.edu
Hal Sandick	sandick@vnet.ibm.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Martin Schulman	schulman@smtp.sprint.com
Isil Sebuktekin	isil@nevin.bellcore.com

Michael See	mikesee@vnet.ibm.com
Kanan Shah	kshah@cmf.nrl.navy.mil
Satya Sharma	ssharma@chang.austin.ibm.com
Vincent Shekher	vin@sps.mot.com
Ming Sheu	msheu@vnet.ibm.com
Uttam Shikarpur	uttam@zk3.dec.com
Chi Shue	chi@casc.com
W. David Sincoskie	sincos@bellcore.com
Andrew Smith	asmith@synoptics.com
Tae Song	tae@novell.com
Steve Suzuki	steve@fet.com
George Swallow	gswallow@bbn.com
Matsuaki Terada	tera@sdl.hitachi.co.jp
Susan Thomson	set@bellcore.com
Hoe Trinh	htrinh@vnet.ibm.com
Keisuke Uehara	kei@cs.uec.ac.jp
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Thomas Walsh	tomw@kalpana.com
Chuck Warlick	chuck.warlick@pscni.nasa.gov
Guy Wells	guyw@uswest.com
Douglas Williams	dougw@vnet.ibm.com
Honda Wu	honda@nat.com
Liang Wu	ltwu@bellcore.com
Shinichi Yoshida	yoshida@sumitomo.com
Jessica Yu	jyy@merit.edu
Chin Yuan	cxyuan@pacbell.com
Mauro Zallocco	mdz@netlink.com
Weiping Zhao	zhao@nacsis.ac.jp

2.2.6 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

"PPP LCP Extensions", 09/07/1993, W. Simpson <draft-ietf-pppext-lcpext-04.txt>

"PPP over ISDN", 10/14/1993, W. Simpson <draft-ietf-pppext-isdn-03.txt>

"PPP in Frame Relay", 10/07/1993, W. Simpson <draft-ietf-pppext-frame-relay-02.txt>

"PPP over SONET/SDH", 09/22/1993, W. Simpson <draft-ietf-pppext-sonet-01.txt>

"PPP in X.25", 10/07/1993, W. Simpson <draft-ietf-pppext-x25-02.txt>

"PPP Bridging Control Protocol (BCP)", 11/12/1993, F. Baker, R. Bowen <draft-ietf-pppext-for-bridging-02.txt>

"The PPP Multilink Control Protocol (MCP)", 11/22/1993, K. Sklower, D. Carr <draft-ietf-pppext-multilink-04.txt>

"The PPP NetBIOS Frames Control Protocol (NBFCP)", 12/21/1993, T. Dimitri <draft-ietf-pppext-netbios-fcp-03.txt>

"PPP Reliable Transmission", 10/06/1993, D. Rand <draft-ietf-pppext-reliable-00.txt>

"PPP Stacker LZS Compression Protocol", 10/20/1993, R. Lutz <draft-ietf-pppext-stacker-00.txt>

"The PPP Compression Control Protocol (CCP)", 12/22/1993, D. Rand <draft-ietf-pppext-compression-03.txt>

"PPP Gandalf FZA Compression Protocol", 10/26/1993, D. Carr <draft-ietf-pppext-gandalf-00.txt>

"PPP Hewlett-Packard Packet-by-Packet Compression (HP PPC) Protocol", 10/29/1993, J. Petty <draft-ietf-pppext-hpppc-00.txt>

"PPP Predictor Compression Protocol", 12/15/1993, D. Rand <draft-ietf-pppext-predictor-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"

RFC 1547 "Requirements for an Internet Standard Point-to-Point Protocol"

RFC 1548 "The Point-to-Point Protocol (PPP)"

RFC 1549	"PPP in HDLC Framing"
RFC 1552	"The PPP Internetwork Packet Exchange Control Protocol (IPXCP)"
RFC 1553	"Compressing IPX Headers Over WAN Media (CIPX)"

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

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

Two documents were referred, without discussion, to the IESG for consideration as Proposed Standards.

- "PPP over ISDN" (draft-ietf-pppext-isdn-03.txt)
- "PPP over SONETE/SDH" (draft-ietf-pppext-sonet-01.txt)

The following Drafts generated quite a bit of discussion. *Editor's Note: Details of the discussion are available via FTP or mail server from the remote directories as /ietf/pppext/pppext-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.*

- "PPP over X.25" (draft-ietf-pppext-x25-02.txt)

  There was some discussion on whether certain language should be changed in the document. It was decided that no revisions were required and the document will be recommended to the IESG for consideration as a Proposed Standard.

- "PPP in Frame Relay" (draft-ietf-pppext-frame-relay-02.txt)

  The discussion of the Frame Relay document led to a recommendation that a new sentence would clarify the requirement that a system re-negotiate if it sees an encapsulation it was not expecting. Several options were proposed and the one favored by the largest number of attendees was Option 2, "If the negotiations are performed on a medium that has a default encapsulation, default to the media's preferred encapsulation type. Provide an LCP option to go back to PPP (0xCF) encapsulation."

  Given this option, it is recommended that the single sentence be added to draft-ietf-pppext-frame-relay-02.txt, and the resulting draft-ietf-pppext-frame-relay-03.txt be considered by the IESG as a Proposed Standard.

The obvious place to put this option is "PPP LCP Extensions" (draft-ietf-pppext-lcpext-04.txt), but it contains other work that has been waiting and needs to be moved forward. Therefore, the recommendation is that draft-ietf-pppext-lcpext-04.txt be considered by the IESG as a Proposed Standard, and another document will be drawn up describing the LCP option for negotiation of encapsulations.

> [A note from the PPPEXT Chair: It is not clear that the group reached an effective consensus concerning the default encapsulation, or that this consensus represents the many members of the PPP Working Group who were not in the meeting. It was stated clearly and unanimously conceded in the meeting that

the indeterminate interaction with RFC 1490 systems is only of concern if the default data encapsulation is 1490-style; if the negotiation results in the use of the PPP encapsulation, and given the renegotiation on receipt of the other encapsulation, there is no ambiguity. The members of IPLPDN present in the meeting stated that they found the ambiguity acceptable because it enabled them to not change their micro code for their routers, to which the counter-argument was made that to continue using the 1490 encapsulation they need only not negotiate the indicated NCP. The chair observes that there is also a backward compatibility issue; by the time the working group agrees on the LCP option and publishes a document, there will assuredly be compliant PPP/Frame Relay implementations fielded, which will be using the 0xCF data encapsulation it recommends. The chair also notes, without prompting from the members of the working group, that it is as easy for one political camp to negotiate the option as it is for the other, so the argument that the default *must* be to use 1490 encapsulation after the NCP has been negotiated appears weak. The chair further notes that the PPP encapsulation inside a compressed or multi-link data stream is (by specification) the PPP encapsulation without any address/control field, requiring Frame Relay system to recognize the encapsulation anyway if they use the PPP features that they wish to import.

The chair notes that he has sought throughout this debate to mediate a strong disagreement between two working groups, and give each what they wish out of it. The objective facts seem to suggest that the LCP option should negotiate the use of a non-PPP encapsulation after the negotiation of the NCP, as the use of the PPP encapsulation is provably correct and the other—a point freely admitted by the proponents of the other position—is not. This is the most important attribute of all, and should, in his opinion, drive the debate to its conclusion.

The chair's recommendation (to be freely and spiritedly debated by all who wish) is that the document describing the option should be drawn up by Bill Simpson, indicating that the default encapsulation is the provably correct PPP encapsulation, but that the other is negotiable. The updated PPP/Frame Relay document and the document describing this LCP option should become Proposed Standards together.]

Day 2 - Further Document Review

Dave Rand presented the "PPP Reliable Transmission" document, (draft-ietf-pppext-reliable-00.txt). After some discussion, the document was recommended for consideration by the IESG as a Proposed Standard.

Dave also presented "The PPP Compression Control Protocol (CCP)", (draft-ietf-pppext-compression-01.txt). Numerous changes were recommended by the working group, separating the "Predictor" algorithm into a separate document, and modifying the structure of the CCP options. These will be edited into a new draft, which will be posted in the Internet-Drafts directory for discussion. It is anticipated that this work can be sent to the IESG before year end.

Rich Bowen then presented the updated "PPP Bridging Control Protocol (BCP)" document (draft-ietf-pppext-for-bridging-01.txt). Minor revisions were suggested. It is anticipated that this will go to the IESG by year end.

Thomas Dimitri presented his NETBEUI/PPP proposal, "The PPP NetBIOS Frames Control Protocol (NBFCP)" (draft-ietf-pppext-netbios-fcp-00.txt). This was cut short due to time constraints and will be taken to the list.

Keith Sklower discussed "The PPP Multilink Control Protocol (MCP)" (draft-ietf-pppext-multilink-02.txt), that he had mailed to the list just before the IETF meeting. This discussion continued with key players during lunch. He will post the draft (draft-ietf-pppext-multilink-03.txt) for discussion; it is anticipated that this work will be ready for IESG consideration by year end.

The chair had planned to discuss the plan for the PPPEXT Working Group for the coming two years, but was unable to do so due to lack of time. This matter will be taken to the list.

Attendees

Andy Adams	ala@merit.edu
James Allard	jallard@microsoft.com
Fred Baker	fbaker@acc.com
Rich Bowen	rkb@ralvm11.vnet.ibm.com
Caralyn Brown	cbrown@wellfleet.com
Steve Buchko	stevebu@newbridge.com
David Carr	dcarr@gandalf.ca
Cheng Chen	chen@accessworks.com
Chris Chiotasso	chris@lightstream.com
George Clapp	clapp@ameris.ameritech.com
Thomas Coradetti	tomc@digibd.com
Jonathan Didner	jonb@bangate.compaq.com
Thomas Dimitri	tommyd@microsoft.com
Robert Downs	bdowns@combinet.com
Craig Fox	craig@ftp.com
Richard Fox	rfox@metricom.com
John Gawf	gawf@compatible.com
Shawn Gillam	shawn@timonware.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.lkg.dec.com
Joel Halpern	jmh@network.com

Ronald Jacoby	rj@sgi.com
B.V. Jagadeesh	bvj@novell.com
Jan-Olof Jemnemo	jan-olof.jemnemo@farsta.trab.se
David Kaufman	dek@magna.telco.com
Robert Lutz	
Andrew Malis	malis@maelstrom.timeplex.com
Glenn McGregor	ghm@lloyd.com
William Miskovetz	misko@cisco.com
Dennis Morris	morris@altair.disa.mil
Andy Nicholson	andyni@microsoft.com
Todd Palgut	todd@nei.com
Eric Peterson	elpeterson@eng.xyplex.com
James Philippou	japhilippou@eng.xyplex.com
Venkat Prasad	vsp@3com.com
David Rand	dave_rand@novell.com
Kenneth Rehbehn	kjr@netrix.com
Allen Rochkind	Allen_Rochkind@3com.com
Robert Roden	roden@roden.enet.dec.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Paul Serice	serice@cos.com
Satya Sharma	ssharma@chang.austin.ibm.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Timon Sloane	timon@timonware.com
Steve Suzuki	steve@fet.com
Thomas Walsh	tomw@kalpana.com
James Watt	james@newbridge.com
Bradley Wilson	wilson@ftp.com
Honda Wu	honda@nat.com
Mauro Zallocco	mdz@netlink.com

Point to Point Working Group

Fred Baker, Chair

In the IESG's In Basket Now

- LCP Extensions pppext-lcpext-04.txt -> Proposed

Let's move on these

- PPP/ISDN pppext-isdn-03.txt -> Proposed
- PPP/X.25 pppext-x25-02.txt -> Proposed
- PPP/SONET pppext-sonet-01.txt -> Proposed
- PPP/FR pppext-frame-relay-02.txt -> Proposed

PPP/FR pppext-frame-relay-02.txt -> Proposed

- Interaction with RFC 1490 is indeterminate in some failure cases
- To make it determinate, add a sentence clarifying that receipt of a data frame in the wrong encapsulation forces a renegotiation of LCP
- New option to the LCP Extensions document

New option to the LCP Extensions document

- select data protocol encapsulation
- default is "native" encapsulation for interface type
  - Frame Relay uses 1490 as "native"
  - X.25 uses RFC 1356 encapsulation as "native"
  - HDLC and Asynchronous links have no alternative, use PPP encapsulation
- option may negotiate use of PPP alternative encapsulation

Current Work

- Reliable PPP Links
- PPP Message Compression
- Bridging On PPP
- NETBEUI On PPP
- Multiple PPP Links

Reliable PPP Links

pppext-reliable-00.txt, Dave Rand, Novell

- Changes Address/Control to Numbered Mode (LAPB)
- LCP Configuration Option:
  - Window, Address
- Does not recommend ISO Multilink
- LAPB Parameter defaults
- Implementations of PPP/LAPB

Implementations of PPP/LAPB

- 3COM
- Novell
- Gandalf

PPP Message Compression

pppext-compression-01.txt, Dave Rand, Novell

- Compression Algorithms Documented
- Issues in Compression Document
- Resolutions

Compression Algorithms Documented

- pppext-gandalf-00.txt (FZA)
- pppext-hpppc-00.txt (HP Compression)
- pppext-stacker-00.txt (Stac Compression)
- Others being defined
  - Microsoft
  - UNIX Compress

Issues in Compression Document

comments on the list

- Option Numbering and Structure
- Behavior of Reject
- Behavior of NAK
- Predictor Compression Algorithm

Resolutions

- separate "predictor" into a separate document
- negotiate "I will keep no history"
- text regarding licensing decompression.
- code for LZW decompressor in main CCP document?
- option format
- sequence of configuration messages

option format

- standard algorithms have an option type per algorithm (1..254)
- encoding of option is type/length/value
- one option for all OUIs, essentially = current #2, 0 -> type
- 255 is undefined
- encoding of sub-options is type/length/value
- list OUI options separately, follow 1331 procedure

sequence of configuration messages

- receiver offers acceptable encodings
- reject case
  - sender rejects a subset of those
  - receiver offers ONE encoding among those that remain
- if I can do all:
  - ACK accepts first algorithm listed

Bridging On PPP

pppext-for-bridging-01.txt, Rich Bowen, IBM

- Updates text in document
- Changes to options

Updates text in document

- Explanatory text of Source Routing updated
- Sample code moved to an appendix

Changes to options

- Existing Options Clarified/Updated
- MAC-Address Negotiation
- Spanning-Tree-Protocol Negotiation

Existing Options Clarified/Updated

- Clarified Source Routing Text
- MAC-Support
  - Option negotiation procedure per RFC 1334
- Tinygram-Compression
  - no change
- LAN-Identification
  - no change, but much additional text courtesy Network Systems

Clarified Source Routing Text

- Bridge-Identification (half bridge model)
- Line-Identification (whole bridge model)
- default is line-identified
- in half bridge, sender modifies RIF

MAC-Address Negotiation

- define or announce my MAC Address on this interface

Spanning-Tree-Protocol Negotiation

- 802.1(d) BPDU
- 802.1(g) BPDU
- IBM BPDU
- DEC STP

NETBEUI On PPP
pppext-netbios-fcp-01.txt, Thomas Dimitri, Microsoft

Multiple PPP Links
pppext-multilink-02.txt, Keith Sklower, UC Berkeley

- New Draft
- Remaining issues

New Draft

- removes:
  - reset-on-loss
  - maximum-completed-received-sequence
- retain the sequenced-delivery option.
- drop Maximum Receive Reconstructed Unit option?

Remaining issues

- if LAPB is underneath, may have out of sequence packets one could send
- therefore allow labelled resequenced packets

Way Forward

Objective: Bring all work to Full Standard inside two years

- Link Control Protocol
- Security
- IPCP
- Bridging
- Novell IPX
- AppleTalk RFC 1378
- OSI RFC 1377
- DECNET RFC 1376

Link Control Protocol

Bill Simpson

- HDLC Framing
  - Revision of RFC 1331
  - Draft Standard
- LCP
  - Revision of RFC 1331
  - Draft Standard
- LCP Extensions
- LQM RFC 1333
- MIB RFC 1471

Security

Bill Simpson

- Authentication
  - RFC 1334
- MIB RFC 1472

IPCP

Glenn McGreggor, Lloyd & Associates

- Document
  - RFC 1332
- MIB RFC 1473
- updates to RFC? Usage Document?

Bridging

Rich Bowen, IBM

- Document
  - (Son-of-1220)
- MIB (RFC 1474)

Novell IPX
Mark Lewis, Telebit

- IPXWAN - Novell NCP
  - Informational
- IPXCP - IPX NCP
  - Awaiting RFC #
- CIPX - Compressed IPX
  - Awaiting RFC #

AppleTalk RFC 1378
Brad Parker

- Brad indicated that there were some changes he wanted to make

OSI RFC 1377
Dave Katz, Cisco

- implementations include at least:
  - Cisco
  - 3COM
  - Wellfleet

DECNET RFC 1376
Steve Senum, Network Systems

- implementations include at least:
  - Cisco
  - 3COM
  - Wellfleet
  - ACC
  - Network Systems

NBF CP

- NBF CP is the PPP Control Protocol for the NBF and NetBEUI NetBIOS protocols.
  - Negotiates Addition of NetBIOS Names
  - Negotiates Multicast-Filtering
  - Provides Peer Information

A PPP NetBIOS Architecture

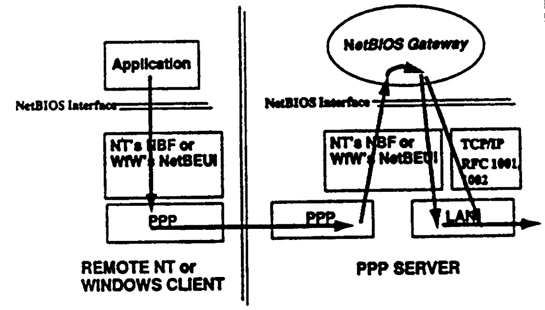

How NetBIOS Names Are Addded

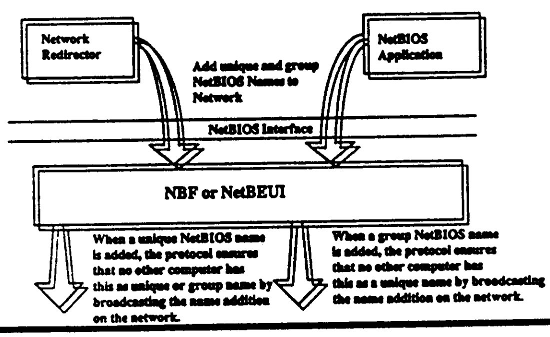

How NBF CP Negotiates NetBIOS Name Additions

- PPP Peer dialing in already has unique and group NetBIOS names added.
- The remote PPP peer must ensure that there are no name conflicts by attempting to add all unique and group names to the local network.
- If the remote PPP peer attempts to add NetBIOS names, the NetBIOS name return codes are sent in the Configure-Ack frame, else a Configure-Nak is sent.
- Based by which names the remote PPP peer was able to add, the local PPP peer may terminate NBF CP by sending an NBF CP Terminate-Request.

Some NetBIOS Final Return Codes when adding NetBIOS names:

00H	Good return
03H	Invalid Name
0DH	Duplicate name in local name table
0EH	Name table is full
15H	Name not found, cannot specify "*" or 00H
16H	Name in use by remote adapter
21H	Interface busy
22H	Too many commands outstanding
25H	Reserve name specified

What Else NBF CP Negotiates

Peer Information - Largely used for informational purposes only. This field is advisory in nature and will be most likely be used for User Interface or Logging purposes.

- Peer Class (Gateway, Bridge, Client only)
- Peer Version Number
- Peer Name (NetBIOS unique machine name)

Multicast-Filtering - Used to define the handling of multicast packets.

- Multicast-Forward-Period - Used to filter multicast traffic. Everytime a computer attempts to add a NetBIOS name or sends a multicast or group datagram, a multicast packet will be sent.
- Priority - If the remote peer cannot miss any multicast packets because it uses apps which use multicast datagrams (or other multicast packets) this feature can be used to set the multicast packet priority high.

2.2.7 Router Requirements (RREQ)

Charter

Chair(s)

Philip Almquist: `almquist@jessica.stanford.edu`

Mailing Lists

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

Description of Working Group

The Router Requirements Working Group has the goal of rewriting the existing Router Requirements RFC, RFC 1009, and a) bringing it up to the organizational and requirement explicitness levels of the Host Requirements 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", 01/03/1994, F. Kastenholz <draft-ietf-rreq-iprouters-require-00.txt>

Request For Comments

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

RFC 1354 "IP Forwarding Table MIB"

2.3 IP: Next Generation Area

Directors:

- Scott Bradner: sob@harvard.edu
- Allison Mankin: mankin@cmf.nrl.navy.mil

Area Summary reported by Allison Mankin/NRL

The IPNG Area co-Directors gave a plenary presentation on their plan for the IPng decision process and introduced the IPng directorate.

Frank Solensky gave an introduction to the ALE BOF and the three IPng proposals gave status reports during the IPng plenary session; Peter Ford gave a status report on TUBA, Steve Deering gave an overview of SIPP, and Rob Ullmann gave an overview of TP/IX (also known as CATNIP).

Address Lifetime Expectations BOF (ALE)

Phill Gross gave an update to the presentation he and Dennis Ferguson prepared for INET '93 describing the growth of the Internet (in terms of both assigned addresses and connected networks) and presented some recommendations for increasing the efficiency of how IP addresses are deployed. A lively discussion ensued.

The working group will be formed, combining resources with CIDR deployment. The emphasis will be on the measurement and projections, evaluating the potential impact of recommendations rather than formulating recommendations itself.

There is also a pressing need to collect more information; all known projections are based on incomplete data.

P. Internet Protocol Working Group (PIP) and Simple Internet Protocol Working Group (SIP)

The PIP and SIP Working Groups have combined their efforts and the working groups will be merged into a new working group called Simple Internet Protocol Plus (SIPP). The two working groups met in two combined sessions co-chaired by Steve Deering, Paul Francis, and Bob Hinden.

At the first session Steve Deering presented an overview of the SIP/PIP Merger. This included the motivation behind the merger, benefits of the merger, and described the new features of SIPP. The purpose of the merger is to keep the simplicity and transition features of SIP and to benefit from the advanced routing capabilities of Pip—while making them easier to use and to understand.

Following this Paul Francis presented the SIPP routing and addressing. This included a description of address sequences and how they are used for mobility, provider selection, and extended addressing. Ramesh Govindan presented detailed examples of these usages of SIPP address sequences.

A overview of the new IPAE draft was given by Bob Gilligan. He gave a short overview of IPAE, and discussed and resolved several open issues.

Bill Simpson presented the current state of his work on SIPP neighbor discovery. It focuses on a "where are you" and "I am Here" functions with optional extensions for additional functionality.

During the second session Rob Coltun presented his proposal for a version of OSPF for SIPP. The group concluded that he should focus on just extending OSPF to support 64-bit addresses and defer the work to add additional levels of hierarchy. The latter work should be presented to the Open Shortest Path First IGP Working Group (OSPF).

Sue Thompson presented her proposal for DNS changes to support SIPP. The group concluded that this was the correct approach for SIPP DNS records.

Jim Bound presented his thoughts on the changes required to the Dynamic Host Configuration Protocol. There was an extended discussion which resulted in general agreement that auto configuration was a key part of any IPng.

Paul Francis presented a proposal for provider based address assignment. After an interesting discussion, the group agreed to proceed with this approach.

TCP/UDP Over CLNP-Addressed Networks Working Group (TUBA)

Dave Marlow reported that the CLNP Multicast work has made progress in ISO. Changes exist in addressing, CLNP, ES-IS and the network service definition. Group addressing is a full standard, other changes are in ballot at this time.

Ross Callon discussed the revised NSAP Addressing Guidelines document and took an action item to make the document somewhat less "backbone-centric."

CLNP mobility was discussed. Mark Knopper briefly described CDPD, a specification for cellular mobile data service from a consortium of cellular carriers. It uses CLNP as the primary protocol, and provides IP service using IP-over-CLNP encapsulation. The mobility protocol is similar to ongoing work in the IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP). The group discussed whether or not it should be proactive, or wait for the MOBILEIP Working Group to settle. Yakov Rekhter and Dave Piscitello agreed to recast the mobile IP document in terms of CLNP and publish it as an Internet-Draft.

Yakov Rekhter described his work on a method for transparently adding options to CLNP. It codes which options are required to be processed by routers and/or hosts, even when the option is otherwise unrecognized. He also described work on strong versus weak QoS forwarding.

Dave Katz spoke about the outcome of the Extensions to OSI for use in the Internet BOF (OSIEXTND) that was held in Amsterdam. The net effect of IESG policies is that the work will progress within the TUBA Working Group. Dave Katz then presented an extension to the standard dynamic NSAP address assignment function, which would allow the end system to suggest a system ID for itself.

Peter Ford presented his draft document on the Dual Stack Transition plan. It is an "inside out" approach that begins with infrastructure deployment. It was pointed out that this transition framework needs to be completed as soon as possible.

TCP/UDP Over CLNP-Addressed Networks Working Group (TUBA) and TP/IX Working Group (TPIX)

The TUBA Working Group met in joint session with the the TPIX Working Group. TPIX then continued on to a separate session in the same room.

Robert Ullmann presented his new proposal Common Architecture For The New Internet (CATNIP). The new proposal is based on RFC 1475. CATNIP is designed to use header compression by including a flow cache ID or "handle" in its header. It also uses a NSAP style of addressing. The joint meeting was held to explore commonality between CATNIP and TUBA proposals.

The group came up with the following list of milestones:

- Submit the CATNIP proposal as an Internet-Draft
- Rewrite the TPIX Working Group charter to realign it with the new proposal
- Possibly rename the TPIX Working Group to the CATNIP Working Group

CURRENT MEETING REPORT

Reported by Frank Solensky/FTP Software, Inc.

Minutes of the Address Lifetime Expectations BOF (ALE)

Phill Gross gave an updated version of the presentation he and Dennis Ferguson prepared for INET '93. It included an extensive analysis and projections of the growth of the Internet and also provided an estimate on how efficiently IP addresses were being assigned. The presentation concluded with several ideas on how the address space utilization could be improved after CIDR has been deployed:

- Reclaiming IP network numbers which are assigned but not connected to the Internet.
- Tougher address assignment policies.
- Encouraging connected networks to renumber into a smaller portion of either their assigned net number (freeing up the rest of that net number for reassignment) or within their provider's range of addresses (removing the need to announce the original net number between providers).

The ensuing discussion was invaluable for setting the direction of the ongoing analysis. Some of the points that came out of that discussion were:

- A better feel is needed for the accuracy of the data that has been collected thus far. For example, it was discovered during the discussion that some of the regional carriers have not been reporting IP address assignments back to the NIC since there was not any place on their templates to provide this information.
- The IP address assignment and routing table size problems must be kept separate from each other at all times. These are two different problems that have different factors driving them.
- It is necessary to carefully avoid combining the data collected under different policy ranges into a single trend line.
- Any graph produced that extends a trend line into the future must always carry the caveat that it is based on historic information. It cannot be predicted how future technologies or applications might further impact the growth of the Internet.
- Routing table statistics should be collected from a wider variety of sources, since the Internet can no longer be thought of as having a single backbone.

- It is necessary to evaluate and gather data about the impact (e.g. cost, expected gains) of any policy changes before the group makes specific recommendations.

- All of the providers must buy into any recommendations made that affect them. For example, if most but not all network providers suggest renumbering to their customers, the entire benefit could be lost if one "bad guy" provider decides to get new customers by allowing them to connect using their old numbers.

Frank Solensky also presented some of his more recent findings on the effect of CIDR on the routing table size (this also appeared in Tony Li's technical presentation on 'CIDR Status' earlier in the day) and the trend lines for the proportion of the total address space that announced Class B and Class C net numbers consume. He also cautioned that the switch from announcing single Class B to multiple Class C network numbers has occurred too recently for the trend line to be considered reliable.

The creation of the ALE Working Group charter was deferred as an action for the mailing list. Discussions are now wrapping up and the charter should be completed shortly.

Attendees

Andy Adams	ala@merit.edu
James Allard	jallard@microsoft.com
Vadim Antonov	avg@icm1.icp.net
Susie Armstrong	susie@mentat.com
William Barns	barns@gateway.mitre.org
Tony Bates	tony@ripe.net
Nutan Behki	nebhki@newbridge.com
David Borman	dab@cray.com
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Jim Bound	bound@zk3.dec.com
Scott Bradner	sob@harvard.edu
David Bridgham	dab@epilogue.com
Ronald Broersma	ron@nosc.mil
Steve Buchko	stevebu@newbridge.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	rcallon@wellfleet.com
Peter Cameron	cameron@xylint.co.uk
George Chang	gkc@ctt.bellcore.com
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
J. Noel Chiappa	jnc@lcs.mit.edu
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
Richard Colella	colella@nist.gov

Matt Crawford	crawdad@fncent.fnal.gov
John Curran	jcurran@nic.near.net
James Davin	davin@thumper.bellcore.com
Michael Davis	mike@dss.com
Chuck de Sostoa	chuckd@cup.hp.com
Donald Eastlake	dee@skidrow.lkg.dec.com
Havard Eidnes	havard.eidnes@runit.sintef.no
Robert Enger	enger@seka.reston.ans.net
Dario Ercole	Dario.Ercole@cselt.stet.it
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Carlos Fernandez	carlos@plk.af.mil
Dale Finkelson	dmf@westie.mid.net
Eric Fleischman	ericf@atc.boeing.com
Peter Ford	peter@goshawk.lanl.gov
Dan Frommer	dan@isv.dec.com
Vince Fuller	vaf@barrnet.net
John Gawf	gawf@compatible.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Daniel Hansen	dan@ngc.com
Herluf Hansen	hha@tbit.dk
Dimitry Haskin	dhaskin@wellfleet.com
Eugene Hastings	hastings@psc.edu
Denise Heagerty	denise@dxcoms.cern.ch
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Matt Hood	hood@nsipo.nasa.gov
Kathy Huber	khuber@wellfleet.com
Phil Irey	pirey@relay.nswc.navy.mil
Ronald Jacoby	rj@sgi.com
Dale Johnson	dsj@merit.edu
Matthew Jonson	jonson@ddn.af.mil
Merike Kaeo	mkaeo@cisco.com
Jeanine Kamerdze	kamerdze@nsipo.nasa.gov
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Sean Kennedy	liam@nic.near.net
Edwin King	eek@atc.boeing.com
John Krawczyk	jkrawczy@wellfleet.com
Sundar Kuttalingam	sundark@wiltel.com
Olav Kvittem	Olav.Kvittem@uninett.no
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se
Thang Lu	tlu@mcimail.com
Dan Magorian	magorian@ni.umd.edu

Allison Mankin	mankin@cmf.nrl.navy.mil
David Marlow	dmarlow@relay.nswc.navy.mil
Matt Mathis	mathis@psc.edu
Jun Matsukata	jm@eng.isas.ac.jp
Michael McLay	mclay@eeel.nist.gov
Stephen Miller	smiller@bbn.com
Pushpendra Mohta	pushp@cerf.net
Dennis Morris	morris@altair.disa.mil
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@eng.umd.edu
Michael Patton	map@bbn.com
Alex Reijnierse	a.a.l.reijnierse@research.ptt.nl
Allen Rochkind	Allen_Rochkind@3com.com
Robert Roden	roden@roden.enet.dec.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Shawn Routhier	sar@epilogue.com
Hal Sandick	sandick@vnet.ibm.com
Erik Sherk	sherk@sura.net
Ming Sheu	msheu@vnet.ibm.com
Uttam Shikarpur	uttam@zk3.dec.com
Frank Solensky	solensky@ftp.com
John Stewart	jstewart@cnri.reston.va.us
Bernhard Stockman	boss@ebone.net
Vladimir Sukonnik	sukonnik@process.com
Steve Suzuki	steve@fet.com
Larry Tepper	ltepper@compatible.com
Matsuaki Terada	tera@sdl.hitachi.co.jp
Marten Terpstra	marten@ripe.net
Richard Thomas	rjthomas@bnr.ca
Claudio Topolcic	topolcic@cnri.reston.va.us
Raymond Vega	rvega@cicese.mx
Thomas Walsh	tomw@kalpana.com
William Warner	warner@ohio.gov
Guy Wells	guyw@uswest.com
Chris Wheeler	cwheeler@cac.washington.edu
Cathy Wittbrodt	cjw@barrnet.net
Jane Wojcik	jwojcik@bbn.com
Jessica Yu	jyy@merit.edu

A Look at Some Data

and

a Few Early Observations

(An Update to PG's INET '93 Talk)

Phill Gross
Dennis Ferguson

August 16, 1993

1

ACKNOWLEDGEMENTS
AND
SOURCES OF DATA

- GSI NIC - Scott Williamson and Mark
- Merit - Dale Johnson, JoAnn Ward, Steve Widmayer, Enke Chen, Susan Horvath
- SRI-Zone - Mark Lottor
- RIPE - Marten Terpstra
- Scott Bradner
- The IAB, especially Yakov Rehkter and Elise Gerich

2

LOOKING MORE CLOSELY AT THE DATA

- We need a better understanding of how the IP address space is being utilized, eg,
  - What is percent of total address space assigned?
  - How much is "connected" to the Internet (ie, IP reachable)?
  - How efficiently is it being utilized?
  - What is the average "host density" per assigned network?
  - What is the current growth rate?
  - How has growth rate changed over time?
  - How can assignment policies affect the growth rate?
- We should be looking for how much room we have for growth, and be very careful about projecting timeframes.
- Projections are still an important tool, but we need to be mindful that projections are only accurate if there is no change in the current growth rates!

3

WHAT PARAMETERS CONTROL GROWTH?

- Its really hosts that we want to allow to grow. In fact, one could argue that it is host growth that is the free variable and that the number of networks grow to accommodate host growth. Therefore, let's look at some of the parameters that constitute host growth.
- Its simple to calculate the maximum number of hosts, H, that the IP number space could accommodate (from a weighted calculation of the total number of class A, B, and C networks). However, several parameters affect the actual number that we will see in the Internet.
- First, only the networks that have been assigned contribute to the total hosts we see in the Internet. So, the total host number must be adjusted by the percentage of assigned networks, a.
- Next, only the networks that are connected to the Internet, e.g., registered in the NSFnet policy routing database, contribute to the total hosts we see in the Internet. Therefore, the total host number must be adjusted by the percentage of assigned networks that are in the PRDB, c.
- Finally, we need to consider the efficiency of utilization of the network numbers. So far, we've only talked about network numbers. Now we must multiple by an average percentage host utilization, e, of the assigned, connected networks to reach the total number of hosts in the Internet.
- The result is the following expression

  #_Internet_hosts = H * a * c * e

- We will now take a look at the various data we have available, with goal of getting a measure for these parameters, especially their rate of growth and how to control them.

4

OVERVIEW OF THE GRAPHS

1. "Division of IP Address Space" -- Most of you already understand the information on this slide, but it's still always impressive to me to see how much space we have left in unused Class A's. (p.13)

2. "Growth of the NSFnet Policy Routing Database" -- Going back even to Arpanet and Milnet days, we've often looked at the routing tables to get a feel for Internet growth. But, that is not nearly the whole story, and we need to look at the data in other ways for a better understanding. (p.14)

3. "Categories of Address Utilization" -- IP addresses are assigned and used in various categories. The next slide gives a more detailed definition of those categories. The chart shows the growth in those categories. Even this is not fine-grained enough, though. We need to look at the growth of individual address classes, both by class and by classless percentages. (p.15)

The next three charts look at growth by network address class.

4. "Class B and Class C Assignments" -- This chart shows growth of Class B and C address assignments. Notice the knee in both curves in early 1992, in which C assignments increase and B assignments slow down. This corresponds roughly to when the assignment policies became stricter. (p.16)

5. "Class B and C Growth for Assigned and Configured" -- Extends the previous chart by adding the data for growth in Class B and C "Configured" nets. Notice that, in early 1992, roughly at the same knee as above, the number of C's in the NSFnet PRDB overtakes the B's for the first time. This chart is a good example of the need to include global data. If we really want to understand the growth in actual usage of the net numbers (e.g., in the PRDB or in routing tables), we need to include global data (e.g., RIPE). (p.17)

5

"CATEGORIES" OF ADDRESS UTILIZATION

- "NIC Block Allocations" -- The IP address space is administered by GSI NIC, under guidance from IANA. The NIC parcels out "Block Allocations" and delegates assignment responsibility to regional address authorities, which make both individual address assignments and further delegations of block allocations. (Source of data on "Block Allocations": GSI NIC)

- "Individual Address Assignments" -- Regional address authorities (e.g., RIPE, network service providers, and GSI NIC itself) parcel out individual network number assignments to end users from their "Block Allocations". All delegated address authorities must agree to report all assignments back to the NIC. The quality of the NIC data depends on timely and reliable reporting. (Source of aggregated data on individual assignments: GSI NIC)

- "Configured" in NSFnet Policy Routing Database -- In the U.S., Merit administers routing for the NSFnet. In order to enter network numbers in the NSFnet routing tables, users register with Merit. These network numbers are first entered into the NSFnet Policy Routing database (PRDB). Merit refers to these as the "configured" networks. There is usually a time lag before the network numbers show up in the NSFnet routing tables. For this and the next category, we need to understand how other global data can be included. (Source of "Configured" PRDB data: Merit)

- "Announced" in NSFnet Routing Tables -- At some point, the "configured" network numbers will start showing up in the NSFnet routing tables (usually learned via routing protocols). Merit now tracks the size and contents of the routing tables on a routine daily basis. There is a fair amount of day-to-day fluctuation in the actual routing tables. Merit has been tracking the "announced" networks for about the last year. This is an example of new tracking due to new demands. (Source: Merit)

6

OVERVIEW OF THE GRAPHS (CON'T)

6. "Growth in NSFnet 'Configured' and 'Announced' Networks" -- Completing the tour of looking at the different categories of address utilization by address class, this chart shows the NSFnet "configured" and "announced" networks. Looking horizontally along the 6, 8, and 10,000 host gridlines, you'll notice that there is an approximate 4.5 month lag between the time that the PRDB reaches a certain size and the time that the actual routing table reaches the same size. Notice also that the ratio of "configured" to "announced" networks has remained roughly around 1.4. (p.18)

The next three charts look at growth classlessly.

7. "Percentage of IP Address Space Assigned" -- Dennis Ferguson first suggested looking at the assignments as a percent utilization of the address space. I read his email to Big-Internet on the weekend before INET, and I added a version of this slide. This chart shows the total percent assignment of the address space, plus the contribution of each address class. Notice the very large percentage taken by A's! Also, notice the very small percentage contributed by C's, even though C's now account for a very large number of the networks assigned. This demonstrates to me how much better we could utilize the address space if we used it classlessly. It also tells me we should try to reclaim some of those A's! (p.19)

8. "Percentage of IP Address Space Configured in the NSFnet PRDB" -- This chart uses the same approach as the last chart, but it looks at the percentage of IP address space now in the NSFnet PRDB. (p.20)

9. "Comparison of IP Address Space Percentage Assigned and Configured in PRDB" -- This chart simply compares the results from the previous two charts. (p.21)

10. "Host Growth (Measured by SRI-Zone)" -- Mark Lottor, bless him, has been measuring host growth for years, and now we find how valuable his work is. Looking at the data for the last 12 months show exponential growth with a doubling time of roughly every 14 months. (p.22)

7

OVERVIEW OF THE GRAPHS (CON'T)

11. "Percent Host Utilization of Addresses in the NSFnet Routing Tables" -- This chart attempts to show the host density of the current Internet networks. It uses the SRI-Zone host measurements and the Merit address class breakdown for the PRDB "configured" nets. Also, we observe that (because A's are soooo big) the percentage utilization of the B and C nets is probably closer to the way folks actually utilize their networks. This is the value we will use in our later analyses. This chart needs some more refinement - it should use the address class breakdown of the "announced" nets (i.e., the nets actually in the routing table at the time that the SRI-Zone measurements were made), it should use the global "announced" nets, not just NSFnet "announced" nets (since SRI-Zone is a global measurement, and many sites (e.g., with security gateways) do not allow SRI-Zone to do a host count. We attempt to account for the difference between the "configured" and "announced" nets by adjusting the PRDB data by 1.4. In later versions, we will try to account for the other inaccuracies. However, it is still pretty surprising to the estimated host utilization so low -- 1% for B's and C's, and 37% overall!! (p.23)

12. "Hosts per Domain" -- The SRI-Zone program counts both hosts and domains. This chart shows a straight forward vision of the host count by domain count to show the number of hosts per domain for each SRI-Zone measurement. Interestingly enough, it shows the measurements to be fairly constant (roughly between 30-40 hosts per domain) for the 3.5 year period from July 1988 to late 1991, before it begins to show any growth. By mid-1992, the measurement jumps up to over 60 hosts per domain. Perhaps this increase is related to the change in address assignment policy at around this time. If so, it may be evidence that we can positively affect the host density with changes in assignment policy. Using the various data in the spreadsheet, we could also try looking at, for example, 'domains per autonomous system number' or 'sub-domain per domain' as other measures of the density of address utilization. (p.24)

8

OVERVIEW OF THE GRAPHS (CONCLUDED)

13. "Projection of Class B and C Assignments" -- Finally, the punchline -- at current assignment rates, how much longer do we have before we exhaust the address space? As Jeanne Dixon will gladly tell you, predictions are fraught with danger, because there are many issues that can affect future growth rates (such as NT shipping with TCP/IP). However, let's observe that these type of projections should "not" be used for predicting the future anyway! It's only meant to show what "could" happen if we keep growing at the current rate, with no other changes. So, these types of projections should be used to scare the pejebbers out of us, so that we work real hard to make whatever changes are necessary to make sure the prediction does not come true. In other words, the value of projections is the opposite of a "self-fulfilling prophecy". If the projection gives us bad news, we hope to make it a "self-defeating prophecy". (p.25)

This chart takes the last 12 months of assignment data (roughly, after the aforementioned policy changes), and uses a straight exponential projection. Notice that the chart shows real data prior to Jan 1993, and projected data after Jan 1993. It shows both B's and C's running out in early 1997. That's not enough time! So we have a lot of work to do turn this into a "self-defeating prophecy". This chart has not yet been updated since INET '93. A preliminary quick look, using 2 more months of data, seems to indicate the exhaustion date pushed back to around mid-1997. We have a lot more work to do on these types of projections. We also plan to look at the change of the exhaustion date over time (by making a series of projections at, say, every 3 months). This would show us if the change is smooth or jumping all over the map.

Plus, we need to look at the projections classlessly. The next slide makes a start in that direction. (p.26)

9

HOW CAN WE CONTROL THE GROWTH PARAMETERS?

- We earlier derived the expression:

  #_Internet_hosts = H * a * c * e

Where

H = maximum number of hosts that the IP number space can accommodate.
a = the percentage of assigned networks.
c = the percentage of assigned networks that are connected to the Internet, i.e., configured in the NSFnet PRDB.
e = average percentage host utilization of the assigned, connected nets.

- Notice that we have a measure of the host growth, provided by SRI-Zone, and that the doubling rate was roughly 14 months.
- We used this host information to estimate (using the above expression) the host utilization at roughly 1% (for B's and C's).
- We also saw that the percentage of assigned networks was currently around 37% and the percentage of assigned networks that are connected to the Internet (i.e., configured in the NSFnet PRDB) was roughly 18%.
- Now, let's work backwards to see how much room we have for growth in the address space.
- If we assume that 16% is a reasonable target for the host utilization, "e", then we have room for 4 doublings from its current value of 1%.
- If we assume that we want to ultimately set a policy that all assigned addresses are connected to the Internet, then the target for "c" is 100%. This allows 2 doublings from its current value of 18%.
- Finally, assuming that we achieve 100% assignment, then "a" has room for ~1.5 doublings from its current value of 37%.
- Taken together, we may have room for roughly 9 doublings of the host growth before we hit the wall. If the current host growth rate of 14 months remains constant, that would give us about 10 years of growth -- with very careful management!
- We now need to look at how we can control these parameters, e.g., with new policies and better tracking.

10

ON BETTER UTILIZATION OF THE ADDRESS SPACE

It seems clear to me that 1) there is still sizable, untapped address space available, if we can recover/reclaim it and then use it effectively, and 2) assignment policies can/should be used to encourage better utilization of the address space (e.g., to increase the host density). Keeping in mind the previous analysis, we should try to develop policies and procedures that encourage the most effective use of the address space. For example:

- We need to move to classless addressing as soon as possible, so that we can use the untapped address space most effectively.
- We need to figure out how to utilize the as-yet untapped 63 Class A's. That's almost 1/4 of the total address space!
- We should make a stronger effort to reclaim unused addresses (especially, those 27 Class A's that are assigned but are not configured in the NSFnet Policy Routing Database. That's almost 1/8 of the total address space!). Most of the assigned addresses are "not" in the Internet routing tables. So, let's think the unthinkable - reclaim and renumber some current assignments! (Especially the A's)
- We need to encourage much denser host utilization of the address space. For example:
  - We need variable length subnet masks (including IGP's to support it and a manual on how to use them).
  - We need even tougher address assignment policies to encourage denser utilization. (e.g., Even tighter on Class B's, and "no" more Class A assignments. Fourteen A's have been given out since March 1991!).

We still need to look more closely at the data, but from what we've seen so far we feel that, with some hard work, we can carefully manage our growth (with reclamation, CIDR, classless addressing, and tighter policies on assignment) to buy back enough time to deploy IPng before we run out of IPv4 addresses.

11

SUMMARY

- New Policies
  - Even tougher assignment policies to encourage denser host utilization.
  - Future assignments should have a "host utilization" target.
  - Assign numbers only to those nets that will be connected to the Internet.
  - Reclaim unused network addresses (especially A's).
  - Consider reclaiming some currently assigned net numbers (eg, some Class A's, and non-connected networks).
- New Technology
  - Deploy classless addressing ASAP!
  - Make renumbering easier!
  - Need variable length subnet masks and "how-to" manual on how to use them.
- Education and Social Engineering
  - Conservation and resource management!
  - Renumbering may be a hassle, but .....
  - More efficient host utilization
  - Non-connected nets don't get IP numbers

12

DIVISION OF IP ADDRESS SPACE

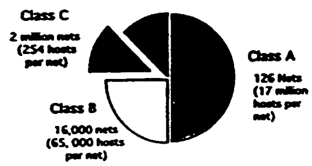

13

GROWTH OF THE NSFNET POLICY ROUTING DATABASE

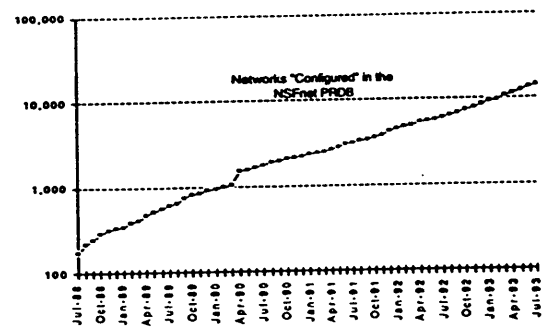

14

"CATEGORIES" OF ADDRESS UTILIZATION

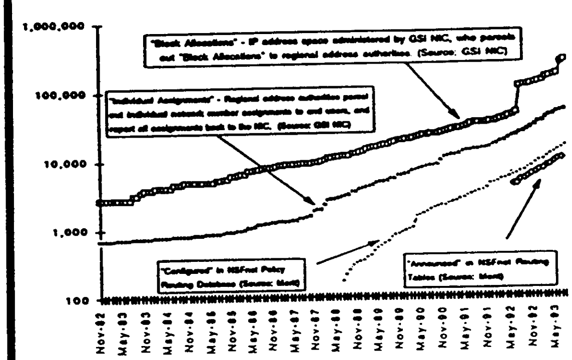

15

CLASS B AND CLASS C ASSIGNMENTS

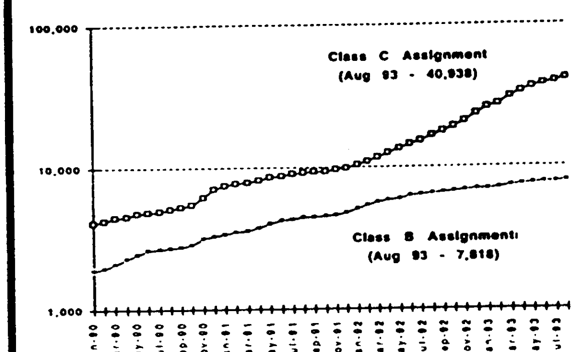

16

CLASS B AND C GROWTH FOR ASSIGNED AND "CONFIGURED"

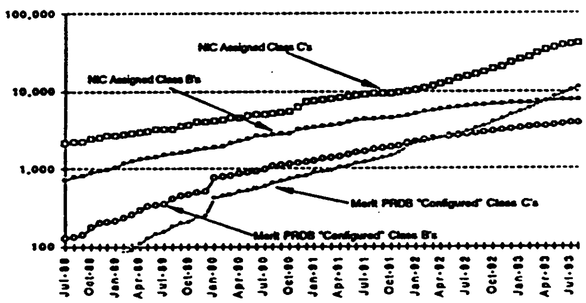

17

GROWTH IN NSFNET "CONFIGURED" AND "ANNOUNCED" NETWORKS

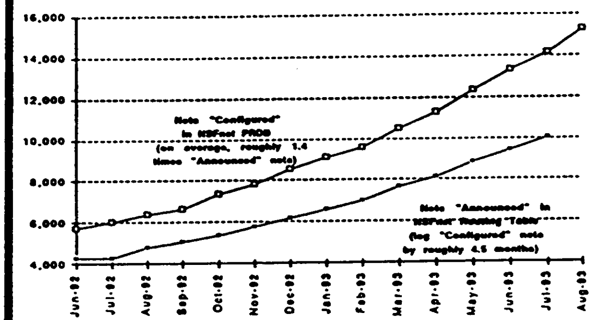

18

PERCENTAGE OF IP ADDRESS SPACE ASSIGNED

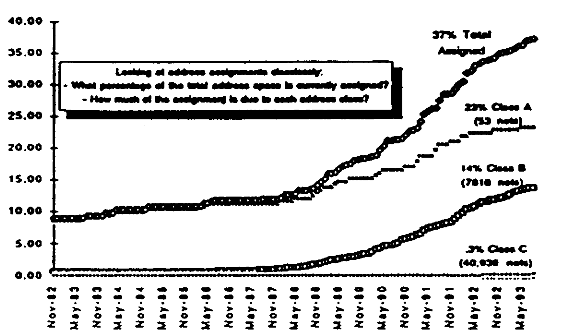

19

PERCENTAGE OF IP ADDRESS SPACE "CONFIGURED" IN NSFNET PRDB

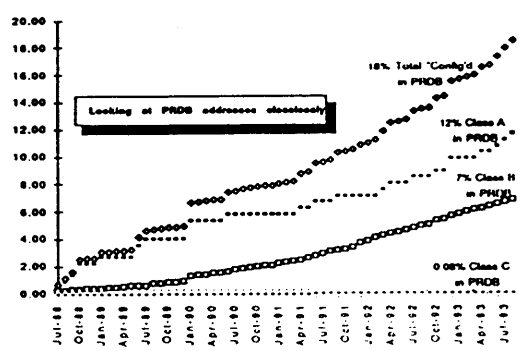

20

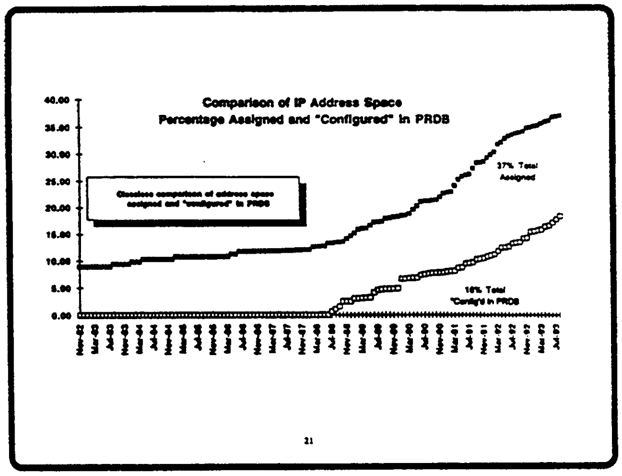
Comparison of IP Address Space
Percentage Assigned and "Configured" in PRDB
37% Total
Assigned
18% Total
"Config'd in PRDB
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.00
0.00
21

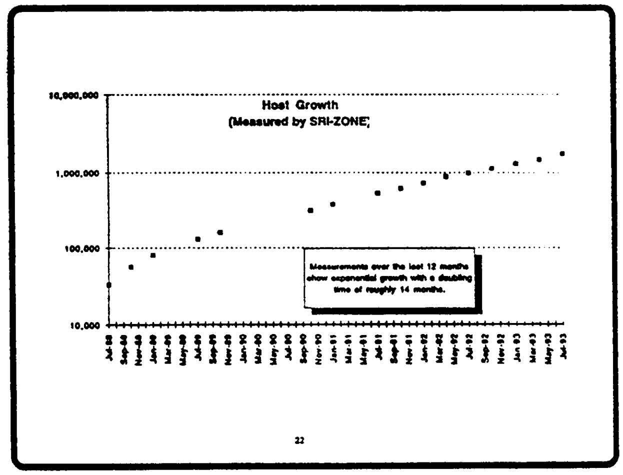
Host Growth
(Measured by SRI-ZONE)
Measurements over the last 12 months show exponential growth with a doubling time of roughly 14 months.
10,000,000
1,000,000
100,000
10,000
22

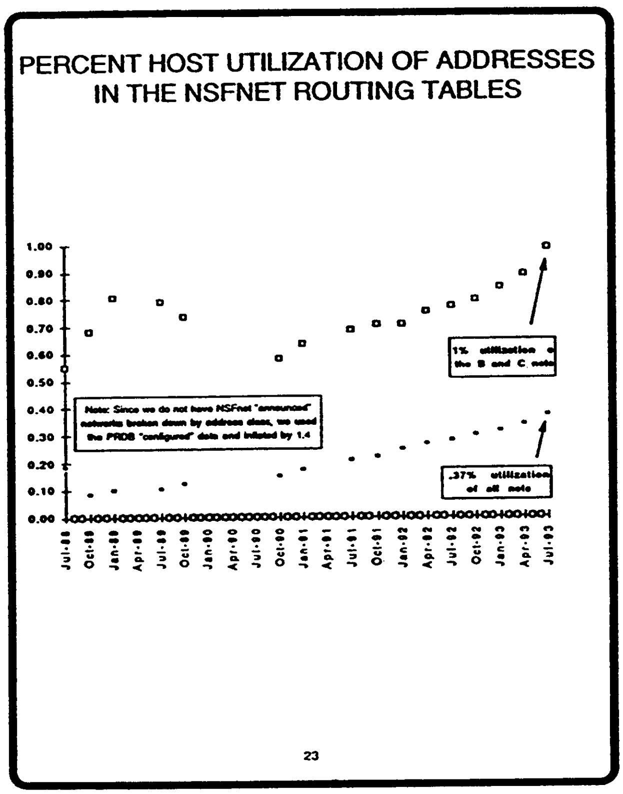
PERCENT HOST UTILIZATION OF ADDRESSES IN THE NSFNET ROUTING TABLES
Note: Since we do not have NSFnet "announced" networks broken down by address class, we used the PRDB "configured" data and inflated by 1.4
1% utilization of the B and C nets
.37% utilization of all nets
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
Jul-88
Oct-88
Jan-89
Apr-89
Jul-89
Oct-89
Jan-90
Apr-90
Jul-90
Oct-90
Jan-91
Apr-91
Jul-91
Oct-91
Jan-92
Apr-92
Jul-92
Oct-92
Jan-93
Apr-93
Jul-93
23

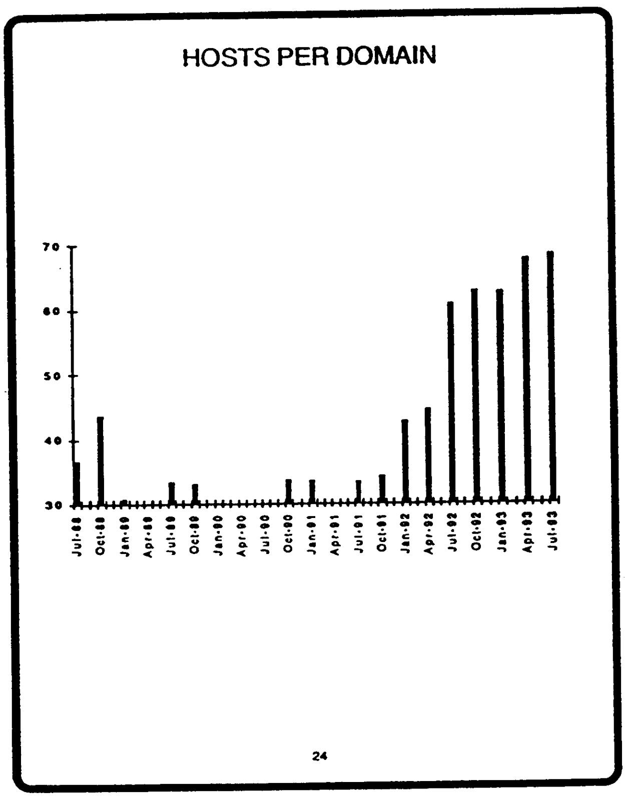
HOSTS PER DOMAIN
70
60
50
40
30
Jul-88
Oct-88
Jan-89
Apr-89
Jul-89
Oct-89
Jan-90
Apr-90
Jul-90
Oct-90
Jan-91
Apr-91
Jul-91
Oct-91
Jan-92
Apr-92
Jul-92
Oct-92
Jan-93
Apr-93
Jul-93
24

PROJECTON OF CLASS B AND C ASSIGNMENTS

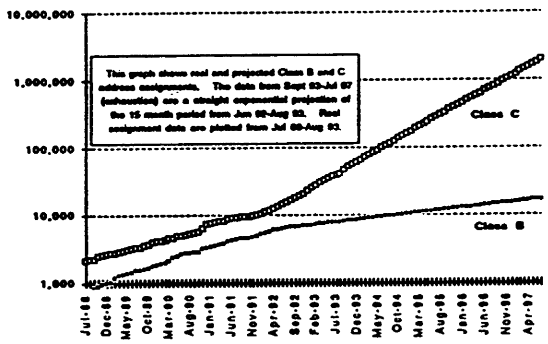

25

PROJECTION OF PERCENTAGE ADDRESS ASSIGNMENTS

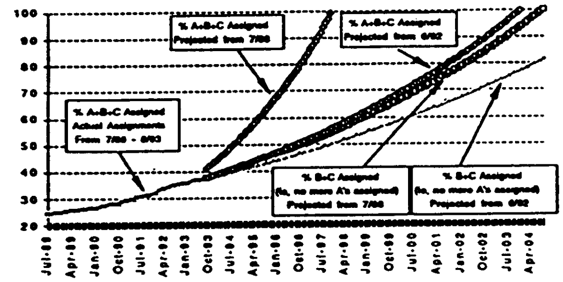

26

2.3.1 P. Internet Protocol (PIP)

Charter

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 IPng 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-4, OSPF, ARP, DNS, etc.); specifying PIP 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 IPng, 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-4 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.

2.3.2 Simple Internet Protocol (SIP)

Charter

Chair(s)

Steve Deering: `deering@parc.xerox.com`
Robert Hinden: `hinden@eng.sun.com`

Mailing Lists

General Discussion: `sip@caldera.usc.edu`
To Subscribe: `sip-request@caldera.usc.edu`
Archive:

Description of Working Group

SIP is a candidate for IPng. 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-bit 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 IPng 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

"SIP-RIP", 06/29/1993, G. Malkin, C. Huitema <draft-ietf-sip-rip-01.txt>

"SIPP Program Interfaces for BSD Systems", 12/21/1993, R. Gilligan <draft-ietf-sip-bsd-api-01.txt>

"Administrative Allocation of the 64-bit Number Space", 04/19/1993, W. Simpson <draft-ietf-sip-64bit-plan-00.txt>

"SIPP Neighbor Discovery", 12/06/1993, W. Simpson <draft-ietf-sip-discovery-03.txt>

"SIP addresses in the domain name service Specifications", 06/11/1993, C. Huitema <draft-ietf-sip-dnss-00.txt>

"Simple Internet Protocol Plus (SIPP): Overview of Routing and Addressing Extensions to SIP", 10/06/1993, S. Deering, P. Francis, R. Govindan <draft-ietf-sip-overview-00.txt>

"Extensions to DNS to support SIPP", 10/28/1993, C. Huitema, S. Thomson <draft-ietf-sip-sippdns-00.txt>

"IPAE: The SIPP Interoperability and Transition Mechanism", 11/16/1993, R. Gilligan, E. Nordmark, R. Hinden <draft-ietf-sip-ipae-transition-00.txt>

CURRENT MEETING REPORT

Reported by Robert Hinden/Sun Microsystems

Minutes of the Joint Sessions of the SIP and PIP Working Groups

These minutes are based on the notes taken by Christian Huitema and Bob Hinden.

The Simple Internet Protocol Working Group (SIP) and the P. Internet Protocol Working Group held two joint sessions. The first session was on Monday, November 1. The second session was held on November 4. Both sessions were carried on the Internet Multicast.

The agenda distributed prior to the meeting was reviewed and updated for the meeting.

SIPP Merger Overview (Steve Deering)

The purpose of the merger is to keep the simplicity and transition features of SIP and the advanced routing capabilities of Pip—while making them easier to use and to understand. The mailing lists have been merged, and Bob Hinden is writing a charter for the merged group.

This has resulted in some changes in the specifications, and in some terminologies. The changed terms are:

SIP → SIPP
system → node
anyone address → cluster address
Source route header → Routing header

The new terminology:

- The uniqueness scope of an address; for example the uniqueness scope of the loopback address is just one single node.
- The routing scope of an address, which is generally global to the Internet, but can sometimes be restricted e.g., for a "local use address."

Routing scope is always less than uniqueness scope, but not necessarily equal to it.

SIPP Overview and Issues (Steve Deering)

The address semantics have changed. Addresses identify nodes or set of nodes, not interfaces. A node may have several addresses, which may, in some instances, be tied to an interface.

The packet format has not changed, except for the "reserved" field which is now called "flow label." The 64-bit addresses are still composed of an IP address and a 32-bit prefix. The 64-bit SIPP address space is 10 million times larger than the global telephone number space.

The address formats are:

```
- classic: prefix, customer ID, node ID.
    -----------------------------------------
    |c| provider ID | customer ID | node ID |
    -----------------------------------------

- cluster:
    -----------------------------------------
    |c| provider ID | 0...................0 |
    -----------------------------------------

- local use address:
    -----------------------------------------
    | 0..0 | subnet |   IEEE 802 address    |
    -----------------------------------------

- multicast address:
    -----------------------------------------
    | 1..1 | flags + | group ID             |
    |      | scope   |                      |
    -----------------------------------------
```

The addresses are "provider oriented." The current SIP addressing drafts are obsolete. New SIPP versions will be submitted.

Options are encoded as a sequence of headers. SIPP options currently defined are fragmentation, routing and hop-by-hop options. Options for end-to-end security and flow set-up are under development. Options are not limited to 40 bytes like IP.

The format of the routing header is:

```
---------------------------------------------
| Payload | Number of | Next     | Reserved |
|         | Addresses | Address |          |
---------------------------------------------
|                Reserved                   |
|                                           |
---------------------------------------------
|                                           |
+               Address[0]                  +
|                                           |
---------------------------------------------
|                                           |
+               Address[1]                  +
|                                           |
---------------------------------------------
|                                           |
.                                           .
.                  ....                     .
.                                           .
|                                           |
---------------------------------------------
|                                           |
+               Address[n]                  +
|                                           |
---------------------------------------------
```

The minimum packet length has not changed. The routing header uses 64-bit chunks rather than the 16-bit chunks of Pip. Paul Francis mentioned that the advantage of this approach was "simplicity of handling." The addresses have their own routing scope, which relieves the need for the "routing contexts" which were present in Pip.

Noel Chiappa observes that the routing header is more traditional source routing rather than Pip "flows." Paul Francis said this was incorrect and that the Pip routing was not intended as flows.

The 28 bits of the flow-label will be structured according to one of two possible formats:

```
----------------------------------------
| DP |            0             | TOS|
----------------------------------------

----------------------------------------
| DP |          flow-ID                |
----------------------------------------
```

- 4 bits of "drop priority"
- 4-bit TOS is traditional IP type of service
- 24-bit flow-ID is a pseudo random number chosen by source to identify special flow state along path

The reason that the flow-ID is random, based on an idea from Dave Clark, is that it makes it easy to use a subset of it (bit slice) as a "hash code" for access to a flow table within the routers. To a question on TOS, it is observed that this really is a heritage from IPv4, although current experience in IPv4 networks is rather bleak. There was considerable discussion leading to the suggestion to drop IPv4 TOS.

SIPP Routing and Addressing (Paul Francis, Ramesh Govindan)

Paul Francis presented the use of the routing header. All packets are identified with 64-bit addresses which are unique with the scope, but may need additional 64-bit addresses to complement an insufficient routing scope. There is also a need for mobile hosts, or when special policies are required.

The SIPP addresses are contiguous bit-wise maskable (similar to IP with CIDR). This poses conditions for extended addresses:

- Single hierarchy element cannot straddle 64-bit boundary.
- Top and bottom 64 bits have to be both globally unique; one could perhaps release this requirement for "middle" addresses.

Currently SIPP assumes hierarchical provider addresses.

The cluster address is similar to an "anycast address," i.e., it addresses any of the routers sharing a prefix. If a packet arrives from "outside," it is accepted by the first node that matches the prefix; if from within the node, it is accepted by the first router that operates at "that prefix length." In the current state of the art, they will have to be "hand configured."

Examples of such addresses are:

- Provider.0: accepted by first router in provider network, used for provider selection.

- Provider.subnet.0: can be used for mobility support.

The local use addresses provide an 8-bit fixed prefix and an 8-bit "subnet number" in complement to the 48-bit IEEE-802 address. The local use addresses can be used over a multi-subnet site. It could be used exclusively for a site not connected to the Internet.

The address sequence has to be "manipulated" by the hosts. This is really what the merger with Pip is all about. Note that the SIPP header format did not change in the merger.

Hosts should be able to:

- Represent their own address as a sequence, not just a single 64-bit address.
- Reverse an address sequence.

If hosts do this from the start, new semantics can be added to the Internet, for example extended addresses, without having to update any internet hosts.

The group mentions that there should be a minimum size specification, e.g., "at least three components." This applies to local configuration, nodes should be able to process arbitrarily long routing headers. Similarly, a limit is needed for DNS configuration (size of record) and for "reverse look up" in the DNS (depth of the tree). Also, the "error behavior" should be specified – what should be done if the host receives a packet that it cannot understand.

Paul then presented the literal notation for the source route mechanism: <X, Y, Z>. Two kinds of address sequence have been defined: source capable and not source capable. For example, a multicast address is not "source capable": it cannot be used as a source address in a packet.

Suppose a sequence <S0, S1, .. , Si, Dj, Dj-1, .. D0>, i.e., the source chain then the destination chain. In most cases the chain will have exactly two elements <S, D>. This was only true in Pip for local communications. Paul presented the mechanism for building and reversing source routes, and mentioned the open issue: whether routes should be stored in the internet program, in the transport or in the application.

Ramesh Govindan presented different examples of sophisticated routing using the SIPP routing header. This included:

- Basic routing involving only the DNS. Sequence has two elements, reversal is trivial.

- Selection of the first hop provider. Sequence has three elements; change of provider within the association life time is easy.

- Item with "extended addresses," with four elements in the sequence.
- Examples are also given for multicast, including source routing prior to multicasting.
- Multicast is also possible with extended addresses: this allows recipients to reply to the source address.
- Mobility examples are also given: the address sequence includes the identification of the "base station." Note that the "mobile cluster" scenario is not presented! Address extension can also be used for auto-configuration:
  1. Hosts creates a "local wire" address.
  2. Host will receive a local cluster address, e.g., by receiving advertisement. It can combine his hardware address with this prefix, to form either a 64-bit address if the prefix is short enough, or an extended address otherwise.

Several members of the group question the "automatic reversal" of source routes in the case of "provider selection." There are clearly several degrees of liberty at this stage.

IPAE Specification Overview and Issues (Bob Gilligan/Erik Nordmark)

A new specification has been written by Bob Gilligan, Erik Nordmark and Bob Hinden. This is based on the original specification by Dave Crocker, and one year of work and discussion. The components of the specification are:

- Encapsulation within IPv4 for "tunnelling."
- 64-bit SIPP addressing scheme is compatible with IPv4 plan:

```
  6  6            3 3            0
  3  2            2 1            0
+---+-------------+--------------+
| c | Site Prefix | IPv4 address |
+---+-------------+--------------+
```

  The "c" bit explains whether the host is SIPP capable or not.
- Host algorithms for direct interoperability with IPv4 hosts.
- Translation agent between SIPP and IPv4.

Bill Simpson questioned the change of vocabulary from "commonwealth" to "site"—as commonwealth implied a larger kind of object. Steve Deering believes no name for these objects is really needed. Christian Huitema noted the need for a conventional 32-bit prefix, removing the need for "mapping tables" as long as hosts are capable of IP routing. John Curran mentioned the relation between site table and provider IDs: if one changes provider, then one changes prefixes, thus one has to change the "mapping table." The upper 32 bits carry an assumption about provider connectivity. The picture has changed a lot since the advent of CIDR; if CIDR really solves the routing table explosion, then the "mapping table" is not necessary. As Steve Deering mentions, the group really hates the mapping tables.

Jim Bound mentioned the complexity of transition for a host, and suggested that the group emphasize the inherent simplicity of the 64-bit approach.

A list of remaining IPAE issues came out when revising the specification.

Editor's Note: A detailed list of remaining IPAE issues is available via FTP or mail server from the remote directories as /ietf/sip/sip-pip-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Erik Nordmark presented the problems of keeping state when "tunnelling" is used:

- ICMP packet too big: Need to memorize the tunnel MTU, for either immediate transcription.
- ICMP TTL exceeded: Need to memorize the tunnels TTL,
- ICMP "unreachable": Signals an incorrect tunnel.

These "states" should really be "soft state," i.e., updated cautiously. The SIPP design helps the error handling, as the initial hop limit was present in the first bytes of the packet. This helps computing the "exact length" of the tunnel.

The state can be discarded for garbage collection (reduce the memory requirement) and also for detecting improvements – for example if a remote router suddenly becomes reachable. The MTU increases will regularly be probed by the source, so the absence of remote ICMP may be an indication of the absence of problem.

Neighbor Discovery/ARP (Bill Simpson)

The protocol has been renamed "neighbor discovery" after the merging. It has two packets: "where are you" stating the address looked for, and "I am here," with variable parameters. All packets include a "media type" and "MAC address" parameter, so that one does not need ARP.

Bill questioned the need for further usage of the "change prefix" parameter, which is used to broadcast "changes of providers." This is now well done, with prefix length, old address and new address.

Another questioned feature was the passing of information about other routers and other subnets—use discovery as a router protocol, or at least as a replacement for "OSPF hellos." The particular format of this "routing information" is hotly debated; in particular it is suggested to separate information on the router address and information on the "connected subnets." For each subnet, there are "preferences" and "priority," as well as a "zone" used for local addresses, and "MRU" indicating the maximum packet length used by the routers. The utility of several fields, or even the very utility of this parameter, is debated:

- MRU is generally understood as "not needed."
- The parameters taken from OSPF and IS-IS should go away.
- "Zone" is an inappropriate name for "local scope subnets," which should just be passed as particular subnets.

The "system heard" parameter is essential for support of eliminating the "hidden transmitter" problem. For each system heard, this pass various parameters: quality of reception, advertisement number, etc. This seems too complex to many listeners.

Steve Deering requested the removal of the "service advertisements." Bill also presented "transit informations" and "redirects." Further discussion is clearly needed!

SIPP OSPF (Rob Coltun)

Rob proposes the acronym "OSPPF": bigger addresses, more protocols. For carrying big addresses, one needs to:

- Provide "link state ID" independent of address. Currently, an LSA is identified by [Router ID, LS-ID], where LS-ID represents the "network number." A 32-bit locally unique ID will be used in OSPPF.

- Advertisement will have to carry long address in addition to LS-ID.

The schema of the LSA is:

```
----------------------
| Advertising router |
----------------------
| Link State ID      |
----------------------
| Type               |
----------------------
|                    |
| Address            |
|                    |
----------------------
| Router ID          |
----------------------
| ...                |
----------------------
```

There is agreement that the "advertising router" should be a 64-bit field; in general, routers should be identified by their 64-bit identifying address. The LSA is identified by the combination of advertising router and LS-ID; the LS-ID has to be unique within the router, i.e. can be a random 32-bit number. It is not even necessary to keep the same number during different "instantiations," e.g., after a reboot, as the old segments will either be replaced or fade away. Indeed, this implies that the LS-ID cannot be overloaded.

For the big addresses, one has to carry a length field (in bytes) and the number of significant bits; thus it makes sense to also carry a "type" field, which enables for running other protocols in parallel:

```
-----------------------------
| Type | Len  | Mask size   |
-----------------------------
|         Address           |
-----------------------------
|                           |
-----------------------------
```

The "type" field is used to specify e.g., IP or SIPP, which means that OSPPF has dual protocol capability.

Rob then addressed the "hierarchical" problem. Two levels are probably enough (200 routers per area imply 40,000 routers). It is easy to do a multiple level version, e.g., to accommodate regionals which want to integrate their clients as OSPF areas, and also because inter domain routing requires a lot of work. There is however a general agreement that such developments should be discussed in the OSPF group, and that the SIPP version should really be a straight forward transcription of OSPF to 64-bit addresses.

SIPP Service Interfaces and DNS Changes (Sue Thomson)

Sue Thompson presented the changes to the DNS for storing address sequences and for supporting the transition. These are:

- A new "ASEQ" record, a sequence of 64-bit elements, which does not cause additional processing.

- A new "inverse look-up name," which was defined similarly to that of the initial SIP, and used a PTR query. There is however a consensus on a "per octet" break up that seems more rational given the "bit mask" nature of the address. This will be represented as a sequence of hex tokens, without leading zeros.

Jim Bound would like the DNS interfaces to strip the upper parts of the address sequence when they are not necessary. This will have to be specified in the routing architecture.

There are two transition issues:

1. Whether resolvers should return A records if no ASEQ address is present. According to Sue, resolvers will have to ask for both ASEQ and A.

2. Whether the additional section should only include A records, or also ASEQ records. Decision is that if the query is received from a SIPP host, then ASEQ should also be returned.

Sue is going to implement the specification in bind 4.9.

Auto Configuration and DHCP (Jim Bound)

The DHCP protocol is very straightforward. DHCP is sitting in the application layer, so has to traverse the entire stack; after a simple "connection" exchange, the client is returned a set of configuration information, e.g., an address. In some cases, databases have to be updated. Steve Deering mentioned that dynamic updates of the DNS are not really required; one might as well preallocate name and address types.

John Wroklavski mentioned that auto-configuration is the "single most important" design part of IPng; it should work in a large set of environments. Jim Bound mentioned that DHCP can really be used without problem, and that making a SIPP option is really straightforward.

Ohta asked whether SIPP/DHCP will have "relay agents." In fact, we don't need them as SIPP stations can very easily use a hardware address. Thus, the group will be able to use multicast to find the DHCP server, including with diameters larger than 1 (outside the local net): there is no need for relays, routers do the job easily. Paul Francis proposed to write a specific document explaining how network layer mechanisms can be used to help auto-configuration, but also for discovering DNS servers, gopher servers, etc. Jim insisted that we have to be concerned by automatic configuration of the DNS, i.e., register automatically IP address and DNS name bindings.

Jim Bound will prepare a "64-bit" version of DHCP.

Address Assignment Issues (Paul Francis)

Given the difficulties of managing geographic addresses, there is agreement that only "provider" addresses should be used in the short term. The immediate assignment is:

```
-----------------------------
|1| 31 bits   | 32 bits     |
|C| something | IP address  |
-----------------------------
```

Detail of IP address under CIDR:

```
-----------------------------------------------------
| Provider ID | Subscriber ID | subnet ID | Node ID |
-----------------------------------------------------
```

Without CIDR, the address is:

```
-----------------------------------------------------
|    network number           | subnet ID | Node ID |
-----------------------------------------------------
```

These addresses will be a "legacy" of the pre-CIDR era. Provider, subscriber, subnet and host is a good hierarchy; but eventually growth will force us beyond 32 bits. Thus, at least the provider ID should be pushed into the higher 31 bits of the SIPP address.

The proposal is to:

- Push provider part in upper 31 bits.
- Leave room below provider for subscriber and "subProvider" parts.
- Leave room above provider ID for contingencies.

This gives the following structure:

```
------------------------------------------------------------------
| 8 bits |    24 bits    |  m bits          | p bits     | 32-m-p  |
------------------------------------------------------------------
| C 0..0 | provider Id | subscriber Id | subnet ID | Node ID |
------------------------------------------------------------------
```

The provider ID will be assigned "from the left," which means that they are followed by a number of zeros, which allow for future growth of the "subscriber ID" part. There was a general consensus to proceed with this plan for address assignment.

Attendees

Kenneth Albanese	albanese@icp.net
Steve Alexander	stevea@lachman.com
James Allard	jallard@microsoft.com
Susie Armstrong	susie@mentat.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Anders Baardsgaard	anders@cc.uit.no
William Barns	barns@gateway.mitre.org
Stephen Batsell	batsell@itd.nrl.navy.mil
Nutan Behki	nebhki@newbridge.com
Steven Bellovin	smb@research.att.com
Tom Benkart	teb@acc.com
Ram Bhide	ram@nat.com
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Jim Bound	bound@zk3.dec.com
Scott Bradner	sob@harvard.edu
Monroe Bridges	monroe@cup.hp.com
Ronald Broersma	ron@nosc.mil
Al Broscius	broscius@bellcore.com
Ross Callon	rcallon@wellfleet.com
Peter Cameron	cameron@xylint.co.uk
George Chang	gkc@ctt.bellcore.com
David Conrad	davidc@iij.ad.jp
Matt Crawford	crawdad@fncent.fnal.gov

John Curran	jcurran@nic.near.net
Michael Davis	mike@dss.com
Chuck de Sostoa	chuckd@cup.hp.com
Stephen Deering	deering@parc.xerox.com
Avri Doria	avri@locus.com
Donald Eastlake	dee@skidrow.lkg.dec.com
Havard Eidnes	havard.eidnes@runit.sintef.no
Robert Enger	enger@seka.reston.ans.net
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Carlos Fernandez	carlos@plk.af.mil
Eric Fleischman	ericf@atc.boeing.com
Richard Fox	rfox@metricom.com
Paul Francis	Francis@thumper.bellcore.com
Atanu Ghosh	atanu@cs.ucl.ac.uk
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Ramesh Govindan	rxg@thumper.bellcore.com
Jari Hamalainen	jah@rctre.nokia.com
Herluf Hansen	hha@tbit.dk
Dimitry Haskin	dhaskin@wellfleet.com
Marc Hasson	marc@mentat.com
Robert Hinden	hinden@eng.sun.com
Kathy Huber	khuber@wellfleet.com
Christian Huitema	Christian.Huitema@sophia.inria.fr
Phil Irey	pirey@relay.nswc.navy.mil
Kevin Jackson	kjackson@concord.com
Ronald Jacoby	rj@sgi.com
Dale Johnson	dsj@merit.edu
Timo Jokiaho	timo.jokiaho@ntc.nokia.com
Rick Jones	raj@cup.hp.com
Akira Kato	kato@wide.ad.jp
Elizabeth Kaufman	kaufman@biomded.med.yale.edu
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Edwin King	eek@atc.boeing.com
Andrew Knutsen	andrewk@sco.com
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
Kanchei Loa	loa@sps.mot.com
Allison Mankin	mankin@cmf.nrl.navy.mil
Peram Marimuthu	peram@wg.com
Greg Minshall	minshall@wc.novell.com
Randy Miyazaki	randy@lantron.com
Andrew Myles	andrew@mpce.mg.edu.au
Rina Nathaniel	rina@rnd-gate.rad.co.il
Sath Nelakonda	sath@lachman.com

Erik Nordmark	nordmark@eng.sun.com
Masataka Ohta	mohta@cc.titech.ac.jp
Steve Parker	sparker@ossi.com
Ismat Pasha	ipasha@icm1.icp.net
Michael Patton	map@bbn.com
Charles Perkins	perk@watson.ibm.com
Eric Peterson	elpeterson@eng.xyplex.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Michal Rozenthal	michal@fibronics.co.il
Dallas Scott	scott@fluky.mitre.org
Isil Sebuktekin	isil@nevin.bellcore.com
Vincent Shekher	vin@sps.mot.com
Ming Sheu	msheu@vnet.ibm.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Frank Solensky	solensky@ftp.com
James Solomon	solomon@comm.mot.com
Tae Song	tae@novell.com
Michael Speer	michael.speer@sun.com
Vladimir Sukonnik	sukonnik@process.com
Steve Suzuki	steve@fet.com
Susan Thomson	set@bellcore.com
Akihiro Tominaga	tomy@sfc.wide.ad.jp
Thuan Tran	thuan@xylogics.com
Hoe Trinh	htrinh@vnet.ibm.com
Keisuke Uehara	kei@cs.uec.ac.jp
Chuck Warlick	chuck.warlick@pscni.nasa.gov
Chris Wheeler	cwheeler@cac.washington.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
Jane Wojcik	jwojcik@bbn.com
David Woodgate	David.Woodgate@its.csiro.au
Liang Wu	ltwu@bellcore.com
Jean Yao	yao@cup.hp.com
Shinichi Yoshida	yoshida@sumitomo.com
Jessica Yu	jyy@merit.edu
Weiping Zhao	zhao@nacsis.ac.jp

SIPP Overview

Steve Deering
Xerox Palo Alto Research Center
deering@parc.xerox.com

IETF 28
Houston, Texas
November 1, 1993

Goals of SIP / Pip Merger

- keep simplicity & transition features of SIP
- encorporate Pip's flexible & powerful forwarding mechanism, while making easier to use & understand
- provide platform for introduction of further enhancements to the Internet's common protocol layer
- combine resources from Pip and SIP efforts
- reduce the number of IPng contenders

Changes to SIP Spec

changed / new terminology:

SIP —> SIPP
system —> node
anyone addresses —> cluster addresses
Source Route Header —> Routing Header
uniqueness scope of an address
routing scope of an address

Changes to SIP Spec (cont.)

changes to address semantics:

- addresses identify nodes or sets of nodes, not interfaces
- nodes may have multiple addresses
- addresses said to be "bound" to those interfaces, **if any**, that are in routability scope of that address

The SIP Header

32 bits

Vers.	Flow Label		
Payload Length		Protocol	Hop Limit
Source Address			
Destination Address			

24 bytes long, 4 more than minimal IP

SIP Unicast Address Format

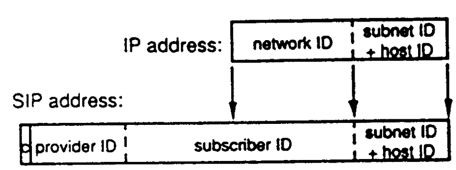

64-bit address space is 10 million times larger than the global phone numbering space

Other SIP Address Formats

cluster address:

local-use address:

multicast:

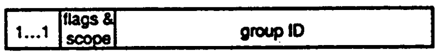

Internal Structure of the Flow Label (tentative)

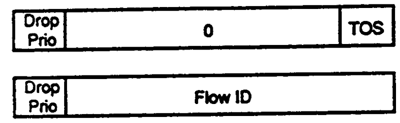

- 4-bit **Drop Priority** intended for layered media encodings
- 4-bit **TOS** is traditional IP Type of Service
- 24-bit **Flow ID** is pseudo-random number chosen by source to identify special flow state along path

Fast Mapping from Flow ID to State

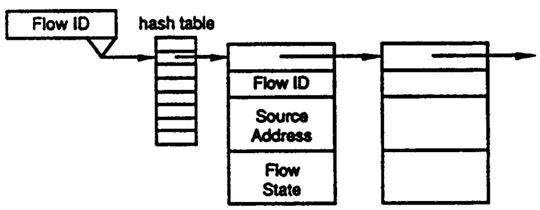

- idea due to Clark
- subfield of random Flow ID acts as pre-computed hash
- must compare full Flow ID & Source Address from packet with state record

SIP Options

- coded as optional headers:

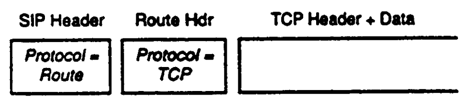

- currently defined:
  Fragment, Routing, Hop-by-Hop Options
- under development: End-to-End Security, Flow Setup
- not limited to 40 bytes like IP

Bellcore
Houston IETF
SIPP WG Session
Frame 2 of 36

SIPP Address

1. Uniquely identify node (or set of nodes)
2. Specify location of addressed node(s), to facilitate routing

Bellcore
Houston IETF
SIPP WG Session
Frame 3 of 36

SIPP Address

- Identify nodes, not interfaces
  - Though can be assigned on a per-interface basis
- Have a certain "routing scope"
  - Topological region over which nodes have sufficient routing info to route the packet
  - Most SIPP Address have global routing scope
  - Some are local, for instance, for bootstrapping.....

Bellcore
Houston IETF
SIPP WG Session
Frame 4 of 36

SIPP Addresses in SIPP Headers

- Every SIPP packet has two Identifying Addresses
  - Identifying source and destination of packet
  - Used for transport-layer pseudo-header checksum
- Identifying Addresses may or may not have sufficient location info
- If insufficient, additional addresses may be included in packet
  - Using optional SIPP Routing Header
  - May be viewed as high-order extensions of the Identifying Addresses
  - Complete vector of addresses called an Address Sequence

Bellcore
Houston IETF
SIPP WG Session
Frame 5 of 36

Examples of Address Sequences

- Mobile node
  - Attached somewhere other than where Identifying Address indicates
- Extended Address
  - 64 bits someday too small to assign globally-routable addresses to all nodes
  - In this case, lower 64 bits (Identifying Address) still uniquely identifies node globally

Bellcore
Houston IETF
SIPP WG Session
Frame 6 of 36

Unicast Addresses

- Contiguous bit-wise maskable
  - Similar to IP addresses under CIDR
- Use of Address Sequence places some constraints
  - Single "address hierarchy element" cannot straddle 64-bit boundary
  - Top and bottom (64-bit) addresses must be globally unique
    - Probably all address should be globally unique
  - Each address must be routable without examining previous addresses in Address Sequence

Bellcore
Houston IETF
SIPP WG Session
Frame 7 of 36

Unicast Address Assignment

```
|1|       n bits        |        m bits         |  p bits   | 63-n-m-p|
+-+---------------------+-----------------------+-----------+---------+
|C|     provider ID     |     subscriber ID     | subnet ID | node ID |
+-+---------------------+-----------------------+-----------+---------+
```

- C-bit has same use as with SIP
- Initial assignments are provider-based
  - Discussion of unicast address assignment later.......

Bellcore

SIPP Unicast Addresses: Cluster Addresses

- Unicast address of form <prefix><zero>
- Results in packet being routed to one of a group of nodes that has that prefix for its own address
  - Thus, a kind of anycast
- If packet sent from outside the group
  - Packet accepted by first node that has that prefix
- If packet sent from inside the group
  - Packet transmitted up the hierarchy until it reaches a router that is acting "at the hierarchy level" of the prefix

Bellcore

SIPP Unicast Addresses: Initial Cluster Addresses

- Provider.0
  - Accepted by first router in provider network
  - Used for provider selection, and more generally for provider-level source routes
- Provider.Subscriber.Subnet.0
  - Accepted by first router on subnet
  - Used for auto-configuration
  - Used for mobility where mobility-scope is the subnet
    - Provider.Subscriber.0 could perhaps be used for mobility where mobility-scope is the subscriber network

Bellcore

SIPP Unicast Addresses: Cluster Address Configuration

- Currently must be hand-configured
- Similar to hand configuration of the prefix that a router will originate in its router advertisements
- While cluster address is a general concept, use of cluster address should probably be designed on a per-function basis
  - Host use of cluster address (i.e., provider selection, mobility, boot-strapping, etc.)
  - Host discovery of appropriate cluster address
  - Router configuration of cluster address

Bellcore

SIPP Unicast Addresses: Local-Use Address

```
| 8 bits | 8 bits  |                 48 bits                     |
+--------+---------+---------------------------------------------+
|01111101|subnet ID|               link address                  |
+--------+---------+---------------------------------------------+
```

- SIPP node forms a SIPP address from its own link address
- Only guaranteed to be locally unique
- Used for local communications
  - Site not connected to Internet
  - Temporarily, for auto-configuration

Bellcore

Other Address Formats

- Multicast, Unspecified, Loopback, Multicast, All Nodes, All Hosts, All Routers
  - Same as current SIP spec

Bellcore

Address Sequence Handling by Nodes

- General rules for manipulating Address Sequences required to make them useful
- Nodes must be able to represent their own addresses as an Address Sequence
- Nodes must be able to "reverse" an address sequence
- If nodes can do this, new semantics can be added to the Internet (such as extended addresses) without necessarily having to update all nodes

Notation

- Literal notation of SIPP Routing Header mechanism looks like this:

```
          source address   dest address   Routing Header
          --------------   ------------   ------------
initial         S                A          >B  D
   next         S                B           A >D
  final         S                D           A  B >
```

- This is an awkward notation
- Equivalent notation is:

```
initial    S,*A, B, D
   next    S, A,*B, D
  final    S, A, B,*D
```


Node's Own Address Sequence

- Series of (64-bit) addresses: <Si, Si-1, Si-2, ..., S0>
- S0 = low-order address and Identifying Address
- Si = high-order address
- High-order address not necessarily hierarchically above low-order address
  - Rather, high-order address comes first in a vector of addresses
- For simplicity, a host can assume some small maximum number of addresses in sequence, for instance, 3 addresses

Node's Own Address Sequence: Two Kinds

- Two kinds of Address Sequences:
  - Source-capable
  - Not source-capable
- Source-capable Address Sequence can be used as a source address
- Unicast addresses are source-capable
- Multicast addresses are non-source-capable
- Both can be considered by a node to be its "own" address

Route Sequence

- Complete sequence of addresses in a SIPP header
- May contain many things:
  - Source address sequence
  - Destination address sequence
  - Policy route
  - Mobile-host base station, etc.
- But from a simple node's perspective, Route Sequence contains only source and destination address sequence:
  - <S0, S1, ..., Si, Dj, Dj-1, ..., D0>

Route Sequence

<S0, S1, ..., Si, Dj, Dj-1, ..., D0>

- For received packet, destination address is the node's own address, source address is everything else
- Source Address Sequence followed by destination Address Sequence
  - Source Address Sequence encoded low-order address first
  - Destination Address Sequence encoded high-order address first
- Common address sequence is <S0, D0>
  - No Routing Header

Route Sequence

- Information to maintain for "association" with correspondent node:
  - Source and destination Identifying Addresses for entire association
  - Source and destination Address Sequences currently in use
- Node that starts association:
  - Learns correspondent node's Address Sequences in normal way (DNS, user typing it in, etc.)
  - Adds one of its own source-capable Address Sequence
- Other node (the one that didn't start the association):
  - Extract its own address from the Route Sequence
  - Treats whatever is left as the correspondent node's Address Sequence

Thus, to Reverse a Route Sequence:

- strips off and stores its own address sequence from the tail of the route sequence of the received packet
- reverses the order of the remaining elements of the route sequence, and places them on the tail of the route sequence of the returned packet
- prepends a valid source-capable address sequence to the route sequence
- sets the active address to be the first address of the destination address sequence

Example Topology:

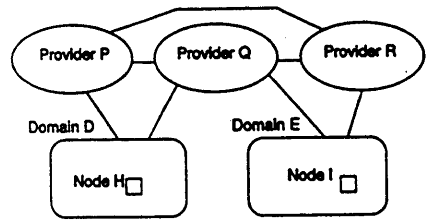

Simple (Non-Extended Addresses)

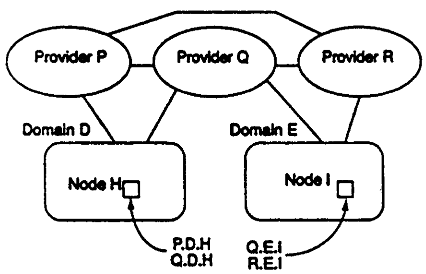

Simple (Non-Extended Addresses)

- H initiates association with I
  - Queries DNS, learns 2 addresses
- H chooses Q.E.I (best match with own addresses)

Route sequence at sender H: Q.D.H, *Q.E.I

Reversed route sequence at receiver I: Q.E.I, *Q.D.H

Simple Addresses with Provider Selection

- Previous example a kind of provider selection (but simple, no explicit instructions)
- Assume H wishes to use provider P:

Route sequence at sender H: P.D.H, *P.0, Q.E.I

Advanced route sequence at provider P router: P.D.H, P.0, *Q.E.I

Reversed route sequence at receiver I: Q.E.I, *P.0, P.D.H

- Assume I wishes to return packets via provider Q:

Alternative reversed route sequence: Q.E.I, *Q.0, P.0, P.D.H

Extended Addresses

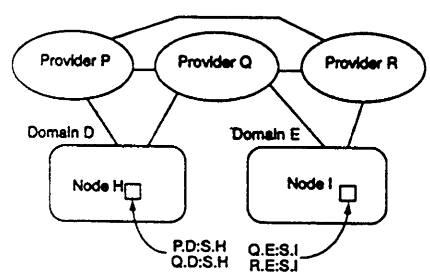

Extended Addresses

- Packet from H to I:

Route sequence at sender H: S.H, Q.D, *Q.E, S.I

Advanced route sequence at router in Domain E: S.H, Q.D, Q.E, *S.I

Reversed route sequence at receiver I: S.I, Q.E, *Q.D, S.H

Multicast

- Multicast address space defined as in SIP spec
  - Single 64-bit address
- None-the-less, general model is that host treats multicast address as an Address Sequence
  - Useful for unicast-then-multicast types of multicast
  - CBT
  - Remote-but-scoped multicast
  - etc.

Multicast Example: Extended Addresses

- Node H with addresses P.D:S.H and Q.D:S.H transmitting to multicast group with address M:

Route sequence at sender H: S.H, Q.D, *M

Reversed route sequence at receiver I: S.I, Q.E, *Q.D, S.H

- Multicast Forwarding Extension:

source address	dest address	routing header
S.H	M	>Q.D

- Router must peek into Routing Header for source-rooted multicast
  - If extended addresses used

Mobility Example With Extended Addresses

- H is mobile host, addresses P.D:S.H and Q.D:S.H
- I is its correspondent host, addresses Q.E:S.I and R.E:S.I

Route sequence from I to H: S.I, Q.E, *Q.D, S.H

- H moves to base station with address D.d

Route sequence from H to I after move: S.H, D.d, *Q.E, S.I

Reversed route sequence from I to H after move: S.I, Q.E, *D.d, S.H

Host Auto-configuration

- Host can construct a temporary address
- Use it to talk to an address server of some sort
- Then obtain a (more) permanent address
- Four scenarios
  - Router is or is not on the host's local link
  - Host can or cannot contact a configuration server

Host Auto-configuration

- Host first creates a Local-Use Identifying Address
  - Routing scope is just local wire
- Using this, the host discovers a router
- Router advertises its own address, plus its subnet cluster address
- Uses the router's address to create its own address
  - New address has same routing scope as router's address
- Uses this address to talk to address server (if necessary)

Host Auto-configuration: Forming Temporary Address

- If <zero> part of router's cluster address larger than host's link address
  - Host forms <clusterPrefix><linkAddress> for it's own address
  - Could be a Local-Use address or a global address
  - This address appropriate for permanent use if global or if Local-Use and contains unique ID
- Otherwise,
  - Host forms address sequence of C:L
  - C is router's subnet cluster address
  - L is host's Local-Use address
  - Can use this address permanently if L contains unique ID

Address Assignment

- Both geographical and provider-rooted assignments have certain benefits
- Perhaps ultimately having both types of assignment available is best
- At this time, geographical assignments seem problematic
  - Does any assignment authority have enough authority?
  - Is there enough organization among providers to achieve connectivity (physical or otherwise) within geographic areas?
- Provider-rooted assignments also problematic
  - Increased address administration within private networks
- At least we can move ahead with provider-rooted for now

A few thoughts on provider-rooted address assignment

- Immediate assignment of SIPP addresses:

```
|1|    31 bits    |     32 bits     |
+-+---------------+-----------------+
|C| something     |  IP Address     |
+-+---------------+-----------------+
```

- Detail of IP address under CIDR:

```
|    n bits       |    m bits     | p bits  |32-n-m-p|
+-----------------+---------------+---------+--------+
| provider ID     |subscriber ID  |subnet ID|node ID |
+-----------------+---------------+---------+--------+
```

Provider-rooted Assignment

- provider . subscriber . subnet . host is a good hierarchy, for now
- Eventually, growth will force us beyond 32 bits
- Thus, will want to push <u>at least</u> provider ID into higher 31 bits of SIPP address
- May or may not want to push subscriber part into higher 31 bits
- May or may not want to create subProvider layer of hierarchy
  - provider . subProvider . subscriber . subnet . host
  - It is up to individual providers
  - If created, may or may not want to push subProvider part into higher 31 bits

Provider-rooted Addresses

- Put provider part in higher 31 bits
- Leave room below provider for subscriber or subProvider parts
- Leave room above provider ID for unexpected contingencies

```
|1|            31 bits             |    m bits     | p bits   | 32-m-p |
+-+--------------------------------+---------------+----------+--------+
|C|          provider ID           |subscriber ID  |subnet ID |node ID |
+-+--------------------------------+---------------+----------+--------+
                                   |corresponds to current IP address|

| 8 bits |      24 bits            |
+--------+-------------------------+
|C0000000| provider ID, assigned   |
|        | "from the left"         |
+--------+-------------------------+
```

OSPPF

- Big Addresses
- Multiple Levels Of Hierarchy

Big Addresses

- Redefining LS ID
  - Currently LS ID Represents Network Number
  - LSA Identified By [Router ID, LS ID] Tuple
  - New Definition: 32-bit Locally Assigned LSA ID
    - Advertisements Will Have To Carry Around Big Address In Addition To LS ID
    - Need LSDB Lookup By LS ID and Address

- Router ID Remains 32-bits
  - Was Unique ID Or IP Address Can Now Be Unique ID Or Low 32-bits of Big Address
- Big Address
  - Mask Field Replaced By [type, len, sig_bits]
  - Address Is Padded To Word Boundary

Type	Len	Sig Bits
Address		

Hierarchy

- 2 Levels Is Probably Big Enough
  - 40,000 Routers - 200 Routers Per Area
- Why Multiple Levels?
  - Regionals Can Have Subscribers Run As Sub Areas
    - Sub Areas Can Be Further Divided
  - Inter-Domain Routing Will Require Alot Of Work

- One Solution For Support Of Real Time Traffic In Integrated Services Network Is To Peridocally Send Metrics So Traffic Can Adapt
- Multiple Levels Of Hierarchy Localizes Frequent Updates

Overview Of Hierarchy

- Expand On Existing OSPF Concepts
- All Levels Have Parent-Child Relationship
  - Parent = Backbone, Child = Non-Backbone Area
- Aggregation Only Occurs Between Adjacent Levels
- Child Levels Can Only Advertise Routes Through Parent

- Cross-Child Relationships Look Like External Connections Which Are Less Preferred
- LSA Flooding Has Scope
  - Information Hiding Controlled By Source
  - Force Aggregation As LSAs Are Flooded Up Levels

Topology

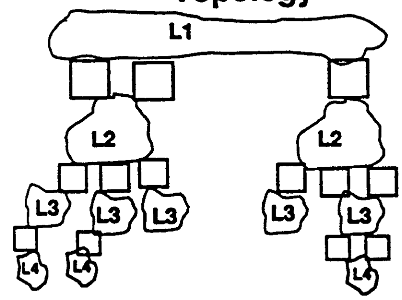

DNS and SIPP

Susan Thomson

IETF DNS Working Group
November, 1993

Modifications Necessary

- SIPP addresses
  - new address resource record (ASEQ)
  - new inverse lookup domain (SIP-ADDR.ARPA)
- IP/SIPP Transition
  - modify RRs that return addresses in additional section (NS, MX, MB)

Type ASEQ Resource Record

- contains address sequence in contiguous 64-bit fields
- regular lookup uses domain name, ASEQ query
- ASEQ query causes no additional section processing
- inverse lookup uses address sequence, PTR query
- inverse lookup name is a decimal and hex string

IP/SIPP Transition Issues

- not known whether hosts have IP or SIPP addresses
  - resolvers/application (libraries) must query for both
- only A RRs returned in additional section of NS, MB and MX records
  - extend definitions to include ASEQs as well

SIPP

DHCP and Autoconfiguration

Topics

1. Framework of Host
2. DHCP Protocol Overview
3. DHCP Changes with SIPP
4. Autoconfiguration Requires DHCP
5. DHCP Database Issue
6. Thought Experiments

A Defacto Host TCP/IP Architecture

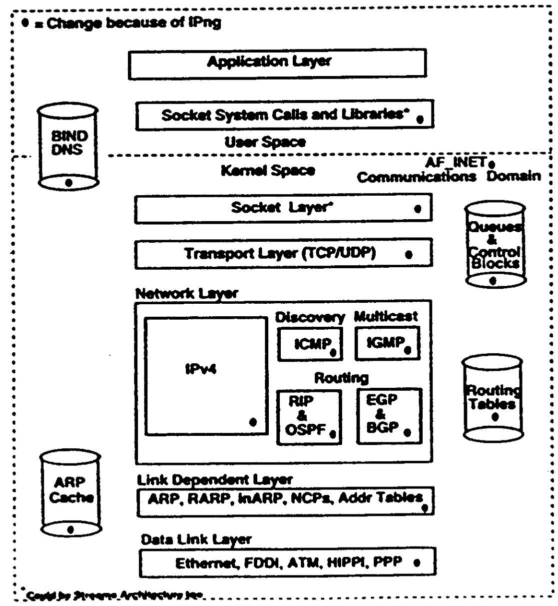

IPng Directorate - Jim Bound (Digital)

2.3.3 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, TCP, and 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.
Done	Post the TP/IX Protocol and the RAP protocol as Experimental RFCs.
Done	Hold Working Group meeting to discuss additional definitions. Prepare criteria to be met prior to standardization.
Done	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

"Initial AD Assignment Plan", 06/30/1993, R. Ullmann <draft-ietf-tpix-adplan-01.txt>

"Transit Policy Routing in TP/IX", 06/30/1993, R. Ullmann <draft-ietf-tpix-transit-01.txt>

"TCP version 7 options", 06/30/1993, R. Ullmann <draft-ietf-tpix-tcpopt-00.txt>

"Common Architecture Technology for Next-generation Internet Protocol", 01/06/1994, R. Ullmann <draft-ietf-tpix-catnip-base-01.txt>

Request For Comments

RFC 1475 "TP/IX: The Next Internet"

RFC 1476 "RAP: Internet Route Access Protocol"

CURRENT MEETING REPORT

Reported by Vladimir Sukonnik/Process Software Corporation

Minutes of the TP/IX Working Group (TPIX)

The TPIX Working Group met jointly with the TUBA Working Group and also met in a second independent session. The agenda for that meeting was:

- Introduce and discuss CATNIP
- Review TP/IX charter and name

The meeting started with Ross Callon introducing the concept of the Forward Cache Identifier, or handle, to be used to speed up processing in routers. A downstream router may send an ICMP message offering an FCI for a particular source, destination, and type of service. The source may then use the FCI in its packet instead of fully specified source and destination addresses. Using the FCI will achieve two goals: smaller packet size and faster processing in the router.

CATNIP

Robert Ullmann introduced CATNIP. CATNIP is a revision of the TP/IX proposal. The TP/IX packet has been extended to include a Forward Cache Identifier and NSAP-style source and destination addresses. Using a NSAP-style address, CATNIP could be used to represent IPv4, IPX and OSI protocols. CATNIP could also be used to connect CLNP systems to IPv4 and IPX systems.

Several people suggested that placing the Forward Cache Identifier in the first longword of the packet may speed up processing. Robert pointed out two reasons for not doing so. First, the first byte must be reserved for NLPID field. Second, it was observed that any current or future processor will be loading at least 64 bits in parallel anyway. It is also easier to make FCI fixed field rather than a variable size.

TCP used to be part of the TP/IX proposal. The working group felt that it would be better to separate TCP (and UDP) from the CATNIP proposal and have them addressed as separate issues.

The working group decided to remove RAP (RFC 1476) from its charter, to be developed separately. It was noted by Dave Katz and others that IS-IS will work fine with CATNIP, as will the IDRP and other methods; the existing OSI routing could be used for CATNIP's addressing scheme without any changes.

A concern was raised that the selection field in the CATNIP header may not be long enough. It seems, however, that 16 bits is a reasonable size for fields that assign one code point to each of a set of protocols (at least below application layer). It's hard to see us designing more than 65000 transport layer protocols.

The TP/IX working group also defined milestones for Seattle:

- Rob will add additional details and publish CATNIP as an Internet-Draft.
- Rob will write the white paper requested by the IPng directorate as soon as the outline is available.
- Vladimir will rewrite TP/IX charter and work with Scott Bradner on renaming the working group to CATNIP.
- The group will plan on meeting jointly with TUBA for one session in Seattle, and will continue to coordinate efforts to find as much common ground as possible.

Attendees

Nick Alfano	alfano@mpr.ca
Vadim Antonov	avg@icm1.icp.net
Susie Armstrong	susie@mentat.com
Jules Aronson	aronson@nlm.nih.gov
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Scott Bradner	sob@harvard.edu
Glen Cairns	cairns@mprgate.mpr.ca
Ross Callon	rcallon@wellfleet.com
Peter Cameron	cameron@xylint.co.uk
George Chang	gkc@ctt.bellcore.com
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
Richard Colella	colella@nist.gov
Michael Collins	collins@es.net
David Conrad	davidc@iij.ad.jp
Matt Crawford	crawdad@fncent.fnal.gov
John Curran	jcurran@nic.near.net
Michael Davis	mike@dss.com
Chuck de Sostoa	chuckd@cup.hp.com
Stephen Deering	deering@parc.xerox.com
Avri Doria	avri@locus.com
Robert Fink	rlfink@lbl.gov
Eric Fleischman	ericf@atc.boeing.com

Peter Furniss	p.furniss@ulcc.ac.uk
Eugene Geer	ewg@cc.bellcore.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Chris Gunner	gunner@dsmail.lkg.dec.com
Denise Heagerty	denise@dxcoms.cern.ch
Phil Irey	pirey@relay.nswc.navy.mil
Kevin Jackson	kjackson@concord.com
David Jacobson	dnjake@vnet.ibm.com
Dave Katz	dkatz@cisco.com
Elizabeth Kaufman	kaufman@biomded.med.yale.edu
Edwin King	eek@atc.boeing.com
Jian Li	jian@rice.edu
Kanchei Loa	loa@sps.mot.com
E. Paul Love	epl@sdsc.edu
Tracy Mallory	tracym@3com.com
David Marlow	dmarlow@relay.nswc.navy.mil
Jun Matsukata	jm@eng.isas.ac.jp
Doug Montgomery	dougm@osi.ncsl.nist.gov
Dan Nordell	
Erik Nordmark	nordmark@eng.sun.com
William Palter	palter@tgv.com
Andrew Partan	asp@uunet.uu.net
Radia Perlman	perlman@novell.com
Eric Peterson	elpeterson@eng.xyplex.com
David Piscitello	wk04464@worldlink.com
James Quigley	Quigley@cup.hp.com
Martin Schulman	schulman@smtp.sprint.com
Vincent Shekher	vin@sps.mot.com
Erik Sherk	sherk@sura.net
Uttam Shikarpur	uttam@zk3.dec.com
Keith Sklower	sklower@cs.berkeley.edu
Frank Solensky	solensky@ftp.com
Vladimir Sukonnik	sukonnik@process.com
Steve Suzuki	steve@fet.com
Larry Tepper	ltepper@compatible.com
Richard Thomas	rjthomas@bnr.ca
Robert Ullmann	ariel@world.std.com
William Warner	warner@ohio.gov
Chris Wheeler	cwheeler@cac.washington.edu
Gerry White	gerry@lancity.com
Cathy Wittbrodt	cjw@barrnet.net
David Woodgate	David.Woodgate@its.csiro.au

2.3.4 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 specifies their encapsulation 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 OSI IDRP for IP Over 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 RFC 1347 reflecting lessons earned in the Working Group deliberation.

Request For Comments

RFC 1526	"Assignment of System Identifiers for TUBA/CLNP Hosts"
RFC 1561	"Use of ISO CLNP in TUBA Environments"

CURRENT MEETING REPORT

Reported by Mark Knopper/Merit

Minutes of the TCP/UDP over CLNP-Addressed Networks Working Group (TUBA)

The meeting was called to order by Mark Knopper and Peter Ford, co-chairs. Dave Katz volunteered to act as recording secretary. Mark Knopper presented the agenda.

CLNP Multicast

Dave Marlow reported on the progression of CLNP Multicast work at the recent ISO SC6 meeting in Seoul. There are currently four documents in progression—changes to the network addressing addendum, extensions to the CLNP and ES-IS protocols, and a change to the network service definition. Documentation of these changes will be released as an Internet-Draft.

The Group Network Addressing addendum has progressed to Full Standard status. This document describes the syntax of multicast NSAP addresses.

The extensions to CLNP have been issued for Draft Amendment (DAM) ballot. This is the final balloting stage before full standardization (similar to Draft Internet Standard status). The changes include a new packet type (so that unicast-only routers do not try to default-forward multicast packets, which could create packet-exploder loops) and two types of scope control.

The ES-IS extensions have also been issued for DAM ballot. These changes provide for the ability of end systems to join multicast groups, as well as for dynamic mapping between network and subnetwork multicast addresses.

The change to the Network Service definition adds multicast capability to the abstract network service.

There has been no significant work on multicast routing, as ANSI is looking to work with the IETF for technical contribution in this area. Possibilities include MOSPF-like extensions to IS-IS, CBT, and ESL.

NSAP Addressing Guidelines Document

Ross Callon reported on the status of the NSAP Addressing Guidelines document. A new version has been made available as an Internet-Draft. Changes in the new version consist primarily of updated document references. Several people mentioned that they thought the current document was overly "backbone-centric," as the Internet routing hierarchy has changed considerably since the document was originally written. An action item was taken to recast this section of the document.

CLNP Projects

Yakov Rekhter gave a brief overview of two projects just getting started in ISO concerning CLNP. The first describes a coding method for option types that provides a hook for transparently adding options. The code indicates which options are required to be processed by routers and/or hosts, even when the option is otherwise unrecognized.

The second project extends the Quality of Service (QoS) option to provide a bit to request strong (vs. weak) QoS forwarding, as well as a bit to say whether or not the requested OoS was delivered along the entire path when weak QoS forwarding is in effect.

CLNP Mobility

A discussion then ensued on the subject of CLNP mobility. Mark Knopper briefly described CDPD, a specification for cellular mobile data service from a consortium of cellular carriers. The system uses CLNP as the primary protocol, and provides IP service using IP-over-CLNP encapsulation. The mobility protocol is quite similar to ongoing work in the Mobile IP Working Group. The group discussed whether or not it should be proactive, or wait for the Mobile IP Working Group to settle. Yakov Rekhter and Dave Piscitello agreed to recast the mobile IP document in terms of CLNP and make it available as an Internet-Draft.

Extensions to OSI for Use in the Internet BOF (OSIEXTND)

Dave Katz briefly updated the group on the status of OSIEXTND. The IESG has issued a formal statement limiting official IETF participation to ongoing "OSI-related work." No new working groups in this area will be chartered until such time as the liaison issue between ISO and ISOC has been settled, or until six months pass. What was not clearly stated at the time was that this action does not apply to anything that could be construed as in support of TUBA. As the majority of the items of work proposed at the OSIEXTND BOF are directly related to TUBA, the net effect is that work will progress in existing working groups.

To that end, Dave Katz then presented a proposed extension to ES-IS to add functionality to the standard dynamic NSAP address assignment function. The extension would allow the end system to suggest a system ID to the entity assigning the addresses, which would then fill in the remainder of the NSAP address if it so chose. A second contribution describing the overall dynamic address assignment mechanism was also presented. Both documents will be made available as Internet-Drafts.

Dual Stack Transition

Peter Ford presented his draft document on the Dual Stack Transition plan. It is an "inside out" approach that begins with infrastructure deployment. It was pointed out that this transition framework needed to be completed as soon as possible. The document is available as an Internet-Draft.

Attendees

Nick Alfano	alfano@mpr.ca
Susie Armstrong	susie@mentat.com
Richard Binder	rbinder@cnri.reston.va.us
Rebecca Bostwick	bostwick@es.net
Scott Bradner	sob@harvard.edu
Arthurine Breckenridge	arbreck@sandia.gov
Glen Cairns	cairns@mprgate.mpr.ca
Ross Callon	rcallon@wellfleet.com
George Chang	gkc@ctt.bellcore.com
John Chang	jrc@uswest.com
Enke Chen	enke@merit.edu
Richard Colella	colella@nist.gov
Michael Conn	4387451@mcimail.com
Dino Farinacci	dino@cisco.com
Robert Fink	rlfink@lbl.gov
Peter Ford	peter@goshawk.lanl.gov
Peter Furniss	p.furniss@ulcc.ac.uk
Eugene Geer	ewg@cc.bellcore.com
Tony Genovese	genovese@es.net
Herluf Hansen	hha@tbit.dk
Susan Hares	skh@merit.edu
Denise Heagerty	denise@dxcoms.cern.ch
Kathy Huber	khuber@wellfleet.com
David Jacobson	dnjake@vnet.ibm.com
Dave Katz	dkatz@cisco.com
Mark Knopper	mak@merit.edu
Walter Lazear	lazear@gateway.mitre.org
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@cmf.nrl.navy.mil
David Marlow	dmarlow@relay.nswc.navy.mil
Matt Mathis	mathis@psc.edu
Greg Minshall	minshall@wc.novell.com
Doug Montgomery	dougm@osi.ncsl.nist.gov
Robert Moose	rmoose@gateway.mitre.org

Dennis Morris	morris@altair.disa.mil
Dan Nordell	
Radia Perlman	perlman@novell.com
David Piscitello	wk04464@worldlink.com
Alex Reijnierse	a.a.l.reijnierse@research.ptt.nl
Yakov Rekhter	yakov@watson.ibm.com
Dallas Scott	scott@fluky.mitre.org
Erik Sherk	sherk@sura.net
Vladimir Sukonnik	sukonnik@process.com
Larry Tepper	ltepper@compatible.com
Michael Thatcher	thatcher@rahul.net
Richard Thomas	rjthomas@bnr.ca
Robert Ullmann	ariel@world.std.com
John Veizades	veizades@ftp.com
William Warner	warner@ohio.gov
Gerry White	gerry@lancity.com
Scott Williamson	scottw@nic.ddn.mil
Cathy Wittbrodt	cjw@barrnet.net
Jane Wojcik	jwojcik@bbn.com
Jessica Yu	jyy@merit.edu

CURRENT MEETING REPORT

Reported by Mark Knopper/Merit

Minutes of the Joint Session of the TUBA and TPIX Working Groups

The two groups met jointly during the second scheduled TUBA session, primarily to discuss the CATNIP proposal. Several other TUBA items remained to be discussed after the first meeting.

Ross Callon began by introducing his ideas on CLNP header compression and flow setup, in relation to the CATNIP ideas. This was followed by a presentation of the CATNIP paper by Rob Ullman.

Editor's Note: Details from each of these presentations and a summary of discussion is available via FTP or mail server from the remote directories as /ietf/tuba/tuba-tpix-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

TUBA Transition Plan (Dave Piscitello)

Dave worked with Tracy Mallory and Jim Bound to develop an outline of a transition plan.

Editor's Note: A transcript of Dave's slides is available via FTP or mail server from the remote directories as /ietf/tuba/tuba-tpix-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

This outline was well received as a good start to a transition plan paper. Some of the points included in the transcript were comments added from the attendees. It was also suggested that the transition plan paper be very clear about where changes are need to hosts as distinguished from routers.

ISO Liaison (Peter Ford)

Peter gave an overview of the current status of the liaison between ISOC and ISO.

Vint Cerf and Jack Houldsworth had discussions at the last IETF. ISOC recently forwarded two letters to ISO—these are Internet-Drafts. Also there is an Internet-Draft, "Liaison between Internet and other Standardization Agencies," by Christian Huitema on this topic.

For TUBA, the issue is change control. Lyman Chapin is working on document distribution. RFC 1310 bis contains language about ceding all copyright control to IETF from ISO. Document review and comments are encouraged. The document can be found in the Internet-Drafts directory. One issue is how can the IETF take documents from other standards bodies into the Internet Standards process?

Concerning the Memorandum of Understanding between ISO and ISOC, Peter felt that convergence in the network layer should be suggested. Also there should be an address change control for the base network protocol document. The SC6 contribution is in line with this.

Peter Furniss is a SC21 member. Both groups claim to be more open than the other. ISO did not understand how IETF/ISOC process works and comes to consensus. Scott Bradner pointed out that RFC 1310 is a description of our process and can be used to help communicate to other groups.

It was discussed that either ISO can retain change control and IETF can have official liaison to ISO; or the IETF can take a clone of CLNP and diverge (with report back to ISO).

CLNP Routing in Europe (Alex Reijnierse)

Alex presented a connectivity diagram of the CLNP Internet from the European (Europanet) perspective.

2.4 Network Management Area

Director:

- Marshall Rose: `mrose.iesg@dbc.mtview.ca.us`

Area Summary reported by Marshall Rose/Dover Beach Consulting

A working group is either active or inactive. Active working groups have charters to develop documents. Inactive working groups have no charter – typically because they have completed their previous charter. These inactive working groups (and their mailing lists) serve as a forum for implementors. When a standards-track document produced by a working group is ready for further evaluation or new documents are appropriate, the working group is re-chartered accordingly.

Due to a lack of senior technical resources in the Network Management area, there is a moratorium on new working groups for the 1993 calendar year. At the beginning of 1994, this policy will be revisited. Until then, only things of the utmost urgency will be receive any consideration for possible working group activation. This policy does not apply to working groups currently inactive, awaiting re-activation due to standards-track activity.

ATM MIB Working Group (ATOMMIB)

The working group is active, editing "Definitions of Managed Objects for ATM Management Version 4.0" (draft-ietf-atommib-atm-01.txt) and "Definitions of Managed Objects for the SONET/SDH Interface Type" (draft-ietf-atommib-sonet-01.txt).

The SONET draft is completed, but a few issues remain on the ATM draft. Completion is expected by early December 1993.

Bridge MIB Working Group (BRIDGE)

The working group is inactive, awaiting the next stage for RFC 1493 and RFC 1525. BRIDGE is eligible to re-activate in April 1994.

Character MIB Working Group (CHARMIB)

The working group is active, waiting for the Interfaces MIB Working Group (IFMIB) to complete the "Evolution of the Interfaces Group of MIB-II" Internet-Draft before proceeding with the evaluation of RFCs 1316-1318 (Proposed Standards) with respect to the standards track.

DECnet Phase IV MIB Working Group (DECNETIV)

The working group is active, waiting for an IESG action on "DECnet Phase IV MIB Extensions" (draft-ietf-decnetiv-mibext-03.txt). The IESG is considering whether to elevate this document to Draft Standard status.

FDDI MIB Working Group (FDDIMIB)

The working group is inactive, awaiting the next stage for RFC 1285 and RFC 1512. FDDIMIB is eligible to re-activate in March 1994.

Frame Relay Service MIB Working Group (FRNETMIB)

The working group is active, editing "Definitions of Managed Objects for Frame Relay Service" (draft-ietf-frnetmib-fr-04.txt). Completion of the Frame Relay Service MIB is expected by mid-November 1993.

Host Resources MIB Working Group (HOSTMIB)

The working group is inactive, awaiting the next state for RFC 1514 (Proposed Standard). HOSTMIB is eligible to re-activate in March 1994.

IEEE 802.3 Hub MIB Working Group (HUBMIB)

The working group is inactive, awaiting the next stage for RFC 1515 (Proposed Standard) and RFC 1516 (Draft Standard). HUBMIB is eligible to re-activate in March 1994.

Interfaces MIB Working Group (IFMIB)

The working group is active, editing "Evolution of the Interfaces Group of MIB-II" (draft-ietf-ifmib-evolution-04.txt) and "Management Information Base for Management of Network Connections" (draft-ietf-ifmib-conntable-00.txt).

Both the Interfaces Evolution and Connection Table MIBS are virtually completed. Wrap-up is expected by mid-November 1993. Following this, the working group is chartered to evaluate RFCs 1229, 1231, 1304 (Proposed Standards), and 1398 (Draft Standard) with respect to the standards track.

Mail and Directory Management Working Group (MADMAN)

The working group is active, waiting for editorial updates. The Network Management Directorate (NMDIR) completed its evaluation of "Network Services Monitoring MIB" (draft-ietf-madman-networkmib-05.txt), "Mail Monitoring MIB" (draft-ietf-madman-mtamib-06.txt), and "X.500 Directory Monitoring MIB" (draft-ietf-madman-dsa-mib-05.txt). The editors are producing new drafts.

Modem Management Working Group (MODEMMGT)

The working group is active, reviewing "Modem MIB" (draft-ietf-modemmgt-mdmmib-00.txt). The group reviewed its first draft, but concensus is lacking. MODEMMGT is also late in meeting its milestones. The area director has given the working group until January 31, 1994 to reach concensus on a completed draft; otherwise, the working group will be terminated.

Remote LAN Monitoring Working Group (RMONMIB)

The working group is active, evaluating RFC 1271 (Proposed Standard) with respect to the standards track.

SNA DLC Services MIB Working Group (SNADLC)

The working group is active, editing "Definitions of Managed Objects for SNA Data Link Control: SDLC" (draft-ietf-snadlc-sdlc-mib-00.txt). The SDLC MIB is expected to be complete by mid-December 1993.

SNA NAU Services MIB Working Group (SNANAU)

The working group is active, editing "Definitions of Managed Objects for SNA NAUs" (draft-ietf-snanau-snamib-00.txt). The NAU MIB is expected to be complete by mid-December 1993.

SNMP Version 2 Working Group (SNMPV2)

The working group is inactive, awaiting the next stage for RFCs 1441-1452 (Proposed Standards). SNMPV2 is eligible to re-activate in October 1994.

Token Ring Remote Monitoring Working Group (TRMON)

The working group is inactive, awaiting the next state for RFC 1513 (Proposed Standard). TRMON is eligible to re-activate in March 1994, but will not do so. Instead, the RMONMIB Working Group will be tasked to evaluate RFC 1513.

DS1/DS3 Working Group (TRUNKMIB)

The working group is inactive, awaiting the next stage for RFCs 1406 and 1407 (Proposed Standard). TRUNKMIB is eligible to re-activate now. However, it will remain inactive until the beginning of 1994.

Uninterruptible Power Supply Working Group (UPSMIB)

The working group is active, editing "UPS Management Information Base" (draft-ietf-upsmib-00.txt).

UPSMIB is several months overdue based on milestones in its charter. The area director has given the working group until November 26, 1993 to reach concensus on a completed draft; otherwise, the working group will be terminated.

X.25 Management Information Base Working Group (X25MIB)

The working group is inactive, awaiting the next stage for RFCs 1381-1382 (Proposed Standards). X25MIB is eligible to re-activate now. However, it will remain inactive until the beginning of 1994.

CURRENT MEETING REPORT

Reported by Joel Gyllenskog/Hewlett-Packard

Minutes of the Printer MIB BOF (PRINTMIB)

A meeting of the PrintMIB BOF was held Wednesday, November 3, 1993 from 1600-1800 at the IETF meeting in Houston. The meeting was chaired by Joel Gyllenskog of Hewlett-Packard. Marshall Rose, the Network Management Area Director was present.

Of the twenty-three whose names were on the roster, all but five asked to have their names added to the mailing list.

The agenda, as posted on the Internet, was presented. Joel gave a brief history of some of the changes in printing on LANs over the past five years. He provided rationale as to why the printer manufacturers have an interest in having their printers be manageable in a network environment. He talked about other efforts that have occurred or that are underway including the DMTF printer team. This group had representatives of thirteen different printer manufacturers at their October meeting, along with representatives from other companies with an interest in the use of printers. The goal of the group is to specify a set of objects that can be used to manage printers and to have that set be the same for the DeskTop as for the network. This desire is shared by members of COSE and DSIS.

A discussion of the potential number of objects that would fall into a standard printer MIB was led by Marshall. The sense of the group was that there will probably be on the order of from forty to sixty objects.

Steve Waldbusser led a discussion of existing MIBs and how they may be used by printers. This includes the Host MIB, MIB II, and the Character MIB. Steve has indicated a willingness to provide technical oversight to the group.

The idea was put forward that fonts, spooling, and print job management are not part of the set that will be included in the print MIB. These areas are interesting and important, but their inclusion would add significantly to the effort of the group and decrease the likelihood of a timely completion of the MIB.

The areas that were discussed that the group felt could profitably be included in a print MIB included: Print Engine, Interpreters, Media, Input Sources, and Output Destinations.

Those present indicated the importance of using the net to allow broad input and participation.

A discussion of the time schedule proposed resulted in a recommendation that the charter specify completion by the July IETF meeting, but that the group work to accomplish the task by May.

The consensus of those in attendance was that a proposed charter be given to the Area Director with the recommendation that a Print MIB Working Group be established.

Attendees

Jim Barnes	barnes@xylogics.com
Bill Cash	cash@bangote.compaq.com
Manuel Diaz	diaz@davidsys.com
Jonathan Didner	jonb@bangate.compaq.com
Christine Gressley	gressley@uiuc.edu
Joel Gyllenskog	jgyllens@hpdmd48.boi.hp.com
Melanie Humphrey	msh@uiuc.edu
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@ftp.com
Charley Kline	cvk@uiuc.edu
Andrew Knutsen	andrewk@sco.com
David Lapp	lapp@waterloo.hp.com
Larry Masinter	masinter@parc.xerox.com
Jon Saperia	saperia@zko.dec.com
Michael Scanlon	scanlon@ftp.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@timonware.com
Bob Stewart	rlstewart@eng.xyplex.com
Dean Throop	throop@dg-rtp.dg.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Steve Zilles	szilles@adobe.com

2.4.1 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

Done	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.

Internet-Drafts

"Definitions of Managed Objects for the SONET/SDH Interface Type", 01/03/1994, Tracy Brown, Kaj Tesink <draft-ietf-atommib-sonet-04.txt>

"Definitions of Managed Objects for ATM Management Version 4.0", 12/22/1993, M. Ahmed, K. Tesink <draft-ietf-atommib-atm-03.txt>

CURRENT MEETING REPORT

Reported by Kaj Tesink/Bellcore

Minutes of the ATM MIB Working Group (ATOMMIB)

SONET MIB

The current Internet-Draft, ietf-draft-atommib-sonet-01.txt, is considered complete. The meeting participants decided unanimously to recommend this Internet-Draft to the area director for further processing by the Network Management Area Directorate (NMDIR) and the IESG as a Proposed Standard.

ATM MIB

- Use of ifTable for the ATM Level

  The mapping in the current Internet-Draft was agreed upon with a minor amendment to the mapping of ifOutErrors.

- Use of ifTable for AAL5

  A new approach was adopted. The detailed text will be worked out on the mailing list. The approach treats AAL5 entities in switches as connected with the switch through a virtual interface; in hosts the AAL5 level is directly stacked on the ATM level. The approach also requires a small AAL5-specific table in the ATM MIB for AAL5 error counters. Explanatory text on this subject will also be included in the specification.

- Modeling of ATM Connections

  The small group tasked at the Amsterdam meeting has identified two approaches. Ted Brunner identified the approaches:

  1. Proposal Bob - Similar to what is contained in the current ATM MIB and Frame Relay MIB Internet-Drafts.

  2. Proposal Dave - Takes a common approach for hosts, switches, and services.

  The meeting participants decided to proceed with an approach, dubbed "Henrietta", that combines the strong points of Dave and Bob. The current MIB will be amended as appropriate. Detailed text as to how to do orderly connection setup will also be included.

- Any Other ATM MIB Issues

  The proposal, to combine the separate VPL and VCL tables into a single table, and the VCC and VPC tables into a single table, was reviewed. This proposal was not adopted.

- Status of the ATM MIB

  Beyond the details mentioned earlier, no other open issues where identified. It was decided unanimously to recommend this Internet-Draft to the area director for further processing by NMDIR and the IESG as a Proposed Standard after these details have been fixed on the mailing list. The target is to wrap up before the end of the year (in accordance with the charter).

- ATM Management Beyond the ATM MIB

  Some interest had been expressed to pursue management of switched virtual circuits (SVC). However, it was decided not to pursue this at this time. Consequently, the working group can be inactivated after the ATM MIB and SONET MIB have been progressed as Proposed Standards. The mailing list will remain active for the exchange of implementation experience or any other ATM/SONET management discussion.

Generic Connection Table

Ken Rodeman gave a presentation on a generic approach to the management of virtual connections, suggesting a common approach for frame relay and ATM. The generic approach would serve as an umbrella over connection tables that are specific to frame relay or ATM. The contents of the specific tables would not be affected by adoption of the generic approach. Rather, the specific approach would simplify the overall management of connections. Discussion of this topic was deferred to the Interfaces MIB Working Group.

Attendees

Masuma Ahmed	mxa@mail.bellcore.com
Bart Berger	bart_berger@3com.com
Caralyn Brown	cbrown@wellfleet.com
Tracy Brown	tacox@mail.bellcore.com
Theodore Brunner	tob@thumper.bellcore.com
Ted Chang	tchang@atg.wiltel.com
Chris Chiotasso	chris@lightstream.com
Blair Copland	copland@unt.edu
Kurt Dobbins	dobbins@ctron.com
Christopher Dorsey	dorsey@es.net

David Engel	david@ods.com
Dario Ercole	Dario.Ercole@cselt.stet.it
Michael Erlinger	mike@jarthur.claremont.edu
William Fardy	billf@frontier.com
Steve Feldman	feldman@mfsdatanet.com
Robert Fenoglio	fenoglio@vnet.ibm.com
Mark Garrett	mwg@faline.bellcore.com
Steve Garritano	steveg@kalpana.com
Robert Grow	bob@xlnt.com
Stuart Hale	stu_hale@vnet.ibm.com
Augusto Ham	ahamm@atg.wiltel.com
Daniel Hansen	danh@ngc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Mike Holloway	mikeh@newbridge.com
Refael Horev	horev@lannet.com
Frank Kastenholz	kasten@ftp.com
Mark Kepke	mak@fc.hp.com
Cheryl Krupczak	cheryl@empiretech.com
Ted Kuo	tik@vnet.ibm.com
Andrew Malis	malis@maelstrom.timeplex.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Steven Onishi	sonishi@wellfleet.com
Zbigniew Opalka	zopalka@agile.com
David Perkins	dperkins@synoptics.com
Kenneth Rodemann	krr@qsun.att.com
Hal Sandick	sandick@vnet.ibm.com
Michael Scanlon	scanlon@ftp.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Kanan Shah	kshah@cmf.nrl.navy.mil
Satya Sharma	ssharma@chang.austin.ibm.com
Uttam Shikarpur	uttam@zk3.dec.com
Timon Sloane	timon@timonware.com
Kaj Tesink	kaj@cc.bellcore.com
Panos Tsigaridas	Tsigaridas@fokus.gmd.de
Thomas Walsh	tomw@kalpana.com
Alice Wang	alice.wang@eng.sun.com
Honda Wu	honda@nat.com

2.4.2 Character MIB (CHARMIB)

Charter

Chair(s)

Bob Stewart: rlstewart@eng.xyplex.com

Mailing Lists

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

Description of Working Group

The Character MIB Working Group is chartered to 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.
Done	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"

2.4.3 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

Done	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", 01/06/1994, T. Brown <draft-ietf-frnetmib-fr-07.txt>

"Service Management Architecture for Virtual Connection Services", 07/02/1993, K. Rodemann <draft-ietf-frnetmib-virtual-sma-01.txt>

CURRENT MEETING REPORT

Reported by James Watt/Newbridge Networks Corporation

Minutes of the Frame Relay Service MIB Working Group (FRNETMIB)

Agenda

- Administrivia
- Comments from the chair of the FRFTC Working Group
- Discussion of non-Connection model issues
- Discussion of Connection model Issues
- Wrap Up

Comments from the Chair of the FRFTC Working Group

Andy Malis summarized the latest meeting of the corresponding working group in the Frame Relay Forum Technical Committee (FRFTC). They held a walk-through of the MIB with a small number of people and were satisfied with the current state.

Frame Relay MIB

The current Internet-Draft (`draft-ietf-frnetmib-04.txt`) will be complete pending:

- A few editorial changes
- Any changes to the connection model to align it with the ATOMMIB connection model

Once these items are finished, the group agreed that the updated Internet-Draft should be forwarded to the Area Director for processing by the Network Management Directorate (NMDIR) and submitted to the IESG for consideration as a Proposed Standard.

Frame Relay and ATM MIB Connection Model Issues

Subsequent discussions in the ATM MIB (ATOMMIB), Frame Relay Service MIB (FRNET-MIB) and Interfaces MIB (IFMIB) Working Groups led to the adoption of a new connection model for both the ATM and FRNET MIBs.

Conclusion

The updated Internet-Draft will be circulated on the mailing list by 22 November 1993 for comment. Given the recommendation of those present in Houston, there are no known reasons to prevent this document from being forwarded.

Attendees

Tracy Brown	tacox@mail.bellcore.com
Chris Chiotasso	chris@lightstream.com
Christopher Dorsey	dorsey@es.net
Stuart Hale	stu_hale@vnet.ibm.com
John Hanratty	jhanratty@agile.com
Jeff Hilgeman	jeffh@apertus.com
Mike Holloway	mikeh@newbridge.com
Deirdre Kostick	dck2@mail.bellcore.com
Cheryl Krupczak	cheryl@empiretech.com
Joseph Liu	jliu@atg.wiltel.com
Andrew Malis	malis@maelstrom.timeplex.com
Peram Marimuthu	peram@wg.com
Rina Nathaniel	rina@rnd-gate.rad.co.il
Orly Nicklass	orly@radmail.rad.co.il
Tom Nisbet	nisbet@fbsw.tt.com
Shannon Nix	sdn@netlink.com
Steven Onishi	sonishi@wellfleet.com
Kenneth Rehbehn	kjr@netrix.com
Kenneth Rodemann	krr@qsun.att.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Chi Shue	chi@casc.com
Kaj Tesink	kaj@cc.bellcore.com
Michael Thatcher	thatcher@rahul.net
Alice Wang	alice.wang@eng.sun.com
James Watt	james@newbridge.com
Honda Wu	honda@nat.com
Mauro Zallocco	mdz@netlink.com

2.4.4 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: thumper.bellcore.com:pub/tob/ifmib

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 RFC 1229 (the interface-extensions MIB), RFC 1231 (the token-ring MIB), RFC 1304 (the SMDS MIB), and RFC 1398 (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 RFC 1398, 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 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

Done	Post the interface layering document as an Internet-Draft.
Sep 1993	Submit the interface 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", 12/17/1993, K. McCloghrie, F. Kastenholz <draft-ietf-ifmib-evolution-07.txt>

"Management Information Base for Management of Network Connections", 10/21/1993, K. Rodemann <draft-ietf-ifmib-conntable-00.txt>

CURRENT MEETING REPORT

Reported by Theodore Brunner/Bellcore

Minutes of the Interfaces MIB Working Group (IFMIB)

Discussion focussed on the "Evolution of the Interfaces Group of MIB II" Internet-Draft of October 20, 1993. The discussion reviewed issues raised on the mailing list and was the final forum for concerns about the Internet-Draft before it was forwarded from the IFMIB Working Group to the Area Director for consideration as a Proposed Standard.

Evolution of the Interfaces Group of MIB II Internet-Draft

Editor's Note: Summaries of issues discussed are available via FTP or mail server from the remote directories as /ietf/ifmib/ifmib-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Generic Connection Table

An initial presentation was made on the motivation for the current document, and a number of questions were raised, but not fully answered. There is concern as to whether this would delay the ATM and Frame Relay MIBs nearing completion now.

The following questions were raised:

- Will there be benefits to all constituencies from such a generic model (ATM/Frame Relay or CNM/Device Management)?
- What is the meaning of a connection AdminStatus versus media specific AdminStatus?
- Is cnTable mandatory?
- Does cnTable contain any additional information to what is in the media specific tables?
- Is there a difference between CNM and device cross connect?

There is a desire to pursue an effort on generic connection tables, but there is no desire to delay existing efforts; the ATM and Frame Relay MIBs will proceed as scheduled. A generic connection effort will start, and be formulated as an independent addition to those efforts. It is presumed that the questions raised at this meeting will be addressed as part of that effort.

Attendees

Masuma Ahmed	mxa@mail.bellcore.com
Fred Baker	fbaker@acc.com
Jim Barnes	barnes@xylogics.com
Bart Berger	bart_berger@3com.com
Andy Bierman	abierman@synoptics.com
Tracy Brown	tacox@mail.bellcore.com
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Jeff Case	case@cs.utk.edu
Chris Chiotasso	chris@lightstream.com
Manuel Diaz	diaz@davidsys.com
Jonathan Didner	jonb@bangate.compaq.com
Kurt Dobbins	dobbins@ctron.com
Christopher Dorsey	dorsey@es.net
David Engel	david@ods.com
William Fardy	billf@frontier.com
Steve Feldman	feldman@mfsdatanet.com
Robert Fenoglio	fenoglio@vnet.ibm.com
Christine Gressley	gressley@uiuc.edu
Stuart Hale	stu_hale@vnet.ibm.com
Mike Holloway	mikeh@newbridge.com
John Hopprich	hopprich@davidsys.com
Refael Horev	horev@lannet.com
Melanie Humphrey	msh@uiuc.edu
Frank Kastenholz	kasten@ftp.com
Mark Kepke	mak@fc.hp.com
Michael Kornegay	mlk@bir.com
Deirdre Kostick	dck2@mail.bellcore.com
Cheryl Krupczak	cheryl@empiretech.com
Joseph Liu	jliu@atg.wiltel.com
Andrew Malis	malis@maelstrom.timeplex.com
Peram Marimuthu	peram@wg.com
Evan McGinnis	bem@3com.com
Steve McRobert	steve.mcrobert@amd.com
Rina Nathaniel	rina@rnd-gate.rad.co.il
Sath Nelakonda	sath@lachman.com
Orly Nicklass	orly@radmail.rad.co.il
Tom Nisbet	nisbet@fbsw.tt.com
Shannon Nix	sdn@netlink.com
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
Zbigniew Opalka	zopalka@agile.com
Eric Peterson	elpeterson@eng.xyplex.com

Kenneth Rehbehn	kjr@netrix.com
Kenneth Rodemann	krr@qsun.att.com
Michal Rozenthal	michal@fibronics.co.il
Jon Saperia	saperia@zko.dec.com
Michael Scanlon	scanlon@ftp.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Chi Shue	chi@casc.com
Timon Sloane	timon@timonware.com
Andrew Smith	asmith@synoptics.com
Robert Snyder	snyder@cisco.com
Louis Steinberg	louiss@vnet.ibm.com
Bob Stewart	rlstewart@eng.xyplex.com
Adam Stolinski	stolinsk@cerf.net
Kaj Tesink	kaj@cc.bellcore.com
Michael Thatcher	thatcher@rahul.net
Dean Throop	throop@dg-rtp.dg.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Alice Wang	alice.wang@eng.sun.com
James Watt	james@newbridge.com
Evan Wetstone	evanw@vnet.ibm.com
Peter Wilson	peter_wilson@3mail.3com.com

2.4.5 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 14's work on an OSI CMIS/CMIP object definition for V series DCEs entitled "Managed Object Template for V-Series DCE's."

Goals and Milestones

Done	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.

Internet-Drafts

"Modem MIB", 10/26/1993, L. Brown, R. Roysten, S. Waldbusser <draft-ietf-modemmgt-mdmmib-00.txt>

CURRENT MEETING REPORT

Reported by Mark Lewis/Telebit Corporation

Minutes of the Modem Management Working Group (MODEMMGT)

Summary

The third meeting of MODEMMGT was attended by twenty or so people. The Network Management Area Director noted that the group was behind schedule and would be given until January 31, 1994 to reach consensus or the working group would be disbanded.

The group discussed various ways to meet this deadline and agreed to pursue a phased approach. The group will focus on a MIB of core objects and defer other objects to a MIB of extensions.

It was suggested that the group reconsider the work done by the CCITT (now ITU) known as V.im (now V.58). It was agreed that V.58 would be considered a super-set of objects from which the set of core objects would be taken. Objects could be renamed and be explicitly mapped to the corresponding V.58 object with notes in the description. Changes in structure would be made only where justified.

The structure of the working group's draft of the core objects was analyzed. It was agreed that it would be reduced to five groups of objects. The objects within each group will be refined by the authors. The grouping of core objects is as follows:

mdmIDGroup	Identity of the modem
mdmLineInterfaceGroup	Configuration and state of line interface
mdmDTEInterfaceGroup	Configuration and state of DTE interface
mdmCallControlGroup	Call control configuration and state of last call
mdmStatisticsGroup	Statistics of the modem

Internet-Draft Discussion

Editor's Note: An itemized list of document changes is available via FTP or mail server from the remote directories as /ietf/modemmgt/modemmgt-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The Next Step

The five groups of core objects will be divided among the authors for further editing. Each group will be treated individually on the mailing list. The groups must be carefully evaluated and extensively discussed on the mailing list. Thanks to those who have contributed time and expertise to this project.

Attendees

Jim Barnes	barnes@xylogics.com
Jeff Case	case@cs.utk.edu
Bill Cash	cash@bangote.compaq.com
Alan Clark	aclark@hayes.com
Manuel Diaz	diaz@davidsys.com
Christine Gressley	gressley@uiuc.edu
Cheryl Krupczak	cheryl@empiretech.com
Olav Kvittem	Olav.Kvittem@uninett.no
Mark Lewis	Mark.S.Lewis@telebit.com
Peram Marimuthu	peram@wg.com
Orly Nicklass	orly@radmail.rad.co.il
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
Gregory Pearson	
Jim Rees	Jim.Rees@umich.edu
Allen Rochkind	Allen_Rochkind@3com.com
Marshall T. Rose	mrose@dbc.mtview.ca.us
Michael Scanlon	scanlon@ftp.com
Chris Shaw	cshaw@banyan.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Honda Wu	honda@nat.com

2.4.6 Remote LAN Monitoring (RMONMIB)

Charter

Chair(s)

Mike Erlinger: `mike@jarthur.claremont.edu`

Mailing Lists

General Discussion: `rmonmib@jarthur.claremont.edu`
To Subscribe: `rmonmib-request@jarthur.claremont.edu`
Archive: `jarthur.claremont.edu:/pub/rmon`

Description of Working Group

The RMON Working Group is chartered to prepare a recommendation to the IESG evaluating RFC 1271 (the RMON MIB) with respect to the standards track.

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 1271 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 1271, and be designated a draft standard (if only minor changes are made);

- that a draft prepared by the working group, should replace RFC 1271, and be designated a proposed standard (if major changes or feature enhancements are made); or,

- that RFC 1271 should be designated as historic (if this technology is problematic).

Goals and Milestones

Done	Re-activation of WG, call for discussion of experiences.
Done	Meet at IETF to classify and evaluate experiences.
Apr 1994	Submit recommendation, possibly with new draft, to IESG.

Request For Comments

RFC 1271	"Remote Network Monitoring Management Information Base"

CURRENT MEETING REPORT

Reported by Edward Alcoff/Network Application Technology and Michael Erlinger/Harvey Mudd College

Minutes of the Remote LAN Monitoring Working Group (RMONMIB)

Agenda - Monday's Session

- Presentation of new charter.
- Discussion of experiences that may affect RFC 1271 changes.
- Discussion of the four advancement options for RFC 1271.
- Consensus on the particular option to be pursued for RFC 1271.
- Discussion of areas of RFC 1271 that should be modified.

The chair presented the new charter:

> The RMON Working Group is chartered to prepare a recommendation to the IESG evaluating RFC 1271 (the RMON MIB) with respect to the standards track.
>
> The recommendation will document implementation, interoperability, and deployment experience. If this experience suggest that changes should be made to the document, a new draft may be prepared. The recommendation will report one of four outcomes:
>
> 1. That RFC 1271 should be advanced from proposed to draft status, without changes (if no problems are found);
> 2. That a draft prepared by the working group, should replace RFC 1271, and be designated a Draft Standard (if only minor changes are made);
> 3. That a draft prepared by the working group, should replace RFC 1271, and be designated a Proposed Standard (if major changes or feature enhancements are made); or,
> 4. That RFC 1271 should be designated as historic (if this technology is problematic).

After some discussion, the consensus was that a draft prepared by the working group should replace RFC 1271 and be designated a Draft Standard, with minor changes to be made. Work on version 2 of RMON was delayed until the Spring IETF, to allow RFC 1271 to progress through the standards track. The RMON mailing list would also be polled for consensus on this strategy.

Steve McRobert stated that the EtherStats group is incorrectly specified, with regards to dribble bits. Steve Waldbusser agreed and said that RMON implementors were developing RMON the way it made sense to and not the way the RFC specified. McRobert had posted several other items with regards to the EtherStats group and Waldbusser said that fixing them should be a relatively easy task. The Chair said that he would bring the information on this matter to the second session of the RMON working group for discussion.

The RMON working group has also been tasked to write up RMON interoperability issues and information with regard to RMON implementation experience. Steve Waldbusser said that he would help coordinate this effort. Bob Stewart also suggested that the working group start a new features list for consideration for the next version of RMON. The chair then solicited extensions to the RMON that have been implemented by the vendors. This request will also be passed to the RMON mailing list.

The chair then presented a list of fourteen areas for change to RFC 1271 to the meeting and the working group added three more for discussion.

Editor's Note: An itemized list of changes is available via FTP or mail server from the remote directories as /ietf/rmonmib/rmonmib-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The floor was then opened to general questions and contributions.

Thursday's Session

The Thursday meeting was initially dedicated to discussion of the AMD (Ian Crayford and Steve McRobert) concerns with the Ether Stats table.

By the time of the meeting these issues had been resolved by Steve Waldbusser and Steve McRobert. Basically, RMON implementations were doing the 'right thing', but the RFC text was unclear.

The agreed-upon changes were:

- Remove the incorrect definition of alignment errors.
- Define the term "bad packets" that is used frequently.
- Mention that the collisions object is naturally dependent on the position of the probe in the network.

One of Steve McRobert's issues that consensus could not be reached on was that the RMON's usage of the term jabbers was different than the 802.3 definition.

Two possible solutions were proposed:

1. Deprecate the current object (and object ID) and re-create another with the right name.

2. Add text to the description field that says: "Note that this is not the same as 802.3's definition of a jabber."

Consensus on this issue will be sought on the mailing list.

A broad discussion on RMON related to silicon implementation ensued. Two approaches materialized:

1. Wholesale modification of the current RMON specification, and
2. Keeping the current specification stable while acknowledging that RMON II will seriously consider hardware implementation issues, and therefore may not remain compatible with the current RMON. The working group agreed to the second strategy.

One particular concern that was discussed for silicon implementations is that no performance gains can be achieved for filtering when the acceptType is set to acceptFailed. After some discussion it was identified as a general problem with formulas in "Sum of Products" form, and that outlawing them is probably not the right solution given that these are useful for a variety of filtering applications. The suggestion was made that RMON applications could warn the user that the SOP form selected when setting acceptType to acceptFailed can be very inefficient.

Attendees

Edward Alcoff	oldera@nat.com
Jim Barnes	barnes@xylogics.com
Bart Berger	bart_berger@3com.com
Ram Bhide	ram@nat.com
Andy Bierman	abierman@synoptics.com
Jia-bing Cheng	cheng@ralvm6.vnet.ibm.com
Chris Chiotasso	chris@lightstream.com
Frank Ciotti	frankc@telxon.com
Blair Copland	copland@unt.edu
Manuel Diaz	diaz@davidsys.com
Jonathan Didner	jonb@bangate.compaq.com
David Engel	david@ods.com
Michael Erlinger	mike@jarthur.claremont.edu
William Fardy	billf@frontier.com
Steve Garritano	steveg@kalpana.com
Christine Gressley	gressley@uiuc.edu

Robert Grow	bob@xlnt.com
Stuart Hale	stu_hale@vnet.ibm.com
Daniel Hansen	danh@ngc.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
Kevin Jackson	kjackson@concord.com
Mark Kepke	mak@fc.hp.com
Kenneth Key	key@snmp.com
Michael Kornegay	mlk@bir.com
Cheryl Krupczak	cheryl@empiretech.com
William Kwan	kwan@rabbit.com
Dave Livingston	squirrel@vnet.net
Carl Madison	carl_madison@3mail.3com.com
Peram Marimuthu	peram@wg.com
Evan McGinnis	bem@3com.com
Steve McRobert	steve.mcrobert@amd.com
Tom Nisbet	nisbet@fbsw.tt.com
Steven Onishi	sonishi@wellfleet.com
David Perkins	dperkins@synoptics.com
Jon Saperia	saperia@zko.dec.com
Michael Scanlon	scanlon@ftp.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@timonware.com
Bob Stewart	rlstewart@eng.xyplex.com
Adam Stolinski	stolinsk@cerf.net
Kaj Tesink	kaj@cc.bellcore.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Thomas Walsh	tomw@kalpana.com
Alice Wang	alice.wang@eng.sun.com
Peter Wilson	peter_wilson@3mail.3com.com
Paul Woodruff	pwoodruff@synoptics.com
Henry Yip	natadm!henry@uunet.uu.net

2.4.7 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, 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

Done	Mailing List discussion of vendor proprietary MIBs.
Done	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.

Internet-Drafts

"Definitions of Managed Objects for SNA Data Link Control: SDLC", 01/03/1994, J. Hilgeman, S. Nix, A. Bartkey <draft-ietf-snadlc-sdlc-mib-01.txt>

CURRENT MEETING REPORT

Reported by Jeff Hilgeman/Apertus Technologies

Minutes of the SNA DLC Services MIB Working Group (SNADLC)

The SNADLC Working Group met at the 28th IETF to continue working on the SDLC MIB. There were twelve people in attendance.

The majority of the time was spent going through the MIB objects and further clarifying their type, range, description, etc. Much of the effort during this working session was in the area of consistency check. There will be a follow up posting to snadlcmib detailing the changes in the MIB.

SDLC MIB

The group has focused exclusively on SDLC for the present time. Other DLCs that were dropped at the formation of the working group were discussed (e.g., QLLC and channel).

The group had originally determined that it would not make the December 1993 due date for delivery of the SDLC MIB and that it would shoot for early 1994. This was unacceptable to the Network Management Area Director, Marshall Rose. The group must deliver the final version of the Internet-Draft in December. This fact, coupled with the fact there are no meetings currently planned between now and December, means the group has to use the network to get the work done. The MIB editor has agreed to have a new version of the SNA SDLC MIB every Thursday starting next week.

They will be submitted to the Internet-Draft administrator at CNRI for file name assignment and distribution to the Internet-Draft directories but they will also be placed onto the cisco FTP server each Thursday (probably in the evening, west coast time) for speed of delivery to working group members.

In order to have a new revision of the Internet-Draft each Thursday evening, all proposed changes to the draft must be submitted to the mailing list by 4:00 p.m. PST Thursday. If there are no changes to the SDLC MIB for a particular week, a notice to that effect will be posted to snadlcmib. Regardless of whether there is a new revision of the draft available for a given week, the latest revision can always be found at `ftp.cisco.com` as draft-ietf-snadlc-sdlc-mib.txt. If you have trouble accessing this file, send e-mail directly to `wclark@cisco.com`.

LLC2 MIB

The group had no plans to look at the LLC2 MIB. There was some discussion of deferring this to a future working group, though Ken Key proposed that the existing SNADLC Working Group develop a new charter and a new timeline specifically for LLC2. In other words, a new working group is not necessary but a redirection of the existing group would do.

MIB Architecture

The SDLC MIB sits beneath the SNA NAU Services MIB and is referenced by it. There is a tight coupling between the SDLC MIB and RS-232 MIB since many of the port attributes for SDLC are addressed by RS-232. Extensions to the RS-232 MIB will be proposed for attributes that are unique to the SDLC environment (e.g., NRZI, controlled RTS, etc.). (This will be the subject of a follow on posting to the snadlcmib mailing list.)

There are two managed entities in the MIB Architecture:

1. Physical ports
2. Logical link stations

There are three management tables:

1. admin table
2. oper table
3. stats table

The primary and secondary links are supported.

It was mentioned that there are five tables rather than the expected six. This is because port statistics are covered in the RS-232 MIB.

Proposed Changes to the SDLC MIB Internet-Draft

Work on the MIB during the meeting included:

- Error code definitions
- Include branch number
- Conformance statements
- Others
  - Statement in the MIB regarding the subject of "rationale" behind some of the objects. This is so that a network administrator can actually use the MIB to perform network management operations (what a concept!).

- Discussion of the need for dynamic row creation for ports and link stations.
- Conformance of the SDLC MIB's SMI to the syntax of SNMPv2 rather than SNMPv1

MIB Changes

This topic will be covered in detail in a follow on posting to the snadlcmib mailing list.

Editor's Note: Highlights of the proposed changes to the MIB are available via FTP or mail server from the remote directories as /ietf/snadlc/snadlc-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Changes to the MIB resulting from this IETF meeting will be available by Monday, November 8.

Action Items

Several action items resulted from the meeting.

Wayne Clark	Investigate the use of SNMPv2's reference clauses within the SDLC MIB. The latest FDDI MIB (RFC 1512) to be used as an example.
	Check with Alan Bartky of Sync Research on the need for sdlcLSAdminXid.
	E-mail the list of RS-232 requirements to snadlcmib for review. This is in preparation for making the list available to the RS-232 MIB working group
	Submit changes to draft-ietf-snadlc-sdlc-mib.txt by Monday, November 8. .
Ken Key	Investigate the applicability of row creation for SDLC link stations and make a recommendation to the working group.
Jeff Hilgeman	Query snadlcmib to see if the timestamps in the MIB are useful. (Objects in question are sdlcLSOperCreateTime, etc.)
Shannon Nix	Investigate whether there is any overlap between the timestamps and those in the ifTable.
	Email recommendations on MIB object name changes so as not to clash with NCP-defined names.

Attendees

Michael Allen	moallen@ralvm6.vnet.ibm.com
Jia-bing Cheng	cheng@ralvm6.vnet.ibm.com
Wayne Clark	wclark@cisco.com
Cleve Graves	cvg@oc.com
Jeff Hilgeman	jeffh@apertus.com
Kenneth Key	key@snmp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
William Kwan	kwan@rabbit.com
Shannon Nix	sdn@netlink.com
Jon Penner	jjp@bscs.uucp
Barbara Sterling	bjs@mcdata.com
Mauro Zallocco	mdz@netlink.com

SNA DLC Services MIB - Working Group

IETF 28 Meeting of the SNA DLC WG
Tuesday, November 2, 1993.

Agenda

Introductions

Working Group Background

SDLC MIB Review

Indentify Open Issues

Discuss Next Steps

28th IETF Conference, Houston, Nov. 2, 1993

SNA DLC Services MIB - Working Group

Working Group Background

o Original Charter

o Modified Goals

o Current Status

28th IETF Conference, Houston, Nov. 2, 1993

SNA DLC Services MIB - Working Group

Working Group Background

Modified Goals

o Will not make 12/93 deliver of SDLC MIB
(shooting for early '94)

o No plans to look at LLC2 MIB.
(deferred to future WG)

28th IETF Conference, Houston, Nov. 2, 1993

SNA DLC Services MIB - Working Group

SDLC MIB Review

o MIB Architecture

o Proposed changes to I-D

o Detailed review of I-D
(refer to MIB handout)

28th IETF Conference, Houston, Nov. 2, 1993

SNA DLC Services MIB - Working Group

MIB Architecture

- o Two managed entities
  - a. Physical ports
  - b. Logical link stations
- o Three management tables
  - a. Admin table
  - b. Oper table
  - c. Stats table
- o Primary and secondary links supported.

28th IETF Conference, Houston, Nov. 2, 1993

SNA DLC Services MIB - Working Group

Identify Open Issues

28th IETF Conference, Houston, Nov. 2, 1993

2.4.8 SNA NAU Services MIB (SNANAU)

Charter

Chair(s)

Zbigniew Kielczewski: zbig@eicon.qc.ca
Deirdre Kostick: dck2@mail.bellcore.com

Mailing Lists

General Discussion: snanaumib@thumper.bellcore.com
To Subscribe: snanaumib-request@thumper.bellcore.com
Archive: thumper.bellcore.com:pub/tob/snanaumib

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.
Done	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.

Internet-Drafts

"Definitions of Managed Objects for SNA NAUs", 12/23/1993, Z. Kielczewski, D. Kostick, K. Shih <draft-ietf-snanau-snamib-02.txt>

CURRENT MEETING REPORT

Reported by Deirdre Kostick/Bellcore

Minutes of the SNA NAU Services MIB Working Group (SNANAU)

The SNANAU Working Group met on November 3, 1993 to review the revised SNA NAU MIB Draft.

New State Model for SNA Nodes

The working group identified a "state" model for SNA Nodes to clarify how objects proposed in the snaNodeAdminTable and snaNodeOperTable would be used to control SNA nodes. Dr. SNMP (Jeff Case) gave some greatly appreciated guidance. Diagrams 1 and 2 (attached at the end of the notes) summarize the new model and usage of RowStatus, OperStatus and AdminStatus.

The current draft of the SNA NAU MIB shall be updated to reflect this model.

Other MIB Revisions

Editor's Note: A detailed list of changes to the MIB is available via FTP or mail server from the remote directories as /ietf/snanau/snanau-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Follow-up Review

The working group still needs to complete the detailed review of the remaining tables. Definitions need to be cleaned-up. During the meeting, the working group did not have time to review comments on the Session, Link, LU tables, etc.

Working Group members are encouraged to make comments on the mailing list preferably with detailed additions, wording, changes, etc.

A conference call will be scheduled (November 8) to continue the MIB review. Results of the call will be posted to the mailing list.

The group has a December deadline to complete the MIB work. There is still *a lot* of work to do to make this date.

Conformance/Compliance Statements

Zbigniew Kielczewski handed out a proposed first cut at conformance/compliance statements. These will be updated and posted to the list.

Future SNA MIB Work

Additional SNA MIB work in 1994 may be approved by the Network Management Area Director, *if* and *only if*, the current SNA NAU MIB is completed on schedule—by December, 1993.

The candidate work item for 1994 is to develop APPC-related management objects.

Work to identify APPN-related management objects could be pursued in parallel, for example, by the AIW.

Next Version

The next version of the MIB will be posted to the mailing list by November 12. Zbigniew, Kitty, Robin, and Deirdre volunteered to assist with editing.

Diagram 1: Node State Model

```
         ----------------
         |              |
         | DOES NOT     |
         | EXIST        |
         |              |
         ---------------
                | ^
        create  | |delete
                V |
         ----------------
         |              |
         | INACTIVE     |
         |              |<---------------|
         |              |                |
         ---------------                 |
                | ^                      |
    activate    | |                      |
                V |                      |
         ----------------                |
         |              |                |
 --------| WAITING      |                |
|        |              |                |
|        |              |                |
|        ---------------                 |
|               | ^                      |
|               | |                      |
|               V |                      |
|        ----------------                |
|        |              |                |
|        | ACTIVE       |                |
|        |              |                |
|        |              |                |
|        ---------------                 |
|               | graceful stop or       |
|               | stop force             |
|               V                        |
|        ----------------                |
|        |              |                |
-------->| STOPPING     |----------------|
         |              |
         |              |
         ---------------
```

Diagram 2: Using the Proposed snaNodeAdmin and snaNodeOper Tables

Admin Table Objects		Oper Table Objects
RowStatus	AdminStatus	OperStatus
(no row entry)	N/A	N/A
To create a row representing a node instance, set RowStatus = active	inactive	agent automatically creates a corresponding entry in the snaNodeOperTable with OperStatus=inactive
active	To activate node, set AdminStatus = active	agent changes node's OperStatus= waiting
active	active	node changes to "active"
active	to stop node, set Shutdown Method to appropriate value & AdminStatus = inactive	stopping or inactive depending on Shutdown Method
invalid see note#1	inactive	inactive

Notes:

1. To delete a row, OperStatus and AdminStatus must both be = inactive. Row deletion can be NMS or agent initiated:
   (a) The NMS sets RowStatus=invalid, or
   (b) The agent detects that a row is in the "under creation" state for greater that some default period, e.g., 5mins
   (c) After reboot, the row with RowStatus=inactive will not be included in the table

2. A new object, shutdownMethod, is needed. shutdownMethod should be a R/W object with the following values:
   (a) Graceful
   (b) Forced
   (c) Rol - request on line?
   (d) Other

   The value of shutdownMethod is only valid for the instance that the "inactive" button is pushed.

Attendees

Michael Allen	moallen@ralvm6.vnet.ibm.com
Jeff Case	case@cs.utk.edu
Cleve Graves	cvg@oc.com
Jeff Hilgeman	jeffh@apertus.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Deirdre Kostick	dck2@mail.bellcore.com
William Kwan	kwan@rabbit.com
Shannon Nix	sdn@netlink.com
David Perkins	dperkins@synoptics.com
Eddie Renoux	elr0262@newsit2.mcdata.com
Kitty Shih	kmshih@novell.com
Mauro Zallocco	mdz@netlink.com

2.4.9 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) controlling 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.
Done	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.

Internet-Drafts

"UPS Management Information Base", 12/21/1993, J. Case <draft-ietf-upsmib-03.txt>

CURRENT MEETING REPORT

Reported by Jeff Case/University of Tennessee

Minutes of the Uninterruptible Power Supply Working Group (UPSMIB)

The UPS MIB Working Group met on Friday, November 5th, 1993, at the 28th Meeting of the Internet Engineering Task Force, in Houston, Texas, USA.

The group conducted an intense review and re-write of the "UPS Management Information Base" Internet-Draft (draft-ietf-upsmib-01.txt). Nearly every object changed in one way or another in the marathon meeting which began at 9 a.m., and concluded about 5 p.m., as scheduled and announced via the mailing list.

A new document reflecting the output of the meeting was posted to the mailing list soon after the meeting. After further discussion and revision on the mailing list, it was submitted as a new Internet-Draft, superceding draft-ietf-upsmib-01.txt.

The consensus in the Houston meeting was strong, and the work of the group is nearly finished. It is almost certain that the Houston meeting will be the working group's last meeting before it's charter is completed and the working group will go dormant until the document is up for revision or new work items are chartered.

Attendees

Tom Brennan	brennan@exide.com
Jonathan Didner	jonb@bangate.compaq.com
Roger Draper	rdraper@cerf.net
Bill Elliot	70760.3022@compuserve.com
Tim Grams	73537.2601@compuserve.com
Brian Hammill	brianh@exide.com
Michael Kornegay	mlk@bir.com
Lawren Markle	72170.460@compuserve.com
Gordon Monk	ghanka@metronet.com
Ronald Pitt	57313.577@compuserve.com
Jean-Bernard Schmitt	jbs@vnet.ibm.com
Jack Stiekema	jack@victron.nl
Adam Stolinski	stolinsk@cerf.net
Greg Wilterdink	71302.3624@compuserve.com
Brian Young	byoung@macc.wisc.edu
Terry Zumwalt	71151.2550@compuserve.com

2.5 Operational Requirements Area

Director

- Scott Bradner: sob@harvard.edu

Area Summary reported by Scott Bradner/Harvard University

Remote Printing on Global Facsimile Devices BOF (TPCINT)

Technical, operational, and financial basis for operation of remote printing services in the `tpc.int` subdomain were discussed. Brief discussion of financial models of operation ensued.

BGP Deployment and Application Working Group (BGPDEPL) and CIDR Deployment BOF (CIDRD)

Discussion topics:

- Virtual CIDR test network
- Software implementation and interoperability test check points
- CIDR deployment plan check point and coordination
- Aggregation registry
- Block assignment guidelines
- Class A usage with CIDR
- Variable-length subnets usage and guidelines
- Renumbering

The Internet wide CIDR/BGP4 test network for network operators was described. Some testing aggregate routes are being exchanged on the test network. Network operators are encouraged to join and perform their tests. There was discussion of aggregate registry which is proposed to register relevant information about aggregate routes. Presentations of policies and procedure of (sub)allocating CIDR block addresses were given by NIC, RIPE and a regional (Barrnet). These people will work on a document which will provide guidelines to such task.

The mechanisms of how to use the remaining IP address space more efficiently with CIDR were also explored. These include Class A space usage with CIDR, variable-length subnets usage and renumbering. The general thinking is that well-documented procedures and methods are the key. Tools help a great deal. Documents of using Variable-length subnets and renumbering procedure will be produced by the group.

Benchmarking Methodology Working Group (BMWG)

Scott Bradner opened the meeting stating that Jim McQuaid will be the new BMWG chair. Bradner gave an overview of the working group activities and briefly reviewed the content of the two working documents. He mentioned that both documents are available from `hsdndev.harvard.edu` in the `pub/bmwg` directory.

The subject of adding modems to the document and methods for testing bridges and routers was brought up. It was thought that adding modems would be straightforward.

The group then discussed the charter at some length, discussing aspects of the work which was done and of the work which needed to be addressed in the future. The revised "goals and milestones" section of the charter will reflect the future work of the working group.

A survey of research and articles on benchmarking methodologies will be conducted and reported at the next meeting.

Generic Internet Service Description Working Group (GISD)

After discussion, the general consensus appears to be that the work done at the two previous BOFs (then known as GISS) was basically correct. No additional aspects were asked for except for the possible addition of a description of training.

The main action was for an "Aspects Guideline" document to be produced giving detail of how to submit aspects and also containing an index of the aspects that still need to be drafted.

Network Joint Management Working Group (NJM) and Network Status Reports (NETSTAT)

Status reports were received from:

- Bill Manning about Sesquinet
- Michael Patton about the Defense Simulation Internet (DSI)
- Jeff Burgan about the NASA Science Internet
- Jordan Becker about ANS
- Mike O'Dell about Alternet
- Scott Bradner about CoREN
- Bernhard Stockman about the Ebone

CURRENT MEETING REPORT

Minutes of the CIDR Deployment BOF (CIDRD)

The minutes of the joint BGPDEPL/CIDRD session follow the BGPDEPL charter. The attendee list below is from the joint session.

Attendees

Susie Armstrong	susie@mentat.com
Jules Aronson	aronson@nlm.nih.gov
William Barns	barns@gateway.mitre.org
Tony Bates	tony@ripe.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Jim Bound	bound@zk3.dec.com
Rich Bowen	rkb@ralvm11.vnet.ibm.com
Scott Bradner	sob@harvard.edu
Al Broscius	broscius@bellcore.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Enke Chen	enke@merit.edu
Henry Clark	henryc@oar.net
Michael Collins	collins@es.net
Rob Coltun	rcoltun@ni.umd.edu
John Curran	jcurran@nic.near.net
Michael Davis	mike@dss.com
Taso Devetzis	devetzis@bellcore.com
Christopher Dorsey	dorsey@es.net
Tom Easterday	tom@cic.net
Havard Eidnes	havard.eidnes@runit.sintef.no
Robert Enger	enger@seka.reston.ans.net
Stefan Fassbender	stf@easi.net
Dennis Ferguson	dennis@ans.net
Robert Fink	rlfink@lbl.gov
Dale Finkelson	dmf@westie.mid.net
Peter Ford	peter@goshawk.lanl.gov
Catherine Foulston	cathyf@rice.edu
Vince Fuller	vaf@barrnet.net
Vincent Gebes	vgebes@sys.attjens.co.jp
Herluf Hansen	hha@tbit.dk
Denise Heagerty	denise@dxcoms.cern.ch
Matt Hood	hood@nsipo.nasa.gov
David Jacobson	dnjake@vnet.ibm.com
Dale Johnson	dsj@merit.edu

Matthew Jonson	jonson@ddn.af.mil
Janine Kamerdze	kamerdze@nsipo.nasa.gov
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Sean Kennedy	liam@nic.near.net
Charley Kline	cvk@uiuc.edu
Mark Kosters	markk@internic.net
John Krawczyk	jkrawczy@wellfleet.com
Walter Lazear	lazear@gateway.mitre.org
Tony Li	tli@cisco.com
Robin Littlefield	robin@wellfleet.com
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se
Bill Manning	bmanning@rice.edu
Glenn Mansfield	glenn@aic.co.jp
Jun Matsukata	jm@eng.isas.ac.jp
Stephen Miller	smiller@bbn.com
Pushpendra Mohta	pushp@cerf.net
Dennis Morris	morris@altair.disa.mil
Jun Murai	jun@wide.ad.jp
Michael O'Dell	mo@uunet.uu.net
Vijayaragavan Pandian	vjp@proteon.com
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@eng.umd.edu
Michael Patton	map@bbn.com
David Piscitello	wk04464@worldlink.com
Kenneth Rehbehn	kjr@netrix.com
Yakov Rekhter	yakov@watson.ibm.com
Isil Sebuktekin	isil@nevin.bellcore.com
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Timon Sloane	timon@timonware.com
Frank Solensky	solensky@ftp.com
Bernhard Stockman	boss@ebone.net
Larry Tepper	ltepper@compatible.com
Marten Terpstra	marten@ripe.net
Claudio Topolcic	topolcic@cnri.reston.va.us
Jerry Toporek	jt@mentat.com
Paul Traina	pst@cisco.com
Keisuke Uehara	kei@cs.uec.ac.jp
William Warner	warner@ohio.gov
Chris Wheeler	cwheeler@cac.washington.edu
Jane Wojcik	jwojcik@bbn.com
David Woodgate	David.Woodgate@its.csiro.au
Jessica Yu	jyy@merit.edu

CURRENT MEETING REPORT

Reported by Carl Malamud/Internet Multicasting Service

Minutes of the Remote Printing on Global Facsimile Devices BOF (TPCINT)

The BOF on "Remote Printing on Global Facsimile Devices" was sucessfully conducted at the 28th IETF with no subpoenas or other disruptions. We discussed the technical, operational, and financial basis for operation of remote printing services in the `tpc.int` subdomain.

A brief discussion of financial models of operation ensued and the BOF was concluded. There seems to be no need to repeat the BOF with the current material, but as services in the `tpc.int` subdomain progress further discussion may be warranted.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Scott Bradner	sob@harvard.edu
Tom Easterday	tom@cic.net
Roger Fajman	raf@cu.nih.gov
Qin Fang	qin_fang@unc.edu
Jari Hamalainen	jah@rctre.nokia.com
Timo Jokiaho	timo.jokiaho@ntc.nokia.com
Phil Karn	karn@qualcomm.com
Brendan Kehoe	brendan@zen.org
David Kristol	dmk@allegra.att.com
Carl Malamud	carl@malamud.com
Wayne McDilda	wayne@dir.texas.gov
Mark Needleman	mhn@stubbs.ucop.edu
Bill Norton	wbn@merit.edu
Michael O'Dell	mo@uunet.uu.net
Marshall T. Rose	mrose@dbc.mtview.ca.us
Henning Schulzrinne	hgs@research.att.com
John Stewart	jstewart@cnri.reston.va.us
Akihiro Tominaga	tomy@sfc.wide.ad.jp

Outreach and Integration

A Public Service of the TPC.INT Subdomain

Carl Malamud (carl@radio.com)

Marshall T. Rose (mrose@dbc.mtview.ca.us)

The Basic Problem

- We Prefer General-Purpose Devices
- Still Need for Special-Purpose Devices
- Fax, TDD, Pager, Telephone, Voicemail
- Question: Can We Integrate Them?

Case 1: Fax

- Communicate With Person Who Doesn't Use Email
- Need Local PostScript Printer (e.g., Hotel)
- Want Single Interface for Email/Fax
- Want to Make Cost-Effective Use of Telecommunications

The DNS Hack

- TPC.INT Subdomain
- Full IDDD Phone Number, Reversed

  +1 (415) 968-2510

  0.1.5.2.8.6.9.5.1.4.1.tpc.int
- MX Record To Server

  *.5.1.4.1.tpc.int MX 10 dbc.mtview.ca.us.

More DNS Hacks

- BIND Hack for Multiple Servers
- DNS Enforces Fair Competition
- Sinkhole for Uncovered Areas

  *.6.1.4.1.tpc.int MX 10 sinkhole.town.hall.org.
- Friendly Message Back to User
- Log Number, Weekly Message to List

Services

- Service 1: Remote Printing
- Local Part: Identifies the Service
- Two Variants: MIME-Enabled and MIME-Disabled
- Message Contents

MIME-Enabled

- Explicit Cover Sheet (application/remote-printing)

  remote-printer@0.1.5.2.8.6.9.5.1.4.1.tpc.int
- MIME Body Part Contains Cover Sheet Information

MIME-Disabled

- Embed Cover Sheet Inside Local Part

  remote-printer.ATOM@0.1.5.2.8.6.9.5.1.4.1.tpc.int
- ATOM: RFC 822 Defined
- Best Bet: [A-Za-z]
- Special Characters:

  "/" = [CR], "_" = [" "]

Message Contents

- MIME-Disabled: ASCII Text
- MIME-Enabled
  - text/plain
  - message/rfc822
  - application/postscript
  - image/tiff
  - multipart

Remote Printer Operators

- Denial of Access
- Privacy
- Scope of Coverage
- Model of Operation

Denial of Access

- Deny Based on Sender, Not Recipient
- Control Over Recipient is in DNS
- Typical Policy 0:
  - x times or y minutes in z time

Privacy

- Interior of Message is *PRIVATE!!*
- Limited Log for Audit/Debug
- No Disclosure That Identifies Individuals
- No Mailing Lists

Scope of Coverage

- Highly Dynamic
- Personal, Organizational
- Can Get Exclusive Access to Own Area
- Non-exclusive Access for Neighborhoods, Regions

Models of Operation

- Does This Cost Money?
- How Do I Pay for It?
- Several Models Can Cooperate
- Local Library Model
- Community Newspaper Model
- Corner Grocery Model

Local Library Model

- Your Marginal Cost Low
- Frequency of Use Anticipated To Be Low
- Provide Community and/or Organizational Service
- Ex: Univ. of Michigan

Community Newspaper Model

- Cover Area as a Profit (Cost) Center
- Ads on Acks, Cover Page
- 1/3 of Cover Page
- 250 Bytes on Ack

How it Works

- Accounts, Contact File
- Tariff Table
- Ad Directory
- Can Assign Ads by Telco Prefix, DNS Suffix
- Logs, Reports Automatic

Corner Grocery Model

- Policy 0: Service to Network
- Policy 1: Increased Service
- Out-of-Band Formation of Calling Circles
- Loose Authentication, Control
- Allow People to "Fax Home" From Hotel

Software

- Flexfax, DBC Glue
- PMDF
- IsoFax, DBC Glue
- NeXT, Other Platforms

Is This Legal?

- Yes.
- Really.
- Falls Within Scope of CCITT D.1
- Not Bypass: “Smartpass”

Why?

- More Cost-Effective
- Emphasize Importance of General-Purpose Infrastructure
- Demonstrate Model of Small, Dynamic Entrepreneurs
- Great Research Problem

Coverage in Parts of ...

- Australia, New Zealand
- U.S.
- Netherlands, Denmark

Anticipated Coverage

- Nordic Countries, Ireland
- More US, Australia, New Zealand
- Tokyo as WIDE Experiment
- Many Others Investigating

Needed Coverage

- Universities
- Corporations
- IP Service Providers
- Can Just Serve Your IP Customers for Start

Administration of Namespace

- Cooperative of Service Providers
- Administered as Public Trust
- Must Run Service to be Enfranchised

TPC.INT Board

- Works on Case Law Basis
- Only Solve Real Problems, Not Hypothetical
- Has Never Met, May Never Meet
- Selected by Members of the Cooperative

Current Board

- Rob Blokzijl, NIKHEF
- Geoff Huston, AARNET
- Carl Malamud, Internet Multicasting Service
- Jun Murai, WIDE
- Marshall T. Rose, Dover Beach Consulting

Next Steps

- Authenticated Message Exchange?
- Options Market for Ads?
- TDD? Pager? Voicemail?
- Need Coders, Not Goers
- Plenty of Ideas: We Need Makefiles

Important Points

- Think Small: Many Small Cells is the Key
- DNS is Dynamic: Can Always Contract Area
- Load on Machines is Negligible
- All You Need is 1 Modem, Small Chunk of a Machine
- This is not Rocket Science!

For Further Reading

- RFC 1530 - TPC.INT General Principles
- RFC 1528 - Remote Printing Technical
- RFC 1529 - Remote Printing Administrative
- FAQ

Mailing Lists, Mail Addresses

- FAQ: tpc-faq@town.hall.org
- Current Coverage: tpc-coverage@town.hall.org
- Mailing List: tpc-rp@aarnet.edu.au
- Administration: tpc-admin@town.hall.org
- Board: tpc-policy@town.hall.org

2.5.1 BGP Deployment and Application (BGPDEPL)

Charter

Chair(s)

Jessica Yu: jyy@merit.edu

Mailing Lists

General Discussion: bgpd@merit.edu
To Subscribe: bgpd-request@merit.edu
Archive: 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.

Request For Comments

RFC 1482	"Aggregation Support in the NSFNET Policy Routing Database"

CURRENT MEETING REPORT

Reported by Peter Ford/LANL

Minutes of the Joint Session of the BGPDEPL Working Group and CIDRD BOF

The BGP/CIDR deployment meeting was held on November 3, 1993 and was chaired by Jessica Yu and Vince Fuller. The first order of business was a brief status report on BGP-4 implementations:

- ANS (Guy Almes): The ANS test network mid-November. Deployment in the ANS production network during December.

- cisco (Paul Traina): In beta. Get the image from `ftp.cisco.com`. Please join the beta list; mail to `pst@cisco.com`.

- Wellfleet (John Krawczyk): Full product in 8.0 by Spring/Summer 1994 test version in by January/February.

- 3Com (Tracy Mallory): Beta available.

- BBN: Under development.

- Europanet testing in progress, deployment by end of the year.

- Rainbow Bridge (Rob Coltun): Status?

Peter Lothberg and Andrew Partan reported on their BGP-4 test network. It is a virtual test network which is accessible to anyone who wished to participate. The current players are predominantly cisco-based, and many use cisco GRE tunnels to obtain connectivity with the test network. There are currently 15K IP networks sloshing around.

The following participants are on the test network and ANS is expected to be soon.

- 3Com
- Alternet
- cisco
- Ebone
- ESnet
- ICMnet
- IIJ
- NEARnet
- RIPE NCC
- STUPI

Peter Lothberg reports that he has converted the EBONE over to using BGP-4 (nine routers). Static aggregates have been injected into EBONE from the regionals and passed over to other regionals and the ICM. He also reported that the ICM system has also been cut over to BGP-4 (four routers). Andrew reported a similar cut over of the Alternet routing system.

It was noted that the current BGP-4 code is beta code and one has to carefully test their current configuration and operation prior to deploying this code in full operation. Join the beta list at cisco for more details.

Andrew reported that Alternet uses IGRP within their system and Peter Lothberg reports that he is using IS-IS.

Merit Routing Registry - Dale Johnson

From the network operators' point of view, there is a need to be able to validate the aggregate routes received via CIDR. One approach is to register aggregate routes in a database with its creator AS information and the contact information of the AS could be obtained from various existing databases such as the InterNIC, RIPE and Merit. Merit, RIPE and the InterNIC will work together on this.

Another approach mentioned at the meeting is to use BGP's AGGREGATOR field to carry such information.

Guidelines for Block Assignment

Marten Terpstra gave a presentation on Guidelines for block assignment. The RIPE people have worked with the European network community to build a distributed operational model for Internet Registries (IRs). They currently work with a model of: Global IR, Regional IR, Local IR. The RIPE NCC allocates addresses to Local IRs based on the following guidelines:

- Get two year estimate of address usage. Make sure blocks are CIDRable. CIDR enforced to end sites. Can claim unused reserve block.

- The RIPE NCC has been assigning the CIDR way, since June 1992. The procedures are documented in RIPE 72 which can be obtained via anonymous FTP or Gopher from `ns.ripe.net`.

- The RIPE DNS scheme for 193.in-addr.arpa is to try to delegate zones to providers. RIPE 84 documents policy (customer shifts, preserve mapping, etc.).

- Dennis Ferguson asked about the current efficiency of use of IP address space from the European side. It is estimated to be 3.8%.

- RIPE only allocates for Europe and would like to advocate that regional registries do the same.
- Marten noted that assignments do not really count as allocated unless they submit detailed information to the network, ensuring that the RIPE NCC is kept up-to-date.

BARRNet Allocation of Addresses

Jessica asked Vince to report on how Barrnet allocates addresses. Vince indicated that this is a manpower intensive process. They sit down with customers to get an estimate for two years out, then they work on a subnetting scheme and do a crystal ball gaze.

InterNIC Allocation of Addresses

Mark Kosters of the InterNIC reviewed how they allocate addresses, which was similar in spirit to the Barrnet and the RIPE NCC. Tony Bates, Mark Kosters and Vince volunteered to write an allocation policy document which can be used as guidance for providers.

The InterNIC will be doing block in-addrs.

When the InterNIC allocates a CIDR block to a provider, assignments are requested to be sent back to the InterNIC who will register the assignments to the InterNIC database. The Shared WHOIS Project (SWIP) is working on getting many of the InterNIC-related electronic updates done in a timely manner. The RIPE NCC does all their updates electronically and are happy with the results to date.

There were several suggestions from the working group to the InterNIC with regard to policy.

It is felt that the InterNIC needs to do more preallocation of blocks for IRs that are to be delegated.

Marten suggested the InterNIC look into swapping CIDRable Cs for Bs that are already allocated, but not heavily utilized.

"Greening of the Internet"

Vince led a discussion on "Greening of the Internet."

The outline of the talk was:

- Class A usage with CIDR
- Subnets and CIDR
- Renumbering issues and tools
- Politics of CIDR—block sizes, provider responsibility
- CIDR Analysis discussion and question and answer

Class A Usage with CIDR has a few small problems:

- There is a DNS issue, which relates to how providers delegate the in-addr namespace.
- Dumb multihomed host problem. Using older BSD systems as routers. It is felt that if a system can not handle variable length subnet masks they are obsolete.

VLSM needs to be better documented, and their use, made simpler by better tools and education. It is observed that most sites simply use 8-bit subnets since they are the easiest thing to read, use, etc.

Charley Kline who is in charge of networks at UIUC, described how they allocate subnets. Following methods described in RFC 1219, and implementing software to help administer the address space, UIUC has been able to manage 13K hosts and 324 subnets in a single Class B network. Charley illustrated the methodology using binary trees.

Tom Easterday and Charley Kline volunteered to work on a document describing the use of VLSM and better utilization of subnets in a single address block. Havard Eidnes agreed to help and offered the use of his INET '93 paper as a starting point. Havard's paper is available in the INET '93 proceedings via anonymous FTP and Gopher from `cnri.reston.va.us`.

There was a unanimous vote in favor of using the IP addr/len syntax for representing prefixes.

Scott Bradner stated that it is important for the ALE and CIDRD Working Groups to establish goals and objectives for address space usage.

Vince presented a "pain and anguish" slide which went into the issue of renumbering sites. The discussion focused on better uses of the already allocated Class A network addresses. Vince pointed out that CIDR does not require one to renumber when a site leaves one provider for another, but to maintain a minimal state of routing information is a good idea. To support this activity we need tools and documentation on renumbering. The use of DHCP can help reduce the effort in converting from one IP address block to another.

Barrnet has helped to renumber several sites and has used the following procedure:

- Add new DNS NS addresses at the InterNIC.
- Add new addresses on the primary server, wait for propagation.
- Reconfigure the network.
- Delete old addresses from the primary server.
- Delete old NS addresses at the InterNIC.

Barrnet customers have not had problems with renumbering, provided the customers are given a good set of instructions. Transitions must be gradual if they are to work. Secondary addresses are needed to facilitate transition, and most routers support this.

Yakov Rekhter volunteered to discuss dynamic updates of DNS with the Domain Name Systems Working Group (DNS).

Several other ideas to conserve address space came up during discussion including ARP being changed to be like ES-IS for IP, dynamic prefix updating, etc.

Tony Li put up a prototype charter of the proposed ALE Working Group. The primary purpose is to watch utilization numbers.

Andrew Partan used Alternet data to show how much CIDR can buy you. The number of networks from AS701 shrinks from 2100 to 650 today.

There was a brief report on how big an Internet can be routed today:

- 28-29K routes in a 16 MB cisco.
- 25K routes in the ANS routers.

Tony Bates volunteered to help continue the monitoring mode. Tony Bates, Marten, David Conrad and Vince will document how to better use address space within sites.

Attendees

Susie Armstrong	susie@mentat.com
Jules Aronson	aronson@nlm.nih.gov
William Barns	barns@gateway.mitre.org
Tony Bates	tony@ripe.net
Erik-Jan Bos	erik-jan.bos@surfnet.nl
Rebecca Bostwick	bostwick@es.net
Jim Bound	bound@zk3.dec.com
Rich Bowen	rkb@ralvm11.vnet.ibm.com
Scott Bradner	sob@harvard.edu
Al Broscius	broscius@bellcore.com
Jeffrey Burgan	jeff@nsipo.nasa.gov

Enke Chen	enke@merit.edu
Henry Clark	henryc@oar.net
Michael Collins	collins@es.net
Rob Coltun	rcoltun@ni.umd.edu
John Curran	jcurran@nic.near.net
Michael Davis	mike@dss.com
Taso Devetzis	devetzis@bellcore.com
Christopher Dorsey	dorsey@es.net
Tom Easterday	tom@cic.net
Havard Eidnes	havard.eidnes@runit.sintef.no
Robert Enger	enger@seka.reston.ans.net
Stefan Fassbender	stf@easi.net
Dennis Ferguson	dennis@ans.net
Robert Fink	rlfink@lbl.gov
Dale Finkelson	dmf@westie.mid.net
Peter Ford	peter@goshawk.lanl.gov
Catherine Foulston	cathyf@rice.edu
Vince Fuller	vaf@barrnet.net
Vincent Gebes	vgebes@sys.attjens.co.jp
Herluf Hansen	hha@tbit.dk
Denise Heagerty	denise@dxcoms.cern.ch
Matt Hood	hood@nsipo.nasa.gov
David Jacobson	dnjake@vnet.ibm.com
Dale Johnson	dsj@merit.edu
Matthew Jonson	jonson@ddn.af.mil
Jeanine Kamerdze	kamerdze@nsipo.nasa.gov
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Sean Kennedy	liam@nic.near.net
Charley Kline	cvk@uiuc.edu
Mark Kosters	markk@internic.net
John Krawczyk	jkrawczy@wellfleet.com
Walter Lazear	lazear@gateway.mitre.org
Tony Li	tli@cisco.com
Robin Littlefield	robin@wellfleet.com
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se
Bill Manning	bmanning@rice.edu
Glenn Mansfield	glenn@aic.co.jp
Jun Matsukata	jm@eng.isas.ac.jp
Stephen Miller	smiller@bbn.com
Pushpendra Mohta	pushp@cerf.net
Dennis Morris	morris@altair.disa.mil
Jun Murai	jun@wide.ad.jp
Michael O'Dell	mo@uunet.uu.net

Vijayaragavan Pandian	vjp@proteon.com
Andrew Partan	asp@uunet.uu.net
Brad Passwaters	bjp@eng.umd.edu
Michael Patton	map@bbn.com
David Piscitello	wk04464@worldlink.com
Kenneth Rehbehn	kjr@netrix.com
Yakov Rekhter	yakov@watson.ibm.com
Isil Sebuktekin	isil@nevin.bellcore.com
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Timon Sloane	timon@timonware.com
Frank Solensky	solensky@ftp.com
Bernhard Stockman	boss@ebone.net
Larry Tepper	ltepper@compatible.com
Marten Terpstra	marten@ripe.net
Claudio Topolcic	topolcic@cnri.reston.va.us
Jerry Toporek	jt@mentat.com
Paul Traina	pst@cisco.com
Keisuke Uehara	kei@cs.uec.ac.jp
William Warner	warner@ohio.gov
Chris Wheeler	cwheeler@cac.washington.edu
Jane Wojcik	jwojcik@bbn.com
David Woodgate	David.Woodgate@its.csiro.au
Jessica Yu	jyy@merit.edu

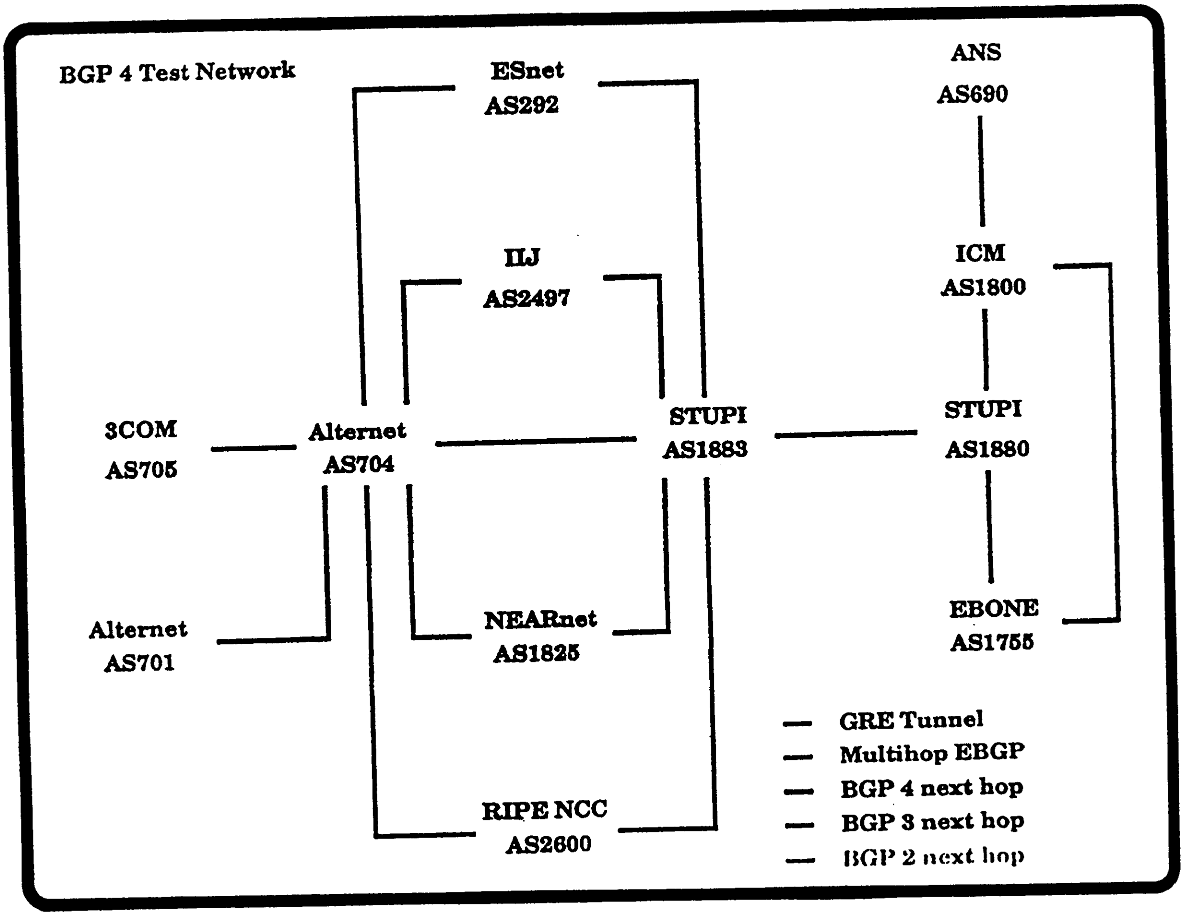
BGP 4 Test Network
ESnet
AS292
ANS
AS690
IIJ
AS2497
ICM
AS1800
3COM
AS705
Alternet
AS704
STUPI
AS1883
STUPI
AS1880
Alternet
AS701
NEARnet
AS1825
EBONE
AS1755
RIPE NCC
AS2600
GRE Tunnel
Multihop EBGP
BGP 4 next hop
BGP 3 next hop
BGP 2 next hop

Internet Registry

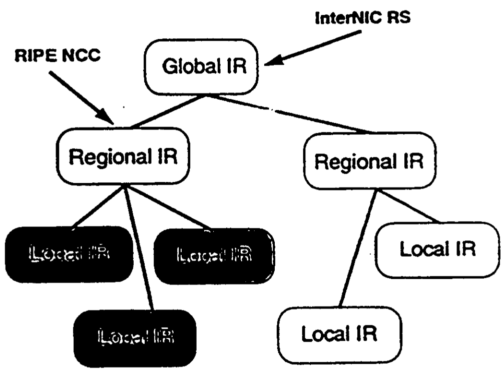

We have:

m 193.0.0.0 mask 255.0.0.0
m 194.0.0.0 mask 255.0.0.0

m couple of Bs at a time

Procedures

m blocks of 256 class Cs to provider and non-provider registries
m reserve enough blocks for their 2 year estimates
m make sure even assigned and reserved blocks are CIDRable
m force registries to assign in CIDR blocks to their customers
m can claim unused reserved blocks back
m any customer assignments over 32 Cs must be OK'ed by the NCC
m NCC does all B assignments

Procedures

m documented in ripe-72
m registries have to understand and adhere to ripe-72 to get any address space
m if they don't ...

m ripe-72 needs update, will be done before end of year

Status

(2nd November, 1993)

m 193.x.y
  m 82 blocks reserved
  m 173 blocks delegated
  m 19204 nets assigned
  m 1738 nets routed (on Euro-RS)
m 194.x.y
  m 12 blocks reserved
  m 16 blocks delegated
  m 131 nets assigned
  m 4 nets routed (on Euro-RS)

2.5.2 Benchmarking Methodology (BMWG)

Charter

Chair(s)

Jim McQuaid: `mcquaid@wg.com`

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"

CURRENT MEETING REPORT

Reported by Jim McQuaid/Wandel & Goltermann Technologies

Minutes of the Benchmarking Methodology Working Group (BMWG)

Twenty-eight people attended the meeting.

Scott Bradner opened the meeting with an announcement that Jim McQuaid is the new BMWG Working Group Chair. Scott then gave an overview of the working group's activities and briefly reviewed the contents of the two working documents. He mentioned that both documents are available from `hsdndev.harvard.edu` in the `pub/bmwg` directory.

The subject of adding modems and methods for testing bridges and routers to the document was brought up. It was thought that adding modems would be straightforward.

The attendees then discussed the charter at some length, reviewing aspects of the work which was done and the work which needed to be addressed in the future. The revised "goals and milestones" section of the charter will be changed to reflect the future work of the BMWG Working Group.

A survey of articles on benchmarking methodologies as well as research being done in this area will be conducted and reported on at the next meeting.

Attendees

Anders Baardsgaad	anders@cc.uit.no
Tom Benkart	teb@acc.com
Jon Boone	boone@psc.edu
Peter Cameron	cameron@xylint.co.uk
John Chang	jrc@uswest.com
Robert Fink	rlfink@lbl.gov
Kathy Huber	khuber@wellfleet.com
Rick Jones	raj@cup.hp.com
Merike Kaeo	mkaeo@cisco.com
Sean Kennedy	liam@nic.near.net
Edwin King	eek@atc.boeing.com
Andrew Knutsen	andrewk@sco.com
Dan Magorian	magorian@ni.umd.edu
Jim McQuaid	mcquaid@wg.com
Stephen Miller	smiller@bbn.com
Dennis Morris	morris@altair.disa.mil
William Palter	palter@tgv.com
Brad Passwaters	bjp@eng.umd.edu

Eddie Renoux	elr0262@newsit2.mcdata.com
Steven Richardson	sjr@merit.edu
Allen Rochkind	Allen_Rochkind@3com.com
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Vladimir Sukonnik	sukonnik@process.com
Larry Tepper	ltepper@compatible.com
Dean Throop	throop@dg-rtp.dg.com
Jerry Toporek	jt@mentat.com
Cathy Wittbrodt	cjw@barrnet.net

2.5.3 Generic Internet Service Description (GISD)

Charter

Chair(s)

Daniel Karrenberg: daniel@ripe.net
Tony Bates: tony@ripe.net

Mailing Lists

General Discussion: giss-wg@ripe.net
To Subscribe: giss-wg-request@ripe.net
Archive:

Description of Working Group

GISD collects short descriptions of Internet service aspects. Internet service in GISD means the interaction of Internet service providers among themselves and with their users. GISD aims to provide a common frame of reference and vocabulary to talk about an Internet service. For each aspect of the Internet service, it describes different options for service provision in use in the current Internet. GISD is merely descriptive and does not proscribe or mandate. The GISD document is intended to be a living document, collecting the work of many contributors.

The GISD Working Group will update and revise the GISD document to assist network service providers in a better understanding and description of what Interent service means.

- Update and revise the GISD document that lists the areas and aspects of interest to TCP/IP network service providers.

- Identify additional GISD areas and aspects appropriate to GISD.

- Identify areas of overlap with other IETF working groups.

- Create a reference document of GISD terms.

- Establish procedures to ensure the ongoing maintenance of the document and identify an organization willing to do it.

Goals and Milestones

Nov 1993 Review current GISD draft and add any additional areas and aspects felt essential.

Mar 1994 Draft of GISD will be prepared. Submit as Internet-Draft.

Jul 1994 Follow-up with final amendments and submit to RFC Editor for publication as an FYI RFC.

CURRENT MEETING REPORT

Reported by Tony Bates/RIPE NCC

Minutes of the Generic Internet Service Description Working Group (GISD)

After the agenda was presented, Tony Bates gave a general overview of GISD. The overview prompted a short discussion on focus including the following topics:

- Is it meant for users ? No
- Is it meant as a checklist ? No
- Is it a service profile ? No, but could be in the future

An overview of the GISD structure was also given showing how GISD aspects are documented. Some concern was raised about the use of minimal/common/maximal in the sense that this could cause some classification of service providers (SPs). However, the general view was that in the the context of GISD, it should be possible to word this in such a way so as not to make this happen. The idea behind this is to show the options rather than to categorise; occasionally it is useful to show different options.

The issue of "who or what constitutes an SP" was raised. The point is that anyone can call themselves a service provider but this is in fact not the issue of GISD. The point behind GISD is to write a document so people know what services an SP can potentially provide and how SPs should interact with each other regarding these services.

Some people also GISD as a direct template (i.e. a "tick the box" type of document). Again, this is not the intention of GISD. However, it could be possible at a later date to produce a template using the terminology and list of aspects detailed in GISD.

Tony Bates gave a overview of the aspects themselves.

It was noted that "training" is becoming more of an issue for SPs to provide to their customers and this could an aspect within the "Information and coordination" area.

A status of the current draft was given. Thirty-eight aspects have been defined as a result of two previous IETF BOFs on the subject. Thus far only eight aspects had been drafted, and the intention is to get members of the GISD group to draft some of these aspects. An action was placed on Tony to produce a short GISD aspect guideline document giving details of the format required and an index of possible aspects still needed to be drafted. The intention would be to work on the areas one at a time.

A basic overview of the process:

```
        Guideline Document (Format + index)
               |
               |
              \./
             Areas ----------------> Volunteers (select an aspect)
              /.\                       |
               |                        |
               |                        |
               +--------------------Review
               |
              \./
           All 6 areas drafted
```

A related idea was when sending the index of aspects out, to seek good candidates for a certain topic outside of the working group. Several people committed to writing aspects once they had seen the guideline document. The meeting concluded with the general consensus that the document should be possible to come together for review by Seattle and final copy by Toronto.

Attendees

Vikas Aggarwal	vikas@jvnc.net
Tony Bates	tony@ripe.net
Steven Blair	sblair@us.dell.com
Scott Bradner	sob@harvard.edu
Henry Clark	henryc@oar.net
David Conrad	davidc@iij.ad.jp
Vince Fuller	vaf@barrnet.net
Eugene Hastings	hastings@psc.edu
Matt Hood	hood@nsipo.nasa.gov
Jeanine Kamerdze	kamerdze@nsipo.nasa.gov
Scott Kaplan	scott@ftp.com
Sean Kennedy	liam@nic.near.net
Kim Long	klong@sura.net
Dan Magorian	magorian@ni.umd.edu
Bill Manning	bmanning@rice.edu
Glenn Mansfield	glenn@aic.co.jp
Stephen Miller	smiller@bbn.com
Pushpendra Mohta	pushp@cerf.net
Scott Paisley	paisley@central.bldrdoc.gov
Andrew Partan	asp@uunet.uu.net

Brad Passwaters	bjp@eng.umd.edu
Marsha Perrott	perrott@prep.net
Henry Sinnreich	hsinnreich@mcimail.com
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
John Veizades	veizades@ftp.com
Evan Wetstone	evanw@vnet.ibm.com
Chris Wheeler	cwheeler@cac.washington.edu

"GENERIC INTERNET SERVICE DESCRIPTION" (gisd)

Agenda

- Administrivia
- Brief overview of GISD
  - General Introduction
  - Status Report
- Discussion of FAQ
- Review of Areas and Aspects
- Call for volunteers
- AOB

GISD Overview

- Aimed at Service Providers
- Collects short descriptions of Internet Service aspects
- Tries to make it easier to talk about Internet Service
- Not a mandating Document
- Is needed

Areas

- Access
- Generic
- Connectivity
- Operations
- Information Provision and Coordination
- Security

Structure

```
GISS
    Area
        Aspect
            What
            Why
            Options
                (Minimal)
                (Common)
                (Maximal)
                (Regional)
            See Also
            Soap Box
            References
            Contributors
        ...
    ...
```

Current Status ?

- Third GISD draft available
  ftp.ripe.net:ripe/docs/ripe-drafts/giss.{ps,txt}
- Decided upon
  - Scope
  - Structure
  - GISS Areas
  - GISS Aspects
- Still needs a lot more work
  - Many aspects still to be completed
  - Contributors
  - Areas and Aspects to be reviewed

2.5.4 Network Joint Management (NJM)

Charter

Chair(s)

Gene Hastings: hastings@psc.edu

Mailing Lists

General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu
Archive:

Description of Working Group

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the Internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the agenda of this group will include the IAB and its task forces, and groups within FARNET. In particular FARNET has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing.

Goals and Milestones

None specified

CURRENT MEETING REPORT

Reported by Eugene Hastings/Pittsburgh Supercomputing Center

Minutes of the Joint Sessions of the NJM and NETSTAT Working Groups

Presentation slides from the NSI and ANSnet status reports can be found in Section 3 of the proceedings.

SESQUINET - Bill Manning

Sesquinet has a staff of three, and recently installed and moved to an FDDI DMZ for their ANS connection. Usage is growing by 15% per month.

Defense Simulation Internet - Mike Patton

The Defense Simulation Internet (DSI) is used to connect the defense industry and perform mission-specific experiments (like encrypted video). It has a T1 cross-country backbone.

NASA Science Internet - Jeff Burgan

It is the 10th Anniversary of the DNS RFC! A T3 has been installed from Goddard Space Flight Center (GSFC) to SURAnet, and another T3 from GSFC to Naval Research Laboratory (NRL) (using cisco). Routed traffic is both IP and Phase IV DECnet, with some CLNP. NSI provides mail and protocol interoperability services. Maps are available via anonymous FTP from `nsipo:nsi/maps/*`. Recent accomplishments include:

- The United Kingdom fat pipe has been transitioned to ICM/Sprint.
- NASA now has a 256kb dedicated link to ULCC (London).
- Australia (AARnet) link has been upgraded to T1.
- Japanese links have been reterminated to FIX-West (from Hawaii).
  - WIDE - 192kb
  - TISN - upgraded 128kb to 512kb
- Korea was upgraded to 256kb.
- Hong Kong link has been upgraded to 128kb.
- FIX-West transitioned to FDDI (but ESnet still Ethernet) Mbone router and DNS servers live on the Ethernet.
- NASA/NSF support for Antarctic connectivity: Upgrade to T1 (IP 512k, balance is video) to McMurdo Station.

Future plans include:

- NSI Backbone upgrades using inverse multiplexers.
  - NASA internal telecomm does multiplexed T3, but cannot offer channel breakout at greater than E1, so NSI is using inverse multiplexers to recombine lower speed channels to HSSI.
  - Digital Link is smart about failed links, even about error per second threshold, and deduces composite clocks.

- General purpose Internet connectivity to Russia via IKI (Russian Space Research Institute) at 256kb is scheduled for circa January 1994. There will be about ten additional sites added, using analog point-to-point links. The connectivity is basically confined to the Moscow area. NASA went through COCOM to approve equipment and Internet service. NASA must retain physical and configuration control, but routes will be exported outside of NSI.

- NSI expects to install a link to Argentina Space Institute at 64Kb BGP deployment! They are looking at multicast routing using MOSPF. Proteon supports MOSPF/DVMRP coupling. You can open a tunnel to a Proteon router and have it do multicast.

- ATM: The DoE/NASA procurement has been re-cycled. Proposals are being evaluated. NASA has gotten their own permission to procure ATM service and hopes to have ATM, at the latest by mid 1994.

ANS - Jordan Becker

The busiest ANS link is Chicago to Cleveland (in PPS). There was a spike in the external route flap last week. Andrew Partan reported that an ignorant or unfriendly site sent UDP packets to every sequential host in the address space. Since this forces a new router cache entry with each successive packet, it can make a router (empirically determined to include cisco) overflow routing cache. The source, albeit unconfirmed, is believed to be an address in Taiwan. Alternet routers got full caches, memory leaks caused dropped BGP sessions. As the source started with low class A networks, the flaps moved from provider to provider, hitting Alternet at network 7.

Alternet - Mike O'Dell

Alternet is now running on an ATM Backbone, sourced by MFS DataNet. The physical topology is a large ring, with cut-through paths, interconnecting Newbridge ATM switches. All switches are interconnected with path diverse dual T3 links. The lines are for redundancy, with a box that does hot fallback between T3s (the ATM switch only sees one at a time).

Alternet now has ATM in seven major hubs, plus San Jose. Alternet's access to this ATM fabric is via a 10Mb AUI connector, behind a learning half-bridge. MFS DataNet takes care of the internals. (It will be some time before Alternet considers a native ATM connection—the cisco ATM Applique is said to cost $30k!) Since one can in principal do FDDI to NewBridge switch, it will be a while before other connections are needed.

MFS DataNet provides a virtual private network on fabric with multiple customers. (Alternet is the largest customer willing to be publicized—there are other, private users of MFS DataNet in the banking and insurance industries.) DataNet owns the switches, but buys the circuits from the regular suspects: WilTel, etc.

Concerning management, there is a list of criteria (bandwidth, delay, etc.) for MFS operations. Alternet's Demark is the AUI. MFS has just announced DataNet service to London. An overall map is available online: `ftp.uu.net:uunet-info`. Alternet staff reports satisfaction with the caliber of MFS Telco people and DataNet data communications people.

COREN - Scott Bradner

Carrier negotiations are still underway, anticipated to close within a couple of weeks. When a carrier is selected, substantial workouts will be performed. There is an (unspecified) fallback plan if a selected carrier is unsuitable. COREN's view of the carrier network is still under discussion.

COREN's interface to the LD carrier is SMDS initially, and is planned to migrate to ATM. Its DMZ is parallel Ethernet and FDDI.

COREN has also put out an RFI on NOC services, and put out an invitation to participate in equipment evaluation. cisco and Wellfleet and two vendors of CSU/DSU equipment have signed on so far. COREN has set up a testbed, with BGP4.

COREN has eight regionals as its founders, plus a number of undisclosed committed or potential subscribers.

COREN is working in parallel to the RIPE routing database effort, and is starting to work with Merit on transition issues.

EBONE - Bernhard Stockman

In the last year, there has been rapid expansion into eastern Europe. Countries near Austria are now connected to Vienna hub. Warsaw (currently connected to Vienna and Stockholm) may become another BB site, feeding Baltic.

Within the last year, all of the transatlantic lines have been connected to a single point, the Global Internet Exchange (GIX) in Washington, DC. (Andrew Partan observed that the GIX Ethernet is close to saturation.)

The EBONE is now running BGP4 on all EBONE Border Systems (EBS), and will not do proxy aggregation for stability reasons.

EBONE route policy filters are performed by means of comparing an administrative database with a live routing table.

Connections to other providers include:

- EUNET
- EPNET
- SPAN
- Unisource Business Systems (UBS)

  This is a joint venture of Dutch, Swedish, and Swiss PTTs, stimulated by a call for tender for Pan-European X.25.

- European MultiProtocol Backbone (EMPB)

  DANTE, Inc., formed by the European research and education community is reselling EMPB as Europanet, in a bundle with X.400 and X.500 services. There is a 1Mb connection between EBONE and EMPB in Amsterdam. Because of the overlap between communities, there are expected to be organizations leaving EBONE for EMPB. Trying to extend the logical boundary of the GIX to Stockholm and Paris to support interconnections. (MAC layer extensions, etc.)

Attendees

Vikas Aggarwal	vikas@jvnc.net
Jordan Becker	becker@ans.net
Bart Berger	bart_berger@3com.com
Rebecca Bostwick	bostwick@es.net
Scott Bradner	sob@harvard.edu
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
David Conrad	davidc@iij.ad.jp
Christopher Dorsey	dorsey@es.net
Tom Easterday	tom@cic.net
Roger Fajman	raf@cu.nih.gov
Stefan Fassbender	stf@easi.net
Dale Finkelson	dmf@westie.mid.net
Catherine Foulston	cathyf@rice.edu

Vince Fuller	vaf@barrnet.net
Steve Garritano	steveg@kalpana.com
Farrell Gerbode	farrell@rice.edu
Herluf Hansen	hha@tbit.dk
Eugene Hastings	hastings@psc.edu
Dale Johnson	dsj@merit.edu
Jeanine Kamerdze	kamerdze@nsipo.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Jian Li	jian@rice.edu
Kim Long	klong@sura.net
Glenn Mansfield	glenn@aic.co.jp
Stephen Miller	smiller@bbn.com
Pushpendra Mohta	pushp@cerf.net
Michael O'Dell	mo@uunet.uu.net
Andrew Partan	asp@uunet.uu.net
Michael Patton	map@bbn.com
Marsha Perrott	perrott@prep.net
Martin Schulman	schulman@smtp.sprint.com
Erik Sherk	sherk@sura.net
Louis Steinberg	louiss@vnet.ibm.com
Bernhard Stockman	boss@ebone.net
Claudio Topolcic	topolcic@cnri.reston.va.us
Evan Wetstone	evanw@vnet.ibm.com
Chris Wheeler	cwheeler@cac.washington.edu
Cathy Wittbrodt	cjw@barrnet.net
Richard Woundy	rwoundy@vnet.ibm.com
Jessica Yu	jyy@merit.edu

2.5.5 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.
Done	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.

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", 04/23/1993, S. Hares, C. Wittbrodt <draft-ietf-noop-echo-02.txt>

"Essential Tools for the OSI Internet", 06/07/1993, S. Hares, C. Wittbrodt <draft-ietf-noop-tools-03.txt>

2.5.6 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"

2.5.7 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")"

2.6 Routing Area

Director:

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun Microsystems

New Internet Routing and Addressing Architecture BOF (NIMROD)

The group reviewed the current draft working group charter and the latest proposed terminology list. General satisfaction was expressed with the current state of both.

Discussion then moved on to some of the open architectural issues. Among the points discussed were:

- Can areas overlap?
- Are abstraction levels identified explicitly?
- Do the nodes in the graph of the network represent interfaces or routers/networks?
- Do interfaces have locators?
- Are the labels which elements of locators globally unique?
- Do locators grow up, down, and can they be expanded in the middle?
- Are partial locators possible?
- Do routers have locators?
- Do we have separate namespaces for interfaces and endpoints?
- What is the smallest thing which can be an endpoint?
- Do we have a hop-by-hop mode, or just source routed packets and flows?
- Do we retain the EGP/IGP split?
- When do we tackle multicast?

The following action items were decided on:

- The meetings at the next IETF should be scheduled for Tuesday and Wednesday mornings if possible.
- All new open issues raised during the working group meeting are to be sent to the working group mailing list.
- The chair will include the new points, re-sort the list into priority order, add a new category of "local" for issues, and resubmit.
- A document showing the outcome of the discussions on the open items will be prepared and sent to the list.
- A moderated list discussion will take on remaining open issues.
- Scheduling a Boston interim meeting will be investigated.
- The working group agreed to have a draft of the architecture RFC, prepared by the end of January 1994, for final examination at the March IETF.

Border Gateway Protocol Working Group (BGP) and OSI IDRP for IP Over IP Working Group (IPIDRP)

The BGP and IPIDRP Working Groups met jointly. All outstanding technical issues with the BGP-4 protocol were resolved. The resulting changes will be incorporated in the appropriate documents, and the documents will be submitted as Internet-Drafts before Thanksgiving with the purpose of advancing BGP-4 to a Proposed Standard. The group also discussed IDRP status and several future enhancements to BGP/IDRP, including domain partition repair and router servers.

Inter-Domain Multicast Routing Working Group (IDMR)

The two PIM documents (PIM = Protocol Independent Multicast, formerly ESL), dense and sparse modes, were presented and discussed. Though some details about the phase shift between sparse and dense mode need to be worked out, the general consensus of the group is that the multiple scaling modes approach is desirable. Implementation of PIM will continue.

No work was done on CBT, but a status report was given describing CBT's state of implementation (almost done). There is still interest in CBT as valuable work, either as a potential alternative to PIM (if PIM proves overly difficult), or as an Experimental Protocol.

The group decided to propose a new name and charter to better reflect that the focus is no longer strictly inter-domain, but rather scaling versus quality in general. Paul Francis will generate the proposal.

Inter-Domain Policy Routing Working Group (IDPR)

The IDPR working group met for one session during this IETF. It spent the majority of the time discussing what is being called IDPR version 2. Version 2 contains support for multicast and multipath routing as well as policy-based resource allocation. The gated implementation of version 2 will begin its testing phase next month. In the early spring, an Internet-Draft will be produced describing the changes to the IDPR version 1 protocols to support this functionality.

The group also received a presentation (via videotape) on the "Routing by Preference" work of Yuko Murayama and colleagues, and we plan to discuss this more on the mailing list.

At the request of the Routing Area Director, the IDPR working group will conclude with this IETF. The group will restart when either an additional independent implementation of IDPR version 1 can be submitted for Draft Standard or when the Internet Draft specification of version 2 is complete. In the meantime, the mailing list will remain open.

Also, there are two new Internet-Drafts, both updated versions of existing documents. One is the MIB and one is the DNS modifications for IDPR. We plan to submit the MIB for consideratin as a Proposed Standard.

IP Routing for Wireless/Mobile Hosts (MOBILEIP)

The MOBILEIP Working Group held an interim meeting on the 9th and 10th of September in Summit, New Jersey. The two day meeting was quite productive. We agreed on a basic model for how mobile-ip works. We then discussed the various messages and information that would need to be passed between the various entities. We selected an editor for the working group document—Charles Kunzinger from IBM. (Charlie was previously editor of the ISO IDRP effort.)

The MOBILEIP Working Group met twice at the 28th IETF. Charlie Kunzinger gave a tutorial introduction to the first draft document he has produced. The group then reviewed this draft and also reviewed the work of three other members of the working group (who have formed an alliance; before they had between them four or five documents, and now only one).

The group plans to have a firmer draft by the end of the year. There are plans for another interim meeting in January. We hope to have a draft specification by the Seattle IETF (and maybe even an implementation or two).

IS-IS for IP Internets Working Group (ISIS)

The ISIS Working Group meet for one session. The major topic discussed was multicast support in ISIS. Three types of multicast were identified: "anycast" for the nearest service location, dense multicast, and sparse multicast. The first two could be supported by ISIS while sparse multicast is best done by some multicast tree approach. This work needs to be brought to the attention of the IDMR Working Group.

The working group also discussed the IPX and Appletalk integration scheme (available as an Internet-Draft) and Novell's NLSP protocol which was derived from ISIS. The group drew up a list of work items, some of which would require enhancing the protocol as defined in the latest ISIS Internet-Draft. Incorporating these changes would probably require defining a new version of the ISIS protocol.

Open Shortest Path First IGP Working Group (OSPF)

The OSPF Working Group met on Wednesday, November 3. The following items were discussed:

- Status Overview
- OSPF Scaling Issues - "Ringing It Out At The Next Level"
- On-Demand Circuit Proposal
- NSSA Implementation And Status
- MIB Changes And Status

RIP Version II Working Group (RIPV2)

The RIP-2 Protocol Internet-Draft was approved by the working group for submission for consideration as a Draft Standard to replace RFC 1388. The MIB was similarly approved to replace RFC 1389.

There are two new implementations of RIP-2, bringing the total to four. Details on the implementations will be provided in a revision of the RIP-2 Protocol Analysis which will be done this month.

The Demand Circuit Routing Internet-Draft by Gerry Meyer was approved for submission for consideration as a Proposed Standard. The Protocol Analysis Internet-Draft will be submitted as an Informational RFC.

Consideration of the SIPP-RIP draft, particularly the Loop Detection algorithm, was postponed until RIP-2 has been accepted as a Draft Standard (so as not to affect that effort). Discussion of the algorithm will be started next month on the ietf-rip mailing list and will be discussed in detail in Seattle.

Routing over Large Clouds Working Group (ROLC)

The ROLC Working Group met for two sessions. The first session had a brief review of the charter, and a discussion of the assumptions about media and network topology. The group briefly discussed the IS-IS over NBMA and RIP over demand circuit documents. There were some issues raised, which will be carried back to the relevant working groups.

The second session was devoted to a discussion of two documents. The discussion of the Braden/Postel/Rekhter architectural document raised a number of issues. There was definite support from this working group for the general purpose and approach. The group consensus was that certain solutions less favored in the document (query/response mechanisms) were important tools.

The group then reviewed the details of the NHRP proposal. It discussed the behavior in the normal case, and the responsiveness to changes in underlying routing. One major flaw which could produce loops was pointed out. An approach to the solution was also suggested. It will be necessary for the group to work more on this issue. There was also the suggestion that we adopt a solution which only works in the absence of address aggregation within the large cloud. The solution and its applicability will be discussed on the e-mail list, while discussion of the more general case continues.

Source Demand Routing Working Group (SDR)

The working group performed a protocol walk-through of the SDR document, and found that only editorial changes were needed. The working group will be reviewing these changes shortly, and submitting the specification for approval as an Experimental RFC.

The working group held brief discussions about route selection and efficient mapping of packets to SDRP routes. Progress on other working group issues was somewhat lacking. Due to personal emergencies, several key members of the working group were not able to attend.

CURRENT MEETING REPORT

Reported by Noel Chiappa

Minutes of the New Internet Routing and Addressing Architecture BOF (NIMROD)

The Nimrod BOF met on Thursday, November 4, 1993. The discussion was lead by Noel Chiappa. Isidro Castineyra, co-chair, took notes on the discussion.

Agenda

- Agenda bashing.
- Review of proposed charter.
- Review of existing and proposed new terminology.
- Debate on some items from "open architectural issues" list.
- Work plan for immediate future.

No changes to the agenda were proposed. Also, there were no comments on the charter and the terminology listing. This was an introductory meeting intended to start the group's work, as such it consisted of the discussion of basic open issues. The rest of these minutes record the discussion on the open issues and the work plan agreed to.

Open Issues Discussion

- Can clusters overlap?

  The argument was made that overlapping clusters would be necessary for re-organization of cluster boundaries to provide a better abstraction hierarchy as the physical topology changed. In this situation, interoperation and updating would be much easier if both the old structure and the new could co-exist for a while. Once this mechanism—overlapping clusters—is available, it could be used for other— unspecified—means.

  It was also pointed out that overlapping clusters will result in endpoints possibly having multiple locators, this could be (mis?)-used for biasing the route generation mechanism. Some people favored this, saying that having multiple locators allowed clients to select which one gave the desired routing behavior. Others maintained that this was exactly the wrong way to do policy, and the locator should simply uniquely name the location of the endpoint, and preferred that other mechanisms—within the routing component—be defined for the purpose of policy, route optimization, etc. Route suffixes, as proposed by David Clark, are one example of such a mechanism.

  It was argued that overlapping clusters would make difficult the enforcement of transit policies. An alternative mechanism to overlapping clusters, to allow re-organization, would be to have multiple hierarchies at different levels. If a simpler re-organization mechanism could be found, overlapping clusters might be unnecessary, resulting in a simpler architecture.

- Are abstraction levels identified explicitly?

  It was argued that explicit levels would prevent growth of the network map at different levels of the network. (In some sense, this is the same question as "Do locators grow up, down, and can they be expanded in the middle?")

  In other words, if an endpoint were located at A.B.C.D.E (to invent a representation of a multi-level hierarchical locator), and cluster A.B.C became too large, so that it had to be split up into C1 ... CN, (resulting in locators of the form A.B.C.C5.D.E), this process would be made more difficult if the cluster A.B.C.D was known to be at the fourth level (counting from the top; the equivalent is A.B being at the fourth level, if counting from the bottom).

  It was also argued that if locators are given from the top, explicit levels are not necessary. (Another way to put this is "Are partial locators possible?") On the other hand, if the locators can grow on the top end (as the network expands, say), a locator which used to start at the top level no longer does so. Since these old locators are likely to be around for a while after a new level is added, some way has to be found to deal with them.

- Are the labels of locators globally unique?

  This question is obviously related to the previous question of partial locators. If the label of each element in a locator is globally unique, it is not necessary to specify which context (i.e., location in the abstraction hierarchy) to use to interpret any partial locator.

  It was pointed out that globally unique labels, while theoretically attractive, would make locators very long. The consensus was that this was probably not necessary.

- Do we have a hop-by-hop mode, or just source routed packets and flows?

  It was argued that a hop-by-hop mode is, in a sense, inherent in a hierarchical network, because intermediate points might have to supply additional route detail not contained in the original source route, when this has been generated using a map without the necessary detail. Such detail might have been unobtainable, if a cluster has an information-hiding policy which prevents *any* information about the internal topology of that cluster from going outside the cluster.

  Strictly speaking, this does not have to be handled by a hop-by-hop mode, since the entry point into the closed area could generate the rest of the path on entry, and either add it to the flow path (for a flow setup), or the source route in the packet (for a source-routed packet). However, such a cluster could run hop-by-hop mode inside the cluster without anyone outside being any the wiser. (In fact, Nimrod imagines that exactly such an operational mode will be used during Nimrod deployment, to handle areas of non-converted old-style routing.)

However, this does not fully answer the original question, since a hop-by-hop mode would mean that all routers in the system have to support such a mechanism, not just those in closed areas. The question really is "How little detail can a source give in a source route?" If the minimum source route consists of only the destination locator, then the system does have to support hop-by-hop mode, or at least something which looks a lot like it, in the sense that the source just labels the packet with the ultimate destination, and lets the routers work out how to get the packet there.

- Do we retain the EGP/IGP split?

  The consensus was that the EGP/IGP split cannot be eliminated, as a given cluster that does not give out its internal organization can always operate internally using any routing architecture it wishes, as pointed out above. However, the notion of a single defined level which is "the" EGP/IGP boundary does appear to be counterproductive.

- When do we tackle multicast?

  It was suggested that multicast should be made the fundamental mode, with unicast as a special case of multicast. It was also pointed out that multicast affects only route generation and forwarding, the other components of routing—i.e., network connectivity representation, map distribution, etc.—are independent of the existence of multicast.

- Do the nodes in the graph representation of the network represent interfaces or routers/networks?

  This debate went on for a while, but no definite conclusion was reached. Those in favor of the former pointed out that it provided the most flexibility, and avoided situations like the difficulty of modeling a router which fell on an administrative boundary. Those in favor of the latter pointed out that interfaces and routers are the basic physical constituents of the network, and the map needed to be able to model them in a way that was both efficient (i.e., not in a way that needed N^2 arcs to model the internal connectivity of a network or a router) and easy to understand (since we need to build a system that many, many people will need to be able to work with).

- What is the smallest thing which can be a cluster?

  This point is obviously closely related to the one above. There were arguments in favor of interfaces, in favor of routers, and in favor of networks.

- Do routers have locators?

  Some think that routers can have locators, but, depending on the level of abstraction, these might not be available.

  The problem with routers having locators is that if a router is connected to two widely separated points in the abstraction hierarchy, which branch of the abstraction hierarchy do you place the router in? Alternatively, you can provide it with a locator which is at the same level as that at which the two branches join, but if there are many such routers, this may present a problem. Yet another alternative is to assign such a router several locators, one for each place where it is connected, but if this is done, perhaps it makes more sense to think of the locators as naming the interfaces, not the router.

  A related question is "Can we tell by looking at a locator whether it names an interface, a network, a router, or a cluster?"

- Do we have separate namespaces for interfaces and endpoints?

  Mobile endpoints are easier to handle if the endpoint has a name which stays constant while it moves. It is hard to see how to provide the latter without having a separate, non-topologically oriented, namespace for endpoints.

  The question then becomes "Do the topologically oriented names (i.e., locators) name endpoints or interfaces?" This is related to the question above. If an endpoint is in a host which has two widely separated interfaces, exactly the same set of options are available for dealing with the situation.

Action Items

The following action items were decided on:

- We will try to schedule the next IETF meetings for Tuesday and Wednesday morning.

- All new open issues raised during the working group meeting are to be sent to the working group mailing list.

- The chair will include the new points, re-sort the list into priority order, add a new category of "local" for issues, and resubmit.

- A document showing the outcome of the discussions on the open items will be prepared and sent to the list.

- A moderated list discussion will take on remaining open issues.

- Scheduling a Boston interim meeting will be investigated.
- The working group agreed to have a draft of the architecture RFC prepared by the end of January, 1994, for final examination at the March IETF.

Attendees

Nick Alfano	alfano@mpr.ca
Susie Armstrong	susie@mentat.com
Jim Barnes	barnes@xylogics.com
Jim Beers	Jim.Beers@cornell.edu
Nutan Behki	nebhki@newbridge.com
Ram Bhide	ram@nat.com
Jim Bound	bound@zk3.dec.com
Monroe Bridges	monroe@cup.hp.com
David Bridgham	dab@epilogue.com
Steve Buchko	stevebu@newbridge.com
Ross Callon	rcallon@wellfleet.com
Ken Carlberg	Carlberg@cseic.saic.com
Isidro Castineyra	isidro@bbn.com
J. Noel Chiappa	jnc@lcs.mit.edu
Matt Crawford	crawdad@fncent.fnal.gov
Michael Davis	mike@dss.com
Chuck de Sostoa	chuckd@cup.hp.com
Avri Doria	avri@locus.com
Havard Eidnes	havard.eidnes@runit.sintef.no
William Fenner	fenner@cmf.nrl.navy.mil
Eric Fleischman	ericf@atc.boeing.com
Dan Frommer	dan@isv.dec.com
Eugene Geer	ewg@cc.bellcore.com
Atanu Ghosh	atanu@cs.ucl.ac.uk
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Ramesh Govindan	rxg@thumper.bellcore.com
Regina Hain	rrosales@bbn.com
Dimitry Haskin	dhaskin@wellfleet.com
Marc Hasson	marc@mentat.com
Kathy Huber	khuber@wellfleet.com
Phil Irey	pirey@relay.nswc.navy.mil
David Johnson	dbj@cs.cmu.edu
Matthew Jonson	jonson@ddn.af.mil
Frank Kastenholz	kasten@ftp.com
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Lee Kilpatrick	leekil@bbn.com
Stev Knowles	stev@ftp.com

Sundar Kuttalingam	sundark@wiltel.com
David Marlow	dmarlow@relay.nswc.navy.mil
Jun Matsukata	jm@eng.isas.ac.jp
Wayne McDilda	wayne@dir.texas.gov
Greg Minshall	minshall@wc.novell.com
Randy Miyazaki	randy@lantron.com
Sath Nelakonda	sath@lachman.com
Vijayaragavan Pandian	vjp@proteon.com
Laura Pate	pate@gateway.mitre.org
Michael Patton	map@bbn.com
Eric Peterson	elpeterson@eng.xyplex.com
Ram Ramanathan	ramanath@bbn.com
Eddie Renoux	elr0262@newsit2.mcdata.com
Robert Roden	roden@roden.enet.dec.com
Shawn Routhier	sar@epilogue.com
Michal Rozenthal	michal@fibronics.co.il
Hal Sandick	sandick@vnet.ibm.com
Martin Schulman	schulman@smtp.sprint.com
Isil Sebuktekin	isil@nevin.bellcore.com
Michael See	mikesee@vnet.ibm.com
Frank Solensky	solensky@ftp.com
Karen Sollins	sollins@lcs.mit.edu
Tae Song	tae@novell.com
Martha Steenstrup	msteenst@bbn.com
John Stewart	jstewart@cnri.reston.va.us
Vladimir Sukonnik	sukonnik@process.com
Larry Tepper	ltepper@compatible.com
Michael Thatcher	thatcher@rahul.net
Dean Throop	throop@dg-rtp.dg.com
Panos Tsigaridas	Tsigaridas@fokus.gmd.de
Keisuke Uehara	kei@cs.uec.ac.jp
Taehwan Weon	weon@cosmos.kaist.ac.kr
Gerry White	gerry@lancity.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Jane Wojcik	jwojcik@bbn.com
John Wroclawski	jtw@lcs.mit.edu
Weiping Zhao	zhao@nacsis.ac.jp

2.6.1 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

"A Border Gateway Protocol 4 (BGP-4)", 12/23/1993, Y. Rekhter, T. Li <draft-ietf-bgp-bgp4-07.txt>

"Definitions of Managed Objects for the Border Gateway Protocol (Version 4)", 12/01/1993, S. Willis, J. Burruss, J. Chu <draft-ietf-bgp-mibv4-04.txt>

"BGP4/IDRP for IP—OSPF Interaction", 12/14/1993, K. Varadhan, S. Hares, Y. Rekhter <draft-ietf-bgp-bgp4ospf-interact-03.txt>

"Application of the Border Gateway Protocol in the Internet", 11/17/1993, Y. Rekhter, P. Gross <draft-ietf-bgp-application-03.txt>

"Application of the Border Gateway Protocol and IDRP for IP in the Internet", 10/18/1993, Y. Rekhter, S. Hares <draft-ietf-bgp-idrp-usage-00.txt>

"BGP-4 protocol document roadmap and implementation experience", 01/03/1994, P. Traina <draft-ietf-bgp-bgp4-implement-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

Reported by David Conrad/Internet Initiative Japan

Minutes of the Joint Session of the BGP and IPIDRP Working Groups

BGP4 Unresolved Issues - Dennis Ferguson

The following issues were discussed:

- Decisions have to be made regarding choosing a next hop forwarding address
- What should be done when there is no IGP route to the forwarding address

It was observed that the tie breaking rules can be directly derived from these two decisions.

With respect to choosing the next hop forwarding address, there are two options: using NEXT_HOP and using neighbor address.

The advantages to using NEXT_HOP for the forwarding address:

- Better routing when there are alternative paths to the DMZ
- Allows use of IBGP route servers
- If you don't care about third party NEXT_HOP, it is cheaper to not set the NEXT_HOP to a local address (not permitted by current spec)

A disadvantage of using NEXT_HOP is that the DMZ address must be propagated into the IGP before a third party NEXT_HOP can be advertised.

Advantages to the use of neighbor's address for the forwarding address are that there is less confusion about whether the DMZ needs to be propagated into the IGP or not, and that cisco BGP3 did it this way.

It was discussed that NEXT_HOP means an unstable IGP may result in retracted routes, and that the handling of IGP instability should be addressed in the specification. Other points presented are that Europeans have NEXT_HOP as a requirement and that NEXT_HOP gives better routing decisions.

There was discussion on what to do when there is no IGP route to the forwarding address. The option of not selecting a route for which you don't have an IGP route was examined as well as the option of blackholing traffic when there is no IGP route.

The advantages of not selecting a route are that an IGP cost for tie-breaking always exists, fallback routes can be used, and black holes are not readvertised. The advantage of blackholing is that it is easy.

Comments made with respect to the options available include the observation that either is interoperable and that there is sympathy for people blackholing as an easy solution when trying to get other things working.

The attendees decided to use NEXT_HOP as next hop forwarding address and not to select a route for which you don't have an IGP route.

Enhancements to AGGREGATOR - Paul Traina

Currently the AGGREGATOR path attribute contains the AS of aggregator. The attendees had decided to add an ASCII string, but subsequent discussions on the BGP mailing list reversed this decision.

The final decision on the AGGREGATOR path attribute is to add an IP address (in addition to the AS number), and indicate in the protocol specifications that this attribute is "highly recommended" for implementation.

LOCAL_PREF - Dimitry Haskin

It was pointed out that there are inconsistencies between the various BGP documents with respect to treating LOCAL_PREF. In the BGP4 Protocol specification higher value → higher preference, while in the BGP4 Usage document lower value → higher preference.

It was observed that BGP4 is unlike all other protocols (except BGP3) on the issue of preference. This could cause transition problems.

Yakov Rekhter volunteered to check the BGP/OSPF interaction document and insure higher value means higher preference (the protocol document is correct). The BGP4 documents will be clarified on this issue as well.

Erroneous NEXT_HOP - Tony Li

The subject of the discussion is how to handle the case when the NEXT_HOP value is wrong. It can be ignored but logged, or a non-fatal notification can be sent to the host generating the bad NEXT_HOP. It was decided that it would be ignored but logged since the other option would result in too much change to the specification. The notification option will wait until the next version.

BGP4 MIB - Andrew Partan

Possible improvements to the BGP4 MIB were discussed. More variables should be added that would be useful from an operational standpoint.

The meeting participants reached the following decisions:

- Define a MIB variable that contains elapsed time since the last BGP peering session establishment/termination. This variable is defined on a per peer basis. If the session was never established, this variable contains the elapsed time since the peer was configured.
- Define a MIB variable that contains elapsed time since the last UPDATE received. This variable is defined on a per peer basis.
- Combine internal and external BGP neighbors MIB tables together.

IDRP Status - Yakov Rekhter

IDRP reached full International Standard in October 1993. The document is available via anonymous FTP from `merit.edu` in PostScript (`/pub/iso/iso10747.ps[.Z]`) or ASCII (`/pub/iso/idrprfc.txt`).

It is expected that the document will be issued as an RFC as well.

IDRP Implementation - David Jacobson

Yakov Rekhter, Rob Coltun and David Jacobson participated in the implementation. It is a standalone IDRP that supports integrated IP and ISO routing. It can run over IP or CLNP, and is loosely coupled to GateD.

The implementation supports the following functions:

- Basic Transport
- Empty RIB-Att
- Confederations
- Policy
- Aggregation

It is expected that by the end of 1993,, the code will be completed and some testing will be performed. More internal system tests on the code will take place in early 1994. In late winter or early spring the code will be available for interoperability testing. The code will be given to NSF, and NSF will decide on the distribution of the code.

Domain Partition Repair - Dennis Ferguson

BGP does not currently handle partition healing like EGP does.

Two types of routing loops were defined:

- Permanent - routing protocol not required to send updates to terminate the loop
- Transient - routing protocol will send update which terminates loop

BGP routing loops were discussed. It was observed that BGP routing loops are always transient, and that the BGP specification chooses 1-cycle loop termination in all cases. If BGP allowed n-cycle loop termination with $n > 1$, partitions may be healable. The cost of setting $n > 1$ is that it can lead to transient loops that require a large number of updates to terminate.

The following changes to the document are needed to support Domain partition repair:

- In section 6.3 remove the check that an AS appears in the AS path only
- In section 9.3 remove the constraint against using a route with the local AS in the path
- Modify the aggregation procedures such that multiple occurences of an AS in the path of a route being aggregated are reflected in the aggregate path
- Modify the constraint in section 5.1.3 on advertising your neighbor's address as the next hop

The following comments were made during the discussion:

- Implementing domain partition repair could impose some addressing constraints (e.g. class As)
- Implementing domain partition repair requires removal of the AS-PATH check in section 6.3; however, removal of this check has no negative impact on the protocol
- Implementing domain partition repair will work only in presence of contiguous sequence of BGP-4 speakers. Passing routes that went through a partition repair to BGP-3 would result in terminating BGP peering with a BGP-3 speaker
- Implementing domain partition repair by removing ASs from the AS path is very dangerous

It was decided that the check in section 6.3 will be removed and that it should be verified that ATOMIC_AGGREGATE reduces the number of bits.

Advancing BGP4 to a Proposed Standard - Yakov Rekhter

Yakov Rekhter will cleanup LOCAL_PREF issues in usage documentation, changes to the BGP4 specifications will be done by Yakov Rekhter and Tony Li by Thanksgiving, and the BGP4 MIB will be updated by John Chu by Thanksgiving.

Selecting an Indirect Provider - Yakov Rekhter

The scheme is described in the Internet-Draft, `draft-rekhter-select-provider-00.txt`. It was discussed that with tunneling, even experienced users can run into trouble. It was also noted that more manageable mechanisms than tunnels are needed.

Route Server - Tony Li

Currently IBGP must be fully meshed. An alternative is to have an IBGP route server. Route servers would be fully meshed. An IBGP route server would constrain the amount of configuration and IBGP connections. The upper bound would be the number of border routers. Route server traffic would get all changes. Packet routing would be decoupled from data flow. To implement an IBGP router server would require an algorithm and protocol to elect designated route server.

Attendees

William Barns	barns@gateway.mitre.org
Stephen Batsell	batsell@itd.nrl.navy.mil
Rebecca Bostwick	bostwick@es.net
Al Broscius	broscius@bellcore.com
Randy Bush	randy@psg.com
Enke Chen	enke@merit.edu
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Christopher Dorsey	dorsey@es.net
Dennis Ferguson	dennis@ans.net
Carlos Fernandez	carlos@plk.af.mil
Vince Fuller	vaf@barrnet.net
Vincent Gebes	vgebes@sys.attjens.co.jp
Herluf Hansen	hha@tbit.dk
Susan Hares	skh@merit.edu
Dimitry Haskin	dhaskin@wellfleet.com
Denise Heagerty	denise@dxcoms.cern.ch
Robert Hinden	hinden@eng.sun.com

David Jacobson	dnjake@vnet.ibm.com
Dale Johnson	dsj@merit.edu
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Sean Kennedy	liam@nic.near.net
John Krawczyk	jkrawczy@wellfleet.com
Charles Kunzinger	kunzinger@vnet.ibm.com
Tony Li	tli@cisco.com
Robin Littlefield	robin@wellfleet.com
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se
Thang Lu	tlu@mcimail.com
Doug Montgomery	dougm@osi.ncsl.nist.gov
Robert Moose	rmoose@gateway.mitre.org
Dennis Morris	morris@altair.disa.mil
Sandra Murphy	murphy@tis.com
Vijayaragavan Pandian	vjp@proteon.com
Andrew Partan	asp@uunet.uu.net
Alex Reijnierse	a.a.l.reijnierse@research.ptt.nl
Yakov Rekhter	yakov@watson.ibm.com
Steven Richardson	sjr@merit.edu
Greg Ruth	gruth@gte.com
Dallas Scott	scott@fluky.mitre.org
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Paul Traina	pst@cisco.com
Chris Wheeler	cwheeler@cac.washington.edu
Cathy Wittbrodt	cjw@barrnet.net
Jessica Yu	jyy@merit.edu
Mary Jo Zukoski	maryjo@gateway.mitre.org

IBGP Hack Forwarding and Tie Breaking

Dennis Ferguson

Advanced Network & Services

The Decisions

- We know the immediate next hop for an IBGP route needs to be determined by picking a "forwarding address" from the route's information, and then looking for a route to the "forwarding address" in your IGP information. The next hop(s) associated with the latter should be used for the IBGP route.
- We know bad things will happen if the route you are forwarding through is not an IGP route. There needs to be a statement of what happens if the route to the forwarding address is not an IGP route, or is not known.
- It is understood that the IGP costs used for tiebreaking should be those associated with the IGP route to the forwarding address. This allows everyone to make a consistant decision.
- There are two candidates for the forwarding address, the NEXT_HOP of the IBGP route and the address of the router which sent you the IBGP route. This needs to be decided.
- There are two possibilities for what happens when you have no IGP route to the forwarding address, selecting the route but blackholing traffic to the destination, or avoiding selecting such routes. This needs to be decided.

The Forwarding Address

- The advantages of using the NEXT_HOP address as a forwarding address are (as I understand them):
  - Using the NEXT_HOP gives you somewhat better routing in situations where you have several alternative paths to reach the DMZ network, the best of which don't pass through the router which sent you the IBGP route.
  - Using the NEXT_HOP makes it possible to implement IBGP route servers, rather than having to maintain the full mesh of IBGP connections. It also makes it conceivable that one could design an IBGP flooding algorithm where IBGP speakers only need to talk IBGP to their direct neighbours. I don't know if this is useful or not.
  - If you don't care able third party NEXT_HOPs it is cheaper not to just set the NEXT_HOP to a local address. The current spec doesn't permit this.
- The disadvantages of using the NEXT_HOP address as a forwarding address are:
  - It requires that the DMZ address be propagated into the IGP before a third party NEXT_HOP can be advertised. This requires text in the draft to specify the the IBGP sends set the NEXT_HOP to a local address if they can't determine that the DMZ is being advertised.
- The advantages of using the neighbour address are:
  - Less confusion about whether the DMZ address needs to be propagated into the IGP or not.
  - Cisco BGP3 did it this way.

Missing IGP Forwarding Route

- The advantages of not selecting a route for which you don't have an IGP route to forward through are:
  - You always have an IGP cost for tie-breaking.
  - You get to use fallback routes which are available if you don't select non-working IBGP routes.
  - You don't readvertise black holes to neighbours
- The advantage of black holing traffic when there is no IGP route is:
  - Some implementations may find this easier to do. Such implementations may do this no matter whay the spec says.

AS Partition Healing and BGP

Dennis Ferguson

Advanced Network & Services

AS Partition Healing

- The Border Gateway Protocol provides loop-free routing in arbitrary topologies, with arbitrary configurations. In comparison its predecessor, EGP2, could only provide loop-free routing if the topology were constrained to the "core model".
- EGP2 supported healing of partitioned autonomous systems without routing loops, as long as all partition fragments retained connectivity to the "core". With BGP we lost the ability to support healing of AS partitions via inter-AS paths, because of the loop termination criterion chosen for BGP, except in limited circumstances (e.g. default routing).
- AS partition healing via inter-AS paths can be useful. The fact that BGP can't support it is a drawback.
- In fact BGP could support AS partition healing, while retaining the guarantee that routing loops will always be terminated, if the routing loop termination condition were modified to accomodate this. In particular, it should be recognized that the following statement from section 9.3 of the BGP4 draft is not always correct:

  – If the local AS appears in the AS path of the new route being considered, then that new route cannot be viewed as better than any other route. If such a route were ever used, a routing loop would result.

Routing Loops

- For the purposes of this discussion, a *routing loop* forms when a route uses a route which leads to a second router which is using a route which leads back to the first.
- A *permanent* routing loop occurs when a routing loop has formed but the routing protocol instances running on the routers involved in the loop are not required to send protocol updates which will terminate the loop.
- A *transient* routing loop occurs when a routing loop has formed, but in response the routing protocol instances running on the routers involved in the loop are required to send protocol updates which will terminate the loop.
- EGP2 would allow permanent routing loops to form if readvertisement of routes was not limited to the "core model" topology. BGP, like most interior routing protocols, cannot prevent transient routing loops from forming, but guarantees that when a transient routing loop has formed the protocol will be required to send updates which will eventually terminate the loop.

Policy Representation

- The examples following use a policy representation of the general form:

```
import proto <proto> aspath <regex> preference <pref> ;
```

- <proto> can be *bgp*, *ibgp* or *igp*.
- <regex> is a regular expression over the AS path. For example the following <regex> matches all routes with AS 690 in the path:

```
.* 690 .*
```

- <pref> is an integer value for the preference (lower values more preferred), or "none" if the route should not be used.
- The following examples also assume the constraint against using a route with the local AS in the path does not exist.

An Example of a BGP Loop

Consider the following topology, where X is a network route:

```
              X
              |
          +-------+
          |       |
          |   A   |
          |       |
          +-------+
          /       \
         /         \
        /           \
+-------+           +-------+
|       |           |       |
|   B   +-----------+   C   |
|       |           |       |
+-------+           +-------+
```

- **Autonomous system A has the following policy:**

```
import proto igp preference 10 ;
import proto bgp|ibgp aspath .* C .* preference 100 ;
import proto bgp|ibgp aspath .* B .* preference 110 ;
```

- **Autonomous system B has the following policy:**

```
import proto igp preference 10 ;
import proto bgp|ibgp aspath .* A .* preference 100 ;
import proto bgp|ibgp aspath .* C .* preference 110 ;
```

- **Autonomous system C has the following policy:**

```
import proto igp preference 10 ;
import proto bgp|ibgp aspath .* B .* preference 100 ;
import proto bgp|ibgp aspath .* A .* preference 110 ;
```

An Example Continued

- Note that AS B is using a route to X with the path
  `< A >`.
- Note that AS C is using a route to X with the path
  `< B A >`
- Consider what happens if A loses its route to X:

```
          +-------+
          |       |
          |   A   |
          |       |
          +-------+
          /       \
         /         \
        /           \
+-------+           +-------+
|       |           |       |
|   B   +-----------+   C   |
|       |           |       |
+-------+           +-------+
```

- **A picks its preferred fallback to X, through C. This route has the path:**
  `< C B A >`
- **A readvertises its choice to B. B now has the path:**
  `< A C B A >`
- **B readvertises its choice to C. C now has the path:**
  `< B A C B A >`
- **C readvertises its choice to A. A now has the path:**
  `< C B A C B A >`
- ... and so on.

Things To Observe About BGP Loops

- Routing loops which form in BGP are always transient. That is, when a loop has formed BGP is required to send protocol messages to update neighbours, whose BGPs in turn are required to send protocol messages to update their neighbours, and so on, as long as the loop persists.
- **The problem with having no constraint is that there is no termination condition for the routing loop. It is a "permanent" transient.**
- **A loop may be detected by looking for "cycles" in the AS path, or alternatively by having each router which handles the route look for its own AS in the path.**
- **A condition sufficient to terminate all loops is that a router use no route in which the local AS appears *N* times.**
- **The BGP spec choses *N* to be 1 (i.e. 1-cycle loop termination) in all cases. This has the effect of minimizing the types of route decisions which can lead to transient loops, as well as minimizing the length of any transient which occurs.**
- **The question which remains is, are there circumstances where using an *N*-cycle loop termination condition, with *N*>1, can provide useful results? In particular, will using a route with your AS in the path once always lead to routing loops?**
- **Also notice that one doesn't need the routing loop constraint built into BGP. For the 3 AS's involved the following configuration would have implemented 1-cycle loop termination:**

```
For A: import proto bgp aspath .* A .* preference none ;
For B: import proto bgp aspath .* B .* preference none ;
For C: import proto bgp aspath .* C .* preference none ;
```

Another Example: Partition Healing

- Consider the following topology:

```
  Other ASes                        Other ASes
      |                                  |
+-----+----------------------------------+-----+
|                                              |
|                   AS 200                     |
|                                              |
+-----+----------------------------------+-----+
      |                                  |
      |                                  |
+-----+----------------------------------+-----+
|                                              |
|                   AS 690                     |
|                                              |
+---------------------+------------------------+
                      |
                 Other ASes
```

- **Suppose AS 690 has the following configuration:**

```
# prefer interior routes
import proto igp preference 10 ;

# route loop detection.  One-cycle termination
import proto bgp aspath .* 690 .* preference none ;

# routes received directly from external local
# peers are preferred
import proto bgp aspath . .* preference 100 ;
import proto ibgp aspath . .* preference 110 ;
```

Another Example Continued

- Same topology:

```
   Other ASes                        Other ASes
      |                                  |
+-----+----------------------------------+-----+
|                                              |
|                  AS 200                      |
|                                              |
+-----+-----------------------------------+----+
      |                                   |
      |                                   |
+-----+-----------------------------------+----+
|                                              |
|                  AS 690                      |
|                                              |
+----------------------+-----------------------+
                       |
                   Other ASes
```

- Suppose AS 200 has the following configuration:

```
# basic 2-cycle loop termination
import proto bgp aspath .* 200 .* 200 .* preference none;

import proto igp preference 10 ;

# accept AS-200-originated routes from AS 690, but with a less
# attractive preference than via internal paths
import proto bgp aspath 690 200 preference 20 ;

# external routes received by other routers in AS have at
# least one AS in the path, and are preference 110.  Also
# accept such routes from AS 690 at the same preference
# (shortest AS path breaks preference ties)
import proto ibgp aspath . .* preference 110 ;
import proto bgp aspath 690 200 . .* preference 110 ;

# for anything other than the above, one-cycle loop terminatic
import proto bgp aspath .* 200 .* preference none ;

# external routes received directly from local peers are
# preference 100
import proto bgp preference 100 ;
```

Another Example Yet Again

- Consider what happens if AS 200 partitions:

```
   Other ASes                        Other ASes
      |                                  |
+-----+-------------+  +-----------------+----+
|                   |  |                      |
|      AS 200       |  |      AS 200          |
|                   |  |                      |
+-----+-------------+  +-----------------+----+
      |                                  |
      |                                  |
+-----+----------------------------------+----+
|                                             |
|                  AS 690                     |
|                                             |
+----------------------+----------------------+
                       |
                   Other ASes
```

- AS 200 internal routing will heal itself via AS 690.
- Both partitions of AS 200 will retain reachability to the rest of the Internet, with routes possibly passing though both partitions.
- Loop termination for AS 200 through AS's other than 690 is 1-cycle, like the standard.
- Loop termination through AS 690 is 1.5-cycle. All loops will be terminated here too.
- If AS 690 avoids sending AS 200 routes back to AS 200 through the same border they were received from, routing will be as free from loops as with 1-cycle termination.
- The comment in the spec is wrong. This is an example of where using routes with the local AS in the path does not necessarily lead to routing loops.

How Many Cycles?

- A routing loop will form if a router selects a route which leads back to itself. Note that this is "router", not "autonomous system".
- Selecting a route with the local AS in the path will certainly lead to a loop if the AS consists of only one router. If the AS consists of m routers then setting $N>m$ will provide no useful results.
- The most appropriate setting for N is the number of AS partition fragments one would like to heal. Only setting $N>1$ for those routes most likely to provide AS partition healing will improve performance.
- The big cost of setting $N>1$ is that it can potentially lead to transient loops which require a large number of updates to terminate. These updates will be visible globally, and will cost everyone's routing. Indiscriminant use of multicycle loop termination would not have a pleasant effect.
- Careful use of such a facility, however, could allow a knowledgable user familiar with the topology to configure AS partition healing via known paths with no additional exposure to routing loops at all. The tradeoff is that giving these knobs to a user who knows what he is doing can potentially allow users who don't know what they are doing to do terrible things.

Changes to the Draft

- In section 6.3, remove the check that an AS appears in the AS path only once.
- In section 9.3, remove the constraint against using a route with the local AS in the path and replace it with something like:

  "If the number of occurances of the local AS in the AS path is greater than or equal to the AS cycle limit the route must not be used. The AS cycle limit must be set to 1 by default, can optionally be increased by configuration on a per-peer or per-route basis, but must never exceed the number of AS partition fragments the topology can reasonably sustain".

- Modify the aggregation procedures such that multiple occurances of an AS in the path of a route being aggregated are reflected in the aggregate's path.
- Modify the constraint in section 5.1.3 on advertising your neighbour's address as the next hop to read something like:

  "A BGP speaker must mark as unusable routes advertised by a neighbour which have one of the receiving speaker's interface addresses as a next hop."

IDRP Implementation Status

Dave Jacobson (IBM Corp.)

Schedule

- By Year End
  - Code Complete + Some Testing
- Early Winter
  - Internal System Test
- Late Winter - Early Spring
  - Available for Interoperability Testing

Participants

- Yakov Rekhter

  IBM Research
- Rob Coltun

  RainbowBridge Communications
- David Jacobson

  IBM Network Routing Systems

Prototype Implementation

- Standalone IDRP
- Integrated IP and ISO Routing
- Runs over IP or CLNP
- Loosely Coupled To GateD
- Functions
  - Basic Transport
  - Empty RIB-Att
  - Confederations
  - Policy
  - Aggregation

Selecting an Indirect Provider

John Scudder jgs@merit.edu
Yakov Rekhter yakov@watson.ibm.com
Jessica Yu jyy@merit.edu

Current Inter-Domain Routing

- Internet as a set of arbitrarily interconnected ASes (Domains).
- Taxonomy of ASes:
  - Direct Service Providers (e.g. Merit/MichNet)
  - Indirect Service Providers (e.g. ANS)
  - Service Subscribers (e.g. U-M)
- BGP/IDRP for Inter-Domain Routing
- Hop-by-hop forwarding (with well-known problems)

Limitations of Current Routing

- Choices of routes available to a subscriber are limited by routes selected by its direct providers.

Example:

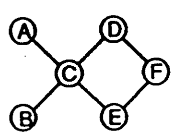

the Routing Fish

- C has to select either D or E as its next-hop domain to reach F.
- A and B are restricted to use C's choices.
- Unimplementable in current routing: C prefers to reach F via D, but B prefers E (i.e. B wants to use E as its *indirect provider*.)

So What's the Problem?

- How does a subscriber learn what routes its indirect provider has? (If the direct provider hasn't already selected those routes.)
- If a subscriber wishes to route through its indirect provider, how does it ensure that its packets are routed the way it wants? ("Consistent forwarding")

Short Answers

- Learning routes: Use BGP/IDRP
- Consistent forwarding: Use tunneling

Learning Routes From Indirect Provider

- Establish direct BGP connection with indirect provider
  - By removing "common subnet" restriction for external neighbors (requires some programming), or
  - Peer through a tunnel (can do this now)
- Routes learned this way ("indirect routes") are indistinguishable for purposes of route selection/distribution from routes learned from a directly-connected provider ("direct routes").

Providing Consistent Forwarding

- Hop-by-hop won't work — need to override direct provider's choices.
- Encapsulate (tunnel) to the indirect provider's BIS.
- In principle, any encapsulation will do (GRE, SDRP, EON...)

Refinements

- Use NEXT_HOP to denote some other decapsulator — this is just an extension of current NEXT_HOP semantics.
- Use AS_PATH of the route to the indirect provider to ascertain actual path to the provider's BIS (but wait to hear about potential pitfalls of this).

Mechanics

With off-the shelf routers:

- Participating routers (i.e. subscriber and indirect provider) set up a tunnel between themselves.
  - cisco 9.21 can do this, maybe others too
- Peer and forward packets through this hand-configured tunnel
  - Tunnel is just a logical point-to-point link
- No need to relax "same subnet" rule since routing protocol sees a directly-connected router on the other end of the tunnel.
- May do bad things when the topology changes (more on this later).

With (more) refinements:

- Relax "same subnet" rule
- Change next-hop semantics a bit
  - Now: "Forward packets via this next hop."
  - New: "If on a common subnet, forward via this next hop, else encapsulate to this next hop."
- Requires change to BGP and forwarding engine.
- Removes requirement to manually configure tunnel.

Example

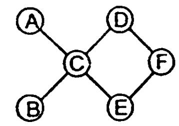

Policies:

Subscriber	Preferred path to F	Fallback to F
A (indirect)	D	E
B (indirect)	E	D
C (direct)	D	E

- A and C have same policy, so works with plain old BGP
- B establishes tunnel to E and prefers routes received via that session

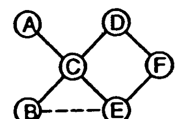

logical topology

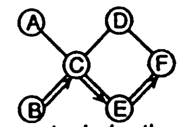

physical path

- Fallback: C-E goes down, B sees E's AS path change to (C, D, F, E) so it tears down tunnel to E and falls back to normal routing (i.e. using D as "indirect provider")

Perils and Pitfalls

- Stupid routes
- Tunnels within tunnels (within tunnels...)

Stupid Routes

This is fine:

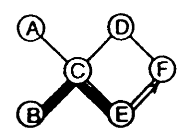

But it turns gross when C-E goes down:

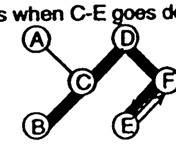

- The problem is that dynamic routing maintains the tunnel to E even when it becomes stupid: Tunnels are too robust.

Possible solution:

- B monitors AS path to E: If AS path to E ever includes F (or gets too long, or changes, or... local policy choice) then tear down tunnel and peering session.
- This requires that AS path to indirect provider reflect the real topology, so it can't be tunneled (solution for this in a minute)

Conclusions

- Provide new functionality with off-the-shelf components:
  - No new routing protocols
  - Any encapsulation suffices
  - Can be deployed today with cisco 9.21 (and others?)
  - Somewhat brittle without deployment rules/ coordination or small changes to protocols
- Only small changes required to make things work "right"
- Similar to "long-distance carrier" selection in telephony
- Can directly apply to other network layer protocols (CLNP, IPng?)
- Not intended to solve AI problem
- Need input: Real-life need for this?

For more info:

draft-rekhter-select-providers-00.txt

2.6.2 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@ossi.com
To Subscribe: mobile-ip-request@ossi.com
Archive: loki.ossi.com:/pub/mobile-ip/

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.

INTERIM MEETING REPORT

Reported by Kannan Alagappan/Digital Equipment Corporation

Minutes of the IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

The MOBILEIP Working Group convened for an interim meeting on September 9 and 10 in Newark, NJ. This group is charted to develop or adopt architectures and protocols to support mobility within the Internet.

In general, the two day meeting was productive. The group reached some agreement on the major architectural issues and terminology.

The goals of the meeting were to generate a group draft, appoint an editor, and diffuse egos.

CDPD Overview

Mark Knopper presented a brief overview of the Wireless Data Market and the CDPD architecture, protocols, services. According to one source, 27% of the market will be for personal communications and 40% for mobile office. CDPD is developing open specifications for air protocol (secured), carrier interoperability, and network functionality ("ISO terminology in specification, but really IP"). A-Interface (airlink) between mobile end systems and mobile database stations, E-Interface (external network) between CDPD network and external world, and I-Interface (inter-service provider) between other CDPD service providers networks.

Volumes 3 and 5 of the specification are relevant to this working group. The MNRP protocol is derived from ES-IS and seems to be between the mobile and visitor agent. The MNLP protocol is from the visitor agent to the home agent. The MDLP protocol is for cell switching between mobile data base stations.

Mark also handed out a paper on the CDPD Engineering Plan for IP Address Allocation, draft 1.1. This paper proposes an allocation plan for IP addresses, and includes a justification and some discussion of the architecture and routing issues for the CDPD network.

User/Functional Requirements

John Penners presented his requirement analysis for Mobile IP. John described hard requirements and soft requirements. The group agreed that our solution should not preclude support for mobile segments, but mobile segments are not a hard requirement.

After some discussion, two fundamental user requirements along with a few additional soft user requirements were decided upon:

1. A mobile host shall be capable of continuing to communicate using the same IP address, after it has been disconnected from the Internet and reconnected at a different point.

2. A mobile host shall be capable of interoperating with existing hosts, routers, and services.

Additional soft user requirements:

1. Not weaken IP security. The general feeling is that there is none now. The marketing requirement is that users do not feel that Mobile IP significantly reduces their present security.

2. A Mobile host should be able to participate in IP multicast groups.

3. There should be a means of hiding mobile location information from correspondent hosts.

Most of the other requirements in John's list were grouped as criteria for evaluating our solution. Greg Bruell rearranged John's list based on a hierarchical approach with a weighted model. These are metrics by which the group will judge its solution.

- Robustness
  - Fault Isolation - the ability to isolate faults created by mobile users should be considered for both individual and group behavior.
  - Lost Packet Operation - protocols involved in supporting mobility should be able to maintain correct operations in the presence of loss of packets.
  - Robustness - support for mobile computing should provide sufficient robustness.
  - Failure Modes - failure modes, and specifically behavior in presence of partioned internet should be carefully evaluated.

- Scalability
  - Distributed Burden - a scheme for mobile computing should be sufficiently flexible with respect to its capabilities of re-distributing the burden associated with supporting mobile computing between various entities within an internet.
  - Incremental Overhead - the incremental overhead of supporting mobile computing should reflect the number of entities that benefit from mobile computing.

- Changes to Infrastructure - a scheme that supports mobile computing shall assume that changes that involve most of the components of the existing infrastructure are infeasible.
- Scalability and Robustness - scalability needs to be complemented with robustness and fault isolation.

- Security
  - Privacy of Location - a solution to mobile computing should be able to allow selective suppression of location information.
  - Security - any scheme for supporting mobile computing shall not adversely impact available security mechanisms.

- Multicast/Broadcast
  - Multicast Applications - The support for mobile computing should allow multicast applications. ability for a mobile host to join a multicast group, send and receive multicast messages must be addressed.

- Use of Resources
  - Minimize Network Resources - a scheme that supports mobile computing should attempt to minimize the use of the networking resources (e.g., bandwidth, memory on routers, CPU on routers) that are required to deal with mobility related issues.
  - Additional Equipment - a scheme that supports mobile computing should attempt to minimize the amount of additional equipment needed.
  - Cost of Resources - in addition to minimizing network resources, any scheme used to support mobile computing should be cognizant of the cost of these resources.

- Level of Mobility
  - Multiple Mobile Host - a solution to mobile computing shall be able to deal with mobile segments that contain one or more hosts.
  - Multiple Levels of Mobility - a mobile computing solution shall be able to handle multiple levels of mobility.
  - Off-line Mobility - a mobile computing solution must not prevent upper layers from achieving off-line mobility while a host becomes disconnected from the rest of an internet for a prolonged period of time.

Next, the group defined functional requirements. After some discussion on the top-down design methodology, the group decided on three functional requirements:

1. Establish and dissolve association with an attachment point.
2. Tunnel packets to a mobile host.
3. Inform other entities of mobile location.

Alan Quirt described a short-term and long-term view for mobile IP:

- Short-term (~2 years)
  - Develop a solution that essentially works.
  - Some broken IP problems.
  - Mostly plug-in, dial-in model.

- Longer term (~5 years)
  - Everything works.
  - New IPng.
  - True wireless mobility.

Architecture/Terminology

Editor's Note: Details of the discussions under this heading are available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93sep.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

- Tunneling Discussion (Encapsulation vs. Options)

- Dogleg Routing Elimination Discussion

  - There are important security issues with trying to eliminate dogleg routing. Hosts need to authenticate redirect messages for MHs.

  - Some people generally said that they would like to see the working group first produce an Internet-Draft based on dogleg routing. Once we have more experience, we can add dogleg elimination or optimal routing. Another comment was if we only wanted a solution with dogleg routing, we could have solved the problem two years ago.

  - A vote was taken : Would you support an Internet-Draft that does not address dogleg routing elimination, but only addresses the basic user requirements for mobile-ip? Yes - 9, No - 4 (a few abstained).

  - Another vote was taken : Would you support an RFC that does not address dogleg routing elimination, but only addresses the basic user requirements for mobile-ip? Yes - 8, No - 5 (a few abstained).

- Based on this tentative vote, the group decided to focus the rest of the day's discussion on getting a simple mobile IP design.

- Beaconing/Registration Discussion

  There was discussion on allowing multiple COAAs (having one COA served by multiple routers). For example, a set of routers on a subnet can act as a COAA for visiting MHs. It was agreed that a COAA should not proxy ARP for guest MHs.
  A registration proposal was discussed similar to CDPD, where an MH sends a registration message to a COAA. The COAA registers the MH with the MH's HAA, and the HAA returns a registration ack/nack message to the COAA. The COAA returns an ack/nack message to the MH. This simple protocol is designed to minimize the MH to COAA traffic. However it requires trust between the COAA and HAA.

Dogleg Eliminators Gave a Simple Dogleg Elimination Proposal

Alan Quirt and Andrew Myles put up a slide each with an analysis of dogleg elimination.

Editor's Note: Additonal information on this subject is available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93sep.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Terminology

The group agreed on the following terminology.

- Mobile Host
- 1Correspondent Host
- Ignorant Host
- Home Subnet
- Foreign Subnet
- Home Agent (was HAA/Location Server)
- Foreign Agent (was COAA/Base Station)
- Triangle Routing (was Dogleg Routing)
- Care-Of-Address (Address of Foreign Agent)

The group needs to define the following terms.

- Weak Security
- Tunnel (v)

Document Editor

Charlie Kunzinger has volunteered as editor. He does not have a stake in any proposals. He is an experienced editor and tends to have a short turn around time.

Instructions for Liaison activities (802.11)

Charlie Perkins is the liaison between The MOBILEIP Working Group and the IEEE 802.11 subcommittee. It would be useful if the 802.11 could provide an indication of MAC address when a MH switches cells. Also, if 802.11 can provide cell arrival signals and cell departure signals we may be able to exploit them.

Attendees

Kannan Alagappan	kannan@dsmail.lkg.dec.com
Gregory Bruell	gob@wellfleet.com
Stephen Deering	deering@parc.xerox.com
Daniel Duchamp	djd@cs.columbia.edu
Jari Hamalainen	jah@rctre.nokia.com
David Johnson	dbj@cs.cmu.edu
Phil Karn	karn@qualcomm.com
Mark Knopper	mak@merit.edu
Charles Kunzinger	kunzinger@vnet.ibm.com
Brian Marsh	marsh@mitl.com
Greg Minshall	minshall@wc.novell.com
Andrew Myles	andrew@mpce.mg.edu.au
Steve Parker	sparker@ossi.com
John Penners	jpenners@advtech.uswest.com
Charles Perkins	perk@watson.ibm.com
Alan Quirt	aquirt@bnr.ca
Yakov Rekhter	yakov@watson.ibm.com
Shyhtsun Felix Wu	wu@cs.columbia.edu
Ruixi Yuan	yuan@syl.dl.nec.com

CURRENT MEETING REPORT

Reported by Greg Minshall/Novell

Minutes of the IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

Thanks to Pierre Dupont for taking notes for these minutes.

Greg Minshall provided opening remarks and a brief history of the MOBILEIP Working Group.

Charlie Kunzinger gave a short presentation on the current Mobile IP Draft. A question and answer session followed the presentation.

Editor's Note: The question and answer portion of this section is available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

IMHP Draft

Andew Myles gave a presentation on the IMHP draft. Topics included:

- A definition of the MH, FA and HA elements.
- The HA configuration (i.e., HA is not necessarily a router).
- A new element, the cache agent, which keeps track of [MH, FA] bindings.
- Security (rationale for weak security).
- Home subnet communication (performance requirements, routing options).
- Notification to the prior FA.

On this final point it was mentioned that notification to the prior FA must be fast so that it does not become a black hole for packets. The protocol should allow the new FA to accept packets from the prior FA before the MH is authorized to use the new FA. The MH must inform the prior FA as soon as it moves to a new FA. A period of questions and answers followed.

Editor's Note: The question and answer portion of this section is available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Outstanding Issues

Charlie Kunzinger presented a list of outstanding issues for discussion.

- Encapsulation method. Generic or Home-grown?

  We need at least one required method. Steve Deering argued against negotiation. Tony Li mentioned there already exists an Internet-Draft on encapsulation (Generic Routing Encapsulation). Dave Johnson stated that it had a large overhead and may not be compatible with ICMP (in terms of header size). Yakov Rekhter stated that GRE was already implemented and being deployed. Steve Deering stated that generic encapsulation can be used with a reason encoding (e.g., Mobile IP host). Greg Minshall recommended that the group continue discussion on the mailing list and pick an encapsulation method later.

- Foreign Agent receives forwarded message to MH for which it has no binding. What does it do with the message? This issue was discussed at the last session.

- Should address fields be expanded to include address type and length?

  Steve said that it may depend on how often packets are sent. Dave said the protocol is IP specific, address must fit into 64 ICMP bits and Tony recommends addresses be TLV fields to support multi protocols (e.g., Mobile appletalk). No consensus was reached.

- Do we need to control the number or frequency of registration requests?

  A discussion followed on whether to allow MH to register in multiple cells (i.e., with more than one FA) and have HA duplicate messages to both FAs. Steve suggested that protocol should not disallow this, but recommended it be deferred to the advanced functionality issue list. This issue was left unresolved.

- Is there a need for a retransmission timer on a registration request by the MH?

  It was suggested that the MH be allowed to retransmit a request and that the FA could respond with an in-progress message if it is awaiting a response from the HA on a previous request for the MH.

- State diagrams in draft document?

  This will be included in the next revision.

- Should the protocol allow a hierarchy of HA?

Should not preclude this option in draft.

- Can TOS bit in IP header be used to identify mobile hosts?

  Dave stated that RFC 1122 suggests this is not possible.

- Why can an FA terminate service to an MH? Also, HA can deregister MH.

  It was suggested that there is no need to include FA to MH deregistration since it will time out eventually.

- Several comments were made on the style, packet format and byte alignment in the draft.

- Should ICMP or UDP be used for registration protocol?

  After some discussion, a poll was taken on the preferred method and UDP was selected by a majority of those responding.

- Weak security: definition needs to be included in the draft.

- To what degree do we break the subnet model?

  This is similar to the problem with large public data networks (e.g., ATM). Yakov volunteered to communicate to the IAB how Mobile IP will break the subnet model (and write an Internet-Draft?).

Cache Agent Model

A discussion on the pros and cons of the intermediate Cache Agent model followed, with no consensus being reached on how to proceed. Some argued it should be left out of the initial draft while others argued the group should continue with plans to merge IMHP into the draft.

Documentation and Implementation Milestones

The group needs a specification which can be used to implement test systems (would like the specification before Christmas). Charlie will continue work as the document editor.

Interim Meeting

An Interim meeting of the Mobile IP Working Group was proposed for January at Xerox PARC. It was suggested that implementors and specification writers convene for two days.

Attendees

Kannan Alagappan	kannan@dsmail.enet.dec.com
Kenneth Albanese	albanese@icp.net
Nick Alfano	alfano@mpr.ca
Stephen Batsell	batsell@itd.nrl.navy.mil
Tom Benkart	teb@acc.com
Mark Beyer	beyer_mark@tandem.com
Ram Bhide	ram@nat.com
Steven Blair	sblair@us.dell.com
Jon Boone	boone@psc.edu
Monroe Bridges	monroe@cup.hp.com
Glen Cairns	cairns@mprgate.mpr.ca
Ken Carlberg	Carlberg@cseic.saic.com
Lida Carrier	lida@apple.com
Bill Cash	cash@bangote.compaq.com
Bilal Chinoy	bac@sdsc.edu
Frank Ciotti	frankc@telxon.com
David Clark	ddc@lcs.mit.edu
Thomas Coradetti	tomc@digibd.com
Stephen Deering	deering@parc.xerox.com
Thomas Dimitri	tommyd@microsoft.com
Waychi Doo	wcd@berlioz.nsc.com
Avri Doria	avri@locus.com
Robert Downs	bdowns@combinet.com
Pierre Dupont	dupont@mdd.comm.mot.com
Julio Escobar	jescobar@bbn.com
Craig Fox	craig@ftp.com
Richard Fox	rfox@metricom.com
John Garrett	jwg@garage.att.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Ramesh Govindan	rxg@thumper.bellcore.com
Darren Griffiths	dag@ossi.com
Robert Grow	bob@xlnt.com
Regina Hain	rrosales@bbn.com
Jari Hamalainen	jah@rctre.nokia.com
Marc Hasson	marc@mentat.com
Cornelius Healy	con@icp.net

Juha Heinanen	juha.heinanen@datanet.tele.fi
Kathryn Hill	khill@newbridge.com
Robert Hinden	hinden@eng.sun.com
Kevin Jackson	kjackson@concord.com
David Jacobson	dnjake@vnet.ibm.com
B.V. Jagadeesh	bvj@novell.com
David Johnson	dbj@cs.cmu.edu
Timo Jokiaho	timo.jokiaho@ntc.nokia.com
Rick Jones	raj@cup.hp.com
Elizabeth Kaufman	kaufman@biomded.med.yale.edu
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Mark Knopper	mak@merit.edu
Tony Li	tli@cisco.com
Tracy Mallory	tracym@3com.com
Wayne McDilda	wayne@dir.texas.gov
Marjo Mercado	marjo@cup.hp.com
Greg Minshall	minshall@wc.novell.com
William Miskovetz	misko@cisco.com
Randy Miyazaki	randy@lantron.com
Robert Moose	rmoose@gateway.mitre.org
Sandra Murphy	murphy@tis.com
Andrew Myles	andrew@mpce.mg.edu.au
Erik Nordmark	nordmark@eng.sun.com
Masataka Ohta	mohta@cc.titech.ac.jp
Todd Palgut	todd@nei.com
Steve Parker	sparker@ossi.com
Ismat Pasha	ipasha@icm1.icp.net
John Penners	jpenners@advtech.uswest.com
Charles Perkins	perk@watson.ibm.com
Wayne Peters	waynep@telxon.com
Ram Ramanathan	ramanath@bbn.com
Jim Rees	Jim.Rees@umich.edu
Yakov Rekhter	yakov@watson.ibm.com
Mike Ritter	mwritter@applelink.apple.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Greg Ruth	gruth@gte.com
Richard Schmalgemeier	rgs@merit.edu
Martin Schulman	schulman@smtp.sprint.com
Dallas Scott	scott@fluky.mitre.org
Isil Sebuktekin	isil@nevin.bellcore.com
Michael See	mikesee@vnet.ibm.com
Satya Sharma	ssharma@chang.austin.ibm.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Henry Sinnreich	hsinnreich@mcimail.com
James Solomon	solomon@comm.mot.com

Michael St. Johns	stjohns@arpa.mil
Martha Steenstrup	msteenst@bbn.com
Robert Stevens	robs@join.com
David Stine	dsa@cisco.com
John Tavs	tavs@vnet.ibm.com
Fumio Teraoka	tera@csl.sony.co.jp
Susan Thomson	set@bellcore.com
Akihiro Tominaga	tomy@sfc.wide.ad.jp
Paul Traina	pst@cisco.com
Hoe Trinh	htrinh@vnet.ibm.com
Keisuke Uehara	kei@cs.uec.ac.jp
John Veizades	veizades@ftp.com
Gerry White	gerry@lancity.com
Steve Willens	steve@livingston.com
Bradley Wilson	wilson@ftp.com
David Woodgate	David.Woodgate@its.csiro.au
Richard Woundy	rwoundy@vnet.ibm.com
Honda Wu	honda@nat.com
Jean Yao	yao@cup.hp.com
Weiping Zhao	zhao@nacsis.ac.jp

Charles A. Kunzinger
IBM Corporation

Basic Registration Model

- Given: mobile host knows IP Address of prospective agent
- Given: mobile host knows IP Address of its Home Agent

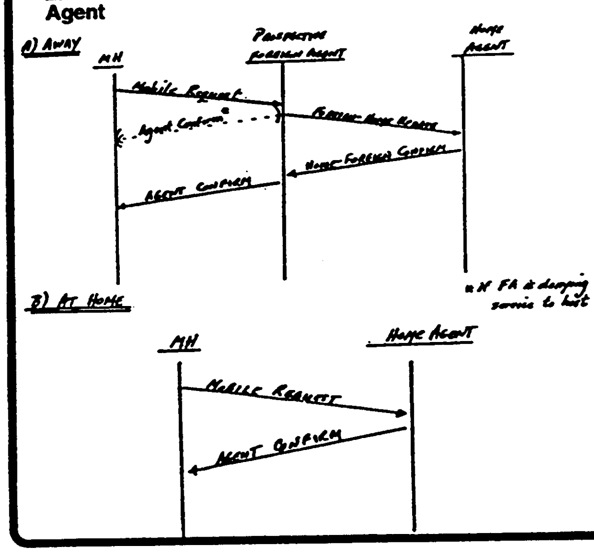

Keeping a Registration Alive

- Home Service Lifetime:
  - value set by Home Agent for a given mobility binding
  - If Home Agent has not received a re-registration request from Foreign Agent before expiration, then Home Agent will de-register—that is, delete the mobility binding
- Foreign Service Lifetime:
  - value set by Foreign Agent for a given mobile host
  - If Foreign Agent has not received a re-registration request from mobile host before its expiration, then Foreign Agent will stop serving that host and will inform the Home Agent

Notes

- mobile host knows both lifetimes
- re-registration
  - mobile host issues Mobile Request using same Registration Number
  - Foreign Agent forwards Foreign-Home Update to Home Agent using same Registration Number as in Mobile Request

Notifying A Prior Foreign Agent

Home Agent must notify the Foreign Agent named in an active mobility binding whenever that binding is deactivated at initiation of Home Agent—for example, because Home Service Lifetime expired. Message is "Home-Foreign Confirm" with appropriate operation code.

** No notification back when Foreign Agent requests de-registration **

New Foreign Agent may notify the prior Foreign Agent named in Mobile Request for a new registration. Message is "Prior Foreign Update." For duration of "Forwarding Hold Time," the old Foreign Agent may forward (via encapsulation) any packets that it receives for the mobile host.

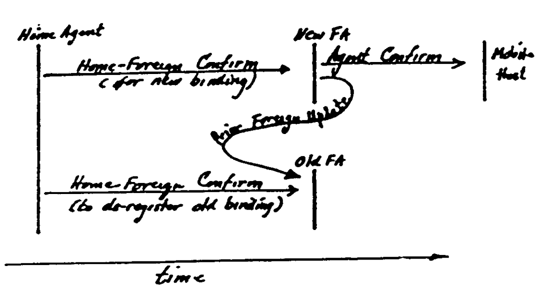

Message Summary

Sender \ Receiver	Mobile Host	Home Agent	Current Foreign Agent	Prior Foreign Agent
Mobile Host			Mobile Request	
Home Agent	Agent Confirm		Home-Foreign Confirm	Home-Foreign Confirm
Current Foreign Agent	Agent Confirm	Foreign-Home Update		Prior-Foreign Update
Prior Foreign Agent				

Items For Discussion

1. Encapsulation - should we specify a method? if so, "home-grown" or existing?

2. Foreign Agent - gets packet addressed to itself, decapsulates, finds no current binding. what to do?

   - if unexpired forwarding Hold Timer for destination mobile host, encapsulate and send to current Foreign Agent

   - else, forward inner packet normally (it should return to its Home Agent)

   *** (?) notify "encapsulator" of bad binding (?) ***

3. Expand address fields - beyond 4 bytes

4. Controls on number/frequency of registration requests

   - count (?)
   - timers (?)
   - max rate to given Foreign Agent (?)

5. How is a registration known to have failed (in absence of explicit response)?

   - timer
   - "Pending" message from Foreign Agent

Items for Discussion (continued)

6. FA-HA loop

   Cause: Home Agent (HA) thinks <Foreign Agent (FA),Host> is valid binding, but FA does not

   Symptom: FA gets packet, decapsulates, sees no binding, so sends to mobile host -> arrives at HA, Home Agent encapsulates, sends to FA...

   **** Should Foreign Agent send "Bad Binding" to the encapsulator?

7. Want state diagrams

8. Agent hierarchy - a "Summarizing Agent" for a complete Routing Domain

9. Can we mark a data packet as having come from a "mobile-aware" system?

10. Beaconing

    - limit on rate?
    - at IP or DLC layer?
    - how to carry "multicast of subnetwork scope" in IP packet?

11. Ad hoc network - IP layer or LAN (i.e., 802.11 style)

Security Information Flow

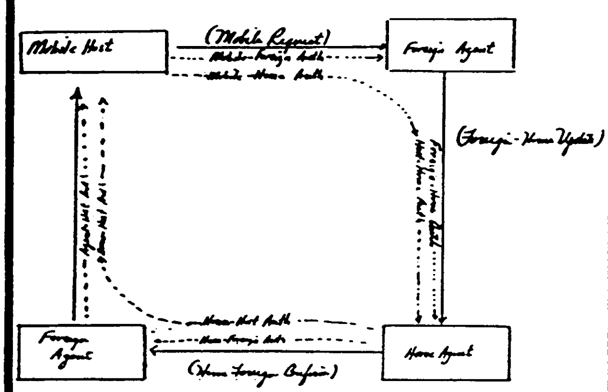

2.6.3 IS-IS for IP Internets (ISIS)

Charter

Chair(s)

Ross Callon: rcallon@wellfleet.com
Chris Gunner: gunner@dsmail.lkg.dec.com

Mailing Lists

General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu
Archive:

Description of Working Group

The 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

"Further Integration of IS-IS; Appletalk, IPX, and Other Protocols", 06/25/1993, R. Perlman, C. Gunner <draft-ietf-isis-atipx-00.txt>

"Routing over Nonbroadcast Multiaccess Links", 07/07/1993, R. Perlman, C. Gunner <draft-ietf-isis-nbma-00.txt>

"Multiple Levels of Hierarchy with IS-IS", 08/09/1993, R. Perlman, C. Gunner <draft-ietf-isis-multilevel-routing-00.txt>

Request For Comments

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

CURRENT MEETING REPORT

Reported by Dave Katz/cisco Systems

Minutes of the IS-IS for IP Internets Working Group (ISIS)

CLNP Multicast

Radia Perlman gave a presentation on several methods of achieving CLNP multicast functionality without significantly changing unicast routing.

One use of multicast is for service location. This could be achieved in CLNP by using a well-known system ID for each service and advertising it via unicast routing (from multiple locations). This is also known as "anycast."

In the very branchy case, one could use NSAP addresses with the multicast bit set in the system ID. Then a single spanning tree per area could be created for all multicast destinations and data packets could be delivered to all subnetworks in the area along this spanning tree. Since it is very branchy, no pruning mechanisms are necessary (making it simple). Inter-area forwarding could be accomplished by simply sending one packet per destination area (normal unicast routing would take place until the packet arrives in its destination area), or by creating a special area address that means "all areas," calculating a single spanning tree over all areas, and delivering the packet to the entire domain.

For very sparse multicast with a limited number of participants per group, a scheme such as CBT would be appropriate. One suggested optimization was to tunnel data packets between the routers at the branching points; this eliminates the need to set up state through the non-branching routers.

IS-IS Over Non-Broadcast Multiaccess (NBMA) Networks

An Internet-Draft has been published on IS-IS over non-broadcast multiaccess (NBMA) networks; the Routing Over Large Clouds Working Group (ROLC) was due to take up the subject during this IETF meeting.

The Integrated IS-IS Specification

The latest version of the Integrated IS-IS specification Internet-Draft (the changes to which were essentially reference updates) has expired. The protocol cannot continue along the standards track until RFCs describing operational experience and an analysis of the protocol are published. Chris Gunner is working on those documents.

Integrated IS-IS for IPX and Appletalk

The ISIS Working Group charter needs to be revisited in order that it be broadened to include other topics, such as Integrated IS-IS for IPX and Appletalk.

Radia gave a presentation on Integrated IS-IS for IPX and Appletalk. The significant issues include encapsulation (in order to carry data packets through islands that do not provide native forwarding service for the protocol), metric translation, route propagation between areas (and between protocols), clustering issues for protocols with small address spaces (such as the creation of addressing domains for Appletalk), the use of tunnels to carry routing information between areas when the level 2 subdomain does not understand the protocol, and various issues specific to the protocols (zones, services, etc.). An Internet-Draft has been published on the subject.

Radia described the new Novell link state protocol for routing IPX, Netware Link Services Protocol (NLSP). NLSP is basically equivalent to IS-IS, and uses compatible packet formats. There are several minor improvements, most of which could be realized in IS-IS without actually changing the protocol definition. There are also several changes deemed necessary for using the protocol in an IPX environment such as the inclusion of service advertisements, and running the protocol over IPX rather than over the data link (though it was observed by some present that this is unnecessary). Several people expressed the view that it would be possible to implement IS-IS and NLSP from a single code base without too much difficulty.

IS-IS Enhancements

Discussion turned to enhancements and changes that the group might like to pursue with IS-IS.

Editor's Note: A list of these enhancements and changes is available via FTP or mail server from the remote directories as /ietf/isis/isis-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Dave Katz agreed to write up a draft for increasing the LSP number, pseudonode ID, and link metric fields, as well as a proposal for how to transition to new versions of the various packets (necessary because the changes are incompatible with the existing packet formats).

Attendees

Kenneth Albanese	albanese@icp.net
William Barns	barns@gateway.mitre.org
Ross Callon	rcallon@wellfleet.com
Chris Gunner	gunner@dsmail.lkg.dec.com
Herluf Hansen	hha@tbit.dk

Susan Hares	skh@merit.edu
Marc Hasson	marc@mentat.com
Denise Heagerty	denise@dxcoms.cern.ch
Cornelius Healy	con@icp.net
Robert Hinden	hinden@eng.sun.com
Phil Irey	pirey@relay.nswc.navy.mil
Matthew Jonson	jonson@ddn.af.mil
Akira Kato	kato@wide.ad.jp
Dave Katz	dkatz@cisco.com
Tony Li	tli@cisco.com
David Marlow	dmarlow@relay.nswc.navy.mil
Randy Miyazaki	randy@lantron.com
Doug Montgomery	dougm@osi.ncsl.nist.gov
Dennis Morris	morris@altair.disa.mil
Ismat Pasha	ipasha@icm1.icp.net
Alex Reijnierse	a.a.l.reijnierse@research.ptt.nl
Allen Rochkind	Allen_Rochkind@3com.com
Greg Ruth	gruth@gte.com
Dallas Scott	scott@fluky.mitre.org
Frank Solensky	solensky@ftp.com
Thuan Tran	thuan@xylogics.com
Scott Williamson	scottw@nic.ddn.mil
Jean Yao	yao@cup.hp.com

2.6.4 Inter-Domain Multicast Routing (IDMR)

Charter

Chair(s)

Tony Ballardie: A.Ballardie@cs.ucl.ac.uk
Paul Francis: Francis@thumper.bellcore.com

Mailing Lists

General Discussion: idmr@cs.ucl.ac.uk
To Subscribe: idmr-request@cs.ucl.ac.uk
Archive: cs.ucl.ac.uk:/darpa/idmr-archive.Z

Description of Working Group

Existing inter-domain multicast routing protocols are not scalable to a large internetwork containing very large numbers of active wide-area groups. The purpose of the IDMR Working Group, therefore, is to discuss proposed inter-domain multicast routing protocols, and put forward one (or a hybrid of several/all) as a Proposed Standard protocol to the IESG.

Several proposals have been made to date, including Core-Based Tree (CBT) multicasting, Core-Based Join (CBJ) multicasting, and Scalable Reverse Path Multicasting (SRPM). Some of the above have yet to be reviewed.

Goals and Milestones

Done	Post the Core Based Trees architecture as an Internet-Draft.
Done	Meet at IETF. All proposals must be submitted by this date. Discuss all proposals which have been submitted.
Dec 1993	Submit the Core Based Trees architecture Internet-Draft to the IESG to be published as an Informational RFC.
Jul 1994	Meet at IETF. Discuss security issues with respect to the proposed protocol(s).
Aug 1994	Post an Internet-Draft for a single protocol (which may be one of the proposals, or a combination of proposals), and an Internet-Draft serving as a protocol analysis document for that protocol (as required by RFC 1264).
Jan 1995	Submit the single protocol to the IESG as a Proposed Standard.
Mar 1995	Post an Internet-Draft for an IDMR MIB.
Jul 1995	Submit the IDMR MIB Internet-Draft to the IESG as a Proposed Standard.

Internet-Drafts

"IGMP Router Extensions for Routing to Sparse Multicast-Groups", 10/28/1993,
S. Deering, D. Estrin, D. Farinacci <draft-ietf-idmr-igmp-sparse-00.txt>

"IGMP Router Extensions for Routing to Dense Multicast-Groups", 10/28/1993,
S. Deering, D. Estrin, D. Farinacci <draft-ietf-idmr-igmp-dense-00.txt>

CURRENT MEETING REPORT

Reported by Paul Francis/Bellcore and Dino Farinacci/cisco Systems

Minutes of the Inter-Domain Multicast Routing Working Group (IDMR)

The IDMR Working Group met during the morning sessions on Tuesday and Wednesday.

Summary

- The two PIM documents (PIM = Protocol Independent Multicast, formerly ESL), dense and sparse modes, were presented and discussed. Though some details about the phase shift between sparse and dense mode need working out, the general consensus of the group is that the multiple scaling modes approach is desirable. Implementation and development of PIM will continue.

- No work was done on CBT, but a status report was given describing CBT's state of implementation (almost done). There is still interest in CBT as valuable work, either as a potential alternative to PIM (if PIM proves overly difficult), or as an experimental protocol.

- The group decided to propose a new name and charter to better reflect that the focus is no longer strictly inter-domain, but rather scaling versus quality in general. Paul Francis will generate the proposal.

First Session

The meeting started with a brief presentation by Deborah Estrin of sparse mode PIM (PIM-S). The main purpose of this presentation was to set the context for a subsequent presentation of dense-mode PIM (PIM-D) by Dino Faranaci. Dino's presentation was followed by a more detailed presentation of PIM-S by Deborah. Deborah's presentation focused primarily on the differences between the current and previous specifications.

Both specifications were well received. Dino stated that he will start an implementation of PIM-D in the next few weeks. Deborah will continue implementation and simulation of PIM-S.

There were two general concerns with this work. One was that there is little need for PIM-D in the absence of PIM-S, given that PIM-D does not do much beyond what current multicast (DVMRP and MOSPF) already do. The other concern was the phase shifting mechanism between PIM-S and PIM-D. This concern was not based on scepticism that it could work so much as simply the lack of specification and experience with it. Work and implementation on this aspect of PIM will continue.

Dino recorded the following specific comments made during the presentations:

- Dense-mode ESL comments
  - Issue with asymmetric routing even in dense-mode.
  - Interest in not requiring unicast and multicast boundary congruence.
  - Brought up issues on what policy decision for phase conversion.
  - Worried about lost packets - how long will we have black holes.
- Sparse-mode ESL comments
  - Steve Deering commented that it scales well for the number of members per group, but it is not known if it can scale based on the number of groups. Concerning sources, he feels that clearing is a concern but unicast aggregation helps. Going from shared tree to per source trees reduces linearly.
  - Van Jacobson feels that source entries can be group independent.
  - Scott Brim is concerned about masks in reachability messages.
  - Deborah Estrin is concerned about switching from shared trees to per source trees. Dave Clark suggested using an interface specific case.
  - Eric Nordmark, before the second meeting, brought up the issue of a looping problem in sparse-mode (packets appear on the Ethernet twice, then start looping).

```
          RP
         / \
        /   \
       ^     v
       A     B
       ^     v
       |     |
    -------------
        ^      v
        |      |
        C      D
       /        \
    SRC          Rec
```

These two presentations were followed by a general discussion of scaling issues in multicast. There are many types of multicast applications, ranging from groups with a single or small number of sources and a high data volume (video broadcast or conferencing) to groups with potentially many senders and receivers, but with very occasional traffic, (such as a news group). Because of this wide range of applications, a range of techniques for dealing with scaling, and a means of dynamically moving from one technique to another, is required. It was generally felt that the current work goes a long way towards improving scaling according to the number of sources, but that scaling according to the number of groups is a major problem. Van suggested the use of a tree shared by multiple groups as one approach to scaling according to the number of groups.

Second Session

The agenda for this session was to cover two specific ideas for improved scaling by source. One was Deborah's idea of the use of source masks in PIM-S. The other was Van's idea of having border routers proxy for sources outside of a routing domain so that scaling in the domain is according to the number of border routers instead of the number of true sources.

As Deborah was unable to attend, the presentation of source masks was given by Dino. It was generally felt that source masks increased complexity and introduced a scaling problem of their own, without significantly improving the scaling problem it tries to address. However, since Deborah was not there to defend it, the issue remains open.

This was followed by a presentation of the proxy idea by Van. Since this presentation was not prepared in advance, and since all the details had not been worked out, no final consensus on the value of the idea could be reached. Van agreed to be responsible for seeing that a specification of the idea is written up.

Attendees

Susie Armstrong	susie@mentat.com
Jim Barnes	barnes@xylogics.com
Tony Bates	tony@ripe.net
Stephen Batsell	batsell@itd.nrl.navy.mil
Nutan Behki	nebhki@newbridge.com
Rebecca Bostwick	bostwick@es.net
Robert Braden	braden@isi.edu
Scott Brim	Scott_Brim@cornell.edu
Ronald Broersma	ron@nosc.mil
Glen Cairns	cairns@mprgate.mpr.ca
Ken Carlberg	Carlberg@cseic.saic.com
Stephen Casner	casner@isi.edu
Isidro Castineyra	isidro@bbn.com
John Chang	jrc@uswest.com
Ping Chen	ping@ping2.aux.apple.com
J. Noel Chiappa	jnc@lcs.mit.edu

David Clark	ddc@lcs.mit.edu
Michael Collins	collins@es.net
James Davin	davin@thumper.bellcore.com
Stephen Deering	deering@parc.xerox.com
Julio Escobar	jescobar@bbn.com
Deborah Estrin	estrin@usc.edu
Dino Farinacci	dino@cisco.com
William Fenner	fenner@cmf.nrl.navy.mil
Dennis Ferguson	dennis@ans.net
Robert Fink	rlfink@lbl.gov
Eric Fleischman	ericf@atc.boeing.com
Paul Francis	Francis@thumper.bellcore.com
Atanu Ghosh	atanu@cs.ucl.ac.uk
Ramesh Govindan	rxg@thumper.bellcore.com
Regina Hain	rrosales@bbn.com
Susan Hares	skh@merit.edu
Dimitry Haskin	dhaskin@wellfleet.com
Shai Herzog	herzog@catarina.usc.edu
Christian Huitema	Christian.Huitema@sophia.inria.fr
Phil Irey	pirey@relay.nswc.navy.mil
Van Jacobson	van@ee.lbl.gov
Merike Kaeo	mkaeo@cisco.com
Charley Kline	cvk@uiuc.edu
Peter Lothberg	roll@stupi.se
Tracy Mallory	tracym@3com.com
David Marlow	dmarlow@relay.nswc.navy.mil
Jun Matsukata	jm@eng.isas.ac.jp
Donald Merritt	don@arl.army.mil
Greg Minshall	minshall@wc.novell.com
Randy Miyazaki	randy@lantron.com
Doug Montgomery	dougm@osi.ncsl.nist.gov
Robert Moose	rmoose@gateway.mitre.org
Sath Nelakonda	sath@lachman.com
Erik Nordmark	nordmark@eng.sun.com
Vijayaragavan Pandian	vjp@proteon.com
Andrew Partan	asp@uunet.uu.net
Radia Perlman	perlman@novell.com
Eric Peterson	elpeterson@eng.xyplex.com
Ram Ramanathan	ramanath@bbn.com
Yakov Rekhter	yakov@watson.ibm.com
Allyn Romanow	allyn.romanow@eng.sun.com
Greg Ruth	gruth@gte.com
Henning Schulzrinne	hgs@research.att.com
Isil Sebuktekin	isil@nevin.bellcore.com
Ming Sheu	msheu@vnet.ibm.com

Keith Sklower	sklower@cs.berkeley.edu
Andrew Smith	asmith@synoptics.com
Martha Steenstrup	msteenst@bbn.com
John Stewart	jstewart@cnri.reston.va.us
Bernhard Stockman	boss@ebone.net
Larry Tepper	ltepper@compatible.com
Fumio Teraoka	tera@csl.sony.co.jp
Akihiro Tominaga	tomy@sfc.wide.ad.jp
John Veizades	veizades@ftp.com
Alice Wang	alice.wang@eng.sun.com
Chris Wheeler	cwheeler@cac.washington.edu
John Wroclawski	jtw@lcs.mit.edu
Mary Jo Zukoski	maryjo@gateway.mitre.org

2.6.5 Inter-Domain Policy Routing (IDPR)

Charter

Chair(s)

Martha Steenstrup: `msteenst@bbn.com`

Mailing Lists

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

Description of Working Group

The Inter-Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Done	Submit the IDPR Specification to the IESG as a Proposed Standard.

Internet-Drafts

"Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)", 10/29/1993, R.A. Woodburn <draft-ietf-idpr-mib-03.txt>

Request For Comments

RFC 1126	"Goals and functional requirements for inter-autonomous system routing"
RFC 1477	"IDPR as a Proposed Standard"
RFC 1478	"An Architecture for Inter-Domain Policy Routing"
RFC 1479	"Inter-Domain Policy Routing Protocol Specification: Version 1"

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

Minutes of the Inter-Domain Policy Routing Working Group (IDPR)

Summary

The IDPR Working Group met for a single session at the November IETF in Houston. The time was divided among a discussion of IDPR Version 2, an excellent video on "Routing by Preference" from Yuko Murayama, and IDPR Working Group business.

IDPR Version 2 provides policy-based multicast and multipath routing as well as policy-based resource allocation. To support the features of Version 2, the IDPR protocols required only minor changes. Specifically, the path control protocol, route generation procedure, and routing information distribution protocol underwent minor modifications. An Internet-Draft describing IDPR Version 2 will be available early in 1994.

The IDPR Version 1 software for route generation and path control required major redesign for Version 2. Ram Ramanathan and Regina Rosales Hain described the software simplifications and flexibility resulting from their redesigns. Version 2 software will begin rigorous testing in December. To obtain a copy of the Version 1 or Version 2 software, please send e-mail to `idpr-wg-request@bbn.com`.

Routing by preference, as presented by Yuko Murayama of the WIDE project in Japan, involves selecting routes that factor in both source and destination preferences, such as which domains are unfavorable and the type of route preferred through unfavorable domains. The example of "preference" presented in the video was to take the shortest (minimum delay) paths through unfavorable domains. Interested parties are encouraged to read the paper on routing by preference submitted to the list by Yuko, and to write to her about suggestions for solutions to the problem presented. The IDPR discussion list should be included in responses as the problem is of interest to the group.

There are two updated Internet-Drafts available concerning IDPR: one is the MIB and the other is the DNS support for IDPR. Comments should be submitted to the authors and to the IDPR discussion list. The IDPR Working Group Chair would like the MIB submitted as a Proposed Standard.

The IDPR Working Group will suspend for the time being. We have accomplished all that we can without an independent Version 1 implementation. We plan to resume when there exists an independent implementation of IDPR Version 1 or when the Version 2 Internet-Draft becomes available, whichever occurs first.

In the meantime, the working group mailing list will remain open to handle the unfinished business of the updated Internet-Drafts and to discuss issues relevant to IDPR.

Attendees

Nick Alfano	alfano@mpr.ca
Robert Austein	sra@epilogue.com
Anders Baardsgaard	anders@cc.uit.no
Ram Bhide	ram@nat.com
David Bridgham	dab@epilogue.com
David Conrad	davidc@iij.ad.jp
Christopher Dorsey	dorsey@es.net
Havard Eidnes	havard.eidnes@runit.sintef.no
Steve Feldman	feldman@mfsdatanet.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Regina Hain	rrosales@bbn.com
Kathryn Hill	khill@newbridge.com
Jeanine Kamerdze	kamerdze@nsipo.nasa.gov
Frank Kastenholz	kasten@ftp.com
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Jian Li	jian@rice.edu
Sandra Murphy	murphy@tis.com
Ram Ramanathan	ramanath@bbn.com
Greg Ruth	gruth@gte.com
Tae Song	tae@novell.com
Martha Steenstrup	msteenst@bbn.com
Gerry White	gerry@lancity.com

2.6.6 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", 07/26/1993, J. Moy <draft-ietf-mospf-multicast-04.txt, .ps>

"MOSPF: Analysis and Experience", 07/26/1993, J. Moy <draft-ietf-mospf-analysis-02.txt>

Request For Comments

RFC 1469 "IP Multicast over Token-Ring Local Area Networks"

2.6.7 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 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 with 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 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 the 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", 11/08/1993, S. Hares <draft-ietf-ipidrp-sip-01.txt>

CURRENT MEETING REPORT

Minutes of the OSI IDRP for IP Over IP Working Group (IPIDRP)

The minutes of the joint BGP/IPIDRP session follow the BGP charter. The attendee list below is from the joint session.

Attendees

William Barns	barns@gateway.mitre.org
Stephen Batsell	batsell@itd.nrl.navy.mil
Rebecca Bostwick	bostwick@es.net
Al Broscius	broscius@bellcore.com
Randy Bush	randy@psg.com
Enke Chen	enke@merit.edu
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Christopher Dorsey	dorsey@es.net
Dennis Ferguson	dennis@ans.net
Carlos Fernandez	carlos@plk.af.mil
Vince Fuller	vaf@barrnet.net
Vincent Gebes	vgebes@sys.attjens.co.jp
Herluf Hansen	hha@tbit.dk
Susan Hares	skh@merit.edu
Dimitry Haskin	dhaskin@wellfleet.com
Denise Heagerty	denise@dxcoms.cern.ch
Robert Hinden	hinden@eng.sun.com
David Jacobson	dnjake@vnet.ibm.com
Dale Johnson	dsj@merit.edu
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Sean Kennedy	liam@nic.near.net
John Krawczyk	jkrawczy@wellfleet.com
Charles Kunzinger	kunzinger@vnet.ibm.com
Tony Li	tli@cisco.com
Robin Littlefield	robin@wellfleet.com
Kim Long	klong@sura.net
Peter Lothberg	roll@stupi.se
Thang Lu	tlu@mcimail.com
Doug Montgomery	dougm@osi.ncsl.nist.gov
Robert Moose	rmoose@gateway.mitre.org
Dennis Morris	morris@altair.disa.mil
Sandra Murphy	murphy@tis.com
Vijayaragavan Pandian	vjp@proteon.com

Andrew Partan	asp@uunet.uu.net
Alex Reijnierse	a.a.l.reijnierse@research.ptt.nl
Yakov Rekhter	yakov@watson.ibm.com
Steven Richardson	sjr@merit.edu
Greg Ruth	gruth@gte.com
Dallas Scott	scott@fluky.mitre.org
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Paul Traina	pst@cisco.com
Chris Wheeler	cwheeler@cac.washington.edu
Cathy Wittbrodt	cjw@barrnet.net
Jessica Yu	jyy@merit.edu
Mary Jo Zukoski	maryjo@gateway.mitre.org

2.6.8 Open Shortest Path First IGP (OSPF)

Charter

Chair(s)
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.

Done Gather operational experience with the OSPF protocol and submit the document as an informational RFC.

Internet-Drafts

"The OSPF NSSA Option", 10/20/1993, R. Coltun, V. Fuller <draft-ietf-ospf-nssa-option-01.txt>

"OSPF Version 2 Management Information Base", 11/15/1993, F. Baker, R. Coltun <draft-ietf-ospf-mib-01.txt>

"OSPF Version 2", 09/20/1993, J. Moy <draft-ietf-ospf-version2-04.txt, .ps>

"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 Rob Coltun/RainbowBridge Communications

Minutes of the Open Shortest Path First IGP Working Group (OSPF)

Status Overview

The group discussed the status of the specification and of implementations. The new version of the specification is now a Draft Standard. There are several implementations of the new specification available. In the next round of OSPF interoperability there will be some testing of backward compatibility issues.

The COS testing lab is coming along. There will be a round of testing this month that will repeat some of tests for the vendors that weren't at the last round. Early next year there will be a week of OSPF stress testing.

OSPF Scaling Issues - "Ringing It Out At The Next Level"

The group reviewed a list of implementation issues that have started to appear as it starts to scale OSPF to large networks. This list will be used in part as a basis for the next round of testing at COS.

Fred Baker mentioned that ACC has one network running close to 200 routers in the same area with no stress on the routers.

On-Demand Circuit Proposal

We reviewed John Moy's on-demand circuit proposal. The major issue is the lack of backward compatibility. A few suggestions were made on how to make it compatible with existing implementations.

NSSA Implementation and Status

So far there is one implementation that will be released within the next few months in several products. The draft has been submitted to become a Proposed Standard.

MIB Changes and Status

Many excuses were given for the delay of the MIB being republished (mostly by Rob Coltun). The MIB has been updated to support NSSA and MOSPF. The group now has to do the SNMPV2 upgrades.

Attendees

Anders Baardsgaard	anders@cc.uit.no
Fred Baker	fbaker@acc.com
Nutan Behki	nebhki@newbridge.com
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Tom Easterday	tom@cic.net
Dennis Ferguson	dennis@ans.net
Vince Fuller	vaf@barrnet.net
John Gawf	gawf@compatible.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Kathryn Hill	khill@newbridge.com
Robert Hinden	hinden@eng.sun.com
Kathy Huber	khuber@wellfleet.com
Jan-Olof Jemnemo	jan-olof.jemnemo@farsta.trab.se
Akira Kato	kato@wide.ad.jp
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Charley Kline	cvk@uiuc.edu
John Krawczyk	jkrawczy@wellfleet.com
Robin Littlefield	robin@wellfleet.com
Thang Lu	tlu@mcimail.com
Dennis Morris	morris@altair.disa.mil
Zbigniew Opalka	zopalka@agile.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Paul Serice	serice@cos.com
Tae Song	tae@novell.com
George Swallow	gswallow@bbn.com
Thomas Walsh	tomw@kalpana.com

Agenda

- Status Overview
- OSPF Scaling Issues - "Ringing It Out At The Next Level"
- On-Demand Circuit Proposal
- NSSA Implementation And Status
- MIB Changes And Status

Status Overview

- New Spec Is Complete
  - At Least 4 Implementations (several products)
  - CIDR Fixes Most Important
- COS Lab Is Set Up
  - 1st Of Basic Tests Complete
  - Scaling Tests Early Next Year

- A Number Of Vendors Supporting Large Topologies

OSPF Scaling Issues

- Number Of Routers In A Single Area
  - Dijkstra Times Dominated By Routing Table Access
  - 50ms For 100 Routers Worst Case (First Time)
- Adjacencies And Buffer Mgmt Issues
  - Forming Multiple Adjacencies At Once
    - Have To Take Snapshot Of LSDB For Each

- Brain-Dead Neighbors - When To Tear Down Adjacency
  - May Keep Sending Hellos But No-one's Home
  - Buffers Are Held In Limbo
  - Convergence Time Slowed
- Rate Issues
  - Ack Timer
    - Slow Acks Hold Buffers And Slows Convergence

  - Number Of Updates Per Time Period
- Priority Handling Of Hellos
  - High Throughput Shouldn't Take Down Adjacencies
- Can Use Simulator For Tests
  - Generate Large Topologies
  - Topology Changes
  - Rate Tests
- Implementation Experience And Hints RFC

On-Demand Circuits

- Problem: Switched Circuits Cost Per Packet But...
  - Hellos Are Sent Every 30 Seconds
  - LSDB Refresh Every 30 Minutes But LSA Refresh Is A Steady Trickle
- What We Have:
  - When Link Comes Up Databases Get In Sync Immediately

- Modifying Hellos:
  - Switched Links Are Configured As Such
  - After Initial LSDB Has Been Exchanged Link Is Presumed To Be Up - No Hellos Sent
- Flooding Modifications:
  - LSAs Only Sent On Switched Link If LSA Changes
    - Check Is Done For Dijkstra Already
  - Inhibit LSA Aging

  - LSAs Received Over Switched Circuit Set DONT_AGE (Age + 0x4000)
- Lost Functionality:
  - LSA Refreshes Go Away (Low Grade Protection)
  - Databases On Both Side Of Link Look Out Of Sync
    - Checksums, Sequence Numbers Different
  - Orphaned LSAs May Remain Forever
- May Contain Switch Circuits Within An Area

NSSA Update

- Draft Moving Forward Proposed Optional Standard
- 1 Implementation So Far
  - Will Be In 5 Products Released In Next Few Months

MIB Update

- Need To Update To SNMP2
- Net Range Deprecated - Aggregate Added (includes netmask)
  - Supports CIDR And NSSA
- Multicast Support Added

2.6.9 RIP Version II (RIPV2)

Charter

Chair(s)

Gary Malkin: gmalkin@xylogics.com

Mailing Lists

General Discussion: ietf-rip@xylogics.com
To Subscribe: ietf-rip-request@xylogics.com
Archive: xylogics.com:gmalkin/rip/rip-arc

Description of Working Group

RIP Version 2 and the Version 2 MIB was approved as a Proposed Standard in January 1993. They were published as RFC 1388 and RFC 1389. 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 RFC 1264 "Routing Protocol Criteria."

This group is chartered to prepare revisions of RFC 1388, RIP Version 2, RFC 1389, the RIP Version 2 MIB, and RFC 1387, 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.

Done — Hold working group meetings to review RIP Version 2 implementations and make any changes needed to the specifications.

Done — Review the RIP over Demand Circuits Internet-Draft.

Done — Submit the RIP over Demand Circuits to the IESG for consideration as a Proposed Standard.

Done — Post as an Internet-Draft a report describing the implementation and operational experience of the RIP v2 protocol in accordance with the RFC 1264 "Routing Protocol Criteria."

Mar 1994 — Submit the RIP Version 2 protocol to the IESG for consideration as a Draft Standard.

Internet-Drafts

"RIP Version 2 Carrying Additional Information", 10/01/1993, G. Malkin <draft-ietf-ripv2-protocol-00.txt>

"RIP Version 2 MIB Extension", 10/21/1993, G. Malkin, F. Baker <draft-ietf-ripv2-mibext2-00.txt>

"RIP Version 2 Protocol Analysis", 12/08/1993, G. Malkin <draft-ietf-ripv2-protocol-analysis-00.txt>

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"

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the RIP Version II Working Group (RIPV2)

Agenda

- Administrivia
- Review Protocol Internet-Draft
- Review MIB Internet-Draft
- Review implementation experience
- Review the Demand Circuit Routing Internet-Draft
- Review the SIPP-RIP Internet-Draft
- Summary of decisions and actions

The RIP-2 Protocol Internet-Draft was approved for submission for consideration as a Draft Standard to replace RFC 1388. The MIB was similarly approved to replace RFC 1389.

There are two new implementations of RIP-2, bringing the total to four. Details on the implementations will be provided in a revision of the RIP-2 Protocol Analysis which will be done this month.

The Demand Circuit Routing Internet-Draft by Gerry Meyer was approved to submit for consideration as a Proposed Standard. The Protocol Analysis Internet-Draft will be submitted as an Informational RFC.

Consideration of the SIPP-RIP Draft, particularly the Loop Detection algorithm, was postponed until RIP-2 has been accepted as a Draft Standard (so as not to affect that effort). Discussion of the algorithm will be started next month on the ietf-rip mailing list and will be discussed in detail in Seattle.

Attendees

Steve Alexander	stevea@lachman.com
Stephen Batsell	batsell@itd.nrl.navy.mil
Akira Kato	kato@wide.ad.jp
Gary Malkin	gmalkin@xylogics.com
Sath Nelakonda	sath@lachman.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Michal Rozenthal	michal@fibronics.co.il
Tae Song	tae@novell.com
Paul Traina	pst@cisco.com
Gerry White	gerry@lancity.com

2.6.10 Routing over Large Clouds (ROLC)

Charter

Chair(s)

Joel Halpern: jmh@network.com

Mailing Lists

General Discussion: rolc@network.com
To Subscribe: rolc-request@network.com
Archive: nsco.network.com:~/rolc/rolc.arc

Description of Working Group

Summary: This group is created to analyse and propose solutions to those problems that arise when trying to perform IP routing over large "shared media" networks. Examples of these networks include SMDS, Frame Relay, X.25, and ATM.

Definition: Internetwork Layer: To avoid confusion with multiple meanings of "network" layer, we will use the term "Internetwork" layer to unambiguously refer to that layer at which IP runs. This is the layer at which IP routing functions. This is also the layer at which CLNP, Decnet, ... all run.

Large Cloud: A collection of "end-points", be that routers or hosts, connected over a fabric such that communication can be established, in the absence of policy restrictions, between any two such entities. This communication within a cloud takes place using addressing and capabilities below the "InterNetwork" layer.

The connectivity may or may not require circuit setup before communication. Such a collection is considered large if it is infeasible for all routing entities on such a "cloud" to maintain "adjacencies" with all others. Examples include, but are not limited to, ATM, Frame Relay, SMDS, and X.25 public services.

Description of Working Group:

The group will investigate the operation of IP routing protocols and services over "Large Clouds". Whenever possible, solutions shall be applicable to a range of "cloud" services. That is, the goal is a single solution applicable to multiple kinds of large "clouds", be they public or private, and independent of the specific technology used to realize the "cloud" (even a very large bridged ethernet). It is also an objective that solutions, where possible, apply to network layer protocols other than IP.

The problems the group will cover are:

A) The architectural implications of allowing direct communication between entities which do not share a common IP NET number. The group will also entertain proposals on the use of a common IP net number. If (as many believe) it is infeasible, an effort to document the difficulties will be made.

B) The routing/information protocol required to allow direct communication between two entities which were not directly exchanging routing information. This will include address resolution. The solution must couple closely to routing. It must take into account realistic connectivity policies.

C) Operation of existing protocols between peers on such clouds. Are any changes necessary or desirable? If changes are required, they will be proposed to the relevant working group.

D) Consideration of how policy restrictions and constraints (such as access control and policy-based routing paths) affect A, B, and C.

The group will also review the applicability of the work to ISDN and POTS. These technologies have a prima-facia difference, in that the number of simultaneous connections is much smaller. The implications of this for routing and relaying at the internetwork layer will need to be explored further.

Goals and Milestones

Done	Kick off meeting of group
Done	Release initial proposal for Problem B
Done	Release Internet-Draft based on discussion of proposal
Done	Meet at Houston IETF. Discuss outstanding drafts : ROLC, RIP over Demand Circuits (coordinate w/RIPV2 WG), IS-IS (coordinate w/ISIS WG), and OSPF (coordinate w/OSPF WG) and produce minuted description of what specific work is expected from the WG.
Dec 1993	Re-issue Internet-Draft on Problem B
Dec 1993	Release draft on problems using common numbering over all of a large cloud.
Apr 1994	Meet at Seattle IETF and review changes to proposals
Apr 1994	Meet at Seattle IETF. Prepare and discuss draft "analysis document"
Jul 1994	Submit base ROLC document to IESG as a Proposed Standard
Jul 1994	Submit companion analysis document to IESG

Internet-Drafts

"NBMA Next Hop Resolution Protocol (NHRP)", 10/15/1993, J. Heinanen, R. Govindan <draft-ietf-rolc-nhrp-00.txt>

CURRENT MEETING REPORT

Reported by Joel Halpern/Network Systems Corporation

Minutes of the Routing Over Large Clouds Working Group (ROLC)

Agenda - Tuesday's Session

- Charter review
- Scope of work and approaches
- Modeling, assumptions, and requirements
- Overview of ongoing work on IS-IS over NBMA networks
- Description of RIP and BGP over demand circuits

A definition of a "large cloud" was presented. The definition was taken directly from the charter. It was noted that the group summary uses the term "shared media" which some people found confusing. However, the formal definition does not, so no change was actually made to the charter. In discussion, it was agreed that a large cloud could be a broadcast network, but it was not necessarily so. Also, a cloud would normally be transitive (i.e., A↔B and B↔C connectivity implies A↔C connectivity), but special cases could arise (e.g., because of policy constraints). It includes connection-less large clouds (e.g., SMDS, or a large bridged Ethernet network) and connection-oriented large clouds with signaling (e.g., ATM, Frame-Relay with signaling, or X.25). It was suggested that, in the connection-oriented case, each entity connected to the cloud must be able to have a certain minimum number of connections (e.g., the extreme case where an entity can have only one connection open at a time is not a large cloud). Thus, POTS and possibly N-ISDN do not qualify as large clouds. (VC management was mentioned as a factor in POTS/ISDN.) It was noted that, while not large clouds, POTS/ISDN needs to be dealt with, and should borrow from this work. They do fall within the charter, but will need separate attention.

Discussion of the charter highlighted the need for the working group to strive for a general-purpose solution applicable to all types of large cloud. The solution will consider internetwork-layer(s) over, rather than between, large clouds, but will not prohibit such interworking.

Today's problems with routing over such large clouds were listed as:

1. The ability of two entities attached to the same large cloud to communicate directly when they do not have a common IP network number, in respect to both:
   (a) The operations of existing protocols between entities attached to a large cloud, and
   (b) The assumptions of routing/information protocols concerning paths between entities attached to large clouds.

2. Policy restrictions and constraints.

The operation of routing protocols over large clouds is likely to involve the aggregation of routing information. For example, a large cloud with 5000 attached routers has to have aggregation. It was agreed that the working group should aim to solve the more complex case of having MULTIPLE levels of aggregation.

There was discussion of the complication that, in the abstract, one wants to optimize the entire end-to-end routing path, not just the hops across the cloud. It was agreed that while the more general solution was desirable, this group would concentrate on the intra-cloud optimization. The further complication of trying to allow for actual "costs" for the paths across the cloud were discussed. That was felt to be more than the group could tackle.

One of the items necessary to achieve the groups goals is the relaxation of the constraints on direct communication between addresses on different IP (sub-)networks (see RFC 1122).

NBMA Networks

Chris Gunner made a presentation on the work being defined for NBMA networks for IS-IS. This involves the use of a Designated Router and both Data-Redirects and Hello-Redirects, with NBMA-addresses being extracted from inside NSAP addresses. The latter is problematic for ATM because the NBMA-address for ATM has the same syntax/structure as an (CLNP) NSAP, and thus cannot be embedded in the (CLNP) NSAP. IS-IS NBMA will reduce the number of hops across the cloud to one per IS-IS area. Thus, a ROLC solution is still needed above the IS-IS NBMA in order to obtain a single hop across multiple areas. Comments on the use of Redirects included:

- Problems with knowing that routes have changed.
- The security issue of knowing that Redirects are authentic.
- The need for timers.
- Redirects are less per-packet overhead than the short-cut routing approach of including the address of the entry-point into the cloud in the headers of each packet.
- Redirects are invoked by data packets, as opposed to the use of a separate query-response interaction (c.f., NHRP) in which the data packets and control packets can take different paths.

It was agreed that the working group needs to have a Requirements document to list both what is needed and what is not needed.

"Routing over Demand Circuits - RIP"

There was discussion of the "Routing over Demand Circuits - RIP" Internet-Draft, which seeks to avoid the need for N^2 connections between RIP entities wishing to be peers. When the exchange of routing information reaches a stable state the circuits between peers are terminated, and each peer assumes that while a circuit is down, the information contained in the last RIP update remains valid. It was observed that BGP is looking at something similar in terms of not invalidating information when it brings a demand circuit down, and not requiring keep-alives in such circumstances to maintain information validity It was suggested that a possible race-condition exists with 3rd-party announcements.

Agenda - Thursday's Session

- Discussion of draft-braden-shared-media-00.txt ('Braden draft')
- Discussion of ietf-rolc-nhrp-00.txt ('NHRP draft')
- Continued discussion of RIP over demand circuits
- Discussion of additional work
- Recruiting of editors

Joel Halpern opened the meeting, and presented the agenda. It was announced that Yakov Rekhter would present the "Braden draft" in Robert Braden's absence.

"Braden Draft"

Yakov Rekhter gave an overview of the "Braden draft"; he stated that the intention of the authors was only to stimulate discussion within the IETF, and not to make any specific proposals. The draft discusses the limitations of the current IP subnet model with respect to 'shared media' networks—i.e., networks such as ATM, Frame Relay, etc., where it is possible to have multiple subnets defined on the same medium. The current subnet model allows for direct connectivity between systems on the same medium, only if the nodes are within the same subnet—'short cut' or direct routes are precluded. This is the same problem that the ROLC Working Group proposes to solve.

The paper proposes four possible solutions to this problem:

1. Hop-by-hop redirection
2. Extended routing protocols
3. Proxy ARP mechanisms
4. Route query protocols

Yakov noted that any solution to the problem needed to meet certain criteria, including:

- Interoperability. Modified hosts and routers must interoperate with unmodified nodes.
- Practicality. Minimal software changes should be required.
- Security
- Robustness. The new scheme must be robust against errors in software, configuration, or transmission.

There was general agreement on these criteria.

There was extensive discussion of the limitations of the current model, and of the various solutions. It was noted that there were circumstances where direct routes were not desirable, or where policy constraints might preclude direct routes (e.g., to maintain firewalls, etc.). There was also some discussion about whether it was in fact optimal to have direct routes, but the consensus appeared to be that it was desirable to always have access to (and generally use) the direct path.

There was extensive and wide ranging discussion about the specific proposals, as well other issues raised by the discussion.

With respect to proposal hop-by-hop redirection, the number of re-directs needs to be limited, since some hosts might be unable or unwilling to set up direct routes. Yakov also noted that changes were required in the host software to allow them to accept redirects from routers on different subnets (refer to draft). He also discussed the Extended ARP mechanism described in the draft, whereby the redirect also contains the shared media address of the redirect router, to facilitate the ARP process.

It was noted that the extended routing protocols proposal can be viewed as an optimization of the hop-by-hop proposal, in that extended routing protocols allow a single redirect from the the first router in the path, since this has information, obtained through the extended routing protocols, about the final router, rather than having multiple redirects from each router in the route.

A problem identified with the use of redirects was that they would not work when direct routes were needed between two routers in the shared medium network (i.e., the hosts were outside the network, and could not use the information). Routers cannot (and should not, it was noted; it was generally agreed that host routes were a bad idea) listen to redirects. It was agreed that this might require that third party routing information be passed, as in BGP or EGP. Other questions included:

- What happens in the presence of aggregation? Only get a direct route to the point of aggregation.

- Are direct routes optimal? Not necessarily - only get optimal path within the routing domain, not end to end.

- Does the routing information flow across the same path as the data? Not necessarily.

- What happens to existing router to router connections if the routing information ceases? Not clear—it may be best to take all paths down. It was also noted that the problem of route partition within the cloud cannot generally be solved.

John Garrett also noted that a virtue of these proposals is that they use routing to solve a routing problem, rather than some other mechanism such as ARP (as in the NHRP proposal). It was noted that this approach may be safer, and may fit better with policy considerations.

Joel stated that he felt that the redirect proposals had limitations, but that they were worthy of consideration.

There was no discussion about the proxy ARP mechanisms proposal, since John Garrett stated that, contrary to the assertion in the "Braden draft," directed ARP had no relation at all to the model presented in the paper. Similarly, Joel noted that NHRP is much more like the route query protocols proposal rather than the proxy ARP proposal, as suggested by the draft.

There was then an extensive discussion about the route query proposal in general, and about the NHRP proposal in particular. Joel presented an overview of NHRP, noting that it proposed to use route queries in place of, or in addition to, the first packet (i.e., data forwarding could go on through the default routers, while the direct route was being found). The route respose uses the same path as the route queries, and 'cuts through' the route hierarchy. The complications of this scheme arise from its interaction with routing.

Yakov noted that this scheme can also be used to cut through from router to router, not only between hosts, since routers can send route queries.

Points raised in the discussion:

- How will policy restrictions be supported. Not all policies may be able to be supported.

- Will it be possible to discover autonomous system paths? The route response records router addresses, so it could also record AS path information.

- Why not have the egress router send the route response directly to the ingress router/host? This could be done, but requires a stronger trust model (i.e. it is more secure for response to follow same path as query).

- What if the end router cannot, or will not, accept a direct connection (policy issues, lack of connection space, etc.)? NHRP backs up from the final router to the furthest intermediate node (may be none) that is willing to accept a connection. This node could also then attempt a further cut through.

- Why is this different from router redirect? Route query is controlled by the source, hence it can better control and understand the context of the reply; this is more robust and may require less state.

- What triggers the route query? This is a complex question, and needs further thought. Maybe a query is always sent if the source is on a cloud?

- What about staleness of the route information? Periodic checks of the routes are needed.

At this point, Yakov completed his presentation by noting that it was very important to preserve backwards compatibility, and to minimize the impact on the infrastructure. He concluded, however, by stating that the traditional subnet model was causing lots of problems, not just with shared media networks like ATM, but also with mobile IP, etc. Perhaps it is time to abandon the model?

There was much support at the meeting for this sentiment. It was noted that the "Braden draft" only addressed half the current limitations of the subnet model (i.e., direct routes), but did not address the increasing problems of configuring hosts with subnet masks. Many felt that it was not appropriate for hosts to have to tackle such issues.

Joel asked whether the ROLC group should work with the "Braden draft" to generate a document to submit to the IESG to argue the case for abandoning the subnet model. Yakov responded that the MOBILEIP Working Group was already working on such a document.

NHRP Discussion

Juha Heinanen gave a more detailed overview of the NHRP proposal, and lead a discussion of issues about it. He also introduced and thanked his co-author, Ramesh Govindan. Juha noted that the NHRP route request is forwarded between next hop servers (NHS), which COULD also be next hop routers; not all routers need be next hop servers, however.

The NH servers would span the same administrative and routing hierarchy of the router network, so that end-to-end routes can be found. The NH servers have permanent connections between themselves (i.e., by PVCs, or by having configured addresses of the adjacent NH servers). He noted that this configuration information was required since a single cloud network could support multiple DISJOINT (logical) NBMAs.

Joel noted that the NH servers COULD be co-resident with the classical ARP servers.

It was noted that the NHRP proposal was still tied to the subnet model since it required "mask and match" to determine whether the last hop had been reached. It was proposed that the entire model should be abandoned, but Joel noted that this would require, in order to determine whether the last hop had been reached, that there be some kind of 'hello' or registration protocol, as with ES-IS or the classical ARP model.

Noel Chiappa stated that it would be acceptable for the router to use mask and match, as long as the hosts did not need to; Joel agreed, and noted that NHRP was at least an enabler for getting rid of the subnet model. He also noted that one virtue of the NHRP proposal was that it separated the local registration problem from that of direct route discovery—hosts could use a variety of mechanisms for the local ARP process, including redirects or simply ARPing for everything.

There was also discussion about whether the NHRP proposal should be made network layer independent, or whether it should be re-written for each protocol. There was much discussion of this topic, with the final consensus being that it made sense to spend *some* time trying to make the proposal generic, so that it would be 'easy to steal the technique'. It was noted that address registration was protocol dependent, and there was also a request that the IP specific binding be made explicit, in order to facilitate interoperability, and because the IETF only has de jure authority over IP.

Juha noted that the latest proposal has added a route record capability, in order to allow hosts to seek connections along the path to the destination, if the final router was not willing to make a direct connection. He added that not all NH servers need be routers (e.g., some could be route servers), and that only NH servers that were willing and able to forward packets need record their addresses. This comment was in the context of recording in the backwards direction. If we record going forwards, then each record must indicate whether that entity is willing to forward packets. All entities must be recorded, so that the path can be used for the response.

- Should the route be recorded going forward or backward? Going forward, since the path may not be symmetric (e.g. routing may be asymmetric, and the response should go the same way the request went.)

- What happens if packets are being forwarded (hop by hop) at the same time as the route query? Get a cascade of route queries, hence need to decide when to send a route query. Joel suggested that connection IDs should be cached and incremented for the number of packets forwarded, and that a route query should only be sent once a threshold has been exceeded. Others suggested that only the initiating host/router should send a query, but it was noted that it was very hard to determine if a router is the first hop. Another suggestion was that the host should set a bit in the PDY which would be cleared by the first hop router. No clear consensus was reached.

NHRP and Routing Protocols

There was then an extensive discussion of the interaction between NHRP and routing protocols. In particular, much discussion centered around what may happen if routes change (e.g., a better path opens up, link goes down, etc.), and whether this may lead to loops. It was noted that this problem was particularly bad in the case where the route changes occurs outside the NBMA network, but may affect a direct route. (This diagram and problem were presented by John Garrett, courtesy of work he had done on Directed ARP and Shortcut Routing.)

The discussion revolved around the following network diagram, which shows packets being sent from host H1 to host H2, both outside the shared media network, but through a direct route from R1 to R4:

```
R7 ------------- R6 -------------------R5
I                                      I
I                                      I
I     |________________________|       I
I     | Shared Media Network   |       I
I     |                        |       I
I H3--|    ooooooooooooooooooo |       I
I     |    o                  o|       I
I     |____R1____R2_____R3____R4--------I
I          I     I             I
I          I     I             I L1
I----------I     R8            I
           I     I        ]--------[
       ]-------[ I           I    I
         I       R9----------I    I
         I                        H2
         H1

Hn       - Host
Rn       - Router
ooooo    - direct (short cut) route
]----[   - Non NBMA network
```

It was stated that a problem arises if link L1 goes down - because R4 still has a path to R1, through (R5, R6, etc.), it will forward packets down that path, while R1, still having a path to R4 (the direct route), will simply loop them back to R4—*assuming* that no routing information is sent down the direct route. There was some discussion about which routing protocols would actually not detect this loop, but it was generally agreed that the problem could arise with some protocols, at least, assuming particular values of route metrics, etc.

Tony Li stated that this problem implies that the ends of the direct connection needed to be told if the current path becomes less optimal, even for apparently unrelated changes. He proposed that a host level IDRP adjacency ('mini-IDRP') be formed between the host and the first hop router, to solve this problem (details of this were apparently given in an e-mail message to the list some months ago).

It was noted, however, that it may be necessary to detail the circumstances under which the connection needs to be changed. Noel also noted that there has to be some mechanism to allow easy identification of which flows might be affected. He proposed that this was a clear argument for flows, but this did not meet with widespread approval. It was agreed that lots of state information might need to be kept.

Joel stated that the fundamental question was what changes should be noted, and who should notice them—i.e., whether or not it was sufficient for only the two end points of the direct route to notice the route change or not. In order to reduce the 'churning' of connections it was noted that only changes within a given level of route aggregation should cause a change within that level.

There was no clear resolution of these issues. John Garrett stated that the real problem was that NHRP violates the fundamental premise of (current) routing, in that the router uses a path (the one found by NHRP) different from the one it learned by routing. Routing does not talk about this path and therefore can produce inconsistency. There was agreement that this was indeed a source of difficulty. There was no agreement as to whether John's direct ARP solution, no solution at all, or a minimal exchange of routing information across the direct path were the best solution. Joel Halpern did note that the Direct ARP solution did not work with aggregation, which was an agreed ROLC requirement.

John responded that there was a need to write down the set of criteria and requirements for the ROLC work, since the aggregation requirement, for instance, was not stated in the ROLC scope. Joel agreed.

Attendees

Masuma Ahmed	mxa@mail.bellcore.com
Anthony Alles	aalles@cisco.com
Susie Armstrong	susie@mentat.com
William Barns	barns@gateway.mitre.org
Jim Beers	Jim.Beers@cornell.edu
Nutan Behki	nebhki@newbridge.com
Tom Benkart	teb@acc.com
Scott Brim	Scott_Brim@cornell.edu
Caralyn Brown	cbrown@wellfleet.com
Steve Buchko	stevebu@newbridge.com
Glen Cairns	cairns@mprgate.mpr.ca
Ross Callon	rcallon@wellfleet.com
Lida Carrier	lida@apple.com
John Chang	jrc@uswest.com
J. Noel Chiappa	jnc@lcs.mit.edu
George Clapp	clapp@ameris.ameritech.com
Robert Cole	rgc@qsun.att.com

Michael Collins	collins@es.net
Rob Coltun	rcoltun@ni.umd.edu
Thomas Coradetti	tomc@digibd.com
Matt Crawford	crawdad@fncent.fnal.gov
Michael Davis	mike@dss.com
Thomas Dimitri	tommyd@microsoft.com
Waychi Doo	wcd@berlioz.nsc.com
Ed Ellesson	ellesson@vnet.ibm.com
Robert Enger	enger@seka.reston.ans.net
Dario Ercole	Dario.Ercole@cselt.stet.it
William Fenner	fenner@cmf.nrl.navy.mil
Dennis Ferguson	dennis@ans.net
James Forster	forster@cisco.com
Craig Fox	craig@ftp.com
Dan Frommer	dan@isv.dec.com
John Garrett	jwg@garage.att.com
John Gawf	gawf@compatible.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Eugene Geer	ewg@cc.bellcore.com
Shawn Gillam	shawn@timonware.com
Fengmin Gong	gong@concert.net
Ramesh Govindan	rxg@thumper.bellcore.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Chris Gunner	gunner@dsmail.lkg.dec.com
Joel Halpern	jmh@network.com
John Hanratty	jhanratty@agile.com
Dimitry Haskin	dhaskin@wellfleet.com
Marc Hasson	marc@mentat.com
Ken Hayward	Ken.Hayward@bnr.ca
Denise Heagerty	denise@dxcoms.cern.ch
Juha Heinanen	juha.heinanen@datanet.tele.fi
Kathryn Hill	khill@newbridge.com
Robert Hinden	hinden@eng.sun.com
Refael Horev	horev@lannet.com
Kathy Huber	khuber@wellfleet.com
Melanie Humphrey	msh@uiuc.edu
David Jacobson	dnjake@vnet.ibm.com
Ronald Jacoby	rj@sgi.com
B.V. Jagadeesh	bvj@novell.com
Jan-Olof Jemnemo	jan-olof.jemnemo@farsta.trab.se
Merike Kaeo	mkaeo@cisco.com
Akira Kato	kato@wide.ad.jp
Yasuhiro Katsube	katsube@mail.bellcore.com
Hiroshi Kawazoe	kawazoe@trl.ibm.co.jp
Ted Kuo	tik@vnet.ibm.com

Sundar Kuttalingam	sundark@wiltel.com
Mark Laubach	laubach@hpl.hp.com
Tony Li	tli@cisco.com
Robin Littlefield	robin@wellfleet.com
Kanchei Loa	loa@sps.mot.com
Thang Lu	tlu@mcimail.com
Bryan Lyles	lyles@parc.xerox.com
Dan Magorian	magorian@ni.umd.edu
Tracy Mallory	tracym@3com.com
David Marlow	dmarlow@relay.nswc.navy.mil
Jun Matsukata	jm@eng.isas.ac.jp
Keith McCloghrie	kzm@hls.com
Donald Merritt	don@arl.army.mil
Dennis Morris	morris@altair.disa.mil
Robert Moskowitz	3858921@mcimail.com
Sath Nelakonda	sath@lachman.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Zbigniew Opalka	zopalka@agile.com
Maryann Perez	perez@cmf.nrl.navy.mil
Charles Perkins	perk@watson.ibm.com
Drew Perkins	ddp@fore.com
Radia Perlman	perlman@novell.com
James Philippou	japhilippou@eng.xyplex.com
Ram Ramanathan	ramanath@bbn.com
Kenneth Rehbehn	kjr@netrix.com
Yakov Rekhter	yakov@watson.ibm.com
Allen Rochkind	Allen_Rochkind@3com.com
Robert Roden	roden@roden.enet.dec.com
Benny Rodrig	brodrig@rnd-gate.rad.co.il
Shawn Routhier	sar@epilogue.com
Michal Rozenthal	michal@fibronics.co.il
Greg Ruth	gruth@gte.com
Timothy Salo	tjs@msc.edu
Hal Sandick	sandick@vnet.ibm.com
Isil Sebuktekin	isil@nevin.bellcore.com
Michael See	mikesee@vnet.ibm.com
Paul Serice	serice@cos.com
Satya Sharma	ssharma@chang.austin.ibm.com
Vincent Shekher	vin@sps.mot.com
Ming Sheu	msheu@vnet.ibm.com
Uttam Shikarpur	uttam@zk3.dec.com
W. David Sincoskie	sincos@bellcore.com
Keith Sklower	sklower@cs.berkeley.edu
Andrew Smith	asmith@synoptics.com
Tae Song	tae@novell.com

Martha Steenstrup	msteenst@bbn.com
Steve Suzuki	steve@fet.com
George Swallow	gswallow@bbn.com
Susan Thomson	set@bellcore.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Thomas Walsh	tomw@kalpana.com
Guy Wells	guyw@uswest.com
Douglas Williams	dougw@vnet.ibm.com
David Woodgate	David.Woodgate@its.csiro.au
Jessica Yu	jyy@merit.edu
Chin Yuan	cxyuan@pacbell.com

2.6.11 Source Demand Routing (SDR)

Charter

Chair(s)

Deborah Estrin: estrin@usc.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 SDRP (Source Demand Routing Protocol) as an inter-domain routing protocol capability in conjunction with IDRP and BGP inter-domain routing protocols. The purpose of SDR is to support source-initiated selection of inter-domain 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, and 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.
Done	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.

Internet-Drafts

"Source Demand Routing Policy Language", 06/21/1993, T. Li <draft-ietf-sdr-pl-00.txt>

"Source Demand Routing: Route Setup", 06/23/1993, D. Estrin, D. Zappala, T. Li <draft-ietf-sdr-route-setup-00.txt>

"BGP SDRP_SPEAKERS Attribute", 09/13/1993, K. Varadhan <draft-ietf-sdr-speakers-attribute-00.txt>

CURRENT MEETING REPORT

Reported by Deborah Estrin/Information Sciences Institute

Minutes of the Source Demand Routing Working Group (SDR)

Tony Li started the first meeting with a walk through of the current specification to collect comments. There was minimal discussion. The specification appears stable.

USC is working on their second implementation which will be available for distribution in November or December. A first prototype release was made available during the summer.

The remainder of the first session was spent discussing the agenda for the second day's meeting and some future issues, in particular, the relationship between route setup and reservation setup, and acquiring information to compute SDRP routes not made available by BGP/IDRP.

The second day's discussion addressed:

- Construction of SDRP routes

  The issue of how to evolve this functionality over time and what is a useful starting point was discussed as well as the issue of what is needed longer term: query/search techniques, extract proposal for using IDRP/BGP.

- Indirect provider extensions to IDRP

  Prior to the meeting, Yakov Rekhter proposed using SDRP to extend IDRP/BGP functionality to allow overriding of source-specific preferences. However Yakov had to attend the ROLC meeting in Braden's absence and was unable to present this in detail.

- Mapping packets to SDRP routes

  This must be done at the first SDRP speaker using local criteria and packet information. There are questions about how to specify the criteria, policy, and how to apply it efficiently so as not to downgrade performance for SDRP and non-SDRP packets.

- Open issues

  Design issues that are most open and in need of attention are longer term techniques for inter-domain and intra-domain SDRP route construction, managing feedback from SDRP route failures for use in route selection, and the possible use of anycast addresses.

Open usage/deployment issues are about the best places to start deploying SDRP (e.g., in conjunction with upcoming real time experiments, perhaps?), use for other tunneling requirements, and the need to dispel beliefs that encapsulation can never be efficient.

- Volunteers were requested to review the setup specification, finish the usage draft (Sue Hares and Peter Ford are currently responsible for this), participate in experiments (Merit and USC are currently participating), and participate in route construction design work.

- A status report on the current implementation noted that the two unimplemented features were MTU discovery and policy checking by transit domains.

Several of these issues were not discussed in adequate detail because of Deborah Estrin's sudden departure in response to the LA fires. We will return to them via e-mail discussions and during the March IETF.

Attendees

J. Noel Chiappa	jnc@lcs.mit.edu
Osmund deSouza	osmund.desouza@att.com
Deborah Estrin	estrin@usc.edu
Christian Huitema	Christian.Huitema@sophia.inria.fr
Tony Li	tli@cisco.com
Yakov Rekhter	yakov@watson.ibm.com

2.7 Security Area

Director:

- 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 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.

Following the SAAG minutes is a status report for each of the working groups 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.

Security Area Advisory Group (SAAG)

During the monday afternoon meeting, Steve Crocker led a discussion on Internet security architecture issues. Topics included application support for security, transport/network layer support for security, identification of zones of trust, and firewalls.

The discussion included the following points.

- Are firewalls the best (or at least one of the best) approaches to security? Consider that applications today expect to be end-to-end, i.e., in general they are not firewall friendly. If firewalls are a direction of the IETF/Internet, then there should be a statement of this principle so that we build all future protocols with it in mind.

- The PSRG is working on a security architecture document. It explicitly addresses end-to-end security services and mechanisms, while a hybrid approach involving firewalls appears to be more common. In any case, a discussion of firewalls would be useful.

- Firewalls were compared to the installation of burglar alarms in homes. In particular, most people install alarm systems in response to a burglarly, as opposed to installing them as a preventative measure. Is it possible that if environments (or protocols) protected themselves they might not need a firewall? Firewalls are a "convenient" security measure to install in the short-term.

In addition to the architecture discussion, Steve Kent presented a status of the PSRG security architecture document, Bill Simpson expressed a desire for additional authentication mechanisms in PPP, in particular Kerberos, and it was noted that although Triple-DES will be discussed within the PEM working group, it is applicable in many contexts and may need broader exposure.

Authorization and Access Control Working Group (AAC)

The AAC Working Group discussed a revised framework for representing privilege attributes and restrictions for distributed authorization credentials and access control list entries. The new framework addressed concerns raised at the Amsterdam IETF. Attributes are now assigned to one of three classes: privilege attributes, restrictions, and aggregates. The contents of the security context used as input to the authorization API was discussed next. The security context should be filled in initially by the authentication mechanism (e.g., GSSAPI) and might be subsequently augmented by other security mechanisms. The security context could include verified authentication and authorization information, and might separately specify unverified information and delegated credentials. The form of the arguments and return values of the authorization API was discussed next and will be further refined. To close, the group discussed the drafting of a document to provide network application developers with guidelines for supporting authorization in their systems.

Common Authentication Technology Working Group (CAT)

The CAT Working Group met for two sessions in Houston. The status of ongoing activities was reviewed, including a reworked GSS-API implementation for Kerberos V5 beta 3; this implementation, and an Internet-Draft describing its GSS-API mechanism characteristics and token formats, are scheduled to become available later this year. Some interface clarifications and extensions (e.g., a new GSS_Inquire_context primitive) were discussed as inputs to Internet-Draft successors to RFCs 1508/1509, targeting inclusion in eventual Draft Standard versions to supplant those RFCs and comprise a "Version 2" GSS-API. Related topics to be discussed further on the mailing list include multi-mechanism credential management and error reporting. Piers McMahon gave a presentation on SESAME's multi-mechanism implementation, and distributed a paper for comment. Sam Sjogren and Steve Lunt led a discussion on the FTP Security Internet-Draft, to be updated shortly and to be used as the basis for an interoperability test (using Kerberos V4 technology) planned for March 1994. Representatives from the NASREQ Working Group described their currently-contemplated architecture, as input to determining how the CAT Working Group and technology might support their needs. Ran Atkinson gave a presentation on the Internet Authentication Guidelines Internet-Draft, receiving and soliciting comment from the working group.

Internet Protocol Security Protocol Working Group (IPSEC)

There are several known experimental implementations of IP Security, two of which demonstrated their approach during the Houston IETF. Demonstrations of preliminary interoperable implementations is targeted for the next IETF. Key management is still an open issue. The implementations include:

- I-NLSP - Rob Glenn
- swIPe - Phil Karn
- KeyRing - Rob Hagens
- TANDU/Cryptonette - Charlie Kaufman
- LAN Guardian - Mike O'Dell
- SP3 - Paul Lambert

Jim Zmuda and Phil Karn gave demonstrations of their implementations. Jim's implementation was based on the ISO 11577 specification for Network Layer Security Protocol (NLSP) and used the NLSP specification of a Security Exchange Protocol (SAEP) for key management. The implementation demonstrated by Phil Karn was based on Phil's KA9Q software running on a portable computer (80386 based). This demonstration ran between Houston and Phil's home. Key management was based on the manual entry of DES key variables.

Network Access Server Requirements Working Group (NASREQ)

A very brief presentation of distributed authentication was presented as a possible future subject for the working group to consider. The possibility of changing the charter was discussed, and the following elements were described as a possible direction:

- Finish the NAS Requirements document and submit it for consideration as an Informational RFC following the Seattle IETF.

- Revise the RADIUS protocol definition and submit it for consideration as an RFC after review at the Seattle IETF.

- Move KAP/PKAP to the Point-to-Point Protocol Extensions Working Group (PPPEXT) and/or to a working group in the Security Area.

- Focus the attention of the group on distributed authentication in support of shared dialin between organizations. This will likely have other implications and should have significant support from security area folks to be successful.

Privacy-Enhanced Electronic Mail Working Group (PEM)

The meeting covered implementation status reports, discussion of potential electronic notary services, PEM-MIME integration, certificate servers, and triple-DES.

Only MIT and TIS implementation efforts were represented, but other implementations are proceeding in various countries.

Dave Solo provided a presentation on various "notary-style" validation services for non-repudiation (see slides following the PEM minutes): simple time stamping, enhanced non-repudiation, document registration, archival signature validation, assurance issues, validation of other attributes.

Two MIME-PEM designs now exist:

- MIME-PEM "lite" (Jeff Schiller)
- MIME-PEM "full-bodied" (Steve Crocker)

These two proposals have different consequences and reflect some divergence in goals within the PEM and mail communities.

Christian Huitema (INRIA) proposed a certificate server proposal for PEM to facilitate retrieval of certificates and CRLs with locally managed, simple databases. The index for search is the user's mailbox name. This calls for operators of the hosts that provide the user's mailbox to provide this responder facility. However, mail services such as CompuServ and MCIMail are unlikely to provide this service. A new record type may need to be created to allow indirection to other than the user's actual mailbox provider. Also, this proposal is based on TCP, but not all prospective PEM users are reachable by TCP, e.g., users of non-IP nets or firewall. There was a suggestion to add this facility to FINGER instead, to minimize firewall problems. It was proposed that email-based access should be baseline, with real-time access an optional additional service.

Triple-DES was discussed briefly. At issue is the best method of using triple-DES in CBC mode. Burt Kaliski had circulated a summary of his findings, but he was not present for discussion. This remains an open topic.

TELNET Working Group (TELNET) - Applications Area

There is a draft specification combining both authentication and encryption security services. Implementations of the specification will be present for the next IETF.

In addition, the draft document specifying the use of Kerberos Version 5 as a TELNET authentication option has expired. It will be resurrected.

Router Requirements Working Group (RREQ) - Internet Area

Due to lack of progress on the documents this work item was officially closed as a Security Area concern. If the working group is reconstituted to continue work then the security area will re-open this work item. In the time since the Houston IETF, this document has received renewed attention within the IESG and is now under the control of the Internet area.

Domain Name System Working Group (DNS) - Service Applications Area

The DNS Security sub-group of the DNS working group met to identify the threats, security services, and requirements of interest to the DNS. The requirements will be distributed to the mailing list for discussion until November 30, 1993. After that time, strawman proposals may be distributed until January 31, 1993. The group will evaluate all proposals with the goal of creating one proposal at the next IETF.

It was decided to create a DNS security working group. In parallel with the activities above a charter will be drafted for review and submission to the IESG.

Export Control Issues

There was some consensus that there exists rumors that at least some of the rules regarding export may change soon. This work item exists to track this activity. (In the time since the Houston IETF meeting, no changes in United States policy have been announced.)

IP: The Next Generation

An IPng Directorate has been formed to track and evaluate the IPng proposals. Steve Bellovin is a member of the directorate representing the security area.

Mobile IP Security

There is a draft document describing what is identified as a weak security mechanism and how it can be used until IP security is available to provide strong security. Phil Karn will distribute this description to SAAG for review and comment.

Random Number Generation Issues

A revised Internet-Draft was distributed for comment. It has been published as an Informational RFC.

Routing Security Plan

This work item was re-assigned to Sandy Murphy at the Houston IETF. She will prepare a draft document summarizing the authentication and integrity issues in routing for the next IETF.

2.7.1 Authorization and Access Control (AAC)

Charter

Chair(s)

Clifford Neuman: bcn@isi.edu

Mailing Lists

General Discussion: ietf-aac@isi.edu
To Subscribe: ietf-aac-request@isi.edu
Archive: prospero.isi.edu:~/pub/aac/*

Description of Working Group

The goal of the Authorization and Access Control Working Group is to develop guidelines and an Application Program Interface (API) through which network accessible applications can uniformly specify access control information. This API will allow applications to make access control decisions when clients are not local users, might not be members of a common organization, and often not known to the service or application in advance.

Several authentication mechanisms are in place on the Internet, but most applications are written with local applications in mind and no guidelines exist for supporting authorization and access control based on the output of such authentication mechanisms. The CAT working group developed the GSS-API, a common API to support authentication. The AAC Working Group will develop a common API that accepts the identity of a client (perhaps the output of the GSS-API), a reference to an object to be accessed, and optionally an indication of the operation to be performed. The API will return a list of authorized operations or a yes/no answer that can be easily used by the application.

A second, longer term purpose of the working group will be to examine evolving mechanisms and architectures for authorization in distributed systems and to establish criteria which enable interworking of confidence and trust across systems. The working group will develop additional goals an milestones related to this purpose and will submit a revised charter once the appropriate goals and milestones are determined. To the extent possible this additional work will encourage evolution toward credential formats that more readily allow support for or translation across multiple mechanisms.

Goals and Milestones

Done	Submit charter and milestones for approval.
Done	Meet at the Columbus IETF to identify common characteristics of evolving distributed authorization mechanisms and begin discussion of approaches for interoperability across mechanisms.

Jun 1993 Post draft API as an Internet-Draft.

Jun 1993 Post an Internet-Draft of the guidelines for authorization and access control for network accessible applications.

Aug 1993 Submit the AAC guidelines document for approval as an Informational RFC.

Jan 1994 Submit the AAC API for consideration as an Experimental RFC.

CURRENT MEETING REPORT

Reported by B. Clifford Neuman/Information Sciences Institute and Piers McMahon/ICL

Minutes of the Authorization and Access Control Working Group (AAC)

The 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`.

Agenda

- Presentation of a revised list of restrictions and privilege attributes needed by applications and existing security systems, and a proposed method for representing them.

- Discussion of the information maintained in the security context and where it should come from. The security context maintains information about the user that is used to make authorization decisions.

- Discussion of the intended use of the authorization attributes by applications, ACL formats for an authorization API, and discussion of an API to provide a simple interface for application developers.

- Discussion of guidelines for adding authorization to network accessible applications. These guidelines should be written and released initially as an Internet-Draft, and eventually as an Informational RFC.

The purpose of the the initial work on a framework to represent security attributes and restrictions is to provide a uniform framework for both distributed authorization and local authorization. Local authorization will be based on access control lists. Distributed authorization will be based on proxies. Both have common elements. These elements (or attributes) may be stored on access control list entries on security servers, carried in credentials, and evaluated on the end system the same way they would be evaluated if stored on an access control list maintained by the end system itself.

Presentation of a Revised List of Attributes and Restrictions

At the Amsterdam IETF, there was concern that the representation of all rights conveyed by a set of credentials as restrictions on rights possessed by the grantor of those credentials was confusing. A revised framework for representing rights and restrictions on those rights was sent to the mailing list and presented at the meeting. In the Amsterdam discussion, the use of the term certificate was also confusing. The revised framework uses the term proxy which is less confusing, though more closely tied to one particular proposed implementation of the framework.

In the revised framework for representing privileges, privileges are conveyed by proxies. Proxies enumerate positive rights they convey, but these rights are limited by the rights available to the principal that signed the proxy (the grantor). Some proxies specify that all rights available to the grantor are conveyed, supporting unrestricted delegation or authentication forwarding. Proxies may be further restricted to limit the rights that are conveyed. Certain restrictions require the presentation of additional proxies called endorsements before the original proxy may be used. A proxy and all required endorsements together form a proxy chain (a proxy that does not require endorsements constitutes a chain by itself). Endorsement may apply additional restrictions. The restrictions specified by each proxy in a chain are applied, resulting in a more restricted set of rights. A principal may collect multiple independent proxy chains, and the rights granted by each chain are added providing additional rights.

- **Attributes and Restrictions**

A structured set of authorization attributes are associated with each proxy. Each attribute in the set is typed, and the interpretation of the data associated with the attribute is determined by its type. Some types of attributes are specific to specific servers. A flags field associated with each kind of attribute encodes the appropriate behavior if the interpretation of the attributed is not known. The two possible behaviors are rejection of the request, or ignoring the attribute.

Attributes are further divided into three classes: privilege attributes, restrictions, and aggregates. The class is also encoded in the flags field.

Privilege attributes identify some operation that is permitted, or assert identifying information such as group information or user identifiers that grant additional rights to the principal that presents a proxy. Privilege attributes may be placed in a proxy only at the time it is created. They may not be added subsequently, and they may not appear in endorsements. Further, in order for a privilege attribute to apply, the principal granting the proxy must possess the ability to grant such an attribute. (Though privilege attributes grant rights, they are interpreted as explicitly enumerating the rights that are to be granted by a proxy and all other rights are presumed not to be granted by the proxy).

Restrictions specifically remove rights from those conveyed by a proxy, or they place additional constraints on when, how, where, or by whom a proxy may be exercised. Restriction may be present when a proxy is initially created, they may be added subsequently and they may be present in endorsements. Restrictions that are present in a required endorsement are applied in addition to the restrictions present in the proxy that is endorsed.

Aggregate attributes may be positive (privilege attributes) or negative (restrictions). The rules about where each may appear apply to aggregates as well. An aggregate attribute encapsulates a set of attributes and specifies how they are to be interpreted. One form of aggregate limits the application of a set of attributes to a particular end server or a particular application.

- **Discussion**

After the initial presentation several questions were raised. Among the questions was how one determines which principals are authorized to issue proxies granting particular rights. The end server will still have to maintain access control lists for this information, but these ACLs only contain the names of the principals that grant rights for an operation, rather than those authorized to perform the operation. Presumably this information changes less frequently, and is more compact.

Checking these ACLs would be handled by the authorization API, so the client does not have to worry about it directly. The contents of this ACL might be specified by a configuration option for the local system. This check might be done when establishing a security context for a request or a login session, on the remote machine. Ted Ts'o commented that the information about the grantor of a privilege should be carried along as extra information in the security context even after verification, so that an application that has special requirements can check it.

Another issue raised was concern about placing too much of the burden on the application programmer to interpret such proxies and proxy chains. In fact, the interpretation would be handled by two APIs, one of which would be specific to the distributed authorization method in use, and one would be the authorization API which was discussed later in the meeting. This latter API will take several arguments, and answer yes or no indicating whether an operation is to be allowed.

- **Enumeration of Attributes and Restrictions**

An initial set of positive and negative attributes was presented and it was discussed how they fit into the revised framework. Some of the attributes are application specific, and other application specific attributes can be defined.

The positive privilege attributes can only be grated in an initial proxy and can not be added subsequently, to augment rights. In enumerating an initial set, a goal was to cover the rights needed and used by various distributed authorization mechanisms including ECMA, DCE, and restricted proxies.

Among the positive attributes were:

LOCAL-UID	set UID for local system
GROUP-MEMBERSHIP	enumerate local groups to be used locally
DCE-PAC	globally unique groups and uuids
AUTHORIZED	encodes list of objects and rights for objects
QUOTA	like authorized, but specifies numeric limit

Other suggested positive attributes include:

ROLE	to include rights of individual in particular role
ALL	which might be needed to mean all rights of grantor

For the AUTHORIZED attribute, the encoding of the identified objects and the rights on those objects would be application specific, and opaque to other parts of the system.

Among the negative attributes were:

FOR-USE-BY-PRINCIPAL	who may use a proxy, or compound principal
FOR-USE-BY-GROUP	only members of group may use proxy
ACCEPT-N-TIMES	can only be exercised N times on a given server
VALID-TIME-OF-DAY	e.g. can only be used between 9AM and 5PM
VALID-PERIOD	expires after, not good until
NETMASK	application specific, used by network access server
FOR-USE-ON-SERVER	identifies specific server where proxy is valid
FOR-USE-FROM	local terminal, secure area, not sure how to implement

Other suggestions for negative attributes included a means to restrict the day of the week (for example, only Monday through Friday). John Linn made the observation that given that time of day is a one time decision, once logged in, one can do whatever one wants. This can be addressed if the application checks periodically, or notes the end time during the initial check and requires reauthorization at that time. It can be made easier for applications to note the end of the authorization period if an expiration time were returned as an additional value by the authorization API. Some applications would use this expiration time, while others might not. Some systems, AFS in particular, do note the expiration of credentials. Finally, though login has a one time authorization check, other operations like file accesses might not.

It was also pointed out that the negative attributes did not include the ability to exclude enumerated rights. This might be useful so that one can list a broad class of positive attributes, then enumerate exceptions as negative attributes. Thus, the restriction can be added:

EXCLUDE	excludes list of objects and rights for objects

The encoding of the objects and rights for EXCLUDE would be application specific.

It was suggested that there be a way to combine positive attributes from separate proxies into a single proxy. Unfortunately, this is clearly not possible when the rights conferred are granted by different principals. When granted by a the same principal, a proxy can be granted to enumerates multiple rights. Combining them after multiple proxies have been issued, however, requires the reissuance of the combined proxy by the grantor.

There was some discussion of how the attributes are related to the authorization API. This discussion appears later in the minutes. Due to the limited time for the meeting, the aggregate attributes were not discussed at the meeting, but are described in the message sent out to the mailing list.

Discussion of Information Maintained in Security Context

Piers McMahon noted that authorization decisions would be made on two types of application servers:

1. Servers which handle multiple requests within a single process (e.g. DCE RPC)
2. Per-request daemons (e.g.: TELNET)

He concluded that the authorization security context in the authorization API must be able to refer to either a security context (for the former case) or a delegated credential (for the latter case).

In order to make the authorization API simpler, he also suggested that the GSS-API gss_accept_sec_context should always return a credential which could then be used consistently to represent the authorization context (in an analogous way to the DCE login context) even if delegation was not enabled for this context.

Ted Ts'o warned that this might inappropriately overload the GSS-API credential, as such non-delegated "credentials" could not be used to initiate contexts. In response, Piers observed that this was also true for acceptor usage credentials, and suggested that the "usage" of the credential could be perhaps extended to include an "accepted security context" type in addition to initiator and acceptor. After some further exchanges, it was agreed that further discussion on this topic should be carried on in the Common Authentication Technology Working Group (CAT).

John Linn commented that it might be appropriate to add a primitive inquire_security_context, to to query information from the security context, and perhaps a similar primitive to add information to the security context after it is initially established.

Discussion of the Authorization API

A one-page handout was distributed that outlined the arguments to the check_authorization function in the authorization API. Some of the goals in the design of the API were that it be simple to use for simple applications, but extensible and also easily usable by applications that have additional constraints and more advanced requirements for authorization.

The arguments to the check authorization function are described below:

```
answer = check_authorization(det_answer,/* Detailed answer (out)     */
                             sc,        /* Security context          */
                             target,    /* Object to be manipulated  */
                             operation, /* Operation to be performed */
                             parameters)/* Modifiers to request      */
```

In the interest of providing a simple interface to most applications, the authorization API returns an answer of yes (0), no (2), or maybe (1). Simple applications would treat this as yes, or not yes. Further, the first argument to the API is an out parameter containing extra information that may be used by more advanced applications if the answer is a maybe. This structure might also contain information about when the authorization so granted is due to expire.

The form of the security context, the second argument, was already discussed. Ideally, information to be contained in the security context will include the following elements:

- verified authentication information - gssapi, uuids, groups
- unverified authentication information - to be checked when needed
- verified and unverified authorization information - proxies, etc
- delegated authorization credentials - to be used with other servers

It should be possible to add to the security context subsequent to its initial creation. This might allow lazy verification of authorization credentials (i.e. do not verify them until they are needed for an operation), as well as requests for additional credentials from clients for certain operations.

The third argument identifies the target of the operation for which authorization is to be checked. This would typically be a null terminated string with the name of the object. The namespace from which the name is drawn can be local to an application. Names for different applications can be made distinct by specifying a namespace identifier in the parameters argument described later.

There was a question raised about ACL management. In particular, it was not clear where there was one manager for all the ACLs, or whether the application stores the ACL itself. In fact, both models should be supported. For the latter model, one might want to pass in the ACL as an argument, rather than the name of the object which would then be looked up by the ACL manager. One way to support this is using an additional flag in the parameters argument that would indicated that the target argument is a pointer to the actual ACL.

There was a suggestion to use object identifiers to identify the target, but they are really not needed. If you have them, you can use them as one of the name spaces, but you will not always have object identifiers as the names of objects in an application. Only if the application uses OIDs, would this be the name of the targets. A separate name space for OIDs is easily supported in the parameters.

The fourth argument is a pointer to a bit vector identifying the operations for which authorization is being checked, or the bit vector itself if less than 32 bits. How the field is interpreted depends on flags in the parameter argument. The parameter argument also tells how the bits are to be interpreted so that privilege bits from one application (meaning certain rights) are not confused with those from other applications that might mean something different.

The final argument is the parameters argument. This argument describes the behavior of the API. It is expected that the argument will remain fixed for a particular application. It defines how the other arguments are to be interpreted. It identifies the name space of the object names and the form of the other arguments. The same structure would be passed on all calls to check authorization.

- **Access Control List Entries**

The access control lists used by the API are similar to those in common use today. An access control list will be associated with each object to be protected. Entries in the list will identify the principals or groups authorized access by that entry. The principal and group specified will be matched against the principal and group identifiers in the security context by the authorization API. The rights specifically authorized by the entry will be specified in the entry as a tagged bit vector. The tag indicates how the bits in the bit vector are to be interpreted.

The ACLs will be extended beyond those in common use today in that each entry will have an optional additional list of restrictions (the negative attributes described earlier in the minutes). This list of restrictions places additional constraints on the access granted. For example, it might allow access to be granted only between 9AM and 5PM, or the access might expire at a particular date and time. Restrictions in an ACL entry might also encode which intermediaries are allowed to be involved when authenticating principals that use the entry.

Simple applications will never see these restrictions. Instead, they will be evaluated by the authorization API itself, and the yes/no answer seen by the application will already reflect the result of that evaluation.

- **Use of the Authorization API by Advance Applications**

More advanced applications will rely on additional information returned in the detailed answer argument. Among the information returned will be an expiration time for access so granted if such a restriction in the ACL entry is present or if the expiration time of any authentication or authorization credentials is known.

Further, the authorization API can return an answer of maybe, indicating that restrictions (negative attributes) were present in authorization credentials or access control list entries that could not be interpreted directly by the authorization API. This is likely to be the case if application specific restrictions were used.

When such an answer is returned, the unresolved restrictions are returned to the application in the detailed answer structure. This list will then be checked by the application using a fairly well defined procedure, plugging local checks into boilerplate restriction checking code that will be provided.

The network access server will provide a good example of how such checks are to be performed, and a description of the use of the authorization API by the network access server should be included in documents describing the mechanism.

Discussion of Guidelines Document

There was a brief discussion about the development of a set of guidelines for application developers on how to support fine grained authorization for network accessible applications. This document would be released initially as an Internet-Draft and then as an Information RFC.

The advice would be basic. For example, advising the developer to first look through the application and identify the objects to be protected, and the granularity of the objects to be protected: is it a file, is it something more general, or more specific. The document would then work through how one would use the API which is being developed.

The document should also discuss design implications of particular choices, especially with respect to how long authorization continues (i.e. whether you check the returned expiration time), and whether authorizations checks are per session, or per operation, as well as issues related to the granularity of the objects that are protected.

The document would give an example of the use of the authorization API for a simple application that just uses the yes/no answer, and for a more advanced application such as the network access server that makes use of application specific restrictions.

How to Proceed

The meeting concluded with a brief discussion of how to proceed. Work items include refining the authorization API and developing real code, coming up with a revised definition of the security context, revising the list of authorization attributes, and writing the guidelines document. Piers will continue his work on the security context in part here, and in part in the CAT Working Group. Cliff will work with Piers on the security context, and will work on the other work items. John Vollbrecht will work on the section of the guidelines document describing how the network access server uses the authorization API.

Attendees

Andy Adams	ala@merit.edu
Garrett Alexander	gda@tycho.ncsc.mil
Nick Alfano	alfano@mpr.ca
Alireza Bahreman	bahreman@bellcore.com
Steven Bellovin	smb@research.att.com
Larry Blunk	ljb@merit.edu
Antonio Fernandez	afa@thumper.bellcore.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Jisoo Geiter	geiter@mitre.org
Mei-Jean Goh	goh@mpr.ca
Chris Gorsuch	chrisg@lobby.ti.com
Marco Hernandez	marco@cren.net
Charlie Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
David Kristol	dmk@allegra.att.com
Gordon Lee	gordon@ftp.com
John Linn	linn@security.ov.com
Kanchei Loa	loa@sps.mot.com
Steven Lunt	lunt@bellcore.com
Piers McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Michael Michnikov	mbmg@mitre.org
Bob Morgan	morgan@networking.stanford.edu
Robert Moskowitz	3858921@mcimail.com
Clifford Neuman	bcn@isi.edu

Allan Rubens	acr@merit.edu
Richard Schmalgemeier	rgs@merit.edu
Wolfgang Schneider	schneiw@darmstadt.gmd.de
Vincent Shekher	vin@sps.mot.com
Sam Sjogren	sjogren@tgv.com
Dave Solo	solo@bbn.com
Don Stephenson	don.stephenson@sun.com
Barbara Sterling	bjs@mcdata.com
Vladimir Sukonnik	sukonnik@process.com
Dean Throop	throop@dg-rtp.dg.com
Jerry Toporek	jt@mentat.com
Theodore Ts'o	tytso@mit.edu
John Vollbrecht	jrv@merit.edu
Peter Yee	yee@atlas.arc.nasa.gov

2.7.2 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 US Department of Defense. CIPSO will be designed to provide labeling for the commercial, US civilian and non-US 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.

2.7.3 Common Authentication Technology (CAT)

Charter

Chair(s)

John Linn: linn@security.ov.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.
Done	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.
Done	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.
Done	Submit service interface specification to to the IESG for consideration as a Proposed Standard.

Internet-Drafts

"FTP Security Extensions", 11/15/1993, S. Lunt <draft-ietf-cat-ftpsec-04.txt>

Request For Comments

RFC 1507	"DASS - Distributed Authentication Security Service"
RFC 1508	"Generic Security Service Application Program Interface"
RFC 1509	"Generic Security Service API : C-bindings"
RFC 1510	"The Kerberos Network Authentication Service (V5)"
RFC 1511	"Common Authentication Technology Overview"

CURRENT MEETING REPORT

Reported by John Linn/OpenVision and Sam Sjogren/TGV

Minutes of the Common Authentication Technology Working Group (CAT)

Overview

The CAT Working Group met for two sessions in Houston. The status of ongoing activities was reviewed, including a reworked GSS-API implementation for Kerberos V5 beta 3. This distribution, and an Internet-Draft describing its GSS-API mechanism characteristics and token formats, should be available later this year. Some interface clarifications and extensions (e.g., a new GSS_Inquire_context primitive) were discussed as inputs to Internet-Draft successors to RFCs 1508 and 1509, targeting inclusion in eventual Draft Standard versions to supplant those RFCs and comprise a "Version 2" GSS-API. Related topics to be discussed on the mailing list include multi-mechanism credential management and error reporting. Piers McMahon gave a presentation on SESAME's multi-mechanism implementation, and distributed a paper for comment. Sam Sjogren and Steve Lunt led a discussion on the FTP Security Internet-Draft, to be updated shortly and to be used as the basis for an interoperability test (using Kerberos V4 technology) planned for March 1994. Representatives from the Network Access Server Requirements Working Group (NASREQ) described their currently contemplated architecture as input to determining how the CAT Working Group and technology might support their needs. Ran Atkinson gave a presentation on the Internet Authentication Guidelines Internet-Draft, receiving and soliciting comment from the attendees.

Status Review

Ted Ts'o reported that two independent implementors are reworking the GSS-API implementation for Kerberos V5; it is expected that the result of one of these activities will be incorporated into Kerberos V5 beta 3, to be available as a redistributable release in December. (This step will replace and obsolete the "alpha quality" GSS-API in Kerberos V5 beta 2.) Detailed documentation, including token formats for the mechanism, is being prepared and will be included in an Internet-Draft which John Linn stated would also be distributed in December.

No effort on a Kerberos V4 GSS-API implementation is known. Ted Ts'o offered to review and contribute to a design specification for KV4 GSS-API if anyone wishes to drive this activity.

Piers McMahon provided hardcopy of a memo he drafted describing a framework for GSS-API extensions targeted for POSIX environments, and solicited comments.

John Linn reported that the ongoing liaison with the X/Open Security Working Group is progressing well, and that the technical content of RFCs 1508 and 1509 is incorporated in a document currently undergoing X/Open Company Review for publication as an X/Open Preliminary Specification.

Interface Extensions and Refinements

The procedure through which any changes and extensions to be made to 1508 and 1509 would be reflected and characterized was discussed. The current RFCs were declared as constituting "GSS-API Version 1," and successor Internet-Drafts will be generated enroute to become "GSS-API Version 2" (aka GSSV2).

The GSS_Inquire_context() primitive, as discussed on the mailing list, was accepted as an addition for GSSV2. Renaming of the per-message primitives, per X/Open terminology request and as also discussed previously on the mailing list, was accepted as a change for GSSV2.

The attendees discussed the issue of credential acquisition in a multi-mechanism environment, including a proposal made to the mailing list for definition of a new GSS_Add_cred() primitive to be used in preference to the current GSS_Acquire_cred(). Since GSS_Acquire_cred(), like other GSS-API calls, returns only a single pair of major_status and minor_status values, its use in a multi-mechanism environment cannot return specific information about each of the supported mechanisms for which credentials may or may not have been successfully acquired.

Several attendees observed the fact that the need to disambiguate minor_status values is primarily of interest to callers embodying knowledge of mechanism-specific characteristics and needing to make decisions based on those characteristics, a class of callers which attendees sought to minimize. Despite this fact, some mechanism-cognizant callers (or callers seeking to display meaningful minor_status indications to their clients) will certainly exist, and it's appropriate to consider how they could be better served for GSSV2.

In addition to the prior GSS_Add_cred() proposal, it was observed that callers requiring unambiguous per-mechanism status information could use the current GSS_Acquire_cred(), explicitly specifying a single mechanism per invocation, at the cost of losing the convenience of multi-mechanism credentials. [Though not cited in meeting discussion, the GSS_Indicate_mechs() primitive provides the necessary data for a caller to perform this iteration.] Following some discussion, John Linn accepted the action of further summarizing options and tradeoffs in a message to the mailing list.

The level of portability to be supported by GSS-API mechanisms was discussed, and it was agreed to take feasible and apparent measures in the interests of supporting object-level portability across different implementations. Specifically, the forthcoming Kerberos V5 mechanism Internet-Draft should define a set of common minor_status values to be used by implementors of the mechanism, though additional minor_status codes specific to particular implementations are also possible. Further, it was agreed that the "gssapi.h" header file at the end of RFC 1509 should be considered part of the standard, noting that refinements and additional elements (e.g., type definitions for name representations having broader scope than a single mechanism) might be incorporated for GSSV2 and that particular implementations would likely append their own, implementation-specific definitions over and above gssapi.h.

The attendees discussed a prior request to incorporate a form of per-message protection which would provide confidentiality without integrity, but did not elect to incorporate such a facility.

Presentation on Multi-Mechanism Issues

Piers McMahon (ICL) gave a presentation entitled "GSS-API IN A MULTI-MECHANISM ENVIRONMENT," covering four topics: (1) problem domain, (2) architecture of GSS-API implementation by SESAME, (3) API implications, and (4) approaches to credential acquisition.

Regarding the problem domain, Piers observed the following: (1a) today's (and probably tomorrow's) problem is heterogeneity, (1b) users expect single sign-on, and (1c) administrators expect single point of user registration. Regarding API implications, he cited internal structure constructs designed to separate the GSS-API layer from individual underlying mechanisms: internal APIs to deposit/append/clear credential elements, credential element/security context specialization features (function vectors, data), and use of a common cryptographic support facility (CSF) for all mechanisms. In terms of external elements (those visible to GSS-API callers), he noted that it was necessary to provide a single common view of timeouts and other attributes, spanning all underlying mechanisms. Regarding credential acquisition and related login functions, he cited three concepts: multiple login, use of shared data elements relevant to more than a single mechanism, and an "access manager" which would use, e.g., passwords or credentials for one mechanism as a basis for which to acquire credentials of another mechanism on behalf of a user.

Architecture of GSS-API Implementation by Sesame

```
    Principal --owns-> Credential

                           |
                      contains one or more
                           V

                    Credential Element

                           ^
                      inherits from
                           |

          +---------- SESAME Credential Element
       initiate/
        accept
          |          Security Context --------uses---> Cryptographic
          |                                            Support
          |                ^                           Facility
          |           inherits from
          |                |
          |
          +---------> SESAME Security Context
```

Internet Authentication Guidelines Draft

Ran Atkinson gave a presentation on the Internet-Draft he had co-authored with Neil Haller (`draft-haller-auth-requirements-01.txt`), to solicit comments from the IETF community. The document distinguishes four classes of authentication: none, disclosing (subject to passive attacks), non-disclosing (subject to active, but not to passive attacks), and strong (resistant to passive and active attacks); its pragmatic motivation is to encourage migration to at least the non-disclosing level. While this taxonomy was accepted as useful and primary, it was noted that technologies could also be distinguished on other grounds: human-oriented versus machine-oriented, orientation to point-to-point versus distributed system usage, and requirements for shared secrets. It was recommended that the document retain a consistent and specific focus on authentication, and that tutorial material be separated from commentary and opinion.

It was noted that some of the content overlapped with sections of the Internet Security Architecture document being prepared by the PSRG; Jeff Schiller commented that he believed the documents were generally aligned, but that some work would be needed in order to assure that terminology definitions were consistent. As a particular example, Ran noted that he had observed U. S. Defense Department usage of the term "digital signature" as referring to integrity protection without non-repudiation, a form of usage inconsistent with much of the literature.

Network Access Security Requirements (NASREQ)

John Vollbrecht attended a portion of the CAT meeting in order to inform attendees on NASREQ's environment and concerns, to solicit comment, and to explore possible areas of overlap between the groups. Review of anticipated design documents was solicited.

The NASREQ environment includes Network Access Clients (NACs, typically PCs) accessing Network Access Servers (NASs) via dial-up. It is planned that the NASs will communicate with authentication servers across a network, perhaps indirectly by way of "helper" devices. PPP is used across the dial-up link, presently with PAP and CHAP but with new KAP (Kerberos), PKAP (public key) and SCAP (smart card) authentication schemes contemplated but not yet documented; a brief explanatory memo will be distributed shortly. The "RADIUS" protocol is being considered as a basis for interaction between NASs and authentication servers. Mobility support, enabling users to connect to NASs in foreign domains (with multiple intermediary helpers between the access point and a user's home NAS) is desired and introduces inter-domain trust considerations.

Two authentication types are currently distinguished within user records: "UNIX" (password-level) and "Kerberos" (in which a Kerberos server is involved in the process of authenticating a user for access to network resources via a NAS). It was suggested that Derek Atkins' MIT thesis on use of Kerberos in a dial-up environment represents an alternate approach worthy of consideration. GSS-API might be useful as a means to protect traffic between NASs and helpers and/or authentication servers, but its current underlying mechanisms are not oriented to operation across a dial-up link where clients lack independent access to authentication servers.

FTP Security

Since the Amsterdam IETF meeting the FTP security Internet-Draft (the current version is `draft-ietf-cat-ftpsec-03.txt`) had been changed by the author, Steve Lunt, to reflect discussions at that meeting. The changes were:

- Principal name fallback (use "rcmd" if "ftp" doesn't exist)

  This would allow maximal flexibility for an administrator to restrict an FTP server and the environment it runs in, while allowing for simplification of administration by not requiring the configuration of new principals if a site wished to just use the "rcmd" principal which they would already have for use by Kerberized R-Services and Telnet and the like. Any restrictions on what principal must be used and other configuration issues would be implementation and site specific.

- Changed GSS_Safe to GSS_Seal with conf_flag == False
- Changed GSS_Verify to GSS_Unseal
- Changed GSS_Seal to GSS_Seal with conf_flag == True
- Changed the mailing list from `ftp-wg@tgv.com` to `cat-ietf@mit.edu`.

  There is a mail reflector at TGV which will remain in existence indefinitely. So, mail sent to `ftp-wg@tgv.com` will merely get forwarded to `cat-ietf@mit.edu`. It is recommended, however, that the cat-ietf address be used.

There were two outstanding protocol length issues which were introduced for discussion leading to closure. First, the issue of the length of buffers allowed for protected file transfers, and second, the lack of restriction on the length of base-64 encodings.

It is desirable to have a finite buffer size used for protected file transfers, as there may be situations in which a system would need to read an entire buffer into memory before being able to operate on it, and some systems have far more finite memory resources than others. However, specifying some arbitrarily small buffer size could have an impact on performance and even functionality. It was decided that a negotiation would be added to the protocol which would allow the client and server to agree on a buffer size. Steve will add the specification of this to the document.

A base-64 encoding may be of arbitrary length. The binary authentication data that is encoded may be of arbitrary length. Although a line-wrapping scheme could be specified that would wrap lines and thereby limit the clear-text line length while allowing the arbitrarily long binary data, it was not felt that there was any need to do that. The document will be modified to note that the base-64 encoded data lines can be arbitrarily long without line-wrapping being used.

The issue of a fall-back targ_name for the GSS-API specification for FTP was not resolved. The name "SERVICE:ftp@hostname" is currently specified, but it is unclear what would be a more common name to fall back to (as with the "ftp" to "rcmd" fallback in the Kerberos V4 specification). This issue will be resolved via e-mail.

A request for adding state diagrams was made. This will be satisfied in a future revision of the document.

Interoperability Bakeoffs

Sam Sjogren led a discussion which proposed the idea of having "virtual Bakeoffs" between IETF meetings to motivate implementation of standards being worked on and interoperability testing of those implementations. A tentative date of the week of 14 March 1993 will be the target for a virtual bakeoff of the FTP Security work.

Attendees

Garrett Alexander	gda@tycho.ncsc.mil
Nick Alfano	alfano@mpr.ca
Alireza Bahreman	bahreman@bellcore.com
Luc Boulianne	lucb@cs.mcgill.ca
Chuck de Sostoa	chuckd@cup.hp.com
Antonio Fernandez	afa@thumper.bellcore.com
Steve Garritano	steveg@kalpana.com
Jisoo Geiter	geiter@mitre.org
Chris Gorsuch	chrisg@lobby.ti.com
Marco Hernandez	marco@cren.net
Charlie Kaufman	kaufman@zk3.dec.com
Walter Lazear	lazear@gateway.mitre.org
Gordon Lee	gordon@ftp.com
John Linn	linn@security.ov.com
Kanchei Loa	loa@sps.mot.com
Steven Lunt	lunt@bellcore.com
Chip Matthes	chip@delphi.com
Wayne McDilda	wayne@dir.texas.gov
Piers McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Michael Michnikov	mbmg@mitre.org
Bob Morgan	morgan@networking.stanford.edu
Clifford Neuman	bcn@isi.edu
Rakesh Patel	rapatel@pilot.njin.net
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
Wolfgang Schneider	schneiw@darmstadt.gmd.de
Vincent Shekher	vin@sps.mot.com
Sam Sjogren	sjogren@tgv.com
Frank Solensky	solensky@ftp.com
Dave Solo	solo@bbn.com
Don Stephenson	don.stephenson@sun.com
Jerry Toporek	jt@mentat.com
Theodore Ts'o	tytso@mit.edu
John Veizades	veizades@ftp.com
John Vollbrecht	jrv@merit.edu

2.7.4 Internet Protocol Security Protocol (IPSEC)

Charter

Chair(s)

Al Hoover: hoover@ans.net
Paul Lambert: paul_lambert@email.mot.com

Mailing Lists

General Discussion: ipsec@ans.net
To Subscribe: ipsec-request@ans.net
Archive: ftp.ans.net:~/pub/archive/ipsec

Description of Working Group

Rapid advances in communication technology have accentuated the need for security in the Internet. The IP Security Protocol Working Group (IPSEC) will develop mechanisms to protect client protocols of IP. A security protocol in the network layer will be developed to provide cryptographic security services that will flexibly support combinations of authentication, integrity, access control, and confidentiality. The protocol formats for the IP Security Protocol (IPSP) will be independent of the cryptographic algorithm. The preliminary goals will specifically pursue host-to-host security followed by subnet-to-subnet and host-to-subnet topologies.

Protocol and cryptographic techniques will also be developed to support the key management requirements of the network layer security. The key management will be specified as an application layer protocol that is independent of the lower layer security protocol. The protocol will initially support public key based techniques. Flexibility in the protocol will allow eventual support of Key Distribution Center (KDC - such as Kerberos) and manual distribution approaches.

Goals and Milestones

Jun 1993 Post as an Internet-Draft the IP Security Protocol.

Jul 1993 Post as an Interenet-Draft the specification for Internet key management.

Nov 1993 Report on pilot implementation of the IP Security Protocol. Update Protocol as needed.

Mar 1994 Report on pilot implementation of the Internet Key Management Protocol. Update Internet-Draft as needed.

Jul 1994 Submit the IP Security Protocol to the IESG for consideration as a Proposed Standard.

Jul 1994 Submit the Internet Key Management Protocol to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Paul Lambert/Motorola

Minutes of the Internet Protocol Security Protocol Working Group (IPSEC)

Many thanks to Tom Benkart who served as recording secretary for these meetings.

The IPSEC Working Group met twice during the Twenty-Eighth IETF. On Tuesday, November 2 the IPSEC working group met and discussed the IP Security Protocol (IPSP). On Wednesday, November 3 the working group held a demonstration of two IPSP implementations and discussed the key management requirements of IPSEC.

Working Group Status

Paul Lambert began the meeting with a review of the charter and a working group status report.

- The working group is behind schedule.
- Only I-NLSP has been submitted as an Internet-Draft.
- It is important to track the IPng candidates.

Since I-NLSP is the only Internet-Draft so far, it may be used as the template for a document from the entire group. That does not mean that its concepts would be adopted, just that it provides the working group with a starting point.

Richard Thomas stated that the NLSP specification may be available electronically soon. It is on the list of documents to be made available from ISO.

IPSP Requirements Review

Agreement must be reached on the requirements before the contending proposals can be evaluated. Requirements were discussed in Amsterdam, but the decisions were not considered final due to the small number of participants. Since the minutes from that meeting were not published, the working group as a whole did not have a chance to comment. The following comments reflect points from the Amsterdam and Houston meetings.

Editor's Note: A list of issues that were discussed is available via FTP or mail server from the remote directories as /ietf/ipsec/ipsec-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Fragmentation was the major item addressed and remains an open issue. Protocols such as NFS send maximum-sized UDP datagrams, and the encapsulation done by IPSP in ISs frequently results in additional fragmentation. MTU Discovery can be a solution (provided the routers account for the IPSP encapsulation in their ICMP messages), but MTU Discovery is not commonly used today. NFS 3 runs over TCP, so this might not be so large an issue when that version is available.

Reassembly is not required in IPSP as long as layering is maintained. For example, in the case of IPSP between two ISs, reassembly is handled by the normal IP processing since the added IP header specifies the remote IS as the IP destination.

Jim Zmuda and Bill Simpson volunteered to write a requirements document based on the discussions.

Experimental IPSP Implementation Review

The following implementations were reviewed:

- I-NLSP - Rob Glenn
- swIPe - Phil Karn
- KeyRing - Rob Hagens
- TANDU/Cryptonette - Charlie Kaufman
- LAN Guardian - Mike O'Dell
- SP3 - Paul Lambert

Editor's Note: A review of each implementation is available via FTP or mail server from the remote directories as /ietf/ipsec/ipsec-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The question of supporting multiple PDU formats as part of the IPSP specification is an open issue. Arguments in favor are that different media types will need different formats to be efficient, and that ES-ES IPSP could do TCP-UDP over IPSP instead of IP over IP encapsulation. Arguments opposed are that the added complexity will make IPSP more difficult to specify and implement. One possible approach is to specify a negotiation mechanism with defaults (like PPP).

IPSP Implementation Demonstrations

Jim Zmuda and Phil Karn gave demonstrations of their implementations. Jim's implementation was based on the ISO 11577 specification for Network Layer Security Protocol (NLSP) and used the NLSP specification of a Security Exchange Protocol (SAEP) for key management. The implementation demonstrated by Phil Karn was based on Phil's KA9Q software running on a portable computer (80386 based). This demonstration ran between Houston and Phil's home. Key management was based on the manual entry of DES key variables.

Key Management

Topics of discussion included the following:

- What is key management and what is the group's charter for key management?
- What work already exists that might be taken advantage of?
- Who would serve as liaisons for some of this existing work?

Editor's Note: Itemized lists of answers to each of the above questions is available via FTP or mail server from the remote directories as /ietf/ipsec/ipsec-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Key Management Requirements

This meeting was the first attempt to list the requirements for a KMP. The requirements fall into two categories - Peer-to-Peer Exchanges and Security Management.

Editor's Note: A list of requirements for each catagory as well as a list of issues concerning these requirements is available via FTP or mail server from the remote directories as /ietf/atm/atm-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Existing hosts have to be able to take advantage of IPSP services in routers without any change to the host (i.e., IPSP is transparent to non-IPSP end systems).

Dave Solo volunteered to write a requirements document for IPSEC Key Management.

Concern was expressed about supporting public keys first in the IPSEC goals because of possible delays from patent issues (a lesson learned by PEM). Having a standard API for communicating keys from the KMP entity to IPSP would facilitate support for private/shared keys.

Existing Key Management Implementation Presentations

Rob Hagens gave a presentation on KeyMan product. Following are attributes of the product:

- The routers have a pre-configured list of peer entities.
- A Key Encryption Key is used in the current Beta release; a new version using public/private keys is in development.
- Traffic key lifetimes are dynamically specified.

- Different TEKs are used in each direction; setup can be done in two messages, but four are normally used.
- All PDUs have the same format (48 bytes).
- Public keys are currently locally stored (manual distribution).
- TEK generation does not use Diffie-Hellman, so recorded traffic could be decrypted if the private keys are ever learned.
- If a router crashes, it establishes new SAs; the other routers discard the old SA when a new setup is requested.

Jim Zmuda gave a presentation of key management in the NetLock product. It uses SAEP and a trusted Certification Authority on at least one of the systems. A Diffie-Hellman exchange is used to generate a secret for shared keys. The secret is also encrypted with the sender's private key for authentication.

Attendees

Garrett Alexander	gda@tycho.ncsc.mil
Randall Atkinson	atkinson@itd.nrl.navy.mil
Alireza Bahreman	bahreman@bellcore.com
Steven Bellovin	smb@research.att.com
Tom Benkart	teb@acc.com
Larry Blunk	ljb@merit.edu
Ken Carlberg	Carlberg@cseic.saic.com
Cheng Chen	chen@accessworks.com
Richard Colella	colella@nist.gov
Stephen Crocker	crocker@tis.com
Waychi Doo	wcd@berlioz.nsc.com
Robert Downs	bdowns@combinet.com
Donald Eastlake	dee@skidrow.lkg.dec.com
Antonio Fernandez	afa@thumper.bellcore.com
Richard Fox	rfox@metricom.com
Dan Frommer	dan@isv.dec.com
Vincent Gebes	vgebes@sys.attjens.co.jp
Jisoo Geiter	geiter@mitre.org
Robert Glenn	glenn@osi.ncsl.nist.gov
Mei-Jean Goh	goh@mpr.ca
Chris Gorsuch	chrisg@lobby.ti.com
Phillip Gross	pgross@ans.net
Robert Hagens	hagens@ans.net
Phil Karn	karn@qualcomm.com
Charlie Kaufman	kaufman@zk3.dec.com
Elizabeth Kaufman	kaufman@biomded.med.yale.edu

Stephen Kent	kent@bbn.com
Edwin King	eek@atc.boeing.com
Michael Kornegay	mlk@bir.com
David Kristol	dmk@allegra.att.com
Paul Lambert	paul_lambert@email.mot.com
John Linn	linn@security.ov.com
Brian Lloyd	brian@lloyd.com
Kanchei Loa	loa@sps.mot.com
Steven Lunt	lunt@bellcore.com
Gary Malkin	gmalkin@xylogics.com
Glenn McGregor	ghm@lloyd.com
Michael Michnikov	mbmg@mitre.org
Greg Minshall	minshall@wc.novell.com
Sandra Murphy	murphy@tis.com
Andrew Myles	andrew@mpce.mg.edu.au
Michael O'Dell	mo@uunet.uu.net
Steve Parker	sparker@ossi.com
Henning Schulzrinne	hgs@research.att.com
Vincent Shekher	vin@sps.mot.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Frank Solensky	solensky@ftp.com
Dave Solo	solo@bbn.com
James Solomon	solomon@comm.mot.com
Michael St. Johns	stjohns@arpa.mil
Don Stephenson	don.stephenson@sun.com
Richard Thomas	rjthomas@bnr.ca
Susan Thomson	set@bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Jerry Toporek	jt@mentat.com
Paul Traina	pst@cisco.com
Theodore Ts'o	tytso@mit.edu
Anthony Valle	valle@huntsville.sparta.com
Chuck Warlick	chuck.warlick@pscni.nasa.gov
David Woodgate	David.Woodgate@its.csiro.au
Peter Yee	yee@atlas.arc.nasa.gov
James Zmuda	zmuda@mls.hac.com

IP Security (IPSEC) Working Group at the Twenty Eight Internet Engineering Task Force (IETF)

Meeting 1 of 2 - Tuesday November 2, 1993

Paul A. Lambert
Motorola, Inc.
+1-602-441-3646
Paul_Lambert@email.mot.com

Agenda

TUESDAY, November 2, 1993
0930-1200 IP Security Protocol (IPSP)

WEDNESDAY, November 3, 1993
0930-1200 Key Management for IP Security and Demonstrations of Experimental IP Security Implementations

Mailing lists:

General Discussion:	ipsec@ans.net
To Subscribe:	ipsec-request@ans.net
Archive:	ftp.ans.net:~/pub/archive/ipsec

Why Do We Need Security?

Because there are real threats! These threats fall in the following categories:

- Masquerade/Impersonation
- Unauthorized behavior and access
- Leakage of Information
- Integrity threats including:
- Message Sequencing (Replay, Pre-Play, Delay)
- Modification of Information
- Repudiation (Denial of Origin, Denial of Delivery)

Security is an "enabling" technology.

- Strong security provides the "trust" in an electronic communication system that will convince users to replace existing paper paradigms.
- Strong security is required to protect the communication infrastructure from abuse.
- Security will eventually be integrated into communication systems so that it is pervasive and transparent to the users.

Applicable Standards

- IEEE Standard for Interoperable LAN Security (SILS IEEE 802.10B)
- IEEE 802.10C Key Management
- ANSI X3T5.7 Generic Upper Layer Security
- IETF - Common Authentication Technologies
- IETF - Privacy Enhanced Mail
- X.509 - Security for the Directory Service
- X.411 Security for Electronic Messaging
- ANSI X9 Wholesale and Retail Key Management Standards
- CALS & EDI
- NIST Digital Signature Standard (DSS)
- RSA Signature Specifications
- ANSI X3S3.3 (Network Layer Security Protocol - NLSP)
- CCITT T1.P1 (Universal Personal Communications)
- ISO SC6 and SC21
- and others

Why Do We Need Security?

Because there are real threats! These threats fall in the following categories:

- Masquerade/Impersonation
- Unauthorized behavior and access
- Leakage of Information
- Integrity threats including:
- Message Sequencing (Replay, Pre-Play, Delay)
- Modification of Information
- Repudiation (Denial of Origin, Denial of Delivery)

Security is an "enabling" technology.

- Strong security provides the "trust" in an electronic communication system that will convince users to replace existing paper paradigms.
- Strong security is required to protect the communication infrastructure from abuse.
- Security will eventually be integrated into communication systems so that it is pervasive and transparent to the users.

Security Mechanisims

- Authentication Exchanges
- Encipherment (symmetric, or asymmetric)
- Access Control
- Data Integrity
- Traffic Control
- Digital Signature
- Routing Control
- Notarization
- Security Labels
- Event Detection
- Security Audit Trail
- Time Stamping
- Security Recovery

The mechanisms above can be used to provide security services.

Standards for Lower Layer Security

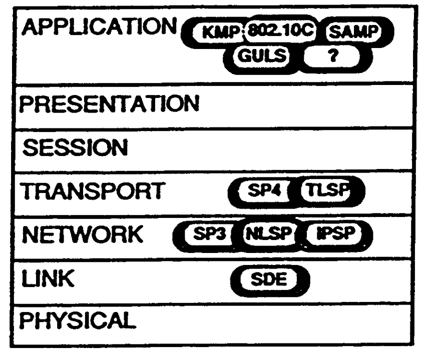

Model for SP3 and Network Layer Security Protocol

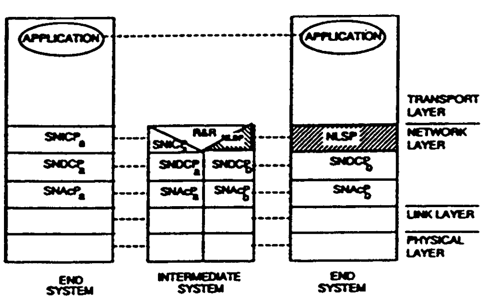

Comparison of Services

ISO 7498/2 Service	ISO 7498/2 Service Recommendations by Reference Model Layer 1	2	3	4	5	6	7	Lower Layer Security Protocol: Layer 2 SDE	Layer 3 NLSP	Layer 4 TLSP
Peer Entity Authentication	•	•	Y	Y	•	•	Y	•(1)	•(1,2)	•(1)
Data Origin Authentication	•	•	Y	Y	•	•	Y	Y	Y	Y
Access Control Service	•	•	Y	Y	•	•	Y	Y	Y	Y
Connectionless Confidentiality	Y	Y	Y	Y	•	Y	Y	Y	Y	Y
Connection Confidentiality	•	Y	Y	Y	•	Y	Y	•	•(2)	Y
Selective Field Confidentiality	•	•	•	Y	•	Y	Y	•	•	•
Traffic Flow Confidentiality	Y	•	Y	Y	•	•	Y	•	•(2)	•
Connection Integrity with Recovery	•	•	•	Y	•	•	Y	•(3)	•(3)	Y
Connection Integrity without Recovery	•	•	Y	Y	•	•	Y	•	Y	Y
Selective Field Connection Integrity	•	•	•	•	•	•	Y	•	•	•
Connectionless Integrity	•	•	Y	Y	•	•	Y	Y	Y	Y
Selective Field Connectionless Integrity	•	•	•	•	•	•	Y	•	•	•
Non-repudiation, Origin	•	•	•	•	•	•	Y	•	•	•
Non-repudiation, Delivery	•	•	•	•	•	•	Y	•	•	•
Communication Capabilities										
"End-to-End" Security								•	Y	Y
Intermediate System Security; bridge or router like security systems								Y	Y	•
Labels for Access Control								•	Y	Y
Security for non-ISO Communications (TCP/IP)								Y	Y	•
Protection for Routing Protocols: ES-IS, IDRP, ARP, EGP, etc.								Y	Y	•

Charter - Key Management

Protocol and cryptographic techniques will also be developed to support the key management requirements of the network layer security. The key management will be specified as an application layer protocol that is independent of the lower layer security protocol. The protocol will initially support public key based techniques. Flexibility in the protocol will allow eventual support of Key Distribution Center (KDC - such as Kerberos) and manual distribution approaches.

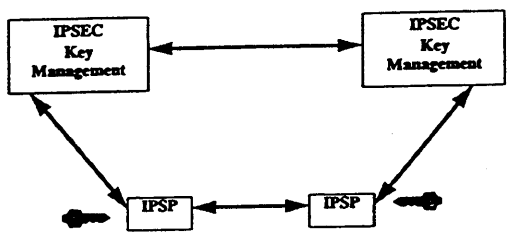

Existing Work

Key Management and Authentication Specifications

- SDNS KMP
- IEEE 802.10C
- GULS
- PEM
- X.509
- X9.17
- SAMP
- SAEP
- Kerberos
- others

Other Applicable Work

- SNMP vII Authentication
- PPP Authentication
- CATS - GSSAPI
- others

Key Management Requirements

Peer-to-Peer Exchange

- Authentication Mechanism / Algorithm Negotiation
- Peer-Entity Authentication
- Key Establishment
- Security Association Negotiation - Establish context key usage.
- Termination of Security Association

Security Management

- Certificate Distribution
- CRL Distribution
- Centralized Key Distribution
- Access Control Attributes

Other Requirements

Phil Karn
karn@qualcomm.com

- Experimental IP Security Protocol in KA9Q NOS package for DOS
- Manual Key Management at present
- Encryption and / or Authentication

Packet Formats:

IP	IPSEC	TCP, UDP, etc	DATA

IP	IPSEC	IP	TCP,	DATA

<< note - other hand drawn figures not included >>

Observations

- WIth a "statefule" IPSEC, 1 standard security header is not needed for all applications! Can negotiate one of many formats as part of key exchange.
- How to control use of security protocol and to assure application?
- Compression along with encryption?
- Overhead is noticeable, especially with X over SLIP (no TCP nagle algorithm)
- Lower latency encryption
- Things that break: traceroute, VJ header compression , ...

IPSP Requirements *Issues*

HUGHES

1. Encapsulation protocol versus yet another IP Option.
2. Limited Coupling between "Key Management" protocol and IPSP
3. ~~Use of TLV encoded fields~~
4. Overall IPSP format:
5. Restrictions on encoding of header fields and the two headers:
6. SAID
7. Sequence numbering
8. Flags field
9. Protocol
10. ICV
11. Explicit or implicit per PDU Labels? *Implicit Per PDU labels*
12. ~~Label field~~
13. ~~Label field format~~
14. ~~Labels and SAID field~~ RELY ON MTU DISCOVERY on F X NST-
15. Peek-through
16. Multicast
17. Fragmentation *- IPSP interaction with existing fragmentation should be covered under this subject, for the IS case...*

**Item 6 would be in the clear header. Items 7 through 9, would be in the protected header.*

James E. Zmuda

WHAT IS I-NLSP?

K. Robert Glenn (N.I.S.T.)

November 2, 1993

Background

- NLSP (ISO11577) - Two Modes:
  - NLSP-CL (Connectionless Mode),
  - NLSP-CO (Connection Oriented Mode).

- I-NLSP == NLSP-CL + *Enhancements*

NLSP-CL

- Simple Encapsulation/Decapsulation Protocol,

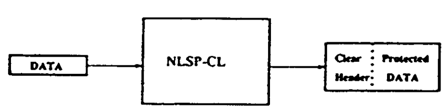

Encapsulation Process

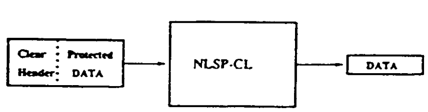

Decapsulation Process

NLSP-CL - (Cont.)

- Provides Five Security Services:
  - **Data Origin Authentication,**
    - **Using Encryption/Integrity Keys.**
  - **Access Control,**
    - **Labels,**
    - **Security Association.**
  - **Traffic Flow Confidentiality,**
    - **Padding,**
    - **Address Hiding.**
  - **Connectionless Confidentiality,**
    - **Encryption/Decryption.**
  - **Connectionless Integrity.**
    - **ICV Generation/Check.**

NLSP-CL - (Cont.)

- Requires Access to Security Attributes That Identify A Security Association,
- Independant of Cryptographic Algorithms/Mechanisms,
- Generic External Interfaces,
  - Designed to solely provide a connectionless network security service.

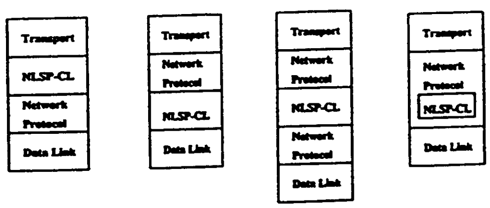

NIST — K. R. Glenn

NLSP-CL - (Cont.)

- Generic TLV Encoded Packet Format,

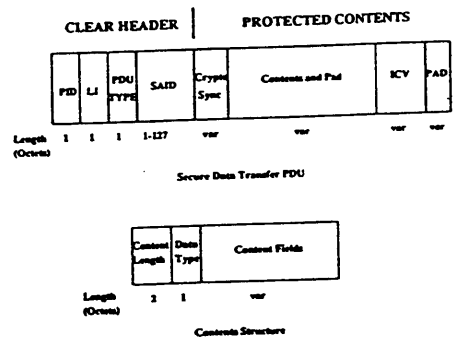

NIST — K. R. Glenn

I-NLSP

- Integrated to provide protection for both CLNP and IP.

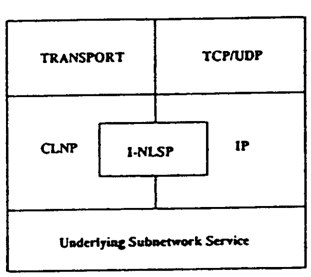

NIST — K. R. Glenn

I-NLSP (Cont.)

- Highly Integrated Interface
  - Packets encapsulated as IP/CLNP packets and forwarded over Normal IP/CLNP Networks,
  - Avoids duplication of IP/CLNP functionality.
  - Security Decision Performed by I-NLSP

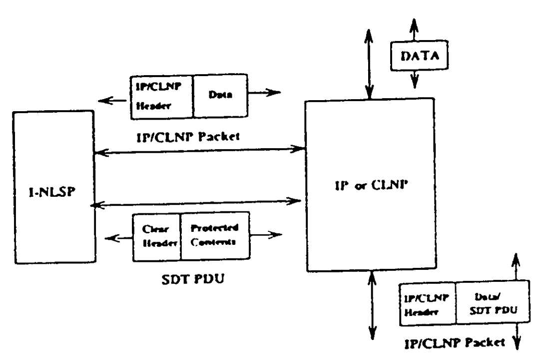

NIST — K. R. Glenn

I-NLSP (Cont.)

- Specialized SDT PDU Format.
  - Word Boundary Issues,
  - TLV Encoding Issues,
  - SDT PDU Future/Private Extensions.

CLEAR HEADER | PROTECTED CONTENTS

	PID	LI	PDU TYPE	SAID	Crypto Sync	Contents and Pad	ICV	PAD
Length (Octets)	1	1	1	1-127	var	var	var	var

Secure Data Transfer PDU

	Content Length	Data Type (x010 0001 IP PACKET) (x010 0010 CLNP PACKET)	IP/CLNP PACKET	PAD
Length (Octets)	2	1	var	var

Possible Revised Contents Format

I-NLSP SUMMARY

- Simple Encapsulation/Decapsulation Protocol,
- Provides Five Security Services,
- Requires Access to Security Attributes That Identify A Security Association,
- Independant of Cryptographic Algorithms/Mechanisms,
- Integrated to provide protection for both CLNP and IP.
- Highly Integrated Interface.
- Specialized SDT PDU Format.

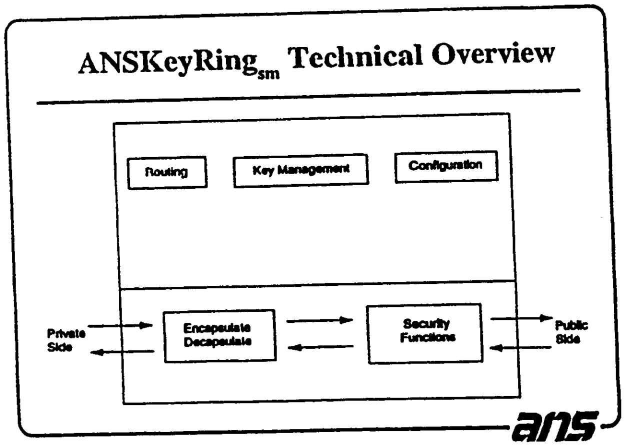
ANSKeyRing$_{sm}$ Technical Overview
Routing
Key Management
Configuration
Private Side
Encapsulate Decapsulate
Security Functions
Public Side

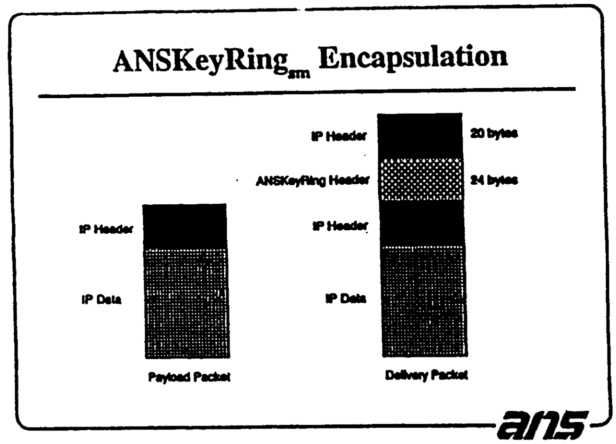
ANSKeyRing$_{sm}$ Encapsulation
IP Header
IP Data
Payload Packet
IP Header
20 bytes
ANSKeyRing Header
24 bytes
IP Header
IP Data
Delivery Packet

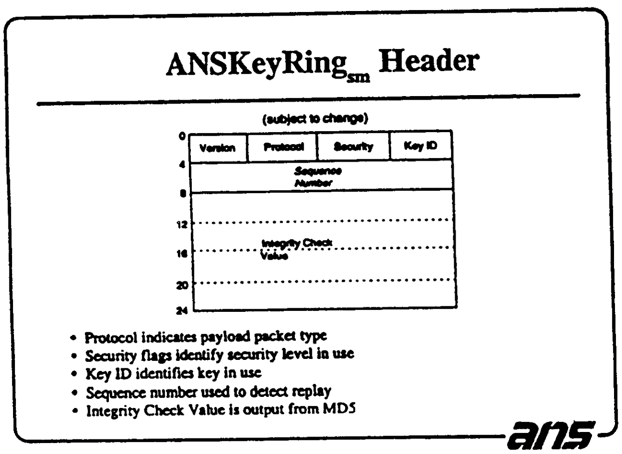
ANSKeyRing$_{sm}$ Header
(subject to change)
0
Version
Protocol
Security
Key ID
4
Sequence Number
8
12
16
Integrity Check Value
20
24
• Protocol indicates payload packet type
• Security flags identify security level in use
• Key ID identifies key in use
• Sequence number used to detect replay
• Integrity Check Value is output from MD5

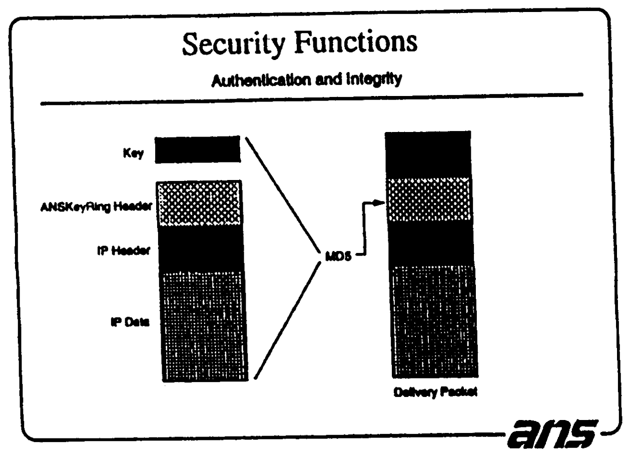
Security Functions
Authentication and Integrity
Key
ANSKeyRing Header
IP Header
IP Data
MD5
Delivery Packet

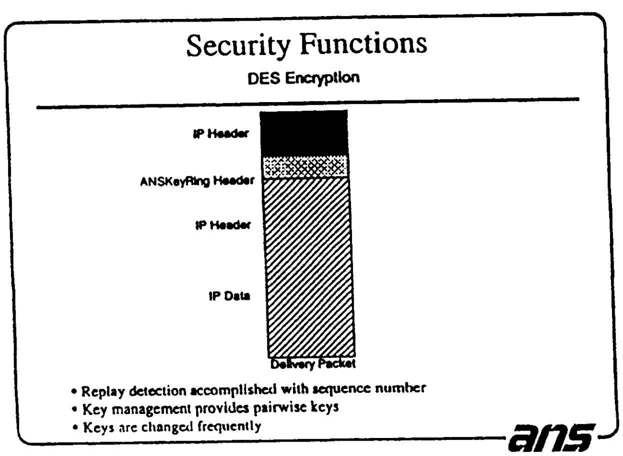
Security Functions
DES Encryption
IP Header
ANSKeyRing Header
IP Header
IP Data
Delivery Packet
• Replay detection accomplished with sequence number
• Key management provides pairwise keys
• Keys are changed frequently

"KeyMan" Key Management Protocol

- Automatic two-party negotiation of secret keys
  - we currently use 56-bit DES keys
- Keys can be rolled over at any desired frequency
  - we normally use 2 hour maximum key lifetime
- Negotiation of the secret keys is protected using public/private key cryptography

Applications for the KeyMan Protocol

- To set up and maintain secure encapsulating IP tunnels

  Secure Closed User Groups can be formed using a network of these tunnels built on the public Internet
- To set up interactive, secure host-to-host sessions

  Could be done using a secure socket type, which puts session-specific secret keys in place when the connect is done

Communications Architecture for KeyMan

- The KeyMan protocol is currently run over TCP
- Could also be run over a light-weight transport
  - simple "stop and wait" flow control, retransmission
  - KeyMan packets sent along with transport connect request and connect confirm packets
- 2 KeyMan packet exchanges (1 round-trip) needed if a trusted transport connection is maintained between key negotiations
- For the most general case, 4 KeyMan packet exchanges are required (2 round-trips) once the transport connection is set up. 4 packets are needed to meet the following constraints:
  - for replay protection, each party sends a randomly created session key and doesn't trust what he hears from his peer until that session ID is returned to him.
  - encrypt keys are not put into use until the installation of the corresponding decrypt key is confirmed.

KeyMan Packet Format

- The same packet format is used for all KeyMan transmissions
- The packet size could be optimized using different packet types, but following public/private key encryption all packets still end up the size of the key modulus
- Each tunnel uses two secret DES key pairs, one for each direction of traffic

KeyMan Packet Format

```
-------------------------------------------------------
| protocol ID | version no. | PDU type  |     ttl     |
-------------------------------------------------------
| initiator's session ID                              |
-------------------------------------------------------
| requestor's session ID                              |
-------------------------------------------------------
| ekey ID     | dkey ID     |  ekey                   |
-------------------------------------------------------
|  ekey                                               |
-------------------------------------------------------
|  ekey                     |  dkey                   |
-------------------------------------------------------
|  dkey                                               |
-------------------------------------------------------
|  dkey                     | integrity check value   |
-------------------------------------------------------
|  integrity check value                              |
-------------------------------------------------------
|  integrity check value                              |
-------------------------------------------------------
|  integrity check value                              |
-------------------------------------------------------
|  integrity check value    | dkey state  | pad       |
-------------------------------------------------------
```

KeyMan Protocol Exchanges

- Protocol echanges to negotiate bi-directional key pairs

(*EKEY* is INITIATOR's encrypt key, *DKEY* is INITIATOR's decrypt key)

```
INITIATOR                                              RESPONDER

create EKEY, and i_session ID
send REQUEST PDU       ------------------------->
                                                 return DKEY, and both session IDs
                       <-------------------------  send RESPONSE PDU
check EKEY, and i_session ID
install DKEY
send CONF1 PDU         ------------------------->
                                                 check KEYs and session IDs
                                                 install both KEYs
                       <-------------------------  send CONF2 PDU
install EKEY
```

Securing the Transmitted Data

Sender:

1. Run MD5 over the data and add the ICV to the packet (integrity)
2. encrypt with sender's private key (source authenticaion)
3. encrypt with receiver's public key (privacy)

Receiver:

1. decrypt with receiver's private key
2. decrypt with senders public key
3. run MD5 over the data (zeroing the ICV) and compare the result with the transmitted ICV

ans

Public Key Lookup

Public key lookup can be incorporated into the key management protocol:

- Each peer stores a certificate containing his public key
- Requesting peer sends a request for certificate along with a random session ID and a copy of his certificate
- Responder returns his certificate and a random session key of his own
- Encrypted key request and key confirm packets then negotiate the key pair

ans

Packet Format

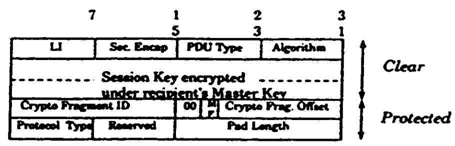

- Clear part
  - Algorithm
  - Encrypted Key
- Protected part: Frag./Reassembly to reduce chance of fragmentation by IP
- Trailer: DES padding, ICV

Charlie Kaufman

Requirements From Protocol to Allow a H/W Implementation:

- ICV at the end of the packet
- Large key-ID (to allow the use of an encrypted key)
- Single algorithm for ICV/encryption (e.g. DES)
- Fragmentation/reassembly mechanism

Charlie Kaufman

2.7.5 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 (NASs) 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, PCs 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", 07/02/1993, J. Vollbrecht, A. Rubens, G. McGregor <draft-ietf-nasreq-nasrequirements-01.txt>

CURRENT MEETING REPORT

Reported by John Vollbrecht/Merit Network and Allan Rubens/Merit Network

Minutes of the Network Access Server Requirements Working Group (NASREQ)

The NASREQ Working Group met on Tuesday, November 2. There was a brief review of the rationale for NAS/helper separation and Steve Willens walked through the proposed RADIUS protocol document that could be used to support this separation.

Steve provided copies of the document which will be updated and submitted as an Internet-Draft. There was a lot of discussion about the document. The general consensus was that it was a good idea to have such a protocol, that the protocol met a number of needs, and it should eventually be submitted for consideration as an RFC. Some of the issues raised were:

- Security:

  An MD5 hashing algorithm is used to hide the password. It was suggested that this might not be a good mechanism, and that it might not be exportable. It is not known where to get answers to these issues.

  Secrets shared between NAS and RADIUS server are configured rather than obtained from a authentication server. It was suggested was that this could be done either way, depending on whether the NAS is able to do Kerberos.

- Extensibility:

  A lot of discussion concerned whether parameters should be identified with ASCII strings or numeric IDs. This discussion will presumably continue on the mailing list.

- TCP versus UDP:

  A suggestion was made that the protocol should be built on TCP rather than UDP. This will be considered more on the mailing list, but consensus seemed to favor TCP.

- Downloadable filters:

  Filters should be dynamically settable.

- Other:

  The text of the document needs to clarify which attributes belong together, which are sent by the NAS, and which are returned by the RADIUS server.

  May want to be able to send an arbitrary string to be interpreted by the command interpreter in the NAS.

A very brief presentation of distributed authentication was presented as a possible future subject for the working group to consider. This was discussed further in the Security Area Advisory Group (SAAG) meeting on Thursday and we agreed to have this discussion at the first SAAG meeting in Seattle.

We discussed changing the charter of the group and the following elements were described as a possible direction:

- Finish the NAS Requirements document and submit it for consideration as an Informational RFC following the Seattle IETF. We need volunteers to work on pieces of the document.
- Revise the RADIUS protocol definition and submit it for consideration as an RFC after review at the Seattle IETF.
- Move KAP/PKAP to the Point-to-Point Protocol Extensions Working Group (PPPEXT) and/or to a working group in the Security Area. The group that it might go to in the Security Area is under discussion.
- Focus the attention of the group on distributed authentication in support of shared dialin between organizations. This will likely have other implications and should have significant support from security area folks to be successful.

Attendees

Nick Alfano	alfano@mpr.ca
Jim Barnes	barnes@xylogics.com
Larry Blunk	ljb@merit.edu
Cheng Chen	chen@accessworks.com
Blair Copland	copland@unt.edu
Robert Downs	bdowns@combinet.com
Antonio Fernandez	afa@thumper.bellcore.com
Jisoo Geiter	geiter@mitre.org
Mei-Jean Goh	goh@mpr.ca
Chris Gorsuch	chrisg@lobby.ti.com
Marco Hernandez	marco@cren.net
Matt Hood	hood@nsipo.nasa.gov
John Linn	linn@security.ov.com
Brian Lloyd	brian@lloyd.com
Glenn McGregor	ghm@lloyd.com
Piers McMahon	p.v.mcmahon@rea0803.wins.icl.co.uk
Michael Michnikov	mbmg@mitre.org
Bob Morgan	morgan@networking.stanford.edu
Michael O'Dell	mo@uunet.uu.net
Răkesh Patel	rapatel@pilot.njin.net

Allan Rubens	acr@merit.edu
William Simpson	Bill.Simpson@um.cc.umich.edu
Dave Solo	solo@bbn.com
Don Stephenson	don.stephenson@sun.com
Theodore Ts'o	tytso@mit.edu
Raymond Vega	rvega@cicese.mx
John Vollbrecht	jrv@merit.edu
Steve Willens	steve@livingston.com

2.7.6 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 issued 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", 10/26/1993, S. Crocker, N. Freed, J. Galvin <draft-ietf-pem-mime-03.txt>

"An Alternative PEM MIME Integration", 10/26/1993, J. Schiller <draft-ietf-pem-mime-alternative-00.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)

A quick poll of attendees indicated about 60-70% are new, i.e., have not previously attended a PEM Working Group meeting. A number of MIME Working Group members attended because of the MIME-PEM topic.

Implementation Status Reports

MIT: Jeff Schiller reported that the status was mostly unchanged from Amsterdam; anonymous FTP availability at MIT for a Mac implementation, integrated with TechMail messaging software (uses POP3 server); expected new release in a few weeks, with bug fixes and some new user interface features; an X-windows version will become available later.

TIS: Jim Galvin reported that version 6.1 was released on October 29; available via anonymous FTP from TIS for United States and Canadian users; an RFC 1421 implementation. A United Kingdom-developed version is under way at the TIS United Kingdom office, targeted for PCs.

No other PEM implementation representatives were present.

Electronic Notary Services

Dave Solo provided a presentation on various "notary-style" validation services for non-repudiation (see slides following the minutes):

- Simple time stamping
- Enhanced non-repudiation
- Document registration
- Archival signature validation
- Assurance issues
- Validation of other attributes

The group observed that non-repudiation with proof of submission and/or delivery are not viable services in much of the Internet because the submission and delivery agents are usually under the administrative control of the originator and recipient (or their respective organizations). Only if one has a timestamp notary which acts as a mail forwarder does one recapture the proof of submission notion, but that makes the notary an element of the MHS, and requires double enveloping by the originator (to direct the original message to the notary/forwarder).

MIME-PEM (The Saga Continues)

About 50% of attendees have read the Internet-Drafts issued on this topic last week.

- MIME-PEM "lite" (Jeff Schiller)

  This design requires very minimal changes to existing PEM and is capable of accommodating simple MIME-PEM constructs. It utilizes the Content-Domain construct to identify the payload of the message as requiring this enhanced form of processing. The goal is to add MIME capability to existing PEM implementations without substantial delay. INRIA and Bellcore report that they have implemented this design and found it quite simple. Some attendees note that most Macintosh clients don't have MIME and this approach minimizes the effort required for a (simple) PEM- MIME system. There was a suggestion to modify the current proposal to employ an "application/text" content type for MIC- CLEAR messages and use "application/1421" for MIC-ONLY and ENCRYPTED.

- MIME-PEM "full-bodied" (Steve Crocker)

  The goal is a design that preserves maximal functionality for PEM and MIME users, all MIME content types, etc. There is no backward compatibility goal. It does away with MIC-CLEAR and MIC-ONLY distinction, because MIME content transfer encoding addresses that requirement. It separates PEM header information from the message body. One major change from the previous design is use of "application/quoted-mime-entity" to make the PEM body opaque, protecting the body against MIME gateway transformations. Another change is the separation of certificate and CRLs into a separate body part. The constructs for encryption and signature are capable of being nested in either order, for forwarding signed, encrypted messages. Constructs allow for sending certificate chains and/or CRL chains plus use of the same facility for CRL and certificate queries.

[Working group Chair observations after the meeting: The thrust of the first proposal is preservation of the investment in current PEM implementations designed to operate with (vanilla) 822 messages, while extending these implementations to support MIME constructs in the simplest possible fashion. This proposal also addresses the processing of PEM-MIME messages by MIME mail readers that do not provide integrated support for PEM and by PEM implementations that do not provide integrated MIME support. The proposal is extremely simple to implement and is backwards compatible with the current PEM design, e.g., it makes use of the Content-Domain header construct to identify a MIME content. The second proposal represents a fairly radical departure from RFC 1421, essentially re-engineering PEM for the MIME environment, in order to provide more flexible security services for (extensible) MIME UAs. As a result, this design is not backward compatible with current 1421 implementations.

The path being pursued here, through these two proposals, does not converge in a single PEM implementation serving both basic 822 and MIME UAs. This is an unfortunate outcome, but it is the result of a long period of work by a number of individuals in both the PEM and MIME working groups. When MIME replaces basic 822 as the ubiquitous e-mail protocol throughout the Internet and the other networks that are e-mail-connected to the Internet, then the second proposal probably becomes the obvious choice, due to its increased functionality. However, prior to that time, the group will pursue dual approaches that accommodate distinct subscriber groups within the MIME-PEM user community.]

A Certificate Server Proposal for PEM

This proposal, presented by Christian Huitema, is designed to facilitate retrieval of certificates and CRLs with locally managed, simple databases. Index for search is the user's mailbox name. This calls for operators of the hosts that provide the user's mailbox to provide this responder facility. However, mail services such as CompuServ and MCIMail are unlikely to provide this service. There may be a need to create a new record type to allow indirection to other than the user's actual mailbox provider. Also, this proposal is based on TCP, but not all prospective PEM users are reachable by TCP, e.g., users of non-IP nets or firewall. A suggestion was made to add this facility to FINGER instead, to minimize firewall problems? Suggest e-mail-based access should be baseline, with real-time access an optional additional service.

Triple DES

Burt Kalaski was not available to lead discussion at this meeting so the topic was defered. This is still an important topic but the group is awaiting publication (later this year) of an analysis which is purported to reach conclusions at odds with those of the analysis prepared by Burt. In the mean time, all interested parties are encouraged to read the analysis posted to the PEM-DEV list during the last week of October.

Attendees

Garrett Alexander	gda@tycho.ncsc.mil
Alireza Bahreman	bahreman@bellcore.com
Stephen Crocker	crocker@tis.com
Donald Eastlake	dee@skidrow.lkg.dec.com
Erik Fair	fair@apple.com
Qin Fang	qin_fang@unc.edu
Antonio Fernandez	afa@thumper.bellcore.com
James Fielding	jamesf@arl.army.mil
James Galvin	galvin@tis.com

Jisoo Geiter	geiter@mitre.org
Chris Gorsuch	chrisg@lobby.ti.com
Terry Gray	gray@cac.washington.edu
Christian Huitema	Christian.Huitema@sophia.inria.fr
Neil Katin	katin@eng.sun.com
Charlie Kaufman	kaufman@zk3.dec.com
Stephen Kent	kent@bbn.com
Paul Lambert	paul_lambert@email.mot.com
John Linn	linn@security.ov.com
Steven Lunt	lunt@bellcore.com
Chip Matthes	chip@delphi.com
Wayne McDilda	wayne@dir.texas.gov
Michael Michnikov	mbmg@mitre.org
Keith Moore	moore@cs.utk.edu
Sandra Murphy	murphy@tis.com
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Jeffrey Schiller	jis@mit.edu
Wolfgang Schneider	schneiw@darmstadt.gmd.de
Dave Solo	solo@bbn.com
Theodore Ts'o	tytso@mit.edu
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
David Woodgate	David.Woodgate@its.csiro.au
Peter Yee	yee@atlas.arc.nasa.gov

Electronic Signature Timestamp Server Issues

David Solo
solo@bbn.com
617-873-6261

Possible Services

- Simple timestamping
- Enhanced Non-repudiation
- Document registration
- Archival signature validation
- Validation of other attributes

Timestamping

- Association of a timestamp with a document
- Granularity of timestamp
- What is it bound to?
- Synchronization/time source issues
- Authentication of time

Archival Signatures

- Goal is to be able to validate signatures far into the future
- Problems include
  - windows between CRLs
  - subsequent compromise of a user's key
  - subsequent compromise of a CA's key
  - preservation of applicable certification path and CRLs
  - relationship between expiration and revocation
  - association with applicable authorization information
- Need to support different assurance/risk tolerance goals

Assurance Considerations

- Validation of signatures
- Validity of time information
- Subsequent compromise of server's key
- Back-dating/forward-dating
- Malicious server
- Conspiring server and user
- Relationship to IPRA, PCA, CA
- Approaches
  - Risk assessment (different assurance levels)
  - Independent correlation
  - Operational practices
  - Mechanism enhancements

Functionality

- Request
  - desired services
  - plaintext message
  - ciphertext message
  - message hash
- Server may
  - apply timestamp
  - validate user's signature and certification path
  - validate authorizations
- Response
  - timestamp (or enhanced timestamp)
  - signature/authorization validation information
  - archival signature information
- Access via
  - electronic mail
  - other options

Issues

- Definition of services
- Specification of request/response syntax and semantics
- Specification of processing rules
- Time management approach
- Resolution of revocation window issues
- Assurance approaches

2.8 Service Applications Area

Director:

- David Crocker: dcrocker@mordor.stanford.edu

Area Summary reported by David Crocker/Silicon Graphics

The Service Applications Area encompasses support capabilities for system configuration and query and for structured inter-process communication.

Mail-Based File Distribution BOF (MAILFTP)

This BOF was tasked with discussing interest in developing a mail-based file transfer capability. It grew out of an initial effort by Marko Kaittola. The MAILFTP BOF started with a brainstorming on problems in the area. Identified issues were file name, size, type and structure, application type, gateways, bulk distribution, update mechanisms and file system hierarchies. A number of tools and protocols were identified which address some of the problems listed. Limited user need and expert time lead to the conclusion not to start a working group to solve the open problems in a coherent way.

Domain Name System Working Group (DNS)

The DNS Working Group discussed the status of the DNS MIB documents, the status of the DNS Security effort, the formation of the proposed DNS security working group, and the status of current SIPP DNS efforts. The group heard presentations from Ed King on DNS Vendor Requirements, and from Masataka Ohta on Dynamic Host Configuration. The Service Applications Area Director, Dave Crocker, also discussed a plan to conclude the current working group with the publication of the DNS MIB, and start the DNS Security effort as a newly-constituted working group. He explained that this was part of the general IETF effort to limit use of on-going, open-ended working groups, instead favoring ones with specific focus.

Minimal OSI Upper-Layers Working Group (THINOSI)

The THINOSI Working Group is pursuing use of that portion of the OSI stack which is above Transport and below specific user applications. It encompasses Session, Presentation, ACSE and ROSE. The THINOSI meeting was small but effective. As discussed on the mailing list, the present cookbook will be split into two documents. The profile aspects (i.e., subsetting of the OSI upper-layers as used by the basic communication applications) would be split out as a small document that just cites the OSI standards and other stable open documents. This would be targeted on the standards track. The rest of the cookbook (most of it) would be progressed as Informational. Discussion also showed that as well as the THINOSI target of reducing the processing needed to deal with the standard OSI protocols, there was a requirement in many circles for reduced bandwidth (i.e. a different protocol) to support some of the OSI upper-layer functions—many people had assumed this was what THINOSI was doing. The THINOSI mailing list will be opened for discussion on this, at least until an appropriate forum is established.

Service Location Protocol Working Group (SVRLOC)

The SVRLOC Working Group is developing a resource location protocol. The working group held two sessions at this IETF. The first went into the details of the Internet-Draft that has been submitted by the chairs. There were few comments about the wire specification as described except for issues about the size of several fields and the scaling issues of the protocol, and the use of the directory agent in the protocol. These issues were addressed in the afternoon session which was more design oriented. A proposal for directory agent discovery and use was worked out and will be written up by the working group chair. The issue of directory agent and service discovery out of the campus was punted as future work.

2.8.1 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 Working 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 Support for IDPR", 10/26/1993, R. Austein <draft-ietf-dns-idpr-02.txt>

"DNS Server MIB Extensions", 07/08/1993, R. Austein, J. Saperia <draft-ietf-dns-server-mib-01.txt>

"DNS Resolver MIB Extensions", 07/19/1993, R. Austein, J. Saperia <draft-ietf-dns-resolver-mib-01.txt>

"Incremental Transfer and Fast Convergence in DNS", 12/13/1993, A. Kumar, S. Hotz, J. Postel <draft-ietf-dns-ixfr-01.txt>

Request For Comments

RFC 1480 "The US Domain"

RFC 1536 "Common DNS Implementation Errors and Suggested Fixes."

RFC 1537 "Common DNS Data File Configuration Error"

CURRENT MEETING REPORT

Reported by Rob Austein/Epilogue Technology

Minutes of the Domain Name System Working Group (DNS)

Documents

Three new DNS-related Informational RFCs have come out recently. RFC 1535 (also known as "the EDU.COM emergency RFC") details problems with a widely-used but ill-advised DNS search heuristic, and proposes a solution. RFC 1536 details common DNS implementation errors, with proposed solutions; this document was accepted as a DNS Working Group project at the 27th IETF (Amsterdam), completed, and accepted on the mailing list. RFC 1537 details common DNS configuration errors; while it was never formally accepted as a DNS Working Group document, it was reviewed by the working group members. These three RFCs are closely related and cross-reference each other, so, on advice of the RFC Editor, the DNS Working Group Chair approved "fast track" publication of these documents on behalf of the Working Group. If anybody has serious problems with these documents, blame it on the Chair.

Dave Perkins reported on the current status of the DNS MIB documents on behalf of the Network Management Area Directorate (NMDIR). Basically, there are no remaining hard problems, just some remaining detail work. One of the authors, Rob Austein, has received a detailed list of remaining concerns, none of which appear to be show-stoppers. It should be possible to get these documents out the door before the 29th IETF in Seattle. Dave pointed out two design issues that are not objections but of which he thinks the DNS Working Group should be aware:

1. Due to SNMP protocol limitations, the length limits on DNS names used as indices to "conceptual tables" in the MIBs will be shorter than the DNS name length limit of 255 octets. Based on analysis of current usage, this should not be a problem, so we'll flag it with a warning statement in the document but otherwise leave it alone.
2. The most recent versions of the documents (not yet released as Internet-Drafts) use the SNMPv2 SMI rather than the SNMPv1 SMI, in order to clear up some problems with unsigned 32-bit integers. NMDIR wants to be sure that the DNS Working Group realizes that this means only SNMPv2-capable implementations will be able to use these MIBs.

DNS Security Sub-Group

James Galvin gave a report on the meeting held by the DNS Security "sub-group" (a spin off from the DNS Working Group at the 26th IETF in Columbus).

Editor's Note: The minutes of the DNS Security meeting are available via FTP or mail server from the remote directories as /ietf/dns/dns-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The DNS Security effort will be spun off as a separate working group in the Service Applications Area (SAP), as soon as James can get the charter approved. The DNS Security mailing list is `dns-security@tis.com`; requests to subscribe should be sent to `dns-security-request@tis.com`.

Discussion of the incremental zone transfer protocol (draft-ietf-dns-ixfr-00.txt) was deferred because none of the authors were present at the meeting. Comments on this draft should be sent to the authors and/or the Namedroppers mailing list.

DNS Efforts to Support SIPP

Sue Thomson gave a brief report on current DNS efforts to support SIPP (the merger of the SIP and PIP proposals). See the latest version of the Internet-Draft, draft-ietf-sip-sippdns-nn.txt, for details.

DNS Reliability Issues - Boeing

Ed King gave a presentation on DNS reliability issues in Boeing's production environment. Ed has to support DNS on a corporate network with thousands of subnets and DNS software from many vendors in a production environment that never shuts down and where an interruption to DNS services due to a power hit can leave hundreds of engineers sitting idle waiting for their workstations to finish booting. Much of the problem is that each vendor has their own slightly different (and often more than slightly broken) interface between DNS, local host tables, and the vendor's own pet name resolution mechanism. Replacing or repairing all the DNS software in an environment isn't economically feasible, so the most constructive approach seems to be to write a "DNS Requirements" document to use as a reference when pressuring vendors to fix their DNS implementations. The DNS portion of the Host Requirements documents (RFC 1123 section 6.1) and the newly published DNS "Common Errors" Informational RFCs are good starting points, but companies like Boeing need a document that has the force of a standard and that goes into more detail on interface design issues than Host Requirements does.

No definite decision was reached as a result of Ed's presentation, but watch Namedroppers for further discussion and probably a call to form a working group.

DNS Support for DHC and Mobile Hosts

Masataka Ohta gave a presentation on a possible way to implement some of the DNS support needed for dynamic host configuration and mobile hosts. The presentation went into more detail than there is room for in these minutes, so expect to see a summary of this on the Namedroppers list.

The Future of the DNS Working Group

Dave Crocker spoke about the future of the DNS Working Group. As has been discussed at previous meetings, the DNS Working Group as currently organized doesn't really fit well into the current IETF organizational framework. Accordingly, Dave asks that DNS reorganize itself more along the current IETF pattern. The proposal is to move the "permanent" functions of the DNS Working Group (DNS oversight within the IETF, mostly) into the SAP Area Directorate, that Dave will be forming "Real Soon Now," while reincarnating specific closed-ended tasks as separate working groups within the SAP Area. The SAP Area Directorate will hold open meetings at regular intervals, so that there will still be a forum for overall DNS design work. For formal purposes, the current DNS Working Group will probably be retroactively construed as having been the DNS MIB Working Group, and will be closed down as soon as the DNS MIB documents hit the streets. As a practical matter, and in the Chair's opinion, the current DNS Working Group will effectively reconstitute itself as the attendees of the DNS portion of the SAP Area Directorate open meetings. Dave expects to have the reorganization completed by the 29th IETF in Seattle.

The discussion that followed Dave's statement made it clear that there are people with strong feelings on both sides of this issue (keep the DNS Working Group as it is versus reorganize per Dave's plan). Unless somebody feels strongly enough about this to make a formal appeal, the reorganization will probably go through.

Attendees

Steve Alexander	stevea@lachman.com
Garrett Alexander	gda@tycho.ncsc.mil
Robert Austein	sra@epilogue.com
Anders Baardsgaad	anders@cc.uit.no
Alireza Bahreman	bahreman@bellcore.com
William Barns	barns@gateway.mitre.org
Stephen Crocker	crocker@tis.com
Donald Eastlake	dee@skidrow.lkg.dec.com
Havard Eidnes	havard.eidnes@runit.sintef.no
Erik Fair	fair@apple.com
Roger Fajman	raf@cu.nih.gov

Patrik Faltstrom	paf@nada.kth.se
Antonio Fernandez	afa@thumper.bellcore.com
James Fielding	jamesf@arl.army.mil
James Galvin	galvin@tis.com
Chris Gorsuch	chrisg@lobby.ti.com
Ronald Jacoby	rj@sgi.com
Rick Jones	raj@cup.hp.com
Charlie Kaufman	kaufman@zk3.dec.com
Elizabeth Kaufman	kaufman@biomded.med.yale.edu
Stephen Kent	kent@bbn.com
Edwin King	eek@atc.boeing.com
Paul Lambert	paul_lambert@email.mot.com
Walter Lazear	lazear@gateway.mitre.org
Lars-Johan Liman	liman@ebone.net
John Linn	linn@security.ov.com
Jun Matsukata	jm@eng.isas.ac.jp
Paul Mockapetris	pvm@darpa.mil
Sath Nelakonda	sath@lachman.com
Masataka Ohta	mohta@cc.titech.ac.jp
Michael Patton	map@bbn.com
Jon Postel	postel@isi.edu
Jeffrey Schiller	jis@mit.edu
Richard Schmalgemeier	rgs@merit.edu
Michael St. Johns	stjohns@arpa.mil
John Stewart	jstewart@cnri.reston.va.us
Theodore Ts'o	tytso@mit.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
David Woodgate	David.Woodgate@its.csiro.au
Weiping Zhao	zhao@nacsis.ac.jp

2.8.2 MHS-DS (MHSDS)

Charter

Chair(s)

Kevin Jordan: `Kevin.E.Jordan@cdc.com`
Harald Alvestrand: `Harald.T.Alvestrand@uninett.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 RFC 822 use of Directory and Directory support for RFC 822 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

"A simple profile for MHS use of Directory", 07/08/1993, S. Kille <draft-ietf-mhsds-mhsprofile-03.txt, .ps>

"Representing Tables and Subtrees in the Directory", 07/07/1993, S. Kille <draft-ietf-mhsds-subtrees-03.txt, .ps>

"Representing the O/R Address hierarchy in the Directory Information Tree", 07/07/1993, S. Kille <draft-ietf-mhsds-infotree-03.txt, .ps>

"Use of the Directory to support mapping between X.400 and RFC 822 Addresses", 07/07/1993, S. Kille <draft-ietf-mhsds-supmapping-03.txt, .ps>

"Use of the Directory to support routing for RFC 822 and related protocols", 07/08/1993, S. Kille <draft-ietf-mhsds-822dir-03.txt, .ps>

"MHS use of Directory to support MHS Routing", 07/19/1993, Steve Kille <draft-ietf-mhsds-routdirectory-03.txt>

"MHS use of Directory to support MHS Content Conversion", 07/08/1993, S. Kille <draft-ietf-mhsds-convert-01.txt, .ps>

"Introducing Project Long Bud Internet Pilot Project for the Deployment of X.500 Directory Information in Support of X.400 Routing", 06/21/1993, H. Alvestrand, K. Jordan, S. Langlois <draft-ietf-mhsds-long-bud-intro-00.txt>

2.8.3 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.

Internet-Drafts

"Octet sequences for upper-layer OSI to support basic communications applications", 11/29/1993, P. Furniss <draft-ietf-thinosi-cookbook-02.txt>

"Use of upper-layer OSI protocols to support basic communications applications", 11/29/1993, P. Furniss <draft-ietf-thinosi-profile-00.txt>

CURRENT MEETING REPORT

Reported by Peter Furniss/Consultant

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

Status and Content of Upper-layer Cookbook

The question of whether the cookbook was a parallel specification of the OSI protocols or a mixture of profile and implementation advice was finalised, following earlier e-mail discussions.

It transpired that the Area Director, Dave Crocker, and several others had originally thought that the group was specifying an alternative protocol to provide the OSI upper-layer functions, and they had been surprised to discover the protocol was the same, or a subset of the standard protocols.

The group accepted Dave's view that there could be no such thing as "re-specification" of a protocol—there was only one defining text. Anything which restated, without modification, what was in the original specification was really an implementation guide. As such it could become an Informational RFC, but not a standards track document.

However, the cookbook also subsets the standard protocols, and in this respect is similar to the profiles produced by the OSI "regional workshops"—OIW, EWOS and AOW. Such subsetting is protocol specification, and Dave said would be suitable for standards track if the base documents are:

- The result of an open process.
- Stable (for Proposed status); or published standard (for IETF Standard status).

Dave suggested the model of the SMI RFC, which cites the OSI ASN.1 Basic Encoding Rules and defines a subset of them, could be taken as an example of the style.

Peter Furniss had already produced a first attempt at separating the profile-like aspects of the cookbook, treating the OSI base standards and the two general profiles (Common Upper-Layer Requirements - Part 1 and Part 3) as the cited documents. Both the ISO and regional workshop processes meet the openness requirement. However, although CULR Part 1 (general requirements) is stable, and about to begin draft International Standardised Profile ballot, Part 3 is not stable and is some way from international ballot. It was agreed that Peter would expand the profile, citing just the base standards and CULR-1 and reproducing the CULR-3 restrictions. It includes a few further restrictions that are not currently in CULR-3. Peter would submit this as an Internet-Draft. Peter will also revise the cookbook again, referring to the other document.

Charter Revision

The decisions above require revision of the charter to reflect what the group is actually doing. It was agreed that since the "thinDAP" work has not progressed, and in any case would be an Informational RFC, it should be dropped from the work plan.

Peter will work up a draft revision and post it to the list, (really).

Application-specific Mapping Documents

There was no clear view on the possibility of defining application- specific mapping documents (e.g., how to use Z39.50 with the cookbook). Peter will explain this idea to the mailing list.

Reduced-OSI

Following the realisation that many people had expected, and wanted, the group to investigate alternative, lighter, protocols, it was agreed the mailing list would be opened to provide an (interim) forum for discussion of this. Walter Lazear had identified at least 7 different groups (mostly .gov or .mil) interested in this, and at least two others were represented (electric power companies) or known of (civil air-ground). Walter had a one-page summary of this to hand out. The concentration would be on trying to establish what the requirements really were—i.e., which parts of OSI function were still wanted.

This widening of the list will not be formally part of the working group work plan. It is just taking advantage of a mailing list that people may have thought was doing it anyway!

Parallel Documents

- Jim Quigley reported that CULR-3 had been revised further, and the OIW ULSIG were concentrating on getting the compliance/conformance terminology sorted out.

- The X/Open XTI/mOSI specification is still about to be published as a preliminary specification.

- Jim Quigley reported that ITU-T SG7 was planning to make the cookbook into a Recommendation.

Implementations

- Terry Sullivan (Florida Center for Library Automation) released his "tosi" implementation three days previously

- Peter has started extending the X/osi code to a more general thinosi implementation, with XTI/mOSI as the upper interface. He hopes to have it available around the end of the year.

Next Time

The group does not expect to meet in Seattle.

Attendees

David Crocker	dcrocker@mordor.stanford.edu
Richard desJardins	desjardi@eos.nasa.gov
Peter Furniss	p.furniss@ulcc.ac.uk
Walter Lazear	lazear@gateway.mitre.org
Mark Needleman	mhn@stubbs.ucop.edu
Dan Nordell	
James Quigley	Quigley@cup.hp.com
Brien Wheeler	blw@mitre.org

2.8.4 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 internets. 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.

2.8.5 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 for consideration for standards status in the Internet community.

This working group has a number of specific objectives: to provide an Internet-Draft 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 for publication as an RFC.
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"

2.8.6 Service Location Protocol (SVRLOC)

Charter

Chair(s)

John Veizades: veizades@ftp.com
Scott Kaplan: scott@ftp.com

Mailing Lists

General Discussion: srv-location@ftp.com
To Subscribe: srv-location-request@ftp.com
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

"Service Location Protocol", 10/19/1993, J. Veizades, S. Kaplan <draft-ietf-svrloc-protocol-02.txt>

CURRENT MEETING REPORT

Reported by John Veizades/FTP

Minutes of the Service Location Protocol Working Group (SVRLOC)

The Service Location Working Group held two sessions at this IETF meeting. The first session focussed on the Internet-Draft submitted by the chairs, "Service Location Protocol."

This discussion was divided into several areas:

- The Base Protocol Specification
- Authentication
- Predicate Language
- Directory Agents

The base protocol specification had several issues in the size of several fields. The locale field was extended to a 16-bit value with the need to look at any international specifications that define language specifiers. (The chair found that ISO 639 defines many language specifiers and that the mime working group is working on extending this work to include dialects of spoken languages increasing the size of this field to four characters.)

The group made the recommendation that the specification writers look at the work that the Common Authentication Technology Working Group (CAT) is doing as a way of defining the authentication types and also to look at the GSS API work. The authentication length field was increased to 16 bits.

The predicate language was hashed out, and the addition of wild card methods on strings was added. The conditional operators that were accidentally left out of the specification were also added.

The afternoon session was devoted to the discussion of the directory agent interaction in the protocol. The issues that needed to be resolved were that the directory agent is the point of scaling in the protocol and that it is necessary that the directory agents need to solve all the scaling issues.

The following protocol was defined for directory agents:

- Directory agents have a concept of a scope that they are responsible to support. The scope is a text string. Samples of a scope would be "engineering" or "marketing." During a directory agent discovery, the directory agent passes to the user agent the scopes that are available on the network and the scope(s) that they are acting for. The user agent needs to find the directory agent that supports the scope they are searching in. This is done by sending a service query with the directory agent and the scope that is interested in the appropriate directory agent responds.

- Directory agents need to find themselves and exchange the scopes that they know of.

- Service agents need to register with all directory agents that support the scope they have chosen to be in.

The next question is, "How do directory agents and sites advertise themselves in the Internet at large?" The proposals that were considered were:

- Advertising in a directory service like DNS or X.500, and
- Advertising in a resource discovery service like gopher or WWW.

The discussion continued with the statement that sites may want to advertise their service in several taxonomy domains, these include a white pages system indexed by the public name of the advertiser and/or a yellow pages type service indexed by service provided, geographical location, etc.

The chairs will post an updated draft by the end of the year. Implementation of this protocol is proceeding at FTP Software and other implementors are being sought out.

Attendees

Andy Adams	ala@merit.edu
Steve Alexander	stevea@lachman.com
James Allard	jallard@microsoft.com
Frank Ciotti	frankc@telxon.com
Chuck de Sostoa	chuckd@cup.hp.com
Mei-Jean Goh	goh@mpr.ca
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Kevin Jackson	kjackson@concord.com
Scott Kaplan	scott@ftp.com
Gordon Lee	gordon@ftp.com
Glenn Mansfield	glenn@aic.co.jp
Matt Mathis	mathis@psc.edu
Mike Ritter	mwritter@applelink.apple.com
Michael Safly	mike.safly@msfc.nasa.gov
Martin Schulman	schulman@smtp.sprint.com
Robert Stevens	robs@join.com
Larry Tepper	ltepper@compatible.com
John Veizades	veizades@ftp.com
Walter Wimer	walter.wimer@andrew.cmu.edu

2.8.7 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", 03/01/1993, Fred Glover <draft-ietf-tnfs-spec-03.txt>

2.9 Transport Area

Director:

- Allison Mankin: `mankin@cmf.nrl.navy.mil`

Area Summary

The following Transport Area working groups met in Houston:

- Audio/Video Transport Working Group (AVT)
- Multiparty Multimedia Session Control Working Group (MMUSIC)

Three BOFs under the Transport Area were held in Houston:

- Realtime Packet Forwarding and Admission Control BOF (REALTIME)
- RSVP - Resource Reservation Setup Protocol BOF (RSVP)
- TCP Multiplexing BOF (TMUX)

CURRENT MEETING REPORT

Reported by Robert Braden/Information Sciences Institute

Minutes of the Real-time Packet Forwarding and Admission Control BOF (REALTIME)

The demand for multimedia communication, and the success of IETF audio/videocasts, will soon create an urgent requirement for resource reservation and control in the Internet. From an architectural viewpoint, this represents a new Internet service model. Such a service model should include, in an integrated fashion, both real-time and link-sharing services along with the traditional best-effort datagram services. Research in DARTnet has developed (a) an integrated service model for the Internet, and (b) a particular set of mechanisms to realize this model.

To provide end-to-end service suitable for realtime applications, the routers must all implement the same service model, although there may be alternative mechanisms. The group therefore proposes that the service model be standardized. This BOF will describe the service model and the realization, and suggest the service model as a candidate for Internet standardization.

The session was opened by Scott Shenker, who noted that Bob Braden, one of the BOF chairs, could not be present due to concerns about his home in the current LA fires.

Scott delivered a talk on the necessity of explicit service models in the Internet. He described the taxonomy of service models that are proposed in the related Internet-Draft. He stressed that this proposal is not for a closed set of service models, but for a growing set of explicit models. He also mentioned dissenting opinions from the loyal opposition, who would prefer an approach without explicit service models or explicit service reservations.

Scott continued by presenting and discussing three fundamental questions.

1. Why do we need a service model? The answer is efficiency. He argued that implementing an explicit mechanism is more effective than expending bandwidth in improving overall service, and this mechanism can only be introduced in the context of a model that defines the objective of the mechanism.

2. Should the service model be explicit or implicit? He advocated that binding between application and service should be determined by the application, outside the network, rather than internally to the network.

3. Is admission control necessary? He argued that, at least for heavy users, occasional blocking is a much more economical approach than vast over-provisioning.

He then took a straw poll which, while informal, suggested that there was not wide dissension to his conclusions to the above questions.

A discussion of service models followed. There was support for the concept of admission control.

David Clark then gave a talk on the way in which routers must be constructed in order to realize the service models discussed earlier. He used the guaranteed real-time service as an example of behavioral characterization of router functionality.

More discussion followed. It was observed that behavioral characterization of functionality is a very difficult intellectual problem, and that it was important that the community not get bogged down in this exercise. We must start to implement and deploy routers and get real experience at the same time we work towards a formal performance characterization.

It was observed that while the presentation had emphasized router requirements, there will also be a need for "subnetwork requirements" or "link level requirements." In the past, the Internet has demanded relatively little of its subnet technology, but these QOS requirements will change that. ATM, in particular, must fit into this architecture in a harmonious way.

The BOF ended at 9:30, due to exhaustion of all parties. There was an informal assessment that a working group ought to be formed.

Attendees

Anthony Alles	aalles@cisco.com
Robert Braden	braden@isi.edu
Scott Brim	Scott_Brim@cornell.edu
Theodore Brunner	tob@thumper.bellcore.com
Stephen Casner	casner@isi.edu
John Chang	jrc@uswest.com
David Clark	ddc@lcs.mit.edu
Ed Ellesson	ellesson@vnet.ibm.com
Roger Fajman	raf@cu.nih.gov
Mark Garrett	mwg@faline.bellcore.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Robert Hinden	hinden@eng.sun.com
Phil Irey	pirey@relay.nswc.navy.mil
Van Jacobson	van@ee.lbl.gov
Ronald Jacoby	rj@sgi.com
Frank Kastenholz	kasten@ftp.com
Mark Laubach	laubach@hpl.hp.com
Bryan Lyles	lyles@parc.xerox.com
Dan Magorian	magorian@ni.umd.edu
Andrew Malis	malis@maelstrom.timeplex.com
David Marlow	dmarlow@relay.nswc.navy.mil
Thomas Maslen	maslen@eng.sun.com

Keith McCloghrie	kzm@hls.com
Greg Minshall	minshall@wc.novell.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Steve Parker	sparker@ossi.com
Drew Perkins	ddp@fore.com
Allyn Romanow	allyn.romanow@eng.sun.com
Hal Sandick	sandick@vnet.ibm.com
Henning Schulzrinne	hgs@research.att.com
Andrew Smith	asmith@synoptics.com
Michael Speer	michael.speer@sun.com
Abel Weinrib	abel@bellcore.com

A New Internet Service Model

Scott Shenker
David Clark
Lixia Zhang

joint work with many others

IETF November 1993

Integrated Services Internet Project

MIT:
- David Clark
- John Wroclawski

USC/ISI:
- Bob Braden
- Deborah Estrin
- Shai Herzog
- Danny Mitzel

Xerox PARC:
- Steve Deering
- Sugih Jamin
- Scott Shenker
- Lixia Zhang

Background

New generation of applications:

- Audio, Video, Teleconferencing, FAX, etc.
- Wide variety of service requirements

The current Internet cannot efficiently support these applications

- Single *best-effort* service class
- Cannot tailor service to application needs

We propose changing the service model!

Why focus on the service model?

Service Model: set of services offered by the network

Service model functions as interface between applications and network

Stable service model allows rapid progress both above and below

- Will outlive particular instantiations
- Most fundamental piece of architecture
- Based on application requirements
  - Inherently speculative

Application Taxonomy and SCZ Service Model

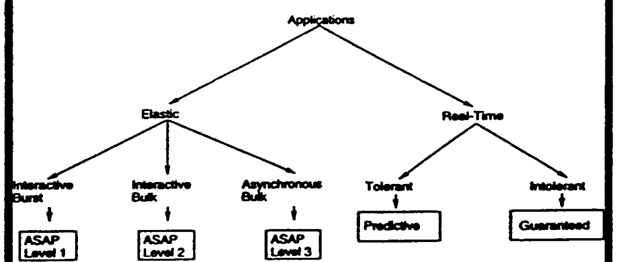

Elastic Applications:

- ASAP: *As soon as possible*

Realtime Applications:

- Intolerant: Guaranteed Service
  - *Worst-case* delay bound
- Tolerant: Predictive Service
  - Delay bound based on predictions

Link Sharing

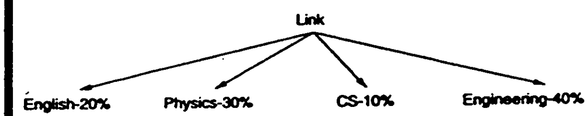

Link-Sharing:

- Controls the relative usage of a link
- Allows unused allocations to be shared
- Designed for collective entities
  - firms, protocols, etc.
- Network management interface
- Sharing can be done hierarchically

A New Internet Service Model: Three Fundamental Questions

- (Why) do we need a new service model?
- Should service be explicitly requested, or implicitly supplied?
- Is admission control necessary?

Why do we need a new service model?

Efficiency!

Definitions for simple model:

- s_i service given to flow i
- $U_i(s_i)$ performance of the application
- Total Value: $V = \sum_i U_i(s_i)$

Richer service model gives higher value of V

- Allows the network to deliver service tailored to application needs

Tradeoff between bandwidth and mechanism

- Amount of bandwidth needed to offset richer service model depends on the U_i
- Our guess is that for the new generation of applications, the rate of return on mechanism is very large.

Should service be explicitly requested or implicitly supplied?

Consider two possibilities

Option # 1: Have categories of traffic:

- e.g., interactive burst, interactive bulk, asynchronous, realtime, ...
- Network determines which category based on port number, gives service accordingly
- *Network in control*

Option # 2: Have categories of service:

- Application/user makes a specific service request
- *User in control*

We advocate option # 2

- mapping between application and service should not reside in the network

Two (Unfortunate) Implications

User incentives must be considered:

- Need some form of incentive system
- Requires accounting infrastructure

Service model must be stable (extensible):

- Driven by need for backward compatibility
- Service model will become standardized (formal or de facto)

Is admission control necessary?

Overloading can be prevented by:

- Overprovisioning
- Using admission control to deny access

Two subquestions:

- Can/should one overprovision everywhere?
- Should we "share-the-pain" or deny access?

Two time-scales:

- Short-term (~five years): speculation
- Long-term (?): sheer fantasy

Can/should one overprovision everywhere?

Short-Term:

- Demand will increase as video becomes standard and fast LAN's are deployed
- Not-Overprovisioned spots will abound
- Not a technical problem but a financial one

Long-Term:

- Phone network is not a good model
- Internet load will be much more variable
- Overprovisioning would lead to drastic underutilization

Common counterargument:

- Blocking must be rare to satisfy customers
- Reply: thick vs. thin markets

	Short-Term	Long-Term
Overprovision?	No	Probably not

Should we "share-the-pain" or deny access?

This is a matter of math, not philosophy!

- Precise mathematical criterion for which design choice produces the greater good
- Caveat: this criterion involves the utility functions of applications

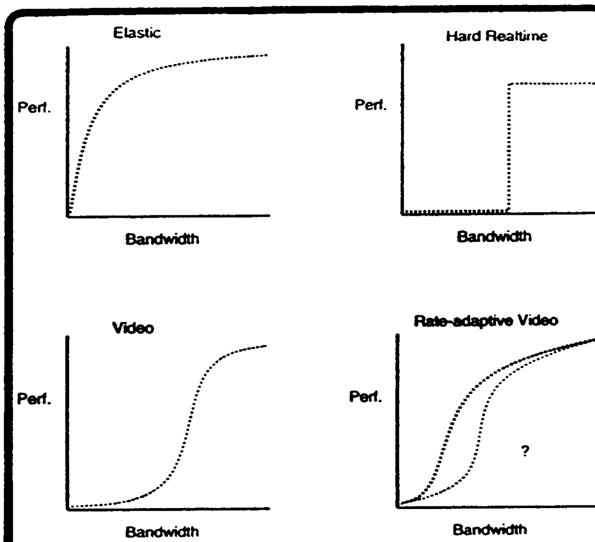

The details of these curves depend on the user, but the general shapes are probably invariant.

	Short-Term	Long-Term
Share-the-pain?	No	Perhaps not

To advocate no admission control, one must claim that:

- No need for migration path
- Most applications will look like elastic

Design Philosophy:

- Let future applications evolve without undue constraints
- Applications don't need to use admission-controlled services
- Market may eventually turn away from admission-controlled services
  - Will have some unused mechanism
- Risk of not offering admission-controlled services is much higher
  - Market will go elsewhere!

Straw Poll

1. We should change/extend the service model.
2. We should let users explicitly request service.
3. We should use admission control for some services.

Translating Global Service Requirements into Local Router Requirements

David Clark etc.

IETF Nov 93

What is the problem?

End-to-end service definitions are neat, but...

To ship product (router or end-node) must specify behavior at device level.
- Need open standards
- Multi-vendor operation required

Two objectives:
- Define what must defined.
- Do not over-define.

Product differentiation depends on parts being undefined....

Is this hard??

Specifying correctness is not too hard:
- TCP state machine

Performance specifications are much harder:
- TCP Slow start congestion avoidance

Two approaches:
1) Specify the objective:
" A sending TCP shall detect and adapt to congestion in such a way that congestive collapse of the network shall not occur."

2) Specify the algorithm.

Specify QOS algorithms?

Some protocols, like RSVP, have a traditional state machine, and need a correctness specification.

Some aspects, like admission control, should not be constrained to just one algorithm.

- Very complex algorithm to optimize bandwidth.
- Very cheap algorithm (just say yes) if links are very lightly loaded.
- Offer more conservative or aggresive options.

This seems like a good place for competitive product differentiation.
- Give examples, but not mandatory algorithms.

How to specify local router requirements

Must define operationally. Wrong time to be abstract.

Options:
- Algorithm (but we do not prefer this approach...)
- Testable behavior.

Our job (the IETF working group)!!
- Define necessary local behavior such that global service is achieved.
- Define tests necessary to confirm this behavior.

Guaranteed service

Global service definition: If source is constrained by token bucket shaper, the maximum delay is bounded.

Local service behavior:
For any burst size B, max delay D, allocated rate R, max packet size L, and outgoing link capacity C:

$$D \leq B/R + L/R(1+R/C)$$

In words: A maximal burst must not have a delay greater than
- The time for the burst to drain out at the overload rate, plus
- A term that accounts for the "phase delay" in getting to the link.

The Integrated Services Internet

Testing the router

To make the foregoing concrete, how could we test a router for conformance?

At a minimum:

- Set up a number of guaranteed flows.
- Deliver a burst on every one at once (the proven worst case)
- Measure the maximum packet delay.
- Repeat for a number of rates and burst sizes.
- Repeat with other services turned on.
- Etc.

Note: we are testing for a bound that is supposed to be absolute. So test should be supplemented with statement from vendor as to why it will always be met.

Slide 7

The Integrated Services Internet

Testing the "phase delay"

Load switch with lots of flows (so that lots of packets arrive)

Inject one packet on test flow.

Measure its delay.

Repeat. Find maximum.

Slide 8

The Integrated Services Internet

Admission control

Must test to see how much guaranteed bandwidth the box will allocate.
- If it accepts lots of flows, it is probably more useful.
- If it accepts very few flows, it may be cheaper.

(It is easy to meet these bounds on an underloaded network.)

Slide 9

The Integrated Services Internet

Approximation

A switch need not be exactly within these bounds.

It could be different, so long as the bound is known.
- Explicit bound is needed for guaranteed service.

For example:
- A box might have a phase delay term within a known constant of L/R(1+R/C).

For example:
- Our scheme used WFQ to isolate the guaranteed flows. Priority might yield tolerable bounds for very light loads.

But if bounds are different, how can they be discovered?

Slide 10

The Integrated Services Internet

Architecture questions

1) Should we specify the delay formulation or let vendor modify it?

If vendor can modify formulation, how complex is it to compute actual bound?
- Vendors get more implementation options.
- We may need distributed protocol to compute the bound. (Ugly...)

2) Does the boundary enforce the burst size?
- If so, router needs to work properly with oversized burst.
- If not, must test for operation in this region.

Slide 11

The Integrated Services Internet

Other services

This presentation used guaranteed service as an example.

We need a similar development for:
- Predictive service.
- Link sharing.
- Hierarchical link sharing.
- Priority elastic.

and so on.

This development for guaranteed service is only a first cut.

Slide 12

Discussion questions

Should vendors reveal their algorithm?
- I vote: vendor's choice

Should we let vendors use a different delay formula?
- I vote no. Stick with one.

Should we assume that the burst size is enforced as part of overall architecture?
- I vote no. It is a more robust answer.

CURRENT MEETING REPORT

Reported by Robet Braden/Information Sciences Institute

Minutes of the Resource Reservation Setup Protocol BOF (RSVP)

This BOF introduced a new protocol, RSVP, designed for setting up resource reservations in the Internet. It is a necessary component of a proposed extension of the Internet architecture to support integrated service, i.e., to support delay-sensitive applications.

Dave Clark introduced the speakers, noting that Bob Braden and Deborah Estrin had returned to LA to deal with the risk of fire to their homes. The talk that Braden was going to give was instead given by John Wroclawski, Scott Shenker and Lixia Zhang. John gave the introduction, Scott discussed the RSVP model, and Lixia Zhang discussed the RSVP protocol itself. John returned to discuss the future direction of the working group.

There were questions confirming the basic paradigm that RSVP lives on top of multicast routing, and that the PATH message serves the purpose of assuring that RSVP can know the reverse path of the multicast route.

Merging, a complex topic, received some clarifying discussion.

Noel Chiappa proposed that a flow identifier would be a better means of classifying packets. There was considerable discussion concerning flow IDs. It was proposed that while a flow ID is a dandy optimization, it was a mistake to make it a requirement.

It was noted that security may hide some of the fields in the packet that might be used for packet classifying, but there must be some field in the packet used to select the proper decryption key, which would equally serve to classify a packet. It was noted that on a fragmented packet, some of the fields may be missing on all but the first fragments. Currently, MTU discovery is being used to avoid this problem. This is consistent with the current trend in the Internet away from fragmentation.

There was discussion of the design decision within RSVP that the receiver rather than the sender should make the reservation. Some situations were proposed in which the sender might be in a better position than the receiver to make the reservation. The distinction was made that while in some cases it may be more direct for the sender to make the actual request, even in those cases it was the receiver that starts the process, and that understands what the request should be.

There was a discussion of "route pinning," which describes the objective of fixing the paths alone which resource reservations have been made, in order to assure that the reservation remains in place. Some members of the audience had concluded that RSVP did not intend to achieve this sort of functionality, which led to the conclusion that the "quality" of the assurance that RSVP would achieve would be less than that of a protocol such as ST-II. Lixia clarified this point, stating that it was the intention of the RSVP designers that the assurance quality of an RSVP guarantee be very robust. However, they were of the opinion that RSVP should not contain mechanism to prevent "route flapping," but that this ought to be a service of the routing protocol. More discussion followed, and this topic was noted for further discussion. It was stressed from the floor that RSVP must architect its behavior on route loss.

The final discussion topic concerned whether RSVP could rapidly respond to network events, since it used timers to preserve its state. The presenters clarified this point, stressing that while RSVP used timers and refresh messages to maintain state, this did not preclude the use of event-driven messages to trigger recovery to such things as lost links. It is in fact the intention of RSVP to use event-driven messages for this purpose.

A proposed working group agenda was presented, and there was general acceptance of forming a working group along those lines.

Attendees

Andy Adams	ala@merit.edu
Steve Alexander	stevea@lachman.com
Anthony Alles	aalles@cisco.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
William Barns	barns@gateway.mitre.org
Tom Benkart	teb@acc.com
Lou Berger	lberger@bbn.com
Luc Boulianne	lucb@cs.mcgill.ca
Robert Braden	braden@isi.edu
Monroe Bridges	monroe@cup.hp.com
David Bridgham	dab@epilogue.com
Al Broscius	broscius@bellcore.com
Steve Buchko	stevebu@newbridge.com
Stephen Casner	casner@isi.edu
Isidro Castineyra	isidro@bbn.com
John Chang	jrc@uswest.com
Ping Chen	ping@ping2.aux.apple.com
J. Noel Chiappa	jnc@lcs.mit.edu
David Clark	ddc@lcs.mit.edu
James Davin	davin@thumper.bellcore.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Luca Delgrossi	luca@ibmpa.awdpa.ibm.com
Taso Devetzis	devetzis@bellcore.com
Ed Ellesson	ellesson@vnet.ibm.com
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
William Fenner	fenner@cmf.nrl.navy.mil

James Forster	forster@cisco.com
Paul Francis	Francis@thumper.bellcore.com
Ron Frederick	frederick@parc.xerox.com
Dan Frommer	dan@isv.dec.com
Mark Garrett	mwg@faline.bellcore.com
Fengmin Gong	gong@concert.net
Ramesh Govindan	rxg@thumper.bellcore.com
Joel Gyllenskog	jgyllens@hpdmd48.boi.hp.com
John Hanratty	jhanratty@agile.com
Dimitry Haskin	dhaskin@wellfleet.com
Ken Hayward	Ken.Hayward@bnr.ca
Juha Heinanen	juha.heinanen@datanet.tele.fi
Shai Herzog	herzog@catarina.usc.edu
Russ Hobby	rdhobby@ucdavis.edu
Van Jacobson	van@ee.lbl.gov
Ronald Jacoby	rj@sgi.com
Matthew Jonson	jonson@ddn.af.mil
Frank Kastenholz	kasten@ftp.com
Yasuhiro Katsube	katsube@mail.bellcore.com
David Kaufman	dek@magna.telco.com
Lee Kilpatrick	leekil@bbn.com
David Kristol	dmk@allegra.att.com
Ted Kuo	tik@vnet.ibm.com
Jim Martin	jim@noc.rutgers.edu
Thomas Maslen	maslen@eng.sun.com
Donald Merritt	don@arl.army.mil
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Vijayaragavan Pandian	vjp@proteon.com
Craig Partridge	craig@bbn.com
Laura Pate	pate@gateway.mitre.org
Maryann Perez	perez@cmf.nrl.navy.mil
Drew Perkins	ddp@fore.com
J. Mark Pullen	mpullen@cs.gmu.edu
Bala Rajagopalan	braja@qsun.att.com
Doris Roland	droland@ghost.darpa.mil
Allyn Romanow	allyn.romanow@eng.sun.com
Shawn Routhier	sar@epilogue.com
Allan Rubens	acr@merit.edu
Hal Sandick	sandick@vnet.ibm.com
Eve Schooler	schooler@isi.edu
Henning Schulzrinne	hgs@research.att.com
Isil Sebuktekin	isil@nevin.bellcore.com
Michael See	mikesee@vnet.ibm.com
Kanan Shah	kshah@cmf.nrl.navy.mil
Scott Shenker	shenker@parc.xerox.com

Ming Sheu	msheu@vnet.ibm.com
W. David Sincoskie	sincos@bellcore.com
Karen Sollins	sollins@lcs.mit.edu
Michael Speer	michael.speer@sun.com
Louis Steinberg	louiss@vnet.ibm.com
John Stewart	jstewart@cnri.reston.va.us
Matsuaki Terada	tera@sdl.hitachi.co.jp
Claudio Topolcic	topolcic@cnri.reston.va.us
Jerry Toporek	jt@mentat.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
Abel Weinrib	abel@bellcore.com
Taehwan Weon	weon@cosmos.kaist.ac.kr
Richard Woundy	rwoundy@vnet.ibm.com
John Wroclawski	jtw@lcs.mit.edu
Jean Yao	yao@cup.hp.com
Shinichi Yoshida	yoshida@sumitomo.com
Jessica Yu	jyy@merit.edu
Lixia Zhang	lixia@parc.xerox.com

RSVP

Resource ReSerVation Protocol

Bob Braden, USC Information Sciences Institute

Houston IETF Meeting
November 3, 1993

Introduction: Outline

- o Role of RSVP
- o History
- o RSVP Assumptions
- o Two Fundamental Ideas

Introduction

Purpose of RSVP:

Set up *resource reservations* in routers and hosts

We believe that resource reservation will be necessary in the Internet, to support many future applications.

The reasons for this belief, and its consequences, will be the subject of tonight's realtime BOF. Here we will assume it, and describe a protocol for setting up reservations.

Role of RSVP

Internet Integrated Services:

- Add: *real-time* service and *link-sharing* to the Internet

Components:

- TRAFFIC CONTROL -- *packet-by-packet*
  - — Packet Scheduling Mechanism
  - — Admission Control
  - — Classifier
- RESERVATION SETUP PROTOCOL -- RSVP

Role of RSVP

ROUTER IMPLEMENTATION MODEL

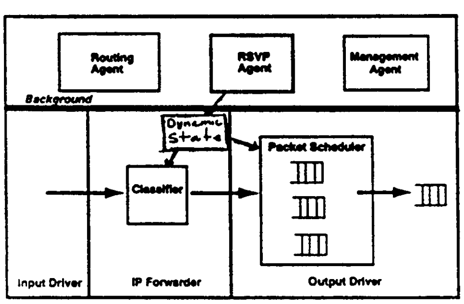

History of RSVP

- Discussions in DARTnet community and End-to-End RG of IRTF
- Lixia Zhang (PARC): tested protocol concepts via simulation

  "RSVP: A New Resource ReSerVation Protocol", Lixia Zhang, Steve Deering, Deborah Estrin, Scott Shenker,and Daniel Zappala, IEEE Network, September 1993.
- Sugih Jamin (USC/PARC): experimental implementation
  - RSVP within modified SunOS 4.1.1 kernel [DARTnet]
  - Kernel included CSZ packet scheduler *[tonight's BOF]*
  - Demonstrated in DARTnet June 8, 1993.

History (cont'd)

- Shai Herzog (ISI): new RSVP implementation
  - Daemon process
  - Demonstrated for ARPA Oct 26, 1993

MANY PEOPLE HAVE CONTRIBUTED TO RSVP EFFORT, including:

Lixia Zhang, Scott Shenker, Deborah Estrin, Dave Clark, Steve Deering, Sugih Jamin, Shai Herzog, Bob Braden, Daniel Zappala, and Craig Partridge

RSVP ASSUMPTIONS

- IP MULTICASTING IS FUNDAMENTAL
  - Generalizing 1->1 to 1->n is not effective design paradigm. Instead, make unicast a special case.
- DECOUPLED FROM ROUTING PROTOCOLS
  - Can use existing routing protocols.
  - Can use future routing protocols -- e.g, CBT or ESL or ...
- MUST SUPPORT HETEROGENEOUS RECEIVERS
  - Important to handle diversity.

RSVP Assumptions (cont'd)

- MUST SCALE TO VERY LARGE NUMBER OF RECEIVERS

  => Frequent group membership changes

  => O(n) or better growth of overhead (n = size of group)
- SUPPORT '*CHANNEL SWITCHING*'
  - Idea from Jon Crowcroft: each receiver can independently select, *within the network*, which subflow it wants to receive.
- TRUTH, JUSTICE, and the INTERNET WAY:
  - robustness, modularity, extensibility, ...

Two Fundamental Ideas in RSVP

- SOFT STATE IN ROUTERS
  - End systems must periodically refresh router state
  - Route change => refresh msg sets up state on new route
- RECEIVERS CONTROL RESERVATIONS
  - Simplicity and directness: receivers know what they want.
  - Matches IP multicasting, which is receiver-oriented.
  - Efficiency: Can combine reservations for different senders.

RSVP Model: Outline

- o Reservation Parameters
- o Sessions
- o Reservation Styles

RSVP Communication Paradigm

RESERVATION message flows from receivers to sender.
- Carries flow spec (the resource request).
- Follows info laid down by path message.

PATH Message flows from sender to receivers.
- Follows unicast/multicast IP routing info.
- Establishes reverse path for RSVP.

RSVP Session

RSVP uses simplex model

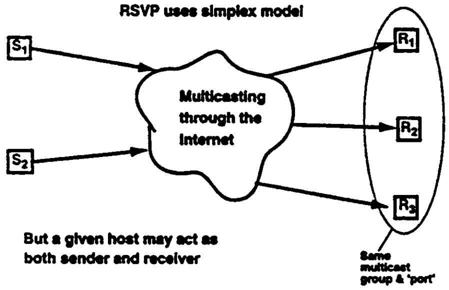

But a given host may act as both sender and receiver

(Assume a given session)

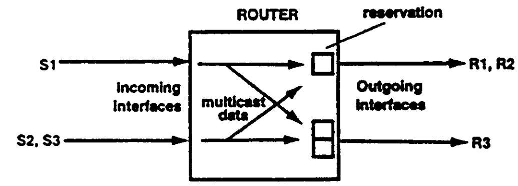

In reality: same physical interface may be both incoming and outgoing interface for RSVP.

Reservation Parameters

- FLOWSPEC
  - Defines quality of service (bandwidth, delay, burstiness...)
  - Opaque to RSVP
- FILTERSPEC
  - Sets Classifier to select subset of data packets, eg

    { Src-IP-addr / Src-port , Dst-IP-addr / Dst-port}
    or
    { Dst-IP-addr / Dst-port } (Src is '*wildcard*')
  - Partially opaque to RSVP

Reservation Parameters

- Engineering issue: general form of Filterspec in RSVP
  - (Address, port): too simple
  - (bit mask and value): limited generality
  - Pseudo-code ?

- RESERVATION STYLE
  - Determines how reservations are made and shared.
  - Currently 3 types, more may be added in future
    - o WILDCARD FILTER
    - o FIXED FILTER
    - o DYNAMIC FILTER

Reservation Styles

- Wildcard Filter (WF) Style

  Filter: { Dst-IP-addr / Dst-port }

  Creates single shared reservation 'pipe', e.g., for audio.

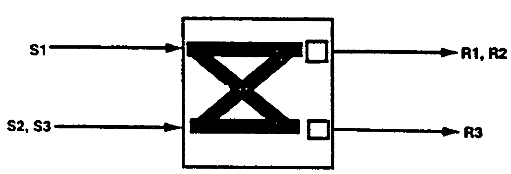

Reservation Styles

Filter (FF) Style

- Each receiver explicitly lists senders it wants to see (=> *filterspecs*)
- Different sender's reservation channels are *not shared*– e.g., video.

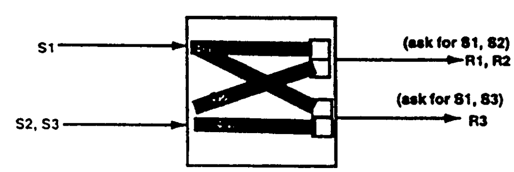

Reservation Styles

Dynamic Filter (DF) Style

Channel-switching style

Reserve n non-shared 'channels' of capacity *(e.g., for video)*

Choose/change selection of particular senders to use these.

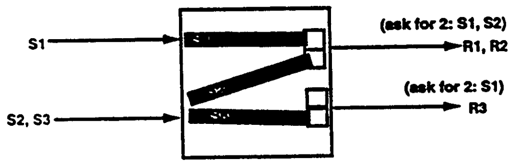

Host Model

Host must:

- Register as a sender (if appropriate)
  - => RSVP starts sending PATH msgs
- Join multicast cast group (if a receiver)
- Request reservation (if a receiver)
  - => RSVP starts sending RESV msgs

Host may:

- Change reservation request at any time

Interfaces

- API
- Traffic Control kernel functions
  - o Add/delete Flowspecs (=> admission control)
  - o Add/delete Filterspecs
  - o Abstract operations on Flowspecs
- Routing
  - o 'Tell me the route to...'
- Forwarding

Interoperation

Want to handle partial RSVP coverage in Internet

– PATH IP datagrams addressed to uni-/multicast destination

– Requires special intercept for Protocol 46 in routers

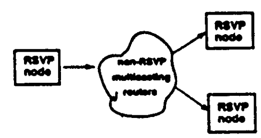

RSVP Protocol: Outline

- o RSVP Messages
- o Merging
- o Managing Soft State
- o API Model
- o Interoperability
- o Interfaces

RSVP Message Types

- RESV
  - – Follow reverse data paths to selected senders
  - – Leave *reservation state* in each router
- PATH
  - – Follow forward data paths
  - – Leave *path state* in each router, to route RESV messages.
- ERR
- TEAR

RSVP MESSAGES

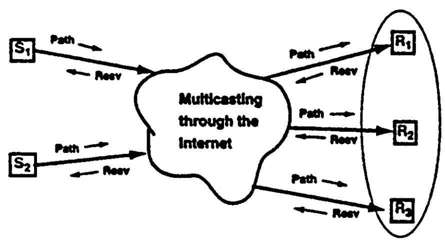

PATH Messages

- Session ID: ***Uni-multicast dest addr + 32-bit number***
- Refresh period, TTL interval
- Previous hop IP address

  Fundamental path state, for routing RESV msgs
- Drop/BE bit

  Controls what happens to session traffic that does not match a filter.
- List of Sender Descriptors: (Template, FlowSpec)

RESV Messages

- Session ID
- Refresh period, TTL interval
- Receiver address

  To establish ownership of receiver-specific reservation state (DF), and to prevent looping RESV messages.
- Reservation Style
- Number of DF reservations
- List of Flow Descriptors: (FilterSpec, FlowSpec)

Merging

To limit RSVP traffic: ***Merge*** PATH, RESV messages at each router.

– Same session, reservation style

Example: Wildcard Filter Style:

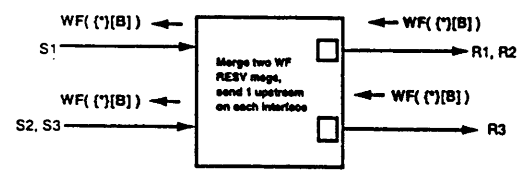

Merging Example

E.g. audiocast: only one WF message arrives at sender!

Example: Fixed Filter Style:

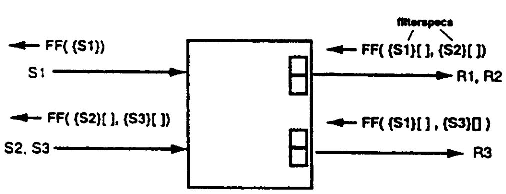

RSVP Messages

- RSVP messages sent as datagrams directly over IP.
- RSVP messages are given some preferred class of service (with other critical network control messages)

Managing Soft State

- Merging requires gathering message state at each node =>
  Node forwards (merged) message at end of the current *refresh period*.
- State that is not refreshed times out after TTL interval.
- TTL interval > 3 * refresh period, to allow loss of individual RSVP messages.

Future

- (Some) open issues
- Working group agenda
- References

WG Open Issues

- User identification/credential field.
- Generalize and bound filterspec capabilities.
- Improve event-driven response behavior.
- Interfaces to rest of system.

Future Routing Protocol Improvements

- Provide route-change signals to RSVP
- Provide routes that are known to support RSVP (TOS routing)
- Source-controlled route selection when setup fails
- Inter-twine routing with route selection/setup

Draft Agenda for RSVP Working Group

Primary goal: Evolve RSVP specification and introduce it into the Internet standards track.

Plan/Milestones:

Nov 93: Begin discussions via email

Jan 94: New draft prepared

Apr 94: Meet at Seattle IETF, discuss latest draft. Decide whether to request publication as a Prototype RFC. Plan strategy for attacking open issues.

July 94: Coordinate with Router Requirements WG.

Nov 94: Submit RSVP Spec for Proposed Standard status.

References

- Zhang, Deering, Estrin and Zappala, RSVP: A New Resource ReSerVation Protocol, IEEE Network, September 1993
- Zhang, Braden, Estrin, Herzog, Jamin, Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification,
  Internet Draft: file draft-braden-rsvp-00.ps
- Mitzel, Estrin, Shenker, Zhang, An Architectural Comparison of ST-II and RSVP, submitted to Infocom 94

CURRENT MEETING REPORT

Reported by Jim Barnes/Xylogics

Minutes of the TCP Multiplexing BOF (TMUX)

Agenda

- Introduction
- How we got here
- TMux Overview and implementation experience
- Issues from the mailing list
- What now

Peter Cameron presented a short overview of the TMux protocol and gave a summary of the implementation experiences to date. TMux has been implemented in Unix System V.4 streams and BSD 4.3 systems. The total number of implementations to date is six. Peter noted that since the interface between the IP and TCP layers is not well defined, implementing a portable TMux module can be difficult. A problem with FTP traffic was also noted when there was a single FTP session. The implementations of TMux now do not attempt to multiplex FTP traffic. Only Telnet and Rlogin data is multiplexed now.

The results from some performance tests simulating multiple Telnet sessions were reviewed. The following discussion resulted in requests for additional performance information including perceived response times for TMux versus non-TMux situations. Peter took the action item to send additional performance numbers to the mailing list.

After the presentation, a general discussion followed. The following significant points were raised:

- Since we want to prevent intermediate nodes from fragmenting and reconstructing TMux frames, the "do not fragment" flag should be set.

- The document needs to include an applicability statement.

- If the TMux implementation begins to see timeouts with exactly one datagram in the packet (that is, there is little traffic to multiplex), TMux should be turned off.

- A packet with the IP OPTIONS field is not a candidate for multiplexing with TMux.

- Check the test implementations to make sure everything that was done to overcome an implementation problem is reflected in the draft document.

Discussion then moved on to consider specific points raised on the mailing list.

- Length field: The consensus is that the length field should be 16 bits.
- Checksums: After considerable discussion, no real consensus was reached, so the checksum field will stay in.

A lengthy message from Don Eastlake was posted to the mailing list just before the BOF. The issues raised and the consensus reached during the BOF were as follows:

- The document is too terminal server centric. The consensus of the BOF attendees was that TMux was a simple solution for a very specific problem. The problem definition should remain tightly focussed.

- Type of Service concerns. TMux should ensure that all datagrams within the multiplexed packet have the same TOS.

- Broadcast packets. Only unicast addresses should be allowed.

- Larger limit on the maximum size of TMux datagrams. The maximum size of 30 will be replaced with information gained during implementation. This max datagram size probably should be configurable.

- Use TMux only in congested situations. Agreed.

- The section on security needs clarification. Agreed.

The attendees were asked if there were any blocking issues that would prevent TMux from being put on the standards track. No one raised any such issue and the consensus was that TMux could be recommended to the IESG for further action.

Attendees

Steve Alexander	stevea@lachman.com
Jim Barnes	barnes@xylogics.com
David Borman	dab@cray.com
Peter Cameron	cameron@xylint.co.uk
Lida Carrier	lida@apple.com
Alan Clegg	abc@concert.net
David Crocker	dcrocker@mordor.stanford.edu
Chuck de Sostoa	chuckd@cup.hp.com
Marc Hasson	marc@mentat.com
Jeff Hilgeman	jeffh@apertus.com
William Kwan	kwan@rabbit.com
David Lapp	lapp@waterloo.hp.com

Walter Lazear	lazear@gateway.mitre.org
Allison Mankin	mankin@cmf.nrl.navy.mil
Matt Mathis	mathis@psc.edu
Marjo Mercado	marjo@cup.hp.com
William Palter	palter@tgv.com
Kanan Shah	kshah@cmf.nrl.navy.mil
Kitty Shih	kmshih@novell.com
Hoe Trinh	htrinh@vnet.ibm.com
John Vollbrecht	jrv@merit.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
Weiping Zhao	zhao@nacsis.ac.jp

TMux Overview

TMux

TECHNICAL OVERVIEW

Pete Cameron
Xylogics International Ltd.
Email: cameron@xylint.co.uk
Phone: +44 908 222112

TMux Overview

Background

- Over the last couple of years, many people have identified a host overhead problem when using large numbers of Telnet and Rlogin sessions into host systems.
- At the Amsterdam IETF in July, a BOF was held on the CMP proposal to overcome this problem by multiplexing Telnet and Rlogin sessions using a protocol layered between the application and TCP.
- Dave Crocker also made a presentation, in which he suggested that the multiplexing should be moved. This is called TMux.
- Xylogics has implemented and evaluated both protocols on it's in-house systems and they give almost identical performance.

TMux Overview

The Proposed solution: TMux

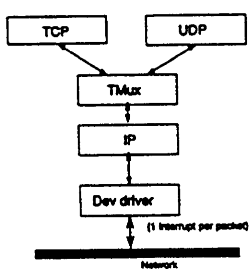

- TMux multiplexes packets from one system to another.
- TMux uses a delay to multiplex all segments into a single packet.
- The number of packets the host has to deal with is greatly reduced, lowering host and network load.

✓ Preliminary performance results are very encouraging.

TMux Overview

Brief details of the protocol

- Each TMux packet is sent as a single IP datagram containing multiple transport segments, each preceded by a short TMux mini-header, ie:
  | IP hdr | TMux hdr | Tport seg | TMux hdr | Tport seg | ...

```
| 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
+-------------------------------+
|   Reserved    |  Length high  |
+-------------------------------+
|          Length low           |
+-------------------------------+
|          Protocol ID          |
+-------------------------------+
|           Checksum            |
+-------------------------------+
```

- Each of these mini-headers contains all the information required to recreate the transport segment when received by the remote end.

- See Internet Draft *draft-cameron-tmux-01.txt* for full details.

TMux Overview

Implementation Overview

- There are now 2 separate implementations for:

  Unix System V.4 streams
  BSD 4.3 Unix.

- These have now been ported to 6 different operating systems from 5 manufacturers, including multi-threaded.
- Both implementations proved to be relatively easy, both versions took about 2 man weeks to code a simple working version. Production ready code is 2000 lines of well commented C code for the V.4 streams version and 1600 lines for the BSD version.

TMux Overview

Problems

- Problem with the protocol found to be that if it is used with a single FTP session, it increases the time taken to transfer the file by about 15% due to the TMux delay (though multiple sessions are faster). The test implementations now do not delay FTP type traffic. This eliminates the problem, but slightly reduces the potential host performance gain when FTP and Telnet traffic is mixed.
- Because the interface between TCP and IP is not defined, unlike the TPI interface between applications and TCP, it is more difficult to make a portable TMux module than it would be to make a portable CMP module.

TMux Overview

Performance Comparisons

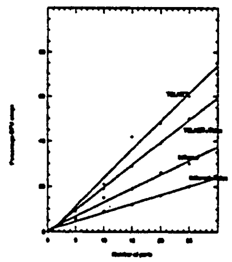

- Tests simulate typists at 5 CPS from single system (Xylogics Annex) to a second system (33MHz SPARC Unix V.4).
- Figures are % host CPU load.
- 24% increase in number of ports supportable for same CPU load for out of kernel protocols (Telnet).
- 55% increase for in-kernel protocol (Xylogics protocol).

TMux Overview

Conclusion

- TMux is simpler than CMP, and gives same performance, so is more suitable for standardisation.
- Because the network software has only to deal with one packet for several ports, the number of interrupts, and hence the host overhead is reduced.
- Test results show TMux gives significant savings on CPU overhead when many sessions are multiplexed from one host to another.

TMux Overview

What next

☻ We would like to propose that TMux moves forward to the standards track.

2.9.1 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:pub/mailing-lists/mail-archive/rem-conf

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", 10/21/1993, H. Schulzrinne <draft-ietf-avt-issues-01.txt, .ps>

"RTP: A Transport Protocol for Real-Time Applications", 10/21/1993, H. Schulzrinne, S. Casner <draft-ietf-avt-rtp-04.txt, .ps>

"Media Encodings", 09/17/1993, H. Schulzrinne <draft-ietf-avt-encodings-02.txt>

"Sample Profile for the Use of RTP for Audio and Video Conferences with Minimal Control", 10/21/1993, H. Schulzrinne <draft-ietf-avt-profile-03.txt>

"Packetization of H.261 video streams", 12/10/1993, T. Turletti, C. Huitema <draft-ietf-avt-video-packet-02.txt>

CURRENT MEETING REPORT

Reported by Steve Casner/USC-ISI

Minutes of the Audio/Video Transport Working Group (AVT)

The AVT Working Group met for only one session at this meeting since the draft specification for the Real-time Transport Protocol (RTP) is nearly completed for submission as an RFC. The emphasis of this session was on implementation experience with the focus shifting to companion specifications for profiles and encodings.

Status of Draft RTP Specification

This group did not meet in Amsterdam, but there has been substantial progress on the RTP specification via e-mail and a teleconference, and a new draft-ietf-avt-rtp-04.txt and .ps has been installed. The specification has been submitted to the Area Director with a request for "IESG Last Call," and is in review by the Directorate.

Steve Casner gave a brief description of the most recent change to the specification, which was the addition of the APP option. This option allows experimental application-specific options to be defined without official registration while avoiding conflicts with other option definitions. See the draft RTP specification for details. A brief description was also given on a proposal from Andrew Cherenson to add an option, not in the main RTP specification but in the audio/video profile, to indicate the mode or state of a participant. The proposed set of states were: active, video frozen (still image), private (listening but not sending), and hold (not listening and not sending).

A good fraction of the attendees at this meeting had read the RTP specification. Comments were solicited both on the specification and on the two options just described, but no comments were offered. However, behind the scenes, some objections have been raised to the classification of RTP as a Proposed Standard and to certain details of the specification. These issues will be discussed further on the mailing list.

Implementation Experience

Ron Frederick from Xerox PARC gave a presentation on his experience with implementing RTP in the network video (nv) program. He reported that overall, the implementation went very cleanly, and that the combination of the sequence number, timestamp and sync bit worked well together. He found the option format easy to generate and parse, but cautioned that the parser must watch out for an illegal option length zero or length greater than the packet length. (The example option parsing code in the appendix to the specification includes these checks.)

The one nuisance Ron found was that the program needs to know if an SSRC option is present to fully identify the sender before the parsing can act upon the other options. This requires parsing the options twice, or storing the information while parsing and then acting upon it at the end. To reduce this nuisance, it was proposed that the specification be modified to require that if an SSRC option is present, it must follow immediately after the fixed header. Since this is the logical place for translators to insert the SSRC option, and since there can be only one, this restriction should cause no difficulties.

David Kristol from AT&T described his work (just beginning) on a quality of service monitor for RTP. It would create a map of the MBONE, and display a measure of the reception quality for each receiver on the map using data obtained from reception reports multicast by the receivers. This would allow a visual determination of bottleneck points. One observation was that the measure of video delay is affected by the use of the same timestamp on all packets of a video frame even though the packets are not transmitted at the same time. A solution is to measure delay only on the first packet of a frame. This illustrates that reception quality measurement may be dependent upon the medium.

Dave also implemented a vat/RTP translator to allow participation in vat audio sessions inside the AT&T firewall. This turned out to be very simple, the only problem being translation of vat's beginning-of-talkspurt flag into RTP's end-of-talkspurt flag. For now, he is just copying the bit and ignoring the distinction.

Encoding Specifications

Frank Kastenholz from FTP Software asked for the addition in the audio/video profile of an 8-bit linear encoding ("L8") and a format code for L8 encoding at 11.025 KHz. This matches the capability of common audio hardware on PC and Mac platforms. It is possible to convert in software to 8-bit mu-law at 8 KHz, but this increases the minimum processing power required to participate. This request was generally agreed upon, and Frank was requested to provide the details to go into the profile. Henning Schulzrinne cautioned that adding a new "standard" encoding places a burden on all implementations to include at least a decoder for it.

Bill Fenner from NRL and Ron Frederick gave presentations on carrying JPEG video over RTP, and on the issues to be addressed in an encoding specification. Although the JPEG specification includes a variety of formats, Ron recommended that we stick with 4:2:2 video format, square pixels (as produced by most of the chips even though CCIR 601 specifies rectangular pixels), a 16x8 block as the minimum coded unit, and progressive scan. Ron also recommended that we use the Q factors defined for C-JPEG and D-JPEG by the Free JPEG Group and use the standard Huffman coding table, though these could be overridden by custom table definitions.

Bill has designed an encoding for JPEG over RTP, and implemented it using the Parallax JPEG hardware. He points out that JPEG frames are large, so they are likely to require segmentation and reassembly. Losing one packet out of a frame will result in frame loss because the Huffman reset mechanism that is part of the standard does not provide enough sequence space for packet-size losses. He also observed that the Q factor does not provide much usable quality range (the picture gets lots uglier without the frame rate increasing as much as one would expect).

The encoding Bill defined uses the same RTP timestamp on all packets of a frame, and the RTP sync bit indicates the last packet of the frame, as usual. In addition, he has defined a small header to go at the beginning of the data in the first packet of a frame. The presence of this header is indicated by the first two bytes being one of the application-specific codes (0xFF 0xE1) provided in the JPEG specification and guaranteed not to appear in the data. This code is followed by two bytes to encode the Q factor, Huffman table index, and some size information. Special values of these indices can be used to indicate that custom quantization and/or Huffman tables will follow. The mechanisms for requesting and/or periodically retransmitting custom tables are still to be decided and tested. There were no major objections to this design other than the suggestion that explicit image width and height factors be included. Bill agreed to produce a first draft specification for JPEG over RTP with assistance from Ron and Fengmin Gong from MCNC.

Video Decoder API

In Columbus we had a good discussion on the feasibility of creating a common interface for software video decoders so that each packet video program can incorporate decoders for many or all of the other programs' native formats to enable interoperation. At this meeting, Ron Frederick gave an update on the decoder API in the nv program in which decoding and rendering of the image data are 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 YUV pixels for a portion of the image. A callback routine is provided to render a rectangular portion of the image after decoding.

Ron identified several open issues that have arisen:

- Is YUV a good choice for color decoding? It allows easy rendering into monochrome or color images, but requires extra processing for encodings that would more naturally use RGB or dithered data. The difficulty is that number of variations in the rendering code is already large to handle variations in pixel depth and ordering. It may not be worthwhile to double or triple this to render from additional input formats.

- It is desirable to enable the use of hardware encoders and/or decoders for increased performance, but what additional hooks are required to fit this into the model? Some answers may come from exploring the options for the SunVideo board Cell-B encoder and for JPEG video using the Parallax board as Bill Fenner has done.

- Should the common code handle resequencing of packets? Previously, nv ignored packet sequencing because packets of the nv encoding can be processed out of order.

Now, nv is processing the sequence numbers to accumulate packet loss information, and could do the resequencing. However, Ron feels that this function should be left to the decode routines because the requirements may not be the same for all encodings, unless we can define as part of the profile an extra level of framing for all the encodings to use.

Other API's may also be needed. Henning suggested that video encoding routines should also be sharable to reduce the effort of writing them. Since nv already separates the frame grab from the encoding, an interface could be explored there. Abel Weinrib also pointed out that we need API's at a higher layer, that of whole media agents to be controlled by different session managers.

Report from IMA Network Focus Group

A the end of the session, we got a report from Thomas Maslen of Sun on the recent first meeting of the IMA Network Focus Group, and on the potential interaction with the IETF AVT and MMUSIC Working Groups' activities. The Interactive Multimedia Association (IMA) is an industry group chartered to develop standards to support multimedia applications. In particular, the Multimedia System Services (MSS) proposal defines an object-oriented architecture for the infrastructure to support multimedia applications.

In a way, the MSS work fits between the AVT and MMUSIC areas. The MSS proposal does not specify media transport mechanisms or protocols. The Network FoG is to address the requirements for network transport in the MSS, and to define network transport interfaces, target environments and protocol profiles to support those requirements. The group will work with other standards groups, including the IETF, to incorporate existing protocols and cooperate on the definition of new ones where needed. At first look, it appears that RTP may be suitable as one of the protocols to be used for transport of real-time media.

Similarly, MSS provides infrastructure for multimedia applications such at teleconferencing, but does not include the applications themselves. Abel pointed out that it does not include higher-level objects like people in its model, nor does it include policies. Therefore, MMUSIC sits above MSS, and the session management mechanisms to be developed in that working group might be used for communication among a set of applications implemented using MSS.

Future Working Group Activity

The session closed with a discussion of future working group activity. As work on the RTP specification is completed, the group's emphasis will shift to profile and encoding specifications. From the point of view of our Area Director, Allison Mankin, it is appropriate for the group to continue work as needed, or to go on hiatus but keep the mailing list (rem-conf) active. Meetings at future IETFs may then be called to address new questions such as the interface between network real-time services and RTP, or when appropriate to advance any of the specifications through the standards process.

Attendees

Andy Adams	ala@merit.edu
Stephen Batsell	batsell@itd.nrl.navy.mil
Tom Benkart	teb@acc.com
Richard Binder	rbinder@cnri.reston.va.us
Ronald Broersma	ron@nosc.mil
Stephen Casner	casner@isi.edu
Ping Chen	ping@ping2.aux.apple.com
Chuck de Sostoa	chuckd@cup.hp.com
Stephen Deering	deering@parc.xerox.com
David Dubois	dad@pacersoft.com
Ed Ellesson	ellesson@vnet.ibm.com
Julio Escobar	jescobar@bbn.com
Roger Fajman	raf@cu.nih.gov
William Fenner	fenner@cmf.nrl.navy.mil
James Fielding	jamesf@arl.army.mil
Robert Fink	rlfink@lbl.gov
Ron Frederick	frederick@parc.xerox.com
Mark Garrett	mwg@faline.bellcore.com
Atanu Ghosh	atanu@cs.ucl.ac.uk
Shawn Gillam	shawn@timonware.com
Robert Gilligan	Bob.Gilligan@Eng.Sun.Com
Fengmin Gong	gong@concert.net
Darren Griffiths	dag@ossi.com
Regina Hain	rrosales@bbn.com
Shai Herzog	herzog@catarina.usc.edu
Phil Irey	pirey@relay.nswc.navy.mil
Rick Jones	raj@cup.hp.com
Frank Kastenholz	kasten@ftp.com
David Kaufman	dek@magna.telco.com
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Charley Kline	cvk@uiuc.edu
Michael Kornegay	mlk@bir.com
David Kristol	dmk@allegra.att.com
Allison Mankin	mankin@cmf.nrl.navy.mil
David Marlow	dmarlow@relay.nswc.navy.mil
Jim Martin	jim@noc.rutgers.edu
Thomas Maslen	maslen@eng.sun.com
Marjo Mercado	marjo@cup.hp.com
Greg Minshall	minshall@wc.novell.com

Dan Nordell	
Marsha Perrott	perrott@prep.net
J. Mark Pullen	mpullen@cs.gmu.edu
Jim Rees	Jim.Rees@umich.edu
Eve Schooler	schooler@isi.edu
Henning Schulzrinne	hgs@research.att.com
Michael Speer	michael.speer@sun.com
John Stewart	jstewart@cnri.reston.va.us
Matsuaki Terada	tera@sdl.hitachi.co.jp
Chuck Warlick	chuck.warlick@pscni.nasa.gov
Abel Weinrib	abel@bellcore.com
Jean Yao	yao@cup.hp.com
Shinichi Yoshida	yoshida@sumitomo.com

2.9.2 Multiparty Multimedia Session Control (MMUSIC)

Charter

Chair(s)

Eve Schooler: `schooler@isi.edu`
Abel Weinrib: `abel@bellcore.com`

Mailing Lists

General Discussion: `confctrl@isi.edu`
To Subscribe: `confctrl-request@isi.edu`
Archive: `venera.isi.edu:~/confctrl/confcrtl.mail`

Description of Working Group

The demand for Internet teleconferencing has arrived, yet an infrastructure to support this demand is barely in place. Multimedia session control, defined as the management and coordination of multiple sessions and their multiple users in multiple media (e.g., audio, video), is one component of the infrastructure. The Multiparty Multimedia Session Control Working Group is chartered to design and specify a protocol to perform these functions.

The protocol will provide negotiation for session membership, underlying communication topology and media configuration. In particular, the protocol will support a user initiating a multimedia multiparty session with other users ("calling" other users) over the Internet by allowing a teleconferencing application on one workstation to explicitly rendezvous with teleconferencing applications running on remote workstations. Defining a standard protocol will enable session-level interoperability between different teleconferencing implementations.

The focus of the working group is to design a session negotiation protocol that is tailored to support tightly-controlled conferences. The MBONE currently carries primarily loosely-controlled sessions, i.e., sessions with little to no interaction among members and with no arbitration facility, security, or coordination of quality-of-service options for time-critical media. Users may learn of available sessions using the "sd" utility or other out of band mechanisms (e.g., email). However, there is clearly also a need for tightly-controlled sessions that provide mechanisms for directly contacting other users to initiate a session and for negotiating conference parameters such as membership, media encodings and encryption keys. In addition, these sessions should support renegotiation during a session, for example to add or delete members or change the media encoding. It is possible that the protocol will, in the limiting case, also support loosely-controlled sessions.

The main goal of the working group will be to specify the session control protocol for use within teleconferencing software over the Internet. The working group will focus on the aspects of the session control problem that are well understood, while keeping an eye on evolving research issues. Toward this end, the working group has made an inventory of existing conferencing systems and their session control protocols. The working group will document the requirements of the existing prototypes as a basis for the protocol development. The working group will iteratively refine the protocol based on implementation and operational experience.

Furthermore, the working group will coordinate with other efforts related to multimedia conferencing, such as directory services for cataloguing users and conferences, the RTP and RTCP protocols developed by the Audio/Video Transport Working Group, resource reservation and management at the network level, and schemes for multicast address allocation.

Goals and Milestones

May 1993	Hold an on-line working group meeting to discuss the conference control framework, the relevant terminology, a functional taxonomy and how different conversational styles place requirements on session protocols.
Jun 1993	Submit the Conference Session Control Protocol to the IESG for consideration as an Experimental Protocol.
Aug 1993	Post an Internet-Draft describing the Session Control Requirements.
Nov 1993	Post an Internet-Draft of the Session Control Protocol.
Mar 1994	Submit a revised Internet-Draft based on implementation experience.

CURRENT MEETING REPORT

Reported by Abel Weinrib/Bellcore

Minutes of the Multiparty Multimedia Session Control Working Group (MMUSIC)

An on-line copy of the minutes and the accompanying slides may be found in the directory `venera.isi.edu:confctrl/minutes` as files `ietf.11.93` and `slides.[a-d].11.93.ps`.

The MMUSIC Working Group met for two sessions at the Houston IETF meeting. The first day was dedicated to a short overview of the goals and context for the working group and a presentation of an algorithm and framework for managing shared session state. The second meeting focused on preliminary ideas as to what might comprise shared session state for a couple of different session types, and three short presentations on related work.

Overview and Framework

Abel Weinrib presented an overview of the goals of the MMUSIC Working Group and discussed the framework for the work. This presentation was basically a review of the work of previous working group meetings; refer to the minutes of those meetings available from the confctrl archives for more detail.

In setting the context for the next two talks, a distinction was made between the "agreement algorithm" and the "session control protocol." The agreement algorithm supports generic control of group membership and enforces correctness and other policies on state shared among the members. This agreement layer "understands" membership and policies, but views the rest of the domain-specific session state as opaque. The session control protocol understands the domain-specific session state, using the services of the agreement protocol to manage the state shared among the members. The session protocol may also use other services in addition to the agreement services, such as services that support soft state sharing and recovery.

Issues that were raised during discussion:

- Where should a "session manager" that terminates a session control protocol reside? Various alternatives are on a workstation (as shown in the framework slide) for one or multiple users, or one per domain that could act as a demultiplexing agent by passing on session control messages for users in that domain to the appropriate place. The second alternative may provide hooks for supporting user mobility and may deal well with security firewalls.
- Should floor control be done through the session control protocol or through some other mechanism?

- Should policies be chosen from a predefined set, or should they be defined in all of their generality by each application? This has implications on interoperability and the complexity of the applications.
- In the framework, resource reservation is separate from session management. The session control protocol is used to propagate a shared view of the state, which includes descriptions of the media streams required by a conference.

An Algorithm for Managing Shared Teleconferencing State

Scott Shenker described some preliminary ideas being developed for expressing policies about how session state can be changed and the degree to which members agree on their views of the state. Policy can be expressed along three dimensions: voting policies, consistency policies, and initiator policies. Voting policy defines which members must agree for a state change to take place. Consistency policies describe how the state seen by different members may differ. Initiator policies set which members may initiate changes to the state. The policy framework provides the vocabulary for concretely describing various session styles.

He then presented an algorithm that supports operations on the shared state while enforcing the policies associated with the session. These operations might be adding a member, changing the policies themselves, or modifying some other domain specific state variable such as an encryption key. The basic mechanism is a group agreement algorithm based on a two-phase commit procedure or correctness.

For additional information on this work there is a rough draft document in the confctrl archives in `docs/agree.ps`. Notice of the availability of more complete drafts of the document will be sent to the confctrl mailing list.

Some points raised in the discussion during and following the talk:

- It was observed that some members of a session may be programs running on computers. The fall-back position of always allowing members to leave a corrupted session may be less useful than for human members who can more easily detect the corruption.
- Critical and non-critical membership allows there to be a core group of members that control the conference and a potentially much larger set of members that can more easily enter and leave.
- This talk is about agreement, not negotiation. The distinction is that there is no support for multiple rounds of proposals and counter-proposals. This could be future work, or could be done at the application level building on top of the basic agreement service.

Session Control Above the Agreement Protocol

Eve Schooler's talk was devoted to the interpretation and usage of the agreement protocol for teleconference session control. Discussion attempted to place the agreement protocol in the context of a traditional protocol stack and to hint at implementation concerns. Examples were given for generic and domain-specific session operations, as well as for the array of potentially interesting state attributes (session-wide, membership-related, or media- and policy-specific). To illustrate the range of sessions that can be constructed from different sets of policies, two example paradigms were presented; one for an open hailing-channel session with little coordination among members, and another for a minimal invitation-only session.

The second half of the presentation focused on several open issues: Tradeoffs between different end-system organizations, addressing issues related to the use of unicast and multicast and to the interaction of media agents and session agents, and alternate techniques for user rendezvous that resemble what is currently in place on the MBone for session directories.

For additional information on this work, there is a very rough draft document in the confctrl archives in `docs/usage.txt`. Notice of the availability of more complete drafts of the document will be sent to the confctrl mailing list.

Some points raised in the discussion:

- Issues of media typing and the addressing of media agents are related to problems that need to be solved for WWW as well as XMosiac naming and MIME mailcap media descriptions.
- It would be nice if session control did not assume that the media used by the conference is necessarily carried over an IP network.

Consensus and Control in Wide-Area Communication

Bala Rajagopalan briefly presented his work on agreement and control of group membership in wide area communications. He also handed out a paper that presents his model and algorithm in more detail; contact him via email for a copy of his paper.

The model allows a group to (eventually) come to consensus on its membership in the presence of unreliable message delivery. His algorithm uses wide area multicast and a coordinator for each partition's "view" of the membership state. Operations on groups include join, leave, delete, reform, merge. One underlying assumption of this work that led to some heated discussion during the meeting is that connectivity is transitive, meaning that if A is connected to B and B is connected to C, then A is connected to C; this assumption may break down during certain failure scenarios in the Internet.

This work appears to be relevant to the concerns of the MMUSIC Working Group. More effort is required to understand how and where it might fit into the MMUSIC charter.

RTCP Implications for MMUSIC

Steve Casner discussed the relationship of RTCP, the "real time control protocol" defined by the Audio/Video Transport Working Group (AVT), to the MMUSIC Working Group effort. RTCP is separate from the RTP protocol (which supports transport of time-critical media streams) and may in the future be replaced by a higher level control protocol, such as the MMUSIC session control protocol. In particular, he described the functions that RTCP currently provides, and discussed other functions that would be useful in supporting an application such as multimedia teleconferencing (see the slides). He concluded that it may make sense to use some part of the RTCP in conjunction with a higher level control protocol.

Session Control Work at BBN

Julio Escobar presented a list of relevant work at BBN that is addressing similar issues to the MMUSIC Working Group. He mentioned Chip Elliott's work on the "sticky" protocol (Chip had actually presented this work at an earlier MMUSIC/CONFCTRL BOF), Lou Berger's simulation exercise management tool, and Walter Milliken's work on resource coordination objects. Julio promised to send additional information on this work to the confctrl mailing list (which he has done).

Attendees

Lou Berger	lberger@bbn.com
David Borman	dab@cray.com
Stephen Casner	casner@isi.edu
Ping Chen	ping@ping2.aux.apple.com
George Clapp	clapp@ameris.ameritech.com
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
David Dubois	dad@pacersoft.com
Ed Ellesson	ellesson@vnet.ibm.com
Julio Escobar	jescobar@bbn.com
William Fenner	fenner@cmf.nrl.navy.mil
James Fielding	jamesf@arl.army.mil
Ron Frederick	frederick@parc.xerox.com
Atanu Ghosh	atanu@cs.ucl.ac.uk
Fengmin Gong	gong@concert.net
John Hanratty	jhanratty@agile.com

Ken Hayward	Ken.Hayward@bnr.ca
Van Jacobson	van@ee.lbl.gov
Yasuhiro Katsube	katsube@mail.bellcore.com
Charley Kline	cvk@uiuc.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Ted Kuo	tik@vnet.ibm.com
Paul Lambert	paul_lambert@email.mot.com
Mark Laubach	laubach@hpl.hp.com
Jim Martin	jim@noc.rutgers.edu
Thomas Maslen	maslen@eng.sun.com
Donald Merritt	don@arl.army.mil
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Laura Pate	pate@gateway.mitre.org
J. Mark Pullen	mpullen@cs.gmu.edu
Bala Rajagopalan	braja@qsun.att.com
Steven Richardson	sjr@merit.edu
Eve Schooler	schooler@isi.edu
Henning Schulzrinne	hgs@research.att.com
Scott Shenker	shenker@parc.xerox.com
Michael Speer	michael.speer@sun.com
John Stewart	jstewart@cnri.reston.va.us
Daniel Swinehart	swinehart.parc@xerox.com
Matsuaki Terada	tera@sdl.hitachi.co.jp
Claudio Topolcic	topolcic@cnri.reston.va.us
Abel Weinrib	abel@bellcore.com
Taehwan Weon	weon@cosmos.kaist.ac.kr
John Wroclawski	jtw@lcs.mit.edu
Shinichi Yoshida	yoshida@sumitomo.com
Lixia Zhang	lixia@parc.xerox.com
Weiping Zhao	zhao@nacsis.ac.jp

Multiparty Multimedia Session Control Working Group (mmusic)

Eve Schooler <schooler@isi.edu>
Abel Weinrib <abel@bellcore.com>

1,2 November 1993

Agenda

Overview of mmusic
- goals
- framework

Managing shared session state
- policy specification
- agreement algorithm

Discussion

Session control above the agreement protocol
- usage of agreement framework

Discussion

Other topics
- an agreement protocol—
- relation to RTCP
- session control work at BBN
- action items
  - accomplishing old ones
  - setting new ones

AW 11/18/93 2

MMusic Overview

Motivation for work

Working group charter
- focus

Milestones (from last time)

8/93 Post an Internet Draft describing the Session Control requirements
11/93 Post an Internet Draft of the Session Control Protocol
3/94 Submit a revised Internet Draft based on implementation experience
6/94 Submit the Session Control protocolÔ77773340 to the IESG for consider as an Experimental Protocol

AW 11/18/93 3

Conference Styles

point-to-point call
impromptu hallway meeting
architecture design review
classroom
panel discussion
seminar
lecture
pay-per-view
TV broadcast

more coordination
more rigid
more overhead
smaller sessions
more static

less coordination
fewer policies
less overhead
larger sessions
more dynamic

AW 11/18/93 4

Framework

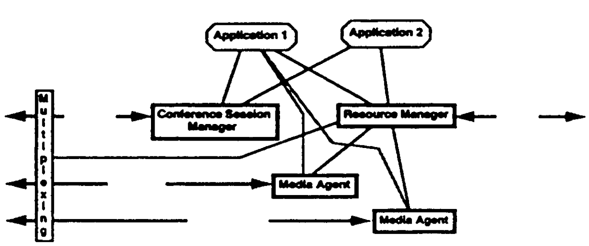

AW 11/18/93 5

State of a Session

1. **Session identifier**
2. **Member list (M)**
   - each member appears only once
   - at least one member (for tight controlled session)
3. **Session policies (P)**
   - voting
   - correctness
   - initiator of operations
4. **Other domain-specific state variables not "understood" by session manager**
   - media encodings and communication topology
   - security: encryption key distribution
   - floor control
   - common name
   - etc.

AW 11/18/93 6

Session Control Protocol Requirements

Distributed session management
- establish, terminate, merge
  - once established, a session has a unique name (for example, to advertise the session in a directory service)
- schedule—negotiate the session in advance

Dynamic state management
- add member, delete member
- modify policies
- modify other state variables

AW 11/18/93 7

Protocol Assumptions

One Session Manager per workstation port
- usually a single one at a well-known port
- a Session Manager may serve multiple users
- a Session Manager is not required to maintain state across failures

There is an of a session
- anyone may propose initiation of a session
- does the initiator have to be a member of the resulting session?

The initiator proposes the initial session state
- the proposal is forwarded to all potential members
- each proposed member knows up-front what it is agreeing to
- the voting policy defines who must agree for session establishment to succeed

AW 11/18/93 8

More Protocol Assumptions

Once established, a member may propose changes to the session state
- only members may make such proposals
  - a non-member that wants to join a session asks a member to propose the required change
- voting policy dictates whether proposal is approved

While a proposal is outstanding, the proposer coordinates the outcome
- once the outcome is determined, the proposer reverts to having no special role

A member may always leave the session without negotiation

AW 11/18/93 9

Other Session Protocol Messages

Query
- sent by a session manager to request the session state from another session manager
- the receiving session manager responds with a status message

Status
- sent by a session manager to share the state of the session
- no reply is required

Session Protocol	
Agreement	State Sharing

AW 11/18/93 10

Underlying Transport Mechanism

Internet realities
- packet loss
- delay variations

Communication building blocks
- UDP vs. TCP
- unreliable vs. reliable "datagrams"
- point-to-point vs. multipoint
  - CCP reliable multipoint
- causal/atomic multicast
  - ISIS group services

Adaptability to delay variations
- adapt peer-to-peer timeouts to dynamic network conditions

AW 11/18/93 11

Outstanding Issues (from last time)

Mobility at session establishment and during the session

Distributed system issues
- keep alive messages using Announce
- state synchronization (instead of strict correctness)
- join by contacting a or
- replace terminate session by using delete all members
- error handling
- serializability conflicts and fault tolerance
- (new) correct handling of timeouts for fault recovery

Transport infrastructure
- UDP, TCP, Sun RPC, VMTP, TP4, ISIS, etc.?

Schedule and preschedule

Floor control

Richness of policies

Other uses?

AW 11/18/93 12

Action Items (from last time)

Documents

- requirements
  - framework
  - terminology
  - conversation styles
  - protocol requirements
- specification

AW 11/18/93 13

An Algorithm for Managing Shared Teleconferencing State

Extremely Preliminary!

Scott Shenker
Abel Weinrib
Eve Schooler

IETF November 1993

Teleconferences

Shared State:

- Membership
- Roles (speaker, next-in-line, chairperson)
- Media encodings, encryption keys, etc.
- Payment model

Human endeavor:

- Deciding on shared state is not just a technical problem
- Decision criterion: policy issue (social/political)
- Means to implement: technical issue

Many different policy options:

- Anarchy, dictatorship, majority rule, ...

Telephone Examples

Two-party phone call:

- No change to membership possible
- Tight control of membership

Party Line:

- Anyone can join
- Loose control of membership

Typical teleconferencing mechanisms:

- Policy embedded in mechanism
- No policies in between the two extremes

Social Example

Country Club Membership:

- Need three sponsors to join
- Can be blackballed by a single member
- Membership policy can be changed by changing the charter of the country club
- General mechanism: voting at meetings

Typical social institutions:

- Intermediate membership policies
- Policy not embedded in mechanism

This is what we want to emulate!

Goals of this Work

1. Express policy
   - Define general family of policies
   - Explore the *tight/loose* spectrum

2. Implement policy
   - Define a single mechanism to implement this family of policies
   - Variant of two-phase commit
     - A member initiates change in shared state
     - Propose > Collect Votes > Announce Results
     - Members exchange messages for voting, setting locks, committing, etc.

Context and Notation

Assume (for now) that

- Messaging is reliable but asynchronous
- All messages from a particular sender to a particular receiver arrive in order

Notation

- S: shared state
  - S_i state variables
  - $S_1 = M$ membership
  - $S_2 = P$ session policies
- O_i: operations on state variable S_i
  - O_i particular operation in O_i
  - General change operation $op = O_i \ldots \circ O_j$
  - $STATE(op)$: all possible state variables changed by op

Expressing Policy

Three dimensions to policy

- Voting policies
- Consistency policies
- Initiator policies

Voting Policies

For every O_i there is some *voting rule* $V(O_i)$

- Takes set of votes {YES, NO, ABSTAIN}
- Returns 0 (Fail) or 1 (Pass)
- Depends on operation, not just on i
- Independent of initiator

For a general change $op = O_i \ldots \circ O_j$

- $V(op) = V(O_i) \ldots \circ V(O_j)$

Null voting rules:

- No votes (other than the initiator's) need be cast

Consistency Policies

Views of shared state can be inconsistent

- Permanent: operations executed in different order
  - Long-lived inconsistency can render teleconference inoperable
- Temporary: not all change messages have been delivered (inevitable)
  - Votes can be cast with different views of shared state

Two Consistency Conditions

Eventual Consistency in E

- Eventual set $E \subseteq S$
- When session is quiescent, all members have identical views of the state variables in E.

Voting Consistency

- *Critical* state C and *Consistent* state I
- All votes cast on changes in I are made with consistent views of C ($\Rightarrow C \subseteq E$)

Choice of I, C, E is statement of consistency policy!

Restrictions

- $I \subseteq E$
- $I = C$ *(will discuss later)*
- Voting consistency independent of operation

Sets

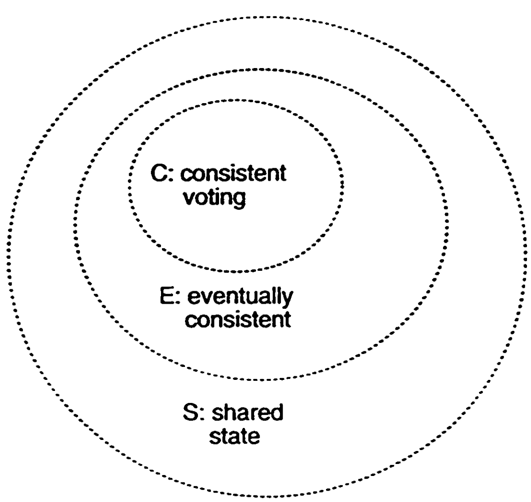

Initiation Policies

Can have general set of policies about which members can initiate changes to various state variables.

These policies define set U_m for each member m

- Changes to variables in U_m can only be initiated by m
- $U_m \cap U_n = \phi$ when $m \neq m$
- $U = \cup_m U_m$
- *Paper wrong*

Unique initiation does not imply unilateral control

Sets

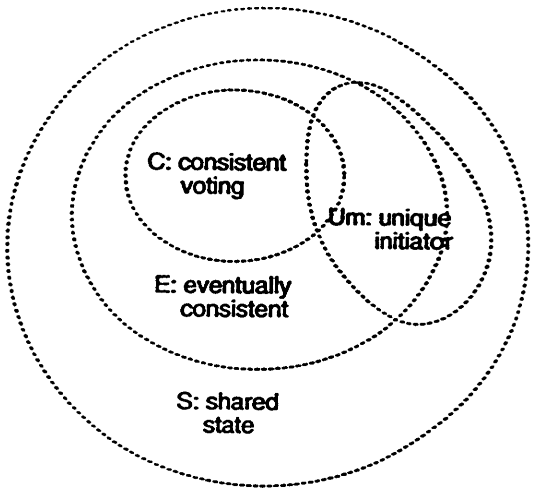

Initiation policies are orthogonal to consistency policies!

Special Role of Membership

If membership is not in C, then we suggest that C is empty

- All members must have a consistent view of C when voting on changes in C
- If membership is not in C, what does *"all members"* mean?

For now, assume that membership is in C

Example

Shared State S

- Membership
- Encryption key
- Common names of members
- Next-to-speak

Policy:

- Membership: in C
  - Adding member: unanimous vote
  - Deleting member: if and only if member wants to leave
- Encryption Key: in $E - C$
  - Null voting
- Common name of m: in $E - C$ and U_m
- Next-to-Speak: not in E
  - Majority rule

Implementing Policy

Variant of two-phase commit

- Dynamic group membership
- Looser correctness conditions
- Spectrum of voting policies

Basic mechanism

- Change initiated by single member
- Propose > Collect Votes > Announce results
- Change assigned unique id
- Members are either busy (locked) or non-busy (unlocked)

Basic Message Functions:

- Poll, Lock, Respond, Announce

Message Types

`Poll(Id: id, Operation: op)`

- Requests vote on proposed operation
- Response is {YES, NO, ABSTAIN}

`Lock(Id: id)`

- Requests exclusive lock
- Response is {OK, BUSY}

`Poll_Lock(Id: id, Operation: op)`

- Requests exclusive lock and vote on proposed operation
- Response is {YES, NO, ABSTAIN, BUSY}

`Response(Id: id, Response: response)`

- Response to `Poll`, `Lock`, or `Poll_Lock` message
- No response

`Announce(Id: id, Operation: op)`

- Operation op is to be applied to state
- No response

`Release(Id: id, Operation: op)`

- Associated lock is released
- No response

`Commit(Id: id, Operation: op)`

- Operation op is to be applied to state
- Associated lock is released
- No response

Message Exchanges

Four Basic Exchange Patterns:

- `Poll_Lock` > `Response` > `Commit` or `Release`
- `Lock` > `Response` > `Commit` or `Release`
- `Poll` > `Response` > `Announce`
- `Announce`

Cross Product of Polling and Locking

- Poll: if $V(op)$ is not null
- Lock: if $STATE(op)$ intersects shaded region

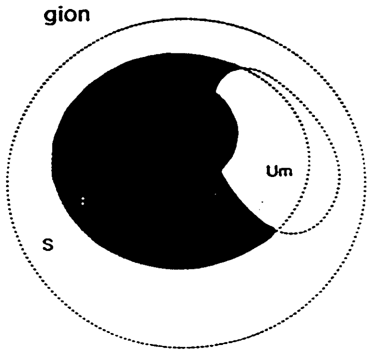

Message Exchanges

Consider some change operation op initiated by member m

Message exchange depends on:

- Does it require a vote?
- $STATE(op) \cap C \neq \phi$?
- $STATE(op) \cap E - C \neq \phi$?
- $STATE(op) \subseteq U_m$?

VOTE and $STATE(op) \cap C \neq \phi$:

- `Poll_Lock` > `Response` > `Commit` or `Release`
- Release is sent if vote fails or if any BUSY response is received
- Locking needed to ensure no two operations are going on simultaneously

Same sequence applies for:

- VOTE, $STATE(op) \cap C = \phi$, $STATE(op) \cap E - C \neq \phi$, $STATE(op) \notin U_m$
- Locking needed to ensure operations are done in same order

NO_VOTE and $STATE(op) \cap C \neq \phi$:

- `Lock` > `Response` > `Commit` or `Release`
- Release is sent if any BUSY response is received

Same sequence applies for:

- NO_VOTE, $STATE(op) \cap C = \phi$, $STATE(op) \cap E - C \neq \phi$, $STATE(op) \notin U_m$

VOTE, $STATE(op) \cap C = \phi$, $STATE(op) \cap E - C \neq \phi$, $STATE(op) \in U_m$:

- `Poll` > `Response` > `Announce`
- No lock needed because unique initiator ensures no variables will be changed by any other operation

Same sequence applies for:

- VOTE, $STATE(op) \subseteq S - E$

NO_VOTE, $STATE(op) \cap C = \phi$, $STATE(op) \cap E - C \neq \phi$, $STATE(op) \in U_m$:

- `Announce`
- Unique initiator means eventual consistency is assured without locks.

Same sequence applies for:

- NO_VOTE, $STATE(op) \subseteq S - E$
- No voting and no consistency conditions allow a unilateral announcing of the new state variables.

Consistency Conditions

Eventual Consistency Whenever there are no outstanding messages, all participants have the same view of E.

Consistent Voting All votes cast involving changes to C are done with identical views of C.

Loose Control of Membership

What happens when membership is not in C?

- Previously argued that C must be empty
- If E empty, everything vacuously OK

For nonempty E

- Our mechanism does not work
- $E \subseteq U$: use periodic **Announces**
  - changes still use `Poll` messages
- Xerox Clearinghouse algorithms
  - after-the-fact timestamp ordering
  - need to keep history on membership
- Both of the above need some method to eventually reach entire membership

Unreliable Transport

Design Choice

- Cope with inconsistencies rather than achieve correctness
- Need hooks for resynchronizing state

Unilateral withdrawal: when all else fails...

- Send **Announce** declaring departure
- Modified consistent voting condition
- Voting is not consistent but *plausible*

Open Issues

- Finer grain concurrency control
- Critical and noncritical membership
- Session initiation
- Voting by incoming members
- Changes to policy
- Responses to protocol errors
- Better and more complete formalism
- Is this useful in building real systems?

IETF Multiparty Multimedia Session Control WG (mmusic)

Abel Weinrib <abel@bellcore.com>
Eve Schooler <schooler@isi.edu>

1, 2 November, 1993

Overview

Part I:
- **Interpretation of agreement protocol:**
  - Short and long term goals
- **Teleconferencing state:**
  - Operations and attributes
- **Example policy mappings:**
  - Known paradigms: sd and mmcc

Part II:
- **Discussion:**
  - Influence of addressing on protocol behavior
  - Directory service or active agent?
    - Loose vs tight control, public vs private models
  - Idea for a user directory service
  - Proposed protocol for explicit sessions

2

Agreement Protocol

- **Framework to describe a diversity of sessions:**
  - Operations that may be performed on sessions
  - Who is qualified to perform those ops
  - Who must approve the proposed ops: *voting rules*
  - The state attributes effected by the op
  - Consistency conditions
- **Aims to support:**
  - Re-usability of messaging components
  - Fluidity from one mode to another
  - Could build negotiation service on top of AP

3

Session Control Stack

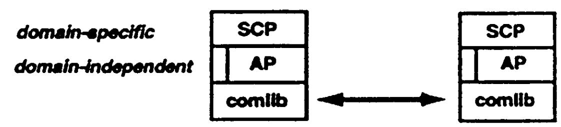

SCP	= session control protocol
AP	= agreement protocol
comlib	= communication library, *e.g, N-way, M-way, 1-to-1*
other	= generic ops not needing agreement

4

Implementing Session Agreements

- **The Session:**
  - Collective of associated peers for control of shared state
- **Session manager:** `manager_id = <host,port>`
  - Endpoint for receipt and initiation of proposals
  - Listens at known locale, acts on behalf of user
- **Proposer:**
  - member who makes a proposal, responsible for outcome
- **Initiator:** `session_id = <user,manager_id,ts>`
  - Uniquely names session
  - Proposes the initial session configuration

5

Operations: $O = \{O_i\}$

- **Distributed session management:**
  - Establish, terminate [, merge, schedule]
- **Dynamic membership management:**
  - Add member, delete member, [modify member]
- **Session policy management:**
  - Set, modify, remove
- **Generic:** Status, query
- **Domain-specific:**
  - Specification of media configurations: capabilities, QoS
  - Distribute privacy info: session encryption key
  - Policies: for floor control, sender vs receiver, etc.

6

Teleconference State: $S = \{S_i\}$

- **Attribute hierarchy: lattice of info**
  - Session-wide characteristics
  - Membership-related attributes
  - Media configuration parameters
  - Policies describing their interaction
- **Specific or general attributes:**
  - session- , membership- , or media-wide
- **Differentiating factors:**
  - Critical for operation or simply informational
  - Maintained *globally* or locally
  - Dynamic or static variables

Policy Attributes: P

- **Attributes that *define* style of session**
- **A policy is a binding between:**
  - $\langle O_i, \mathtt{propose}(O_i), V(O_i), \mathtt{STATE}(O_i)\rangle$
- **Initiator establishes meaningful set**
- **Propose(O_i) and V(O_i):**
  - May be entirely different members
  - Either group may be empty
- **Policy descriptor:**
  - A profile of common policy sets
  - Full matrix of policies not needed
- **Consistency policies:**
  - $S_i \in \Sigma,\ S_i \in C$

Session Attributes

Might Include:

- **Session identifier**
- **Membership list:** `{member_id}`
- **Session address(es):**
  - Inclusion of session key
- **Session information:**
  - Session alias
  - Session start/end time: impact resource reservation?
- **Extensibility...**

Sd Cache Entries

```
n=2148114532 2148114532 740786921
s=** Please don't start a radio session **
i=There is not enough bandwidth to allow more
   than one global radio session at a time.
   See Radio Free Vat session to get a time
   slot.  For further explanation, email
   casner@isi.edu
o=casner@oak.isi.edu
c=224.2.174 191 0 0
m=audio 52858 17126
a=fmt:pcm
```

Membership Attributes

Might Include:

- **Member address:**
  - `member_id = <user@manager_id>`
- **Member information:**
  - Member status (e.g., `On Hold`)
  - Member alias
  - Electronic mail address
  - Geographic location of site
  - Other RTCP fields

Media Attributes

Might Include:

- **Domain-specific payload:**
  - Specific to teleconferencing
- **Carried via session manager exchanges:**
  - Processed by media agents
- **Media list:** `{media_id?}`
- **Media type:** audio, video, groupware
  - Media agent: nevot, ivs, wb
- **Media address**
- **Application specifics:**
  - Encoding, priority, demux schemes, etc.

Policy Matrix: Open Hailing Channel

Oi (Operation)	Propose(Oi) (Who may propose)	V(Oi) (Approval needed)
Session Ops		
Create	ALL	Initiator
Terminate	Initiator	Initiator
Member Ops		
Add	ALL	Proposer
Delete	ALL	Proposer
Modify	M	NULL (U_-)
Policy Ops		
Set	Initiator	Initiator
Modify	NONE	–
Media Ops		
Add	Initiator	Initiator
Modify	NONE	–

Hailing Channel: no E, C, V(O_i)

- **Initiator role simple, yet pivotal: *owner***
  - Session establishment and termination
  - Sets overall session description: static
  - Only one eligible to change it?
- **Dynamic membership:**
  - Unrestricted joins, no notion of invitation
  - Any member may leave at will
  - Member state propagation decentralized and periodic
- **Media is symmetric:** all members snd and recv
- **Two-stage termination:**
  - Termination a directory service op or a session op?
- **All ops – performed globally**

Policy Matrix: Explicit Invite Session

OI (Operation)	Propose(OI) (Who may propose)	V(OI) (Approval needed)
Session Ops		
Create	ALL	M=1
Terminate	{ Initiator { ALL	Initiator ALL
Member Ops		
Add	ALL	M=1
Delete	ALL	Member effected
Modify	M	NULL (Um)
Policy Ops		
Set	Initiator	Initiator
Modify	ALL	ALL
Media Ops		
Add	Initiator	Initiator
Modify	ALL	{ Member effected { ALL

Explicit Invite:

- **Initiator only responsible for session init:**
  - Sets overall session description
  - Anyone eligible to change it, tho some S_i fixed
- **Session termination:**
  - Anyone can perform
- **Dynamic membership: MϵC**
  - Any member may leave at will
  - Sponsorship to invite/join: 1 member
  - scaling: critical and non-critical members
- **Media is symmetric:** all members snd and recv
  - Mid-session changes: QoS, floor ctrl, suspend
- **Most ops globally performed [all voted on]:**
  - Consistency is a function of session formality: privacy
  - Piggyback state on proposals-->Periodic refreshes

Discussion: End-System Architecture

- **Per domain session manager**
  - To mask addressing and firewall complexities
- **Per workstation session manager**
  - Demux on user id
  - Daemon to listen at given port
- **Per workstation, per user**
  - Individualized; fewer levels of indirection
  - Hard to support one port model

Discussion: Addressing Issues

- **Teleconferences: multidimensional sessions**
  - Repeated issue of demultiplexing
  - Multiple sessions, members, media, policies
- **Session address(es):**
  - Unicast or multicast or both
  - Same or different from `manager_id`
  - Inclusion of session key(s)
- **Media address(es):**
  - mcast vs unicast, different from ctrl, port, ttl, key

Discussion: Directory or Agent?

- Differences between session directory and session manager?
  - Should there be?
- Session existence: 0 or 1 member
  - Distinction between session creation/membership and session termination/member deletion
- Treatment of members and non-members
- Tradeoff of join/invite:
  - Each is difficult in other model
  - Public vs private session registry
- Staying in the loop:
  - Mid-session configurability
    - Inter-media (e.g., QoS, floor control)

Discussion: User Directory Service

- Multicast service for user registration
  - In the style of sd
- User registers when directory tool launched
  - Avoids trying to call an unavailable user
  - Tracks multiply-homed users
    - Re-register user based on machine activity
- Scope implications
  - Periodic announcements at different ttl's
  - Announcement interval decreases as ttl increases
  - Derives minimum ttl to reach remote users

Session Directory Wishlist

- Semi-public sessions:
  - Limited entry info in registry
  - Limited access/distribution
- Access to multicast allocation
  - Sharing among all multicast tools
  - Private, semi-public or public
- Decide convention for addr and port usage
  - Single vs multiple addresses
  - Fixed or dynamic ports
- Loose, but private sessions:
  - Distribution of handle to members only
  - Allows late-comers to join

RTCP Implications for MMUSIC

RTP spec says that RTCP is separable and may be replaced by a higher-layer control protocol. What would be required for a protocol produced by MMUSIC to replace RTCP?

- Functions of RTCP
- Functions attributed in spec to higher-layer protocol
- More sophisticated functions

Functions of RTCP

- Define mapping of "format" values to encodings (FMT)
- Report reception quality (packet loss & delay) (QOS)
- Provide mapping between locally unique source identifiers and participant information (SDES)
  - QOS reports return the globally unique ID
  - Activity indication depends upon text names
- Indicate when a participant leaves (BYE)

Functions Attributed to Control Protocol

- Distribute multicast addresses, ports, channel IDs, TTLs
- Negotiate media encodings and parameters
- Arrange a table specifying keys and the selected encryption and digest algorithms, in common for all sources within one channel
- Profile or control protocol may specify:
  - grouping of channels to be covered by same SDES
  - requirement for timestamps to be synchronized (NTP)

More Sophisticated Functions

- Arrange for bridges and translators to be set up and configured where needed for a particular session
  - Communicate local identifier mapping
  - Set up keys as appropriate

► *It may make sense to still use some parts of RTCP in conjunction with a higher-level control protocol — for example, use SDES to map local identifiers, but use a different kind of unique handle belonging to control protocol*

2.9.3 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.

Internet-Drafts

"TCP Extensions for High Performance: An Update", 06/23/1993, R. Braden
<draft-ietf-tcplw-extensions-00.txt>

Request For Comments

RFC 1323	"TCP Extensions for High Performance"

2.10 User Services Area

Director:

- Joyce Reynolds: `jkrey@isi.edu`

Area Summary reported by Joyce Reyolds/Information Sciences Institute

Ten working groups in the User Services Area of the IETF met in Houston, Texas.

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, IDS 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 IDS Working Group reviewed, discussed and/or progressed the following documents:

- "X.500 Implementation Survey" - held up, pending more vendor responses.
- "X.500 Pilot Project Catalog" - held up, pending a query to Paradise.

Coming "soon" (within the next month):

- "WHOIS++ Implementation Catalog"
- "Model for Information Privacy of Directories"
- "Legal Issues for Directories in Europe"
- "Data Management Issues"
- "Overview of Directory Services"

Also discussed was the inclusion of the CSO nameserver protocol in the IDS effort (consensus was to include it), and the more general issue of how to make all these directory services work together (lots of ideas, but no solid conclusions, yet).

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."

Clifford Lynch discussed his paper on using the Z39.50. Margaret St. Pierre discussed the Internet-Draft, "WAIS over Z39.50 1988." This document is being considered as an Informational RFC profiling the use of Z39.50 version 1988 by the traditional WAIS protocol. The idea of quality assurance was discussed. Quality assurance in this context addresses such issues as invalid pointers to data objects, interoperability among the current information systems and the ability to contact information maintenance personnel. A mailing list is now established called `quality@sunsite.unc.edu`. The idea of a data types registry was discussed and Greg Vaudreuil agreed to write a document on using the MIME content type registry in Amsterdam, but that document was never posted to the IIIR list. Applications co-Area Director John Klensin indicated that the document had been denied by the IESG based on the fact that it undermined an existing RFC.

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-12th Grade) schools, public and private, to the Internet, and school networking in general.

After a period for introductions of those attending and a statement of their interest in ISN, the group launched into a discussion of whether it should continue to exist. With a tentative decision to continue activities, a revision of the body of the charter began. This was accomplished, and minor refinements and word-smithing will take place on the list. April Marine reported on her action item to investigate the InterNIC's ability and willingness to maintain a directory of people in primary and secondary school education who are involved in networking. The group then defined a set of milestones which will take the group through March, 1995 and was therefore permitted to go to lunch. April's report was considered in crafting the milestones.

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.

- Documents: The Internet-Draft on international NIC structures and RFC 1302/FYI 12 will be completed/revised by the Seattle meeting.

- NIC-Profiles: This information will be revived, working in cooperation with the InterNIC.

- NISI Future: There was good discussion about the scope and role of NISI including possible future projects. Various plans were discussed, including the possibility of ending the group and having a new group emerge if needed, with a different name and focus. Further discussion will take place on the list followed by action in Seattle.

Network Training Materials Working Group (TRAINMAT)

The Network Training Materials Working Group is chartered to enable the research community to make better use of 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 ongoing work towards a training materials catalog was reviewed. The template for materials was reviewed, with recent changes, suggested by Pete Percival, incorporated. The working group agreed with the newest version. Sample entries developed by a team working with Margaret Isaacs was reviewed, and the working group agreed that the focus should be on training rather than documentation or resource guides, the latter which could be included in a bibliography rather than in the main part of the document. Sample subject headings were reviewed, with the agreement that these might change if documentation and guides were removed from the body of the catalog. Additional volunteers were recruited to complete putting materials from the University of Newcastle catalog into template form, with the goal of adding new materials once this task is completed—about half the original Newcastle entries are already in template form. Updates will be sent to the e-mail list. Discussion also focused on materials which might need to be developed in the training area, particularly the usefulness of videos. The final segment of the session included discussion of efforts by other groups, and a roundtable on what attendees are doing in the training area.

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.

Jim Fullton gave the current status of CNIDR. Jill Foster gave an overview of RARE activities in this area. The working group split into small groups to discuss each section of the NIR status report as a final review. Section 5, mailing lists and gopher are among the major areas changed. It will be sent to the list for final comments. April Marine discussed the "checklist" history. She has put together a simple NIR tool checklist.

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.

The URI Working Group held three sessions in Houston. The first two were dedicated to closing work on the Uniform Resource Locators, which seems to have occurred. The final session worked on Uniform Resource Names. Document drafts have been commissioned to reflect those discussions.

User Documents Revisions Working Group (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 USERDOC2 Working Group had a small but enthusiastic set of attendees due to a conflicting applications group meeting that many of the usual suspects attended. The group set a new record by having everyone in the room volunteer to assist in writing and reviewing documents in progress! A "Not Quite an Internet Draft" bibliography was distributed and discussed which will update RFC 1175. The new document will focus on books, journals and other bibliographies rather than all possible documentation. A section will be developed to cover RFCs and FYIs. The final document is scheduled to be completed before the next IETF. A second document covering Internet connectivity is almost ready for review, and will be distributed to the mailing list. This is a joint project with the ISN Working Group.

User Services Working Group (USWG)

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.

Gary Malkin briefly discussed the "DAWG" (Distribution and Announcement Working Group) idea that has been sitting on USWG's back burner for a while. A BOF will be held at the next IETF to see if there is further interest in this topic. Ann Cooper led a talk and discussion on the US Domain. Jill Foster and Joyce Reynolds reported on the RARE ISUS meetings and the EARN Network Services Conference held in Warsaw, Poland, in which they participated. Jill announced the INET94/JENC5 Call For Papers—User Information Track to the USWG. Jill was asked to run this track, and asked Joyce if she would be co-track leader. There was continued discussion from the Amsterdam IETF on Bill Manning's thoughts about how to "empower" users to utilize and document tools.

WHOIS and Network Information Lookup Service Working Group (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 discussion on the status of the WHOIS Architecture. Chris Weider and Simon Spero led a discussion on the status of the distributed WHOIS++ model and centroids. Chris described changes to the draft WHOIS++ document. Simon Spero described the mechanism for searching a centroid tree from the bottom, up. Jim Fullton described the status of WHOIS ++ Clients. Jim mentioned the use of WHOIS++ in support of networked information retrieval and the type of client development that is occurring as part of other application development. The session concluded with a discussion on the recommendations and modifications to the WHOIS Protocol and a discussion of WHOIS++ Implementations by Joan Gargano.

2.10.1 Integrated Directory Services (IDS)

Charter

Chair(s)

Chris Weider: clw@bunyip.com
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 Working 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 X.500, 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".

Done 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

"X.500 Pilot Projects", 06/15/1993, A. Marine <draft-ietf-ids-pilots-00.txt>

"A Revised Catalog of Available X.500 Implementations", 10/08/1993, A. Getchell, S. Sataluri <draft-ietf-ids-catalog-00.txt>

Request For Comments

RFC 1491 "A Survey of Advanced Usages of X.500"

CURRENT MEETING REPORT

Reported by Tim Howes/University of Michigan

Minutes of the Integrated Directory Services Working Group (IDS)

Review of Previous Minutes

The minutes from the Amsterdam meeting were accepted without change.

Liaison Reports

- North American Directory Forum (NADF)

  Tim Howes served as unofficial liaison to the NADF and reported that the NADF continues its piloting activities and has opened up its membership to user organizations (as well as service provider organizations). The next NADF meeting is in Reston, VA the week of November 8.

- Whois and Network Information Lookup Service Working Group (WNILS)

  Chris Weider reported that the centroids and overall architecture documents have been out for some time and that the WHOIS++ document would be out shortly.

Status Reports

- WHOIS++

  Chris reported there were now 4 implementations of WHOIS++ servers, one of which supports centroids. Client development is proceeding or promised for several client platforms, but none are available yet.

- Lightweight Directory Access Protocol (LDAP)

  Tim reported that version 3.0 of the University of Michigan LDAP distribution was released just before the Houston IETF. It includes a server and several client implementations.

Progress of Documents

Editor's Note: A list of documents and their progress is available via FTP or mail server from the remote directories as /ietf/ids/ids-minutes-93nov.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

New Directory Services

The Computing Services Office (CSO) Nameserver protocol was discussed. In lieu of Joel Cooper from Notre Dame, Chris led the discussion. CSO is currently a local directory service, but the consensus of the group was that there is work to be done in the "common schema and access" area, which makes it appropriate for inclusion in IDS.

Chris will talk to Joel Cooper about getting the CSO documents (or pointers to them) submitted to the list.

AOB

In the CSO discussion it was brought up that a relevant meeting had occurred earlier in the week. Tim gave a brief summary of the purpose and results of the meeting. The National Science Foundation, finding itself inundated with many requests to fund directory service efforts, convened a group of experts in the Internet community for the purpose of advising NSF on how best to allocate its limited funds to these projects.

The consensus of the group was that no single white pages directory service was about to take over the Internet any time soon. Sites have extensive flexibility to run whatever they please, for a variety of reasons. Therefore, the informal recommendation of the group was for NSF to diversify itself and fund a variety of white pages directory service efforts. Furthermore, the group recommended that another area worthy of research is a system to tie together the various disparate white pages directory services. Tim introduced the basics of one such plan which had been presented at the meeting by Marshall Rose. Much interesting discussion ensued.

Next Meeting

The next meeting of the IDS Working Group will be at the March IETF in Seattle, WA.

Attendees

Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Sepideh Boroumand	sepi@aol.com
Luc Boulianne	lucb@cs.mcgill.ca
Glen Cairns	cairns@mprgate.mpr.ca
Peter Deutsch	peterd@bunyip.com
Urs Eppenberger	eppenberger@switch.ch
Qin Fang	qin_fang@unc.edu
Jill Foster	Jill.Foster@newcastle.ac.uk

Jisoo Geiter	geiter@mitre.org
Arlene Getchell	getchell@es.net
Mei-Jean Goh	goh@mpr.ca
Chris Gorsuch	chrisg@lobby.ti.com
Judith Grass	grass@cnri.reston.va.us
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Marco Hernandez	marco@cren.net
Tim Howes	tim@umich.edu
Richard Huber	rvh@ds.internic.net
Barbara Jennings	bjjenni@sandia.gov
Michael Kornegay	mlk@bir.com
John Kunze	jak@violet.berkeley.edu
Ben Levy	seven@ftp.com
Wayne McDilda	wayne@dir.texas.gov
Michael McLay	mclay@eeel.nist.gov
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Rakesh Patel	rapatel@pilot.njin.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Richard Rodgers	rodgers@nlm.nih.gov
Kenneth Rossen	kenr@shl.com
Srinivas Sataluri	sri@internic.net
Richard Schmalgemeier	rgs@merit.edu
Rickard Schoultz	schoultz@sunet.se
Sam Sjogren	sjogren@tgv.com
Mark Smith	mcs@umich.edu
Milan Sova	sova@feld.cvut.cz
Margaret St. Pierre	saint@wais.com
David Staudt	dstaudt@nsf.gov
Chris Weider	clw@bunyip.com
Brien Wheeler	blw@mitre.org
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov
Russ Wright	wright@lbl.gov

2.10.2 Integration of Internet Information Resources (IIIR)

Charter

Chair(s)

Chris Weider: `clw@bunyip.com`
Kevin Gamiel: `kgamiel@cnidr.org`

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", 10/26/1993, C. Weider <draft-ietf-iiir-transponders-01.txt>

"A Vision of an Integrated Internet Information Service", 10/26/1993, C. Weider, P. Deutsch <draft-ietf-iiir-vision-01.txt>

"Hypertext Markup Language (HTML): A Representation of Textual Information and MetaInformation for Retrieval and Interchange", 07/23/1993, T. Berners-Lee, D. Connolly <draft-ietf-iiir-html-01.txt, .ps>

"WAIS over Z39.50-1988", 10/26/1993, M. St. Pierre, J. Fullton, K. Gamiel <draft-ietf-iiir-wais-00.txt>

"Hypertext Transfer Protocol (HTTP) A Stateless Search, Retrieve and Manipulation Protocol", 11/16/1993, T. Berners-Lee <draft-ietf-iiir-http-00.txt, .ps>

CURRENT MEETING REPORT

Reported by Kevin Gamiel/MCNC - CNIDR

Minutes of the Integration of Internet Information Resources Working Group (IIIR)

Z39.50 Over TCP/IP

Clifford Lynch discussed his paper on using the Z39.50 Search and Retrieve protocol directly over TCP as opposed to OSI. The paper was introduced in Amsterdam by Jim Fullton. Members of the group requested that the paper be resubmitted to the list for further review. The paper will be submitted as a candidate for an Informational RFC.

WAIS and Z39.50

Margaret St. Pierre discussed the Internet-Draft "WAIS over Z39.50 1988." The document was submitted to the IIIR list on October 27, 1993 for review. It is being considered as an Informational RFC profiling the use of Z39.50 version 1988 by the traditional WAIS protocol. The document is meant to be an overview of current implementation rather than a rigorous technical specification. It addresses such issues as the history of the WAIS protocol as well as assumptions made by the WAIS community. The group offered no additions or corrections to the document.

Quality Assurance

With the growing numbers of commercial products based on protocols such as Gopher, the need for quality assurance has increased. Quality assurance in this context addresses such issues as invalid pointers to data objects, interoperability among the current information systems and the ability to contact information maintenance personnel. A mailing list has been established named `quality@sunsite.unc.edu`. To subscribe, send a message to `listserv@sunsite.unc.edu` with "subscribe quality <real name>" in the text. If activity on this list is substantial, the group will consider moving activity into IIIR in Seattle. An informal BOF on quality assurance was held by Mitra following this IIIR meeting. The notes from that BOF will be posted to the IIIR list.

Format Types Registry

The idea of a data types registry was discussed in Amsterdam. Greg Vaudreuil agreed to write a document on using the MIME content type registry in Amsterdam, but that document was never posted to the IIIR list. Area Director John Klensin indicated that the document had been denied by the IESG based on the fact that it undermined an existing RFC.

During the second IIIR meeting, Jon Postel announced that a solution would be worked out to allow assignment of a wide variety of new types to the MIME specification. This might include such things as registering a 'Application' type and allowing registry of additional subtypes under this. The discussion will be taken to the list.

The Vision Document

The Vision document was approved for progression as an Informational RFC after appropriate text was added to make it clear that the IIIR Working Group believed this to be a valid vision of the future but that this was not the consensus of all members of the group. Chris Weider stated that the appropriate text would be added and the document would be moved on.

Resource Transponders

The Resource Transponder document was approved for moving to Informational RFC following the same pattern as the Vision document. Again Chris stated that the appropriate text would be added and the document would be moved on.

HTML/HTTP

This topic was deferred to the list due to the absence of Tim Berners-Lee.

Horton Hears a WHOIS

This presentation had been given in the WHOIS and Network Information Lookup Service Working Group (WNILS) meeting to substantially the same audience; and Simon Spero graciously waived it in IIIR to allow more time for other items. Thanks, Simon!

Gopher0

After some discussion, an acceptable compromise was reached to clear up minor problems in RFC 1436. These changes will be made and the RFC resubmitted.

Data Elements

The non-existent Data Elements Working Group reported that the effort to define standard data elements for networked objects was unsuccessful. The work to define standard templates will be done in a different working group. That particular working group has not yet been created, and needs to fit in with the reorganization of the directory services working groups planned for the Directory Services Directorate. More discussion will be taken to the list.

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Jules Aronson	aronson@nlm.nih.gov
Anders Baardsgaard	anders@cc.uit.no
Sepideh Boroumand	sepi@aol.com
Luc Boulianne	lucb@cs.mcgill.ca
Ronald Broersma	ron@nosc.mil
Randy Bush	randy@psg.com
Susan Calcari	calcaris@internic.net
Hallie Carlson	hallie@nsipo.arc.nasa.gov
David Crocker	dcrocker@mordor.stanford.edu
Walt Drummond	drummond@noc.rutgers.edu
Alan Emtage	bajan@bunyip.com
Qin Fang	qin_fang@unc.edu
Carlos Fernandez	carlos@plk.af.mil
Jill Foster	Jill.Foster@newcastle.ac.uk
Paul Francis	Francis@thumper.bellcore.com
Kevin Gamiel	kgamiel@cnidr.org
Arlene Getchell	getchell@es.net
Mei-Jean Goh	goh@mpr.ca
Judith Grass	grass@cnri.reston.va.us
Deborah Hamilton	debbieh@internic.net
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Russ Hobby	rdhobby@ucdavis.edu
Richard Huber	rvh@ds.internic.net
Brendan Kehoe	brendan@zen.org
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Jim Knowles	jknowles@binky.arc.nasa.gov
Mark Kosters	markk@internic.net
John Kunze	jak@violet.berkeley.edu
Paul Lambert	paul_lambert@email.mot.com
Walter Lazear	lazear@gateway.mitre.org
Edward Levinson	elevinson@accurate.com

Ben Levy	seven@ftp.com
Dave Livingston	squirrel@vnet.net
Clifford Lynch	calur@uccmvsa.ucop.edu
April Marine	april@atlas.arc.nasa.gov
Jim Martin	jim@noc.rutgers.edu
Chip Matthes	chip@delphi.com
Mark McCahill	mpm@boombox.micro.umn.edu
Wayne McDilda	wayne@dir.texas.gov
Michael McLay	mclay@eeel.nist.gov
Michael Mealling	michael.mealling@oit.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Keith Moore	moore@cs.utk.edu
Mark Needleman	mhn@stubbs.ucop.edu
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Pete Percival	percival@indiana.edu
Marsha Perrott	perrott@prep.net
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Cecilia Preston	cpreston@info.berkeley.edu
Joyce K. Reynolds	jkrey@isi.edu
Steven Richardson	sjr@merit.edu
Richard Rodgers	rodgers@nlm.nih.gov
Richard Schmalgemeier	rgs@merit.edu
Rickard Schoultz	schoultz@sunet.se
Martin Schulman	schulman@smtp.sprint.com
Henning Schulzrinne	hgs@research.att.com
Mark Smith	mcs@umich.edu
Patricia Smith	psmith@merit.edu
Suzanne Smith	smith@es.net
Karen Sollins	sollins@lcs.mit.edu
Milan Sova	sova@feld.cvut.cz
Simon Spero	ses@unc.edu
Margaret St. Pierre	saint@wais.com
Craig Summerhill	craig@cni.org
Raymond Vega	rvega@cicese.mx
Janet Vratny	janet@apple.com
Chris Weider	clw@bunyip.com
Taehwan Weon	weon@cosmos.kaist.ac.kr
Brien Wheeler	blw@mitre.org
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov

2.10.3 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.

Internet-Drafts

"How to Use Anonymous FTP", 06/15/1993, P. Deutsch, A. Emtage, A. Marine <draft-ietf-iafa-howftp-00.txt>

"Publishing Information on the Internet with Anonymous FTP", 08/17/1993, P. Deutsch, A. Emtage <draft-ietf-iafa-publishing-00.txt>

"Data Element Templates for Internet Information Objects", 08/17/1993, P. Deutsch, A. Emtage <draft-ietf-iafa-templates-00.txt>

2.10.4 Internet School Networking (ISN)

Charter

Chair(s)

Jennifer Sellers: sellers@quest.arc.nasa.gov
Arthur St. George: astgeorg@nsf.gov

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 address relevant issues related to the connection of primary/secondary schools worldwide to the Internet. The key audiences include network service providers and educational policy makers responsible for network access and use. The key areas of focus for this group are advocacy and articulation.

1. Advocacy. The ISN-WG will facilitate dialog between the primary/ secondary education community and the Internet engineering community in order to identify and fulfill the needs of the primary/secondary school community.

2. Articulation. Informed by the group's experience and with input from other IETF working groups, the ISN-WG will articulate solutions to the challenges a school may experience in seeking and gaining a connection to the Internet, as well as the benefits of such a connection. Advantages to Internet connectivity may be articulated by means of pointers to such services as user interfaces, directories, organizations, and training programs, as well as to other resources. Articulation will most often be in the form of periodic documents that address key issues of interest to the school networking community. Representative issues to be addressed by the WG include connectivity models, educational directories, and acceptable use policies.

Goals and Milestones

Mar 1994 Release as an FYI RFC a short document that gives guidance to schools in setting an Acceptable Use Policy.

Mar 1994 Release as an FYI RFC a general document that gives guidance in how to connect to the Internet. Included as an appendix will be models for connectivity that may be of particular interest to schools. This is a joint activity with the User Documents Revisions Working Group.

Jul 1994 Define the information to be included in an online database of educational people involved in networking, recommend a process for collecting and updating the data, and coordinate with a directory services provider to implement the database. Results will be published in an FYI RFC.

Mar 1995 Write a set of two documents, one aimed at connection providers and the other aimed at educational sites, providing guidelines for bringing educational sites online. Included will be a broad definition of connection providers. Interim milestones: 3/94 complete outline; 7/94 first draft; 3/95 review completed.

Internet-Drafts

"FYI on Questions and Answers: Answers to Commonly Asked "Primary and Secondary School Internet User" Questions", 10/05/1993, J. Sellers <draft-ietf-isn-faq-02.txt>

CURRENT MEETING REPORT

Reported by Jennifer Sellers/NASA NREN

Minutes of the Internet School Networking Working Group (ISN)

Thanks to Avri Doria who provided the notes for these minutes.

The meeting began with introductions of those present and an annoucement that the new ISN co-Chairs are Art St. George and Jennifer Sellers. Art was unable attend this meeting.

Jennifer began with some of her own opinions concerning the direction of the group, which included the idea that the group had not completely found its identity yet. She pointed out that while some work had been done on the charter since the last IETF, it was up to the group today to complete the bulk of that work (with refinements *only* to take place on the list after the meeting) in order for the group to continue.

In the discussion that followed, it was brought out that ISN:

- needs to define what the Internet (and perhaps an internet) is to the K-12 community at large,
- needs to help spread the word that there is a defined domain naming structure that should be used by K-12 institutions when attaching to the Internet, and
- could provide a communications media to tell members of the educational community about the opportunities/features that exist on the Internet and how to get access to them.

The group then moved to consideration of the body of the charter. One of the salient features of this discussion was the idea of being a liaison should be prevalent, and the idea of advising should not. IETF working groups can only make recommendations. Another major point discussed was the nature of User Services working groups in general and ISN in particular. It was noted that among other things, the IETF is a second level provider of support. In other words, although many documents may be directed toward and/or reach end-users, our focus within the IETF is to help those providing services to the primary and secondary school community. The charter was widened to reflect the fact that the IETF is an international organization.

Following the work on the body of the charter, April Marine gave a report on some information gathering she had done on the creation of a directory of educational people involved in networking.

This led into a discussion of the milestones, one of which is to define the information to be included in such a directory, and to recommend a process for collecting and updating the data. There are a handful of service providers who may be interested in offering the service, and the milestone includes coordinating with one of them to see that the directory becomes a reality. A number of other milestones were created, and each one had at least one volunteer prepared to carry it through. For more specific information on the wording of the charter and the contents of the milestones, please refer to the ISN charter.

Attendees

Sepideh Boroumand	sepi@aol.com
Arthurine Breckenridge	arbreck@sandia.gov
Hallie Carlson	hallie@nsipo.arc.nasa.gov
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
Ann Cooper	cooper@isi.edu
James Davin	davin@thumper.bellcore.com
Avri Doria	avri@locus.com
Sallie Fellows	sallie@ed.unh.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Alf Hansen	Alf.Hansen@uninett.no
Eugene Hastings	hastings@psc.edu
John Krawczyk	jkrawczy@wellfleet.com
Dave Livingston	squirrel@vnet.net
Dan Magorian	magorian@ni.umd.edu
Bill Manning	bmanning@rice.edu
April Marine	april@atlas.arc.nasa.gov
Marsha Perrott	perrott@prep.net
Joyce K. Reynolds	jkrey@isi.edu
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@quest.arc.nasa.gov
Patricia Smith	psmith@merit.edu
John Vollbrecht	jrv@merit.edu
William Warner	warner@ohio.gov
Gerry White	gerry@lancity.com

2.10.5 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.
Done	Post an Internet-Draft describing NIC interelationships.
Jun 1993	Post an Internet-Draft of a NIC user handoff procedure based on the UCP work.
Done	Post an Internet-Draft describing accessing the nic-profiles data in the X.500 database at Merit.
Dec 1993	Submit the NIC Interelationship document for consideration as an FYI RFC.
Dec 1993	Submit the User Handoff procedures for consideration as an FYI RFC.
Dec 1993	Submit the Nic-Profiles paper for consideration as an FYI RFC.

Internet-Drafts

"Current NIC Interrelationships", 06/28/1993, A. Marine <draft-ietf-nisi-nics-00.txt>

Request For Comments

RFC 1302	"Building a Network Information Services Infrastructure"
RFC 1355	"Privacy and Accuracy Issues in Network Information Center Databases"

CURRENT MEETING REPORT

Reported by April Marine/NASA NAIC

Minutes of the Network Information Services Infrastructure Working Group (NISI)

Thanks, once again, to Marsha Perrott for providing her notes. These minutes follow the agenda items.

Review of Where We Are

The group agreed they were in Houston.

Actually, for benefit of newbies, of which there were several, the group reviewed the Amsterdam meeting (since even several oldbies had trouble making that one!). The discussion primarily centered on how the group got into writing its current document. It was mentioned how in Amsterdam there was agreement to split the one diagram illustrating NIC relationships into three views: US-centric; AP-centric; Euro-centric.

Review "Regional" NIC Diagrams

Does the group agree on these? Discussion was defered.

Plan Final Touches on Relationships Document

The diagrams David Conrad from APNIC had shared were shown. Much discussion *again* on the problems with diagraming (even though the scope of what the group was trying to picture had been narrowed). It was agreed to punt the diagrams and just go with short descriptions on the types of NICs that exist. A registration hierarchy diagram will be included, and more a more complete discussion on the relationship and distinctions between registration services and information services within a NIC will take place. It is possible that a diagram for Europe will be included (but that remains to be seen).

April will try to come up with this with input from the list. Debbie Hamilton, David Conrad, and Scott Paisley will help review.

Discuss Need/Plan to Update RFC 1302/FYI 12

There was basic agreement that this document, "Building a Network Information Services Infrastructure," should be updated. Specific volunteers were April, Debbie, and Dave Livingston. It is not clear that a date was put on this, but the Chair would like to be pretty far along by the March IETF in Seattle.

Summarize Next Steps and Action Items

- Finish NICs document by Seattle (April, Debbie, David Conrad)
- Start update of RFC 1302 by Seattle (April, Debbie, Dave Livingston)

NISI Scope/Liaison Discussion

This was the fun part. Basically, it seems like for the last few sessions, NISI has been operating on more of a "top down" approach than a grassroots "bottom up" approach, and that is not the IETF way. So, the new attendees were asked what they had expected when they walked in, and what had not been talked about that they felt were issues for which they wanted a forum. This was all by way of scoping out what NISI, as a forum for NICs, was best suited for. One idea has been to fold NISI and start up another working group whenever the surge seems to warrant it.

Here are some of the types of things people mentioned they were looking for:

- Role models of "good" NICs including:
  - Guidelines for how to put things online
  - How to set services up in a unified way
- Methods of collecting and disseminating information
- A list of NIC addresses and contact information and ideas of what they had online
- More action on the NIC Forum list
- Guidelines for setting up registration services at a NIC
- Methods of measuring quality of service (except evidently that term can't be used for some political reason :-)
- Cooperating to maintain the information mesh

Some of these ideas are being worked. The NIC Profiles (registry of NICs) idea came up again, so some discussion was had on how to best keep that up to date and make it more accessible. The InterNIC has had some ideas on how to collect and offer this information. Susan Calcari, Pat Smith, and Glenn Mansfield are talking about how to best offer this information, and, if done via X.500, how to get schema registered and recognized.

Jill Foster described how the EuroGopher folks (who are now handling more than just gopher issues) are working to distributed information services within Europe (as a point re: the last bullet above).

It was a good discussion, but no particular item blazed forth as the obvious thing for NISI.

AOB

Susan Calcari described the upcoming (now gone :-) NICFest. She mentioned that there are currently several forums for NIC-types, none of which gets everyone, so there is a need to make sure the forums somehow work together (a la Joyce and Jill keeping each other informed on similar IETF and RARE ISUS activities).

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Sepideh Boroumand	sepi@aol.com
Susan Calcari	calcaris@internic.net
Hallie Carlson	hallie@nsipo.arc.nasa.gov
John Chang	jrc@uswest.com
David Conrad	davidc@iij.ad.jp
Ann Cooper	cooper@isi.edu
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Judith Grass	grass@cnri.reston.va.us
Deborah Hamilton	debbieh@internic.net
Ellen Hoffman	ellen@merit.edu
Barbara Jennings	bjjenni@sandia.gov
Brendan Kehoe	brendan@zen.org
Ben Levy	seven@ftp.com
Dave Livingston	squirrel@vnet.net
Gary Malkin	gmalkin@xylogics.com
Glenn Mansfield	glenn@aic.co.jp
April Marine	april@atlas.arc.nasa.gov
Scott Paisley	paisley@central.bldrdoc.gov
Marsha Perrott	perrott@prep.net
Joyce K. Reynolds	jkrey@isi.edu
Patricia Smith	psmith@merit.edu

2.10.6 Network Training Materials (TRAINMAT)

Charter

Chair(s)

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 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)

Old Business

Minutes of the previous meeting were reviewed. The group thanked Ellen Hoffman, who has resigned as co-chair. A new co-chair is expected to be appointed soon.

Template

Jill Foster presented the latest round in data elements for a template to collect information about training materials. This template has been reviewed by others, including Pete Percival of the Top Node project. Efforts will contine to coordinate this with other projects including NIR. The template has fields for describing the material and optional fields for types such as FTP, Gopher, etc. Jill noted that while such templates may eventually be maintained by professionals such as librarians, the first attempt to use them would be through volunteer efforts. The template was generally approved by the attendees and Jill will post this version to the mailing list.

Catalog

A first attempt to develop the catalog using an earlier version of the template was distributed. This used materials that were already included in the training catalog done by Margaret Isaacs at the University of Newcastle. About fifteen people are working on this effort, testing the template design and verifying entries. Once the initial effort of cataloging the existing materials from Margaret's research is completed, new materials will be sought for inclusion. The results will be published as an RFC.

Subject Headings

Margaret has made a first effort at categorizing the materials, and the group noted that many of the materials included were either documentation or resource guides rather than training materials. After some discussion about scope, it was agreed that such materials should be included as a bibliography but that the catalog should focus on training even if it were shorter that way. A statement at the beginning could be used to define the focus.

Available Materials

Jill raised the issue of what kinds of training materials are needed and where gaps exist. The discussion focused on videos and their usefulness. Jill noted that one trainer had indicated videos were very helpful because students could take them home to study. Updating was done by handouts with the video. Jill suggested discussing the video topic on the mailing list.

Liaison with Other Groups

Jill noted that Susan Harris at Merit was continuing to post relevant materials from the nettrain mailing list to the trainmat group. Jill also described the efforts at Newcastle to develop network training materials for network trainers. These are available for anyone (for non commercial use) from the Newcastle server and include overheads, speaker notes, and handouts. One module covers general networking and the other is on network information retrieval tools. An experimental "training materials Gopher" is being set up at Newcastle. Gopher to `trainmat.ncl.ac.uk` 70.

Subject Resource Guides

Jill raised the topic of tailoring training for specific disciplines, giving an example of the top twenty resources for a particular group. She noted several existing efforts and suggested this as an area for some future work by the group.

Training on the Net

The idea of a hypertext document which could be used for training was discussed. Generally, the issue of interactive training on the network was discussed, and while some efforts are underway, there currently no good source for this. April and Ellen agreed to talk more about developing something.

Roundtable Discussion

The group closed with the traditional roundtable discussion of what each person is doing in terms of training and other efforts in which the group would be interested.

Attendees

Sepideh Boroumand	sepi@aol.com
Hallie Carlson	hallie@nsipo.arc.nasa.gov
Henry Clark	henryc@oar.net
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Ellen Hoffman	ellen@merit.edu
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Brendan Kehoe	brendan@zen.org
April Marine	april@atlas.arc.nasa.gov
Wayne McDilda	wayne@dir.texas.gov
Marsha Perrott	perrott@prep.net
Jennifer Sellers	sellers@quest.arc.nasa.gov
Raymond Vega	rvega@cicese.mx

2.10.7 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 Working 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", 01/06/1994, J. Foster <draft-ietf-nir-status-report-02.txt>

CURRENT MEETING REPORT

Reported by Kevin Gamiel/MCNC - CNIDR

Minutes of the Networked Information Retrieval Working Group (NIR)

Status Report on CNIDR

Jim Fullton gave the current status of CNIDR. CNIDR is continuing freeWAIS server development, the current version being 0.202. A generic search engine application programmer's interface is being developed to allow the integration of any search engine under the WAIS protocol stack. FreeWAIS 1.0, a WAIS server based on Z39.50 1992, is being developed as well. Other projects include a CUSeeMe client for MS Windows and the Global Schoolhouse project.

Status Report on RARE Related Activities

Jill Foster gave an overview of RARE related activities. The NIR Working Group is a joint IETF/RARE working group. The RARE members of the working group met in Warsaw, Poland in October. Anders Gillner (SUNET) reported there on the Eurogopher activities and held a BOF session on recommendations for NEPAs (National (gopher) Entry Point Administrators). The recommendation was that NEPAs coordinate registration of gopher and other information services within their own countries. Discussion included Veronica harvesting at a national level, recommendations for structure of the top level menu, and on sharing management tools.

Other Subgroups of the RARE ISUS (Information Services and User Support) Working Group include MMIS (Multi-Media Information Services) and UNITE (User Network Interface to Everything!). The latter group has undertaken a review of how well certain collections of software for accessing the network work together on various platforms, with a view to identifying what is lacking.

NIR Status Report

The working group split into smaller groups to discuss each section of the NIR status report as a final review. Rick Rodgers from NLM is updating section 5 of the document in light of discussion and will post to Jill for re-integration. Several additional mailing lists were added to the report. There was a question about adding interactive tools such as CUSeeMe and other networked conferencing systems to the report. The group determined such interactive services fall outside the scope of the report and suggested a call for comments on the list.

Maintenance - Originally, the report was to be updated every three months. This was determined to be a difficult task and the group decided in Amsterdam to update only once per year. The Internic has agreed to publish the report under WAIS, Gopher, and World Wide Web. CNIDR will maintain and update the report. Updated templates can be sent via email to nir-update@cnidr.org. Larry Manister offered his time to provide an independent check of say four templates (to be assigned by Jill), if other volunteers did the same.

NIR Evaluation

Kevin Gamiel felt this was one of the areas where the group got carried away and the group needs to re-think what it wants to do here. April Marine discussed the background to the "checklist." This is to be a simple checklist of which clients/servers are available on particular platforms. April had sent a checklist to the mailing list for discussion. The group decided to work on the checklist as a separate document in the interest of moving the current NIR report up the food chain in the next month. However, the checklist should be integrated into the report as an appendix in the future. Changes to April's checklist include: changed word "taxonomy" to "purpose," added "required ancillary tools" attribute and added "stack" attribute. It was noted that the "X.500 Tools Catalog" has a similar template and Brendan Kehoe agreed to look into that catalog and report to the list. A proposal was made that the group should evaluate individual server packages. (UNITE has performed some informal evaluation of client software.) The consensus was that the group should not, as it was a political issue and lists devoted to each tool already serve this purpose.

Review of the Charter

NIR was set up to act as a clearinghouse for information on what is going on in this area. The report was the natural first move on this. There are other things we could do to improve current awareness of what is going on. It might be useful, for example, to post minutes (or at least summaries) of the various related working groups to the NIR list. The group decided to postpone this discussion until the next meeting, by which time the NIR report and checklist should be completed and the mechanisms for updating them should be in place.

Attendees

Farhad Anklesaria	`fxa@boombox.micro.umn.edu`
Sepideh Boroumand	`sepi@aol.com`
Luc Boulianne	`lucb@cs.mcgill.ca`
Susan Calcari	`calcaris@internic.net`
Hallie Carlson	`hallie@nsipo.arc.nasa.gov`

Ann Cooper	cooper@isi.edu
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Jim Fullton	fullton@cnidr.org
Kevin Gamiel	kgamiel@cnidr.org
Arlene Getchell	getchell@es.net
Judith Grass	grass@cnri.reston.va.us
Deborah Hamilton	debbieh@internic.net
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Richard Huber	rvh@ds.internic.net
Brendan Kehoe	brendan@zen.org
Jim Knowles	jknowles@binky.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Ben Levy	seven@ftp.com
Dave Livingston	squirrel@vnet.net
Gary Malkin	gmalkin@xylogics.com
Glenn Mansfield	glenn@aic.co.jp
April Marine	april@atlas.arc.nasa.gov
Larry Masinter	masinter@parc.xerox.com
Mark McCahill	mpm@boombox.micro.umn.edu
Mark Needleman	mhn@stubbs.ucop.edu
Scott Paisley	paisley@central.bldrdoc.gov
Pete Percival	percival@indiana.edu
Marsha Perrott	perrott@prep.net
Joyce K. Reynolds	jkrey@isi.edu
Richard Rodgers	rodgers@nlm.nih.gov
Srinivas Sataluri	sri@internic.net
Milan Sova	sova@feld.cvut.cz
Simon Spero	ses@unc.edu
Margaret St. Pierre	saint@wais.com
Anthony Valle	valle@huntsville.sparta.com
Raymond Vega	rvega@cicese.mx
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov

2.10.8 Uniform Resource Identifiers (URI)

Charter

Chair(s)

Jim Fullton: 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 standard representations of Uniform Resource Locators (URLs) 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, users and network-based tools to uniquely identify specific resources on the Internet, and 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", 07/20/1993, T. Berners-Lee <draft-ietf-uri-url-01.txt, .ps>

"Uniform Resource Names", 10/19/1993, C. Weider, P. Deutsch <draft-ietf-uri-resource-names-01.txt>

CURRENT MEETING REPORT

Reported by Alan Emtage/Bunyip Information Systems

Minutes of the Uniform Resource Identifiers Working Group (URI)

The Uniform Resource Identifiers Working Group held three sessions in Houston. Two were scheduled in advance, and the third was scheduled on-site as an additional session would be necessary to complete the business of the working group. These minutes are separated on a per-session basis. Our thanks to Craig Summerhill for taking notes.

Session 1

Introductory remarks were made by the co-chairs and the minutes of the previous meeting were approved.

Erik Huizer, co-Director of the Applications Area, spoke to clarify certain procedural aspects concerning the running of the working group as well as to comment on some of the discussions which had occurred on the mailing list before the Houston meeting. He made the following points:

1. The URI Working Group falls under the joint administration of the User Services and Applications Areas of the IESG and as such must approach the problems to be solved with both protocol development and the user's perspective in mind.

2. The working group chair has the authority to remind the group that certain topics have already been discussed and to direct members to the mailing list archives and previous minutes of meetings to review the discussion. However the group may still choose to discuss a topic if the issue still exists.

3. Given the nature of the work, discussions revolving around the current installed base, while important, must not be the primary focus of the group. This installed base is very small when compared with expected growth in the information systems area.

4. Rough consensus must be achieved on all documents—this does not, however, mean unanimity. Voting is a gauge and does not necessarily determine if consensus has been obtained.

5. The area directors do not believe that consensus has been reached on the current Uniform Resource Locator (URL) document and would not approve it if submitted.

6. The area directors require the group to produce a companion document to the current URL draft containing a list of functional specifications and requirements. This docu-

ment can then be used to determine if the current URL draft meets the requirements. Similar documents will be required for all UR* protocol specifications.

There was some discussion that the working group would be delayed by having to produce another document. However, the co-Area Director made it clear that while this document did not have to be very long, without it the current URL document would not be recommended for approval by the User Services and Applications Area Directors for acceptance by the full IESG.

Jim Fullton agreed to coordinate efforts on the functional specification draft.

John Kunze made a presentation describing an earlier meeting of some members of the working group. This meeting attempted to iron out some of the problems currently being discussed on the mailing list.

Editor's Note: A detailed list of the problems discussed is available via FTP and mail server from the remote directories as /ietf/uri/uri-minutes-93nov.txt. Refer to Section 1.2 of the proceedings retrieval instructions.

A general discussion on the functional specifications for the URL followed using the points set out by John Kunze. A number of decisions were made.

Editor's Note: A complete list of the decisions made is available via FTP and mail server from the remote directories as /ietf/uri/uri-minutes-93nov.txt. Refer to Section 1.2 of the proceedings retrieval instructions.

A discussion of URN→URL mapping reached no agreement. It was decided that further discussion was required.

Session 2

Jim Fullton presented a document produced by a number of working group members after the previous session, setting out the functional specifications.

Editor's Note: A list of the functional specifications is available via FTP and mail server from the remote directories as /ietf/uri/uri-minutes-93nov.txt. Refer to Section 1.2 of the proceedings retrieval instructions.

Having reached general consensus on the functional specifications, the group reviewed the current URL document for agreement between the two.

Mitra presented a review of the results of the discussion on the mailing list before the Houston meeting. A number of points and decisions were made.

Editor's Note: The results of the discussion are available via FTP and mail server from the remote directories as /ietf/uri/uri-minutes-93nov.txt. Refer to Section 1.2 of the proceedings retrieval instructions.

General consensus was reached as to the agreements on changes to the current draft. The author Tim Berners-Lee will be asked to make the changes.

Discussion on the current URN and URC drafts was started. Larry Masinter suggested that this be postponed until functional specification and requirements documents could be produced.

The functional specifications for the URNs was discussed and the following points were made:

- They should provide persistent unique names for resources.
- They should be able to detect sameness of resources.
- Should they be used in conjunction with a description for a process for mapping or resolving to locators URN→URL?

Session 3

The chair proposed that the session be split into discussions on the functional specifications for the URN and then on the current URN draft.

Peter Deutsch had some comments on the working group process

- We have two groups (IRTF and USRV/APPS) melting into one working group, and we have done a great deal of work in the last year.

- We traditionally do research and development, and we may be doing more R&R (research and release).

He then presented suggested functional specifications for URNs as proposed by a group of working group members.

Editor's Note: A list of the suggested functional specifications for URNs and resulting discussion is available via FTP and mail server from the remote directories as /ietf/uri/uri-minutes-93nov.txt. Refer to Section 1.2 of the proceedings retrieval instructions.

Karen Sollins and Larry Masinter volunteered to write the initial draft of this document.

The group also decided that discussions on the functional specfications of the URCs should be started.

Attendees

Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Farhad Anklesaria	fxa@boombox.micro.umn.edu
Jules Aronson	aronson@nlm.nih.gov
Sepideh Boroumand	sepi@aol.com
Luc Boulianne	lucb@cs.mcgill.ca
Randy Bush	randy@psg.com
Susan Calcari	calcaris@internic.net
Brian Capouch	brian@saintjoe.edu
Hallie Carlson	hallie@nsipo.arc.nasa.gov
James Conklin	jbc@bitnic.educom.edu
John Curran	jcurran@nic.near.net
Peter Deutsch	peterd@bunyip.com
Walt Drummond	drummond@noc.rutgers.edu
Alan Emtage	bajan@bunyip.com
Qin Fang	qin_fang@unc.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
Jim Fullton	fullton@cnidr.org
Kevin Gamiel	kgamiel@cnidr.org
Arlene Getchell	getchell@es.net
Judith Grass	grass@cnri.reston.va.us
Deborah Hamilton	debbieh@internic.net
Russ Hobby	rdhobby@ucdavis.edu
Jeroen Houttuin	houttuin@rare.nl
Richard Huber	rvh@ds.internic.net
Erik Huizer	Erik.Huizer@SURFnet.nl
Barbara Jennings	bjjenni@sandia.gov
Brendan Kehoe	brendan@zen.org
John Klensin	Klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Andrew Knutsen	andrewk@sco.com
John Kunze	jak@violet.berkeley.edu
Edward Levinson	elevinson@accurate.com
Ben Levy	seven@ftp.com
Clifford Lynch	calur@uccmvsa.ucop.edu
April Marine	april@atlas.arc.nasa.gov
Jim Martin	jim@noc.rutgers.edu
Larry Masinter	masinter@parc.xerox.com
Chip Matthes	chip@delphi.com
Mark McCahill	mpm@boombox.micro.umn.edu
Michael McLay	mclay@eeel.nist.gov
Michael Mealling	michael.mealling@oit.gatech.edu
Mitra	mitra@pandora.sf.ca.us

Pushpendra Mohta	pushp@cerf.net
Keith Moore	moore@cs.utk.edu
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Rakesh Patel	rapatel@pilot.njin.net
Pete Percival	percival@indiana.edu
Karen Petraska-Veum	karen.veum@gsfc.nasa.gov
Cecilia Preston	cpreston@info.berkeley.edu
Joyce K. Reynolds	jkrey@isi.edu
Steven Richardson	sjr@merit.edu
Richard Rodgers	rodgers@nlm.nih.gov
Jon Saperia	saperia@zko.dec.com
Srinivas Sataluri	sri@internic.net
Rickard Schoultz	schoultz@sunet.se
Michael Schwartz	schwartz@cs.colorado.edu
Mark Smith	mcs@umich.edu
Patricia Smith	psmith@merit.edu
Karen Sollins	sollins@lcs.mit.edu
Milan Sova	sova@feld.cvut.cz
Simon Spero	ses@unc.edu
John Stewart	jstewart@cnri.reston.va.us
Craig Summerhill	craig@cni.org
Thuan Tran	thuan@xylogics.com
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
Janet Vratny	janet@apple.com
William Weems	wweens@oac.hsc.uth.tmc.edu
Chris Weider	clw@bunyip.com
Taehwan Weon	weon@cosmos.kaist.ac.kr
Jackie Wilson	Jackie.Wilson@msfc.nasa.gov

2.10.9 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 USERDOC2 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" (RFC 1175, 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 USERDOC2 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.
Oct 1993	Post a draft of long bibliography and Getting Connected documents
Nov 1993	Review Getting Connected document (with ISN) at Houston IETF
Nov 1993	Review long bibliography at Houston IETF
Dec 1993	Final review of Internet-Drafts
Jan 1994	Submit Getting Connected document to IESG for consideration as an FYI RFC
Jan 1994	Submit long bibliography to IESG for consideration as an FYI 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)

The USERDOC2 Working Group had a small but enthusiastic set of attendees due to a conflicting applications group meeting. The group set a new record by having everyone in the room volunteer to assist in writing and reviewing documents in progress!

Working Group Status

The meeting began with a review of the charter and listing of the address for the group's mailing list. The charter, which was not in the archives at the last IETF, has been restored and the milestones have been updated.

FYI 19

Current projects were reviewed. FYI 19 (RFC 1463) is getting distributed through many NICs and has generally been successful in reaching Internet novices. Concern was expressed about how to get wider distribution and, in discussion, support was indicated for the proposed distribution and announcement working group discussed in the User Services Working Group (USWG) meeting on Wednesday. The on-line archive of documents which goes with FYI 19's bibliography has continued to be updated. (For information on the archive, send email to `nis-info@nic.merit.edu` with the text: send access.guide). FYI 19 will be reviewed at the next IETF for possible revision and updating.

Connecting to the Internet Document

The joint project with the Internet School Networking Working Group (ISN) for a "Connecting to the Internet" document is ongoing, and a draft should be sent to the list in the next few weeks. If the draft gets a good reception from the list, it will be moved on for consideration as an RFC before the next IETF. If more discussion is needed, the draft will be reviewed when the working group reassembles in Seattle.

Revision of RFC 1175

A "Not Quite an Internet-Draft" bibliography was distributed and discussed which will update RFC 1175 (FYI 3). The new document will focus on books, journals and other bibliographies rather than all possible documentation. The final document is scheduled to be completed before the next IETF. Volunteers were recruited to work on abstracts for the books listed. Additional books need to be added to the first draft and it will be posted to the user-doc list. It was also decided to add a more comprehensive section on RFCs and FYIs. The final version will divide materials into categories to assist users in finding what they need. This is expected to be completed by the next IETF.

Attendees

Sepideh Boroumand	sepi@aol.com
Hallie Carlson	hallie@nsipo.arc.nasa.gov
Ann Cooper	cooper@isi.edu
Ellen Hoffman	ellen@merit.edu
Matt Hood	hood@nsipo.nasa.gov
Lenore Jackson	jackson@nsipo.arc.nasa.gov
Marsha Perrott	perrott@prep.net
Jennifer Sellers	sellers@quest.arc.nasa.gov
Raymond Vega	rvega@cicese.mx

2.10.10 User Services (USWG)

Charter

Chair(s)

Joyce K. Reynolds: jkrey@isi.edu

Mailing Lists

General Discussion: us-wg@nic.near.net
To Subscribe: us-wg-request@nic.near.net
Archive: ftp.near.net:/mail-archives/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); and

- 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

None specified

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/Information Sciences Institute

Minutes of the User Services Working Group (USWG)

This USWG session was multicast for the first time for those folks who were unable to attend the Houston IETF. Joyce Reynolds provided a report on 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.

New FYI RFC publications since the last IETF:

FYI 21	RFC 1491	"A Survey of Advanced Usages of X.500"	July 1993
FYI 17	RFC 1539	"The Tao of IETF - A Guide for New Attendees of the Internet Engineering Task Force"	October 1993

User Services Area Internet-Drafts posted since the last IETF:

draft-ietf-isn-faq-01.txt
draft-ietf-ids-pilots-00.txt
draft-ietf-iafa-howftp-00.txt
draft-ietf-ids-x500-survey-02.txt
draft-ietf-iiir-html-01.txt, ps
draft-ietf-nisi-nics-00.txt
draft-ietf-uri-resource-names-01.txt

Gary Malkin briefly discussed the "DAWG" (Distribution and Announcement Working Group) idea that has been sitting on the USWG's backburner for awhile. A BOF will be held at the next IETF to see if there is further interest in this topic. Ann Cooper led a presentation and discussion on the US Domain. Ann and Jon Postel fielded a "questions and answers" session on this topic. (Note: Ann's slides follow these minutes.)

Jill Foster and Joyce Reynolds reported on the RARE ISUS meetings and the EARN Network Services Conference held in Warsaw, Poland, in which they participated. Jill announced the INET94/JENC5 Call For Papers—User Information Track to the USWG. Jill was asked to run this track, and asked Joyce if she would be co-track leader.

There was continued discussion from the Amsterdam IETF on Bill Manning's thoughts about how to "empower" users to utilize and document tools. The concept about the development of a series of notes that will address the manners and morals of the collective body was addressed. Another topic that was discussed was how to deal with the basics (e.g., How do "I" get attached?). It was deemed more appropriate to address these topics in the USERDOC2 Working Group and/or in additional IETF User Services Area working groups.

At the next USWG session during the Seattle IETF, we will be holding the USWG sessions at the very first part of the week and will continue to multicast the session for those who do not have travel funding to attend the USWG. Joyce has already spoken to Megan Walnut about the scheduling of the this group. The USWG also will be discussing what other avenues that this group can undertake, along with what we have already been doing *very* successfully :-).

GOLD STARS FOR EVERYONE!

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Jules Aronson	aronson@nlm.nih.gov
Sepideh Boroumand	sepi@aol.com
Susan Calcari	calcaris@internic.net
Hallie Carlson	hallie@nsipo.arc.nasa.gov
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
Ann Cooper	cooper@isi.edu
Avri Doria	avri@locus.com
Sallie Fellows	sallie@ed.unh.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Kevin Gamiel	kgamiel@cnidr.org
Judith Grass	grass@cnri.reston.va.us
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Deborah Hamilton	debbieh@internic.net
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Ellen Hoffman	ellen@merit.edu
Melanie Humphrey	msh@uiuc.edu
Brendan Kehoe	brendan@zen.org
Mark Kosters	markk@internic.net
Ben Levy	seven@ftp.com
Dave Livingston	squirrel@vnet.net
Gary Malkin	gmalkin@xylogics.com
April Marine	april@atlas.arc.nasa.gov
Mark McCahill	mpm@boombox.micro.umn.edu
Michael Mealling	michael.mealling@oit.gatech.edu
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Marsha Perrott	perrott@prep.net
David Piscitello	wk04464@worldlink.com

Jon Postel	postel@isi.edu
Cecilia Preston	cpreston@info.berkeley.edu
Joyce K. Reynolds	jkrey@isi.edu
Rickard Schoultz	schoultz@sunet.se
Martin Schulman	schulman@smtp.sprint.com
Jennifer Sellers	sellers@quest.arc.nasa.gov
Patricia Smith	psmith@merit.edu
Milan Sova	sova@feld.cvut.cz
Taehwan Weon	weon@cosmos.kaist.ac.kr
Scott Williamson	scottw@nic.ddn.mil

THE US DOMAIN

ANN COOPER
ISI

3 NOV 93

USWG
IETF

THE US DOMAIN

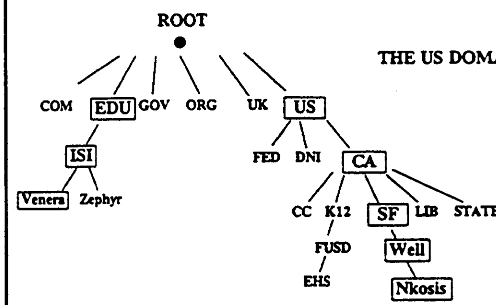

VIEW OF SECOND LEVEL DOMAINS UNDER US

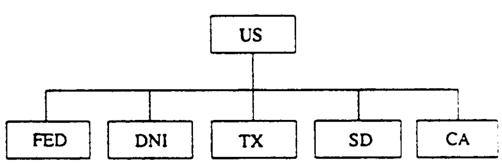

RELATIONSHIP BETWEEN
INTERNIC REGISTRATION SERVICES
AND ISI

- ISI ADMINISTERS US DOMAIN
- UNDER CONTRACT FROM NETWORK SOLUTIONS
- WE WORK CLOSELY WITH INTERNIC SETTING POLICIES.

USC/ INFORMATION SCIENCES INSTITUTE

WHO ADMINISTERS THE US DOMAIN?

- Jon Postel
- Ann Cooper
- US-DOMAIN@ISI.EDU

USC/ INFORMATION SCIENCES INSTITUTE

OUR RESPONSIBILITIES AS US DOMAIN ADMINISTRATOR

- Assignment of all DNS Names ending With ".US"
- Provide Nameservice
- Delegate Branches of Namespace
- Inform Administrators of Updates
- Settle Disputes Among Administrators

USC/ INFORMATION SCIENCES INSTITUTE

US DOMAIN NAMESERVERS

- VENERA.ISI.EDU
- NS.ISI.EDU
- RS.INTERNIC.NET
- NS.UU.NET
- NS.CSL.SRI.COM
- EXCALIBUR.USC.EDU
- MIL.ARL.ARMY.MIL

USC/ INFORMATION SCIENCES INSTITUTE

RESPONSIBILITIES OF ADMINISTRATORS OF DELEGATED DOMAINS

- Knowledgeable technical contact
- Provide two independent name servers
- Accept all applicants on an equal basis.
- Provide timely processing of requests.
- Follow guidelines in RFC 1480
- Use US Domain Template for applicants
- Register IN-Addr PTR and NS records

NETWORK PROVIDERS

- Follow guidelines in RFC 1480.
- Assist applicants in name selection
- Use US Domain Template
- Register schools under K12, CC, and TEC instead of EDU.
- Refer small businesses, and individuals to register under "locality" not COM.
- Encourage state agencies to register in US Domain instead of GOV
- Register libraries, museums, city, county agencies under LIB, MUS, CI, or CO.
- Get updated copy of : "us_domain_delegated_domains "
- Send US Domain application to : <us-domain@isi.edu>

ISI USC/INFORMATION SCIENCES INSTITUTE

LIST OF DELEGATED ZONES

- FTP
  - VENERA.ISI.EDU
  - in-notes/us-domain-delegated.txt
- EMAIL: TO: RFC-INFO@ISI.EDU
  - Help: US_DOMAIN-DELEGATED_DOMAINS

ISI USC/INFORMATION SCIENCES INSTITUTE

DELEGATING WHOLE STATES

- Discover new situations to evaluate
- May need to create new subdomain in best interest for all states
- MUS for Museums, COG for Councils of Governments

ISI USC/INFORMATION SCIENCES INSTITUTE

REGISTERING IN US DOMAIN DOES NOT PROVIDE

- GRANT PERMISSION TO USE INTERNET OR ITS COMPONENT NETWORKS
- ALLOCATE IP ADDRESS
- PROVIDE INTERNET FORWARDING HOSTS FOR NON-INTERNET SITES
- REGISTER NETWORKS WITH INTERNIC

USC/ INFORMATION SCIENCES INSTITUTE

TYPES OF REGISTRATIONS

- DIRECT
  - IP Hosts with IP addresses
    "A" records
  - Non-IP hosts (UUCP)
- DELEGATED
  - Provide Two Nameservers,
  - No MX records pointing to Internet Forwarding Hosts

USC/ INFORMATION SCIENCES INSTITUTE

TYPES OF REGISTRATIONS IN US DOMAIN

- K12 schools and related organizations
- Community Colleges
- Technical Colleges
- Libraries
- Museums
- State Government Agencies
- City Government
- County Government
- Individuals With Computers at Home
- Small businesses / BBS

USC/ INFORMATION SCIENCES INSTITUTE

WHAT ENTITIES GO WHERE

- ONLY FEDERAL AGENCIES IN GOV
- STATE AND LOCAL AGENCIES IN US DOMAIN
- ONLY 4 YEAR COLLEGES IN EDU
- K12, CC, TEC, SCHOOLS IN US DOMAIN
- LIBRARIES, MUSEUMS, IN US DOMAIN
- SMALL BUSINESSES, INDIVIDUALS, BBS IN US DOMAIN

USC/ INFORMATION SCIENCES INSTITUTE

THE CHANGING INTERNET

- SO MUCH FOR GOOD INTENTIONS
- INCONSISTENCY IN WORLD AND INTERNET
- DIFFERENT DOMAIN NAMES
  - LA.GOV VS. STATE.LA.GOV
- EDUCATE STATE OFFICES
- GROWTH IN INTERNET CHANGES BALANCE
  - More registrations under US

USC/ INFORMATION SCIENCES INSTITUTE

BRANCHES UNDER US DOMAIN

- LOCALITY= CITIES, COUNTIES
- K12 = K THRU 12 SCHOOLS
- CC = COMMUNITY COLLEGES
- TEC = TECH/VOC SCHOOLS
- STATE = STATE AGENCIES
- LIB = LIBRARIES
- DNI = DIST. NAT. INST'S.
- FED = FEDERAL AGENCIES
- MUS = MUSEUMS
- GEN = GENERAL (BBS)
- COG = COUNCILS OF GOV

VIEW OF LOCALITY

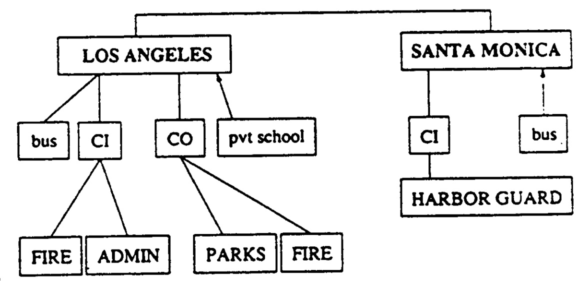

USC/ INFORMATION SCIENCES INSTITUTE

LOCALITY NAMES

- Full city names spelled out unless very well-known abbreviation
- If we are already using abbreviation we will continue
- The Western Union city Mnemonics code. (What went wrong)
- Munroe Falls = MRFS
- Use postal name extensions (See Zip Code directory)

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN LOCALITY

- COLUMBUS.OH.US
  Columbus, Ohio
- SF.CA.US
  San Francisco
- HOU.TX.US
  Houston, Texas
- CI.PHOENIX.AZ.US
  Phoenix city gov't agencies
- CI.SEATTLE.WA.US
  Seattle gov't agencies
- CO.DONA-ANA.NM.US
  Dona-Ana county gov't agencies

SCHOOL VIEW

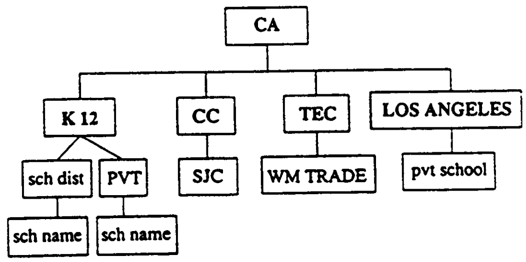

USC/INFORMATION SCIENCES INSTITUTE

TYPES OF THINGS IN K12

- SCHOOLS (PUBLIC AND PRIVATE)
- SCHOOL DISTRICTS
- SCHOOL BOARDS
- SPECIAL EDUCATIONAL SERV. UNITS
- DEPT. OF EDUCATION (STATE, CITY, COUNTY)
- PUBLIC FUNDED SCHOOLS
- STATE AGENCIES CONNECTING K12
- MULTI-DISCIPLINARY RESEARCH ORG.

USC/INFORMATION SCIENCES INSTITUTE

EXAMPLES IN K12

- BOHS.BREA.K12.CA.US
  Brea Olinda High School
- KJRHS.CUPERTINO.CA.US
  Kennedy Jr. High School
- SFUSD.K12.CA.US
  San Francisco Unified School District
- OCDE.K12.CA.US
  Orange County Dept. of Education
- CDE.K12.CA.US
  California Dept. of Education
- CAMS.K12.CA.US
  California Academy of Math and Science
- DOIT.K12.WA.US
  K12 State multi-disciplinary Res Org.
- EXCEED.K12.TN.US
  Consortium providing K12 connectivity

USC/INFORMATION SCIENCES INSTITUTE

"SCHOOL" THINGS THAT CAN GO UNDER OTHER DOMAINS

- STATE DEPARTMENT OF EDUCATION
  - IND.DOE.STATE.IN.US
- A COUNTY DEPT. OF EDUCATION
  - ED.CO.ORANGE.CA.US
- US MILITARY SCHOOLS OVERSEAS
  - K12.DOD.FED.US
- PRIVATE SCHOOLS
  - TOWN.SF.CA.US

USC/ INFORMATION SCIENCES INSTITUTE

TYPES OF THINGS IN CC and TEC

- Community Colleges
- Technical /Vocational Schools
- School libraries, bookstores, campus administrative offices.

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN CC and TEC

JSCC.CC.TN.US

KCCCD.CC.CA.US

TAFT.CC.CA.US

AUGUSTA.TEC.GA.US

STIM.TEC.TN.US

View of
STATE, LIBRARIES, MUSEUMS AND GENERAL AGENCIES

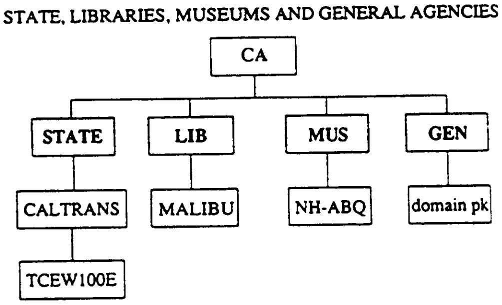

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN STATE

- ARCHIVES.STATE.RI.US
- KDLA.STATE.KY.US
  KY Dept. of Libraries and Archived
- DIS.STATE.KY.US
  KY Dept. of Information Systems
- TC.STATE.KY.US
  KY Dept of Transportation Cabinet
- DOL.STATE.VT.US
  VT Dept. of Libraries
- DESE.STATE.MO.US
  MO Dept of Elementary/Secondary Ed.
- IND.DOE.STATE.IN.US
  IN Dept of Education

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN LIB

- STATE, CITY, COUNTY PUBLIC, PRIVATE AND RESEARCH LIBRARIES

  - KCPL.LIB.MO.US
    Kansas City, Missouri Public Library

  - QUEENS.LIB.NY.US
    Queens Public Library

  - LINC.LIB.IL.US
    Lincoln Integrated Network Consortium

  - MNL.LIB.MA.US
    Minuteman Library Network

  - STATE.LIB.OH.US
    State Libraries of Ohio

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN MUS

- NMMNH-ABQ.MUS.NM.US
  NM Natural History Museum of Albq.

USC/ INFORMATION SCIENCES INSTITUTE

EXAMPLES IN GEN

- DS.GEN.NJ.US

  BBS featuring Usenet News, Mail

DELEGATED DOMAINS UNDER .US

	K12	CC	TEC	STATE	LIB	MUS	GEN
AR	X						
AZ	X			X	X	X	
CA	X	X					
CO	X	X	X		X		
DE	X	X	X	X	X		
FL	X	X	X	X	X		
GA	X		X	X	X		
IA	X	X	X		X		
ID	X	X	X	X			
IL	X	X	X		X		
KY	X	X	X	X	X	X	X
LA	X			X			
ME				X			
MI	X	X	X		X		
MN	X	X	X	X	X	X	
MO	X	X		X	X		
MS	X			X			
NC	X	X		X	X		
ND	X			X			
NE	X			X	X		
NH	X		X				
NJ	X						
NM	X						
NY	X			X			
OH	X	X	X	X	X	X	X
OR	X				X		
PA	X						
RI	X						
SC	X	X	X	X	X		
SD	X			X			
TX	X	X		X	X		
UT	X			X	X		
VA	X		X	X			
WV	X	X	X	X	X	X	X
VT	X			X			

USC/ INFORMATION SCIENCES INSTITUTE

WHERE DO I SEND MY APPLICATION?

- US-DOMAIN@ISI.EDU

USC/ INFORMATION SCIENCES INSTITUTE

WHAT TEMPLATE DO I USE?

- US Domain Template 6/93
- Internic Domain Template

USC/ INFORMATION SCIENCES INSTITUTE

APPLICATIONS SENT TO INTERNIC

- REQUESTS ENDING IN GOV, EDU, ORG
- FORWARDED TO US DOMAIN REGISTRAR
- EDUCATE APPLICANT ON US DOMAIN NAMING STRUCTURE
- DELAYS

USC/ INFORMATION SCIENCES INSTITUTE

WHOIS DATABASE

- WHAT IS BEING REGISTERED
- INTERNIC POLICY
- STATE.OH.US
- K12.OH.US
- COLUMBUS.OH.US
- WHOIS SOFTWARE
    - We need to let Directory Services Information Services WGs know we need a solution.
- RESPONSIBILITY OF DELEGATED SUBDOMAINS

USC/ INFORMATION SCIENCES INSTITUTE

WHERE TO ADD THE IN-ADDR NS RECORD

- The Network number needs to be listed in root servers IN-ADDR zone file on the INTERNIC with a NS delegation to your name server.
- Send msg to (negreg@internic.net)
  - In the Root Server the following record is added:
- For example: 198.146.8
- 8.146.198.IN-ADDR.ARPA NS dsccvx.dscc.tn.us (Nameserver Name)

USC/ INFORMATION SCIENCES INSTITUTE

WHAT YOUR NETWORK ADMINISTRATOR NEEDS TO REGISTER

- The Administrator of your network has to add to his nameserver an IN-ADDR PTR record for your host.
- For example: 198.146.8.11

11.8.146.198.IN-ADDR.ARPA PTR DSCC.CC.TN.US (Hostname)

USC/ INFORMATION SCIENCES INSTITUTE

REFERENCES

- RFC 1480 – "The US Domain" Cooper, A. and J. Postel, USC/ISI December 1992
- "DNS and Bind – Help for UNIX System Administrators," Albitz, P., C Liu, O'Reilly and Associates, Inc., October 1992.
- RFC 1359 – "Connecting to the Internet", ACM SIGUCCS Networking Taskforce, What Connecting Institutions Should Anticipate", August 1992.
- RFC 1033 – "Domain Administrators Operations Guide", Mark Lottor. SRI, November 1987.

2.10.11 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 RFC 954, 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.
Done	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 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", 10/27/1993, C. Weider, J. Fullton, S. Spero <draft-ietf-wnils-whois-02.txt>

"Whois and Network Information Lookup Service Whois++", 07/06/1993, J. Gargano, K. Weiss <draft-ietf-wnils-whois-lookup-00.txt>

CURRENT MEETING REPORT

Reported by Joan Gargano/University of California, Davis

Minutes of the WHOIS and Network Information Lookup Service Working Group (WNILS)

The minutes of July 13, 1993 were unanimously approved without changes.

Status of the WHOIS++ Architecture

Peter Deutsch presented several additions required to the protocol:

- Inclusion of MIME
- Remove the counters for the number of responses in short responses
- Boolean searching

There was some discussion on the current use of WHOIS++ for URN to URL conversions as part of the URI Working Group activities. Peter Deutsch described the use of domain name service to find a top level WHOIS++ server for URN information. It was agreed that this work is helpful to the development of WHOIS++, but unless this work affects the WHOIS++ protocol specifications, further discussions will continue within the other working groups.

Status of the Distributed WHOIS++ Model - Centroids

Chris Weider described changes to the document.

- An X-hierarchy field to provide metainformation for intelligent traversing of the index service has been added to provide topology, geographical and administrative values. Further extensions are needed.

- A mechanism has been added for weighting information for attributes.

- Identifiers at the start and end of attribute information have been added.

- There is now a way to signify any field or any value in a field in a centroid change report.

- A field has been added to designate case sensitivity of string values.

- Simon Spero described the mechanism for searching a centroid tree from the bottom, up.

- Features that are still needed include:
  - Need pointers that get from a top level centroid to the bottom levels, bypassing intermediate centroids.
  - Replication supported by the ability to pass over the entire contents of a centroid, rather than a subset.

Status of WHOIS++ Clients

Jim Fullton mentioned the use of WHOIS++ in support of networked information retrieval, and the type of client development that is occurring as part of other application development. He recommended that the description of clients focus on the work in these areas rather than white pages interfaces. Some discussion followed regarding the continuing need for WHOIS++ as a white pages service.

Status of Recommended Modifications to the WHOIS Protocol

Joan Gargano introduced that this document is available as an Internet-Draft and has not been modified for six months. Final submission as an Informational item was delayed pending stabilization of the architecture document. It does not appear it will require further modification due to protocol development. Discussion focused on the final status of this paper.

Discussion of WHOIS++ Implementations

The following implementations are available:

- Alan Emtage's implementation
  Alan Emtage `bajan@bunyip.com`
  FTP: `ftp.ucdavis.edu`: `/dist/bunyip-whois++-1.0a.tar.Z`

- DUA Interface using LDAP
  Mark Prior `mrp@itd.adelaide.edu.au`
  FTP: `ftp.adelaide.edu.au`: `/pub/whois/whois++beta.tar.Z`

- PERL and dbm
  Rickard Schoultz `schoultz@sunet.se`
  FTP: `othello.admin.kth.se`: `/pub/schoultz/kth-whois++-1.1a.tar.Z`

- PERL 4.036 on SunOS 4.1.xxx
  Martin Hamilton m.t.hamilton@lut.ac.uk
  FTP: genie.lut.ac.uk: /lut-whois++-alpha.tar.Z

There are a few terminal based clients available. Chris Weider offered to compile a list of clients and servers and post them to the ietf-wnils mailing list. There is still a need for graphical user interfaces to WHOIS++ as a directory service to provide wide scale testing.

Update of Goals and Milestones

The following goals need revision:

- 7/31/93 - Submit the WHOIS and Network Information Lookup Service Recommendations document as an Internet-Draft.
- 7/31/93 - Submit the WHOIS++ protocol document as an Internet-Draft.
- 7/31/93 - Submit the "Architecture of the WHOIS++ Index Service" document as a revised Internet-Draft.

It was decided that all three of the working papers will be completed and submitted as Proposed Standards. Protocol work will be frozen for six months to allow for software development. In the meantime the working group will discuss the future direction of WNILS. Areas for discussion include:

- Developing a role for WHOIS++ as a directory service, potentially including work on data elements.
- Continued work on WHOIS++ as a component of other network information retrieval tools.
- Closing the working group after completion of the work on the current protocols.

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Jules Aronson	aronson@nlm.nih.gov
Sepideh Boroumand	sepi@aol.com
Randy Bush	randy@psg.com
Brian Capouch	brian@saintjoe.edu
C. Allan Cargille	allan.cargille@cs.wisc.edu
Hallie Carlson	hallie@nsipo.arc.nasa.gov
Ann Cooper	cooper@isi.edu

Peter Deutsch	peterd@bunyip.com
Alan Emtage	bajan@bunyip.com
Qin Fang	qin_fang@unc.edu
Jill Foster	Jill.Foster@newcastle.ac.uk
Jim Fullton	fullton@cnidr.org
Kevin Gamiel	kgamiel@cnidr.org
Joan Gargano	jcgargano@ucdavis.edu
Judith Grass	grass@cnri.reston.va.us
Deborah Hamilton	debbieh@internic.net
Roland Hedberg	Roland.Hedberg@rc.tudelft.nl
Russ Hobby	rdhobby@ucdavis.edu
Richard Huber	rvh@ds.internic.net
Byonghak Kim	bhkim@cosmos.kaist.ac.kr
Mark Kosters	markk@internic.net
John Kunze	jak@violet.berkeley.edu
Dave Livingston	squirrel@vnet.net
Clifford Lynch	calur@uccmvsa.ucop.edu
April Marine	april@atlas.arc.nasa.gov
Mark McCahill	mpm@boombox.micro.umn.edu
Michael McLay	mclay@eeel.nist.gov
Michael Mealling	michael.mealling@oit.gatech.edu
Keith Moore	moore@cs.utk.edu
Mark Needleman	mhn@stubbs.ucop.edu
Lars-Gunnar Olsson	Lars-Gunnar.Olsson@data.slu.se
Scott Paisley	paisley@central.bldrdoc.gov
Rakesh Patel	rapatel@pilot.njin.net
Pete Percival	percival@indiana.edu
Cecilia Preston	cpreston@info.berkeley.edu
Richard Rodgers	rodgers@nlm.nih.gov
Srinivas Sataluri	sri@internic.net
Richard Schmalgemeier	rgs@merit.edu
Rickard Schoultz	schoultz@sunet.se
Martin Schulman	schulman@smtp.sprint.com
Patricia Smith	psmith@merit.edu
Simon Spero	ses@unc.edu
Margaret St. Pierre	saint@wais.com
Craig Summerhill	craig@cni.org
Anthony Valle	valle@huntsville.sparta.com
William Weems	wweens@oac.hsc.uth.tmc.edu
Chris Weider	clw@bunyip.com
Russ Wright	wright@lbl.gov

Chapter 3

Network Status Briefings

3.1 NASA Science Internet

Reported by Jeffrey Burgan/NASA Ames Research Center

NASA Science Internet Office
Ames Research Center
Information and Communications Systems Division

NASA Science Internet Status Report

28th IETF Meeting
November [illegible], 1993

Jeffrey Burgan
NASA Science Internet Office
jeff@nsipo.nasa.gov

NASA Science Internet Office
Ames Research Center
Information and Communications Systems Division

Recent Activities

- UK "fat-pipe" transitioned to ICM/Sprint
  NASA now has 256Kb dedicated link to ULCC
- Australia (AARNet) link upgraded to T1
- Japanese links reterminated to FIX-West
  - WIDE - 192Kb
  - TISN - 512Kb upgrade
- Korea link upgraded to 256K
- Hong Kong link upgraded to 128K
- FIX-West transitioned to FDDI
- NASA/NSF support for Antarctic connectivity

NASA Science Internet Office
Ames Research Center
Information and Communications Systems Division

Overview

- NSI is a worldwide multiprotocol network serving NASA science
- Built primarily out of Proteon routers
- Uses leased lines at speeds from 9.6 Kbps to 45 Mbps (T3)
- Connects over 150 locations in 14 countries
- Integrated routing with the Internet (TCP/IP) and worldwide DECnet networks
- Centrally managed from a dedicated 24x7 NOC at Ames
- Supports mail and protocol interoperability services

NASA Science Internet Office
Ames Research Center
Information and Communications Systems Division

Future Activities

- NSI backbone upgrades utilizing inverse multiplexers
- General purpose Internet connectivity to Russia via IKI (Russian Space Research Institute) at 256Kb
- Link to Argentina Space Institute at 64Kb
- BGP deployment
- Multicast routing using MOSPF
  Proteon supports MOSPF/DVMRP coupling
- ATM

NASA Science Internet

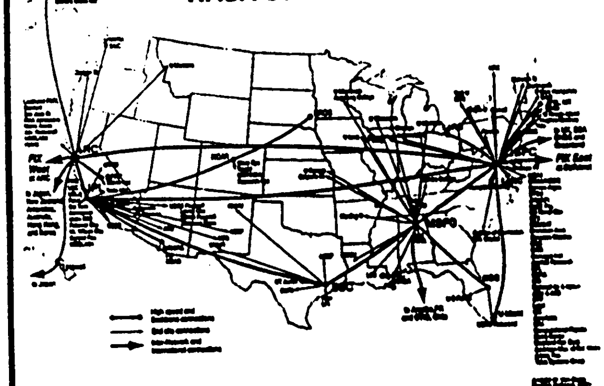

3.2 ANSNet

Reported by Jordan Becker/Advanced Network & Services

ANSNet Update

November 1993 - IETF

Jordan Becker,
Advanced Network & Services
becker@ans.net

ans

Growth in Network Usage

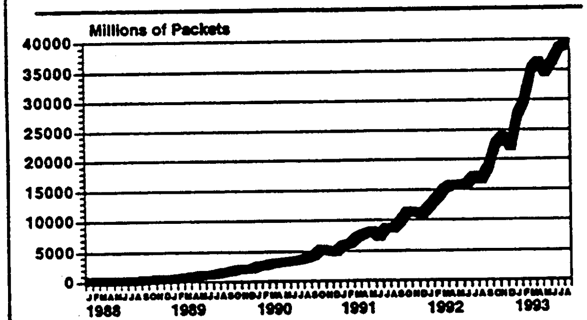

ans

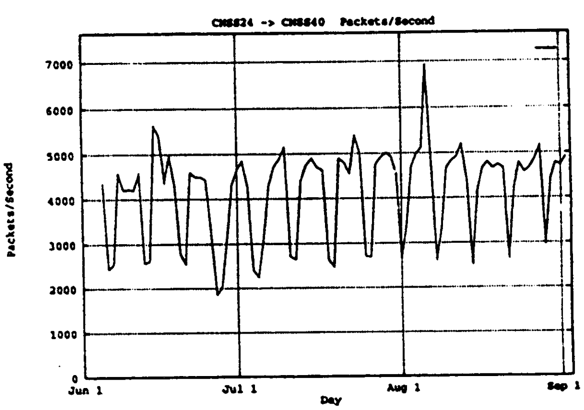

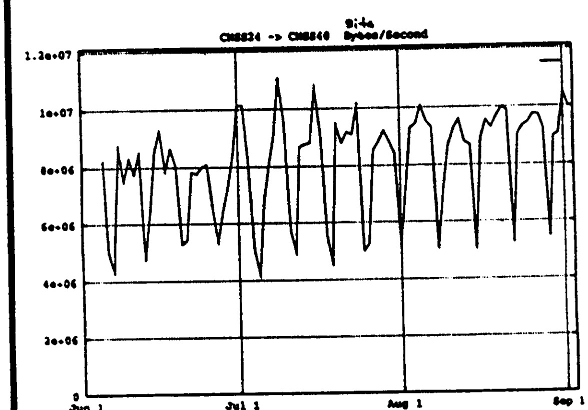

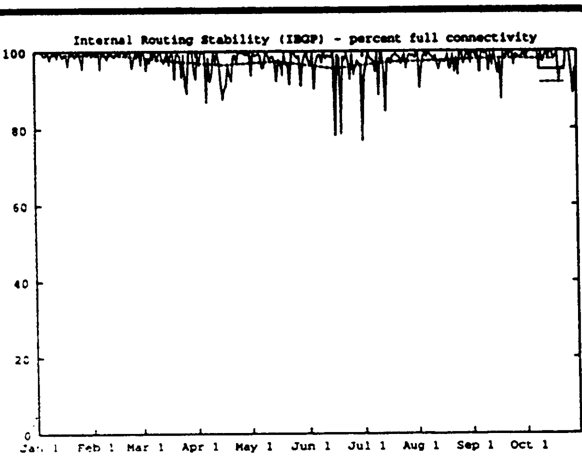

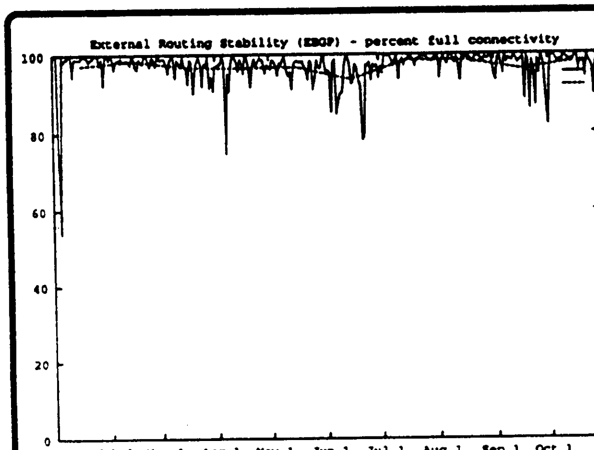

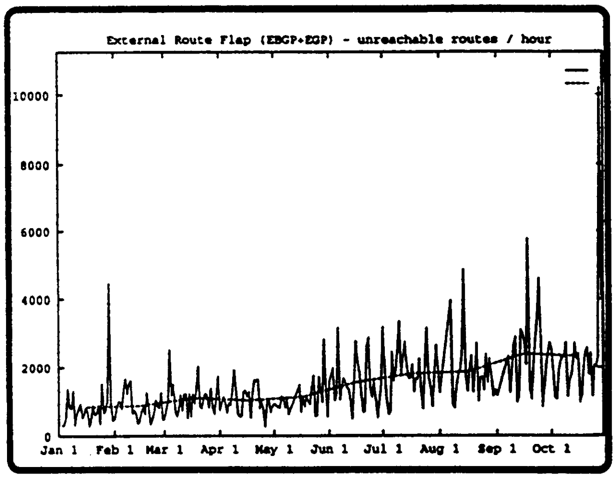

Phase 5 Upgrade Overview

- Full Bandwidth T3
  - New RS960 Adapter
  - New DSUs
- New CNSS Architecture
  - Internal FDDI Ring
  - T3 Bandwidth Manager
- Software Changes
  - Router System Software
  - Routing Protocol Software

ans

RS960 T3 and FDDI Upgrades

- Base Card Upgrade
  - 40KPPS per card @ 250 byte packet size
  - 5 interfaces per RS/6000 Router
  - Memory parity problem fixed
- New HSSI Daughter Card Supports 44.7Mbps
- FDDI interface will run at 30KPPS

ans

New CNSS Architecture

- Common FDDI Ring for CNSS Routers
  - Simplifies CNSS Router Maintenance
  - Lower hopcount
  - Increase aggregate throughput
- Introduce T3Plus BMX45 Bandwidth Manager
  - Increased fault tolerance
  - Increased operational flexibility
- ENSS-CNSS Links to use Larscom Access-T45

ans

New CNSS Architecture

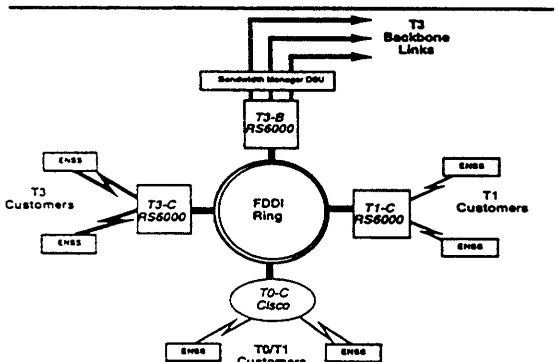

ans

ANSNet Phase-5 Deployment Schedule

Step	Date	ENSS	CNSS	Sites
1	11/06/93	E143	SEJ	Washington University and Seattle POP
		E142	DNJ	University of Utah and Denver POP
		E141	DNJ	University of Colorado and Denver POP
2	11/13/93	E128	HAY	Stanford University and San Fran POP
		E144	HAY	NASA Ames and San Fran POP
		E130	DNG	Argonne Nat'l Labs and Chicago POP
		—	SEJ	2nd visit for SEJ-HAY link
3	11/20/93	E135	DOH	San Diego SuperComp and LA POP
		E179	ABQ	Sandia Labs and Albuquerque POP
		E172	ABQ	Phillips Labs and Albuquerque POP
		E139	HSN	Rice University and Houston POP
		—	HAY	2nd visit for HAY-DOH link
4	12/04/93	E129	CRC	U of Illinois/UC and St. Louis POP
		E140	CRC	U of Lincoln/Nebraska and St. Louis POP
		E138	ATU	Georgia Tech and Atlanta POP
		—	GNJ	Greensboro POP (BMX only)
		—	DNJ	2nd visit for DNJ-CRC link
		—	DNG	2nd visit for DNG-CRC link
		—	HSN	2nd visit for HSN-CRC & HSN-ATU links
5	12/11/93	E131	NOR	U of Michigan and Cleveland POP
		E132	NOR	Pittsburgh SCC and Cleveland POP
		E137	NYU	Princeton University and New York POP
		—	DNG	3rd visit for DNG-NOR link
6	12/18/93	E133	HTF	Cornell University and Hartford POP
		E134	HTF	MIT and Hartford POP
		E136	WTN	U of Maryland and Wash D.C. POP
		E145	WTN	FIX-East and Wash D.C. POP
		E146	WTN	ARPA and Wash D.C. POP
		—	NOR	2nd visit for NOR-HTF link
		—	GNJ	2nd visit for GNJ-WTN link
		—	NYU	2nd visit for NYU-HTH & NYU-WTN links

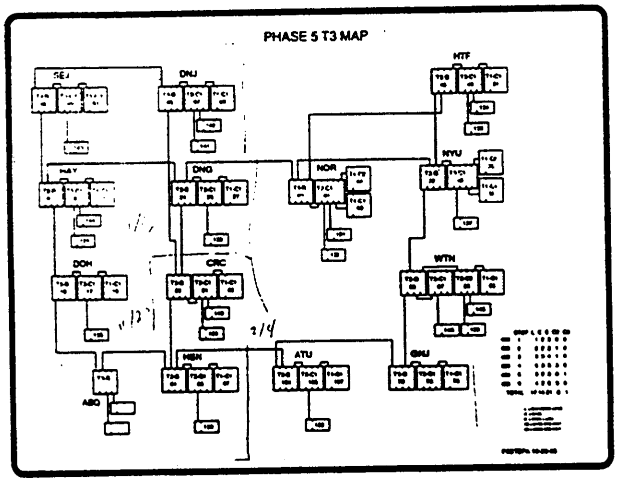

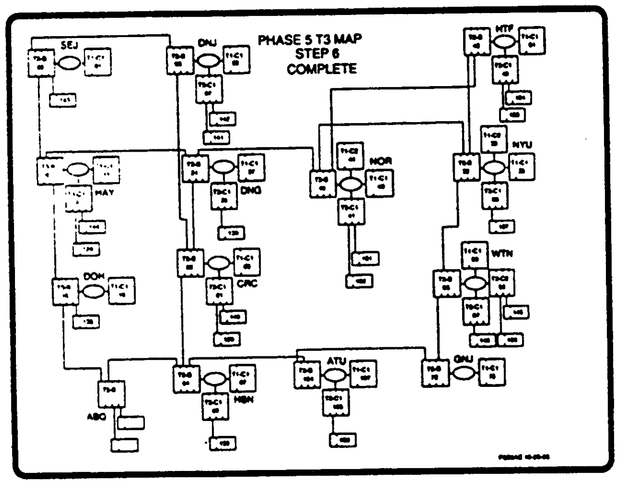

Software Changes

- Router System Software
  - Supports 25,000 Destination Networks
  - Kernel support for CIDR/Supernetting
  - Support for New RS960 adapter
- Initial GateD release (11/~~94~~93) supports
  - SLSP, IBGP,BGP3/4, EGP2 Protocols
  - SLSP, BGP MIBs
- Rcp.routed support for broadcast media

ans

Growth of Configured Networks

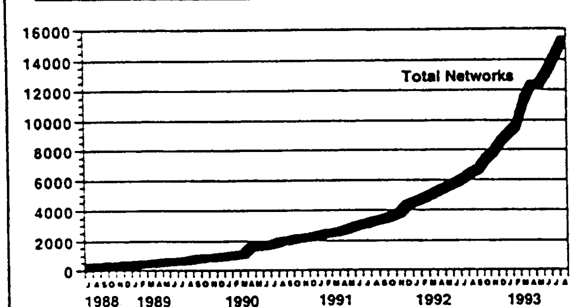

ans

ANSNET/NSFNET Router Table Growth - 9/93

Month	Max. Nets	Rate (%)
Jul-92	4,596	
Aug-92	4,866	5.87%
Sep-92	5,070	4.19%
Oct-92	5,432	7.14%
Nov-92	5,772	6.26%
Dec-92	6,239	8.09%
Jan-93	6,654	6.65%
Feb-93	7,037	5.76%
Mar-93	7,767	10.37%
Apr-93	8,239	6.08%
May-93	8,961	8.76%
Jun-93	9,534	6.39%
Jul-93	10,113	6.07%
Aug-93	10,879	7.57%

Average Monthly Growth Rate: [illegible]

Router Table Growth Forecast

Rate	5.0%	5.5%	6.0%	6.5%	7.0%	7.5%	8.0%	8.5%	9.0%	9.5%
Sep-93	11,423	11,477	11,532	11,586	11,641	[illegible]	11,749	11,804	11,858	11,913
Oct-93	11,994	12,108	12,224	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	13,044
Nov-93	12,594	12,773	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Dec-93	13,223	13,477	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Jan-94	13,885	14,218	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Feb-94	14,579	15,000	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Mar-93	15,308	15,825	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Apr-94	16,073	16,696	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
May-94	16,877	17,614	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Jun-94	17,721	18,583	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Jul-94	18,607	19,606	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Aug-94	19,537	20,683	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

4Q93 Software Upgrades

- T3 Router OSI Support
  - CLNP, ES-IS
- Migration from SLSP to IS-IS (GateD)
- Improved Route Download Performance
- T960 Precedence Queuing
- T960 PPP Support
- Full T960 MIB Support

ans

Chapter 4

IETF Protocol Presentations

4.1 A Report on the Temporary IPng Area

Presented by Scott Bradner/Harvard and Allison Mankin/NRL

Bio: *Scott Bradner has been involved in the design, operation and use of data networks at Harvard University since the early days of the ARPANET. He was involved in the design of the Harvard High-Speed Data Network (HSDN), the Longwood Medical Area network (LMAnet) and the New England Academic and Research network (NEARnet). He is currently chair of the technical committees of LMAnet, NEARnet and CoREN.*

Mr. Bradner is the Operational Requirements Area Director and IPng Area co-Director in the IESG and is a trustee of the Internet Society.

Mr. Bradner is a consultant at the Harvard Office of Information Technology, Network Service Division where he works on the design and development of network-based applications and manages the Network Device Test Lab.

Bio: *Allison Mankin is part of a group at Naval Research Laboratory (NRL) in Washington D.C. that is developing all the pieces needed for ATM network service in the ATDnet, a Federal interagency network coming on-line this spring. She has worked on Blanca as well. She has persistent interests in congestion control and measurement. In addition to being co-Director for the IPng Area, she is the Transport Services Area Director. She earned her Computer Science master's at Northeastern. At various times, she has pursued other topics, including mandarin Chinese and medieval Latin. She is on the Editorial Board of IEEE Network.*

A DIRECTION FOR IPNG

From The IPDecide BOF at the Amsterdam IETF (July 1993)

- "What is the basis for choosing the next generation IP (i.e., what are the technical requirements and decision criteria)."
- "With the advent of CIDR and new, more stringent address assignment policies, are we comfortable that we truly understand the level of urgency?"
- "Should the IETF or the marketplace make the final IPng decision".

1

A DIRECTION FOR IPNG

A consensus at the Amsterdam open IESG plenary (July 1993)

- "The IETF should take active steps toward a technical decision, rather than waiting for the "marketplace" to decide.
- "The IESG has the responsibility for developing an IPng recommendation for the Internet community."
- "The procedures of the recommendation-making process should be open and published well in advance by the IESG."
- "As a part of the process, the IPng WGs may be given new milestones and other guidance to aid the IESG."
- "There should be ample opportunity for community comment prior to final recommendation."

2

A DIRECTION FOR IPNG

- A temporary, ad hoc, "area" to deal specifically with IPng issues
- The charter for this new IESG area is to develop a recommendation for the "next IP". Submitted to the IESG and to the Internet community for review
- The IESG will issue a final IPng recommendation
- All IPng-related working groups will be moved into this new area
- An "Area Directorate" will be formed
- Completely open process
- The IPng area will be headed by
  - •Allison Mankin (NRL)
  - •Scott Bradner (Harvard)

3

CHARTER OF THE IPNG AREA

1. Develop an understanding of the available timeframe.
   - Review Internet growth metrics
   - Understand the time required for the fielding of a new IP
   - Review the expected impact of CIDR address aggregation
2. Make recommendation on the scope for IPng - Scaling issues only or advanced topics also?
3. Develop a clear set of decision criteria for IPng.
4. Based on all the above, make a recommendation on IPng to the IESG.

4

A Report on the Temporary IPng Area

co-chairs

Scott Bradner <sob@harvard.edu>
Allison Mankin <mankin@cmf.nrl.navy.mil>

IPng area structure

- co-chairs
- white papers
- working groups
- directorate
- review board

white papers

- provide simple description of proposals
- provide forum for community to express concerns
- provide communication channel to non-IETF community
- some white papers specifically solicited
  - IPng proposers
  - directorate members (biases & ideas)
  - people active on big-internet
  - IETF and non-IETF researchers in field
- open to all
- reviewed by directorate and review board, can suggest clarifications
  will not block publication
- submitted as informational RFC

white paper requirements

- maximum 10 pages!
- follow an outline to be developed by directorate
  can do only part you want to talk about

  possible contents: executive summary (max 1/2 page), config, admin & operation, extensibility (mobility, flows, policy), scaling, security, timescale, transition & coexistence, action items.
- number sections & paragraphs
  keep numbers, even if only including some sections

IPng process (current plan)

- since startup on 7 Sept. 1993
- policy on IPng standards
  some proposal documents already moved to Experimental
- directorate selection
- ALE WG preliminary charter

---Houston---

- form "requirements" WG
- form "transition, coexistence and testing" WG
- define & form review board

---post Houston---

- directorate defines white paper outline
- solicit white papers
  - about IPng proposals by proposers
  - from wider community about requirements

----Seattle---

- ALE WG presents time frames
- requirements WG presents a range of requirements and time frames at which technology could meet them.
- transition... WG presents status report

----post Seattle---

- ALE timeframe placed against requirements time frame
  Area produces ID detailing process time frames
- each proposal produces white paper describing how
  it meets requirements with their accompanying time frames.

----Toronto---

- area co-chairs produce ID of recommendation
- extended Last Call issued to request community input
- IESG reviews area recommendation in light of community input
- IESG approves recommendation or requests additional effort
- loop on this page if required (sure hope not)
- one protocol advanced to Proposed
- AS produced for selected IPng

area philosophy

- proposals to be made as good as they can be
- proposals to be evaluated on technical basis
- open process
- formal processes for soliciting community input
- includes testing phase
- can't be debugged in-place

proposals to be made as good as they can be

- proposals reviewed for clarity & completeness
- proposals reviewed for technical feasibility
- negative comments should include specific suggestions
- iterative process

proposals to be evaluated on technical basis

- politics not part of decision process
- (but real world is out there)
- owning not required, cloning ok
- review done in context of proposal

open process

- all documents placed in public archives, including www/gopher IPng server
- minutes taken of directorate meetings & made public
- directorate email list <ipng@cmf.nrl.navy.mil>
- directorate email archive published weekly
- open directorate meetings at IETF meetings
  also at Interop if enough interest

formal processes for soliciting community input

- white papers

review process
- directorate
- review board
- extended Last Call

testing process

- part of transition and coexistence WG effort
- strong view that IPng can't be debugged 'in-place'

IPng Working Groups

existing

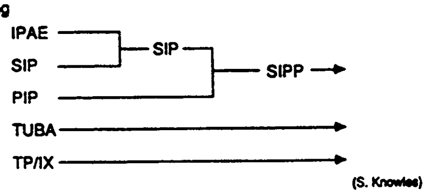

new
- Address Lifetime Expectations (ALE)
- requirements
- transition/coexistence & testing

ALE Working Group

- recognize overlap with existing CIDR efforts
  in many forums
- gather information from all sources
- replicate as little as possible

IPng directorate procedures

- teleconferences
- open meeting at IETF
  starting with the next one
- minutes kept and published
- directorate mailing list archived
- directorate asked to read & respond to big-internet
  (where appropriate)

IPng directorate - background

- technical aces
- must be able to represent self
- must be able to articulate needs of corporation and customers
- must be able to articulate needs of community
- asked that participation be checked with corporation
  "sign on" to process
  no implication that corporation agrees to result
- willingness to help juggle 'hot potato'

IPng directorate - selection process

- solicited names of appropriate people
- from inside and outside of IETF community
- final selection insured depth on current proposals
- solicited advice on list from many
- checked with suggested person :-)
- repeated process for final list

IPng directors & directorate expertise

- security
- routing
- international, national & regional network operation
- large users
- end system manufacturers
- Unix and non-Unix platforms
- router manufacturers
- theoretical research
- protocol architecture

IPng directorate

Steve Bellovin - AT&T	<smb@research.att.com>
Jim Bound - DEC	<bound@zk3.dec.com>
Ross Callon - Wellfleet	<rcallon@wellfleet.com>
Brian Carpenter - CERN	<brian.carpenter@cern.ch>
Dave Clark - MIT	<ddc@lcs.mit.edu >
John Curran - NEARnet	<curran@nic.near.net>
Steve Deering - Xerox	<deering@parc.xerox.com>
Dino Farinacci - Cisco	<dino@cisco.com>
Paul Francis - Bellcore	<francis@thumper.bellcore.com>
Eric Fleischmann - Boeing	<ericf@atc.boeing.com>
Daniel Karrenberg - RARE	<daniel@ripe.net>
Mark Knopper - MERIT	<mak@merit.edu>
Greg Minshall - Novell	<minshall@wc.novell.com>
Paul Mockapetris - ISI	<pvm@arpa.mil (pvm@isi.edu)>
Rob Ullmann - Lotus	<ariel@world.std.com>
Lixia Zhang - Xerox	<lixia@parc.xerox.com>

IPng review board

- facility to get review and advice from larger community
- specific areas of expertise
- industry input - but from a person
- agree to review documents & make specific suggestions

summary

- 1st it gotta work
- please help us debug process
- participate!

4.2 IPng: ALE

Presented by Frank Solensky/FTP Software

Bio: *Frank Solensky recently joined FTP Software as a project manager and has been an IETF participant since 1989. At the Vancouver IETF in 1990, he presented an analysis of the growth of the Internet that warned of the impending address assignment crisis. Frank has been reporting his projections of the size of the NSFNet Policy Routing Database to the big-internet mailing list.*

The presentation focussed on the the agenda of the IPv4 Address Lifetime Estimation BOF (ALE):

- Form a consensus on the aims and purposes of the proposed working group
- Resolve overlaps with other efforts
- Present and evaluate some of the independent analyses performed to date
- Define a working group charter

4.3 IPng: SIPP

Presented by Steve Deering/Xerox

Bio: *Steve Deering is a member of the research staff at the Xerox Palo Alto Research Center (PARC). He has been an active participant in the IETF and IRTF since 1984, and has served as chair of several IETF working groups. His current interests include addressing and routing for very large internets, with support for multicast, mobility, and multi-media services.*

SIPP Update

Steve Deering
Xerox Palo Alto Research Center
deering@parc.xerox.com

IETF 28
Houston, Texas
November 1, 1993

SIP + Pip = SIPP ("SIP Plus")

retains SIP packet format and basic functionality

advanced routing functionality of Pip accomplished by:

(1) optional Routing Header (identical to previous SIP Source Route header)

(2) concise specification of route-reversal rules

(3) enhancement of DNS to return address sequences

Pip vs. SIPP Source Routing

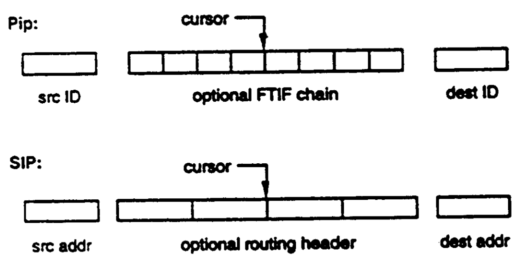

Pip and SIP Source Routing (cont.)

changes from Pip to SIPP:

- src & dest addresses can encode some or all location information => wider scope for packets that do not carry a Routing Header
- routing elements are fixed-length, 64 bits, rather than variable length
- different mechanism for advancing cursor => routing index always in same place in the basic header

Routing Header Functions

- service-provider selection
- policy routing
- host & domain mobility
- auto-configuration
- auto-renumbering
- ID / location separation (when necessary)
- extended addressing (if necessary)
- other?

To Learn More

WG sessions labeled "SIP" or "Pip" are all SIPP meetings

- Monday 1:30–3:30, 4:00–6:00 Austin A+C
- Thursday 1:30–3:30 Austin A+C

first session will have presentations on:

- SIP + Pip Merger
- SIPP Specification
- SIPP Addressing & Routing Overview
- SIPP Transition Plan (IPAE)

4.4 IPng: TPIX

Presented by Robert Ullman/Lotus

Bio: *Robert Ullmann started on VLSI CAD systems at Hamshire Engineering in 1972, and worked in 2D, CAD, and similar things for 10 years. He went "polygon-happy" in 1981, and moved into networking. He worked on X.25, e-mail and 7 different TCP/IP implementations in lots of different operating system environments. He has conducted basic research in data compression, resulting in the LZJU and PCU algorithms.*

The common architecture defined by CATNIP provides a compressed form of the existing network layer protocols IP, CLNP, and IPX. Each compression is defined so that the resulting network protocol data units are identical in format. The fixed part of the compressed format is 16 bytes in length, and may often be the only part transmitted on the subnetwork.

It is possible for a transport layer protocol (such as TCP) to operate properly with one end system using one network layer (e.g. IP version 4) and the other using a different network protocol such as CLNP. All of the existing transport layer protocols used on connectionless mode network services will operate over the common infrastructure without any modification of host software.

The architecture uses cache handles, carried in the fixed part of the network layer header, to provide both rapid identification of the next hop in high performance routing as well as abbreviation of the network header by permitting the addresses to be omitted when a valid cache handle is available.

The addressing is OSI NSAPs, with an AFI (to be assigned) for the Internet Protocol address space, and an ICD under AFI 47 (to be assigned) for the Novell IPX address space. There is no "address translation"; there are simple direct mappings into the NSAP format.

CATNIP

Robert Ullmann
Lotus Development Corporation
1 November 1993

Common Architecture for ...

- OSI CLNP
- Internet Protocol
- Novell IPX

Objectives

- Common fixed header
- Common addressing
- Incremental deployment

Transport Protocols

- OSI TP4 and CLTP
- Internet TCP and UDP
- Novell IPX, SPX, and the NCP

Information

- CATNIP draft
- /pub/tpix/draft-ietf-tpix-catnip-00.ps on world.std.com
- tpix-request@world.std.com
- TP/IX session 13:30 to 15:30 Wednesday
- TP/IX session 16:00 to 18:00 Wednesday

4.5 IPng: TUBA

Presented by Peter Ford/LANL

Bio: *Peter S. Ford is a member of the technical staff at Los Alamos National Laboratory (LANL). He co-chairs the TUBA working group with Mark Knopper at Merit. Mr. Ford currently works for the National Science Foundation on the NSFNET project and on internetworking projects at LANL.*

The TUBA Approach to the Evolution of the Internet

Presented by
Peter S. Ford
Los Alamos National Laboratory
peter@goshawk.lanl.gov

IETF
1 November, 1993
Houston, Texas, USA

TUBA: TCP&UDP with Bigger Addresses

- TUBA transitions IPv4 network layer to CLNP
- TUBA is not a wholesale transition to OSI, just the network layer.
- Use existing transports of Internet Suite (TCP & UDP)
- Use higher level protocols and applications of the Internet Suite. (Telnet, SMTP, FTP, X, Mosaic, etc.)
- CLNP and routing documents by anon FTP:

  merit.edu:pub/iso.

TUBA strengths

- Very Flexible Addressing – NSAPs
  - NSAP – length, type(AFI), value, demux (NSEL)
  - AFI is used to delegate authority for address of value and allocation
  - NSAP Routing and Addressing Plan (RFC 1237bis)
  - Can imbed addresses (IP, IPX) and routing information (ASes, multicast groups, etc.) in NSAPs.
- CLNP is field tested and deployed.

TUBA Transition

- Simple, Dual Stack.
- Inside Out Transition
- Encapsulation to cross uni-stack infrastructures.
- Translation (problematic in the general case)
- Transition plan will work for any IPNG.

TUBA Specific Implementation Experience

- NCSA Telnet (NIST), FTP, Finger for MS-DOS. Freeware.
- NSAP mods for BIND (Cisco)
- NCSA Traceroute (SURAnet).
- BSD Unix (UCB/LANL). Telnet, FTP, finger.
- Cisco: telnet, finger. Uses DNS for NSAPs, EON encapsulation.
- 3 Com: telnet, EON encapsulation.
- Tcpdump (INRIA)

Current Implementations (cont.)

- TBone – EON tunnels.
- SunOS (INRIA). TUBA is Unix User and Kernel mode. User mode also acts as a TUBA network address translator.
- IBM RS 6000.
- Bull.
- All those routers that are CLNP capable.

Map of CLNP infrastructure

CLNP for the Internet

- OSI XNTD work, profiling for Internet use.
- Can profile current CLNP for better alignment of fields and to use shorter addresses.
- Can evolve CLNP if necessary. Should be able to do so in a backward compatible fashion.
- Need spec for Internet multicast using CLNP

TUBA WG: Schedule and Agenda

- Monday, 13.30 & Wednesday, 13.30 (1st hour on Wed. joint with TP/IX)
- Technical topics:
  - OSI XTND (Katz)
  - Dynamic ES configuration (Katz)
  - Multicast (Marlow)
  - Mobility (CDPD?)
  - CATNIP/Hdr Compression (Callon and Ullman)
  - TUBA Transition ID (Piscitello & Ford)
  - DNS issues (Colella & Manning)

TUBA WG Administrative Work

- Create up-to-date document status.
- ISO document status wrt IDs and RFCs. (Chapin)
- Liaison with other groups (e.g. ISO/ISOC)
- TUBA implementation and infrastructure status

4.6 CIDRD Status

Presented by Tony Li/cisco Systems

Bio: *Tony Li has a Bachelor of Science degree in Mathematics from Harvey Mudd College and a Ph.D. in Computer Science from the University of Southern California. He currently wrangles bits for cisco Systems, Inc., specializing in high speed packet switching, interdomain routing protocols and really ugly bugs.*

Classless InterDomain Routing (CIDR) will improve the size of the routing tables in central routers and improve the ability to use address space more efficiently. Even though CIDR has not been deployed, it is possible to estimate the resulting amount of time left before the IP address space is completely assigned.

On CIDR and the Remaining Lifetime of the Internet

Tony Li
cisco Systems, Inc.

Overview

- What are the problems?
- What is CIDR?
- Data
- Other improvements
- Summary

The Problems

- Class B network number exhaustion
- Routing table explosion
- IP address space exhaustion

What is CIDR?

- Mechanism: BGP4, exterior routing based on prefixes. A prefix implies a block of addresses. Obsolete the class A, B, C system.
- Address allocation: Allocate prefixes hierarchically, based on topology. Sites get contiguous block of IP addresses. Block size is a power of two.
- Aggregation: Aggregate (combine prefixes) at topological boundaries.

How does CIDR solve the Problems?

- Aggregation reduces amount of routing information.
- Prefixes allow sites to use multiple class C's or portions of class A's or B's efficiently. Class B no longer required.
- Prefixes allow for more efficient address space utilization, since the allocated address space fits the organization. Utilization is now <<1%.

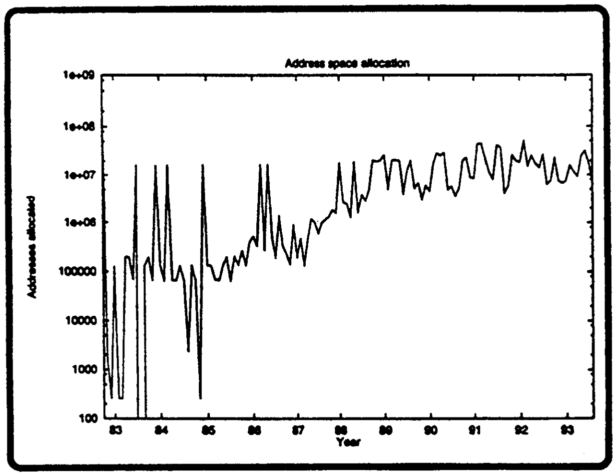
Address space allocation
Addresses allocated
1e+09
1e+08
1e+07
1e+06
100000
10000
1000
100
83 84 85 86 87 88 89 90 91 92 93
Year

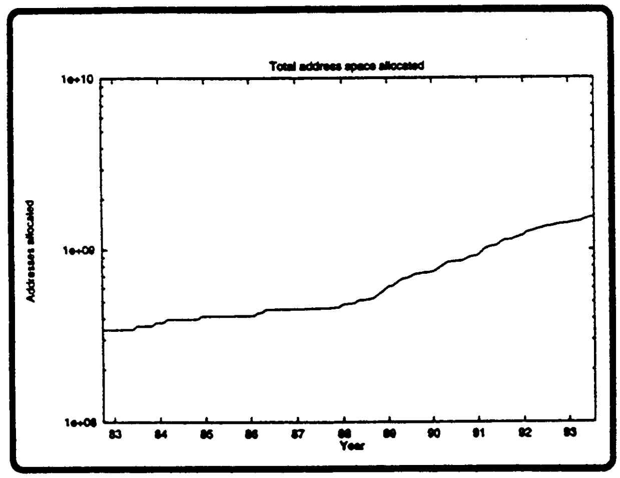
Total address space allocated
Addresses allocated
1e+10
1e+09
1e+08
83 84 85 86 87 88 89 90 91 92 93
Year

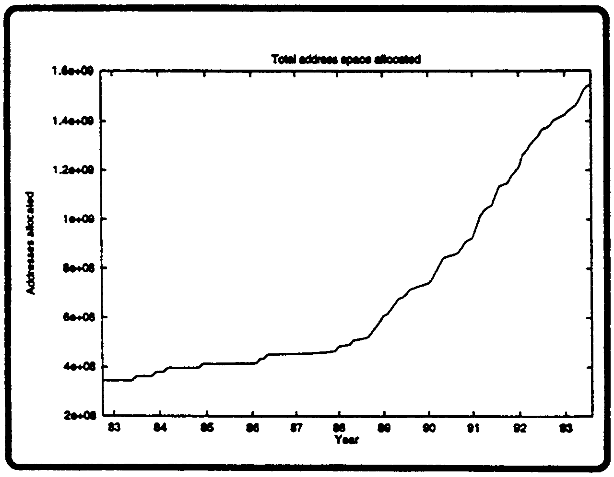
Total address space allocated
Addresses allocated
1.6e+09
1.4e+09
1.2e+09
1e+09
8e+08
6e+08
4e+08
2e+08
83 84 85 86 87 88 89 90 91 92 93
Year

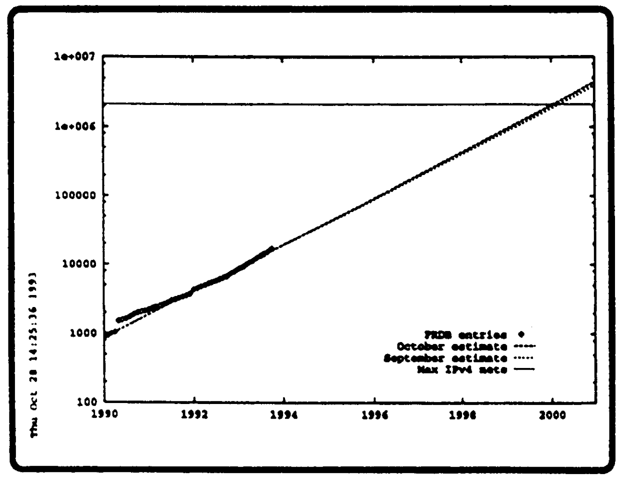
Thu Oct 28 14:25:36 1993
1e+007
1e+006
100000
10000
1000
100
1990 1992 1994 1996 1998 2000
PRDB entries
October estimate
September estimate
Max IPv4 nets

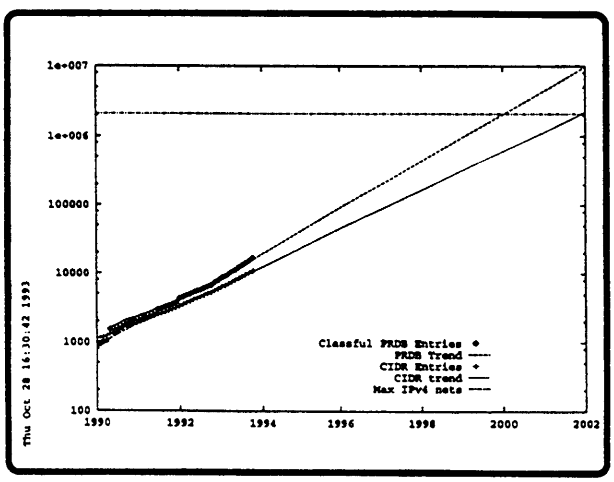
Thu Oct 28 16:30:42 1993
1e+007
1e+006
100000
10000
1000
100
1990 1992 1994 1996 1998 2000 2002
Classful PRDB Entries
PRDB Trend
CIDR Entries
CIDR trend
Max IPv4 nets

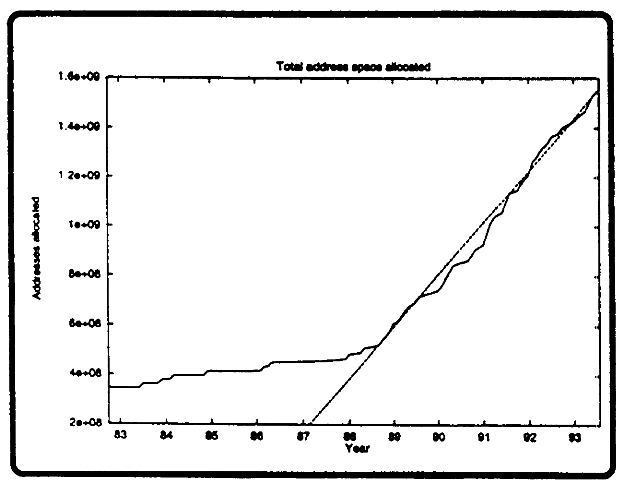
Total address space allocated
Addresses allocated
1.6e+09
1.4e+09
1.2e+09
1e+09
8e+08
6e+08
4e+08
2e+08
83 84 85 86 87 88 89 90 91 92 93
Year

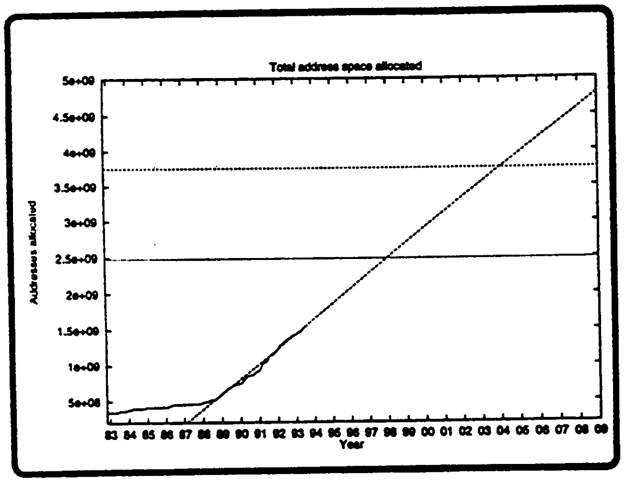
Total address space allocated
Addresses allocated
Year

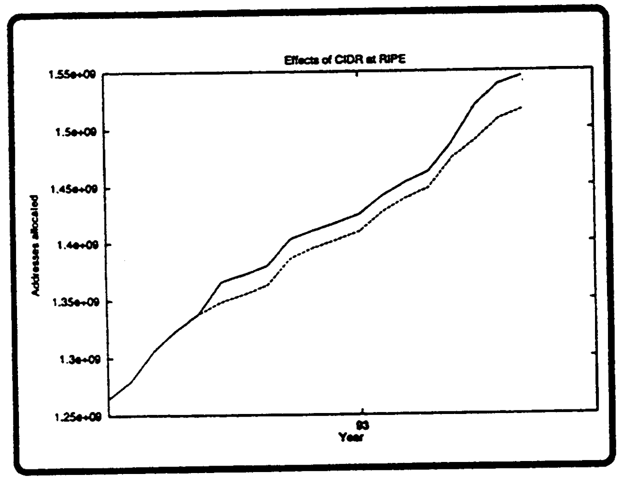
Effects of CIDR at RIPE
Addresses allocated
93
Year

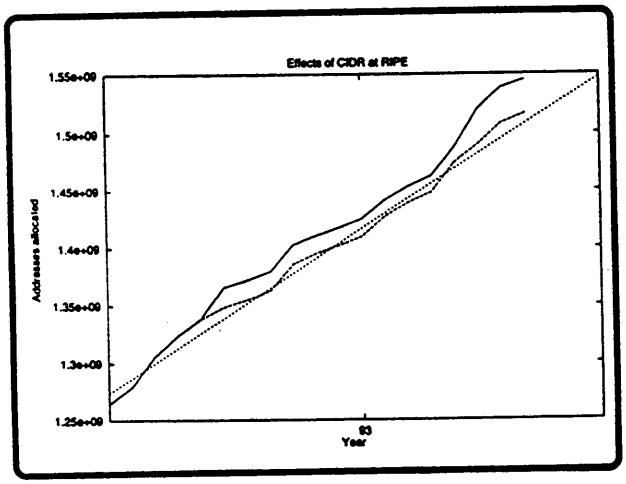
Effects of CIDR at RIPE
Addresses allocated
93
Year

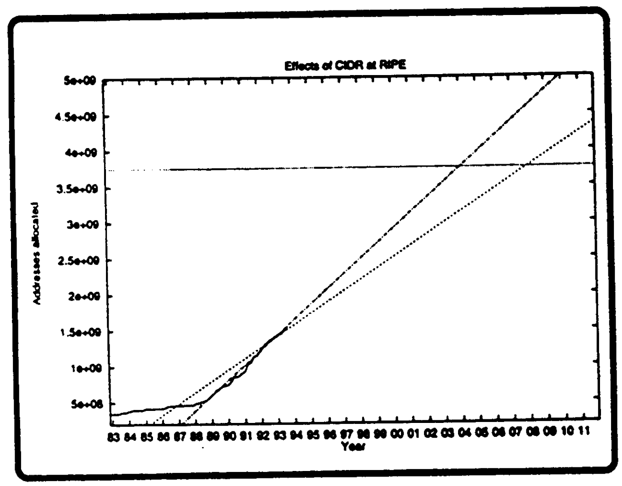
Effects of CIDR at RIPE
Addresses allocated
Year

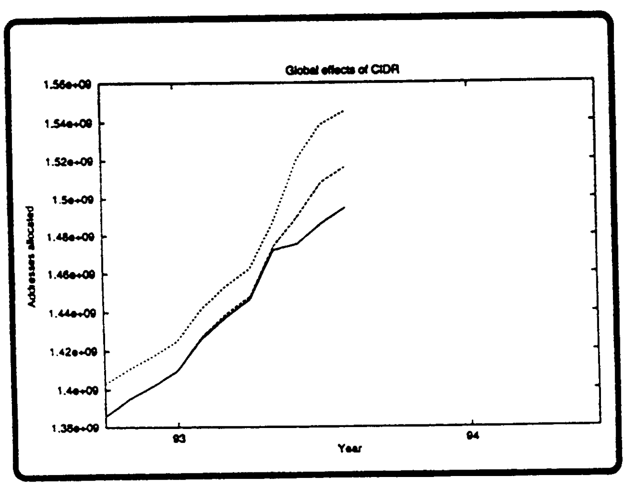
Global effects of CIDR
Addresses allocated
93
94
Year

Other improvements

- Renumbering to use more appropriate sized blocks helps improve utilization. Renumbering technology needs to be deployed.
- Renumbering on moves to preserve aggregation.
- Variable Length Subnet Masks help utilization as subnets can be correctly sized.
- Network Address Translators help utilization by hiding large private IP networks behind small globally-unique address spaces.

Summary

- Without CIDR, the address space is exhausted in ~1998.
- CIDR helps in two ways: allows remaining A's to be used efficiently, improves utilization.
- With CIDR, the address space lasts until ~2008.

Acknowlegements

- Thomas Williams, Colin Kelley - Gnuplot
- Mark Kosters, Internic
- Tony Bates, RIPE
- Paul Traina, cisco
- Frank Solensky, FTP
- Peter S. Ford, Los Alamos
- Hans-Werner Braun, SDSC
- Yakov Rekhter, IBM T.J. Watson
- and many others...

4.7 ST-II

Presented by Craig Partridge/Bolt Beranek and Newman

Bio: *Craig Partridge is a senior research scientist at Bolt Beranek and Newman. He is the editor-in-chief of IEEE Network Magazine, and a consulting assistant professor at Stanford. He is a former member of the IESG, a senior member of the IEEE, and attended his first IETF in 1986.*

ST-II is a protocol that embodies a particular approach to supporting real-time services. To understand ST-II, it is necessary to both understand how the protocol works and the philosophy behind it. This presentation briefly explains how ST-II operates and points out the limitations and strengths of the ST-II approach.

An Overview and Appraisal of ST-II

Craig Partridge

Bolt Beranek and Newman, Inc.

Goals of ST-II

- ST-II is both a protocol and a "philosophy."
- Want a internetwork-layer protocol that supports performance guarantees for one-to-many *streams* as requested by applications.
- ST-II designers believed that such a protocol should operate by reliably maintaining state information inside the network, with error recovery (if any) also the responsibility of the network.
- Further goal was to provide an efficient forwarding mechanism for delivering packets over streams.
- Note: discussion here based on the spec (implementations have various extensions).

How ST-II Works

- To setup a connection, a sender (*origin*) develops a list of all the parties it wants to talk to (*targets*), divides the list up according to the next hop router for the parties, and then sends one CONNECT message to each next hop.
- CONNECT message for a router contains only a list of all targets via that router and a FlowSpec describing the desired type of service.
- On receiving CONNECT, router concurrently starts negotiation of unique hop id (HID) with sender, and splits up list of targets according to next hops, modifies FlowSpec if needed, and sends CONNECT messages to next hops.

How ST-II Works (2)

- When a target gets a CONNECT message, it decides if it wants to accept the connection, and if so, sends an ACCEPT message with a possibly modified FlowSpec.
- ACCEPT messages are aggregated at routers and sent up towards origin. Modified FlowSpecs combined (inconsistencies to be sorted out by origin).
- NOTE: Streams are one-way data streams. Origin can send to targets but not reverse (unless if there's just one target).

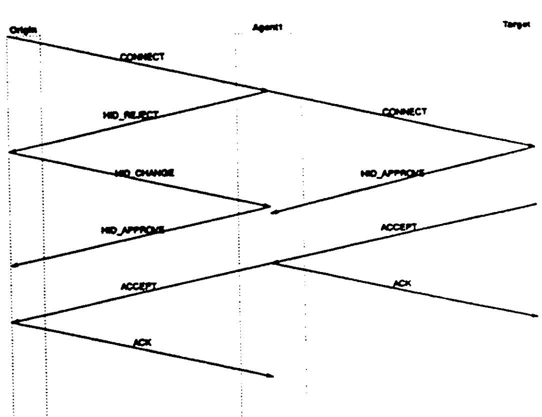

How ST-II Works (3)

- Once stream established, ST-II accepts responsibility for keeping stream up.
- Sender just transmits packets with tiny header (HID, checksum, len and a few control bits) and expects network to forward them with very high reliability.
- ST-II routers ping each other (intervals measured in milliseconds) to make sure they are up. If a router crashes, neighbors may try to heal streams.
- Origin can change stream to add new targets, delete targets, or change FlowSpec. Targets can only disconnect by themselves, otherwise must ask origin to make changes (such as changing QoS).

What ST-II Does and Doesn't Do

- ST-II is simply the protocol that puts state into the network, maintains the state, and defines a packet format that defines how routers access the state.
- Thus ST-II is only part of a larger set of services required for streams.
- In particular, ST-II expects others to specify:
  1. a queueing algorithm (e.g., Fair Queueing);
  2. any routing algorithm;
  3. how to communicate without a stream (must use IP until ST stream established).
- ST-II does specify a flowspec but probably should not.

Strengths of ST-II

- Paths remain fixed after setup, unless network fails.
- Its sender-oriented setup approach is well understood (ATM and X.25 use it too). and works well for small (2 or 3 party) conversations.
- Multiple interoperable subset implementations exist.
- ST-II has been shown to work in a number of modest-sized simulation exercises and demos.
- Matches well with telephony-like applications (small conference calls, etc.).
- It can be replaced without violence to other parts of the architecture (e.g., queueing, routing, and flowspec).

Problems with ST-II As Specified

- ST-II, however, is an imperfect realization of its goals.
- It encourages least-common-denominator homogeneity of service within a flow. If I want color and you want black-and-white video from a stream, I have to take black-and-white, or the origin has to open two streams.
- No state machine is defined in RFC, which can really stymie the implementor. (What to do when simultaneously processing requests to add a target, delete a target, and change the flow spec for the same flow?)
- Method for specifying a set of targets is clumsy (TargetList parameter).
- Too many choices for subset implementations.

Problems with ST-II As Specified (cont)

- Timestamp mechanism negotiation doesn't work.
- There probably isn't enough management information in the header for error recovery.
- Error handling is messy. Dozens of error codes, yet no way to say if an error is fatal. A pointer to where error occurred, but no rules about where pointer should point within a parameter that is bad.
- Flows can be (and) are identified in three different ways
- Four different ways to detect router failure (feeling safe is good but...)

Summary of ST-II

- The major virtues of ST-II are that paths don't change except after a failure and that it supports very small conferences well.
- A serious flaw in ST-II is that the sender manages the stream. Scales poorly in two ways. (1) Limits the number of receivers – to join or leave stream a receiver must contact the sender so sender can change stream. (2) Very difficult to share bandwidth among multiple senders (a Group parameter exists but not how to use it).
- Another problem with sender control is poor responsiveness (one RTT minimum to join or change a stream).
- A concern about ST-II is its huge state space and the complexity of the code it requires to maintain state in the network.

Chapter 5

Technical Presentations

5.1 Using E-mail in Europe: Opportunities and Challenges

Presented by Harald Alvestrand/UNINETT

Bio: *Harald Alvestrand graduated from the Norwegian Institute of Technology in 1984. He worked with X.400 implementations and other communications software for Norsk Data from 1984 to 1989. He worked with UNINETT's e-mail service from 1989 to 1993 and has been the service manager for UNINETT since 1993. Harald has been the head of the RARE working group on mail and messaging since 1992 and has been active in the IETF since 1991.*

The e-mail situation in Europe has several interesting characteristics:

- Multiple directions due to national priorities taking precedence over European priorities.

- Little cooperation or integration between the commercial and academic sectors, even less with the private market.

- Both X.400 and Internet Mail markets are growing very quickly.

This presentation tries to focus on the reasons behind these developments and identify some of the key players, like RARE, the European association of academic networks.

Using E-mail in Europe

Opportunities and challenges

Harald Tveit Alvestrand
<Harald.T.Alvestrand@uninett.no>

What is Europe?

Europe, like the Internet, does not exist.

The US: 200 million people. One language, one law.

Europe: 320 million people. Twenty languages. Thirty laws.

Power

Power in Europe is not centralized with a government.

Sources of power are:

National governments

- Individual agendas
- Looking for advantage in anything

Large corporations

- Usually affiliated with a "home base" country
- Cooperating on whatever seems profitable
- Separate agenda from the governments

International bodies

No more power than the members give them.

Hamstrung by fighting over control.

The EC is still a surprising success!

Languages

- Maltreating a language is an insult
- Most languages can *not* be written in ASCII
- More than 300 *latin* letters, and Greek and Cyrillic
- International organizations must give languages "equal" treatment
- \> Interpreter is safest job in all Europe!

Ø ➡ OE

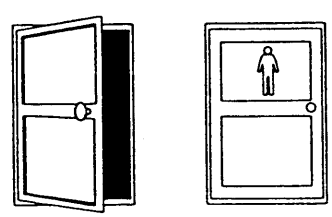

Dør ➡ Doer

How a door becomes a toilet.....

UNINETT

Deregulations

PTTs used to dominate all European countries
They still show it.

PTTs are proud

This means that they think they can dominate a market by stepping into it, and define the way things are done by mandating a standard.

PTTs are large

This means that they have a hard time adjusting

PTTs are powerless

Recent decisions by the European Commission have taken away almost all power held by the PTTs to influence how value added services are being offered by others.

However, they still have the financial and political resources to become major players in most communications-related fields.

UNINETT

E-mail sectors

Academic

- Coasting from network research to research networks
- Fuzzy financing, fuzzier control
- "Get what works" vs "Get what's right"

Commercial

- Internal operations have priority
- External waits for a Standard
- EDI is more important than personal mail

Private

- Triggered by MINITEL in Italy, France
- BBSes in other countries
- EUNet (UNIX users' group) active, but small
- Nobody knows what is going on

UNINETT

Commercial market

X.400 is The Standard

- People are very afraid of taking chances
- R&D arm often has Internet, but who cares?

EDI is The Application

- Personal E-mail is an interesting by-product
- Strategic Alliances are more important than Directories
- E-mail requires revising mission-critical SW
- Mailing lists? What is that?

ADMDs are The Service

- Chain of committment is required
- Someone to sue when things go wrong
- Quality of Service (even though it's bad)

Things are changing

- Internal E-mail widespread
  Linking out becomes common
- "Everyone" wants Internet access
- Fear of The Net decreasing with knowledge

UNINETT

Academic market

Spending: National networks is 10x European, site investment is 10x national networking....

National initiatives

- Differing policies: X.400/OSI vs mixed
- Growing need for Internet for US traffic
- Connection to commercial world chaotic at best, mainly X.400-based

(In Norway, approximately half the external IP traffic goes to the US!)

European coordination

- RARE -> COSINE -> DANTE:
  A viable X.400 infrastructure
- EARN, HEPNET/SPAN
- RIPE: Making IP work
- EBONE: Connect it NOW
- EuropaNET: Connect it with political correctness

UNINETT

THE FUTURE

X.400 will grow

Especially for EDI, it will grow explosively.
People who already have X.400 will use it for E-mail too.

The Internet will grow

On the contrary, Internet is now a commercial commodity.
Some of the sellers are PTT subsidiaries!

Standards will have Æ and α

Those that don't have them will die.

Standards will interwork

If they don't want to, we have to make them!

5.2 Federal Internetworking Requirements Panel (FIRP)

Presented by Richard desJardins/NASA

Bio: *Richard desJardins is a senior systems engineer with NASA Goddard Space Flight Center, and is currently responsible for network management within the Earth Observing System (EOS) Project. He has served as program manager of C3 systems with DARPA, senior systems engineer with CTA Incorporated, and design engineer for advanced control center systems within Goddard. Active in the OSI standardization community for many years, Richard was one of the authors of the Reference Model, has chaired the OSI technical committees for both ANSI and ISO, and currently teaches and writes about OSI/GOSIP/TCP/IP/Internet convergence.*

Richard holds a Master of Science degree in Computer Science from University of Maryland, and degrees in physics and mathematics from The Catholic University of America. He received the NASA Goddard Engineer-of-the-Year (Schneebaum) Award in 1977 for his pioneering contributions to control center design.

FIRP Charter

- Problem: Internetworking and convergence of network protocols, particularly Internet and OSI (and proprietary where appropriate)
- (Panel interprets this problem area broadly rather than narrowly)
- Study short term and long term issues
- Consider international relationships and implications
- Recommend Federal actions
- Seek views of public and private sectors

FIRP Scope

- Identify Federal requirements, evaluate fit of protocols
- Consider coexistence, interoperability and convergence options
- Consider cost to agencies of alternatives
- Consider process for Federal investments in research and development and infrastructure to best effect Federal requirements
- Consider specifications, maintenance, and testing for Federal Information Processing Standards (FIPS)
- Consider procurement and deployment scenarios
- Report recommendations, identify issues outside of scope

FIRP Establishment

- Endorsed by Federal Networking Council (FNC) and Federal Information Resource Management Policy Council (FIRMPoC)
- Office of Management and Budget (OMB) requested NIST to charter the panel
- Panel Members:
  - Diane Fountaine, DOD, Chair
  - Jason Cannon, Treasury
  - Mike Corrigan, GSA
  - Dick desJardins, NASA
  - Walt Houser, VA
  - Bill Hughes, NTIA
  - Milo Medin, NASA
  - Tom Rowlett, DOE
  - Steve Wolff, NSF

FIRP Schedule

- Meetings biweekly October-December 1993
- Initial draft report due by mid-January 1994
- Report will be widely disseminated, announced and freely available on Internet
- 30-day period of comments invited
- Panel will reconvene to consider comments

FIRP Work Structure

- Panel has divided its work into five areas:
  - Requirements Issues
  - International Interoperability and Trade Issues
  - Standards Process Issues
  - Technical and Technology Issues
  - Economic and Cost Issues

Richard desJardins' Personal Opinions About FIRP

- Panel has the right people, mix of views, responsibilities
- Panel members are not simpletons, they're smart about problem
- The problem is not simple, but a good outcome is achievable
- Panel members aim to come up with recommendations for Federal actions that will fix the problem, achieve the goal
- Goal is to achieve effective Federal interoperability, allowing agencies to meet their mission needs at low cost, while taking international implications into account

5.3 Intellectual Property Policy

Presented by Vinton Cerf/CNRI

Bio: *Vint Cerf is President of the Internet Society and Vice President of the Corporation for National Research Initiatives.*

In connection with revising RFC 1310 on the Internet Standards process, a section on intellectual property, including copyright and patent handling, has been prepared. This presentation will outline the proposed treatment of key aspects of standards documentation with regard to copyright and situations in which the technology of the standard(s) is known to be subject to patent claims. In broad outline, the Internet Society, under whose auspices the IETF conducts its standards work and through which the Request for Comments series are published, has attempted to assure that it will be able to authorize distribution and presentation of all standards and standards-related documents in all media (including on-line, paper, CD-ROM and so on) without limitations. At the same time, the Internet Society has tried to fashion its guidelines to require the minimum necessary from each document author (i.e., holder of copy rights by virtue of authorship, editing, or significant contribution) to assure that the Internet Society has clear rights to make such authorizations. Typically this requires only that rights holders offer no-cost, in-perpetuity licenses to the Internet Society for these purposes.

With regard to patented technology, the guidelines seek to assure that anyone making use of Internet Standards will have access to any required patent licenses on a non-discriminatory basis and at reasonable cost. The term "reasonable" is not further defined so as to avoid placing the Internet Society, IAB, IESG, IETF into an impossible position trying to negotiate actual license fees for any such technology. In any event, the IESG always has the option of withdrawing a standards recommendation in the event that patent holders fail to live up the the principles in the guidelines. In general, it is stated that Internet Standards-makers prefer to avoid the use of patented technology but that it is not ruled out, since to do so might prevent the Internet Community from making use of critically important technology in which commonality is essential for Internet use.

Technology which is considered a trade secret is probably, a priori, not a candidate for Internet Standards since most standards require widespread disclosure to be useful.

Internet Standards and Intellectual

Property Rights

November 1993

Vint Cerf

Page 1

Revision 2 of Internet Draft 9/93
"The Internet Standards Process"

1. Contributors grant to ISOC perpetual, non-exclusive, royalty-free, world-wide license to:

a) distribute (+ many other verbs)

b) Make derivative works and distribute

c) authorize others to do so

(ISOC can confirm to requestors rights to reproduce and distribute Internet Standards).

2. Contributors warrant they can grant rights above

3. ISOC has no confidentiality obligation w.r.t. contributions

Page 2

4. Contributors warrant their contributions don't violate other's rights

5. Material submissions become the property of ISOC

When appropriate, written confirmation of terms are to be obtained by IETF Exec Dir or designee (email is ok)

Principal contributors to Standards will be identified by WG chair or editor of document. Only named contributors asked to confirm terms.

All contributors asked to inform the IETF Secretariat of any known proprietary claims in any Standards.

If someone asserts proprietary rights, must inform editor or WG

Page 3

chair. Include person as named contributors (with confirmed terms) or edit to remove the contributions.

Page 4

Principle:

ISOC will not endorse or maintain Standards subject to copyrights, patents, patent applications or other rights without prior written assurances from rights owners that:

1. Any party can get rights to implement the standard and use it under specified, reasonable and non-discriminatory terms. [ISOC won't try to evaluate terms, though]

2. Party giving assurances must assert the right and power to do so.

Page 5

NOTICES ON STANDARDS-TRACK RFCs

1. Standards using proprietary elements will bear notice assuring non-discriminatory, reasonable licensing terms.

2. Note to all interested parties to inform IETF Exec Dir of any proprietary rights claims.

3. If applicable:
"As of <date>, no proprietary claims indications received"

4. Copyright Notice

"Copyright (c) ISOC <year>"

"Permission is granted to reproduce (etc) any Internet Standards material subject to ISOC Copyright, provided credit is

Page 6

given to the source. See IETF Executive Director for questions."

[Note - these notices and provisions allow ISOC to authorize reproduction of Internet Standards without restrictions.]

5. Disclaimer

ISOC disclaims all warranties, including merchantability or fitness for a particular purpose. Nor can ISOC guarantee that the document or the standard does not violate the rights of others.

[Possibility of unknown claims of proprietary rights.]

Page 7

First Live Test of Transfer:

SUN Microsystems NFS, etc.

Documents in draft for comment by end of November.

Page 8

5.4 Status of the ISOC/ISO Liaison MOU

Presented by Vinton Cerf/CNRI

The Internet Society, pursuant to the results of the July open IAB meeting and subsequent discussions in the IETF meetings in Amsterdam, has been continuing its discussions with the Organization for International Standardization. Depending on the state of affairs by the time of the IETF meeting in Houston in November, this presentation will either simply describe the status of the negotiations or present an outline of the draft memorandum of understanding which would be proposed to govern interactions between ISO committees, subcommittees and other standards groups and the various parts of the Internet Standards activities associated with the IAB and the IETF.

STATUS OF ITU LIAISON

Three or four iterations with ITU general secretary, Pekka Tarjanne.

Most recent provided documentation to ITU about ISOC and the IAB and IETF and the draft standards process.

Advised by ITU that request for waiver of approx $60K/year in fees would be taken up May 1994.

These fees are new and a result of the recent re-organization of CCITT as ITU-T. The Secretary-General does not have the authority to waive fees unilaterally.

A number of other organizations are in the same boat.

Page 1

STATUS OF ISO LIAISON

Reference:

Internet Draft "Liaison between Internet and other standardization agencies" June 1993 (Huitema/IAB)

OMNICOM Open Systems Communication Nov 1 1993

(talks about SC6 effort at Seoul meeting in October to set stage for category A liaison with Internet Society. Open: terms and conditions of a memorandum of understanding for joint work and common objectives.

Page 2

Huitema/IAB:

"...agreements...should not affect the IETF process nor [prejudice] in any way the results of the IPng discussions."

1. Recognition of IETF Internet Standards on an international basis improves government access to Internet-based products and services in many countries.

2. ISOC is also interested in establishing liaison relations for purposes of promoting expansion and use of Internet infrastructure world-wide.

Page 3

Brief History

1. Inter-regional Tele-communications Standards Conferences

20 Feb 1990 - Vicksburg, VA (Cerf and Rutkowski)

Sep 1991 - Nice, France (Chapin)

One result: ITU information server on the Internet to make some ITU documents available.

December 12, 1992 Letter to ISO ISOC requests Category A liaison between ISO and ISOC.

In parallel: SC6 proposes Category C arrangement

Page 4

March 1993 IETF meeting. Open IAB meeting discusses proposals and differences between Category A and C relations.

A: Peer to Peer (ISOC to ISO)

C: working group can submit contributions to ISO working groups.

[There is an implicit reciprocity since IETF has mechanisms for accepting contributions from many sources]

Action on Category C proposal deferred subject to Category A resolution.

October 1993 Seoul (SC6 Plenary)

Chapin, Houldsworth et al involved in discussions leading to recommendations from SC6 to JTC

Page 5

1 (the ISO/IEC secretariat) to pursue a Memorandum of Understanding under Category A between ISO/SC6 and ISOC.

Reference: SC 6 N 8420

"Statement of Expected Benefits Regarding Liaison between Internet Society and ISO/IEC JTC1/SC6"

Proposed MOU to include:

1) agreed benefits and objectives

2) procedures for protocol change management

Page 6

SC6 View of Goals and Benefits

-Simplify deployment and operation of global comm infra-structure

-Facilitate growth of infra-structure based on open systems standards

-Reduce complexity and costs by reducing alternatives

-Use scarce expert networking resources in devising protocols

-Promote use of ISO/IEC standards

[Note - these do not represent inputs from IAB, IETF and are not proposed to be taken up without adequate review, discussion and adaptation by IETF and IAB - VGC]

Page 7

JTC1 rules require some statement of benefits from establishing liaisons and above is the SC6 perspective.

Proposed next step is discussion of principles which should govern such agreements from IETF and IAB perspective (cf: Huitema, November 3 IETF BOF).

A thought: it may turn out to be more productive to engage in topic-specific discussions with appropriate ISO Technical and Sub-Committees and then generalize from there, if it seems worthwhile.

If joint discussions of a technical nature are conducted, agreement as to who may do what with the results of the discussion are probably in order.

Page 8

Chapter 6

Workshop for Working Group Chairs

6.1 Workshop for Working Group Chairs

Presented by David Crocker/Silicon Graphics

Attendees

Steve Alexander	stevea@lachman.com
Harald Alvestrand	Harald.T.Alvestrand@uninett.no
Fred Baker	fbaker@acc.com
David Crocker	dcrocker@mordor.stanford.edu
Steve DeJarnett	steve@ibmpa.awdpa.ibm.com
Luca Delgrossi	luca@ibmpa.awdpa.ibm.com
Urs Eppenberger	eppenberger@switch.ch
Michael Erlinger	mike@jarthur.claremont.edu
Jim Fullton	fullton@cnidr.org
Kevin Gamiel	kgamiel@cnidr.org
Joel Gyllenskog	jgyllens@hpdmd48.boi.hp.com
Joel Halpern	jmh@network.com
Marko Kaittola	Marko.Kaittola@funet.fi
John Kunze	jak@violet.berkeley.edu
Paul Lambert	paul_lambert@email.mot.com
Edward Levinson	elevinson@accurate.com
Bill Manning	bmanning@rice.edu
Drew Perkins	ddp@fore.com
Patricia Smith	psmith@merit.edu
Cathy Wittbrodt	cjw@barrnet.net

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)
- IETF WORKING GROUP GUIDELINES AND PROCEDURES (IN PROGRESS)
- ["EVOLVING THE SYSTEM" IN INTERNET SYSTEM HANDBOOK, LYNCH & ROSE, EDS.]
- ["MAKING STANDARDS THE IETF WAY" IN ACM STANDARDSVIEW, SUMMER 1993; REPRINTED IN CONNEXIONS, AUGUST 1993.]

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

D. Crocker, Workshop for Working Group Chairs 7

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)

D. Crocker, Workshop for Working Group Chairs 8

IETF acceptance criteria

Competence:	technically sound
Constituency:	providers & users
Coherence:	clear writing
Consensus:	rough but clear

D. Crocker, Workshop for Working Group Chairs 9

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

D. Crocker, Workshop for Working Group Chairs 10

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

D. Crocker, Workshop for Working Group Chairs 11

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

D. Crocker, Workshop for Working Group Chairs 12

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

D. Crocker, Workshop for Working Group Chairs 13

II. Developmental steps

(7.5) IAB CONFLICT RESOLUTION
- If formal challenge not resolved by IESG

8. RFC PUBLICATION
   - RFC editor has publication criteria

D. Crocker, Workshop for Working Group Chairs 14

III. The inner working group

- THE LIVES OF A CHAIR
- GROUP STYLE
- GROUP ROLES
- DEVELOPMENTAL PHASES (PROBLEM SOLVING 101)
- VENUES
- DEBATE
- CONFLICT MANAGEMENT

D. Crocker, Workshop for Working Group Chairs 15

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

D. Crocker, Workshop for Working Group Chairs 16

Working group style

- FREE-FLOWING
  - Cohesive group
  - Clear purpose
- TIGHTLY-MANAGED
  - Complex topic
  - Group diversity
  - Major differences in philosophy

D. Crocker, Workshop for Working Group Chairs 17

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

D. Crocker, Workshop for Working Group Chairs 18

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.

D. Crocker, Workshop for Working Group Chairs 25

3. Chain of appeal

- WG CHAIR
- AREA DIRECTOR
- IESG (PLENARY)
- IAB

D. Crocker, Workshop for Working Group Chairs 26

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 opinion, but move on
- Ask questions
- Make it happen!

D. Crocker, Workshop for Working Group Chairs 27

Appendix A

Attendees

Andy Adams
Merit Network, Inc.
Phone: 313-763-4897
Fax: 313-747-3745
ala@merit.edu

Vikas Aggarwal
JvNCnet/GES
Phone: 609-897-7307
Fax: 609-897-7310
vikas@jvnc.net

Masuma Ahmed
Bell Communications Research
Phone: 908-758-2515
Fax: 908-758-4192
mxa@mail.bellcore.com

Kannan Alagappan
Digital Equipment Corporation
Phone: 508-486-5953
Fax: 508-486-5279
kannan@dsmail.enet.dec.com

Kenneth Albanese
Sprint
Phone: 703-904-2361
Fax: 703-904-2192
albanese@icp.net

Edward Alcoff
Network Application Technology
Phone: 408-370-4205
Fax: 408-370-4222
oldera@nat.com

Steve Alexander
Lachman Technology Inc.
Phone: 708-505-9555
Fax: 708-505-9574
stevea@lachman.com

Garrett Alexander
Department of Defense
Phone: 301-688-0849
gda@tycho.ncsc.mil

Nick Alfano
MPR Teltech, Ltd.
Phone: 604-293-5342
Fax: 604-293-5787
alfano@mpr.ca

James Allard
Microsoft Corporation
Phone: 206-882-8080
Fax: 206-936-7329
jallard@microsoft.com

Michael Allen
IBM Corporation
Phone: 919-254-4442
Fax: 919-254-5483
moallen@ralvm6.vnet.ibm.com

Anthony Alles
Cisco Systems
Phone: 415-428-5424
Fax: 415-688-7666
aalles@cisco.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

Harald Alvestrand
UNINETT A/S
Phone: +47 7 59 70 94
Fax: +47 7 59 64 50
Harald.T.Alvestrand@uninett.no

Farhad Anklesaria
University of Minnesota
Phone: 612-624-9360
Fax: 612-625-6817
fxa@boombox.micro.umn.edu

Vadim Antonov
US Sprint
Phone: 703-904-2013
Fax: 703-904-2292
avg@icm1.icp.net

Susie Armstrong
Mentat Inc.
Phone: 310-208-2650
Fax: 310-208-3724
susie@mentat.com

Jules Aronson
National Library of Medicine
Phone: 301-496-9300
Fax: 301-496-0673
aronson@nlm.nih.gov

Randall Atkinson
Naval Research Laboratory
atkinson@itd.nrl.navy.mil

Robert Austein
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sra@epilogue.com

Anders Baardsgaard
University of Tromso
Phone: +47 834 4105
Fax: +47 834 4100
anders@cc.uit.no

Alireza Bahreman
Bell Communications Research
Phone: 908-699-7398
Fax: 908-336-2943
bahreman@bellcore.com

Chase Bailey
Efficient Networks
Phone: 214-991-3884
Fax: 214-991-3887
chase@efficient.com

Fred Baker
Advanced Computer Communications
Phone: 805-685-4455
Fax: 805-685-4465
fbaker@acc.com

Stan Barber
Baylor College of Medicine
Phone: 713-798-6042
Fax: 713-798-6675
sob@bcm.tmc.edu

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

Stephen Batsell
U.S. Naval Research Laboratory
Phone: 202-767-3834
Fax: 202-767-1191
batsell@itd.nrl.navy.mil

Charles Bazaar
Emulex Corporation
Phone: 714-668-5416
Fax: 714-668-6819
bazaar@emulex.com

Jordan Becker
Advanced Network and
Services, Inc.
Phone: 914-789-5308
Fax: 914-789-5310
becker@ans.net

Jim Beers
Cornell University
Phone: 607-255-2627
Fax: 607-255-9086
Jim.Beers@cornell.edu

Nutan Behki
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3681
nebhki@newbridge.com

Steven Bellovin
AT&T Bell Laboratories
Phone: 908-582-5886
Fax: 908-582-3063
smb@research.att.com

Tom Benkart
ACC Systems
Phone: 410-290-8100
Fax: 410-290-8106
teb@acc.com

Bart Berger
3Com Corporation
Phone: 408-764-6065
Fax: 408-764-5002
bart_berger@3com.com

Lou Berger
Bolt Beranek and Newman
Phone: 703-284-4651
Fax: 703-284-2766
lberger@bbn.com

Mark Beyer
Tandem Computers Inc.
Phone: 408-285-8660
Fax: 408-285-8369
beyer_mark@tandem.com

Ram Bhide
Network Application Technology
Phone: 408-370-4358
Fax: 408-370-4222
ram@nat.com

Andy Bierman
SynOptics Communications, Inc.
Phone: 408-764-1673
Fax: 408-764-1899
abierman@synoptics.com

Richard Binder
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
rbinder@cnri.reston.va.us

Steven Blair
DELL Computer Corporation
Phone: 512-728-2563
Fax: 512-728-3653
sblair@us.dell.com

Larry Blunk
Merit Network, Inc.
Phone: 313-763-6056
Fax: 313-747-3745
ljb@merit.edu

Jon Boone
Pittsburgh Supercomputing
Center
Phone: 412-268-6959
Fax: 412-268-5832
boone@psc.edu

David Borman
Cray Research, Inc.
Phone: 612-683-5571
Fax: 612-683-5599
dab@cray.com

Sepideh Boroumand
B&B Corporation
Phone: 703-759-0499
Fax: 703-759-0785
sepi@aol.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: 510-665-8282
Fax: 510-422-1482
bostwick@es.net

Luc Boulianne
McGill University
Phone: 514-398-6698
Fax: 514-398-3883
lucb@cs.mcgill.ca

Jim Bound
Digital Equipment Corporation
Phone: 603-881-0400
Fax: 603-881-0120
bound@zk3.dec.com

Rich Bowen
IBM Corporation
Phone: 919-543-9851
rkb@ralvm11.vnet.ibm.com

David Boyes
Rice University
Phone: 713-527-4852
dboyes@rice.edu

Robert Braden
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
braden@isi.edu

Scott Bradner
Harvard University
Phone: 617-495-3864
Fax: 617-495-0914
sob@harvard.edu

Arthurine Breckenridge
Sandia National Laboratories
Phone: 505-845-8390
Fax: 505-844-2067
arbreck@sandia.gov

Dan Brendes
Memorex Telex
Phone: 919-250-6609
Fax: 919-250-6475
brendes@raleng.mtc.com

Tom Brennan
Exide Electronics
Phone: 919-872-3020
Fax: 919-878-2385
brennan@exide.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

Ronald Broersma
NRaD
Phone: 619-553-2293
Fax: 619-553-6811
ron@nosc.mil

Al Broscius
Bell Communications Research
Phone: 201-829-4781
Fax: 201-829-2504
broscius@bellcore.com

Caralyn Brown
Wellfleet Communications, Inc.
Phone: 508-436-3835
Fax: 508-670-8760
cbrown@wellfleet.com

Tracy Brown
Bell Communications Research
Phone: 908-758-2107
Fax: 908-758-4177
tacox@mail.bellcore.com

Theodore Brunner
Bell Communications Research
Phone: 201-829-4578
Fax: 201-829-2504
tob@thumper.bellcore.com

Steve Buchko
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
stevebu@newbridge.com

Jeffrey Burgan
NASA Ames Research Center
Phone: 415-604-5705
Fax: 415-604-0063
jeff@nsipo.nasa.gov

Randy Bush
RGnet
Phone: 503-297-8820
randy@psg.com

Thomas Butts
OpenConnect Systems
Phone: 214-888-0441
Fax: 214-888-0686
tom@oc.com

Glen Cairns
MPR Teltech, Ltd.
Phone: 604-293-6020
Fax: 604-293-5787
cairns@mprgate.mpr.ca

Susan Calcari
InterNIC Information Services
Phone: 619-455-3904
Fax: 619-455-4640
calcaris@internic.net

Ross Callon
Wellfleet Communications, Inc.
Phone: 508-436-3936
Fax: 508-670-8760
rcallon@wellfleet.com

Peter Cameron
Xylogics International, Ltd.
Phone: +44 908 222112
Fax: +44 908 222115
cameron@xylint.co.uk

Brian Capouch
Phone: 219-253-8552
brian@saintjoe.edu

C. Allan Cargille
University of Wisconsin
Phone: 608-262-5084
Fax: 608-262-9777
allan.cargille@cs.wisc.edu

Ken Carlberg
SAIC - TIEO
Phone: 703-715-0281
Fax: 703-749-5997
Carlberg@cseic.saic.com

Hallie Carlson
NASA Science Internet
Phone: 513-320-1328
hallie@nsipo.arc.nasa.gov

Michael Carney
SunSelect
Phone: 508-442-0469
Fax: 508-250-5068
mwc@sun.com

David Carr
Gandalf Data Limited
Phone: 613-723-6500
Fax: 613-226-1717
dcarr@gandalf.ca

Lida Carrier
Apple Computer, Inc.
Phone: 408-974-8325
Fax: 408-974-0892
lida@apple.com

Jeff Case
University of Tennessee
Phone: 615-573-5931
Fax: 615-573-9197
case@cs.utk.edu

Bill Cash
Compaq Computer Corporation
Phone: 713-374-9184
Fax: 713-374-8375
cash@bangote.compaq.com

Stephen Casner
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
casner@isi.edu

Isidro Castineyra
Bolt Beranek and Newman
Phone: 617-873-6233
Fax: 617-873-7776
isidro@bbn.com

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@osi.ncsl.nist.gov

John Chang
US West Advanced Technologies
Phone: 303-541-6255
Fax: 303-541-6250
jrc@uswest.com

Ted Chang
Wiltel
Phone: 713-364-5466
tchang@atg.wiltel.com

A. Lyman Chapin
Bolt Beranek and Newman
Phone: 617-873-3133
Fax: 617-873-6091
lyman@bbn.com

Cheng Chen
AccessWorks Communications
Phone: 908-888-4570
chen@accessworks.com

Enke Chen
Merit Network, Inc.
Phone: 313-936-3335
enke@merit.edu

Ping Chen
Apple Computer, Inc.
Phone: 408-862-3599
Fax: 408-446-2897
ping@ping2.aux.apple.com

Jia-bing Cheng
IBM Corporation
Phone: 919-254-4434
Fax: 919-254-5483
cheng@ralvm6.vnet.ibm.com

J. Noel Chiappa
Phone: 804-898-8183
Fax: 804-898-7663
jnc@lcs.mit.edu

Bilal Chinoy
San Diego Supercomputer Center
Phone: 619-534-8305
Fax: 619-534-5113
bac@sdsc.edu

Chris Chiotasso
LightStream Corporation
Phone: 508-262-1094
Fax: 508-262-1111
chris@lightstream.com

Frank Ciotti
Telxon Corporation
Phone: 216-867-3700
Fax: 216-873-2057
frankc@telxon.com

George Clapp
Ameritech Advanced Data Services
Phone: 703-765-2360
Fax: 703-248-3996
clapp@ameris.ameritech.com

Alan Clark
Hayes Microcomputer Products
Phone: 404-840-9200
Fax: 404-447-0178
aclark@hayes.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: 617-253-6003
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: 415-688-4575
wclark@cisco.com

Alan Clegg
MCNC - CONCERT Network
Phone: 919-248-1120
Fax: 919-248-1405
abc@concert.net

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

Michael Collins
Lawrence Livermore National Laboratory
Phone: 510-422-4018
Fax: 510-422-1482
collins@es.net

Rob Coltun
RainbowBridge Communications
Phone: 301-340-9416
rcoltun@ni.umd.edu

James Conklin
CREN/EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
jbc@bitnic.educom.edu

Michael Conn
MCI
Phone: 214-918-5710
Fax: 214-918-6052
4387451@mcimail.com

David Conrad
Internet Initiative Japan
Phone: +81 3 3580 3781
Fax: +81 3 3580 3782
davidc@iij.ad.jp

Ann Cooper
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
cooper@isi.edu

Blair Copland
University of North Texas
copland@unt.edu

Thomas Coradetti
DigiBoard, Inc.
Phone: 612-943-9020
Fax: 612-943-0803
tomc@digibd.com

Stephen Coya
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
scoya@cnri.reston.va.us

Matt Crawford
Fermi National Accelerator Laboratory
Phone: 708-840-3461
Fax: 708-840-8208
crawdad@fncent.fnal.gov

Mark Crispin
University of Washington
Phone: 206-543-5762
Fax: 206-685-4045
mrc@cac.washington.edu

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

John Curran
Bolt Beranek and Newman NEARnet
Phone: 617-873-4398
Fax: 617-873-5620
jcurran@nic.near.net

James Davin
Bell Communications Research
Phone: 201-829-4798
Fax: 201-829-2504
davin@thumper.bellcore.com

Michael Davis
Penril Datability Networks
Phone: 508-653-5313
Fax: 508-653-6415
mike@dss.com

Michael Davison
FiberCom, Inc.
Phone: 703-342-6700
Fax: 703-342-5961
davison@fibercom.com

Chuck de Sostoa
Hewlett-Packard
Phone: 408-447-2845
Fax: 408-447-3660
chuckd@cup.hp.com

Osmund deSouza
AT&T Bell Laboratories
Phone: 908-949-1393
Fax: 908-949-4673
osmund.desouza@att.com

Stephen 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

Luca Delgrossi
IBM European Networking Center
Phone: +49 6221 594330
Fax: +49 6221 593400
luca@ibmpa.awdpa.ibm.com

Richard desJardins
NASA
Phone: 301-286-8554
Fax: 301-286-1744
desjardi@eos.nasa.gov

Peter Deutsch
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8134
peterd@bunyip.com

Taso Devetzis
Bell Communications Research
Phone: 201-829-4764
Fax: 201-829-2504
devetzis@bellcore.com

Manuel Diaz
DAVID Systems, Inc.
Phone: 408-541-6865
Fax: 408-541-6980
diaz@davidsys.com

Jonathan Didner
Compaq Computer Corporation
Phone: 713-378-7459
Fax: 713-374-8375
jonb@bangate.compaq.com

Thomas Dimitri
Microsoft Corporation
Phone: 206-936-4187
tommyd@microsoft.com

Kurt Dobbins
Cabletron Systems, Inc.
Phone: 603-337-5018
Fax: 603-424-8121
dobbins@ctron.com

Waychi Doo
National Semiconductor Corporation
Phone: 408-721-6887
Fax: 408-721-4115
wcd@berlioz.nsc.com

Avri Doria
Locus Computing Corporation
Phone: 617-229-4980
Fax: 617-229-2969
avri@locus.com

Christopher Dorsey
Lawrence Livermore National Laboratory
Phone: 510-422-4474
Fax: 510-422-0435
dorsey@es.net

Robert Downs
Combinet
Phone: 408-522-9020
Fax: 408-732-5479
bdowns@combinet.com

Roger Draper
Liebert Corporation
Phone: 714-858-5099
rdraper@cerf.net

Ralph Droms
Bucknell University
Phone: 717-524-1801
Fax: 717-524-3760
droms@bucknell.edu

Walt Drummond
Rutgers University
Phone: 908-932-0882
Fax: 908-932-2968
drummond@noc.rutgers.edu

David Dubois
Pacer Software, Inc.
Phone: 508-898-3300
Fax: 508-366-1356
dad@pacersoft.com

Pierre Dupont
Motorola Wireless Data Group
Phone: 604-241-6243
Fax: 604-241-6243
dupont@mdd.comm.mot.com

Tom Easterday
CICNet, Inc.
Phone: 313-998-6285
Fax: 313-998-6105
tom@cic.net

Donald Eastlake
Digital Equipment Corporation
Phone: 508-486-6577
Fax: 508-486-7417
dee@skidrow.lkg.dec.com

Havard Eidnes
Uninett
Phone: +47 73 59 4468
Fax: +47 73 59 1700
havard.eidnes@runit.sintef.no

Ed Ellesson
IBM Corporation
Phone: 919-254-4115
Fax: 919-254-5410
ellesson@vnet.ibm.com

Bill Elliot
ONEAC Corporation
Phone: 708-816-6000
Fax: 708-680-5124
70760.3022@compuserve.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@seka.reston.ans.net

Urs Eppenberger
SWITCH
Phone: +41 1 268 1550
Fax: +41 1 268 1568
eppenberger@switch.ch

Dario Ercole
CSELT
Phone: +39 11 228 5051
Fax: +39 11 228 5069
Dario.Ercole@cselt.stet.it

Michael Erlinger
Harvey Mudd College
Phone: 909-693-2617
Fax: 909-621-8465
mike@jarthur.claremont.edu

Julio Escobar
Bolt Beranek and Newman
Phone: 617-873-4579
Fax: 617-873-3776
jescobar@bbn.com

Deborah Estrin
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
estrin@usc.edu

Erik Fair
Apple Computer, Inc.
Phone: 408-974-1779
Fax: 408-974-3103
fair@apple.com

Roger Fajman
National Institutes of Health
Phone: 301-402-4265
Fax: 301-480-6241
raf@cu.nih.gov

Herbert Falk
Sisco
Phone: 313-254-0020
Fax: 313-254-0053
sisco@mcimail.com

Patrik Faltstrom
KTH - Royal Institute of Technology
Phone: +46 8 7906274
Fax: +46 8 7900930
paf@nada.kth.se

Qin Fang
University of North Carolina at Chapel Hill
Phone: 919-962-9107
Fax: 919-962-5604
qin_fang@unc.edu

William Fardy
Frontier Software Development
Phone: 508-851-8872
billf@frontier.com

Dino Farinacci
cisco Systems
Phone: 415-688-4696
Fax: 415-903-8008
dino@cisco.com

Stefan Fassbender
IBM Corporation
Phone: +49 6221 594 446
Fax: +49 6221 593 300
stf@easi.net

Steve Feldman
MFS Datanet, Inc.
Phone: 408-975-2230
Fax: 408-975-2210
feldman@mfsdatanet.com

Sallie Fellows
New Hampshire State Department
Phone: 603-271-2778
Fax: 603-271-1953
sallie@ed.unh.edu

William Fenner
Naval Research Laboratory
Phone: 202-404-7030
Fax: 202-767-1683
fenner@cmf.nrl.navy.mil

Robert Fenoglio
IBM Corporation
Phone: 919-254-4354
fenoglio@vnet.ibm.com

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

Carlos Fernandez
Phillips Laboratory
Kirkland AFB
Phone: 505-846-1928
Fax: 505-846-0087
carlos@plk.af.mil

James Fielding
U.S. Army Research Laboratory
Phone: 410-278-6664
Fax: 410-278-5075
jamesf@arl.army.mil

Robert Fink
Lawrence Berkeley Laboratory
Phone: 510-486-5692
Fax: 510-486-6363
rlfink@lbl.gov

Dale Finkelson
MIDnet, Inc.
Phone: 402-472-5032
Fax: 402-472-8486
dmf@westie.mid.net

Eric Fleischman
Boeing Computer Support Services
Phone: 206-957-5334
Fax: 206-957-5050
ericf@atc.boeing.com

Sally Floyd
Lawrence Berkeley Laboratory
Phone: 510-486-7518
floyd@ee.lbl.gov

Peter Ford
Los Alamos National Laboratory
Phone: 505-665-0058
Fax: 505-665-7793
peter@goshawk.lanl.gov

James Forster
cisco Systems
Phone: 415-688-7854
Fax: 415-688-8282
forster@cisco.com

Jill Foster
University of Newcastle upon Tyne
Phone: +44 91 222 8250
Fax: +44 91 222 8765
Jill.Foster@newcastle.ac.uk

Catherine Foulston
Rice University
Phone: 713-285-5233
Fax: 713-527-6099
cathyf@rice.edu

Craig Fox
FTP Software, Inc.
Phone: 508-659-6219
craig@ftp.com

Richard Fox
Metricom
Phone: 408-399-8235
Fax: 408-354-1024
rfox@metricom.com

Paul Francis
Bell Communications Research
Phone: 201-829-4484
Fax: 201-829-2504
francis@cactus.ntt.jp

Ron Frederick
Xerox Corporation
Phone: 415-812-4459
Fax: 415-812-4471
frederick@parc.xerox.com

Ned Freed
Innosoft International, Inc.
Phone: 909-624-7907
Fax: 909-621-5319
ned@innosoft.com

Dan Frommer
Digital Equipment Corporation
Phone: +972 2 782 551
Fax: +972 2 782 373
dan@isv.dec.com

Vince Fuller
BARRNet
Phone: 415-723-6860
Fax: 415-723-0908
vaf@barrnet.net

Jim Fullton
MCNC - CNIDR
Phone: 919-248-1911
Fax: 919-248-1101
fullton@cnidr.org

Peter Furniss
Consultant
Phone: +44 81 313 1833
Fax: +44 81 313 1833
p.furniss@ulcc.ac.uk

James Galvin
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
galvin@tis.com

Kevin Gamiel
MCNC - CNIDR
Phone: 919-248-1911
Fax: 919-248-1101
kgamiel@cnidr.org

Joan Gargano
University of California,
Davis
Phone: 916-752-2591
Fax: 916-752-2591
jcgargano@ucdavis.edu

John Garrett
AT&T Bell Laboratories
Phone: 908-580-4719
Fax: 908-580-4721
jwg@garage.att.com

Mark Garrett
Bell Communications Research
Phone: 201-829-4439
Fax: 201-829-2504
mwg@faline.bellcore.com

Steve Garritano
Kalpana, Inc.
Phone: 408-749-1900
Fax: 408-749-1690
steveg@kalpana.com

John Gawf
Compatible Systems Corporation
Phone: 303-444-9532
Fax: 303-444-9595
gawf@compatible.com

Vincent Gebes
AT&T Corporation
Phone: +81 3 5561 2966
Fax: +81 3 3584 0810
vgebes@sys.attjens.co.jp

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

Tony Genovese
Lawrence Livermore National Laboratory
Phone: 415-423-2471
Fax: 415-422-1482
genovese@es.net

Farrell Gerbode
Rice University
Phone: 713-527-4988
Fax: 713-527-6099
farrell@rice.edu

Arlene Getchell
Lawrence Livermore National Laboratory
Phone: 510-423-6349
Fax: 510-422-1482
getchell@es.net

Atanu Ghosh
University College London
Phone: +44 71 387 7050
Fax: +44 71 387 1397
atanu@cs.ucl.ac.uk

Shawn Gillam
TimonWare, Inc.
Phone: 408-727-0699
Fax: 408-727-0785
shawn@timonware.com

Robert Gilligan
Sun Microsystems, Inc.
Phone: 415-336-1012
Fax: 415-336-6015
Bob.Gilligan@Eng.Sun.Com

Robert Glenn
National Institute of Standards and Technology
Phone: 301-975-3667
Fax: 301-590-0932
glenn@osi.ncsl.nist.gov

Mike Goguen
SynOptics Communications, Inc.
Phone: 408-764-1289
Fax: 408-764-1299
goguen@synoptics.com

Mei-Jean Goh
MPR Teltech, Ltd.
Phone: 604-294-1471
Fax: 604-293-5787
goh@mpr.ca

Fengmin Gong
MCNC
Phone: 919-248-9214
Fax: 919-248-1405
gong@concert.net

Chris Gorsuch
Texas Instruments
Phone: 214-575-4740
Fax: 214-575-4755
chrisg@lobby.ti.com

Ramesh Govindan
Bell Communications Research
Phone: 201-829-4406
Fax: 201-829-4504
rxg@thumper.bellcore.com

Tim Grams
International Power Machines
Phone: 214-205-1547
73537.2601@compuserve.com

Judith Grass
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
grass@cnri.reston.va.us

Cleve Graves
OpenConnect Systems, Inc.
Phone: 214-888-0416
Fax: 214-484-6100
cvg@oc.com

Terry Gray
University of Washington
Phone: 206-543-3880
Fax: 206-685-4045
gray@cac.washington.edu

Christine Gressley
University of Illinois
Phone: 217-244-0276
Fax: 217-244-7089
gressley@uiuc.edu

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, Inc.
Phone: 617-821-7333
Fax: 617-821-4218
dan@merlin.dev.cdx.mot.com

Robert Grow
XLNT Design, Inc.
Phone: 619-487-9320
Fax: 619-487-9768
bob@xlnt.com

Chris Gunner
Digital Equipment Corporation
Phone: 508-486-7792
Fax: 508-486-5279
gunner@dsmail.lkg.dec.com

Joel Gyllenskog
Hewlett-Packard
Phone: 208-396-4515
Fax: 208-396-4122
jgyllens@hpdmd48.boi.hp.com

Robert Hagens
Advanced Network and Services, Inc.
Phone: 703-758-7700
Fax: 703-758-7717
hagens@ans.net

Regina Hain
Bolt Beranek and Newman
Phone: 617-873-4652
Fax: 617-873-6091
rrosales@bbn.com

Stuart Hale
IBM Corporation
Phone: 919-543-7812
Fax: 919-543-4517
stu_hale@vnet.ibm.com

Martyne Hallgren
Cornell University
Phone: 607-254-8767
Fax: 607-255-5771
martyne@mitchell.cit.cornell.edu

Joel Halpern
Network Systems Corporation
Phone: 612-424-1606
Fax: 612-424-1736
jmh@network.com

Augusto Ham
Wiltel
Phone: 713-364-4169
ahamm@atg.wiltel.com

Jari Hamalainen
Nokia Mobile Phones
Phone: +358 31 165896
Fax: +358 31 165888
jah@rctre.nokia.com

Deborah Hamilton
AT&T Bell Laboratories
Phone: 908-949-9459
Fax: 908-949-1726
debbieh@internic.net

Brian Hammill
Exide Electronics
Phone: 919-872-3020
Fax: 919-878-2385
brianh@exide.com

John Hanratty
Agile Networks, Inc.
Phone: 508-287-0700
Fax: 508-287-0606
jhanratty@agile.com

Alf Hansen
UNINETT
Phone: +47 7 359 2982
Fax: +47 7 359 6450
Alf.Hansen@uninett.no

Daniel Hansen
Network General Corporation
Phone: 201-616-0986
Fax: 201-616-8311
danh@ngc.com

Herluf Hansen
Telebit Communications
Phone: +45 86 28 81 76
Fax: +45 86 28 81 86
hha@tbit.dk

W.S. Harborth
Houston Associates, Inc.
Phone: 703-812-8730
Fax: 703-812-5990
wharbort@ghost.darpa.mil

Susan Hares
Merit Network, Inc.
Phone: 313-936-2095
Fax: 313-747-3745
skh@merit.edu

Dimitry Haskin
Wellfleet Communications, Inc.
Phone: 508-436-8124
Fax: 508-670-8765
dhaskin@wellfleet.com

Marc Hasson
Mentat Inc.
Phone: 310-208-2650
Fax: 310-208-3724
marc@mentat.com

Eugene Hastings
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
hastings@psc.edu

Ken Hayward
Bell Northern Research
Phone: 613-763-4042
Fax: 613-763-2626
Ken.Hayward@bnr.ca

Denise Heagerty
CERN
Phone: +41 22 767 49 75
Fax: +41 22 767 71 55
denise@dxcoms.cern.ch

Cornelius Healy
Sprint
Phone: 703-904-2257
Fax: 703-904-2262
con@icp.net

Roland Hedberg
Technische Universiteit Delft
Phone: +31 15 785210
Fax: +31 15 783787
Roland.Hedberg@rc.tudelft.nl

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@cren.net

Shai Herzog
Information Sciences Institute
Phone: 310-821-3743
herzog@catarina.usc.edu

Jeff Hilgeman
Apertus Technologies, Inc.
Phone: 612-828-0668
Fax: 612-828-0102
jeffh@apertus.com

Kathryn Hill
Newbridge Networks
Phone: 703-708-5940
Fax: 703-708-5959
khill@newbridge.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: 916-752-0220
rdhobby@ucdavis.edu

Ellen Hoffman
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
ellen@merit.edu

Eric Hoffman
Naval Research Laboratory
Phone: 202-404-7019
Fax: 202-767-1683
hoffman@cmf.nrl.navy.mil

Mike Holloway
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
mikeh@newbridge.com

Matt Hood
NASA Science Internet
Phone: 415-604-4764
Fax: 415-604-0063
hood@nsipo.nasa.gov

John Hopprich
DAVID Systems, Inc.
Phone: 408-541-7459
Fax: 408-541-6980
hopprich@davidsys.com

Refael Horev
LANNET Data
Communications, Ltd.
Phone: +972 3 6458354
Fax: +972 3 5447146
horev@lannet.com

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

Kathy Huber
Wellfleet Communications, Inc.
Phone: 508-436-3791
Fax: 508-670-5747
khuber@wellfleet.com

Richard Huber
InterNIC/AT&T Bell
Laboratories
Phone: 908-949-4524
Fax: 908-949-1726
rvh@ds.internic.net

Jeff Hughes
Hewlett-Packard
Phone: 719-531-4777
Fax: 719-531-4540
jeff@col.hp.com

Christian Huitema
INRIA
Phone: +33 93 65 77 15
Fax: +33 93 65 77 65
Christian.Huitema@sophia.inria.fr

Erik Huizer
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
Erik.Huizer@SURFnet.nl

Melanie Humphrey
University of Illinois
Beckman Institute
Phone: 217-244-1079
Fax: 217-244-8371
msh@uiuc.edu

Farooq Hussain
Sprint Corporation
Phone: 703-904-2186
Fax: 703-904-2680
farooq@icm1.icp.net

Phil Irey
NSWC
Phone: 703-663-1582
Fax: 703-663-8673
pirey@relay.nswc.navy.mil

Kevin Jackson
Concord Communications, Inc.
Phone: 508-460-4646
Fax: 508-481-9772
kjackson@concord.com

Lenore Jackson
NASA Science Internet
Phone: 415-604-0455
Fax: 415-604-0063
jackson@nsipo.arc.nasa.gov

Warren Jackson
University of Toronto
Phone: 416-978-8948
Fax: 416-978-6620
wcj@vm.utcs.utoronto.ca

Ole Jacobsen
ConneXions
Phone: 415-962-2515
Fax: 415-949-1779
ole@interop.com

David Jacobson
IBM Corporation
Phone: 919-254-5590
Fax: 919-254-9515
dnjake@vnet.ibm.com

Van Jacobson
Lawrence Berkeley Laboratory
Phone: 510-486-7519
van@ee.lbl.gov

Ronald Jacoby
Silicon Graphics, Inc.
Phone: 415-390-1371
Fax: 415-969-2314
rj@sgi.com

B.V. Jagadeesh
Novell, Inc.
Phone: 408-473-8789
Fax: 408-946-9761
bvj@novell.com

Jan-Olof Jemnemo
Telia Research
Phone: +46 8 713 1423
Fax: +46 8 724 9243
jan-olof.jemnemo@farsta.trab.se

Barbara Jennings
Sandia National Laboratories
Phone: 505-845-8554
Fax: 505-844-2067
bjjenni@sandia.gov

David Johnson
Carnegie Mellon University
Phone: 412-268-7399
Fax: 412-681-5739
dbj@cs.cmu.edu

Dale Johnson
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
dsj@merit.edu

Timo Jokiaho
Nokia Telecommunications
Cellular Systems
Phone: +358 5104 2230
Fax: +358 5104 2411
timo.jokiaho@ntc.nokia.com

Laurent Joncheray
Merit Network, Inc.
Phone: 313-936-3335
lpj@merit.edu

Rick Jones
Hewlett-Packard
Phone: 408-447-2493
Fax: 408-447-2969
raj@cup.hp.com

Matthew Jonson
United States Air Force
Phone: 205-416-4075
Fax: 205-416-4738
jonson@ddn.af.mil

Merike Kaeo
cisco Systems
Phone: 415-903-7150
Fax: 415-903-7666
mkaeo@cisco.com

Marko Kaittola
FUNET
Phone: +358 31 162 933
Fax: +358 31 162 913
Marko.Kaittola@funet.fi

Jeanine Kamerdze
NASA Science Internet
Phone: 415-604-4764
Fax: 415-604-0063
kamerdze@nsipo.nasa.gov

Scott Kaplan
FTP Software, Inc.
Phone: 415-543-9001
Fax: 415-543-9002
scott@ftp.com

Phil Karn
Qualcomm, Inc.
Phone: 619-658-4359
Fax: 619-658-2113
karn@qualcomm.com

Frank Kastenholz
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
kasten@ftp.com

Neil Katin
SunSoft
Phone: 415-336-3618
Fax: 415-336-5178
katin@eng.sun.com

Akira Kato
The University of Tokyo
Computer Center
Phone: +81 3 3812 2111
Fax: +81 3 5684 7775
kato@wide.ad.jp

Yasuhiro Katsube
Bell Communications Research
Phone: 908-758-2039
Fax: 908-758-4192
katsube@mail.bellcore.com

Dave Katz
cisco Systems
Phone: 415-688-8284
Fax: 415-688-8282
dkatz@cisco.com

Charlie Kaufman
Digital Equipment Corporation
Phone: 603-881-1495
Fax: 603-881-2257
kaufman@zk3.dec.com

David Kaufman
Telco Systems, Inc.
Phone: 617-255-9400
Fax: 617-255-5885
dek@magna.telco.com

Elizabeth Kaufman
Yale University
Phone: 203-432-6679
Fax: 203-432-6165
kaufman@biomded.med.yale.edu

Hiroshi Kawazoe
IBM Japan, Ltd.
Phone: +81 462 73 4936
Fax: +81 462 74 4282
kawazoe@trl.ibm.co.jp

Brendan Kehoe
The Zen Internet Group
brendan@zen.org

Lois Keiper
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
lkeiper@cnri.reston.va.us

Bill Kelly
Auburn University
Phone: 205-844-4512
Fax: 205-844-9680
kellywh@mail.auburn.edu

Sean Kennedy
Bolt Beranek and Newman
Phone: 617-873-3354
Fax: 617-873-5620
liam@nic.near.net

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@snmp.com

Zbigniew Kielczewski
EICON Technology Corporation
Phone: 514-631-2592
Fax: 514-631-3092
zbig@eicon.qc.ca

Steve Kille
ISODE Consortium
Phone: +44 81 332 9091
Fax: +44 81 332 9019
S.Kille@isode.com

Lee Kilpatrick
Bolt Beranek and Newman
Phone: 617-873-4287
Fax: 617-873-6091
leekil@bbn.com

Byonghak Kim
Korea Advanced Institute of Science and Technology
Fax: +82 42 869 3554
bhkim@cosmos.kaist.ac.kr

Edwin King
Boeing Computer Services
Phone: 206-865-7308
Fax: 206-865-1603
eek@atc.boeing.com

John Klensin
United Nations University
Phone: 617-227-8747
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
NASA Ames Research Center
Phone: 415-604-1005
Fax: 415-604-6999
jknowles@binky.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
stev@ftp.com

Andrew Knutsen
The Santa Cruz Operation, Inc.
Phone: 408-427-7538
Fax: 408-458-0811
andrewk@sco.com

Michael Kornegay
VisiSoft/Technology Dynamics
Phone: 404-874-0428
Fax: 404-874-6412
mlk@bir.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@mail.bellcore.com

John Krawczyk
Wellfleet Communications, Inc.
Phone: 508-436-3811
Fax: 508-670-8760
jkrawczy@wellfleet.com

David Kristol
AT&T Bell Laboratories
Phone: 908-582-2250
Fax: 908-582-5192
dmk@allegra.att.com

Cheryl Krupczak
Empire Technologies, Inc.
Phone: 404-350-0107
cheryl@empiretech.com

John Kunze
University of California, Berkeley
Phone: 510-642-1530
Fax: 510-643-5385
jak@violet.berkeley.edu

Charles Kunzinger
IBM Corporation
Phone: 919-254-4142
Fax: 919-254-5410
kunzinger@vnet.ibm.com

Ted Kuo
IBM Corporation
Phone: 919-254-6214
Fax: 919-254-4027
tik@vnet.ibm.com

Sundar Kuttalingam
Wiltel Advanced Technology Group
Phone: 713-364-4132
Fax: 713-367-9036
sundark@wiltel.com

Olav Kvittem
UNINETT
Phone: +47 7359 6981
Fax: +47 7359 6450
Olav.Kvittem@uninett.no

William Kwan
Rabbit Software Corporation
Phone: 215-993-1135
Fax: 215-640-1379
kwan@rabbit.com

Paul Lambert
Motorola, Inc.
Phone: 602-441-3646
Fax: 602-441-8377
paul_lambert@email.mot.com

David Lapp
Hewlett-Packard
Phone: 519-883-3051
Fax: 519-886-8620
lapp@waterloo.hp.com

Mark Laubach
Hewlett-Packard Laboratories
Phone: 415-857-3513
Fax: 415-857-8526
laubach@hpl.hp.com

Walter Lazear
The MITRE Corporation
Phone: 703-883-6515
Fax: 703-883-7142
lazear@gateway.mitre.org

Gordon Lee
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
gordon@ftp.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

Edward Levinson
Accurate Information Systems
Phone: 908-389-5550
Fax: 908-389-5556
elevinson@accurate.com

Ben Levy
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
seven@ftp.com

Mark Lewis
Telebit Corporation
Phone: 408-745-3232
Fax: 408-734-3333
Mark.S.Lewis@telebit.com

Jian Li
Rice University
jian@rice.edu

Tony Li
cisco Systems
Phone: 415-688-8186
Fax: 415-688-4575
tli@cisco.com

Fred Lien
Boeing Computer Services
Phone: 206-957-5311
Fax: 206-957-5050
fred@speedster.ca.boeing.com

Lars-Johan Liman
EBONE NOC
Phone: +46 8 790 65 60
Fax: +46 8 24 11 79
liman@ebone.net

John Linn
OpenVision Technologies
Phone: 617-374-3700
Fax: 617-374-3715
linn@security.ov.com

Robin Littlefield
Wellfleet Communications, Inc.
Phone: 508-436-3812
Fax: 508-670-8760
robin@wellfleet.com

Joseph Liu
Wiltel
Phone: 713-364-4170
Fax: 713-367-0278
jliu@atg.wiltel.com

Dave Livingston
News & Observer Publishing Company
Phone: 919-829-4836
Fax: 919-829-4808
squirrel@vnet.net

Brian Lloyd
B.P. Lloyd & Associates, Inc.
Phone: 916-676-1147
Fax: 916-676-3442
brian@lloyd.com

Kanchei Loa
Motorola, Inc.
Phone: 602-897-5122
Fax: 602-897-3822
loa@sps.mot.com

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

E. Paul Love
Consultant
Phone: 802-254-9087
epl@sdsc.edu

Thang Lu
MCI Telecommunications
Phone: 214-918-6763
Fax: 214-918-6052
tlu@mcimail.com

Steven Lunt
Bell Communications Research
Phone: 908-699-4244
Fax: 908-336-2943
lunt@bellcore.com

Robert Lutz
Stac Electronics
Phone: 619-431-7474
Fax: 619-431-0880

Bryan Lyles
Xerox Corporation
Phone: 415-812-4466
Fax: 415-812-4471
lyles@parc.xerox.com

Clifford Lynch
University of California, Berkeley
Phone: 510-987-0522
Fax: 510-839-3573
calur@uccmvsa.ucop.edu

Carl Madison
3Com Corporation
Phone: 508-836-1735
Fax: 508-393-6934
carl@zeus.st.3com.com

Dan Magorian
University of Maryland
Phone: 301-405-3004
Fax: 301-314-9220
magorian@ni.umd.edu

Carl Malamud
Internet Multicasting Service
Phone: 202-628-2044
Fax: 202-628-2042
carl@malamud.com

Andrew Malis
Ascom Timeplex, Inc.
Phone: 508-266-4522
Fax: 508-264-4999
malis@maelstrom.timeplex.com

Gary Malkin
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-2618
gmalkin@xylogics.com

Tracy Mallory
3Com Corporation
Phone: 408-764-6518
Fax: 408-764-5002
tracym@3com.com

Allison Mankin
Naval Research Laboratory
Phone: 202-404-7030
Fax: 202-767-1683
mankin@cmf.nrl.navy.mil

Bill Manning
Rice University
Phone: 713-285-5415
Fax: 713-527-6099
bmanning@rice.edu

Glenn Mansfield
AIC Systems Laboratory
Phone: +81 22 279 3310
Fax: +81 22 279 3460
glenn@aic.co.jp

Peram Marimuthu
Wandel & Goltermann
Technologies, Inc.
Phone: 919-941-5730
Fax: 919-941-5751
peram@wg.com

April Marine
NASA NAIC
Phone: 415-604-0762
Fax: 415-604-0978
april@atlas.arc.nasa.gov

Lawren Markle
Tripp Lite
Phone: 312-329-1777
72170.460@compuserve.com

David Marlow
NSWC
Phone: 703-663-1675
dmarlow@relay.nswc.navy.mil

Jim Martin
Rutgers University
Phone: 908-932-3719
Fax: 908-932-2968
jim@noc.rutgers.edu

Larry Masinter
Xerox Corporation
Phone: 415-812-4365
Fax: 415-812-4333
masinter@parc.xerox.com

Thomas Maslen
SunSoft
Phone: 415-336-5099
Fax: 415-965-4903
maslen@eng.sun.com

Matt Mathis
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
mathis@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

Chip Matthes
Delphi Internet Services
Phone: 617-491-3342
Fax: 617-491-6642
chip@delphi.com

Mark McCahill
University of Minnesota
Phone: 612-625-1300
Fax: 612-625-6817
mpm@boombox.micro.umn.edu

Keith McCloghrie
Hughes LAN Systems
Phone: 415-966-7934
Fax: 415-966-7838
kzm@hls.com

Wayne McDilda
Department of Information Resources
Phone: 512-475-4715
Fax: 512-475-4759
wayne@dir.texas.gov

Evan McGinnis
3Com Corporation
Phone: 408-764-6064
Fax: 408-764-6134
bem@3com.com

Glenn McGregor
B.P. Lloyd & Associates, Inc.
Phone: 916-676-1147
Fax: 916-676-3442
ghm@lloyd.com

Michael McLay
National Institute of Standards and Technology
Phone: 301-975-4099
Fax: 301-975-4091
mclay@eeel.nist.gov

Piers McMahon
ICL
Phone: +44 734 586211
Fax: +44 734 855106
p.v.mcmahon@rea0803.wins.icl.co.uk

Jim McQuaid
Wandel & Goltermann Technologies, Inc.
Phone: 919-941-5730
Fax: 919-941-5751
mcquaid@wg.com

Steve McRobert
Advanced Micro Devices
Phone: 408-987-2492
Fax: 408-987-2800
steve.mcrobert@amd.com

Michael Mealling
Georgia Institute of Technology
Phone: 404-894-1712
Fax: 404-894-9548
michael.mealling@oit.gatech.edu

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

Michael Michnikov
The MITRE Corporation
Phone: 617-271-2807
Fax: 617-271-3816
mbmg@mitre.org

Stephen Miller
Bolt Beranek and Newman
Phone: 617-873-4153
Fax: 617-873-5620
smiller@bbn.com

Greg Minshall
Novell, Inc.
Phone: 510-975-4507
Fax: 510-947-1238
minshall@wc.novell.com

William Miskovetz
cisco Systems
Phone: 415-688-4682
Fax: 415-688-4575
misko@cisco.com

Mitra
Pandora Systems
Phone: 415-488-0944
Fax: 415-931-7931
mitra@pandora.sf.ca.us

Randy Miyazaki
Lantron Design, Inc.
Phone: 408-452-0787
Fax: 408-452-0788
randy@lantron.com

Paul Mockapetris
USC/ISI
pvm@isi.edu

Pushpendra Mohta
CERFnet
Phone: 619-455-3900
Fax: 619-455-3990
pushp@cerf.net

Gordon Monk
Netx, Inc.
Phone: 214-644-2599
ghanka@metronet.com

Doug Montgomery
National Institute of
Standards and Technology
Phone: 301-975-3630
Fax: 301-590-0932
dougm@osi.ncsl.nist.gov

Keith Moore
University of Tennessee
Phone: 615-974-3735
Fax: 615-974-8296
moore@cs.utk.edu

Robert Moose
The MITRE Corporation
Phone: 703-883-5576
Fax: 703-883-7142
rmoose@gateway.mitre.org

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
morris@altair.disa.mil

Robert Moskowitz
Chrysler Corporation
Phone: 313-758-8212
Fax: 313-758-8173
3858921@mcimail.com

Kevin Mullet
Rice University
kevinm@rice.edu

Jun Murai
Keio University
Phone: +81 466 47 5111
Fax: +81 466 49 1101
jun@wide.ad.jp

Sandra Murphy
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
murphy@tis.com

John Myers
Carnegie Mellon University
Phone: 412-268-2314
Fax: 412-268-5249
jgm+@cmu.edu

Andrew Myles
MacQuarie University
Phone: +61 2 8059071
andrew@mpce.mg.edu.au

Rina Nathaniel
RAD Network Devices, Ltd.
Phone: +972 3 6458556
Fax: +972 3 5447368
rina@rnd-gate.rad.co.il

Mark Needleman
University of California
Phone: 510-987-0530
Fax: 510-839-3573
mhn@stubbs.ucop.edu

Sath Nelakonda
Lachman Technology, Inc.
Phone: 708-505-9555
Fax: 708-505-9574
sath@lachman.com

Clifford Neuman
Information Sciences Institute
Phone: 310-822-1511
bcn@isi.edu

Chris Newman
Carnegie Mellon University
Phone: 412-268-6403
Fax: 412-268-4987
chrisn+@cmu.edu

Andy Nicholson
Microsoft Corporation
Phone: 206-936-3775
Fax: 206-936-7329
andyni@microsoft.com

Orly Nicklass
RAD Data Communications
Phone: +97 23 6459400
Fax: +97 23 6472675
orly@radmail.rad.co.il

Tom Nisbet
Telecommunications Techniques Corporation
Phone: 301-353-1550
Fax: 301-353-1536
nisbet@fbsw.tt.com

Shannon Nix
Netlink, Inc.
Phone: 919-878-8612
Fax: 919-872-2132
sdn@netlink.com

Dan Nordell
NSP Research
Phone: 612-330-5822
Fax: 612-330-6810

Erik Nordmark
Sun Microsystems, Inc.
Phone: 415-336-2788
Fax: 415-960-6974
nordmark@eng.sun.com

Bill Norton
Merit Network, Inc.
Phone: 313-936-2656
Fax: 313-747-3745
wbn@merit.edu

Michael O'Dell
UUNET Technologies, Inc.
Phone: 703-204-8000
Fax: 703-204-8001
mo@uunet.uu.net

Karen O'Donoghue
Naval Surface Warfare Center
Phone: 703-663-1567
Fax: 703-663-8673
kodonog@relay.nswc.navy.mil

Masataka Ohta
Tokyo Institute of Technology
Computer Center
Phone: +81 3 5499 7084
Fax: +81 3 3729 1940
mohta@cc.titech.ac.jp

Lars-Gunnar Olsson
Swedish University of
Agriculture
Phone: +46 18 67 17 76
Fax: +46 18 30 15 33
Lars-Gunnar.Olsson@data.slu.se

Steven Onishi
Wellfleet Communications, Inc.
Phone: 508-436-3816
Fax: 508-670-8760
sonishi@wellfleet.com

Zbigniew Opalka
Agile Networks, Inc.
Phone: 508-287-0700
Fax: 508-287-0606
zopalka@agile.com

Scott Paisley
National Institute of
Standards and Technology
Phone: 303-497-7691
Fax: 303-497-3012
paisley@central.bldrdoc.gov

Todd Palgut
Network Express, Inc.
Phone: 313-761-5051
Fax: 313-995-1114
todd@nei.com

William Palter
TGV, Inc.
Phone: 408-457-5200
Fax: 408-457-5205
palter@tgv.com

Vijayaragavan Pandian
Proteon, Inc.
Phone: 508-898-2800
Fax: 508-898-3174
vjp@proteon.com

Steve Parker
Fujitsu Open Systems Solutions
Phone: 510-652-6200
sparker@ossi.com

Andrew Partan
UUNET Technologies, Inc.
Phone: 703-204-8000
Fax: 703-204-8001
asp@uunet.uu.net

Craig Partridge
Bolt Beranek and Newman
Phone: 415-326-4541
craig@bbn.com

Ismat Pasha
Sprint
Phone: 703-904-2259
Fax: 703-904-2292
ipasha@icm1.icp.net

Brad Passwaters
Bell Atlantic
Phone: 301-236-6221
bjp@eng.umd.edu

Laura Pate
The MITRE Corporation
Phone: 703-883-5400
Fax: 703-883-7142
pate@gateway.mitre.org

Rakesh Patel
Rutgers University
Phone: 908-932-3436
Fax: 908-932-2968
rapatel@pilot.njin.net

Michael Patton
Bolt Beranek and Newman
Phone: 617-873-2737
Fax: 617-873-3315
map@bbn.com

Gregory Pearson
Microcom, Inc.
Phone: 818-366-1052
Fax: 818-831-6760

Jon Penner
DCA, Inc.
Phone: 512-244-3871
Fax: 512-388-7090
jjp@bscs.uucp

John Penners
US West Advanced 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
Naval Research Laboratory
Phone: 202-404-7026
Fax: 202-767-1683
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
Novell, Inc.
Phone: 408-321-1722
Fax: 408-946-9761
perlman@novell.com

Marsha Perrott
PREPnet
Phone: 412-268-7870
Fax: 412-268-7875
perrott@prep.net

Wayne Peters
Telxon Company
Phone: 216-867-3700
Fax: 216-873-2057
waynep@telxon.com

Eric Peterson
Xyplex, Inc.
Phone: 508-952-4743
Fax: 508-952-4702
elpeterson@eng.xyplex.com

Karen Petraska-Veum
NASA/Hughes STX
Phone: 301-286-3348
Fax: 301-286-1777
karen.veum@gsfc.nasa.gov

James Philippou
Xyplex, Inc.
Phone: 508-264-9900
japhilippou@eng.xyplex.com

David Piscitello
Core Competence, Inc.
Phone: 215-830-0692
wk04464@worldlink.com

Ronald Pitt
Network Security Systems, Inc.
Phone: 619-622-8712
Fax: 619-552-9162
57313.577@compuserve.com

Jon Postel
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
postel@isi.edu

Venkat Prasad
3Com Corporation
Phone: 408-764-5226
Fax: 408-764-5003
vsp@3com.com

Cecilia Preston
Consultant
Phone: 510-547-3207
cpreston@info.berkeley.edu

J. Mark Pullen
George Mason University
Phone: 703-993-1538
Fax: 703-993-3729
mpullen@cs.gmu.edu

John Quarterman
Texas Internet Consulting
Phone: 512-451-6176
Fax: 512-450-1436
jsq@tic.com

James Quigley
Hewlett-Packard
Phone: 408-447-3505
Fax: 408-447-2247
Quigley@cup.hp.com

Bala Rajagopalan
AT&T Bell Laboratories
Phone: 908-949-8017
Fax: 908-949-1726
braja@qsun.att.com

Ram Ramanathan
Bolt Beranek and Newman
Phone: 617-873-2736
Fax: 617-873-6089
ramanath@bbn.com

David Rand
Novell, Inc.
Phone: 408-321-1259
Fax: 408-946-9761
dave_rand@novell.com

Jim Rees
University of Michigan
Phone: 313-763-4174
Fax: 313-763-4434
Jim.Rees@umich.edu

Kenneth Rehbehn
Netrix Corporation
Phone: 703-742-6000
Fax: 703-742-4048
kjr@netrix.com

Alex Reijnierse
PTT Research
Phone: +31 70 3325779
Fax: +31 70 3326477
a.a.l.reijnierse@research.ptt.nl

Yakov Rekhter
IBM Corporation
Phone: 914-945-3896
Fax: 914-945-2141
yakov@watson.ibm.com

Eddie Renoux
McDATA Corporation
Phone: 303-460-4198
Fax: 303-465-4996
elr0262@newsit2.mcdata.com

Robert Reschly
US Army Research Laboratory
Phone: 410-278-8612
Fax: 410-278-5077
reschly@brl.mil

Joyce K. Reynolds
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
jkrey@isi.edu

Tony Richards
Sprint
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

Mike Ritter
Apple Computer, Inc.
Phone: 408-862-8088
Fax: 408-862-1159
mwritter@applelink.apple.com

Michael Roberts
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
roberts@educom.edu

Allen Rochkind
3Com Corporation
Phone: 408-764-6014
Fax: 408-764-5002
Allen_Rochkind@3com.com

Kenneth Rodemann
AT&T Bell Laboratories
Phone: 908-949-6229
Fax: 908-949-1726
krr@qsun.att.com

Robert Roden
Digital Equipment Corporation
Phone: +44 734 204476
Fax: +44 734 203133
roden@roden.enet.dec.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 6458568
Fax: +972 3 5447368
brodrig@rnd-gate.rad.co.il

Doris Roland
Houston Associates, Inc.
Phone: 703-527-1970
Fax: 703-527-1973
droland@ghost.darpa.mil

Allyn Romanow
Sun Microsystems, Inc.
Phone: 415-336-1326
Fax: 415-336-4802
allyn.romanow@eng.sun.com

Marshall T. Rose
Dover Beach Consulting, Inc.
Phone: 415-968-1052
Fax: 415-968-2510
mrose@dbc.mtview.ca.us

Kenneth Rossen
SHL Systemhouse
Phone: 508-928-5368
Fax: 508-481-8996
kenr@shl.com

Shawn Routhier
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sar@epilogue.com

Michal Rozenthal
Fibronics, Ltd.
Phone: +972 4 313642
Fax: +972 4 550550
michal@fibronics.co.il

Allan Rubens
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
acr@merit.edu

Greg Ruth
GTE Laboratories, Inc.
Phone: 617-466-2448
Fax: 617-890-9320
gruth@gte.com

Michael Safly
NASA/Boeing Computer Support Services
Phone: 205-544-6445
Fax: 205-544-4618
mike.safly@msfc.nasa.gov

Timothy Salo
Minnesota Supercomputer Center
Phone: 612-337-3555
Fax: 612-337-3400
tjs@msc.edu

Hal Sandick
IBM Corporation
Phone: 919-254-4614
Fax: 919-254-4975
sandick@vnet.ibm.com

Jon Saperia
Digital Equipment Corporation
Phone: 603-881-0480
Fax: 603-881-0120
saperia@zko.dec.com

Srinivas Sataluri
AT&T Bell Laboratories
Phone: 908-949-7782
Fax: 908-949-1726
sri@internic.net

Michael Scanlon
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
scanlon@ftp.com

Jeffrey Schiller
Massachusetts Institute of Technology
Phone: 617-253-0161
Fax: 617-258-8736
jis@mit.edu

Richard Schmalgemeier
Merit Network, Inc.
Phone: 313-936-3000
rgs@merit.edu

Jean-Bernard Schmitt
IBM Corporation, France
Phone: +33 92 11 56 60
Fax: +33 93 24 78 09
jbs@vnet.ibm.com

Wolfgang Schneider
Gesellschaft fuer Mathematik und Datenverarbeitung
Phone: +49 6151 869 719
Fax: +49 6151 869 785
schneiw@darmstadt.gmd.de

Eve Schooler
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
`schooler@isi.edu`

Rickard Schoultz
SUNET/KTHnoc
Phone: +46 8 790 9088
Fax: +46 8 102 510
`schoultz@sunet.se`

Martin Schulman
Sprint International
Phone: 703-707-9313
Fax: 703-689-5176
`schulman@smtp.sprint.com`

Henning Schulzrinne
AT&T Bell Laboratories
Phone: 908-582-2262
Fax: 908-582-5809
`hgs@research.att.com`

Michael Schwartz
University of Colorado, Boulder
Phone: 303-492-3902
Fax: 303-492-0391
`schwartz@cs.colorado.edu`

Dallas Scott
The MITRE Corporation
Phone: 703-883-6700
Fax: 703-883-1364
`scott@fluky.mitre.org`

Isil Sebuktekin
Bell Communications Research
Phone: 201-829-4725
Fax: 201-829-2504
`isil@nevin.bellcore.com`

Michael See
IBM Corporation
Phone: 919-254-3480
Fax: 919-254-0391
`mikesee@vnet.ibm.com`

Jennifer Sellers
NASA NREN
Phone: 202-434-8954
Fax: 202-434-4599
`sellers@quest.arc.nasa.gov`

Paul Serice
Corporation for Open Systems International
Phone: 703-205-2713
Fax: 703-846-8590
`serice@cos.com`

Kanan Shah
Naval Research Laboratory
Phone: 202-404-7029
Fax: 202-767-1683
`kshah@cmf.nrl.navy.mil`

Satya Sharma
IBM Corporation
Phone: 512-388-3396
`ssharma@chang.austin.ibm.com`

Chris Shaw
Banyan Systems, Inc.
Phone: 508-836-2898
Fax: 508-898-3604
`cshaw@banyan.com`

Vincent Shekher
Motorola, Inc.
Phone: 602-897-5378
Fax: 602-897-3822
`vin@sps.mot.com`

Scott Shenker
Xerox Corporation
Phone: 415-812-4840
Fax: 415-812-4471
`shenker@parc.xerox.com`

Erik Sherk
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
`sherk@sura.net`

Ming Sheu
IBM Corporation
Phone: 919-254-4262
Fax: 919-254-6430
msheu@vnet.ibm.com

Kitty Shih
Novell, Inc.
Phone: 408-747-4000
Fax: 408-747-4900
kmshih@novell.com

Uttam Shikarpur
Digital Equipment Corporation
Phone: 603-881-1866
Fax: 603-881-2257
uttam@zk3.dec.com

Chi Shue
Cascade Communications Corporation
Phone: 508-692-2600
Fax: 508-692-9214
chi@casc.com

William Simpson
Computer Systems Consulting Services
Phone: 313-996-0769
Bill.Simpson@um.cc.umich.edu

W. David Sincoskie
Bell Communications Research
Phone: 201-829-4426
Fax: 201-829-2504
sincos@bellcore.com

Henry Sinnreich
MCI Telecommunications
Phone: 214-918-5231
Fax: 214-918-6019
hsinnreich@mcimail.com

Sam Sjogren
TGV, Inc.
Phone: 408-457-5200
Fax: 408-457-5205
sjogren@tgv.com

Keith Sklower
University of California, Berkeley
Phone: 510-642-9587
Fax: 510-642-5775
sklower@cs.berkeley.edu

Timon Sloane
timonWare, Inc.
Phone: 415-727-0699
Fax: 415-727-0785
timon@timonware.com

Andrew Smith
SynOptics Communications, Inc.
Phone: 408-764-1574
Fax: 408-988-5525
asmith@synoptics.com

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: 510-423-9557
Fax: 510-422-1482
smith@es.net

Robert Snyder
cisco Systems
Phone: 415-688-8245
Fax: 415-688-8282
snyder@cisco.com

Frank Solensky
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-659-6105
solensky@ftp.com

Karen Sollins
Massachusetts Institute of Technology
Phone: 617-253-6002
Fax: 617-253-2673
sollins@lcs.mit.edu

Dave Solo
Bolt Beranek and Newman
Phone: 617-873-6261
Fax: 617-873-4086
solo@bbn.com

James Solomon
Motorola, Inc.
Phone: 708-576-2753
Fax: 708-576-3240
solomon@comm.mot.com

Tae Song
Novell, Inc.
Phone: 408-473-8439
Fax: 408-946-9761
tae@novell.com

William Soren Deigaard
Rice University
soren@rice.edu

Milan Sova
Czech Technical University
Phone: +42 2 332 3925
Fax: +42 2 311 1786
sova@feld.cvut.cz

Michael Speer
Sun Microsystems, Inc.
Phone: 415-336-1729
Fax: 415-336-6315
michael.speer@sun.com

Simon Spero
University of North Carolina at Chapel Hill
Phone: 919-962-9107
Fax: 919-962-5604
ses@unc.edu

Michael St. Johns
Advanced Research Projects Agency/CSTO
Phone: 703-696-2271
Fax: 703-696-2202
stjohns@arpa.mil

Margaret St. Pierre
WAIS, Inc.
Phone: 415-327-9247
Fax: 415-327-6513
saint@wais.com

David Staudt
National Science Foundation
Phone: 703-306-1950
dstaudt@nsf.gov

Martha Steenstrup
Bolt Beranek and Newman
Phone: 617-873-3192
Fax: 617-873-3776
msteenst@bbn.com

Louis Steinberg
IBM Corporation
Phone: 203-783-7175
louiss@vnet.ibm.com

Don Stephenson
SunSoft
Phone: 415-336-4366
Fax: 415-336-6015
don.stephenson@sun.com

Barbara Sterling
McDATA Corporation
Phone: 303-460-4211
Fax: 303-465-4996
bjs@mcdata.com

Robert Stevens
Competitive Automation
Phone: 415-321-4006
robs@join.com

Bob Stewart
Xyplex, Inc.
Phone: 508-952-4816
Fax: 508-952-4887
rlstewart@eng.xyplex.com

John Stewart
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
jstewart@cnri.reston.va.us

Jack Stiekema
Victron bv
Phone: +31 50 446284
Fax: +31 50 424107
jack@victron.nl

David Stine
cisco Systems
Phone: 415-688-4681
Fax: 415-688-4575
dsa@cisco.com

Bernhard Stockman
Royal Institute of Technology
Phone: +46 8 790 6519
Fax: +46 8 241 179
boss@ebone.net

Adam Stolinski
Merlin Gerin
Phone: 714-557-1637
Fax: 714-557-7230
stolinsk@cerf.net

Vladimir Sukonnik
Process Software Corporation
Phone: 508-879-6994
Fax: 508-879-0042
sukonnik@process.com

Craig Summerhill
Coalition for Networked Information
Phone: 202-296-5098
Fax: 202-872-0884
craig@cni.org

Steve Suzuki
Furukawa Electric Technologies
Phone: 408-248-4884
Fax: 408-248-8815
steve@fet.com

George Swallow
Bolt Beranek and Newman
Phone: 617-873-3102
Fax: 617-873-6263
gswallow@bbn.com

Daniel Swinehart
Xerox Corporation
Phone: 415-812-4473
Fax: 415-812-4471
swinehart.parc@xerox.com

John Tavs
IBM Corporation
Phone: 919-254-7610
Fax: 919-254-4027
tavs@vnet.ibm.com

Larry Tepper
Compatible Systems Corporation
Phone: 303-444-9532
Fax: 303-444-9595
ltepper@compatible.com

Matsuaki Terada
Hitachi, Ltd.
Phone: +81 44 966 9111
Fax: +81 44 966 6853
tera@sdl.hitachi.co.jp

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

Michael Thatcher
Timonware, Inc.
Phone: 408-776-0275
thatcher@rahul.net

Richard Thomas
Bell Northern Research
Phone: 613-763-7025
Fax: 613-763-8385
rjthomas@bnr.ca

Susan Thomson
Bell Communications Research
Phone: 201-829-4514
Fax: 201-829-2504
set@bellcore.com

Dean Throop
Data General Corporation
Phone: 919-248-6081
Fax: 919-248-5942
throop@dg-rtp.dg.com

Akihiro Tominaga
Keio University
Phone: +81 466 47 5111
Fax: +81 466 49 1101
tomy@sfc.wide.ad.jp

Claudio Topolcic
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
topolcic@cnri.reston.va.us

Jerry Toporek
Mentat Inc.
Phone: 310-208-2650
Fax: 310-208-3724
jt@mentat.com

Paul Traina
cisco Systems
Phone: 415-325-7943
Fax: 415-688-7899
pst@cisco.com

Thuan Tran
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-0294
thuan@xylogics.com

Hoe Trinh
IBM Corporation
Phone: 919-254-6226
htrinh@vnet.ibm.com

Rick Troth
Rice University
Phone: 713-285-5148
Fax: 713-527-6099
troth@rice.edu

Theodore Ts'o
Massachusetts Institute of Technology
Phone: 617-253-8091
Fax: 617-248-8736
tytso@mit.edu

Panos Tsigaridas
GMD - FOKUS
Phone: +49 30 25499232
Fax: +49 30 25499202
Tsigaridas@fokus.gmd.de

Keisuke Uehara
The University of
Electro-Communications
Phone: +81 424 83 2161
Fax: +81 424 82 3055
kei@cs.uec.ac.jp

Robert Ullmann
Lotus Development Corporation
Phone: 617-693-1315
Fax: 617-693-5561
ariel@world.std.com

Anthony Valle
SPARTA, Inc.
Phone: 205-837-5200
Fax: 205-830-0287
valle@huntsville.sparta.com

Dono van-Mierop
3Com Corporation
Phone: 408-764-5722
Fax: 408-764-5002
dono_van_mierop@3mail.3com.com

Gregory Vaudreuil
Tigon Corporation
Phone: 214-733-2722
Fax: 214-733-8666
gvaudre@cnri.reston.va.us

Raymond Vega
CICESE
rvega@cicese.mx

John Veizades
FTP Software, Inc.
Phone: 415-543-9001
Fax: 415-543-9002
veizades@ftp.com

John Vollbrecht
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
jrv@merit.edu

Janet Vratny
Apple Computer
Phone: 408-974-2048
Fax: 408-725-8502
janet@apple.com

Joe Waggoner
Central & South West Services
Phone: 214-777-1579
Fax: 214-777-1033

Steven Waldbusser
Carnegie Mellon University
Phone: 412-268-6628
Fax: 412-268-4987
waldbusser@andrew.cmu.edu

Megan Walnut
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
mwalnut@cnri.reston.va.us

Thomas Walsh
Kalpana, Inc.
Phone: 408-749-1600
tomw@kalpana.com

Alice Wang
Sun Microsystems, Inc.
Phone: 415-336-5394
Fax: 415-336-3156
alice.wang@eng.sun.com

Chuck Warlick
Boeing Computer Support
Services
Phone: 205-544-6270
Fax: 205-544-4618
chuck.warlick@pscni.nasa.gov

William Warner
Division of Computer Services
and Telecommunications
Phone: 614-466-6683
Fax: 614-644-3349
warner@ohio.gov

James Watt
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
james@newbridge.com

William Weems
University of Texas, Houston
Phone: 713-792-5277
Fax: 713-794-5653
wweens@oac.hsc.uth.tmc.edu

Chris Weider
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8134
clw@bunyip.com

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

Guy Wells
US West Advanced Technologies
Phone: 303-541-6759
Fax: 303-541-6276
guyw@uswest.com

Taehwan Weon
Korea Advanced Institute of Science and Technology
Phone: +82 42 869 3554
Fax: +82 42 869 3510
weon@cosmos.kaist.ac.kr

Cory West
Rice University
corywest@rice.edu

Evan Wetstone
IBM Corporation
Phone: 203-783-5439
Fax: 203-783-5491
evanw@vnet.ibm.com

Brien Wheeler
The MITRE Corporation
Phone: 617-271-2567
Fax: 617-271-2352
blw@mitre.org

Chris Wheeler
University of Washington
Phone: 206-543-7741
Fax: 206-685-4044
cwheeler@cac.washington.edu

Gerry White
Applitek Corporation
Phone: 508-475-4050
Fax: 508-475-0550
gerry@lancity.com

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@vnet.ibm.com

Scott Williamson
Network Solutions, Inc.
Phone: 703-742-4820
Fax: 703-742-4811
scottw@nic.ddn.mil

Bradley Wilson
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
wilson@ftp.com

Jackie Wilson
NASA/Boeing Computer Support Services
Phone: 205-544-5309
Fax: 205-544-4618
Jackie.Wilson@msfc.nasa.gov

Peter Wilson
3Com Corporation
Phone: +44 442 278102
peter_wilson@3mail.3com.com

Greg Wilterdink
Alpha Technologies, Inc.
Phone: 206-647-2360
Fax: 206-671-4931
71302.3624@compuserve.com

Walter Wimer
Carnegie Mellon University
Phone: 412-268-6252
Fax: 412-268-4987
walter.wimer@andrew.cmu.edu

Cathy Wittbrodt
BARRNet
Phone: 415-725-5481
Fax: 415-723-0100
cjw@barrnet.net

Jane Wojcik
Bolt Beranek and Newman
Phone: 617-873-6247
Fax: 617-873-4086
jwojcik@bbn.com

David Woodgate
CSIRO
Phone: +61 6 276 6810
Fax: +61 6 276 6536
David.Woodgate@its.csiro.au

Paul Woodruff
SynOptics Communications, Inc.
Phone: 408-764-7417
Fax: 408-764-1299
pwoodruff@synoptics.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

Honda Wu
Network Application Technology
Phone: 408-370-4365
Fax: 408-370-4222
honda@nat.com

Liang Wu
Bell Communications Research
Phone: 201-829-4365
Fax: 201-829-2504
ltwu@bellcore.com

Jean Yao
Hewlett-Packard
Phone: 408-447-2934
Fax: 408-447-3660
yao@cup.hp.com

Peter Yee
NASA Ames Research Center
Phone: 415-604-3812
Fax: 415-604-6999
yee@atlas.arc.nasa.gov

Henry Yip
Network Application Technology
Phone: 408-370-4201
Fax: 408-370-4222
natadm!henry@uunet.uu.net

Shinichi Yoshida
Sumitomo Electric U.S.A., Inc.
Phone: 408-737-8517
Fax: 408-737-0134
yoshida@sumitomo.com

Brian Young
Best Power Technology, Inc.
Phone: 800-356-5794
Fax: 608-565-2221
byoung@macc.wisc.edu

Jessica Yu
Merit Network, Inc.
Phone: 313-936-2655
Fax: 313-747-3745
jyy@merit.edu

Chin Yuan
Pacific Bell
Phone: 510-867-5016
Fax: 510-867-1405
cxyuan@pacbell.com

Mauro Zallocco
Netlink, Inc.
Phone: 919-878-0402
Fax: 919-872-2132
mdz@netlink.com

Lixia Zhang
Xerox Corporation
Phone: 415-812-4415
Fax: 415-812-4471
lixia@parc.xerox.com

Weiping Zhao
National Center for Science
Information Systems
Phone: +81 3 3944 7574
Fax: +81 3 5395 7064
zhao@nacsis.ac.jp

Steve Zilles
Adobe Systems, Inc.
Phone: 415-962-4766
szilles@adobe.com

James Zmuda
Hughes Aircraft Company
Phone: 714-732-7240
Fax: 714-732-1185
zmuda@mls.hac.com

Miriam Zohar
IBM Corporation
Phone: 914-945-2567
zohar@watson.ibm.com

Mary Jo Zukoski
The MITRE Corporation
Phone: 703-883-7587
Fax: 703-883-7142
maryjo@gateway.mitre.org

Terry Zumwalt
International Power Machines
Corporation
Phone: 214-205-1540
Fax: 214-494-2690
71151.2550@compuserve.com

Corporation
for
National
Research
Initiatives

*Research
and
Development
for
the
National
Information
Infrastructure*

1895 Preston White Drive
Suite 100
Reston, VA 22091
USA

Tel: (703) 620-8990
Fax: (703) 620-0913