Proceedings of the Twenty-Fifth Internet Engineering Task Force

U.S. Sprint
Washington, DC
November 16-20, 1992

Corporation for National Research Initiatives

PROCEEDINGS OF THE TWENTY-FIFTH INTERNET ENGINEERING TASK FORCE

U.S. SPRINT
WASHINGTON, DC
November 16-20, 1992

Compiled and Edited by
Megan Davies
Cynthia Clark
Debra Legare

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

This material is based on work supported by the National Science Foundation under Grant. No. NCR 8820945.

ACKNOWLEDGEMENTS

The Twenty-Fifth Internet Engineering Task Force was held at the Hyatt Regency on Capitol Hill, Washington, DC, during the week of November 16th. Attendance figures did not quite reach the numbers in Boston, but were still healthy at 634. Seventy-Nine working groups and Birds of a Feather sessions took place during the course of the week.

Internet connectivity was provided by Sprint who served as our Local Host for this IETF. Gary Wightman and Marty Schulman, along with Joe Apple, Brian Shiflett, Craig Haney and others worked tirelessly to get the terminal room up and running.

Our thanks go to those individuals and organizations listed below for their generous donation of equipment and materials for the terminal room, as well as for those who gave informative presentations throughout the course of the week.

Terminal Room Equipment

Cabletron	SprintLink Hubs
Data General	Workstations
EMCOR	Equipment Cabinet
Silicon Graphics	Personal Iris Workstations
Sun Microsystems	Sun SPARC Workstations

Technical Presentations

Paul Tsuchiya, Peter Ford, Mark Knopper, Steve Deering, Bob Hinden, Steve Crocker, Chuck Cranor, Deborah Estrin, Jerry Rainville, Paul Turner and Cyndi Mills.

Once again thanks go to Debra Legare and Cynthia Clark for their help with these Proceedings and for their individual efforts to ensure a successful meeting. Thanks also to Terry Weigler and Cynthia Matthews for their assistance with the on-site logistics.

Megan Davies/CNRI

Contents

Director's Message **1**

Agenda of the Twenty-Fifth IETF **9**

1 IETF Overview **21**
- 1.1 Future IETF Meeting Sites 25
- 1.2 On Line IETF Information 27
  - 1.2.1 The IETF Directory 28
  - 1.2.2 The Internet-Drafts Directory 29
- 1.3 Guidelines to Authors of Internet-Drafts 31

2 Area and Working Group Reports **33**
- 2.1 Applications Area 35
  - 2.1.1 Internet Mail Extensions (smtpext) 53
  - 2.1.2 Internet Message Extensions (822ext) 61
  - 2.1.3 Network Database (netdata) 63
  - 2.1.4 Network News Transport Protocol (nntp) 81
  - 2.1.5 Network Printing Protocol (npp) 85
  - 2.1.6 TELNET (telnet) 87
- 2.2 Internet Area 91
  - 2.2.1 Dynamic Host Configuration (dhc) 117
  - 2.2.2 IP Address Encapsulation (ipae) 123
  - 2.2.3 IP over AppleTalk (appleip) 137
  - 2.2.4 IP over Asynchronous Transfer Mode (atm) 141
  - 2.2.5 P. Internet Protocol (pip) 155
  - 2.2.6 Point-to-Point Protocol Extensions (pppext) 159
  - 2.2.7 Router Requirements (rreq) 165
  - 2.2.8 Simple Internet Protocol (sip) 167
  - 2.2.9 TCP/UDP over CLNP-addressed Networks (tuba) 173
- 2.3 Network Management Area 177
  - 2.3.1 Bridge MIB (bridge) 191
  - 2.3.2 Chassis MIB (chassis) 193
  - 2.3.3 DS1/DS3 MIB (trunkmib) 199

2.3.4 Ethernet MIB (ethermib) 201
2.3.5 FDDI MIB (fddimib) 203
2.3.6 Host Resources MIB (hostmib) 207
2.3.7 IEEE 802.3 Hub MIB (hubmib) 213
2.3.8 Internet Accounting (acct) 217
2.3.9 SNMP Version 2 (snmpv2) 227
2.3.10 Token Ring Remote Monitoring (trmon) 245
2.3.11 Uninterruptible Power Supply (upsmib) 249
2.3.12 X.25 Management Information Base (x25mib) 253
2.4 OSI Integration Area 255
2.4.1 MHS-DS (mhsds) 265
2.4.2 MIME-MHS Interworking (mimemhs) 273
2.4.3 Network OSI Operations (noop) 275
2.4.4 OSI Directory Services (osids) 279
2.4.5 Office Document Architecture (oda) 283
2.4.6 SNMP over a Multi-protocol Internet (mpsnmp) 285
2.4.7 X.400 Operations (x400ops) 287
2.5 Operational Requirements Area 293
2.5.1 BGP Deployment and Application (bgpdepl) 301
2.5.2 Benchmarking Methodology (bmwg) 307
2.5.3 Network Joint Management (njm) 311
2.5.4 Operational Statistics (opstat) 325
2.5.5 User Connectivity (ucp) 329
2.6 Routing Area 333
2.6.1 Border Gateway Protocol (bgp) 349
2.6.2 IP over Large Public Data Networks (iplpdn) 353
2.6.3 IP Routing for Wireless/Mobile Hosts (mobileip) 359
2.6.4 ISIS for IP Internets (isis) 363
2.6.5 Inter-Domain Policy Routing (idpr) 365
2.6.6 Multicast Extensions to OSPF (mospf) 369
2.6.7 OSI IDRP for IP over IP (ipidrp) 371
2.6.8 Open Shortest Path First IGP (ospf) 381
2.6.9 RIP Version II (ripv2) 389
2.7 Security Area 391
2.7.1 Commercial Internet Protocol Security Option (cipso) 405
2.7.2 Common Authentication Technology (cat) 407
2.7.3 Internet Protocol Security Protocol (ipsec) 413
2.7.4 Network Access Server Requirements (nasreq) 417
2.7.5 Privacy-Enhanced Electronic Mail (pem) 425
2.7.6 SNMP Security (snmpsec) 431
2.7.7 TCP Client Identity Protocol (ident) 437
2.8 Transport and Services Area 439

2.8.1 Audio/Video Transport (avt) 441
2.8.2 Distributed File Systems (dfs) 449
2.8.3 Domain Name System (dns) 451
2.8.4 Service Location Protocol (svrloc) 455
2.8.5 TCP Large Windows (tcplw) 459
2.8.6 Trusted Network File Systems (tnfs) 461
2.9 User Services Area 463
2.9.1 Directory Information Services Infrastructure (disi) 469
2.9.2 Integration of Internet Information Resources (iiir) 473
2.9.3 Internet Anonymous FTP Archives (iafa) 477
2.9.4 Internet School Networking (isn) 483
2.9.5 Internet User Glossary (userglos) 485
2.9.6 NOC-Tool Catalogue Revisions (noctool2) 487
2.9.7 Network Information Services Infrastructure (nisi) 489
2.9.8 Networked Information Retrieval (nir) 495
2.9.9 Uniform Resource Identifiers (uri) 499
2.9.10 User Documents Revisions (userdoc2) 505
2.9.11 User Services (uswg) 509
2.9.12 Whois and Network Information Lookup Service (wnils) . . . 513

3 Network Status Briefings **519**
3.1 EBONE 521
3.2 NSI 525

4 IETF Protocol Presentations **529**
4.1 PIP 531
4.2 Simple Internet Protocol 537
4.3 IPAE 541

5 Technical Presentations **547**
5.1 Report form the POISED Working Group 549
5.2 Connection-Oriented Internet Protocols 553
5.3 Source Demand Routing Protocol Specification (Version 1) 561
5.4 Export Controls on Cryptographic Software 567
5.5 IPX Protocol Suite 571
5.6 Internet Accounting 585

6 Process for Organization of Internet Standards **591**
6.0.1 Process for Organization of Internet Standards (poised) 593

A Attendees **601**

Director's Message

The 25th meeting of the IETF was held in Washington, D.C., from November 16-20, 1992. This meeting was hosted by Sprint, and our thanks and appreciation go out to Gary Wightman, Joe Apple, Marty Schulman, Brian Shiflett, and all the others that helped with the social event and the terminal room. As was obvious to those who arrived on Sunday, there was another group already at the Hyatt, and access to the terminal room was not possible until well after 6:00 p.m. Sunday evening. The folks from Sprint worked around the clock in an effort to set up and configure all the work stations, routers, printers, and the local area network.

While the number of registrations received before the meeting was higher than those received for Cambridge, the number of actual attendees dropped slightly from 677 to 634... still a substantial number. Interestingly, though the number of attendees was less, the number of first time attendees was slightly higher; 214 first timers in Cambridge, 225 first timers in Washington). And the first timers are not one-timers. Of the 214 first timers in Cambridge, 114 attended the IETF meeting in Washington.

Responding to an almost unanimous request from the IETF, a second "announcement-only" IETF mailing list was created for the sole purpose of disseminating information to the IETF (IETF meeting announcements and logistics, Internet-Draft and RFC announcements, etc.). The IETF mailing list (ietf@cnri.reston.va.us) remains the primary unmoderated discussion list.

Following the Cambridge meeting, two efforts were undertaken to address the needs of those attending their first IETF meeting. Working with the Secretariat, Gary Malkin assembled, consolidated, and prepared the "Tao of the IETF, a Guide for New Attendees of the Internet Engineering Task Force." This document went through numerous edits and revisions and was made available as an Internet-Draft prior to the Washington IETF meeting.

The second effort was to prepare a presentation for first time attendees to answer frequently asked questions and provide information on the IETF and IETF meetings. A Newcomer's Orientation, conducted by the IETF Secretariat, was held Sunday afternoon before the registration reception. The orientation was well attended, though not entirely by newcomers...a number of veteran IETF meeting attendees were noticed in the audience. Based on the favorable comments received by those attending the first presentation, particularly from the first timers, the Newcomer's Orientation will be repeated at future meetings.

I am pleased to report that Stev Knowles has joined the IESG, and will serve as a Co-Director of the Internet Area. Stev is filling the slot left open when Noel Chiappa resigned. The Internet Area of the IETF is an active one, and the addition of Stev will provide support to that Area's activities.

There was a feeling of anticipation as folks arrived at the 25th IETF meeting. Since the meeting at Cambridge in July, a great deal of discussion was evident on the mailing lists. One thread focused on IPv7, often referred to as the next generation of IP addressing. The IESG had issued a call for proposals prior to the 24th meeting in Cambridge. Which way would the IETF go? More importantly, what were the options from which to choose? Two new proposals were presented at the Cambridge meeting, and another alternative was being prepared. Finally, who will make the choice?

The Road to IPv7

The IETF meeting began Monday morning with technical presentations on each of the alternatives being submitted to address the problems of growth within the Internet. Presentations made to the IETF during this session were: "TCP with Bigger Addresses" (TUBA), "The 'P' Internet Protocol" (PIP), "Simple Internet Protocol" (SIP), and "IP Address Encapsulation" (IPAE). In addition to the technical presentations, subsequent Working Group meetings were held during the week.

Prior to the Washington IETF meeting, RFC 1380, "IESG Deliberations on Routing and Addressing," was published which included an initial set of criteria against which the various submissions would be evaluated. A second set of criteria were proposed in an Internet-Draft submitted just before the Washington meeting. During the week, a Selection Criteria BOF was convened in an attempt to reach consensus on a single set of criteria for evaluating the proposals, and to solicit suggestions, comments, and concerns from the IETF at large, particularly from network operators.

As is to be expected, discussions will continue, at IETF meetings and on the various mailing lists, both on the proposals themselves and on the selection criteria. Many of the finest minds in the IETF are working on or examining the proposals, and a number of implementations are expected to be running in the Internet during 1993.

POISED Deliberations

The POISED Working Group was created "... to examine the Internet standards process and the responsibilities of the IAB, with attention to the relationship between the IAB and the IETF/IESG." During the week of the Washington meeting, there was a presentation made to the entire IETF, working group meetings, accompanied by a number of continuing discussions. This was the single topic of discussion during the Thursday evening IESG Open Plenary. Due to the special nature of this topic, a special section is included in these Proceedings which focus on the items discussed during the Open Plenary. Working Group minutes and the overheads from the presentation will be included in other sections of the Proceedings.

Multicasting Continues

Once again, the IETF had the valuable services of Steve Casner, Steve Deering, and a host of others who provided the audio- and video-cast from the IETF throughout the week. Prior to the meeting, the "IETF TV Guide" was distributed with the schedule of sessions to be covered on "IETF Channel 1" or "IETF Channel 2." Each channel had three concurrent multicast streams (GSM audio encoding, PCM audio encoding, and video). I don't know how they do it, but the broadcast efforts and capabilities seem to be better at every subsequent IETF meeting.

One final note: at the conclusion of the IESG Open Plenary on Thursday evening, the session ended with the attendees singing Happy Birthday to Megan Davies, the meeting coordinator. I believe this may be the first live broadcast of a musical performance over the Internet!

Future Meetings

The next plenary meeting of the IETF will be held in Columbus, Ohio from March 28 - April 2, 1993 (yes, the IESG Open Plenary will be on April First). This meeting will be co-hosted by OARNet and The Ohio State University.

The 27th meeting of the IETF is being held in Amsterdam from July 12-16, 1993, co-hosted by RARE and SURFnet. This is the first meeting of the IETF to be held outside of North America. Negotiations are still in progress for the meeting facilities and catering arrangements, and there will be a higher fee for this meeting (estimated to be at least $200). More information will be provided as it becomes available.

Beginning in 1993, the IETF Secretariat is assuming responsibility for choosing meeting sites. At some point in the future we will be soliciting volunteers to host IETF meetings, primarily to assist with the terminal room and to serve as a local contact point.

Stephen J. Coya
Executive Director, IETF

IETF Progress Report

Between the IETF meetings in Cambridge and Washington, there were seven new Working Groups created:

1. Uninterruptible Power Supply (upsmib)
2. Networked Information Retrieval (nir)
3. TCP/UDP over CLNP-addressed Networks (tuba)
4. Integration of Internet Information Resources (iiir)
5. Process for Organization of Internet Standards (poised)
6. SNMP Version 2 (snmpv2)
7. Uniform Resource Identifiers (uri)

and four Working Groups that were concluded:

1. Connection IP (cip)
2. Network Fax (netfax)
3. Distributed Scheduling Protocol (chronos)
4. Automated Internet Mailing List Services (list)

Additionally, there were twenty-nine RFCs published since the Cambridge IETF meeting in July, 1992:

RFC	Status	Title
RFC1334	PS	PPP Authentication Protocols
RFC1355	I	Privacy and Accuracy Issues in Network Information Center Databases
RFC1356	PS	Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode
RFC1358	I	Charter of the Internet Architecture Board (IAB)
RFC1359	I	Connecting to the Internet What Connecting Institutions Should Anticipat
RFC1360	S	IAB OFFICIAL PROTOCOL STANDARDS
RFC1361	I	Simple Network Time Protocol (SNTP)
RFC1362	I	Novell IPX Over Various WAN Media (IPXWAN)
RFC1363	E	A Proposed Flow Specification
RFC1364	PS	BGP OSPF Interaction
RFC1365	I	An IP Address Extension Proposal
RFC1366	I	Guidelines for Management of IP Address Space
RFC1367	I	Schedule for IP Address Space Management Guidelines
RFC1368	PS	Definitions of Managed Objects for IEEE 802.3 Repeater Devices
RFC1369	I	Implementation Notes and Experience for The Internet Ethernet MIB
RFC1370	PS	Applicability Statement for OSPF
RFC1371	I	Choosing a "Common IGP" for the IP Internet (The IESG's Recommendation to the IAB)
RFC1372	PS	Telnet Remote Flow Control Option

RFC1373	I	PORTABLE DUAs
RFC1374	PS	IP and ARP on HIPPI
RFC1375	I	Suggestion for New Classes of IP Addresses
RFC1376	PS	The PPP DECnet Phase IV Control Protocol (DNCP)
RFC1377	PS	The PPP OSI Network Layer Control Protocol (OSINLCP)
RFC1378	PS	The PPP AppleTalk Control Protocol (ATCP)
RFC1379	I	Extending TCP for Transactions – Concepts
RFC1380	I	IESG Deliberations on Routing and Addressing
RFC1381	PS	SNMP MIB Extension for X.25 LAPB
RFC1382	PS	SNMP MIB Extension for the X.25 Packet Layer
RFC1385	I	EIP: The Extended Internet Protocol A Framework for M Backward Compatibility

POISED WG Discussion

- Change!
- Selection and accountability
  - Marshall Rose
- Modification of process
  - Lyman Chapin
- Internet Society actions
- Next steps
- Questions and Answers

POISED WG Discussion

Change!

- Selection and accountability are essential
- Change is needed soon; credibility at stake
- Process is also at issue
  - Need to push work down the pyramid
- Review process must be coupled closely to the working groups
- Standards process must be streamlined
- Increased attention to quality

POISED WG Discussion

Selection and Accountability

- Fixed terms for management positions
- Selection process visible
- Active IETF involvement
- Accountability essential
- Put process in place ASAP!

POISED WG Discussion

Process changes

- Consensus reached!
- Move review process and standards decision from IAB to IESG
- Beef up process

Poised Issues

⇒ Delay

⇒ Overload

⇒ Accountability

Changes

- WGs and design teams
- IESG makes standards progression decisions
  - Quality control at WG level
  - Area director(ate) and IESG review and approve WG documents for progression
  - IAB involved only for disputes
- Area directorates

Changes - 2

- IAB roles
  - Arbitrate disputes and appeals
  - Oversight of Internet architecture
  - Review WG charters

    ("front-end" contribution to architectural consistency.)
  - Initiatives (proposals and strawmen)

POISED WG Discussion

Next steps

- Rough consensus
  - Take pulse here
  - Document results; publish on POISED WG list
- Running code
  - Transfer of process responsibility to IESG
  - Identify slots to be filled
  - Select Recruiting and Nomination Team
- Standardize
  - Refine all details
  - Bore the hell out of IETF with repeated reports

POISED WG Discussion

Internet Society actions

- Presentation to ISoc trustees December 10, 1992

  Consensus here => formal request to trustees
- Charter for IAB to be replaced
- ISoc trustees appoints Ombudsman

Agenda of the Twenty-Fifth IETF

(November 16-20, 1992)

MONDAY, November 16, 1992

8:00-9:00 am	IETF Registration and Continental Breakfast
9:00-9:30 am	Introductions
9:30-12:00 noon	Technical Presentations

- Pip Internet Protocol (pip) (Paul Tsuchiya/Bellcore)
- TCP/UDP over CLNP-addressed Networks (tuba) (Peter Ford/LANL and Mark Knopper/Merit)
- Simple Internet Protocol (sip) (Steve Deering/Xerox PARC)
- IP Address Encapsulation (ipae) (Bob Hinden/SUN)

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

APP	Remote Conferencing BOF (remconf) (Jack Drescher/MCNC and Ari Ollikainen/LLNL)
INT	IP over Appletalk WG (appleip) (John Veizades/Apple)
INT	IP over ATM WG (atm) (Bob Hinden/Sun)
MGT	SNMP Version 2 WG (snmpv2) (Bob Stewart/Xyplex)
OPS	Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)
OSI	OSI Directory Services WG (osids) (Steve Hardcastle-Kille/ISODE)
RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)
RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
TSV	Domain Name System WG (dns) (Mike Reilly/DEC)
USV	Network Information Services Infrastructure WG (nisi) (April Marine/SRI and Pat Smith/Merit)

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

INT	Pip Internet Protocol WG (pip) (Paul Tsuchiya/Bellcore)

4:00-6:00 pm	Monday, November 16, 1992 - Afternoon Sessions II (cont'd.)	
	APP	Conferencing Control BOF (confctrl) (Eve Schooler/ISI)
	INT	IP over Appletalk WG (appleip) (John Veizades/Apple)
	OPS	Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)
	OSI	OSI Directory Services WG (osids) (Steve Hardcastle-Kille/ISODE)
	RTG	Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)
	RTG	Inter-Domain Policy Routing WG (idpr) (Martha Steenstrup/BBN)
	RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
	USV	Internet Anonymous FTP Archives WG (iafa) (Peter Deutsch/Bunyip and Alan Emtage/Bunyip)
	SEC	Security Area Advisory Group (saag) (Stephen Crocker/TIS)

TUESDAY, November 17, 1992

8:30-9:00 am	Continental Breakfast
9:00-9:30 am	IETF Technical Presentations

- "Report from the POISED Working Group" (Steve Crocker/TIS)

9:30-12:00 noon — Morning Sessions

APP	Conferencing Control BOF (confctrl) (Eve Schooler/ISI)
APP	Internet SMTP Extensions WG (smtpext) (John Klensin/MIT)
APP	Telnet WG (telnet) (Steve Alexander/INTERACTIVE Systems)
MGT	IEEE 802.3 Hub MIB WG (hubmib) (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
OPS	Network Status Reports (netstat) (Gene Hastings/PSC)
OSI	MHS-DS WG (mhsds) (Kevin Jordan/CDC and Harald Alvestrand/SINTEF DELAB)
RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
RTG	IP Routing for Wireless/Mobile Hosts WG (mobileip) (Steve Deering/Xerox PARC)
RTG	OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)
SEC	IP Security BOF (ipsec) (Steve Crocker/TIS)
USV	User Services WG (uswg) (Joyce Reynolds/ISI)

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

INT	IP over ATM WG (atm) (Bob Hinden/Sun)
MGT	Chassis MIB WG (chassis) (Jeff Case/UTenn and Bob Stewart/Xyplex)
OSI	MHS-DS WG (mhsds) (Kevin Jordan/CDC and Harald Alvestrand/SINTEF DELAB)
RTG	New Internet Routing and Addressing Architecture BOF (nimrod) (Noel Chiappa)
RTG	Open Shortest Path First IGP WG (ospf) (John Moy/Proteon)

1:30-3:30 pm	Tuesday, November 17, 1992 - Afternoon Sessions I (cont'd.)	
	SEC	Common Authentication Technology WG (cat) (John Linn/DEC)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
	TSV	Service Location Protocol WG (svrloc) (John Veizades/Apple)
	USV	Internet User Glossary WG (userglos) (Tracy LaQuey Parker/UTexas and Gary Malkin/Xylogics)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	GEN	Process for Organization of Internet Standards BOF (poised) (Steve Crocker/TIS)
	MGT	Host Resources MIB WG (hostmib) (Steve Waldbusser/CMU)
	OPS	User Connectivity Problems WG (ucp) (Dan Long/BBN)
	RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
	SEC	SNMP Security WG (snmpsec) (James Galvin/TIS and Keith McCloghrie/Hughes)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
	USV	User Documents WG (userdoc2) (Ellen Hoffman/UMich and Lenore Jackson/NASA)
7:30-10:00 pm	Tuesday, November 17, 1992 - Evening Sessions	
	INT	Inter-domain Multicast Routing BOF (idmr) (Tony Ballardie/UCL)
	INT	IP over Fibre Channel BOF (fibreip) (Yakov Rekhter/IBM and Lansing Sloan/LLNL)
	INT	IP over Appletalk WG (appleip) (John Veizades/Apple)
	MGT	Host Resources MIB WG (hostmib) (Steve Waldbusser/CMU)
	OSI	Shared Whois Project BOF (whois) (Sheri Repucci/Merit)

7:30-10:00 pm	Tuesday, November 17, 1992 - Evening Sessions (cont'd.)	
	OSI	X.400 Operations WG (x400ops) (Alf Hansen/SINTEF DELAB)
	RTG	Virtual Circuit Routing BOF (vcrout) (Rob Coltun/Consultant and Marco Sosa/Bellcore)
	SEC	TCP Client Identity Protocol WG (ident) (Mike St. Johns/DOD)
	USV	Training Materials BOF (trainmat) Ellen Hoffman/Merit and Jill Foster/UNewcastle-Upon-Tyne)

WEDNESDAY, November 18, 1992

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	• "An Implementation Model for Connection-Oriented Internet Protocols" (Chuck Cranor/WashU)	
9:30-12:00 noon	Morning Sessions	
	APP	Network News Transport Protocol WG (nntp) (Eliot Lear/Silicon Graphics)
	INT	Point-to-Point Protocol Extensions WG (pppext) (Brian Lloyd/Consultant)
	INT	TCP/UDP over CLNP-addressed Networks WG (tuba) (Peter Ford/LANL and Mark Knopper/Merit) *
	MGT	SNMP Version 2 WG (snmpv2) (Bob Stewart/Xyplex)
	OSI	X.400 Operations WG (x400ops) (Alf Hansen/SINTEF DELAB)
	RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
	RTG	IP Routing for Wireless/Mobile Hosts WG (mobileip) (Steve Deering/Xerox PARC)
	RTG	OSI IDRP for IP over IP WG (ipidrp)* (Sue Hares/Merit)
	USV	Internet School Networking WG (isn) (John Clement/EDUCOM, Connie Stout/TheNet and Art St. George/UNM)
	USV	Networked Information Retrieval WG (nir) (Jill Foster/UNewcastle-Upon-Tyne and George Brett/MCNC)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	Remote Mail Protocol BOF (remmail) (Mark Smith/UMICH)
	INT	IP over ATM WG (atm) (Bob Hinden/Sun)
	MGT	SNMP Version 2 WG (snmpv2) (Bob Stewart/Xyplex)
	OPS	Benchmarking Methodology WG (bmwg) (Scott Bradner/Harvard)

1:30-3:30 pm	Wednesday, November 18, 1992 - Afternoon Sessions I (cont'd.)	
	OSI	Network OSI Operations WG (noop) (Sue Hares/Merit and Cathy Wittbrodt/LLNL)
	RTG	Multicast Extensions to OSPF WG (mospf) (Steve Deering/Xerox PARC)
	RTG	New Internet Routing and Addressing Architecture BOF (nimrod) (Noel Chiappa)
	SEC	Network Access Server Requirements WG (nasreq) (Allan Rubens/Merit)
	USV	Uniform Resource Identifiers WG (uri) (Peter Deutsch/Bunyip and Alan Emtage/Bunyip)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	INT	Dynamic Host Configuration WG (dhc) (Ralph Droms/Bucknell)
	INT	Selection Criteria BOF (select) (Philip Almquist/Consultant)
	MGT	SNMP Version 2 WG (snmpv2) (Bob Stewart/Xyplex)
	OSI	FTP-FTAM Gateway BOF (ftpftam) (Joshua Mindel/Open Networks)
	OSI	Network OSI Operations WG (noop) (Sue Hares/Merit and Cathy Wittbrodt/LLNL)
	RTG	IP over Large Public Data Networks WG (iplpdn) (George Clapp/Ameritech)
	SEC	Privacy-Enhanced Electronic Mail WG (pem) (Steve Kent/BBN)
	TSV	Audio/Video Transport WG (avt) (Stephen Casner/ISI)
	USV	WHOIS and Network Information Lookup Service WG (wnils) (Joan Gargano/UCDavis)

* Joint session of IPIDRP and TUBA

7:30-10:00 pm	Wednesday, November 18, 1992 - Evening Session	
	INT	Simple Internet Protocol BOF (sip) (Steve Deering/Xerox PARC)
	INT	Traceroute BOF (tracerte) (Gary Malkin/Xylogics)
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	OPS	BGP Deployment and Application WG (bgpdepl) (Matt Mathis/PSC)
	OSI	Xwindows over OSI and Skinny Stack OSI BOF (thinosi) (Dave Piscitello/Bellcore)
	SEC	SNMP Security WG (snmpsec) (James Galvin/TIS and Keith McCloghrie/Hughes)
	TSV	DNS II BOF (dns2) (Paul Mockepetris/DARPA)
	USV	Integration of Internet Information Resources WG (iiir) (Chris Weider/Merit)

THURSDAY, November 19, 1992

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	• "Source Demand Routing Protocol Specification (Version 1)" (Deborah Estrin/USC, Tony Li/cisco and Yakov Rekhter/IBM)	
9:30-12:00 noon	Morning Sessions	
	APP	NAPLPS as a MIME Extension BOF (napmime) (David Hughes/OCC and George Johnston/MIT)
	APP	Network Database WG (netdata) (Scott Newman/DEC)
	INT	IP Address Encapsulation WG (ipae) (Dave Crocker/TBO)
	MGT	FDDI MIB WG (fddimib) (Jeff Case/UTenn)
	MGT	Token Ring Remote Monitoring WG (trmon) (Mike Erlinger/Lexcel)
	OPS	Network Joint Management WG (njm) (Gene Hastings/PSC)
	OSI	SNMP Application Monitoring BOF (sam) (Steve Hardcastle-Kille/UCL)
	SEC	IP Security BOF (ipsec) (Steve Crocker/TIS)
	USV	Directory Information Services Infrastructure WG (disi) (Chris Weider/Merit)
Breaks	Coffee available throughout the morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	Internet SMTP Extensions WG (smtpext) (John Klensin/MIT)
	INT	Dynamic Host Configuration WG (dhc) (Ralph Droms/Bucknell)
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	MGT	Token Ring Remote Monitoring WG (trmon) (Mike Erlinger/Lexcel)
	OPS	Operational Area Directorate (orad) (Phill Gross/ANS and Bernhard Stockman/SUNET)

1:30-3:30 pm	Thursday, November 19, 1992 - Afternoon Sessions I (cont'd.)	
	RTG	Source Demand Routing Protocol BOF (sdrp) (Deborah Estrin/USC and Tony Li/cisco)
	SEC	Security Area Advisory Group (saag) (Steve Crocker/TIS)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Technical Presentations	

- "Export Controls on Cryptographic Software" (Jerry Rainville)
- "Other Protocols in the Internet: The IPX Protocol" (Paul Turner/Novell)
- "Internet Accounting" (Cyndi Mills/BBN)

7:30-10:00 pm Open Plenary and IESG

FRIDAY, November 20, 1992

8:30-9:00 am	Continental Breakfast	
9:00-12:00 noon	Morning Sessions	
	MGT	IFIP Electronic Mail Management BOF (emailmgt) (Einar Stefferud/NMA and Paul Brusil/MITRE)
	MGT	Internet Accounting WG (acct) (Cyndi Mills/BBN and Gregory Ruth/BBN)
	MGT	SNMP Version 2 WG (snmpv2) (Bob Stewart/Xyplex)
	SEC	Authorization and Access Control BOF (aac) (Clifford Neuman/ISI)
	GEN	IAB Open Meeting

Key to Abbreviations

APP	Applications	Russ Hobby/UCDavis
GEN	General Interest	
INT	Internet	Philip Almquist
MGT	Network Management	James Davin/Bellcore
OSI	OSI Integration	Erik Huizer/SURFnet and David Piscitello/Bellcore
OPS	Operational Requirements	Bernhard Stockman/SUNET and Phill Gross/ANS
RTG	Routing	Bob Hinden/Sun
SEC	Security	Steve Crocker/TIS
TSV	Transport and Services	Dave Borman/Cray Research
USV	User Services	Joyce K. Reynolds/ISI

Chapter 1

IETF Overview

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

The IETF mission includes:

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

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

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

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

5. Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

Technical activity on any specific topic in the IETF is addressed within working groups. All working groups are organized roughly by function into nine technical areas. Each is led by an Area Director who has primary responsibility for that one area of IETF activity.

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

At the time of the 25th IETF meeting, the current Areas and Directors, which composed the IESG, were:

IETF and IESG Chair:	Phill Gross/ANS
Applications:	Russ Hobby/UC-Davis
Internet:	Philip Almquist/Consultant
	Stev Knowles/ FTP Software
Network Management:	James Davin/ Bellcore
OSI Integration:	Dave Piscitello/Bellcore
	Erik Huizer/SURFnet
Operational Requirements:	Phill Gross/ANS
	Bernhard Stockman/SUNET
Routing:	Robert Hinden/Sun
Security:	Steve Crocker/TIS
Transport and Services:	David Borman/Cray Research
User Services:	Joyce K. Reynolds/ISI
Standards Management:	Dave Crocker/TBO

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

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

The working groups conduct business during plenary meetings of the IETF, during meetings outside of the IETF, and via electronic mail on mailing lists established for each group. The IETF holds 4.5 day plenary sessions three times a year. These plenary sessions are composed of Working Group Sessions, Technical Presentations, Network Status Reports, working group reporting, and an open IESG meeting. A Proceedings of each IETF plenary is published, which includes reports from each Area, each working group, and each Technical Presentation. The Proceedings include a summary of all current standardization activities.

Meeting reports, Charters (which include the working group mailing lists), and general information on current IETF activities are available on-line for anonymous FTP from several Internet hosts including nnsc.nsf.net.

Mailing Lists

Much of the daily work of the IETF is conducted on electronic mailing lists. There are mailing lists for each of the working groups, as well as a general IETF list. Mail on the working group mailing lists is expected to be technically relevant to the working groups supported by that list.

To join a mailing list, send a request to the associated request list. All internet mailing lists have a companion "-request" list. Send requests to join a list to <listname>-request@<listhost>.

Information and logistics about upcoming meetings of the IETF are distributed on the general IETF mailing list. For general inquiries about the IETF, requests should be sent to `ietf-info@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 1993

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

Summer 1993

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

1.2 On Line IETF Information

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

Directory Locations

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

cd ietf
cd internet-drafts

Individual files can then be retrieved using the GET command:

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

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

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

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

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

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

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

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

Europe Address: nic.nordu.net (192.36.148.17)

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

1.2.1 The IETF Directory

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

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

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

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

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

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

1.2.2 The Internet-Drafts Directory

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

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

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

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

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

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

where <author> is the author's name.

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

1.3 Guidelines to Authors of Internet-Drafts

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

Internet-Drafts are not an archival document series. These documents should not be cited or quoted from in any formal document. Unrevised documents placed in the Internet-Drafts Directories have a maximum life of six months. After that time, they must be submitted to the IESG or the RFC editor, or they will be deleted. After a document becomes an RFC, it will be replaced in the Internet-Drafts Directories with an announcement to that effect for an additional six months.

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

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

There are differences between the RFC and Internet-Draft format. The Internet-Drafts are NOT RFC's and are NOT a numbered document series. The words "INTERNET-DRAFT" should appear in the upper left hand corner of the first page. The document should NOT refer to itself as an RFC or a Draft RFC.

The Internet-Draft should neither state nor imply that it is a Proposed Standard. To do so conflicts with the role of the IAB, the RFC Editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the Internet-Draft. All Internet-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 nic.ddn.mil, nnsc.nsf.net, nic.nordu.net, ftp.nisc.sri.com, or munnari.oz.au.

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

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

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

Chapter 2

Area and Working Group Reports

2.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UC Davis

A new goal in the Applications Area is to move toward working groups being unified by guiding architectures. Toward this goal the start of two architectures have been defined. The first is an architecture to define workstation based teleconferencing. The second is a joint effort between the Applications Area and the User Services Area to create an Internet Information Architecture to define a system of protocols to allow support information organization, searching and retrieval.

Conference Control BOF (CONFCTRL)

An impromptu BOF on Conference Control (sometimes referred to as connection or configuration management) was held. Discussions were to understand how such a group might contribute to the remote conferencing architecture effort. It was agreed that there is a need for a a session layer control protocol to perform higher layer functions than the protocol proposed in the AVT Working Group. The beginnings of design criteria for this protocol were identified by determining which functions must be supported. Discussion also focused on the range and capabilities of various session types needing support, the list of outside services to which the protocol will interface, and short-term versus long-term functionality considerations.

NAPLPS Graphics and Character Sets as a MIME BOF (NAPMIME)

This BOF explored interest in the definition of a NAPLPS body part for MIME. There was a demonstration of an NAPLPS system showing how presentation graphics can be transmitted using low bandwidths.

Remote Conferencing BOF (REMCONF)

The Remote Conferencing BOF discussed an architecture for all aspects of workstation based teleconferencing. This includes things like video, audio, shared windows, session setup and management. A separate group was spawned off to focus on session configuration and management. This group will become a working group to continue guidance on the architecture.

Remote Mail Protocol BOF (REMMAIL)

The Remote Mail BOF discussed methods for end-user mail delivery and problems with current protocols such as POP and IMAP. The Group reached consensus on two areas of work for a possible working group. First is to standardize a protocol for central mail repository to work with diskless clients. The second is the email support of laptops and other disconnected machines. Discussion of a working group Charter will be done on the ietf-remmail@umich.edu mailing list.

SMTP Extensions Working Group (SMTPEXT)

The SMTP Extensions Working Group came to closure on a set of documents that answers the concerns brought up from the Last Call of the previous documents. These new documents will soon be submitted by the Working Group for approval to be a Proposed Standard.

Network Database Working Group (NETDATA)

The Network Database Working Group discussed the proposal from SQL Access for doing OSI's RDA directly on a TCP/IP stack. Security was the main technical concern of the Group. However, a more significant hurtle may be the logistical and legal one of being able to put the ISO and X/Open specification on line to create a complete description of the overall protocol.

Network News Transport Protocol Working Group (NNTP)

The NNTP Working Group finished up work on the NNTP v2 document and went on to discuss the requirements for a Network News Reader Protocol (NNRP) that would serve between a news repository and a user agent. Questions came up about how NNRP will relate to mail protocols, how authentication can be done, how to do search mechanisms, and whether NNRP should be an extension of NNTP or be developed independently.

Telnet Working Group (TELNET)

The Telnet Working Group continued the work on authentication and encryption for Telnet sessions.

Internet Information Architecture

The Internet Information Architecture is a start to define a system of protocols to support of information organization, searching and retrieval. Four working groups have been created to address several parts of the overall goal. These working groups are:

- **Networked Information Retrieval Working Group (NIR)**: is cataloging the types of information and information services that currently exist. This defines the starting point for work on the overall architecture.

- **Universal Resource Identifiers Working Group (URI)**: is looking at ways to have unique identifiers for information objects on the Internet. This will allow a person to know that they have found a particular object regardless of how the object is named locally.

- **Integration of Internet Information Resources Working Group (IIIR)**: is looking at the various information search and retrieval protocols, such as Archie, Gopher, WAIS and others, and working toward a common protocol or set of protocols to standardize these functions.

- **Whois and Network Information Lookup Service Working Group (WNILS)**: is looking at how to organize directory information that already exists in various WHOIS servers.

CURRENT MEETING REPORT

Reported by Eve Schooler/ISI and Dean Blackketter/Apple

Minutes of the Conferencing Control BOF (CONFCTRL)

One task of the initial BOF sessions was actually to find a suitable definition for "conference control", since the topic has been bandied about for some time in the Remote Conferencing BOF and the Audio/Video Transport Working Group. By broadly defining multimedia conferencing as collaborations in two dimensions (members and media), the Group was able to define conference control as the management and coordination of (multiple) conference members in (multiple) media.

How does conference control pertain to the ongoing RemConf efforts for an overall remote conferencing architecture, and in particular to the developments in the AVT Working Group of a real-time transport protocol? The Group agreed that there is a need for a session layer control protocol to perform higher layer functions than the protocol proposed in the AVT Working Group. For example, three aspects of conference control might include session, connection and configuration management; session management entails who is involved in a conference, connection management involves the topology of who is seeing whom in each media, and configuration management is the negotiation of differences in end-system capabilities.

The Group identified the beginnings of some design criteria for this protocol. First, it should be kept simple, yet extensible. The Group would like for it to accommodate a range of session styles – beyond the unmoderated sessions already available through vat, dvc, nv et al. It was also recognized that there was a need to separate short-term from long-term functionality goals.

The Group brainstormed about which functions MUST be supported versus those which the Group would like to have supported. It falls out of our definition for conference control that, at minimum, support is needed for both membership and media control. Membership control might include admission policies (such as user identification, user payment, meeting sponsorship), whereas media control might encompass capability descriptions, synchronization policies, and floor control (media focus). In both dimensions, session setup, maintenance and/or modification must be supported.

Other features deemed important but probably of lower priority included security (in the form of authentication and encryption), as well as feedback channels for bandwidth balancing. The Group also listed outside services to which it might expect a conference control protocol to interface: a suite of directory services for cataloguing users, conferences, and shared devices; bandwidth allocation and reservation mechanisms; and a scheme for multicast address allocation. The assumption is that eventually these outside services will be available.

To understand the range of capabilities to support in a conference control protocol, the Group explored the types of sessions that might arise. The wishlist included a continuum of session scenarios (although the picture below only lists a sample from the full range and only crudely approximates an ordering). "Secure" variations on these meetings were also discussed.

```
       impromptu
        hallway
        meetings        classroom        seminar     pay-per-view

  |-------|-------|--------|-------|-------|-------|-------|-------|

pt2pt       arch design        panel          lecture           TV
phone         review/       discussion/                      broadcast
 call     ''quilting bee''  presidential debate
```

Observations made about the spectrum were that there are different types of participation (active and passive), that there are gradations of identification policies (known vs. anonymous participants), that there may be extreme variations in the degree of interconnectivity among participants, etc.

The Group discussed that for simplicity's (and implementation's) sake, there is likely to be a need to select a small number of session types that the protocol should support. A rough breakdown into four general session models was presented:

1. Point-to-point calls.
2. Small, tightly-controlled sessions: N-way interconnectivity.
3. Medium-sized, loosely-controlled sessions: lighter-weight model.
4. Very large, fixed sessions: unidirectional broadcasts.

There was discussion that other standards bodies (CCITT) have explored issues in some aspects of connection control (for B-ISDN). In addition, existing prototype conferencing tools should be examined for leads on tradeoffs regarding conference management.

Attendees

Dean Blackketter	deanb@apple.com
Wo Chang	wchang@nist.gov
Osmund de Souza	osmund.desouza@att.com
Hans Eriksson	hans@sics.se
Don Hoffman	don.hoffman@eng.sun.com

Oliver Jones	oj@pictel.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Bill Manning	bmanning@sesqui.net
Kathleen Nichols	nichols@apple.com
Jim Perchik	perchik@athena.mit.edu
Eve Schooler	schooler@isi.edu
Henning Schulzrinne	hgs@research.att.com
Scott Stein	scotts@apple.com
Thierry Turletti	turletti@sophia.inria.fr
Abel Weinrib	abel@bellcore.com

The Impact of Scaling on a Multimedia Connection Architecture

Eve M. Schooler
USC/Information Sciences Institute
Marina del Rey, CA

November 13, 1992

Acknowledgments

This talk is a synthesis of ideas from many sources: Mainly conversations with colleagues at USC/ISI, discussions within the IETF A/V Transport working group and BOF on Remote Conferencing, and correspondences with colleagues at UMass/AT&T and MCNC, among others....

A Software Architecture for Packet Teleconferencing

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

Flow of Control Information

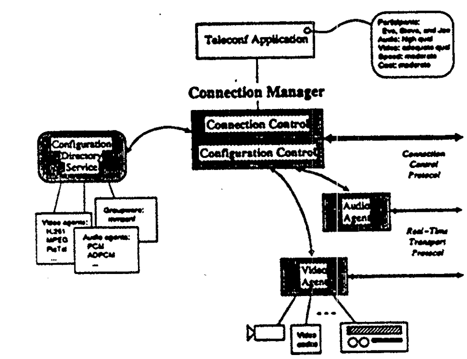

Large-scale Multimedia Teleconferencing

SCALING DIMENSIONS:

1. Very large numbers of participants per conference
2. Large and widely dispersed user population
3. Many concurrent teleconferences

- *The network infrastructure to support widescale packet teleconferencing is not in place.*

Current Teleconferencing Architectures: Axes of Scale

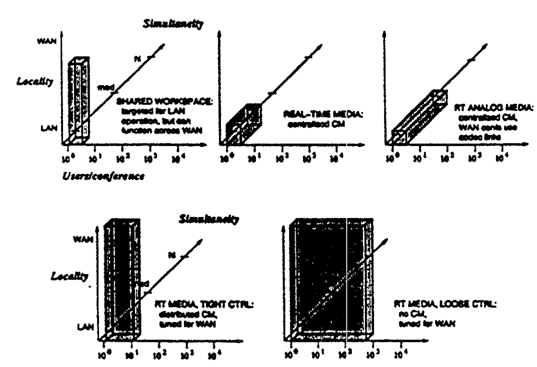

(MMConf)

(Etherphone, Phoenixphone, CAR, CoDesk)

(Touring Machine, Rapport)

(MMCC, DICE)

(IETF tools: vat, ivs, dvc)

The Problem of Scale

How do the requirements of a conferencing architecture and protocols change as we travel along each axis?

Scaling Up in Size of Conferences

- Impromptu Sessions: *ones or a few tens of individuals*
  - Full connectivity among users in all media
  - Flexible negotiation of conferencing parameters
  - Authentication of participants
  - Exchange of data encryption keys
  - Support for subconferencing
- Interactive Seminars: *hundreds of thousands of participants*
  - Too large for *N*–way sharing of either data or control
  - Impromptu feedback channels still needed
  - Support required for dynamic membership
  - Privacy becomes less practical to support

Scaling Up in Size of Conferences (continued)

- Unidirectional "Broadcasts": *10^5, 10^6 and beyond*
  - One–way dissemination of information
  - Sessions pre–arranged or permanent
  - Info for tapping into session might be static
    - *Might be built into receiver, as in TV receiver*

Large Numbers of Geographically Dispersed Users

- Single domain: the local area network
  - Fixed community of user names – homogeneous authentication
  - Often can assume similar configurations at users' systems
- Between domains: proxy agreements
- Inter–domain: WAN operation
  - Age–old issue: how to obtain a unique user address?
  - Heterogeneous end system configurations
  - Less assurances of robustness and timeliness
  - Movement away from centralized designs

Concurrency: Many Simultaneous Sessions

- This axis is not quite as straightforward...

  The number of concurrent sessions is uninteresting taken by itself
  - *What is interesting is how much the sessions overlap*
- Leads to competition for resources
  - *Mostly for bandwidth, but also for addrs, shared MM devices and users*
- Resource discovery needed to locate shared commodities
- Participation management for end–users
  - *Call waiting, forwarding, suspension, merging, subconferencing, etc.*

Implications of Scaling: Key Issues

I. Scalable conference session models

II. Multicast address management

III. Bandwidth reduction

IV. Codification of heterogeneity

V. A suite of directory services

A Scalable Session Model

- Tightly-controlled sessions:
  - A shared global view is actively maintained
    - *Often relies on N-way interconnectivity and reliable communication*
  - Support for flexible negotiation
    - *As N increases, convergence problematic*
- Loosely-controlled sessions:
  - Conference status constructed asynchronously
  - Control messages sent at regular intervals
    - *For large conferences, overhead for periodic communication*
  - Lighter-weight: no session coordination (with other end systems)
    - *No group interaction for authentication, QOS negotiation, etc.*
- Fixed sessions: for very large conferences
    - *Little to no setup, maintenance, or communication among conferees*

Scalable Session Protocol(s)

- Accommodate a range of conference sizes and modes
  - One adaptive session protocol?
  - A family of separate protocols for distinct circumstances?
- Outcome influenced by:
  - Trend for simplicity in Internet standards...
  - The specific crossover points between conferencing types
    - *Where might one protocol's usefulness end? and another's begin?*
  - The characteristics which differentiate models
    (e.g., interconnectivity, how global the session view, degree of dynamics)

Multicast Address Management

- Multicast delivery is instrumental in bandwidth reduction
    - *Fixed multicast address space with dynamic usage*
- Establish a hierarchy of multicast address servers [Schulzrinne]:
  - Partition addresses among all multicast address servers
  - Re-use and borrow addresses
- Integration into connection architecture
  - Per LAN multicast address server
  - Addresses assigned to individual media agents and/or connection managers
  - In private sessions, the connection manager distributes address(es)

Bandwidth Reduction

- Mechanisms needed for reductions at the receiver
    - *May only want or be able to process M of N streams sent*
- *Application-level combination nodes*
  - Functions for mixing, compositing, selection, translation, etc.
  - Likely to be separate from end systems in teleconference
  - Must incorporate into session management, addressing and routing
    - *Increased control complexity and transmission delay*
- *Resource Synthesizer* to assess tradeoff
    - *Works with Connection Manager and Configuration Directory Svc*
- Combination functions reflected in configuration language

Summary

- Presented current connection management architecture
    - *And discussed the limited scaling of most experimental systems*
- Described three critical scaling dimensions
    - *Users/conference, locality and concurrency*
- Identified key components and features for integration:
  - Scalable conference session model and protocol(s)
  - Multicast address management
  - Bandwidth reduction techniques
  - Codification of heterogeneity and of combination functions
  - A suite of directory services

A Scalable Architecture for Personal Teleconferencing

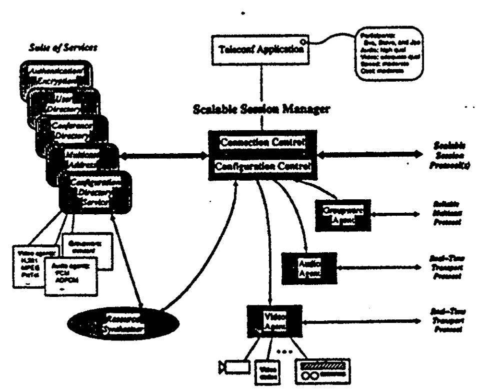

CURRENT MEETING REPORT

Reported by George L. Johnston/MIT

Minutes of the NAPLPS Graphics and Character Sets as a MIME BOF (NAPMIME)

The Chairs began the session by emphasizing the reality that the Internet is, and will continue to be, a highly heterogeneous network, in which some use will involve small systems with pc-level graphics and low-bandwidth connections. On the basis of this reality, they argued that a MIME extension which permits the encoding of pictorial information (including limited animation) and alphanumeric text (including limited animation) with great economy of file size would be very desirable. They asserted that NAPLPS meets these criteria. They introduced Mr. G. Kenneth Holman, Technical Vice President of Microstar Software Ltd., of Nepean, Ontario, Canada, one of the leading developers of NAPLPS software, including the NAPLPS drivers for Prodigy, a videotex service of IBM and Sears, to provide expert information on NAPLPS, including its relation to other international standards.

The Chairs and Mr. Holman described the history of NAPLPS and its status as an international standard. Particularly important is the fact that it is based on the ISO 2022 7 and 8 bit extension standard, which uses escape characters to select in-use tables from a repertoire of such tables. The standard has been extended to include audio and still compressed images (JPEG), as well as sixteen bit characters to represent languages which have such requirements. Dave Hughes demonstrated NAPLPS by means of TeleDraw, an integrated NAPLPS/ASCII terminal emulator, drawing program, and symbol processor (for the creation of DRCS – dynamically redefinable character set provided by the standard) for MS-DOS computers developed by his company.

One individual attending the session complained that he and, he believed, others attending the session felt that they were being sold something. George explained that standard was being introduced with which many were unfamiliar, as it comes from videotex and teletext. He said that it is important to demonstrate the functionality of the standard and that it has been widely implemented.

The same person felt that MIME involves a deliberate decision to avoid ISO 2022, in favor of fully formed characters that can be processed by the party receiving them, and that it favors a multi-part approach instead of pulling everything together, in one file, as NAPLPS does. In response, it was stated that the presentation level approach allows pictorial and character information to be placed in deliberate spatial relationship with each other.

BOF attendees provided helpful suggestions in response to the question of how NAPLPS might be related to MIME in terms of content type. It was stated that one can try to have MIME revised, or go to IANA for registration within an existing content type. The latter approach is clearly the path of least resistance. The content type **image**, with the subtype **naplps**, i.e., **image/naplps**, seems a natural choice for consideration. Less desirable would be **image/x-naplps**.

An alternative to content type **image** would be **application**, i.e., **application/naplps**. A person attending the session stated that a goal of MIME is to do unique labeling. The proponents of NAPLPS as a MIME extension should decide how they want it to be designated. He added that there is little interest in vector graphics. George Johnston reiterated that the Internet is an inhomogeneous network, with some low-end computers and connections, and therefore it is desirable to have a MIME extension which permits the economical encoding of pictorial and character information.

The session concluded with a statement by Ken Holman that he would begin to draft an application to IANA for registration of NAPLPS as content type **image/naplps**.

Attendees

Kay Chang	chang@chang.austin.ibm.com
Letha Dugas	4371362@mcimail.com
Erik Fair	fair@apple.com
Sallie Fellows	sallie%ed@psc.plymouth.edu
Ned Freed	ned@innosoft.com
Thomas Hacker	hacker@citi.umich.edu
Russ Hobby	rdhobby@ucdavis.edu
G. Ken Holman	holman@tmn.com
David Hughes	dave@oldcolo.com
George Johnston	glj@nerus.pfc.mit.edu
John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Charlotte Mooers	mooers@nnsc.nsf.net

CURRENT MEETING REPORT

Reported by Jack Drescher/MCNC and Ari Ollikainen/LLNL

Minutes of the Remote Conferencing BOF (REMCONF)

The goals for the November 16th Meeting were:

- Get a status report and perspective from Russ Hobby on proposals for how REMCONF might be more formally organized within the IETF.
- Review and get consensus on the Draft Multimedia Communication Architecture Paper by Yee-Hsiang Chang. This paper has been on rem-conf for sometime now and some good feedback has been received from a small set of people. A summary report of this feedback will be presented. Hopefully, this will stimulate additional feedback and suggestions on how to resolve some of the issues.
- Start the process of putting together a Catalog of Internet conferencing packages. MCNC has agreed to compile and periodically update the catalog. We'd like to agree on format, access and other logistics. We'd also like to identify initial candidates for the catalog. This can be a departure point for the interoperability discussions set for the Wednesday, November 18th AVT Working Group session.
- Address other issues as time permits.
- Provide a few information item handouts.

Results

I. An MCNC organizational recommendation summary was presented and discussed. *Editor's Note (md): This "summary" is available via ftp under remconf-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Key Discussion Points

1. Russ Hobby reported that approval has been granted to include architectural type work in the scope of IETF activities.
2. There is an urgency to bring resource management out of research and into the IETF now. Future conversations with MIT Computer Science Lab indicate that introduction via BOF could occur at the next IETF in March, 1993. Personal Opinion: Perhaps this could be started in REMCONF.
3. An observation was made that coding may not require a separate Working Group.

II. Yee-Hsiang Chang presented an outline of his paper (pre-supplied) on "An Architectural Approach for Real-Time Multimedia Communications". The Group didn't reach closure on consensus, but several points made were:

1. The layering concept is a good notion, but rigid definitions should be avoided.

2. The conferencing application may be unique in how session level functions invoke multiple services.

3. Ways in which the architecture could be expanded were discussed. These ranged from taking "sub-application slices" and driving them through to see that the functions were there to support them, to picking a few services and "connecting user/provider". It was noted that some services may not need manager relationships.

Eve Schooler volunteered to start a spin-off BOF called Conferencing Control (CONFCTRL).

III. A starter list of conferencing packages was constructed and Tom Sandoski's expanded list and other information about a catalog of offerings is attached. One purpose for doing this was to identify coding implementations that could be further analyzed for possible interoperability purposes. Henning Schulzrinne agreed to gather additional detailed information and a status report from him is included here:

> A first cut at an encodings summary will be part of the suite of Internet-Drafts (future RFCs) to come out within the next few days. Oliver Jones and others volunteered to contribute information on Video Codecs. Discussion will take place through the normal REMCONF channels. After discussions with Steve Casner, the Group decided that this activity was (roughly) within the AVT Charter and there was no need to create more structure at this point.

IV. Programs for the December Packet Video Workshop at MCNC were handed out along with the attached starter list of chipset and codec manufacturers, which will become part of the "Offerings Catalog" mentioned earlier in these Minutes.

It was mentioned that we needed to add card/board level product companies to this list. That will be done.

Attendees

Lou Berger	lberger@bbn.com
Dean Blackketter	deanb@apple.com
Scott Brim	Scott_Brim@cornell.edu
Stephen Casner	casner@isi.edu
Kay Chang	chang@chang.austin.ibm.com
Wo Chang	wchang@nist.gov

Richard Cogger	rhx@cornell.cit.bitnet
Kurt Dobbins	dobbins@ctron.com
Jack Drescher	drescher@concert.net
Tom Easterday	tom@cic.net
Hans Eriksson	hans@sics.se
William Fink	bill@wizard.gsfc.nasa.gov
Jerry Friesen	jafries@sandia.llnl.gov
Robert Gilligan	Bob.Gilligan@eng.sun.com
Russ Hobby	rdhobby@ucdavis.edu
Don Hoffman	don.hoffman@eng.sun.com
Christian Huitema	christian.huitema@sophia.inria.fr
Oliver Jones	oj@pictel.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Christopher Kolb	kolb@psi.com
Paul Lambert	paul_lambert@email.mot.com
Allison Mankin	mankin@gateway.mitre.org
Matt Mathis	mathis@a.psc.edu
Greg Minshall	minshall@wc.novell.com
Michael Newell	mnewell@nsipo.nasa.gov
Ari Ollikainen	ari@es.net
Jim Perchik	perchik@athena.mit.edu
Mike Petry	petry@ni.umd.edu
Allan Rubens	acr@merit.edu
Tom Sandoski	tom@concert.net
Eve Schooler	schooler@isi.edu
Henning Schulzrinne	hgs@research.att.com
Scott Stein	scotts@apple.com
Terrance Sullivan	terrys@newbridge.com
Claudio Topolcic	topolcic@cnri.reston.va.us
Thierry Turletti	turletti@sophia.inria.fr
Janet Vratny	janet@apple.com
Abel Weinrib	abel@bellcore.com
Jeff Young	young@alw.nih.gov
Paul Zawada	Zawada@ncsa.uiuc.edu

An Architectural Approach for the Real-Time Multimedia Communications

Yee-Hsiang Chang
Communications Research
MCNC Center for Communications

email: yhc@concert.net

Outlines

My view of the problem

Our goals

Proposed directions

The Problem

Too many technologies intertwine together.

No effort to allow coherent development at each technology.

What is the role of IETF?

Our Goals

Ease Interoperability

Ease Integration

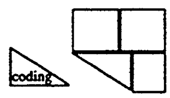

Proposed Directions

Architecture approach -- use an overall picture as a starting point to define the relationship among different technologies.

Functional specifications of each module.

Multimedia Communications Architecture

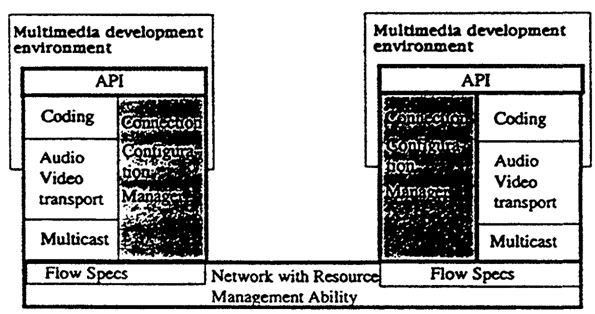

Major technologies lie on: Workstation (HW and SW), Networking, and Coding/Compression.

Protocol Architecture

ISO Layers	*Required Functionality*
Application layer:	API, Conference services
Presentation layer:	Coding
Session layer:	Connection and configuration
Transport layer:	Fast transport protocol for multimedia
Network layer:	Multicast, Resource management
Data link layer:	Resource management

CURRENT MEETING REPORT

Minutes of the Remote Mail Protocol BOF (REMMAIL)

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

Attendees

Bryan Beecher	bryan@umich.edu
Philip Budne	phil@shiva.com
David Conklin	conklin@jvnc.net
Wesley Craig	wes@umich.edu
Chuck Cranor	chuck@maria.wustl.edu
David Crocker	dcrocker@mordor.stanford.edu
Tim Dixon	dixon@rare.nl
Erik Fair	fair@apple.com
Roger Fajman	raf@cu.nih.gov
Ned Freed	ned@innosoft.com
Terry Gray	gray@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
David Katinsky	dmk@rutgers.edu
John Klensin	klensin@infoods.unu.edu
Eliot Lear	lear@sgi.com
Edward Levinson	levinson@pica.army.mil
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Rakesh Patel	patel@noc.rutgers.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Richard Schmalgemeier	rgs@merit.edu
Mark Smith	mcs@umich.edu
Jim Thompson	jim@tadpole.com
Theodore Ts'o	tytso@mit.edu
Stuart Vance	vance@tgv.com

2.1.1 Internet Mail Extensions (smtpext)

Charter

Chair(s):

John Klensin, `klensin@infoods.unu.edu`
Ned Freed, `ned@innosoft.com`

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done	Review the Charter of the Group. Determine if changes to SMTP are necessary. Discuss the needs for backward compatability, and interoperability. This discussion will be held by email.
Done	Discuss the elimination of the 7 bit restrictions in SMTP, and the implications of removing this restriction in terms of interoperation.
Done	Discuss the issues involved with binary transmission. Determine whether a "binary" mode should be pursued, and whether the SMTP line length restriction should be eliminated.
Done	Write a document specifying the changes to SMTP agreed to by the Group. Post as an Internet-Draft.
Done	Review and finalize the SMTP Extensions document.
Done	Submit the SMTP Extensions document as a Proposed Standard.

Internet-Drafts:

"SMTP Service Extensions", 09/02/1992, J. Klensin, N. Freed, E. Stefferud <draft-rose-extensions-06.txt>

"SMTP Service Extension for Message Size Declaration", 09/09/1992, K. Moore, N. Freed, J. Klensin <draft-moore-extension-size-04.txt>

"SMTP Service Extension for 8bit-MIMEtransport", 11/25/1992, J. Klensin, N. Freed, M. Rose <draft-ietf-smtpext-8bit-mime-00.txt>

"Transition of Internet Mail from Just-Send-8 to 8Bit-SMTP/MIME", 11/25/1992, G. Vaudreuil <draft-ietf-smtpext-transition-02.txt>

CURRENT MEETING REPORT

Reported by John Klensin/MIT

Minutes of the SMTP Extensions Working Group (SMTPEXT)

Summary

The Working Group has once again finished its work and is ready to submit rewritten documents to the IESG for Proposed Standard status. Documents reviewed and completed this week include revised versions of the following:

- "SMTP Service Extensions" model
- "SMTP Service Extension for Message Size Declaration"
- "SMTP Service Extension for 8bit-MIMEtransport"

From a protocol standpoint, these documents are substantially equivalent to the one that emerged from the Boston IETF except for the changed keyword model of the "EHLO" command response. The following documents will follow these three in short order:

- A contribution to the MIME effort specifying the logic and conventions for 8bit to 7bit (transport) conversion.

- An informational document describing transitional strategies for existing "8 bit clean" implementations.

- An informational document that contains additional clarification and guidance material needed to support the protocol extensions (most of this material is from the earlier (consolidated) Working Group draft.

The Working Group met twice during this IETF. At the beginning of the first session, the Working Group reviewed new versions of the modular documents developed after the previous last call. These versions, edited by Ned Freed, contained a re-editing to incorporate materials that were still important from the earlier Working Group draft. Significant, and other outstanding, technical issues were then reviewed and decided upon.

- Document format: Three+1 (Service extensions, Size, 8bit + informational) or three+2 (... plus informational and folklore (e.g., using Julian's document as a basis).

  Decision: Multidocument model, not one document, but with the expectation of advancing the three together, i.e., "three documents, one standard".

- Service extensions/EHLO: The key remaining differences between the new proposal and the earlier Working Group one are in the use of keywords, rather than specific verbs, in EHLO and in the use of parameters (where feasible) to existing commands rather than alternate command forms.

Decision: The keyword form is clearly preferable. Given the desire to avoid additional round trips, the increase in complexity of command parsing associated with the parameters is a desirable tradeoff.

- An outstanding question is whether possible future extensions that would be associated with commands that don't accept arguments should be implemented with new commands or with parameters on the old ones.

  Decision: The present Working Group inclination, reflected in the document, is that extensions to parameter-less commands (e.g., DATA should be performed by making new commands. This strategy should be slightly more robust against sloppy implementations. However, this decision can be reviewed when the first such extension is actually proposed.

- If an extended command is issued with more than one set of extension parameters, and the server wishes to indicate that the request was not satisfied (i.e., that there is an error condition), there could be an ambiguity about which of the parameters (or the base command) was at fault. Several possible solutions have been proposed, including using the explanatory text in special ways, creating a series of per-extension error codes (possibly in the current-unused 6yz or 7yz range), or ignoring the issue on the assumption that more detail would encourage attempts to negotiate options.

  Decision: Consistent with tradition and the spirit of RFC1123, things either succeed or fail and we do not provide for tricky negotiation or alternative-seeking. A minimum number of reply codes will be used, implementors may provide textual explanation, but clients should not attempt to take specific action on these.

- SIZE: Change from kilo-octets to bytes, with supporting language.

  Decision: Agreed without dissent.

- Use of a single number versus several numbers (e.g., the old LIMIT).

  Decision: Agreed.

  These two issues were the only apparently-outstanding ones with SIZE and the only substantive differences between the Moore proposal and the original committee draft not covered elsewhere in this notes. SIZE is therefore closed out and ready for forwarding.

- 8bit clean: There was an extended discussion about the existing "8bit clean" vendors and the supporting facilities they needed. It was concluded that the CONVERT/NOCONVERT facilities did them no good and that, if the investment was made to send EHLO, then it was plausible to make the further investment to send MIME.

Decision: The Working Group agreed, following the pre-July draft, that "8bit" implies MIME and that the keywords chosen should reflect this. This change removes the NOCONVERT/ CONVERT/ and MIME keywords from the EHLO response, and eliminates the need for conversion to application/octet-stream and character set "unknown" in the protocol document. A separate, non-standards-track, document will be developed to suggest transition strategies.

- Relaying: RFC1123 attempted to discourage relaying in the Internet. Sending clients in quest of relays who could perform a conversion after receiving a rejection from a target host probably represents bad policy (although there is neither need nor desire to prohibit static determination of conversion gateways). Leaving the "go find a relay" alternative in the text as a means of coping with rejections implies error message complexities that are not worth the trouble.

  Decision: Remove the text that appears to encourage finding a relay if mail cannot be delivered as originally specified.

- MIME-MIME conversions: As things now stand, the text contains several statements about MIME processing that effectively create two-way crossreferences with the MIME document. The earlier Working Group draft resolved this problem by simply insisting that any conversions produce valid MIME, believing that the definitions of "valid MIME" belonged in MIME documents, not in SMTP extensions ones.

  Decision: These text should be removed and replaced by a "convert to valid MIME" statement. Any additional statements about MIME and how to handle it should be made in modifications to the MIME RFC or, if necessary, in non-standards-trace transition document.

- Trace/received syntax: At the start of IETF, the document overloaded the RFC821/822 Received phrase "with" (specified in those RFCs as a transport protocol) to include conversion statements, e.g., "with 8bit-to-base64". This changed the semantics of the 821/822 definition, however subtly. It also produced a significant potential for misunderstanding, as evidenced by the example in the text, e.g., Received: from baiji.dbc.mtview.ca.us by dbc.mtview.ca.us with 8bit-to-base64. It is not clear what this means, since the translation/conversion would normally occur intra-host.

  Decision: A new phrase keyword will be added, "convert", followed by a keyword that will specify the conversion performed in the process of receiving mail and sending it on. This solution also reduces the potential for generating many extra Received lines, which could be problematic for (probably non-conforming) implementations that use the number of Received headers as a trap for mail loops.

- The conversion issue: With the proposed documents, the Working Group appeared to have come full circle to a variation on the so-called "wretched solution" of 18 or so months ago. That approach called for expecting that any MTA that was willing to accept 8bit traffic must be prepared to convert to 7bit [MIME] if needed. This implied the ability to parse MIME and make per-body-part decisions, raising the threshold of effort that must go into such an MTA and forcing inclusion of a facility that would be unneeded if the transition to an entirely 8bit world ever completed. The Working Group agreed to this in San Diego and did not raise it again in Boston, nor was the issue raised during the Last Call discussion /cries of agony. It was, however, suggested that there never was real consensus, just exhaustion, and that the requirement was ultimately spurious, that the only thing accomplished by such a requirement was to insist that an implementation that was unwilling to convert lie about the reason for rejecting the message.

Decision: The document will be revised to indicate a preference for conversion, but to provide for message rejection when conversion was not possible for some reason.

- MXE: Some months ago, the Working Group proposed a DNS extension, MXE, which could be used to identify enhanced SMTP servers prior to opening SMTP connections. This suggestion was forwarded to the DNS Working Group, which has not taken any action on it.

  Decision: the proposal should be withdrawn. Given changes in the extension model, if anything is needed, it might be based on a cross between the EHLO response and the WKS record. Anyone who is convinced that this is important should write a proposal.

The Working Group appears to have reached consensus on the above issues and the form and content of revised documents. After the documents are revised to reflect the decisions outlined above and a brief review has taken place on the mailing list, the documents will once again be recommended to the IESG for processing as a Proposed Standard.

Attendees

Randall Atkinson	atkinson@itd.nrl.navy.mil
Bryan Beecher	bryan@umich.edu
Fred Bohle	fab@interlink.com
Kay Chang	chang@chang.austin.ibm.com
James Conklin	jbc@bitnic.educom.edu
Chuck Cranor	chuck@maria.wustl.edu
Erik Fair	fair@apple.com
Roger Fajman	raf@cu.nih.gov
Ned Freed	ned@innosoft.com
Olafur Gudmundsson	ogud@cs.umd.edu
Marco Hernandez	marco@mh-slip.educom.edu
Russ Hobby	rdhobby@ucdavis.edu
Tim Howes	tim@umich.edu.
Frank Kastenholz	kasten@ftp.com
Neil Katin	neil.katin@eng.sun.com

John Klensin	klensin@infoods.unu.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Eliot Lear	lear@sgi.com
Edward Levinson	levinson@pica.army.mil
Chris Newman	chrisn+@cmu.edu
Michael Patton	map@bbn.com
Marshall Rose	mrose@dbc.mtview.ca.us
Tim Seaver	tas@concert.net
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgras@bitnic.educom.edu
Einar Stefferud	stef@nma.com
Stuart Vance	vance@tgv.com
Gregory Vaudreuil	gvaudre@cnri.reston.va.us

2.1.2 Internet Message Extensions (822ext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done	Review the Charter, and refine the Group's focus. Decide whether this is a worthwhile effort.
Done	Discuss, debate, and choose a framework for the solution. Assign writing assignments, and identify issues to be resolved.
Done	Review exiting writing, resolve outstanding issues, identify new work, and work toward a complete document.
Done	Post a first Internet Draft.
Done	Review and finalize the draft document.
Done	Submit the document as a Proposed Standard.

Internet-Drafts:

"Japanese Character Encoding for Internet Messages", 08/25/1992, Jun Murai, Mark Crispin, Erik van der Poel <draft-ietf-822ext-iso2022jp-02.txt>

Request For Comments:

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

2.1.3 Network Database (netdata)

Charter

Chair(s):

Daisy Shen, `daisy@watson.ibm.com`

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done Review and approve the Charter, making any changes necessary. Examine needs, resources for this network database protocol and define the scope of work. Begin work on a framework for 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 RFC.

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

Internet-Drafts:

"Network Database Protocol", 06/26/1991, Daisy Shen <draft-ietf-netdata-netdata-03.txt>

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

CURRENT MEETING REPORT

Reported by Scott Newman/DEC

Minutes of the Network Database Working Group (NETDATA)

The NETDATA Working Group met for a single session at the November IETF meeting. The meeting was lightly attended, but the interest level was high among most of the attendees. The meeting was chaired by Scott Newman because the regular Chair, Daisy Shen, could not attend.

The entire meeting was spent discussing the approach of the SQL Access Group to remote database access. Scott presented a very quick overview of the SQL Access Group and its efforts. A more technical discussion was facilitated by a longer presentation that provided a more detailed description of ISO RDA, the SQL Access subset and extensions to RDA and SQL Access' approach to using RDA on an IP-based network. (Copies of these presentations are available in Postscript format, and they are expected to appear in the IETF proceedings).

In general, the participants were very supportive of the SQL Access approach. The most significant issues and comments were as follows:

- Currently, none of the specifications required to implement the SQL Access approach are available on-line. SQL Access references X/Open and ISO documents, so the issue is somewhat involved. It is important to address this issue, in order to make the specifications accessible on-line for the Internet community.

- As discussed last month, additional work is required in the area of security. There are several potential ways to address enhancements in this area.

- No one was concerned that the SQL Access approach for mapping to TCP/IP used a direct mapping approach, instead of using RFC 1006. In fact, several participants much preferred the simplicity of the direct mapping approach. There was general feedback that RFC 1006 was not seeing a lot of action at this point in time.

In summary, there is one logistical/legal issue with regard to getting the required specifications on-line, and there was one technical/requirements issue in the area of enhancing security. The Group felt that they would prefer to base their work on SQL Access' work because it has the support of many more companies and organizations than the existing netdata approach, and the specifications are mature. However, it unanimously agreed that due to the meeting's light attendance, and the absence of the regular Chair, there should be further discussion over e-mail. A decision on whether to continue with the existing approach or begin adapting the SQL Access work should be made via e-mail, prior to the next meeting.

The following action items resulted from the meeting:

- Scott and Russ are to arrange for Russ to address the SQL Access Board of Directors (i.e., the managers' meeting) in early 1993.
- Scott will write-up his notes that compare the existing netdata approach with the SQL Access approach. This write-up will be posted to the NETDATA list.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Bryan Beecher	bryan@umich.edu
Letha Dugas	4371362@mcimail.com
Russ Hobby	rdhobby@ucdavis.edu
Scott Newman	newman@broke.enet.dec.com
Richard Schmalgemeier	rgs@merit.edu

SQL Access

Overview

Scott Newman
Database Interoperability Engineering
Digital Equipment Corporation

digital™

Topics

- The SQL Access Group
- SQL Access Interfaces and Specifications
- Summary

digital™

The SQL Access Group

A non-profit corporation comprised of over 40 member companies, including almost all database vendors, tools vendors and some end-users. X/Open works jointly with SQL Access on specifications, including their publication.

digital™

The SQL Access Group

Goals

Work together to accelerate multi-vendor database interoperability and application portability a reality by:

- **Accelerating existing standards efforts**
- **Prototyping multi-vendor interoperability**

digital™

SQL Access Interfaces

- **API — Application Programming Interface**

  An embedded SQL language application programming interface definition.

- **CLI — Call-Level Interface**

  A procedure call-oriented, dynamic SQL interface.

- **FAP — Formats and Protocols**

  A client-server communication protocol for SQL remote database access.

digital™

SQL Access Interfaces

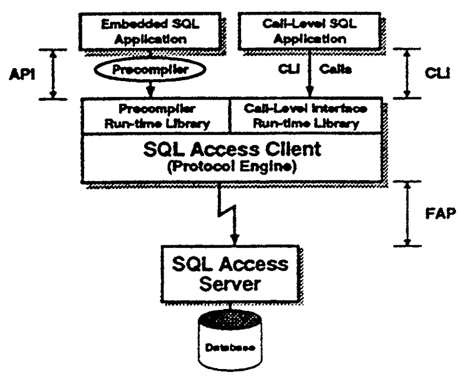

digital™

Application Programming Interface

Specifies the SQL language that is embedded in an application program.

- Based on Entry Level SQL-92 (SQL2), with the following elements from more advanced levels:
  - —Dynamic SQL
  - —Additional Data Types (e.g., VARCHAR)
  - —Schema Information Tables (SITs) subset
  - —Enhanced diagnostic information

digital™

Application Programming Interface

- Client programs can access multiple servers, one-at-a-time
- *SQL-Connection* management statements are used to control client-server associations
  - —Association establishment and tear-down
  - —Server context switching

digital™

Call-Level Interface (CLI)

A procedure call-oriented interface for database access using dynamic SQL.

- Base document jointly submitted by Microsoft, Sybase and Digital. (Many changes since then).
- Removes the API's precompiler requirement.
- Facilitates shrink-wrapped software.
- Forms the "core" of Microsoft's ODBC.

digital™

Formats and Protocols (FAP)

Specifies the formats and protocols used for communication between a SQL Access client and server.

- Based on ISO Remote Database Access (RDA) with the following additions and subtractions:
  - —Several extensions to support API features not supported by ISO RDA.
  - —Does not support *Control Dialogues* for out-of-band cancel and status functions.

digital™

ISO Remote Database Access

- OSI Application Layer International Standard for heterogeneous remote database access
- Recently progressed to International Standard
- Separated into *generic* and *specializations*
  - —ISO 9579-1: Generic RDA
  - —ISO 9579-2: RDA SQL Specialization
- Completely vendor-neutral
  - —Standard SQL
  - —Platform-neutral message encodings

digital™

Formats and Protocols (FAP)

- Currently supports single-phase transaction commitment.
- Supported networks:
  - —Initially ISO OSI, due to ISO RDA origin
  - —TCP/IP mapping draft specification complete

digital™

Current Status

Initial specifications are being finalized for X/Open XPG4.

- API — Final, CAE-level revision for XPG4 available now.
- FAP — Preliminary Specification (near-final) available from X/Open now.
- CLI —Snapshot Specification (first draft) available from X/Open now.
- RDA for TCP/IP —First draft available through X/Open shortly.

digital™

SQL Access

Remote Database Access

Scott Newman
Database Interoperability Engineering
Digital Equipment Corporation

digital™

Topics

- ISO Remote Database Access (RDA)
- SQL Access Formats and Protocols (FAP)
- SQL Access RDA for TCP/IP
- Summary

digital™

ISO Remote Database Access

- OSI Application Layer International Standard for heterogeneous remote database access
- Generic standard with *specializations*
  - — ISO 9579-1: Generic RDA
  - — ISO 9579-2: RDA SQL Specialization
- Generic standard provides a general-purpose mechanism
- Specializations specify database type-specific refinements (e.g., additional rules, constructs)

digital™

RDA Characteristics

- Client-Server model
- Completely vendor-neutral
- Standard SQL: SQL-89 or SQL-92 Entry Level
- Standard error codes
  - —Database errors specified by SQL standards
  - —RDA errors fully specified by RDA standards
- Platform-neutral message syntax and encoding
  - —Uses ASN.1 and Basic Encoding Rules
  - —Completely and unambiguously specified

digital™

RDA Characteristics

- Handling of SQL statements fully specified
  - —General statement processing is specified
  - —Per-statement type handling is specified, including mapping statements to protocol
- Primarily request-response protocol
  - —Some requests have no response, or error responses only for performance reasons
  - —Several request types are "non-blocking" (requests can be issued without waiting)

digital™

RDA Characteristics

- Database language requests
  - —Multiple database requests may be submitted simultaneously
  - —*Repetition count* mechanism allows batching of input and output data values
- Two *Application Contexts:*
  - —*Basic Context* —single-phase commitment (using RDA services)
  - —*TP Context* —two-phase commitment (using ISO TP services)
- Flexible character set usage (per-column basis)

digital™

RDA Architecture

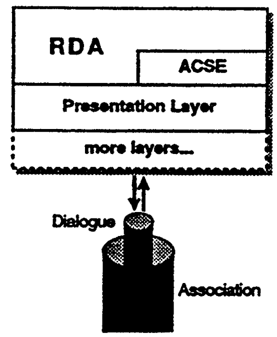

RDA Concepts

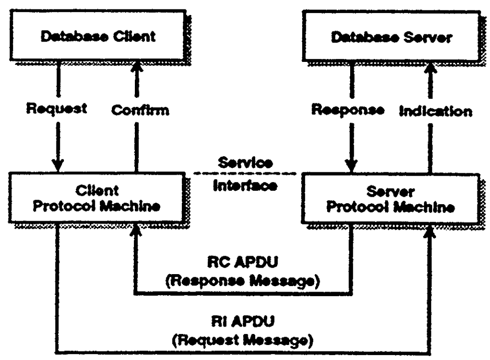

Associations

- The basic, underlying client-server connection
- Managed by ACSE (Association Control Service Element; ISO 8649)
  - — Request primitives for set-up and tear-down
  - — Indication events for association failures
- Associations are shared with ISO TP (if used)
- Security authorization at client-to-server level
- Associations may be re-used for multiple dialogues (amortize set-up costs)

Dialogues

- End-user-level conversation between client and server
- Full-duplex communications path
- One dialogue maps to an underlying association
- Security authorization at end-user level
  - —Authorization information required is specified by individual servers
- Special-purpose *control dialogues* are used for cancel and status functions on operations occurring on another dialogue to the server

RDA Services

- Dialogue management services
  - —*R-Initialize* (confirmed; blocking)
  - —*R-Terminate* (confirmed; blocking)
  - —Both have corresponding req/rsp messages
- Resource handling services
  - —*R-Open* (confirmed; non-blocking)
  - —*R-Close* (confirmed; non-blocking)
  - —Both have corresponding req/rsp messages

RDA Services

- Transaction management services
  - —*R-BeginTransaction* (confirmed if error)
  - —*R-Commit* (confirmed; blocking)
  - —*R-Rollback* (confirmed; blocking)
  - —All have corresponding req/rsp messages
- Control services (in- or out-of-band)
  - —*R-Cancel* (confirmed; non-blocking)
  - —*R-Status* (confirmed; non-blocking)
  - —Both have corresponding req/rsp messages

RDA Services

- Database language services —*R-ExecuteDBL*
  - —"Immediate execution" of database requests
  - —Request primitive parameters
    - * operation id
    - * database handle (from *R-Open*)
    - * SQL statement
    - * argument specification (descriptors)
    - * result specification (descriptors)
    - * argument values
    - * repetition count

digital™

RDA Services

- *R-ExecuteDBL* continued...
  - —Confirm primitive parameters
    - * operation id
    - * completion status information
    - * result specification (descriptors)
    - * result values
- Static SQL statements (specified by RDA)
  - —Dynamic SQL statements (SQL Access)
  - —Statements are processed to replace host variables with ":H"

digital™

RDA Services

- *R-ExecuteDBL* continued...
- Argument and result specifications
  - —Provide metadata information for input and output data values
  - —Can be used to request datatype conversions from the server
  - —Usage depends on executed SQL statement (ranges from required to optional to omitted)
- Both request and response messages
- Multiple requests may be outstanding

digital™

RDA Services

- "Stored execution" services
  - *R-DefineDBL*
  - *R-InvokeDBL*
  - *R-DropDBL*
- Used to define database language that persist for the duration of the dialogue (only)

RDA Protocol

- Service primitives map very closely to actual protocol messages
- Protocol message contents map very closely to actual service primitive parameters
- One additional message is used to synchronize client and server transaction state under certain error conditions
- Failure behaviours are documented
  - Interaction between association/dialogue failure and transaction state is specified

SQL Access Formats and Protocols

Specifies the formats and protocols used for communication between a SQL Access-compliant client and server.

- Uses ISO Remote Database Access (RDA) subset with extensions.
- Currently supports single-phase transaction commitment.
- Supported networks: OSI and TCP/IP

FAP Subset of RDA

The SQL Access Group subsets RDA in order to first, completely specify basic interoperability.

- The following RDA features are not supported:
  - — *Control Dialogues* (used for out-of-band cancel and status)
  - — *R-DefineDBL, R-InvokeDBL and R-DropDBL* (a much-debated performance optimization)
  - — *TP Application Context* (uses ISO TP for distributed transactions)
- All unsupported features are addressed by official RDA subsets (negotiated at start-up).

digital™

FAP Extensions to RDA

The SQL Access Group has specified minor extensions to RDA in order to support SQL features of the API not supported by RDA*.

- The following three SQL features require extensions to RDA:
  - — Dynamic SQL
  - — Character-varying (VARCHAR) datatype
  - — Extended diagnostic information
- All three extensions have been submitted for inclusion in a future RDA Addendum.

* ISO RDA is constrained, by definition, to Entry Level SQL-92 only.

digital™

SQL Access TCP/IP Mapping

- Goals
  - — Maximize potential target platforms
    - * Minimize memory footprint
    - * Avoid dependencies
    - * Maximize user/administrator acceptance
  - — Minimize specification effort; don't re-invent the wheel
  - — Use an approach easily extended to other network technologies

digital™

SQL Access TCP/IP Mapping

- Approach
  - —Map SQL Access FAP directly onto TCP/IP service interface (e.g., sockets, XTI, etc.)
  - —Retain RDA message contents and encoding
  - —Replace ACSE association set-up with a single, simple message exchange
  - —Specified in terms of a mapping from Presentation and ACSE service primitives to TCP/IP service interface

digital™

SQL Access TCP/IP Mapping

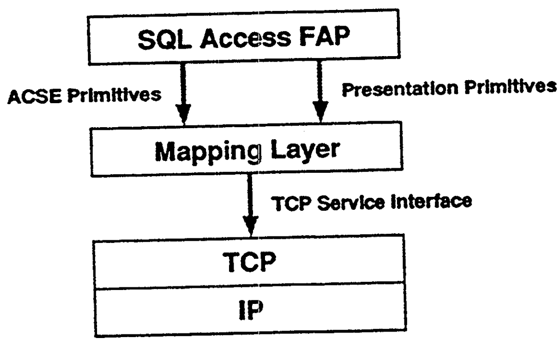

digital™

Summary

- RDA is a stable, complete specification
  - —Vendor- and platform-neutral
  - —Standard SQL language and error codes
  - —SQL statement handling and mapping
  - —Unambiguous and complete message def'ns
  - —Basic and two-phase commitment "modes"
  - —Product of person-decades of work
  - —International standard
  - —Expected to be required by GOSIP and FIPS

digital™

Summary

- SQL Access Group adds:
  - —SQL language extensions in common use
  - —Practical subset of RDA as a first step
  - —Intensive approach using prototyping
  - —Implementor's agreements, such as parameter limits, etc.
  - —Technology-sharing relationship with X/Open
  - —Large body of supporting companies
  - —Mapping for RDA over TCP/IP

digital™

Summary

- RDA over TCP/IP
  - —Simple first step for basic interoperability using the direct mapping approach
  - —Future extensibility using RFC 1006 for more advanced features (e.g., two-phase commit)
- Implementations/Products
  - —Implicitly widely accepted due to ISO, NIST, X/Open and SQL Access
  - —Server product from Digital —client soon
  - —Base technology (source) products from Retix (supports OSI, TCP/IP (direct and RFC1006))

digital™

2.1.4 Network News Transport Protocol (nntp)

Charter

Chair(s):

Eliot Lear, `lear@sgi.com`

Mailing Lists:

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

Description of Working Group:

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

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

(1) User level and site designated authentication methods; (2) Binary transfer capability; (3) Minimization of line turnaround; and (4) Stronger article selection capability.

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

Goals and Milestones:

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

Internet-Drafts:

"Network News Transfer Protocol Version 2: A Protocol for the Stream-Based Transmission of News", 09/30/1991, Eliot Lear <draft-ietf-nntp-news-01.txt, .ps>

CURRENT MEETING REPORT

Reported by Eliot Lear/SGI

Minutes of the Network News Transport Protocol Working Group (NNTP)

This meeting was largely organizational, and fairly short in duration.

There is a new NNTP v2 draft which will be posted early next week. A few tweaks may yet be needed, but it is otherwise done.

With that the Group turned our efforts to reader issues, spending the rest of the meeting on essentially three issues:

1. The relationship between news and mail; does IMAP already have the facilities that would otherwise be required for an NNRP?

2. To what end should a reader protocol be concerned with ACL management?
   - Authentication issues.

3. How general should the search mechanism be? Should discussion trees be handled using it or using something separate? The Group essentially talked about something on this order:

   SEARCH <method> <text>

   A method might be something like BODY-TEXT and text could be something like a quoted string containing "Operational Experience With TUBA".

Issue [1] has essentially been deferred until it can be determined what the REMMAIL Group is doing. Issue [2] has been assigned to Ted Tso. Issue [3] has been assigned to Mel Pleasant. Ted and Mel are going to produce documents for the Working Group by January 20th, going into some detail, with the goal being a combined document for the March time-frame.

There is a forth issue, which is whether the reader stuff should go into NNTP as a set of extensions, or as a new protocol. If we go ahead with a separate protocol, it was stated without objection that we could pursue an experimental track, and upgrade if others like the results enough to implement the protocol. A final decision on this topic doesn't need to be made until March. A separate mailing list will be formed to discuss the reader document(s).

Attendees

Vikas Aggarwal	vikas@jvnc.net
David Conklin	conklin@jvnc.net
Wesley Craig	wes@umich.edu

Chuck Cranor	chuck@maria.wustl.edu
Peter DiCamillo	Peter_DiCamillo@brown.edu
Erik Fair	fair@apple.com
Terry Gray	gray@cac.washington.edu
David Katinsky	dmk@rutgers.edu
Eliot Lear	lear@sgi.com
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Rakesh Patel	patel@noc.rutgers.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Tim Seaver	tas@concert.net
Mark Smith	mcs@umich.edu
Theodore Ts'o	tytso@mit.edu

2.1.5 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 to be considered as standards in the Internet community.

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

Goals and Milestones:

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

Request For Comments:

RFC 1179 "Line Printer Daemon Protocol"

2.1.6 TELNET (telnet)

Charter

Chair(s):

Steve Alexander, `stevea@i88.isc.com`

Mailing Lists:

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

Description of Working Group:

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

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

(2) Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically:

- Environment variable passing - Authentication - Encryption - Compression

(3) Act as a clearing-house for all proposed RFCs that deal with the TELNET protocol.

Goals and Milestones:

Done	Write an environment option.
Done	Post an Internet-Draft describing the authentication option.
Dec 1990	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 Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-04.txt>

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

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

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

Request For Comments:

RFC 1116	"Telnet Linemode option"
RFC 1184	"Telnet Linemode Option"
RFC 1372	"Telnet Remote Flow Control Option"

CURRENT MEETING REPORT

Reported by Steve Alexander/INTERACTIVE Systems

Minutes of the TELNET Working Group (TELNET)

The Telnet Working Group met on November 17th in Washington. We discussed Ted Ts'o's changes to the Kerberos V document, and they were received favorably. Steve Alexander will produce a new draft of the Kerberos V document by year's end.

The Group then discussed delegation of privileges via the authentication mechanism. Most of the discussion centered on whether or not privileges could be delegated in subsets or whether delegation should be all or nothing. No consensus was reached, and Cliff Neuman said he would investigate this further on his own. The remainder of the discussion was about whether or not delegation should be a generic feature of the authentication mechanism or done in mechanism specific ways, since not all authentication protocols support delegation. It was decided that delegation should be handled by the specific mechanisms, and not as a general part of authentication.

Attendees

Steve Alexander	stevea@i88.isc.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Ken Hirata	khirata@emulex.com
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Kent Malave	kent@bach.austin.ibm.com
Louis Mamakos	louie@ni.umd.edu
Clifford Neuman	bcn@isi.edu
Joseph Ramus	ramus@nersc.gov
Jeffrey Schiller	jis@mit.edu
Cris Shuldiner	cws@ftp.com
Sam Sjogren	sjogren@tgv.com
Theodore Ts'o	tytso@mit.edu

2.2 Internet Area

Director(s):

- Philip Almquist: almquist@jessica.stanford.edu
- Stev Knowles: stev@ftp.com

Area Summary reported by Philip Almquist/Consultant

Considerable activity occurred in the Internet Area during this meeting. Eight of the Internet Area's working groups met, and there were an additional four BOF sessions.

The work that undoubtedly attracted the most interest was the continued efforts on proposals to replace the current IP protocol with one which more readily scales to the scope that the Internet is rapidly attaining. Four working groups have been aggressively attacking this problem:

1. IP Address Encapsulation (IPAE)
2. P Internet Protocol (PIP)
3. Simple Internet Protocol (SIP)
4. TCP/UDP over CLNP-addressed Networks (TUBA)

Each of these groups gave a plenary presentation on their progress so far and met during the week to continue to refine their proposals and their documents. In addition, a BOF on Selection Criteria considered the problem of how to best evaluate the strengths and weaknesses of the four proposals.

IP over Fibre Channel (FIBREIP)

A BOF on IP over Fibre Channel met to discuss Yakov Rekhter's Internet-Draft "IP and ARP on Fibre Channel (FC)". The Group felt that only minor changes needed to be made to the document, but elected to defer entering it into the standards process until there is some implementation experience. This work is being closely coordinated with the ANSI Fibre Channel committee.

Inter-Domain Multicast Routing (IDMR)

The IDMR BOF met to discuss dynamic routing of IP multicast datagrams to multicast groups containing members in multiple routing domains.

Selection Criteria (SELECT)

The objective of the Selection Criteria BOF was to develop consensus on a precise statement of the community's goals for a replacement for IP. The goal was to provide a yardstick against which the various proposals could be objectively measured to point up their relative strengths and weaknesses. Needless to say, this goal was far too ambitious to actually be achievable in the single session available.

Dynamic Host Configuration (DHC)

The Dynamic Host Configuration Working Group worked to finish up the set of documents which define the DHCP protocol and its options. The Group then held some preliminary discussions on an additional protocol for coordinating the activities of multiple DHCP servers.

IP over AppleTalk (APPLEIP)

The IP over Appletalk Working Group heard reports on several topics and worked on the IP over Appletalk document and a revision to the Appletalk MIB. The Group expects to wind down its activities during the next couple of meetings due to the development of an appropriate forum (ASIG) for working on Appletalk-related issues outside of the context of the IETF.

IP over Asynchronous Transfer Mode (ATM)

The ATM Working Group finished up its IP over ATM document and intends to submits it as a Proposed Standard. The Working Group also discussed the current state of ATM signaling protocols in CCITT and the ATM Forum, and requirements that Internet protocols impose on ATM signaling. The Group also received a request to establish formal relations with the ATM Forum to facilitate the exchange of protocols and ideas between the two organizations.

Point-to-Point Protocol Extensions (PPPEXT)

The Point-to-point Protocol (PPP) Extensions Working Group worked on finishing up a set of extensions to the Link Layer portion of PPP (LCP). They also discussed some documents concerning IPX over PPP. A subgroup was formed to investigate conformance testing. Part of the meeting was a joint session with the IP over Large Public Data Networks (IPLPDN) Working Group to discuss how PPP mechanisms might be adapted to minimize the number of frames sent (important on WANs which impose per-packet charges).

CURRENT MEETING REPORT

Reported by Lansing Sloan/LLNL

Minutes of the IP over Fibre Channel BOF (FIBREIP)

Agenda

- Introduction to Fibre Channel (Lansing Sloan).
- Review "IP and ARP on Fibre Channel" Internet-Draft (Yakov Rekhter).
- What next?
  - Level of interest.
  - Next steps for document.

Introduction to Fibre Channel

The introduction to Fibre Channel stressed points that influenced "IP and ARP on Fibre Channel."

Fibre Channel (FC) defines several topologies. Some topologies provide many parallel paths, and therefore may not support broadcast and multicast well. This affects address resolution procedures.

Fibre Channel is being defined by ANSI X3T9.3 in a set of (draft) standards. "IP and ARP on Fibre Channel" depends on one of these draft standards, "Fibre Channel – Physical and Signaling Interface (FC-PH)." FC-PH version 3.0 is current. The first ANSI public review of FC-PH ends January 1, 1993.

Some Fibre Channel prototype implementations were shown in November at the Supercomputing '92 conference.

Review "IP and ARP on Fibre Channel"

Fibre Channel has many options, and for interoperability IP must constrain their use appropriately. Some topics are within the scope of the Internet-Draft and all others are outside the scope.

IEEE 802.2 LLC and IEEE SNAP are used for encapsulation (but full support of 802.2 is outside the scope).

Fibre Channel mechanisms ("exchanges") are used in a unidirectional manner. When IP traffic is bi-directional, independent "exchanges" are used for the two directions. Some optimizations may use more than two.

For address resolution, a "hardware address" consists of a 24-bit interface ID and a 64-bit "Initial Process Associator" (the latter may have a null value). A single IP address may map to multiple hardware addresses, to support redundant connections.

The ability to do local address resolution configuring (for bootstrapping and point-to-point links) is required. The ARP server is optional, it has a well-known address, its external behavior is described, its internal behavior is not described. The ARP format is followed; some fields are variable length.

Fibre Channel defines several classes of service, including connection-oriented and datagram modes. A connection maximizes performance for a given pair of interfaces but denies service to other interfaces while the connection lasts. The Draft has some guidelines. Connections should not last longer than 500 milliseconds.

Some Fibre Channel configurations permit non-transitive behavior. The Internet-Draft handles this by defining fully-connected "regions" and assigns distinct IP subnets to each region. No router is required for IP communication within a region. An interface can be in multiple regions and therefore may have multiple IP addresses.

Controversial and/or Unresolved Issues

There was some strong feeling that either the encapsulations should be limited to IP/ARP for efficiency or, alternatively, that IEEE 802.2 XID and TEST functions should be supported. Agreement may have been reached. For now, in any case, the Draft will not change.

There was some discussion whether the Draft should provide more guidance. It now emphasizes interoperability. For now, that will not change.

Better wording for ARP on point-to-point connections is needed.

Hosts can learn the hardware addresses of routers using address resolution, but details were not discussed.

The reason for 500-millisecond connection limits was not discussed.

The Draft does not specify reverse ARP. RARP can be provided, but having hosts pop up in the network may be undesirable.

Decisions (What next?)

Philip Almquist said he thought that the IETF should let ANSI continue with the technical work for now. Attendance was light, and in effect there was no independent review of the Internet-Draft by non-ANSI people. The people with detailed comments all work for companies that attend ANSI X3T9.3 meetings regularly. The IETF wants an IETF Fibre Channel working group but Philip said attendance shows that interest is presently too low.

One suggestion was to fix the Internet-Draft until ANSI is happy and then submit it to IETF for standards processing as a joint BOF/ANSI contribution. However, because the IETF has not had an effective review and because the Draft is not particularly self-contained (it assumes familiarity with quite a bit of Fibre Channel), it is unlikely that the Internet

Architecture Board could effectively read the document and determine if it assures interoperability. Therefore the Draft probably will not be submitted for the IETF standards track until interoperable implementations based on it exist. This will probably happen next summer.

Probably the existing ANSI mail groups "fibre-channel-ext" and "fc-ip-ext" may be used as the IETF mail groups as well, provided that people are not excluded.

A draft MIB for Fibre Channel is expected within a couple of months.

The IETF still wants the final say and change control on IP standards.

Attendees

Philip Almquist	`almquist@jessica.stanford.edu`
Vickie Brown	`brown@osi540sn.gsfc.nasa.gov`
Paul Griffiths	`griff@chang.austin.ibm.com`
Mark Laubach	`laubach@hpl.hp.com`
Drew Perkins	`ddp@andrew.cmu.edu`
Yakov Rekhter	`yakov@watson.ibm.com`
Lansing Sloan	`ljsloan@llnl.gov`
Elizabeth Vanderbeck	`beth@tdcsys2.vnet.ibm.com`
Gerry White	`gerry@lancity.com`

Introduction to Fibre Channel

Lansing Sloan
ljsloan@llnl.gov

Lawrence Livermore National Laboratory

November 17, 1992

IETF "fibreip" Birds of a Feather Session on
IP and ARP on Fibre Channel

I gratefully acknowledge Bryan Cook's assistance and the many viewgraphs that he has kindly provided.

Goals of This Presentation

- Introduce Fibre Channel
- Emphasize points that affect "IP and ARP over Fibre Channel" Internet Draft

Goals of Fibre Channel

Fibre Channel is intended to be a standard that

- offers cost-effective, high-performance communication,
- permits a single interface to access networks and various peripherals,
- permits highly parallel switching for very high overall throughput*,
- works with many existing and future fiber plants, and
- provides high reliability.

* Note: The high parallelism implies that broadcasting and multicasting are difficult, and (among other things) that ARP should not rely on broadcasts or multicasts.

FCS Scope

1. High-performance backbone to support:

- ◆ IP/802.2
- ◆ HIPPI-FP
- ◆ IPI3/SCSI
- ◆ SBCCS (IBM Format 0/1 command sets)

Fibre Channel Base Document

Fibre Channel will be defined by a number of standards. The base standard document for Fibre Channel will be "Fibre Channel — Physical and Signaling Interface," commonly called "FC-PH" or "FCS."

FC-PH is the only Fibre Channel standard required by "IP and ARP on Fibre Channel."

FC-PH

- defines the behavior of an interface (often called an "N_Port"),
- supports connection to a "Fabric," and
- also supports point-to-point connections.

FC-PH defines interface behavior at three levels, called

- FC-0,
- FC-1, and
- FC-2.

FCS Specifies Physical and Signalling Levels

- FC-0 Level
  - Transmitter/receiver types
  - Media types
  - Bit rates
- FC-1 Level
  - 8B/10B transmission code
  - Special characters
  - Synchronization rules
- FC-2 Level
  - Frame Formats
  - Frame Headers (Addresses, Sequence IDs, Exchange IDs)
  - Link Level Control (ACK, Ready, Busy, Reject)
  - Data Flow Control
  - Classes of Service
  - Segmentation, Flow Control, Recovery
  - Multiplexing management
  - Connection management
  - Login/Logout

Specified Signalling and Data Rates

FC-PH specifies signaling rates in FC-0 and is designed to provide the user data rates shown under good circumstances.

FC-0 Instantaneous signaling rate (Mbit/second)	Ideal effective user data rate above FC-2 (MByte/second)
1,062.5	100
531.25	50
265.625	25
132.8125	12.5

Different Instantaneous Data Rates Can Cause "Nontransitive" Behavior

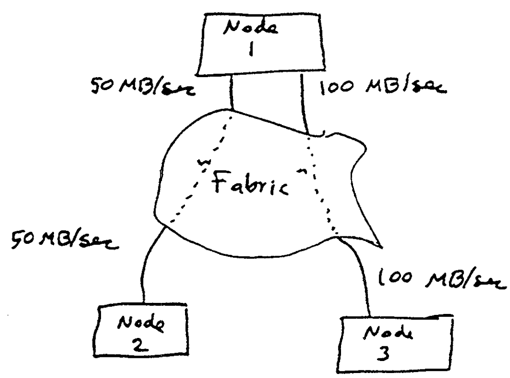

Node 2 can talk to Node 1 (50MB/sec).
Node 1 can talk to Node 3 (100MB/sec).

In some cases Node 2 can NOT talk directly to Node 3.

The Internet Draft discusses "regions" to get around this situation.

FC-0 Overview

- Physical Level - defines:
  - Transmitter/Receiver types
  - Media types
  - Bit rates

Fibre Optic Options

100 MByte/sec 1062.5 Mbit/sec	50 MByte/sec 531.25 Mbit/sec	25 MByte/sec 265.625 Mbit/sec	12.5 MByte/sec 132.8125 Mbit/sec
SM LWLD 1300 nm 2m-10Km	SM LWLD 1300 nm 2m-10Km	SM LWLD 1300 nm 2m-10Km	MM LED 1300 nm 0-1Km
SM LWLD 1300 nm 2m-2Km	MM(50 um) SWLD 780 nm 2m-1Km	SM LWLD 1300 nm 2m-2Km	
		MM(50 um) SWLD 780 nm 2m-2Km	
		MM(62.5 um) LED 1300 nm 0-1Km	

Note: There are also coax and twisted pair options

PART C.

FC-1

- Coding Level - defines:
  - 8B/10B transmission code
  - Special characters
  - Synchronization rules

PART D.

FC-2

Constructs

Following constructs are defined:

1. Frame
2. Sequence
3. Exchange
4. Operation

1. Data is transmitted in Frames:
   - A Frame consists of:
     - Start-of-Frame delimiter
     - Frame Header
     - Optional Headers
     - Payload
     - CRC
     - End-of-Frame delimiter
   - Each frame or group of frames is acknowledged:
     - Flow control
     - Delivery notification

Constructs ...

2. A Sequence is composed of 1-n Frames:
   - Unidirectional stream of frames for an operation
   - Recovery boundary
   - Each Sequence is identified by initiator: SEQ_ID
   - Each frame within a Sequence is numbered: SEQ_CNT

   Note: For example, 1 IP packet

3. An Exchange is composed of 1-n non-concurrent Sequences:
   - One Sequence may be active at a time
   - Exchange is identified by each end: OX_ID, RX_ID

   Note: For example, 1-N IP packets

4. An Operation consists of 1-n Exchanges:
   - May be concurrent
   - Operation is identified by each end: OO_AS, RO_AS

FC-2 Link Level Overview

- Has view of all concurrent sequences/exchanges at single N_Port
- Functions:
  1. Frame Formats
  2. Frame Headers (Addresses, Sequence IDs, Exchange IDs)
  3. Link Level Control (ACK, Ready, Busy, Reject)
  4. Data Flow Control
  5. Classes of Service
  6. Segmentation, Flow Control, Recovery
  7. Multiplexing management
  8. Connection management
  9. Login/Logout

FCS Frame Format

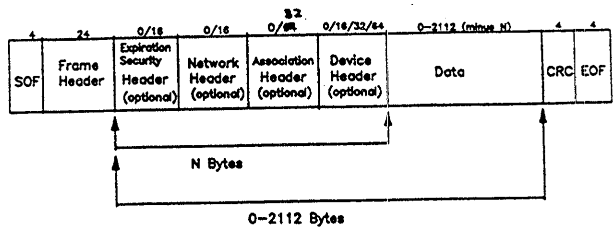

FCS Frame Header

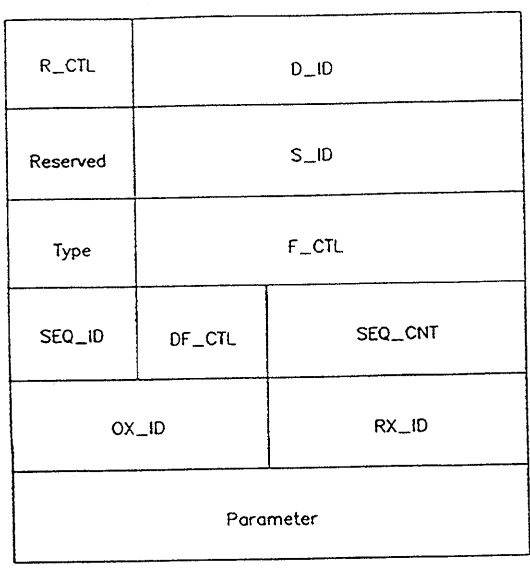

Fibre Channel Sequences and Frames

Fibre Channel interfaces can perform internal segmentation and reassembly.

The atomic unit above FC-2 is the "Information Unit."

- An IP or ARP datagram should be contained in a single Information Unit.
- The Information Unit will map to a single Sequence.

The atomic unit that FC-2 sends across a fiber is a Fibre Channel "frame."

- A Sequence can be fragmented into multiple Fibre Channel frames.
- A Sequence can be almost infinite (4 GBytes).

Note a terminology problem: An IEEE 802 Media Access Control (MAC) frame maps to a Fibre Channel Information Unit, not a Fibre Channel frame.

Fibre Channel Exchanges

Exchanges were designed to identify related Information Units (such as a command to an I/O device and associated data transfers and status replies).

Exchanges are half duplex, since that works well with devices.

Exchanges may be used unidirectionally.

To avoid the complexity of managing half-duplex flow, IP is expected to use unidirectional Exchanges (except during some error recovery situations).

Fibre Channel Operations

IP does not use operations.

Multiplexing Management

1. Operation may consist of multiple exchanges
2. An exchange consists of a single sequence at a time
3. A sequence consists of a uni-directional flow of frames
4. Frames denoted by X_IDs, SEQ_ID
5. Frame position in sequence denoted by SEQ_CNT, Relative Offset
6. Class 1: different operations, same destination (connection)
7. Class 2/3: different operations, different destinations

Login/Logout

1. Part of initialization
2. Each port "logs in" to all other ports, including Fabric
3. Exchange of information, e.g.:
   - Credit
   - Concurrent Sequences
   - Receive data field size
   - Service Options (e.g., which classes)

Fibre Channel Classes of Service

FC-PH defines three classes of service:

- Connection-oriented Circuit-switched ("Class 1"),
- Acknowledged Datagram ("Class 2"), and
- Unacknowledged Datagram ("Class 3").

Each Class of Service supports the same kinds of Exchanges, Information Units, and frames, and thus the same upper layer protocols.

Class 1 raises some service issues.

- Fabrics support Class 1 service by circuit-switching.
- A pair of N_Ports gets very good service to each other instantaneously but poor or no concurrent service to other N_Ports.
- Rules for deciding when to create and remove connections are needed.
- When Class 1 is appropriate, it probably provides the highest throughput.

Some FC-2 functions are not discussed.

- Link Level Control (ACK, Ready, Busy, Reject)
- Data Flow Control
- Segmentation, Flow Control, Recovery

FC-3 Summary

- FC-3 Functions are not yet well defined.
- FC-3 has a view of all operations in the node across all N_Ports.
- Functions include port selection:
  - Non-striping (use a single N_Port)
  - Striping (use mulitple N_Ports for higher bandwidth)
  - Multi-cast (could be useful for some IP functions)
- "IP and ARP on Fibre Channel" does not rely on FC-3 Functions.

PART F.

Fabric

Fabric Overview

- "Switching" - defines:

  1. The topologies used for routing of frames through an FCS network
  2. Concepts and characteristics of the various topologies
  3. Topologies include:
     - Point-to-point
     - Dynamic Switch
     - Broadcast Hub
     - Loop

Fibre Channel Addressing

FC-PH specifies 24-bit addresses ("N_Port Identifiers").

- Fabrics assign 24-bit addresses. Interfaces "log in" to a Fabric to learn their 24-bit addresses.
- A fabric assigns 24-bit addresses to optimize routing within the Fabric.
- (FC-PH specifies other mechanisms to provide long-life identifiers that uniquely identify interfaces.)

FC-PH also allows 64-bit "Process Associators."

- Some Fibre Channel interfaces use Process Associators to quickly locate the process that handles an Information Unit.

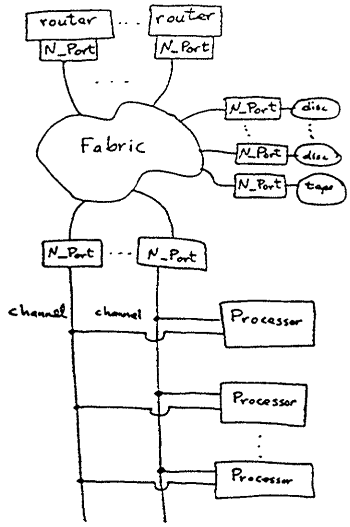

The fabric uses 24-bit addresses (N_Port Identifiers) to route from source N_Port to destination N_Port.

Process Associators may be used by an N_Port to quickly find where a "process" is, including which processor and where within the memory of the processor.

• Note: for redundancy, an IP address should map to several "hardware addresses."

PART G.

Status

X3T9.3 FCS Architecture Schedule

Best guess:

1. Current: FC-PH, 3.0 Available
2. 9/92: Start Public Review (4 months)
3. 1/93: End Public Review
4. 3/93: FC-PH re-issued
5. 4/93: Vote by X3T9 to forward new FC-PH for Public Review
6. 5/93: Start final Public Review (2 months)
7. 7/93: End final Public Review

Overlapped with above:

1. Continue Development
   - Specific Subjects
   - Addenda on base FC-PH
2. FC-3
3. All FC-4s

ANSI X3T9 Develops Interface Standards

X3T9.2 defines SCSI (Small Computer Systems Interface).

X3T9.5 defines FDDI (Fiber Distributed Data Interface).

X3T9.3 defines
- IPI (Intelligent Peripherals Interface).
- HIPPI (HIgh-Performance Parallel Interface), and
- Fibre Channel.

Standards Hierarchy

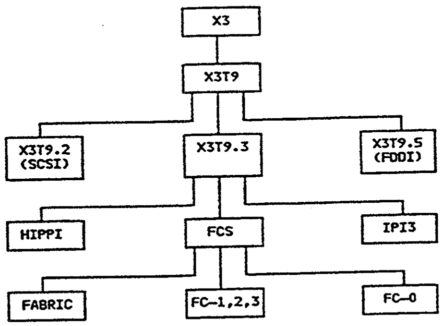

Fibre Channel Standards Efforts

Basic Fibre Channel Standards (ANSI X3T9.3)

- FC-PH (Physical and Signaling Interface)
- FC-EP (Enhanced Physical, includes FC-3)
  These are being considered for FC-EP: multicast, striping (parallel channels), hunt groups, and isochronous service.
- FC-IG (Implementation Guide)

Fabric Standards (ANSI X3T9.3)

- FC-FG (Fabric Generic requirements)
- FC-XS (Cross-Point Switch)
- FC-LT (Low-cost Topology)

Fibre-Channel "FC-4" Standards (various)

- IP and ARP on Fibre Channel (IETF)
- FC-IP: IEEE 802.2 Link Encapsulation (ANSI X3T9.3)
- FCP: SCSI on Fibre Channel (ANSI X3T9.2)
- FC-I3: IPI-3 on Fibre Channel (ANSI X3T9.3)
- FC-FP: HIPPI Framing Protocol on Fibre Channel (ANSI X3T9.3)
- FC-SB: Mainframe Peripherals (ANSI X3T9.3)

Notes:
1. "IP and ARP on Fibre Channel" depends only upon FC-PH.
2. The mnemonic FC-IP was assigned before liaison with IETF revealed that standards related to the Internet Protocol are developed by the IETF.

Status at Lawrence Livermore National Laboratory

- LLNL is the first Fibre Channel customer and is helping with development.
- In 1991 implementors agreed on FC-PH version 1.6 for prototyping. (Note: the Internet Draft corresponds to FC-PH version 3.0, the current version, not to FC-PH version 1.6.)
- Several vendors are developing parts.
- LLNL received prototype parts in August 1992 and is debugging. The parts were designed by Ancor Communications and include Fibre-Channel to VME interfaces and a 16-port crosspoint switch with a signaling rate of 266 Mbit/second.
- The VME interfaces are installed in SUN workstations. LLNL (Seth Abrahams) is writing the driver.

LLNL status as of November 13 follows.

- Telnet, FTP, and TTCP seem to run without causing crashes.
- Some non-TCP traffic does cause crashes.
- TTCP achieves 0.5 MByte/second (two percent of the limit).
- Some of the reasons for slow performance are understood.
- The system is (hopefully) running November 16-20 at the Supercomputer 92 Conference in Minneapolis, concurrent with this IETF meeting.

Why is LLNL Interested in Fibre Channel?

- We want a gigabit/second to the desktop, to visualize complex scientific simulations.
- We want colleagues to be able to interact with each other and with visualized data.
- We want to support many users concurrently. (About 1000 concurrently communicating pairs of interfaces.)
- We want to archive and transfer many large files.
- We want fast communication between heterogeneous special-purpose processors (to achieve speedups greater than linear by letting each processor do what it is specialized for).

Conceptual View of Gigabit Architectures

Functional Awareness	Architecture	Architecture Level	Features	
Access Method Device	IP / 802.2 / SCSI / IPI3 / MI / SBCCS	Upper Level Protocol	Commands Parameters	
SCP		System Architecture	Instructions Interrupts	
CSS	Striping Multicast	FC-3	Common Services Level	
CSS	Frame Formats Frame Headers (Addresses, IDs) LLC(ACK, RDY, BSY, RJT) Service Classes Segmentation/Reassembly Flow Control/Recovery	FC-2	Link Level	
CSS	8/10 Transmission Code Synchronization Loss of signal	FC-1		FC-PH
CSS	MM LED 132 Mbit / MM LED, MM SWLD, SM LWLD 266 Mbit / MM SWLD, SM LWLD 531 Mbit / SM LWLD 1.06 Gbit	FC-0	Connectors Components Speeds	

History of "IP and ARP on Fibre Channel"

- ANSI X3T9.3 decided to do the standard.
- During liaison with IETF, X3T9.3 decided to let IETF do the standard, but X3T9.3 would provide an initial draft.
- Yakov Rekhter began the draft about July 1992.
- ANSI X3T9.3 reviewed the draft
  August (Bellevue, Washington)
  September (Toronto, Ontario)
  October (Ft. Lauderdale, Florida)
- An Internet Draft was submitted in October 1992.
- IETF controls the document now.

On-line Information

Mail Groups

- fibre-channel-ext@think.com
  is used for general Fibre Channel issues including FC-PH.

  To (un)subscribe, send a request to
  fibre-channel-ext-request@think.com

- fc-ip-ext@think.com
  is used for IP (and IEEE 802) over Fibre Channel.

  To (un)subscribe, send a request to
  fc-ip-ext-request@think.com

Note: other specialized fibre channel mail groups are not listed here.

Anonymous FTP

Connect to

nsco.network.com

and explore within the directory

FC

IP and ARP over Fibre Channel (FC)

Yakov Rekhter
T.J. Watson Research Center
IBM Corporation
e-mail: yakov@watson.ibm.com

Design Objectives

- IP level interoperability between conformant implementations

Inside the Scope

- Mechanisms to exchange IP and ARP
- Constraints on FC-2 Frame Header parameters
- IP to N_Port Identifier mapping
- Fair access to node's resources

Outside the Scope

- Everything else:
  - ARP Server solution
  - IP Multicast
  - Network configuration and management
  - IEEE 802 MAC Layer bridging
  - Interaction with other FC-4s running over the same N_Port
  - Full support for IEEE 802.2 LLC

FC-2 Frame Header

- R_CTL field:
  - Routing bits – Device Data
  - Information Category – Unsolicited Data
- TYPE = IEEE 802.2 LLC/SNAP (LLC + SNAP encapsulation)
- Network Header is mandatory
  - Default – carries IP addresses
  - Recipient may ignore the content
- Other Headers (e.g. Association Header) are optional
  - Association Header must be present with a non-null Initial Process Associator

Login Parameters

- Fabric Login and N_Port Login are required to exchanged IP/ARP
- Obtaining N_Port Identifier is outside the scope
- No constraints on
  - Common Login Parameters
  - Parameters for Fabric Login
  - Parameters for N_Port Login

Exchanging IP/ARP packets

- One IP Packet ⟺ One Information Unit
- One ARP Packet ⟺ One Information Unit
- Exchange to pass Information Units
  - Only the Exchange Originator sends IP/ARP packets ⟹
    - * bidirectional traffic requires two Exchanges
  - One or more Exchanges between a given pair of N_Ports

LLC/SNAP Header

- LLC + SNAP encapsulation
- Used with IP over HIPPI, IEEE 802.3, IEEE 802.5
- Single FC-2 TYPE for both IP and ARP

Address Resolution

- < N_Port Identifier, Initial Process Associator > = hardware address
- Support multiple hardware addresses per single IP address
- Local Mapping (required to support)
- ARP Server (optional)
  - Well-known N_Port Identifier – "FFFFFC"
  - A Node registers with an ARP Server after Fabric Login

Fair Access

- Class 1
  - Limiting time of open Class 1 connection (500 milliseconds)
- Resources associated with an Exchange
  - Independent of a particular Class
  - Ability to terminate an Exchange (by either the Originator or the Responder)

MTU

- Single Information Unit – up to ($2^{32} - 1$) octets
- Maximum IP packet size – 65280 octets
  - Consistent with HIPPI
  - Single IP packet into 64 Kbytes buffer with up to 256 octets of overhead

What is a subnet ?

- The concept of a Region
  - Transitivity with respect to connectivity
  - A set of N_Ports such that any N_Port in the set can successfully complete the N_Port Login procedure with any other N_Ports in the set
  - All N_Ports within the set can directly exchange IP/ARP packets with each other
- An N_Port may belong to more than one region
  - Distinct IP address per region

Other issues

- An Exchange per TCP connection – optional
- Upper limit on the duration of a single Class 1 connection – 500 ms
- Class 1 connection for long Information Unit, Class 2 and 3 for the rest

CURRENT MEETING REPORT

Reported by Paul Tsuchiya/Bellcore

Minutes of the Inter-Domain Multicast Routing BOF (IDMR)

Agenda

1. Determine the Charter of the Group.
2. Go over the CBT (Core Based Trees), a proposal for scalable multicast routing.

The first IDMR BOF was held November 17th at the Washington, DC IETF meeting. It was chaired by Tony Ballardie of UCL. Tony amply demonstrated early on in the meeting that the English had best stick to dry humor, and leave the wacky stuff to the Americans (MP notwithstanding).

Concerning Agenda Item 1, it was agreed that there is a need for a new multicast protocol for inter-domain multicast, as the existing schemes do not scale well enough. Therefore, it was decided that:

1. A Working Group should be formed (the IDMR Working Group).
2. The Charter of the Group is to design a standard multicast routing protocol for inter-domain multicast routing.

Though there was no explicit call for consensus, it was assumed that Tony Ballardie would chair the Group, with Paul Tsuchiya as alternate Chair. I assume that it is still possible for people to volunteer to Chair the Group. Also, there was no consensus (for or against) that CBT should serve as the base text for the new IDMR protocol. On the other hand, no other proposals are on the table.

Concerning Agenda Item 2, the following concerns were raised about CBT.

- There was a concern that the current Internet-Draft didn't adequately specify the case where a node receives two joint requests before getting back the first ack.

- The choice of which major core to send terminate request messages to is not specified correctly.

- It was felt that there must be a way for detecting the case where there are two cores with uptree links on the same LAN.

- The format for the core list packet must be worked out, and the drawing in the CBT document is hard to understand.

- There was a concern that the text describing when to send a quit request was not complete.

- The notion of sending a quit some time after receiving a join ack (when changing parents) is no good (should send quit immediately).

- There was a lot of discussion about what to do when the link to the parent goes down. This whole issue needs to be worked out, but there seemed to be a general preference for flushing the whole tree below the break, with everything below subsequently rejoining.

Attendees

Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Tony Bates	t.bates@nosc.ja.net
Scott Brim	Scott_Brim@cornell.edu
Michael Collins	collinsms@es.net
Barbara Denny	denny@erg.sri.com
Hans Eriksson	hans@sics.se
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Shoji Fukutomi	fuku@furukawa.co.jp
Joel Halpern	jmh@network.com
Don Hoffman	don.hoffman@eng.sun.com
Dwight Jamieson	djamies@bnr.ca
Oliver Jones	oj@pictel.com
Paulina Knibbe	knibbe@cisco.com
Jim Knowles	jknowles@binky.arc.nasa.gov
John Krawczyk	jkrawczy@wellfleet.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Jim Perchik	perchik@athena.mit.edu
Thomas Pusateri	pusateri@cs.duke.edu
Benny Rodrig	4373580@mcimail.com
Henning Schulzrinne	hgs@research.att.com
Martha Steenstrup	msteenst@bbn.com
Tang Tang	tt@virginia.edu
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Ioannis Viniotis	candice@ececho.ncsu.edu

CURRENT MEETING REPORT

Reported by Philip Almquist/Consultant

Minutes of the Selection Criteria BOF (SELECT)

The objective of the Selection Criteria BOF was to develop consensus on a precise statement of the community's goals for a replacement for IP. The goal was to provide a yardstick against which the various proposals could be objectively measured to point up their relative strengths and weaknesses. Needless to say, this goal was far too ambitious to actually be achievable in the single session available.

The BOF began with a discussion of two previously written documents that presented goals for the new IP:

- draft-partridge-ipv7-criteria-00.txt (Partridge/Kastenholz)
- draft-iab-ipversion7-00.txt (IAB)

The two sets of goals were compared and contrasted. Craig Partridge and Barry Leiner verbally clarified and presented additional rationale for the goals presented in those papers. Lively discussion ensued, during which the Group modified the list of goals from the Partridge/Kastenholz paper.

In particular, the goals of architectural simplicity and globally unique identification of endpoints from the IAB's list were thought to be important enough to be added. Some participants pointed out that it was also important to add topological flexibility as a goal, noting that we certainly needed to understand whether any of the proposed next generation IP's would preclude topologies currently in use. Others pointed out that, although it may be hard to quantify, the proposals may embody differing amounts of technological risk, and that our criteria needed therefore to address risk. Some felt that performance needed to be a goal. Matt Mathis pointed out that different proposals may differ in how the pain of deployment is allocated among the levels of the networking food chain (backbones, midlevels, campus nets, end users), and emphasized that we are unlikely to successfully deploy any proposal in which some level receives little benefit from the new version of IP yet is expected to shoulder a large chunk of the pain. Finally, a number of people felt that IETF change control and freely available specifications have been critical to the success of the current IP, and that it is therefore important to consider to what extent each of the proposals preserves those features. The goal of providing usage accounting was dropped from the list of goals when it was pointed out that the accounting requirements specified in the Partridge/Kastenholz paper would be trivially met by any proposal.

After reaching near consensus on a list of goals, the Group sought to rank them in terms of importance. No real consensus was attained. The best summarization of this part of the meeting might be that whatever is chosen as an IP replacement must solve the scaling problem and must not be substantially inferior to the current IP in terms of other important attributes (security, manageability, robustness, etc.).

A third task, trying to refine each of the goals on the list into a statement of sufficient detail and precision that the extent to which a particular proposal met the goal could be fairly objectively determined, was skipped due to lack of time.

Craig Partridge and Frank Kastenholz agreed to revise their Internet-Draft to attempt to incorporate the results of the BOF.

Attendees

Roland Acra	acra@cisco.com
Philip Almquist	almquist@jessica.stanford.edu
Nagaraj Arunkumar	nak@3com.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Robert Austein	sra@epilogue.com
Tony Bates	t.bates@nosc.ja.net
Jordan Becker	becker@ans.net
Shiraz Bhanji	bhanji@gateway.mitre.org
Robert Braden	braden@isi.edu
Vickie Brown	brown@osi540sn.gsfc.nasa.gov
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.lkg.dec.com
Ken Carlberg	Carlberg@cseic.saic.com
Brian Carpenter	brian@dxcern.cern.ch
Charles Carvalho	charles@acc.com
Vinton Cerf	vcerf@cnri.reston.va.us
George Chang	gkc@ctt.bellcore.com
A. Lyman Chapin	lyman@bbn.com
Henry Clark	henryc@oar.net
Richard Colella	colella@osi.ncsl.nist.gov
David Crocker	dcrocker@mordor.stanford.edu
John Curran	jcurran@bbn.com
James Davin	davin@bellcore.com
Michael Davis	mad@spirit.clearpoint.com
Steve Deering	deering@parc.xerox.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Tim Dixon	dixon@rare.nl
Kurt Dobbins	dobbins@ctron.com
Jacques Dugast	dugast@issy.cnet.fr
Donald Eastlake	dee@ranger.enet.dec.com
Robert Enger	enger@reston.ans.net
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
William Fink	bill@wizard.gsfc.nasa.gov
Eric Fleischman	ericf@act.boeing.com

Peter Ford	peter@goshawk.lanl.gov
Shoji Fukutomi	fuku@furukawa.co.jp
Vince Fuller	vaf@stanford.edu
Heather Gray	heather@zk3.dec.com
Eugene Hastings	hastings@psc.edu
Robert Hinden	hinden@eng.sun.com
Erik Huizer	huizer@surfnet.nl
Ole Jacobsen	ole@interop.com
David Jacobson	dnjake@vnet.ibm.com
Dan Jordt	danj@nwnet.net
Frank Kastenholz	kasten@ftp.com
Dave Katz	dkatz@cisco.com
Mark Knopper	mak@merit.edu
Stev Knowles	stev@ftp.com
John Larson	jlarson@parc.xerox.com
Barry Leiner	leiner@nsipo.nasa.gov
Fong-Ching Liaw	fong@eng.sun.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Daniel Long	long@nic.near.net
Triet Lu	triet@cseic.saic.com
Gary Malkin	gmalkin@xylogics.com
Tracy Mallory	tracym@3com.com
David Marlow	dmarlow@relay.nswc.navy.mil
Matt Mathis	mathis@a.psc.edu
David Meyer	meyer@oregon.uoregon.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
John Moy	jmoy@proteon.com
Andy Nicholson	droid@cray.com
Erik Nordmark	nordmark@eng.sun.com
Peder Norgaard	pcn@tbit.dk
Andrew Partan	asp@uunet.uu.net
Craig Partridge	craig@bbn.com
Roy Perry	rperry@advtech.uswest.com
Yakov Rekhter	yakov@watson.ibm.com
Joyce K. Reynolds	jkrey@isi.edu
Henry Sanders	henrysa@microsoft.com
John Scudder	jgs@merit.edu
Erik Sherk	sherk@sura.net
Cris Shuldiner	cws@ftp.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@andr.ub.com
John Tavs	tavs@vnet.ibm.com
Marten Terpstra	marten@ripe.net
Richard Thomas	rjthomas@bnr.ca

Stuart Vance	vance@tgv.com
Warren Vik	wmv@i88.isc.com
Curtis Villamizar	curtis@ans.net
A. Lee Wade	wade@nsipo.nasa.gov
William Warner	warner@ohio.gov
James Watt	james@newbridge.com
Von Welch	vwelch@ncsa.uiuc.edu
Guy Wells	guy2@uswest.com
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov
Robert Woodburn	woody@sparta.com

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the Traceroute BOF (TRACERTE)

Agenda

- Discuss the need for a Traceroute protocol.
- Review the Internet-Draft.
- Determine if any additional information should be included.
- Consider alternate proposals (if any).
- Determine if a Working Group is needed.

Unfortunately, I believe I have the honor of chairing the first BOF to be attended by only one person, myself. I therefore declared myself a committee of the whole for the purpose of discussing the items on the Agenda.

The need for a Traceroute protocol is unclear. However, the Internet-Draft was unanimously approved as read.

No alternate proposals were put forth.

A working group is not needed.

It was decided that the Internet-Draft should be submitted for consideration as an Experimental Protocol.

Attendees

Gary Malkin	gmalkin@xylogics.com

Traceroute BOF

Gary Malkin

Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803
(617) 272-8140

gmalkin@xylogics.com


```
Mailing List:           traceroute@xylogics.com
Administrivia:  traceroute-request@xylogics.com
Archives:       xylogics.com:gmalkin/traceroute
Internet Draft:  draft-malkin-traceroute-01.txt
```


IP Option Format

0	2	Number	Length	ID Number
Originator IP Address				
Outbound Hop Count			Return Hop Count	

ICMP Traceroute Message Format

Type	Code	Checksum	
ID Number		Outbound Hop Count	Return Hop Count
Outbound Link Speed			
Outbound Link MTU			

Basic Operation

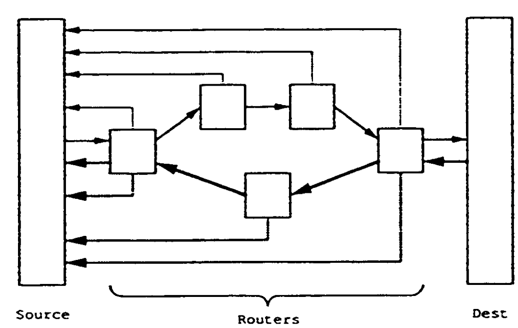

"ping" packet
"ping" response packet
ICMP Traceroute message (for "ping")
ICMP Traceroute message (for "ping" response)

2.2.1 Dynamic Host Configuration (dhc)

Charter

Chair(s):

Ralph Droms, `droms@bucknell.edu`

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

Internet-Drafts:

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

"Dynamic Host Configuration Protocol", 07/09/1991, R. Droms <draft-ietf-dhc-protocol-06.txt, .ps>

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

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

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell

Minutes of the Dynamic Host Configuration Working Group (DHC)

The Dynamic Host Configuration Working Group met twice in Washington. In the first meeting, the Working Group reviewed the state of the protocol specification. Ralph Droms described several recent changes to the specification documents, made in response to the IESG's review. Comments about the changes were posted to the host-conf mailing list and are available from the list archive in sol.cs.bucknell.edu:dhcwg/host-conf-archive. Two additional issues not previously addressed by the IESG were raised by Philip Almquist in an "in-the-hall" meeting: DHCP must permit the server to disallow access to network addresses by unauthorized clients, and DHCP servers should be able to provide client-specific network parameters; i.e., DHCP servers should not be required to provide the same parameters (e.g., DNS server) to all clients on a subnet.

The Group approved of the changes to the specification documents. Once additional changes are made, DHCP will be resubmitted to the IESG for consideration as a Proposed Standard.

There was a brief discussion of BOOTP/DHCP relay agent behavior relating to the insertion of the client's subnet mask in DHCP messages by the relay agent. The Group concluded that DHCP servers must be aware of the network topology and can, therefore, always determine the appropriate subnet mask for a DHCP message. Thus, there is no advantage in allowing relay agents to supply the subnet mask. The Group decided that BOOTP/DHCP relay agents are not allowed to insert a subet mask into BOOTP/DHCP messages. Walt Wimer will modify the BOOTP/DHCP relay agent document to reflect this decision.

The Working Group also discussed backwards compatibility with the use of the 'file' field in BOOTP. DHCP will continue to use the 'file' field as in BOOTP (except where overridden by the 'overload' option [option code 48]). DHCP will also use 'siaddr' to hold the address of the server the DHCP client is to contact for further configuration (e.g., a TFTP server from which the DHCP client may obtain a "boot file"). Walt will modify the BOOTP clarification document and Ralph will modify the DHCP specification to explicitly describe these uses of the 'file' and 'siaddr' fields.

Next, the Group embarked on a lengthy discussion about the use options and the interpretation of some options as "vendor-specific". The concern is that some vendors may have difficulty in obtaining allocation of option numbers from IANA for options that are specific to that vendor. The proposal was to define a "client type" option, and a range of options as "vendor-specific". The "vendor-specific" options would then be interpreted based on the "client type". For example, if a client identified itself as a "Bison Chip Computers" client by including a "client type" option with value "Bison Chip Computers", the "vendor-specific" options would then be interpreted according to "Bison Chip Computers" allocation of option values. Such a mechanism would give individual vendors freedom in allocating options as they desired without having to go to IANA for new options.

The Working Group agreed to define a "client type" option and took the proposal for "vendor-specific options" under advisement. The "client type" option will contain a variable length string of octets, to be interpreted by the server as describing, e.g., the client's manufacturer and configuration. There was a counter-argument to the "vendor-specific options" proposal that fewer than 50 of the available 128 options have been used to date (128-254 are reserved for "site-specific" options), so that the "vendor-specific" option mechanism may not be necessary.

Bob Gilligan suggested some modifications to Walt's BOOTP/DHCP clarification document to explicitly describe the interactions between clients and servers in networks that may have both BOOTP and DHCP servers. In particular, DHCP servers must be configurable to disallow the automatic allocation of network addresses in networks where clients may receive responses from both BOOTP and DHCP servers.

In its second meeting, the Working Group took up the issue of a server-server protocol to automate the replication and reallocation of network address bindings. Greg Minshall presented a specific proposal that would provide redundant allocation, redundant reacquisition of a previously allocated address and distributed extension of an existing lease.

Greg also mentioned the use of SNMP as a configuration tool once DHCP has provided sufficient configuration to the client to allow operation of a transport protocol. Ralph suggested that Steve Deering's work in identifying all of the configurable parameters cited in the Host Requirements documents should be forwarded to the appropriate MIB working groups for their consideration. The Working Group concluded that DHCP should be kept as lightweight as possible, deferring to other configuration mechanisms such as SNMP and TFTP wherever possible.

Attendees

Steve Alexander	stevea@i88.isc.com
James Allard	jallard@microsoft.com
Richard Basch	basch@mit.edu
J. Nevil Brownlee	nevil@aukuni.ac.uz
Matthew Busche	mtb@anchor.ho.att.com
Michael Davison	davison@fibercom.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Jon Dreyer	Jon.Dreyer@east.sun.com
Ralph Droms	droms@bucknell.edu
Robert Gilligan	Bob.Gilligan@eng.sun.com
Nat Howard	nrh@bellcore.com
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@ftp.com
Andrew Knutsen	andrewk@sco.com
Kent Malave	kent@bach.austin.ibm.com
Greg Minshall	minshall@wc.novell.com

Drew Perkins	ddp@andrew.cmu.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Fumio Teraoka	tera@csl.sony.co.jp
Jim Thompson	jim@tadpole.com
Walter Wimer	walter.wimer@andrew.cmu.edu

2.2.2 IP Address Encapsulation (ipae)

Charter

Chair(s):

Dave Crocker, dcrocker@mordor.stanford.edu

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

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

Internet-Drafts:

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

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

INTERIM MEETING REPORT

Reported by Dave Crocker/TBO

Minutes of the IP Address Encapsulation Working Group (IPAE)

This meeting took place on August 27, 1992 via Videoconference and was, essentially, a review session for a number of issues. The one item which was pursued further was a report from Steve Deering about Addressing.

Administrivia

Copies of the current versions of the specifications, Craig Partidge's BSD diffs, and a few other files are now at PARC.

Mike Conn, of MCI, has very gratiously offered to provide a teleconferencing bridge (telephone) for future meetings. This will allow those not able to go to a videoconferencing site to participate over the phone.

Addressing (Steve Deering)

Discussion about geographic-based addressing has gone in the direction of allowing provider-based addressing _also]_, to handle the early stages of the new addressing plan. It appears that to remain strictly geographic will require very considerable complexity inside the datagram routing service, since metropolitan areas are, in no way, guaranteed to have inter-vendor transfer sites (now dubbed 'Metropolitan Internet Exchange' or "MIX".)

There is a need to ensure that the primary addressing authority is independent of any provider.

There is also a continuing concern that the MIX concept requires sharing of customer information between competitors. The retort is that that information is discernible anyhow.

Discussions will continue and the next Addressing meeting will be September 11th.

Side note: The Group feels that the specifications need to be crystal clear about the philosophy that is driving their choices, to facilitate evaluation among the different proposals.

ACTION: (Crocker) upgrade specifications to emphasize end-user friendly and installed base friendly intent of IPAE.

There is intended to be support for "multi-homed" commonwealths. That is, a host may have more than one commonwealth ID. This might facilitate transition issues, such as from vendor-oriented addressing to geographic, but it still requires the ability to add and delete addresses. The question of the way to propagate such information is still open.

IPAE Options

At the previous meeting, the question of providing space for IPAE-level options was discussed and rejected. At this meeting, we reviewed the decision, with no one suggesting it be reversed.

IPAE Border Router Discovery

This was another review topic. In general, most of the techniques that are used to discover an IP first-hop router can be re-used to discover the IPAE first-hop (i.e., border) router. But John Moy suggested use of a fixed, logical address, written into the IPAE specification. This could then trigger an IPAE-ICMP Redirect, when a logical border router gets the first IPAE datagram.

A concern was raised that this scheme would have trouble if the user datagram is fragmented, along the way to the border router, and worse, the fragments traveled to different logical border routers. The conclusion was that fragmentation is relatively rare and this is yet-another strong vote for MTU Discovery. Further, multiple destinations result only from path-splitting or a transient problem. The former is something that can be limited, for the logical address, and the latter is "only" a transient problem.

Miscellaneous

ACTION: (Crocker) The specification needs to better detail the behavior of the _exit_ border routers (the last IPAE hop before the destination host.) More protocol mechanics.

ACTION: (Crocker) The specification should give an example of address handling, as IPAE datagram moves through the Internet.

ICMP

IPAE intends to permit permanent support for unmodified routers, within a commonwealth. This means that routers will be generating current (old- style) ICMP messages, which means ICMP messages without the full (IPAE, global) addresses of the originating host whose action triggered the ICMP datagram. The exit border router (last hop before the router generating the ICMP) has the task of turning the ICMP into an IPAE datagram, though it can't do that if it does not have the full global address of the originating host.

Only three options seem available:

1. Seek to have routers upgraded to generate larger ICMP datagrams, so that they will include the IPAE header from the originating host.

2. Have the Border router throw away ICMPs that it can't convert.

3. Have the Border router perform some sort of record-keeping of IPAE datagrams, so that it can match the returned 64-bits with a full IPAE global address.

The Group discussed these options. After appropriate (and large) amounts of illness-feeling, it was agreed that no other options seemed to exist and all of the listed options were terrible. Options 1 and 2 seemed like the most constructive and practical, with option 3 unlikely.

ACTION: (Champlin) Survey existing router behavior, to determine the size of ICMP datagrams they actually generate, to determine if the theoretical problem is real.

ACTION: (Crocker) Verify Host Requirements statements about ICMP size.

ACTION: (Crocker) Add relevant text to the specification (not Transition document) about this issue, including reasonable options.

INTERIM MEETING REPORT

Reported by Robert Hinden/SUN

Minutes of the IP Address Encapsulation Working Group (IPAE)

This meeting of the IPAE Working Group took place on September 24, 1992 in a Video Conference between Mt. View, California and Lincoln, Massachusetts

1. Reviewed Agenda Items.
2. Reviewed Action Items from Previous Meeting.
3. Reviewed Recent Work.

ICMP

There was a review of Greg Chesson's ICMP extensions and it was decided that an IPAE redirect was necessary, but that special destination unreachable was not needed.

ACTION: Greg will update work and write up mechanism to map ICMP error messages between commonwealths assuming that sender includes IPAE header in ICMP error message.

Addressing

Reviewed work of previous days addressing meeting.

```
Unicast Metro / Provider Address
---------------------------------

 1 1   6         24           32                 32
+-+-+------+------------+----------------+----------------+
| |0|      |   /Metro   | Site           |                |
|0|/| RSVD | City Code  |                |32bit IP Address|
| |1|      |   \Provider| Subscriber     |                |
+-+-+------+------------+----------------+----------------+
 M C

M (Multicast) bit is 0 for Unicast, 1 for Multicast

C (IP Compatibility) bit is 0 for IPAE destination, 1 for IP destination

Multicast (Compatible w/ current IP Multicast)
----------------------------------------------

 1 1   6         24           32                 32
+-+-+------+------------+----------------+----------------+
| | |      |                             |                |
```

```
|1|1| RSVD |          (unused)          |  32bit Class D |
| | |      |                            |    IP Address  |
+-+-+------+-----------+----------------+----------------+
 M C
```

Multicast (New Multicast Format)

```
 1 1   6      8                    80
+-+-+------+-----+----------------------------------------+
| | |      |Flags|                                        |
|1|0| RSVD |  +  |          Multicast Group ID            |
| | |      |Scope|                                        |
+-+-+------+-----+----------------------------------------+
 M C
```

ACTION: Steve Deering will write an IPAE addressing architecture and hold a follow-on addressing sub-group meeting to review writeup.

Commonwealth Router Discovery

Concluded that discovery mechanism proposed by John Moy/Proteon would be fine.

ACTION: Bob Hinden will get IP address assigned for Commonwealth router discovery.

4. Close Outstanding Technical Issues

DNS SUPPORT

The Group discussed where changes to DNS would be necessary and which DNS servers would have to support IPAE at each stage in the transition.

ACTION: Dave Crocker will bring DNS experts into the loop and write up details for a document.

Inter-Commonwealth Routing

The Group discussed and concluded that for provider based addressing, BGP4 with support for 96bit IPAE addresses should be used. Metro based routing some what harder.

ACTION: Steve Deering will write up routing approach.

ACTION: Bob Hinden and Yakov Rekhter will discuss adding support for 96-bit addresses to BGP4.

5. Transition Issues

32-BIT IP Address => Commonwealth Address Mapping

The Group discussed different approaches: Static table, pure DNS, and hybrid consisting of background DNS collection to local server + real time query by Commonwealth Router to local server.

ACTION: Bob Hinden will write up static approach.

6. Implementation Plans and Schedule

   SGI, Proteon, and Sun intend to build prototype IPAE implementations. The minimum test configuration needed was:

```
          |                                          |
   IP-----|                                          |------IP
          |                                          |
          |---------(IP)-----(IPAE)-----(IPAE)------|
          |                                          |
   IPAE---|                                          |----IPAE
          |                                          |
```

   Testing can be done without building any new infrastructure. It can be done using the existing Internet.

   An implementation subgroup will be formed. Bob Gilligan agreed to lead and coordinate.

   ACTION: Bob Gilligan will set up implementors meeting and develop testing plan.

Next Meetings

- October 9th, Video Conference (Mt. View and Lincoln)
- October 15th, Video Conference (Mt. View and Lincoln)

INTERIM MEETING REPORT

Reported by Dave Crocker/TBO

Minutes of the IP Address Encapsulation Working Group (IPAE)

This meeting was held at Sun Microsystems on October 8, 1992.

The major topic of the meeting has been documented in a note sent by Bob Hinden, concerning a change to the IPAE header, to make it be the same as the SIP header. (i.e., make IPAE = IP+SIP, plus transition rules.) This was a somewhat unexpected turn of events, except in hindsight. A number of forces seem to have been moving the Group in this direction and there was a very strong feeling, by the end of the meeting, that this change vastly cleans up the entire scenario for the Internet, giving it the least transition pain and the most amenable longer-term protocol, since it is the closest to current IP AND it has a mechanism for adding new services (via its own mini-layer.)

Other topics discussed included:

- ICMP
- Router Table Size

ICMP

The 64-bit data limit for ICMP continues to be a problem. The Group discussed more about the handling of ICMP messages sent by interior, unmodified routers, which therefore contain only the within-commonwealth IP addresses of the interior router and the border IPAE router and don't have the full IPAE address of the originating host available.

It had generally been believed that this was an unfortunate, but not serious, problem. It was then observed that it _is_ a significant problem for MTU Discovery. The originating host really does need to get the ICMP feedback.

The Group adopted the framework that a commonwealth which does IPAE/IP tunneling – i.e., the interior routers are not IPAE knowledgeable – can be viewed much the same as IP over X.25, with the border routers treating the commonwealth as an underlying datalink environment. Hence, feedback from interior routers is like feedback from interior X.25 packet switches. We would not expect those raw messages to be forwarded back to the originating host.

We would expect the border routers to record the feedback and translate it. In this case, this means that the border router needs to cache MTU information about IP addresses inside its commonwealth. When it gets an IPAE datagram, it needs to check its size against the cache (cache = dest IP addr + MTU) and either fragment the datagram or send back an ICMP Too Big.

Basic language for the specification is: IPAE routers which are the IP recipients of IP/ICMP messages must cache "Can't Fragment" ("Too Big").

Router Table Size

The Group did an extended case analysis of the current and projected sizes for three different router tables: The Source Information Base is the raw stuff that comes in from the routing protocol(s). The Real-Time Table is used for doing that actual data-handling of actual packets. The Policy table is whatever set of contingent rules are needed to turn the first table into the second. Since the Group ran into some nomenclature confusion during the discussion, Dave Crocker has intentially not used more typical terms for the tables.

Note that the IPAE secton is divided into two, since the border routers need to maintain a set of IPAE routing tables as well as a set of IP routing tables (for the commonwealth.)

	SOURCE INFO BASE	REAL-TIME TABLE	POLICY
	(Variable, xmit + storage overhead)	(Variable, compute + storage overhead)	(Static)
Now:	All nets*neighbors	All nets	All nets
IPAE:		(Same as Source Info Base, but without the "* neighbors" component)	All nets (includes the IPAE/IP address map needed during transition while IP addresses are still unique)
IP:	Attached cwlth nets		
IPAE:	CWlth hierarchy, only as needed. (e.g., [all countries + attached metro/provider] * neighbors)		

CURRENT MEETING REPORT

Reported by Dave Crocker/TBO

Minutes of the IP Address Encapsulation Working Group (IPAE)

IPAE seeks to facilitate Internet conversion to a replacement for the current Internet Protocol (IPv4) by attending to transition ease of the installed base. The primary impetus for replacement of IPv4 is its limited address space, containing too few bits and having too little structure for very large-scale global routing tables. IPAE's technical approach is to encapsulate the new addressing information inside old IPv4 headers, so that the new information can transit unmodified networks.

After the Boston (24th) IETF meeting, the IPAE Working Group held a series of face-to-face meetings in Mountain View, California and usually with video conference hookups to Cambridge, Massachusetts. Telephone conferencing bridges also were available. EMail activity during this time was relatively limited.

During the time of these interim meetings, Steve Deering, who also had been participating in IPAE meetings, developed a preliminary specification for a Simplified IP (SIP) which streamlined the classic IP header, including addresses of 64 bits. Over the course of the interim meetings, IPAE then repositioned itself to focus exclusively on the question of transition from IPv4 to the new IP, with detailed specification of the transition focusing on use of SIP as the final stage of deployment.

Working documents of the Working Group are maintained in the at host 'parcftp.xerox.com', in directory 'ip-encaps'. Prior to the Washington IETF meeting, the IPAE specification was updated in the internet-drafts directory of the Internet Repository. Its filename is 'draft-crocker-ip-encaps-01.txt'.

November 19th

The IPAE session divided into a small amount of review about IPAE, some discussion of a few open technical issues, and then an extended presentation (attached) about implementation experience, by Erik Nordmark of Sun and Ron Jacoby of SGI. The Sun implementation is a full IPAE/SIP set of functionality, at the internet layer, and has been tested with SIP-SIP, SIP/IPAE-IPAE/SIP, and SIP-IP interactions. *Editor's Note (md): An ASCII version of the IPAE functional diagram is available via ftp under ipae-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Noteworthy comments about his technical work, on a Unix System V streams base:

- Converting to 64-bit address was somewhat tricky. (Author's Note: Comments from an earlier effort with BSD code suggests that the effort was not trivial, but also was not too difficult, though addresses larger than 96 bits may pose a problem.)

- Adding re-assembly code was easy, since it is identical to IPv4's reassembly. It took about a half a day's effort.

- Initial throughput testing (on prototype code) showed performance about the same as for IPv4. It must be stressed that the testing was by no means thorough, nor was the code tuned.

Open technical issues for IPAE include Domain Name System (DNS), network management and routing enhancements. DNS changes appear to be generic for any addressing upgrade, as do the general requirements for changes to MIB variables, and use of a routing protocol. That is, conversion to larger addresses carries a requirement for a substantial set of changes to Internet components. IPAE (and SIP) appear to impose no special concerns for this.

SIP uses addresses which have an IP address in the lower 32-bits. There was a suggestion that sites add a DNS entry which contains the UPPER 32-bits, so that hosts would have easy access to that information via the DNS. (Author's note: This presumes a direct relationship between IP network addresses and DNS domains, which is not required by the DNS technology.)

In questions from the audience, there was some tension between facilitating transition, versus ensuring a high-quality end-point. There also was a question of the group that would have to absorb the pain of a transition, with the choice being end-users, versus system operators. (Author's note: The term "end-user" is ambiguous, since local-net administrators are end-users, relative to larger service providers. Hence, there is a range of targets for absorbing transition pain.)

Vint Cerf observed that users aren't interested in the question; they simply should not be the ones to do the absorbing.

Implementation of IPAE and SIP

Erik Nordmark
Bob Gilligan

- Goals
- Implementation environment
- Interface changes
- Transport changes
- The routing table
- The integrated internet layer
- Application compatibility
- Problems and non-problems
- Status

Goals

- Host, router and border router functionality
- IPv4, IPAE and SIP
- A single routing table for IPv4 and SIP
- Minimize changes to the transport protocol implementations
- No DNS changes (yet)

Implementation Environment

- SunOS 5.X
- Streams TCP/IP
- Transport/internet interface carries complete IPv4 headers
- TCP and UDP checksum are performed in IP
- Integrated routing table and ARP cache
- Iterative routing table lookup

Interface Changes

Kernel interface:

- Use 8 unused bytes in sockaddr_in structure
- Currently transports check if 32 or 64 bit address

Internal interfaces:

- All addresses are 64 bits
- 32 bit addresses: high-order 32 bits all zero

Transport/internet interface:

- Currently: passes IPv4 + SIP header

Transport Changes

- All addresses are 64 bits in TCP and UDP
- Header templates have IPv4 + SIP headers
- Pseudo-header checksum depends on C bit in remote address
- Connection identification depends on C bit in remote address
- Changes to interactions with the routing table
- Handle four different formats of the headers contained in ICMP error messages

The Routing Table

- Addresses and masks are 64 bits
- One SIP table plus multiple IPv4 tables merged together
- Iterative lookup:

site route:	dest => next hop border router
network route:	border router => next hop gateway
interface route:	next hop => interface

- Results in a fully resolved entry:

fully resolved:	dest => interface, link layer header, [border router IP address]

- IPv4 to SIP address mapping table in routing table:

mapping "route":	IPv4 network => site prefix

Routing Table Example

- Host address: 0000:0002:129.144.248.47
- Lookup destination 0000:0009:36.8.0.11

9:36.8.0.11 =>	site route for 9:0, border router 2:129.150.254.2
2:129.150.254.2 =>	network route for 2:129.150.0.0, next hop 2:129.144.248.1
2:129.144.248.1 =>	interface route for 2:129.144.248.0

- Result: a fully resolved routing entry

9:36.8.0.11 =>	border router IP address 129.150.254.2, link layer header for 2:129.144.248.1

The Internet Layer

- Currently passes around IPv4+SIP headers internally
- SIP/IPAE functions:
  - Encapsulate SIP in IP (all)
  - Decapsulate SIP from IP (all)
  - SIP fragmentation and reassembly (all)
  - SIP forwarding (router)
  - Convert IPv4 to SIP (border router)
  - Convert SIP to IPv4 (border router)
- New tests:
  - Send IPv4 or SIP?
  - Need to encapsulate SIP in IPv4?
  - Received IP version 6?
  - Received IPAE protocol?

Application Compatibility

- Binary and source compatibility
- Separate gethostbyname() routine for 64 bit addresses
- New "address family" for bind, connect and sendto
- Transport/socket layer detect old vs. new application
- TI-RPC applications for free
- "Old" applications cause the host to act as an IPv4 host
- Binding to an IPv4 address - don't care about high-order address bits

Problems and Non-Problems

- 64 bit changes to routing table
- ICMP pseudo-header checksum interaction with fragmentation when border routers translate between IPv4 and SIP fragments
- Generality: IPv4 source route options and SIP source routes
- SIP fragmentation/reassembly: one afternoon
- Size of STREAMS modules (including debugging code):
  IP: +20%, TCP: +5%, UDP: +3%

Performance

- TCP throughput comparison:

Measurement	SunOS 5.1		Prototype	
	IPv4	IPv4	IPAE	SIP
Loopback	19.3 Mbps	18.6 Mbps	-	19.4 Mbps
Ethernet	8.9 Mbps	8.9 Mbps	8.8 Mbps	8.9 Mbps

Status

Working:

- IPv4, IPAE and SIP packet formats
- Host, router and border router
- telnet and ping
- TCP and UDP.
- SIP fragmentation and reassembly
- Translation IPv4 <-> SIP at border router including translating fragments

To do:

- SIP source routing and source routing translation
- SIP redirect
- Performance
- Applications
- DNS

2.2.3 IP over AppleTalk (appleip)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done	Post an Internet-Draft the current set of protocols used to connect Macintoshes to IP internets.
Done	Submit the AppleTalk MIB to the IESG for consideration as a Proposed Standard.
Jan 1993	Submit the IP over Appletalk document to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

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

"AppleTalk Management Information Base II", 12/21/1992, S. Waldbusser, K. Frisa <draft-ietf-appleip-mib2-00.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)

AppleTalk Systems Interoperability Group (ASIG)

Greg Minshall is setting up the ASIG. The IETF can't take on AppleTalk due to size and logistical issues, however, they've given us rooms to talk about AppleTalk only issues. The general ASIG issues include general approaches to routing.

AURP - Alan Oppenheimer

UDP port 387 is the official AURP port. All current Apple prototypes are all still 200. When it ships should use 387.

ATCP - Brad Parker

RFC 1378 gained status as a Proposed Standard three weeks ago. It can be moved to Draft Standard after some comments are resolved. These are related to the various combinations of router to router, end-system to router, and end system to proxy forwarding agent option negotiations in the implementation notes appendix.

MacIP - Chris Ranch

The document was posted to the internet-drafts directory late, but we went through it on a line by line basis anyway. We were able to relegate some implementation specific notes and recommendations from the protocol specification, and move them to the notes and recommendations sections. Another Internet-Draft will be published by the end of the year.

MIB Views- Steve Waldbusser

Steve showed how MIB views can be used to solve the configure time and real time management issue. This was very informative, as there is significant interest in configuration issues.

Other MIB Issues - Karen Frisa

The importance of compatibility with RFC1243 (old consoles should still work) was brought up as a goal for the MIB variable trimming exercise.

The AppleTalk MIB+ was discussed and for the most part closure was reached on all issues. 98 variables were trimmed from 230 and an Internet-Draft should be published shortly.

Karen will evaluate whether it is prudent to deprecate RFC1243 variables in MIB+ or. whether it makes better sense to leave them as they are.

The issues of how to do zone changes with this MIB was raised. Discussion on possible solutions will take place on the mailing list.

Greg Bruell brought up the issue of per port statistics. He reports that gathering statistics in this fashion makes problem diagnostics easier. This would require substantial changes to the MIB at this date and was tabled for further discussion and possibly another version of the MIB.

Attendees

David Ballowe	ballowe@compatible.com
Gregory Bruell	gob@wellfleet.com
Steve Bucey	sabucey@ns.pacbell.com
Philip Budne	phil@shiva.com
Wesley Craig	wes@umich.edu
Peter DiCamillo	Peter_DiCamillo@brown.edu
David Dubois	dad@pacersoft.com
Richard Ford	72510.553@compuserve.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Ken Hirata	khirata@emulex.com
Bob Jeckell	rrj@3com.com
Charley Kline	cvk@uiuc.edu
Grisha Kotlyar	grisha@farallon.com
Joshua Littlefield	josh@cayman.com
Dan Magorian	magorian@ni.umd.edu
Evan McGinnis	bem@3com.com
Greg Merrell	merrell@caldec.enet.dec.com
Daniel Myers	dan@nsd.3com.com
Alan Oppenheimer	oppenheime1@applelink.apple.com
Brad Parker	brad@fcr.com
Christopher Ranch	cranch@novell.com
Mike Ritter	mwritter@applelink.apple.com
Michael Safly	saf@tank1.msfc.nasa.gov
Paul Serice	serice@cos.com

William Simpson	Bill.Simpson@um.cc.umich.edu
Timon Sloane	timon@rahul.net
Robert Snyder	snyder@cisco.com
John Veizades	veizades@apple.com
A. Lee Wade	wade@nsipo.nasa.gov
Jonathan Wenocur	jhw@shiva.com
Gerry White	gerry@lancity.com
Peter Wilson	peter_wilson@3com.com
Kiho Yum	kxy@nsd.3com.com

2.2.4 IP over Asynchronous Transfer Mode (atm)

Charter

Chair(s):

Robert Hinden, hinden@eng.sun.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

Done	First Meeting. Establish detailed goals and milestones for Working Group.

Done — Post an Internet-Draft for a mechanism for IP over ATM. (Multi-Protocol Interconnect over ATM AAL5)

Jan 1993 — Submit the Multi-Protocol Interconnect over ATM AAL5 to the IESG as a Proposed Standard.

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

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

Internet-Drafts:

"Multiprotocol Interconnect over ATM Adaptation Layer 5", 06/12/1992, Juha Heinanen <draft-ietf-atm-multipro-05.txt>

CURRENT MEETING REPORT

Reported by Bob Hinden/Sun

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

Agenda

- Review Latest draft of "Multiprotocol Interconnect over ATM Adaptation Layer 5" by Juha Heinanen.
- Reach a "rough" consensus on one multiplexing approach.
- Talk on Q.93B and ATM Signaling by Fong-Ching Liaw.
- Talk on Internet/ATM Signaling Requirements by Tom Lyon.

Meeting Summary

During the first session on Monday, the Working Group finalized its approach for transporting Internet datagrams over ATM as described in the "Multiprotocol Interconnect over ATM Adaptation Layer 5" Internet-Draft written by Juha Heinanen. The document describes two approaches to encapsulating datagrams in ATM. These are: Virtual Circuity (VC) Based Encapsulation, and Multiplexing Encapsulation.

The current draft of the document describes two approaches to multiplexing encapsulation. The two approaches proposed are roughly equivalent. They differ in the manner used to identify the protocols being encapsulated. One uses NLPID (from Frame Relay) and the other uses LLC's (from 802.x LAN).

A long discussion centered around the question of whether there should be one or two multiplexing approaches. The Working Group voted by a two to one margin that one approach was preferable. After another long discussion a vote was held to select one of the approaches. The LLC approach was the clear winner of this vote. As a result of the these votes the current draft document will be revised to only describe the VC based encapsulation and the LLC multiplexing encapsulation in the main section of the document.

The Working Group also reached a consensus to submit its proposal to the Internet Engineering Steering Group (IESG) for entry into the Internet Standards as a Proposed Standard.

This work is an important step in the work necessary to insure that devices (e.g., hosts, routers, etc.) using Internet protocols (e.g., TCP/IP, ISO CLNP, Appletalk, etc.) can interoperate over ATM networks.

On Tuesday, Fong-Ching Liaw presented a talk on Q.93B and ATM signaling. Her talk described the current state of ATM signaling in CCITT and the ATM Forum. A number of issues were uncovered that relate to signaling parameters necessary for carrying Internet traffic. Fong-Ching Liaw and Brian Lyles agreed to write a draft paper describing these issues and make recommendations regarding how they can be handled. After review by the Working Group, this document will be published as an RFC and sent to the CCITT and the ATM Forum.

During Wednesday's session, Tom Lyon presented a talk on Internet and ATM Signaling Requirements. The talk covered goals for ATM, addressing requirements, header avoidance, connectionless service, multicast, and routing.

Fred Sammartino, the President of the ATM Forum, attended the last day of the Working Group and asked that an official liaison be formed between the ATM Forum and the IETF IP over ATM Working Group. The purpose of this liaison is to facilitate the exchange of protocols and ideas between the two organizations.

Attendees

George Abe	abe@infonet.com
Roland Acra	acra@cisco.com
Randall Atkinson	atkinson@itd.nrl.navy.mil
Floyd Backes	backes@dsmail.enet.dec.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
William Barns	barns@gateway.mitre.org
Ken Benstead	kbenstead@coral.com
Shiraz Bhanji	bhanji@gateway.mitre.org
Dean Blackketter	deanb@apple.com
Fred Bohle	fab@interlink.com
Rebecca Bostwick	bostwick@es.net
Daniel Brennan	dmb@teleoscom.com
Caralyn Brown	cbrown@wellfleet.com
Ross Callon	callon@bigfut.lkg.dec.com
Brian Carpenter	brian@dxcern.cern.ch
Charles Carvalho	charles@acc.com
Brett Chappell	bchappe@relay.nswc.navy.mil
Dilip Chatwani	dilip@synoptics.com
Szusin Chen	szusin.chen@eng.sun.com
Chi Chong	cchong@synoptics.com
Jim Christy	christy@apple.com
George Clapp	clapp@ameris.center.il.ameritech.com
Robert Cole	rgc@qsun.att.com
Michael Collins	collinsms@es.net
James Conklin	jbc@bitnic.educom.edu
Michael Conn	4387451@mcimail.com

Bruce Davie	bsd@bellcore.com
Michael DeAddio	deaddio@thumper.bellcore.com
Barbara Denny	denny@erg.sri.com
Art Dertke	dertke@gateway.mitre.org
Kurt Dobbins	dobbins@ctron.com
Jacques Dugast	dugast@issy.cnet.fr
Robert Enger	enger@reston.ans.net
Steven Fellini	steve@helix.nih.gov
William Fink	bill@wizard.gsfc.nasa.gov
Eric Fleischman	ericf@act.boeing.com
James Forster	forster@cisco.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Mike Goguen	goguen@src.dec.com
Heather Gray	heather@zk3.dec.com
Terry Gray	gray@cac.washington.edu
Paul Griffiths	griff@chang.austin.ibm.com
Tony Hain	alh@es.net
Joel Halpern	jmh@network.com
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Eugene Hastings	hastings@psc.edu
Ken Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Robert Hinden	hinden@eng.sun.com
Jonathan Hsu	brenda@penril.com
Kathleen Huber	khuber@bbn.com
David Husak	dave@synnet.com
Phil Irey	pirey@relay.nswc.navy.mil
David Jacobson	dnjake@vnet.ibm.com
George Kajos	kajos@coral.com
Charley Kline	cvk@uiuc.edu
Andrew Knutsen	andrewk@sco.com
Mark Laubach	laubach@hpl.hp.com
David LeRoy	dleroy@mitchell.cit.cornell.edu
Fong-Ching Liaw	fong@eng.sun.com
David Lin	lind@janus-ccm.zenith.com
Bryan Lyles	lyles@parc.xerox.com
Andrew Malis	malis@bbn.com
Tracy Mallory	tracym@3com.com
Allison Mankin	mankin@gateway.mitre.org
George Marshall	george@adaptive.com
Jun Matsukata	jm@eng.isas.ac.jp
Donald Merritt	don@brl.mil
Kathleen Nichols	nichols@apple.com
Erik Nordmark	nordmark@eng.sun.com

Ari Ollikainen	ari@es.net
Hilarie Orman	ho@cs.arizona.edu
Joe Pagan	jrp@afterlife.ncsc.mil
Craig Partridge	craig@bbn.com
David Piscitello	dave@sabre.bellcore.com
Thomas Pusateri	pusateri@cs.duke.edu
Bala Rajagopalan	braja@qsun.att.com
Greg Ratta	gregra@cc.bellcore.com
Tony Richards	
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Guenter Roeck	roeck@conware.de
Timothy Salo	tjs@msc.edu
Fred Sammartino	fred@eng.sun.com
Jim Scott	scott@kali.enet.dec.com
Allyson Showalter	allyson@nsipo.arc.nasa.gov
Dean Skidmore	deans@ralvmiz.vnet.ibm.com
Keith Sklower	sklower@cs.berkeley.edu
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@andr.ub.com
Marco Sosa	mxs@sabre.bellcore.com
Scott Stein	scotts@apple.com
Terrance Sullivan	terrys@newbridge.com
Sally Tarquinio	sallyt@gateway.mitre.org
Kaj Tesink	kaj@cc.bellcore.com
Alice Wang	alice.wang@eng.sun.com
Von Welch	vwelch@ncsa.uiuc.edu
Guy Wells	guy2@uswest.com
Gerry White	gerry@lancity.com
Peter Will	will@isi.edu
Douglas Williams	dougw@ralvmg.vnet.ibm.com
Kirk Williams	kirk@sbctri.sbc.com
Liang Wu	ltw99@bellcore.com
Chin Yuan	cxyuan@pacbell.com
John Zalubski	zalubski@nic.ddn.mil
Paul Zawada	Zawada@ncsa.uiuc.edu

Q.93B and ATM Signaling

Fong-Ching Liaw
Sun Microsystems
November, 1992

Overview

1. Signaling and Q.93B basics
2. Q.931, Q.93B, and ATM Forum signaling
3. Status of ATM Forum and CCITT

Where do we need signaling ?

- host to switch (or User to Network Interface, UNI)
- switch to switch interface (or ISSI)
- network to network interface (or NNI)

Signaling protocol - Q.93B

- CCITT recommendation for Broadband ISDN first release UNI signaling work
  - international standard
- ATM Forum signaling work baseline document
  - seamless signaling between public and local networks
- A derivative of N-ISDN Q.931 signaling protocol
  - add ATM specific parameters, and remove ATM un-related procedure/parameters

CCITT Q93B signaling capabilities

- point to point connection setup
- out of band signaling, carried in VC number 5

Q.93B protocol

- messages
- procedures (call establishing, clearing, error handling)
- per virtual circuit state machine (~10 states)
- per virtual circuit timers (~ 9 timers)
  - 10 timers on user side (ATM Forum)
  - 9 timers on network side (ATM Forum)

Q.93B message flow (call establishing)

Q.93B Messages
(Call establishing)

- Setup
- Call proceeding
- Connect, ConnectAck

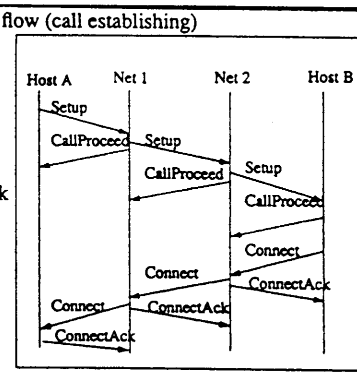

Q.93B Setup Message
- Bear capability
- User cell rate
- AAL parameter
- QoS
- Called party number
- VPI/VCI (O)
- B-LLI (layer 1-3 protocol information)

Q.93B state transition diagram (call establishing)

Q.93B states (call establishing)

- U0 (null)
- U1 (call initiated)
- U3 (outgoing call proceeding)
- U10 (active)

Procedures

- outgoing setup procedure (user)
- incoming setup procedure(user)
- outgoing setup procedure (nwk)
- incoming setup procedure(nwk)

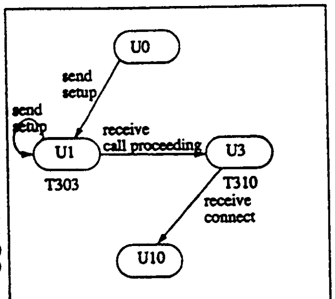

outgoing setup procedure (host A)

Q.93B state transition diagram (call establishing)

Q.93B states (call establishing)

- U0 (null)
- U6 (call present)
- U9 (incoming call proceeding)
- U10 (active)

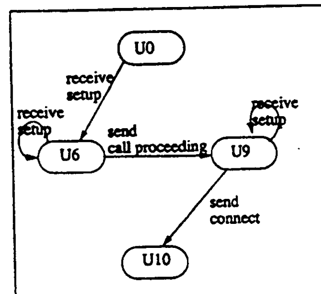

incoming setup procedure (host B)

Q.93B message flow (call clearing)

Q.93B messages (call clearing)

- disconnect
- release
- release complete

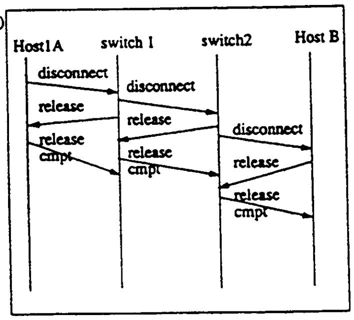

CCITT Q.93B procedure principles

- a point to point protocol
- 3-way handshake for connection setup
- 3-way handshake for connection clearing
- retransmit Setup and Release messages
- duplicate Setup and Release messages are detected by incompatible Q.93B states, duplicated message is discarded.

Q.93B message

- protocol discriminator (1 byte)
- call reference (2-* bytes, max. 15)

  to correlate the messages and the call/connection
- message type (1 byte)
- message length (2 bytes)
- information elements (variable length)

  parameters which convey information between

  requestor and requestee

note: no checksum or sequence number

3.1.11 SETUP: This message is sent by the calling user to the network and by the network to the called user to initiate B-ISDN call and connection establishment.

SETUP message content

Message type: SETUP   Direction: both
Significance: global

Information element	Reference	Direction	Type	Length
Protocol discriminator	4.2	both	M	1
Call reference	4.3	both	M	~~2~~ 3
Message type	4.4	both	M	1
~~Message length~~	~~4.4~~	~~both~~	~~M~~	~~1~~-2
~~Sending complete~~	~~4.5~~	~~both~~	~~O(Note 16)~~	~~1~~-3
~~B-Bearer capability~~	~~4.5~~	~~both~~	~~M~~	~~?~~
3.2 ← ~~Bearer capability~~	~~Q.931~~ 4.6	~~u→n~~	~~O~~	~~4-? (note 15)~~
3.2 ← Progress indicator	4.~~5~~6	both	O(Note 1)	~~2-4~~ 3-5
~~Network-specific facilities~~	~~4.5~~	~~both~~	~~O(Note 2)~~	~~2-*~~
~~Display~~	~~4.5~~	~~n→u~~	~~O(Note 3)~~	~~Note 4~~
End-to-end transit delay	4.5	both	O(Note 5)	~~2-11~~ 3-12
ATM ~~traffic descriptor~~ user cell rate	4.5	both	M	~~4-10~~ 7-11
~~{~~AAL parameters}	4.5	both	~~O~~ M	?
Calling party number	4.5	both	O(Note 6)	~~2~~3-*
Calling party subaddress	4.5	both	O(Note 7)	~~2~~3-2~~3~~4
Called party number	4.5	both	O(Note 8)	~~2~~3-*
Called party subaddress	4.5	both	O(Note 9)	~~2~~3-2~~3~~4
Transit network selection	4.5	u→n	O(Note 10)	~~2~~3-*
~~Repeat indicator~~	~~4.5~~	~~both~~	~~O(Note 11)~~	~~1~~
3.2 ← Low layer compatibility	4.~~5~~6	both	O(Note 12)	~~2-16~~
~~B-High layer compatibility~~	~~4.5~~	~~both~~	~~O~~	~~?~~
3.2 ← High layer compatibility	4.6	both	O(Note 13)	~~2-4~~ 3-5
3.3 ← ~~User-user~~	~~4.5~~	~~both~~	~~O(Note 14)~~	~~2-131~~
Connection Identifier	4.5	both	O	~~2-~~8
~~Message compatibility~~	~~4.5~~	~~both~~	~~M~~	~~?~~

inserted

B-LLI

Note 1- FFS

~~Note 2- Included by the calling user or the network to indicate network-specific facilities information (see Annex E of Recommendation Q.931).~~

~~Note 3- Included if the network provides information that can be presented to the user.~~

~~Note 4- The minimum length is 2 octets; the maximum length is 82 octets.~~

Note 5- May be omitted in the user-to-network direction, if the calling user accepts default values for this quality of service parameter. Always included in the network-to-user direction to indicate cumulative end-to-end transit delay to the called user.

Note 6- May be included by the calling user or the network to identify the calling user.

Note 7- Included in the user-to-network direction when the calling user wants to indicate the calling party subaddress. Included in the network-to-user direction if the calling user included a Calling party subaddress information element in the SETUP message.

Note 8- The Called party number information element is ~~included~~ mandatory by the user to convey called party number information to the network. The Called party number information element is included by the network when called party number information is conveyed to the user.

Q.93B information element

- information element identifier (1 byte)
- information element length (1 byte)
- octet groups

Octet group coding rules
- extension bit at bit 8,
    0 if octet N continue through next byte (Na, Nb ...),
    1 if last byte of the octet group N
- optional octets are marked with asterisks (*)
- and other rules

Bits

8	7 6 5 4 3 2 1	Octet
0	Broadband Low Layer Information 1 0 1 1 1 1 1 Information element identifier	1
0/1	Length of the Broadband LLI contents	2
1 ext	Coding standard | IE Instruction Field: Flag | Res. | Spare | Action Ind.	3
1 ext	0 1 layer 1 id | user information layer 1 protocol	4*
0/1 ext	1 0 layer 2 id | user information layer 2 protocol	5*
0/1 ext	Mode | 0 0 0 Spare | Q.933 use	5a* (Note 5)
1 ext	User specified layer 2 protocol information	5a* (Note 6)
1 ext	Window size (k)	5b* (Note 5)
0/1 ext	1 1 layer 3 id | user information layer 3 protocol	6*
0/1 ext	Mode | 0 0 0 0 0 Spare	6a* (Note 7)
1 ext	User specified layer 3 protocol information	6a* (Note 8)
0/1 ext	0 0 0 spare | Default Packet size	6b* (Note 7)
1 ext	Packet window size	6c* (Note 7)

Figure 4-26/Q.93B
Broadband Low Layer Information

Note 5 - This octet may be present only if octet 5~~7~~ indicates certain acknowledged mode HDLC elements of procedure as indicated in Table 4-15/Q.931

Note 6 - This octet may be present only if octet ~~7~~5 indicates user specified layer 2 protocol.

Note 7 - This octet may be present only if octet 8 6 indicates a layer 3 protocol based on Recommendation X.25, ISO/IEC 8208 or X.223/ISO 8878 as indicated in Table 4-15/Q.931

Note 8 - This octet may be present only if octet ~~8~~6 indicates user specified layer 3 protocol.

Q.93B information elements (Nov. 1992, not complete)

- B-Bear capability (class X, Class C, and Class A) *
- User cell rate (peak cell rate) *
- Called party address
- Called party subaddress
- Calling party address
- Calling party subaddress
- Transit network selection (Carrier selection)
- AAL type (AAL 1-5, assured, non-assured) *
- B-LLI (layer 1-3 protocol information) *
- Connection Identifier (VPI/VCI) *

Note : * are ATM specific information elements

4-31

71

TABLE 4-18/Q.938 (Sheet ~~11 of 12~~ 4 of 5)
~~Low layer compatibility information element~~ B-LLI

- User information layer 3 protocol (octet ~~7~~ 6

ockt no. to be aligned

Bits
5 4 3 2 1

```
0 0 0 1 0  CCITT Recommendation Q.931 (I.451)
0 0 1 0 1  Reserved [Ed.Note: cf. Rec.Q.933]
0 0 1 1 0  CCITT Recommendation X.25, packet layer (Note 2)
0 0 1 1 1  ISO 8208 [41] (X.25 packet level protocol for data
              terminal equipment) (Note 2)
0 1 0 0 0  X.223/ISO 8878 (use of ISO 8208 and CCITT X.25 to provide
              the OSI-CONS) (Note 2)
0 1 0 0 1  ISO 8473 [43] (OSI connectionless mode protocol)
0 1 0 1 0  CCITT Recommendation T.70 [32] minimum network layer
1 0 0 0 0  User specified (Note 1)
```

All other values are reserved.

Note 1 - When this coding is included, octet 7a will include user coding for the user specified layer 3 protocol.
Note 2 - When this coding is included, octets 7a, 7b and 7c with CCITT encoding may be included.

***Editor: text in copy (TD 693/XI, 10/91) unreadable (due to poor reproduction):
____ ISO TR 9577 protocol, identification in the ____ Annex C and D.
***(???)

- Octet 7a for CCITT codings:

Mode of operation (octet 7a)

Bits
7 6

```
0 1  normal packet sequence numbering
1 0  extended packet sequence numbering
```

All other values are reserved.

- Octet 7a for user protocol:

User specified layer 3 protocol information (octet 7a)

The use and coding of octet 7a depends on user defined requirements.

CCITT on-going Q.93B work

- Forward/Backward compatibility
  - ISO application context
  - IE instruction indication octet
- Traffic parameters
- B-LLI codepoint allocation
- Restart procedure
- Status Enquiry procedure
- IE free-ording vs. fix ordering
- And others

ATM Forum signaling work requirements

- simplified point to point connection setup and tear down
- point to multipoint connection setup
- multicast services

  provide multiple sender/receivers service through point to point and point to multipoint connection.
- host self-identification and auto-configuration

  RARP like mechanism, exchange ID and address.
- service access point (SAP) identification

  TCP/UDP port like mechanism to identify the called entity, under discussion.

ATM Forum Q.93B extensions

- point to multipoint connection procedure
    - add new messages and information elements, ready for detailed review
- network recognize multicast ID for multicast services
- addressing structure
    - E.164 in the public network
    - no agreement on what should be in private network

      Proposals discussed in ATM Forum:

      MAC, E.164+MAC, NSAP
- SAP identification
    - no agreement on the mechanism

Route selection support in Q.93B

- called party address IE
- called party subaddress IE
    - E.164 (telephone number)
        - assigned with geographic meaning
    - other addressing
        - need to exchange routing information
- transit network selection IE
    - usually identifies carrier, up to 4 transit nwks
    - provides limited source/policy routing
    - no support for security, route recording
- Virtual circuit routing BOF, Tuesday evening

CCITT	ATM Forum
Q.93B	Q.93B(+-)
SSCOP	Layer 2(?)
AAL(?)	AAL(?)
ATM	ATM

Broadband ISDN UNI signaling stack
(first release)

Schedule and milestones

ATM Forum signaling sub-working group

- release signaling specification on April 1993

ANSI T1S1.2 (B-ISDN)

- approved 12 contributions to CCITT, included new codepoints for B-LLI.

CCITT

- freeze SSCOP on Feb. 1993
- unofficial meeting on Feb. 1993
- release Q.93B specification on Dec. 1993

SSCOP - thumper.bellcore.com

**AALS - thumper.bellcore.com
pub/chw**

Internet/ATM: Signalling Requirements

Tom Lyon

Sun Microsystems Computer Corporation

pugs@eng.sun.com

Goals forATM

- End-to-end ATM usage is desirable, across wide geographies and administrative domains
- Non end-to-end ATM is a fact of life forever - support it well
- Applications, end systems, people don't want to know
- Except they want more performance in the end-to-end case

Addressing

- Addressing and administration are all mixed up
- "Fixing" IP is an orthogonal problem
- Routers do a lot more than route - security, accounting, storm control, etc.
- MAC-level addressing model for ATM means separate routers are still needed
- Integrating ATM with multi-protocol addressing gets end-end performance with good control

Header Avoidance

- For end-end ATM case, large parts of IP and TCP are redundant
- Avoid protocol/header overhead by terminating ATM connections at higher levels of protocol stacks
- Important to keep look, smell, and feel of TCP, but new protocol(s) required
- View connection as a cache of route/protocol decisions

Addressing Requirements

- Type field in every address for multi-protocol addressing (OSI NSAP?)
- Sub-address hack for transiting E.164-only nets
- SAP/port addressing for identifying higher levels in protocol stack
- Address discovery - DHCP?
- Determination of end-to-end-ness

Connectionless

- Connectionless semantics are important, performance needs work
- Certain things will always be better connectionless (service discovery, keep-alives)
- If switch understands network addressing, connectionless forwarding is small extra step
- Need signalling to discover and connect to connectionless service (router)

Multicast

- LAN/IP model is many-to-many;
  ATM model is 1-to-many
- Multicast server(s) provide chokepoint for IP over ATM
- Server/service location & server-to-server issues: dynamic routing/recovery
- Many-to-many works better in connectionless world

Routing

- Dynamic routing is a bigger issue for private networks than for public networks
- Customers unwilling to pay for same level of redundancy; but still want high availability
- Routing, addressing, policy, accounting all mixed up
- Can't "fix it" in ATM - its a higher level problem
- Use existing work for ATM; don't invent a new universe

Other

- Per connection MTU discovery
- Lightweight connections for best-effort QOS

2.2.5 P. Internet Protocol (pip)

Charter

Chair(s):

Paul Tsuchiya, tsuchiya@thumper.bellcore.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

Goals and Milestones:

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

Done Post as an Internet-Draft a description of the Pip Packet Format and Forwarding Engine, the Pip Control Message Protocol (PCMP), the Pip Host Interface Message Protocol, and the Pip MTU Discovery Protocol.

Oct 1992 Post as an Internet-Draft a description of the modifications to BGP IV for Pip, the Modifications to OSPF for Pip, the modifications to DNS for Pip, the modifications to ARP for Pip, the Address assignment in Pip, and the Pip transition strategy.

Done Presentation and review of the PIP specification by the IESG. If acceptable, the first Working Group meeting will be held.

Internet-Drafts:

"Pip Header Processing", 10/30/1992, P. Tsuchiya <draft-ietf-pip-processing-00.txt>

"Pip Objects", 10/30/1992, P. Tsuchyia <draft-ietf-pip-objects-00.txt>

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

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

"Pip Identifiers", 11/03/1992, P. Tsuchiya <draft-ietf-pip-identifiers-00.txt>

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

CURRENT MEETING REPORT

Reported by Paul Tsuchiya/Bellcore

Minutes of P. Internet Protocol Working Group (PIP)

The PIP Working Group was Chaired by Paul Tsuchiya. Except for an announcement by Paul that an implementation of PIP is planned, the whole PIP Working Group meeting was spent covering technical issues. There were many aspects of PIP that were not covered, due to lack of time. It was felt that there should be two sessions devoted to PIP at the next IETF meeting.

Several decisions were made concerning header structure, and these will be reflected in the next version of the PIP header. It was decided that the segmentation function should be moved from the "router options" to the "host options", thus making it a purely end-to-end function. Therefore, routers will not be able to segment, only hosts. If a packet is too large for a particular subnet, then the router will send the host an ICMP-type message.

Some concerns were raised during the meeting. In particular, Joel Halpern pointed out that the Routing Context (RC) and Harding Directive (HD) reformatting required large amounts of memory in the case where both 1) a PIP system had a large number of neighbors, and 2) the RC or HD had a large number of significant bits. Ross Callon raised the concern that PIP might not be able to emulate the NSAP feature of encoding a subnetwork address in the internetwork layer for the purpose of easily binding the internet address to the subnet address. Paul Tsuchiya suggested a couple of ways that this could be accomplished in PIP.

Attendees

Cengiz Alaettinoglu	ca@cs.umd.edu
David Arneson	arneson@ctron.com
Toshiya Asaba	asaba@wide.sfc.keio.ac.jp
Jim Barnes	barnes@xylogics.com
Tony Bates	t.bates@nosc.ja.net
Shiraz Bhanji	bhanji@gateway.mitre.org
Fred Bohle	fab@interlink.com
David Bolen	db31@ans.net
Ross Callon	callon@bigfut.lkg.dec.com
Ken Carlberg	Carlberg@cseic.saic.com
Richard Cogger	rhx@cornell.cit.bitnet
Rob Coltun	rcoltun@ni.umd.edu
Chuck Cranor	chuck@maria.wustl.edu
Bruce Davie	bsd@bellcore.com
Michael Davis	mad@spirit.clearpoint.com
Osmund de Souza	osmund.desouza@att.com
Michael DeAddio	deaddio@thumper.bellcore.com

Tim Dixon	dixon@rare.nl
Jon Dreyer	Jon.Dreyer@east.sun.com
Robert Enger	enger@reston.ans.net
Roger Fajman	raf@cu.nih.gov
Dino Farinacci	dino@cisco.com
William Fink	bill@wizard.gsfc.nasa.gov
Peter Ford	peter@goshawk.lanl.gov
Heather Gray	heather@zk3.dec.com
William Haggerty	haggerty@ctron.com
Joel Halpern	jmh@network.com
Dwight Jamieson	djamies@bnr.ca
Eliot Lear	lear@sgi.com
David LeRoy	dleroy@mitchell.cit.cornell.edu
Tracy Mallory	tracym@3com.com
Dave Monachello	dave@pluto.dss.com
Michael Patton	map@bbn.com
Shawn Routhier	sar@epilogue.com
Assaf Rubissa	asaf@fibhaifa.com
Henry Sanders	henrysa@microsoft.com
Dallas Scott	scott@fluky.mitre.org
Sally Tarquinio	sallyt@gateway.mitre.org
Fumio Teraoka	tera@csl.sony.co.jp
Marten Terpstra	marten@ripe.net
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Luanne Waul	luanne@wwtc.timeplex.com
Douglas Williams	dougw@ralvmg.vnet.ibm.com
Robert Woodburn	woody@sparta.com

2.2.6 Point-to-Point Protocol Extensions (pppext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

None specified

Internet-Drafts:

"The PPP Internetwork Packet Exchange Control Protocol (IPXCP)", 06/10/1992, W. Simpson <draft-ietf-pppext-ipxcp-02.txt>

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

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

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

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

"Compressing IPX Headers Over WAN Media (CIPX)", 12/08/1992, S. Mathur, M. Lewis <draft-ietf-pppext-cipx-00.txt>

"PPP LCP Extensions", 01/08/1993, W. Simpson <draft-ietf-pppext-lcpext-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)"

CURRENT MEETING REPORT

Reported by Brian Lloyd/B.P. Lloyd & Associates

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

Document Status

- New RFCs
  - AppleTalk CP 1378
  - DECNET
  - OSI
  - IPXWAN 1362

- Internet-Drafts
  - LCP Extensions - largely complete
  - Compression - being written, not ready for discussion, compression control protocol.
  - LAPB - being written, not ready for discussion.
    * Numbered mode
    * MLAP procedure
    * Link association
    * Use of RFC 1294 segmentation
  - IPXCP Compromise between Novell desires and other vendors' desires largely complete. Awaiting a few small editorial changes prior to last call.
  - Van Jacobsen's Compression for IPX is out for discussion.
  - Bridging to Draft Standard add "my MAC Address is" option.

Testing Document Subgroup

Two issues: what to test and how to test it. Fred Baker suggested making PT-500 tests available for the purposes; available tests are for synchronous implementations. PPP Consortium could make available to vendors, to make implementations more reliable.

The PPPEXT Working Group and the IPLPDN Working Group met together to discuss using the PPP negotiation mechanisms on PDNs using Frame Relay.

Frame Relay/PPP Parameter Negotiation

- Active PPP Open

  The idea here is that an older Frame Relay system can talk to a newer one by the newer one detecting the other and switching to the old mechanism.

- Agglutination of PPP negotiation parameters (Multiple Option Negotiation)

  Some sort of LCP option saying "I want to do <protocol list > in the default manner". Breaks authentication; NCPs happen after authentication completes. Solution: an LCP configure could be sent as a single datagram carrying a succession of messages.

- Agglutination of PPP negotiation parameters (Multiple Option)

- OUI Acquisition

  Need the necessary OUI for PPP message identification.

- Agglutination of PPP negotiation parameters (Multiple Option)

- OUI/PID values for PPP

  PID should be PPP Protocol Identifier.

- Agglutination of PPP negotiation parameters (Multiple Option)

- Multi-link Transport

  Layered implementation suggested.

Attendees

Fred Baker	fbaker@acc.com
Ken Benstead	kbenstead@coral.com
Daniel Brennan	dmb@teleoscom.com
Caralyn Brown	cbrown@wellfleet.com
Kay Chang	chang@chang.austin.ibm.com
Dean Cheng	dean@sun2.retix.com
Chi Chong	cchong@synoptics.com
Nabil Damouny	nabil@tdd.sj.nec.com
James Forster	forster@cisco.com
Shoji Fukutomi	fuku@furukawa.co.jp
Paul Griffiths	griff@chang.austin.ibm.com
Thomas Hacker	hacker@citi.umich.edu
Frank Heath	heath@cmc.com
Marco Hernandez	marco@mh-slip.educom.edu
Ken Hirata	khirata@emulex.com
Jeff Hughes	jeff@col.hp.com
Bob Jeckell	rrj@3com.com
George Kajos	kajos@coral.com

Mark Lewis	mlewis@telebit.com
David Lin	lind@janus-ccm.zenith.com
Andrew Malis	malis@bbn.com
Julianne Myers	jmyers@network.com
Brad Parker	brad@fcr.com
Drew Perkins	perkins+@cmu.edu
Christopher Ranch	cranch@novell.com
Greg Ratta	gregra@cc.bellcore.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Eva Wang	eva@penril.com
James Watt	james@newbridge.com
Brenda Whitehurst	brenda@penril.com
Ian Wilson	ianw@spider.co.uk
Honda Wu	honda@nat.com

2.2.7 Router Requirements (rreq)

Charter

Chair(s):

Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:

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

Description of Working Group:

The Router Requirements Working Group has the goal of rewriting the existing Router Requirements RFC, RFC-1009, and a) bringing it up to the organizational and requirement explicitness levels of the Host Requirements RFC's, as well as b) including references to more recent work, such as OSPF and BGP.

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

The purposes of this project include:

- Defining what an IP router does in sufficient detail that routers from different vendors are truly interoperable.

- Providing guidance to vendors, implementors, and purchasers of IP routers.

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

Goals and Milestones:

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

Request For Comments:

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

RFC 1354 "IP Forwarding Table MIB"

2.2.8 Simple Internet Protocol (sip)

Charter

Chair(s):

Christian Huitema, christian.huitema@sophia.inria.fr
Steve Deering, deering@parc.xerox.com

Mailing Lists:

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

Description of Working Group:

SIP is another candidate for IPv7. The purpose of the Working Group is to finalize the SIP family of protocol, and to foster the early development and experimentation of this protocol.

There are two major characteristics of the SIP proposal: it is very much a continuation of IP, and it aims at maximum simplicity. A short hand definition of SIP could be "64 bits IP with useless overhead removed".

Following the IP model, SIP uses globally-unique addresses, hierarchically structured for efficient routing. SIP addresses are 64 bits long, which we believe 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 efficently in a very large internet. As a result of this design philosophy, the SIP header is much simpler than the IP header. Simplicity facilitates high-performance implementation and increases the likelihood of correct implementation.

Contrary to several other IPv7 candidates, the SIP effort is focused mostly on the description of the final state, not on the description of the transition. This is due to a coordination with the IPAE working group, which has already engaged an intensive study of transition problems, with SIP in mind as a final state.

Goals and Milestones:

Done Post the complete SIP specification as an Internet-Draft. This specification shall include the header format, the address format, ICMP and IGMP, the fragmentation protocol, the source route protocol, and the the requirements SIP imposes on higher layer protocols and lower later protocols, e.g., ARP.

Jan 1993 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.

Jun 1993 Specify in detail the changes to the routing protocols needed for SIP.

Internet-Drafts:

"Simple Internet Protocol (SIP) Specification", 11/11/1992, S. Deering <draft-deering-sip-00.txt>

CURRENT MEETING REPORT

Reported by Christian Huitema/INRIA

Minutes of the Simple Internet Protocol BOF (SIP)

The Simple Internet Protocol BOF attracted a wide audience. The first part of the meeting was a quick review of the proposed SIP Charter, which was approved by the Group modulo alignment of the milestone dates with the proposed IESG decision schedule. The participants were reminded of the name of the mailing list: <sip-request@caldera.usc.edu> and that preliminary versions of the documents can be obtained by anonymous ftp from "parcftp.xerox.com" in the directories "pub/sip" or "pub/net-research". Related documents on IPAE can be obtained from the same server in the directory "pub/ip-encaps".

The discussion turned next to the SIP specifications, addressing a set of characteristic design points, and in particular some issues that were marked as provisional in the current specification:

- Steve Deering presented a problem posed by the difference between the TCP pseudo header "conceptual layout" and the actual layout of the payload length and type fields in the packets, and asked whether conceptual and physical layout should be aligned. It was observed that the pseudo header remains constant (modulo the packet length) for the duration of the connection, while changing the layout would make the hop count handling in each packet somewhat slower. Moreover, the relation between packet layout and pseudo header will have to remain "conceptual" when options like source routing are used. It was decided not to change the packet layout, but to explain more clearly the pseudo checksum computation rules in the documentation.

- Some Group members questioned the absence of a checksum in the network header. This item had already been debated in the mailing list. The arguments for omitting the checksum will have to be presented in detail in a SIP overview document.

- Some Group members questioned the small size of the payload type field, and the need to provide an extension mechanism, e.g., for student projects. Various solutions were proposed, e.g., to reserve the value "255" for an extension mechanism. The need for a payload type indicating "intermediate options" (to be processed by all routers) was mentioned in the same discussion. An example of a request for such options may be the need of performing "trace route" on a multipoint address. This mechanism will have to be documented in the specification.

- The discussion on "flow-ids" showed that there was no consensus on this point that many members feel it deserves further research, and that the corresponding bits should remain reserved in the initial specification. However, the first implementors reported that the presence of a TOS field similar to that of IPv4 would help the transition process. This field will have to be added in the revised specification.

One of the results of the discussions of the specifications was to outline the need for an "overview" document. The discussion turned then to addressing. Ross Callon objected that the 64 bits SIP addresses were smaller than the 160 bits NSAPs, and therefore could not easily be used to incorporate link layer addressing, e.g., telephone numbers. The discussion showed that the Working Group did not believe that the NSAP size was justified or needed, and that there is virtue in keeping the addresses compact. Steve Deering then presented the "metropolitan" addressing plan. One of the results of the discussion was to outline again the need for more explanations. The overview or the addressing documents should explain how mobility, renumbering and policy routing are supported, based on concrete examples.

Attendees

Cynthia Bagwell	cbagwell@gateway.mitre.org
David Bolen	db31@ans.net
Ross Callon	callon@bigfut.lkg.dec.com
Ken Carlberg	Carlberg@cseic.saic.com
Stephen Casner	casner@isi.edu
Rob Coltun	rcoltun@ni.umd.edu
Michael Conn	4387451@mcimail.com
Chuck Cranor	chuck@maria.wustl.edu
David Crocker	dcrocker@mordor.stanford.edu
Michael Davis	mad@spirit.clearpoint.com
Steve Deering	deering@parc.xerox.com
Barbara Denny	denny@erg.sri.com
Kurt Dobbins	dobbins@ctron.com
Jon Dreyer	Jon.Dreyer@east.sun.com
Ralph Droms	droms@bucknell.edu
Donald Eastlake	dee@ranger.enet.dec.com
Robert Enger	enger@reston.ans.net
William Fink	bill@wizard.gsfc.nasa.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Shoji Fukutomi	fuku@furukawa.co.jp
Robert Gilligan	Bob.Gilligan@eng.sun.com
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Masayoshi Gohara	mg@sinet.ad.jp
Heather Gray	heather@zk3.dec.com
William Haggerty	haggerty@ctron.com
Joel Halpern	jmh@network.com
Robert Hinden	hinden@eng.sun.com
Don Hoffman	don.hoffman@eng.sun.com
Christian Huitema	christian.huitema@sophia.inria.fr
John Ioannidis	ji@cs.columbia.edu
Ronald Jacoby	rj@sgi.com
Charley Kline	cvk@uiuc.edu
Tracy Mallory	tracym@3com.com

Greg Minshall	minshall@wc.novell.com
Dave Monachello	dave@pluto.dss.com
Andy Nicholson	droid@cray.com
Erik Nordmark	nordmark@eng.sun.com
Joseph Ramus	ramus@nersc.gov
Benny Rodrig	4373580@mcimail.com
Henry Sanders	henrysa@microsoft.com
Henning Schulzrinne	hgs@research.att.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Frank Solensky	solensky@andr.ub.com
Tang Tang	tt@virginia.edu
Richard Thomas	rjthomas@bnr.ca
Jim Thompson	jim@tadpole.com
Stuart Vance	vance@tgv.com
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
A. Lee Wade	wade@nsipo.nasa.gov
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Luanne Waul	luanne@wwtc.timeplex.com
Douglas Williams	dougw@ralvmg.vnet.ibm.com
Kirk Williams	kirk@sbctri.sbc.com
Daniel Wilson	dvw@bellcore.com
Robert Woodburn	woody@sparta.com

2.2.9 TCP/UDP over CLNP-addressed Networks (tuba)

Charter

Chair(s):

Mark Knopper, mak@merit.edu
Peter Ford, peter@goshawk.lanl.gov

Mailing Lists:

General Discussion: tuba@lanl.gov
To Subscribe: tuba-request@lanl.gov
Archive:

Description of Working Group:

The TUBA Working Group will work on extending the Internet Protocol suite and architecture by increasing the number of end systems which can be effectively addressed and routed. The TUBA effort will expand the ability to route Internet packets by using addresses which support more hierarchy than the current Internet Protocol (IP) address space. TUBA specifies the continued use of Internet Transport Protocols, in particular TCP and UDP, but encapsulated in ISO 8473 (CLNP) packets. This will allow the continued use of Internet application protocols such as FTP, SMTP, Telnet, etc. An enhancement to the current system is mandatory due to the limitations of the current 32 bit IP addresses. TUBA seeks to upgrade the current system by a transition from the use of the Internet Protocol version 4 to ISO/IEC 8473 (CLNP) and the corresponding large Network Service Access Point address space.

In addition to protocol layering issues and "proof of concept" work, the TUBA approach will place significant emphasis on the engineering and operational requirements of a large, global, multilateral public data network. TUBA will work to maximize interoperatability with the routing and addressing architecture of the global CLNP infrastructure. The TUBA Working Group will work closely with the IETF NOOP and IPRP-for-IP Working Groups to coordinate a viable CLNP based Internet which supports the applications which Internet users depend on such as Telnet, FTP, SMTP, NFS, X, etc. The TUBA Working Group will also work collaboratively with communities which are also using the CLNP protocol, 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 an 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.
Nov 1992	Post as an Internet-Draft a revision to RFC 1347 reflecting lessons learned in the Working Group deliberation.

Internet-Drafts:

"Addressing and End Point Identification, For Use with TUBA", 11/06/1992,
R. Callon <draft-ietf-tuba-address-00.txt, .ps>

CURRENT MEETING REPORT

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

Report not submitted.

Attendees

Roland Acra	acra@cisco.com
Cengiz Alaettinoglu	ca@cs.umd.edu
Philip Almquist	almquist@jessica.stanford.edu
Jules Aronson	aronson@nlm.nih.gov
Nagaraj Arunkumar	nak@3com.com
William Barns	barns@gateway.mitre.org
Bryan Beecher	bryan@umich.edu
Shiraz Bhanji	bhanji@gateway.mitre.org
David Bolen	db31@ans.net
Ross Callon	callon@bigfut.lkg.dec.com
Brian Carpenter	brian@dxcern.cern.ch
George Chang	gkc@ctt.bellcore.com
Henry Clark	henryc@oar.net
Richard Colella	colella@osi.ncsl.nist.gov
Michael Collins	collinsms@es.net
Michael Conn	4387451@mcimail.com
John Dale	jdale@cos.com
Osmund de Souza	osmund.desouza@att.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
David Dubois	dad@pacersoft.com
Jacques Dugast	dugast@issy.cnet.fr
Dino Farinacci	dino@cisco.com
Eric Fleischman	ericf@act.boeing.com
Peter Ford	peter@goshawk.lanl.gov
Shoji Fukutomi	fuku@furukawa.co.jp
Peter Furniss	p.furniss@ulcc.ac.uk
Masayoshi Gohara	mg@sinet.ad.jp
Heather Gray	heather@zk3.dec.com
Thomas Hacker	hacker@citi.umich.edu
William Haggerty	haggerty@ctron.com
Joel Halpern	jmh@network.com
Susan Hares	skh@merit.edu
Tim Howes	tim@umich.edu.
Kathleen Huber	khuber@bbn.com
David Jacobson	dnjake@vnet.ibm.com

Dwight Jamieson	djamies@bnr.ca
Matthew Jonson	jonson@server.af.mil
Merike Kaeo	merike@alw.nih.gov
Akira Kato	kato@wide.sfc.keio.ac.jp
Dave Katz	dkatz@cisco.com
Mark Knopper	mak@merit.edu
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
David Lin	lind@janus-ccm.zenith.com
Tracy Mallory	tracym@3com.com
Bill Manning	bmanning@sesqui.net
Matt Mathis	mathis@a.psc.edu
Jun Matsukata	jm@eng.isas.ac.jp
Donald Merritt	don@brl.mil
David Meyer	meyer@oregon.uoregon.edu
Dave Monachello	dave@pluto.dss.com
Robert Moose	rmoose@gateway.mitre.org
Dennis Morris	morrisd@imo-uvax.disa.mil
Jun Murai	jun@wide.ad.jp
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Andy Nicholson	droid@cray.com
Peder Norgaard	pcn@tbit.dk
Andrew Partan	asp@uunet.uu.net
Laura Pate	pate@gateway.mitre.org
David Piscitello	dave@sabre.bellcore.com
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	abm@tycho.ncsc.mil
Henry Sanders	henrysa@microsoft.com
Dallas Scott	scott@fluky.mitre.org
John Scudder	jgs@merit.edu
Paul Serice	serice@cos.com
Keith Sklower	sklower@cs.berkeley.edu
Frank Solensky	solensky@andr.ub.com
Richard Thomas	rjthomas@bnr.ca
Paul Traina	pst@cisco.com
Iain Wacey	cat@pluto.dss.com
William Warner	warner@ohio.gov
Luanne Waul	luanne@wwtc.timeplex.com
Guy Wells	guy2@uswest.com
Evan Wetstone	evan@rice.edu
Kirk Williams	kirk@sbctri.sbc.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov

2.3 Network Management Area

Director(s):

- James Davin: davin@bellcore.com

Area Summary reported by James Davin/Bellcore

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

IFIP Electronic Mail Management BOF (EMAILMGT)

This BOF session functioned as the second meeting of the IFIP-EMailMgt Task Group. Among the basic points covered were that this is an international task group, not directly chartered within the IETF, that is focused first on pre-standards work, such as requirements determination and abstract modeling for Electronic Mail Management of all kinds of interworking Electronic Mail Systems. The goals and work plan reflect this orientation.

IFIP-EMailMgt plans to call meetings in conjunction with various other Task Force and Workshop meetings around the world to intersect all interested segments of the EMail Industry. In each such meeting, the EMailMgt attendees will adhere to the meeting rules of their host. An IFIP-EMailMgt meeting may be called by any interested group, in conjunction with any meeting venue. Outputs from EMailMgt are directed to whomever may be interested in using them for whatever purposes they may have.

Discussion at this meeting centered on review of an email requirements document and on more general discussion of models for email communication. Plans for the next meeting of this Group in December were also discussed.

SNMP Applications Monitoring BOF (SAM)

This birds-of-a-feather session met to explore interest and possible direction for SNMP instrumentation of email and directory services applications. Many expressed interest in such instrumentation, although few had read the draft specification that had been circulated.

The Group tentatively agreed on a three-pronged strategy that could be the basis for forward progress. A "server MIB" could be drafted by Steve Hardcastle-Kille to instrument generic aspects of connection-oriented applications. A "mail server MIB" could be drafted by Ned Freed to instrument aspects of X.400 email applications. Finally, a "directory server MIB" could be drafted by Steve Hardcastle-Kille to instrument aspects of X.500 directory services.

The Group may seek to become chartered as an IETF working group in the Applications Area. In that case, the relationship between such a working group and the existing Host MIB Working Group would require study and clarification.

Chassis MIB Working Group (CHASSIS)

The Chassis MIB Working Group met briefly at the Washington meeting to continue its deliberations. Of the three objectives in its Charter, the Working Group has focused much of its effort on the first: defining MIB objects that represent the mapping of the logical functions of traditional network devices onto particular, physical hardware resources within the chassis.

The second objective (instrumenting power supplies) has been largely subsumed by the work of the UPS MIB Working Group. The third objective (MIB definitions to support the distillation of aggregates of mangement information) is still viewed as extremely important but less urgent than the first objective.

FDDI MIB Working Group (FDDIMIB)

The FDDI MIB Working Group met to discuss alignment of the existing FDDI MIB (RFC 1285) with version 7.2 of the SMT work recently produced by ANSI. A brief report was received on the standardization status of SMT 7.2 within ANSI. Finalization is expected soon.

The current draft was introduced, including the algorithms used for object inclusion and exclusion and the algorithms used for transforming objects between GDMO and the SNMP/SMI concise MIB format.

The Group made decisions on a number of issues. It selected a strategy for incorporation of PATHConfiguration information into the MIB. It agreed on the units of measure associated with objects of type FddiTime and Timestamp. It adopted a strategy for mapping SMT to SNMP indices.

The editor was directed to prepare a new draft to reflect these decisions and to have it posted to the internet-drafts directories. Discussion of event notifications was deferred until after the MIB design work is completed.

Host Resources MIB Working Group (HOSTMIB)

The Working Group reached consensus that all appropriate substantive issues had been addressed and that the MIB should be submitted to the IESG for publication as an RFC at the Proposed Standard level. In particular, this consensus was that the appropriate amount of effort had been placed in identifying issues, and that each of these issues had been resolved. Any further issues that might be identified after the meeting would be addressed when the MIB is to go to Draft Standard. Bug fixes and clarifications during the final editing process would be agreed upon on the mailing list.

The plan for the document is that Grillo and Waldbusser will make changes based on the results of the meeting and make that document available via FTP and a notification on the mailing list. A week or two of discussion would occur on issues of bug fixes and clarifications.

This would result in editing instructions that would be applied to the MIB and the MIB would then be submitted to the internet-drafts directory and to the IESG.

IEEE 802.3 Hub MIB Working Group (HUBMIB)

This Working Group met in Washington to continue its discussion of the MAU MIB for instrumenting IEEE 802.3 Medium Access Units. The Working Group reached closure on all outstanding issues, and revised text for the MAU MIB will be posted to the mailing list for final approval by the Group. The Group also spent some time discussing implementation experience with the previously standardized Repeater MIB in anticipation of its possible consideration as a Draft Standard.

Internet Accounting Working Group (ACCT)

The Internet Accounting Working Group met Friday morning, November 20th to hear the interim results of work done by Nevil Brownlee of the University of Auckland and review proposed changes to the draft Accounting MIB to improve performance characteristics. Since progress towards a Standard now depends on the implementation of a working MIB, the Working Group has agreed to become dormant pending initial implementation. The Internet Accounting Architecture will be revised so as not to prohibit such changes and will be submitted for publication as an Informational RFC. The Accounting MIB will be published as soon as possible as an Experimental RFC.

SNMP Version 2 Working Group (SNMPV2)

This Working Group met to continue its discussion of SNMP evolution. The Working Group was able to resolve a number of issues that had been raised previously in its discussions: rules on changing OBJECT DESCRIPTOR values, rules on changing distinguished values in enumerations, access to auxiliary objects, the readOnly error status, and the conformance status of the RowStatus textual convention.

In addition, the Working Group addressed a number of new issues that had been raised since its last meeting. These included delayed operations protocol (e.g., due to slow proxy), efficient retrieval of tabular objects, auto-discovery of SNMP agents, non-rectangular tables, numerous SMI issues, unsigned 64-bit SMI types, ASN.1 macros, translation of work in other standards bodies to enumerated INTEGER types, and row-creation.

There was consensus that the final deadline for comments on the SNMPv2 documents should be Friday, December 11. Unless the SNMP Security effort raised new issues, then the documents would be sent forward to the IESG with a recommendation for advancement to the standards-track, pending completion of the work of the SNMP Security Working Group.

Token Ring Remote Monitoring MIB Working Group (TRMON)

This Working Group met for two sessions on Thursday. The major activity was a group by group discussion of the working draft for the Token Ring RMON. The end result was numerous changes for the author. The goal is to incorporate those changes by December 15th. To be followed by a call for consensus as to requesting that Token Ring RMON become a Proposed Standard with RFC status.

There were a couple of cantankerous issues, in particular the order table, for which consensus was reached. It was also stressed that further discussion should take place on the mailing list.

CURRENT MEETING REPORT

Reported by Einar Stefferud/NMA

Minutes of the IFIP Electronic Mail Management BOF (EMAILMGT)

Introduction

This was the Second Meeting of the IFIP-EMailMgt Task Group.

The basis for formation of the IFIP-EMailMgt Joint Task Group of IFIP WG-6.5 and WG-6.6 was reviewed for the benefit of those who were new to the Group. A large portion were new, since this meeting drew many IETF attendees, and the prior meeting was held at OIW in September.

Drawing new participants was one of the primary reasons for organizing an IFIP-EMailMgt BOF meeting at the IETF, and in this we succeeded!

The basic points covered were that this is an international task Group that is focused first on pre-standards work, such as requirements determination and abstract modeling for Electronic Mail Management of all kinds of interworking Electronic Mail Systems. The Charter and the work plan reflect this orientation.

IFIP-EMailMgt plans to call meetings in conjunction with various other Task Force and Workshop meetings around the world to intersect all interested segments of the EMail Industry. In each such meeting, the EMailMgt attendees will adhere to the meeting rules of their host, including payment of attendance fees and provision of meeting reports.

An IFIP-EMailMgt meeting may be called by any interested group, in conjunction with any meeting venue. It is expected that a report of the meeting will be prepared and submitted to the host organization, and to the IFIP-EMailMgt mailing list. The report may include comments on IFIP-EMailMgt work in progress, or may include contributions of any kind, including proposed documents to be progressed for publication as IFIP-EMailMgt output products.

Outputs from EMailMgt are directed to whomever may be interested in using them for whatever purposes they may have.

IFIP-EMailMgt accepts support and participation from any source. Its meetings are entirely open to anyone with an interest in the issues, as is its Mailing list <ifip-emailmgt@ics.uci.edu>. To subscribe, send a request to: <ifip-emailmgt-REQUEST@ics.uci.edu>.

According to the Charter, no decisions may be made in any face-to-face meeting. All decisions will be made openly in the mailing list, using consensus measurement techniques. No specific consensus measurement tools have been selected for this purpose, but we expect to manage it in ad hoc ways as we proceed. Therefore, any meeting can do no more than prepare contributions and proposals to the EMailMgt mailing list.

Meeting Activities

The Charter of the Group was reviewed and a few minor changes were proposed. These will be noted in the Charter for review by the mailing list, along with the process of adoption by the mailing list, which has not yet occurred.

The goals and work plan were reviewed. Given that the work is so near to its beginning, the basic goals (develop requirements, models, and object definitions) were retained and work commenced. The objectives of this meeting were to progress our understanding of the EMailMgt requirements, and to begin work on development of appropriate models. These two areas (Requirements and Models) provide the primary focus for the first phase of our work.

We are not interested in creation of a whole new set of requirements or models. Thus, the work consists of collecting and synthesizing from other work that has been done, or is currently under way.

Some new participants offered to present new work that they are doing, which appears to offer additional prospectives on the overall picture.

What we seem to be finding so far is that existing EMailMgt requirements and modeling work fails to include some aspects that other work does include. Thus our first finding is that we do indeed have some serious work to do.

Email Management should be thought of as a special case of management.

Presentation of Drafts Requirement Document

Emily McCoy is the Design Team Leader for a Requirements Document. She presented the first draft version. (EMGT-92-xxxx), which consisted of a meld of may other documents, which were identified as sources for each paragraph to facilitate tracing concepts back to their roots. Conflicting views were placed in the document to raise discussion points. It contained the following sections:

- General Requirements.
- Requirements from a system administrator view.
- User requirements.

It was agreed that attendees would read the document and discuss it the next day (Thursday, November 19th).

Discussion of Draft Requirements Document

Missing Areas:

Inter-Relay Aspects: There are better versions of Routing management issues (ISO documents authored by Bob Willmott) from Oslo meeting. Should reference MHS-DS documents.

Message Stores User Management Privacy Issues:

- Non-OSI is not clearly stated.
- Need to make sure that it follows all the flavors and needs of Email, but can't be too generic, since that may make it non-usable.
- An editing group was set up to work with Emily on this document.

Another Review of Requirements Document

Emily explained changes; then asked which items were specific and which were generic. It was decided that all items currently in the document were generic! (surprise). There was A LOT (read lots) of discussion when the Group got to user requirements, especially in the area of what users and customers were. There seemed to be a need to differentiate between at least two and potentially three different needs at the user's level.

Remote Mail Management Issues

Context: A set of Email Repository Machines.

Using IMAP, a remote mail access protocol, et al. Management problems are inherently from the user's perspective. (See EMGT-92-25 (Remote email/message MGT). Management includes Boards and News.

Review of Charter

The Group was still having trouble deciding what it is that is management and what this Group is responsible for. How long will mushiness be there until there is firm ground. Can there be a goal (date wise?). It was decided that what the Group is responsible for should be dictated by the Charter.

It was agreed that it should indicate target dates for adoption. Needs a Glossary of items. The editor will rework it where needed, for submission to the mailing list for adoption.

Discussion of Model

Messaging Model: There are two ends (users in most cases).

Some MTA's are in an environment that is larger than themselves. Some are members of other environments. We will call these environments domains. There are protocol enclaves which provide services.

A manager is often confused but is the manager of something. Manager defines management domain. Management domain is what a manager manages. This allows for a manager to manage multiple protocols and platforms in that domain.

Mail Cache, Mail Drop, Queues and Mail Store environments are important. POP server, P7

User ----> MD Manager --------> Service (transport) Manager

This was a good model for the mail system.

Now what is the model for the manager (how to manage it).

1. Manager wants end-to-end connectivity.

2. Message flow. How do you instrument the flows. Requirement is to understand (structure) of the flow.

3. Physical resources.

Some diagrams and pictures were drawn, but they are too hard to render in this report so interested readers are encouraged to join the mailing list and obtain copies of better developed documents.

Harald agreed to develop a more complete version of the model that took shape in the meeting, and publish it on the mailing list before the end of the year. An initial draft will be available for use at the OIW meeting, December 15-17, 1992.

Establishment of Model Design Team

Harald will serve as Modeling Design Team leader.

The Modeling Design Team mailing list is <ifip-tf-model@uninett.no>. To subscribe, send a request to: <ifip-tf-model-request@uninett.no>.

Brief Softswitch Model Presentation

Sheryl Namoglu (Softswitch) offered her understanding of EMGT-92-010 diagram on page 5. Each line below is a different service (or layer).

- User Services (BBS, news, Order processing, etc.)
- User Management (profile of users, etc.)
- Mail Services (Routing, Doc Conv, naming translation, Security, etc.)
- Operating Systems or Transport Service

Presentation of CMU Model

Chris Newman presented the CMU Model. The Model proposed is from the manager's view.

```
client -------------- direct delivery --------------- client

client ------------   message store (semi-direct) ----  client

client -------- mail subscription ---- local gateway ---- client
                                 or global gateway

Many users --- Bulletin boards (or other such) ---  mail service
```

Two radical viewpoints:

1. Elements of management for bb, mail, etc.
2. Services of mail.

Ad Hoc Editing Group

The Requirements Document was reviewed by (John Hawthorne (Rome Research Corp.), Emily McCoy (Mitre), Chris Newman (CMU), Ray Freiwirth (RCI)).

The biggest thing that has to be done to the Requirements Document is folding in other documents.

1. Mailbased servers document.
2. Julian's stuff.
3. Extracted requirements from Ann McLaughlin's MO definitions.
4. Security management document.

Review of Meeting Progress and Future Work Plans

Set Agenda For OIW Meeting in December.

All work beyond completion of Requirements and Modeling Documents is generally on hold because it is dependent on these results. The only exception is that a lot of work is already under way on Definition of CMIS Managed Objects and SNMP MIBs. Thus, current work is organized around these three foci:

1. Requirements
2. Models

3. Object Definitions

Emily's original document is on the net. A Copy of the new document will be placed on the net for review by the time of the OIW meeting.

Harald's first and second Model Document drafts will be placed on the net in time for the OIW meeting.

Paul Brusil will be asked to lead an effort to collect Managed Object and MIB definitions, and then study them to see what can be learned from them. With any luck, they will form a useful base for EMailMgt.

Team Leaders:

Emily McCoy	Requirements
Harald Alvestrand	Models
Paul Brusil	MIB & MO Definition Collection
Future Study	Identify Management Functions

OIW Planned Agenda:

- Administrivia
- Presentations
  - Requirements
  - Model
- Discussions
  - Mapping Requirements/MO
- MO & MIB Collections
- Management Tools and Management Information

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
George Chang	gkc@ctt.bellcore.com
Daniel Fauvarque	dfauvarq@france.sun.com
Ned Freed	ned@innosoft.com
Raphael Freiwirth	5242391@mcimail.com
Terry Gray	gray@cac.washington.edu
Michel Guittet	guittet1@applelink.apple.com
Alf Hansen	Alf.Hansen@delab.sintef.no
John Hawthorne	johnh@tigger.rl.af.mil

Barbara Jennings	bjjenni@sandia.gov
Marko Kaittola	marko.kaittola@funet.fi
Neil Katin	neil.katin@eng.sun.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Edward Levinson	levinson@pica.army.mil
Bob Lynch	lynch@dsteg.dec.com
Emily McCoy	mccoy@gateway.mitre.org
John Myers	jgm+@cmu.edu
Sheryl Namoglu	sfn@softsw.ssw.com
Chris Newman	chrisn+@cmu.edu
Kary Robertson	kr@concord.com
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Jon Saperia	saperia@tcpjon.ogo.dec.com
Michael Sapich	sapich@conware.de
Richard Schmalgemeier	rgs@merit.edu
Chris Shaw	cshaw@banyan.com
John Sherburne	john.sherburne@sprintintl.sprint.com
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de

CURRENT MEETING REPORT

Reported by Harald Alvestrand/SINTEF

Minutes of the SNMP Application Monitoring BOF (SAM)

The Group decided that the current draft was basically a good idea. Very few people had read the draft. But much interest was expressed.

The Group wants to be a working group, and adopted the name "Mail And Directory MANagement" - or MADMAN for short. It should be in the Applications Area, not the OSI Area. Marshall Rose recommended against the Network Management Area, because it is trying to focus on the NETWORK management.

Steve Hardcastle-Kille will write the Charter.

Ned Freed will set up mailing list and archives.

All agreed that the right format was three documents:

1. An "applications" MIB (probably to be renamed a "Server MIB", since it was rather connection type oriented; a spreadsheet would not fit) editor Ned Freed

2. A "Mailserver MIB" Editor Ned Freed.

3. A "Directory server MIB" Editor Steve Hardcastle-Kille
   - There is a limit to be drawn between the (already large) host MIB and this MIB.
   - The exact MIB structure needs polishing.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
David Arneson	arneson@ctron.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Theodore Brunner	tob@thumper.bellcore.com
Stephen Bush	sfb@ncoast.org
Cathy Cunningham	cmc@microcom.com
Cynthia Della Torre	cindy@gateway.mitre.org
Hans Eriksson	hans@sics.se
Daniel Fauvarque	dfauvarq@france.sun.com
Ned Freed	ned@innosoft.com
Ella Gardner	epg@gateway.mitre.org
Michel Guittet	guittet1@applelink.apple.com

Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Gerd Holzhauer	holzhauer1@applelink.apple.com
Marko Kaittola	marko.kaittola@funet.fi
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Edward Levinson	levinson@pica.army.mil
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Rakesh Patel	patel@noc.rutgers.edu
Michael Patton	map@bbn.com
David Perkins	dperkins@synoptics.com
Karen Petraska-Veum	karen@nsisrv.gsfc.nasa.gov
David Piscitello	dave@sabre.bellcore.com
Kary Robertson	kr@concord.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jon Saperia	saperia@lkg.dec.com
Michael Sapich	sapich@conware.de
Dallas Scott	scott@fluky.mitre.org
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@rahul.net
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Cathy Wittbrodt	cjw@nersc.gov
Honda Wu	honda@nat.com
Peter Yee	yee@atlas.arc.nasa.gov

2.3.1 Bridge MIB (bridge)

Charter

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

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

Description of Working Group:

The Bridge MIB Working Group is chartered to define a set of managed objects that instrument devices that conform to the IEEE 802.1 standard for MAC-layer bridges.

This set of objects should be largely compliant with (and even draw from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

The MIB object definitions produced will be for use by SNMP and will be consistent with other SNMP objects, standards, and conventions.

Goals and Milestones:

Done	Publish initial proposal.
Done	Submit an Internet-Draft.
Done	Submit draft for RFC publication.
Mar 1993	Publish a draft revision to RFC 1286 that reflects implementation experience and the result of alignments with IEEE work as an Internet-Draft.
Mar 1993	Publish a draft SNMP MIB that instruments functions specific to source routed bridges as an Internet-Draft.
Apr 1993	Submit a draft MIB for source routing bridge functions to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Definitions of Managed Objects for Bridges", 10/15/1992, E. Decker, P. Langille, A. Rijsinghani <draft-ietf-bridge-objects-01.txt>

Request For Comments:

RFC 1286 "Definitions of Managed Objects for Bridges"

2.3.2 Chassis MIB (chassis)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

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

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

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

Goals and Milestones:

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

Internet-Drafts:

"Definitions of Managed Objects for a Chassis Containing Multiple Logical Network Devices", 01/13/1993, K McCloghrie, D Arneson, M Kaycee <draft-ietf-chassis-mib-00.txt>

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Chassis MIB Working Group (CHASSIS)

Agenda

- Welcome/Introductions/Administrative details
- Chassis Entity Table Issues
  - chassisEntityArgument
  - chassisEntityAdminStatus
  - expository text
- Definition of a Chassis
- Multiple (redundant) Views of the Chassis
- Identify New Issues
- Power Supply
- Environment
- Future Plans

The Chassis MIB Working Group held its third meeting at the IETF meeting in Washington D.C., on Tuesday, November 17, 1992. For this meeting, Jeff Case presided and Bob Stewart recorded. Unlike last time, the Group had done considerable work using the mailing list, but progress broke down due to problems with the overall model for a chassis and how to describe it in a MIB.

Administrative Details

We reviewed the Agenda and made the following changes:

- We added "host as a chassis" to the end of the Agenda, as hosts have slots and Host MIB isn't addressing that. Also added "hosts in a chassis."
- We added implementation model and architectural issues to entity table issues.
- We moved "definition of a chassis" higher.
- We added conformance issues.

The following points were made regarding a suggestion to limit the model to a physical chassis:

- That is a Charter violation.
- We can use a physical chassis as the primary example.
- There are products that are closely-coupled boxes.
- Seven or so companies have products with a proprietary chassis MIB.

- We need to move forward on the specific problem with consideration for the general problem.

Several points were made regarding the model. *Editor's Note (md): A detailed listing of these points is available via ftp under chassis-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Future Plans

- Jeff will prepare an Internet-Draft.
- The Group will discuss the issues on the mailing list.
- The Co-Chairs will be more active on the mailing list.

Attendees

David Arneson	arneson@ctron.com
Jim Barnes	barnes@xylogics.com
Andy Bierman	abierman@synoptics.com
Bill Bowman	bill@hpprsd.mayfield.hp.com
Theodore Brunner	tob@thumper.bellcore.com
Jeff Case	case@cs.utk.edu
Chris Chiotasso	chris@andr.ub.com
Juan Cruz	juan@dss.com
Michael Davison	davison@fibercom.com
Manuel Diaz	diaz@davidsys.com
David Engel	david@ods.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Maria Greene	mngreene@eng.xyplex.com
Pete Grillo	pl0143@mail.psi.net
Ed Heiner	eah@pau.synnet.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
Kevin Jackson	kmj@concord.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Dwain Kinghorn	microsoft!dwaink@uunet.uu.net
Deirdre Kostick	dck2@sabre.bellcore.com
David Lin	lind@janus-ccm.zenith.com
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
William McKenzie	mckenzie@ralvma.vnet.ibm.com
Donna McMaster	mcmaster@synoptics.com
David Minnich	dwm@fibercom.com

Patrick Mullaney	mullaney@ctron.com
Daniel Myers	dan@nsd.3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
David Perkins	dperkins@synoptics.com
Richard Ramos	ramos@mtunm.att.com
Ed Reeder	ereeder@vnet.ibm.com
Sam Roberts	sroberts@farallon.com
Guenter Roeck	roeck@conware.de
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Chris Rozman	chrisr@usr.com
Assaf Rubissa	asaf@fibhaifa.com
Michael Sapich	sapich@conware.de
Michael Scanlon	scanlon@interlan.com
John Seligson	johns@ultra.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@rahul.net
Bob Stewart	rlstewart@eng.xyplex.com
Geoffrey Thompson	thompson@synoptics.com
Peter Wilson	peter_wilson@3com.com
Kiho Yum	kxy@nsd.3com.com

2.3.3 DS1/DS3 MIB (trunkmib)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

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

Internet-Drafts:

"Definitions of Managed Objects for the DS1 and E1 Interface Types", 11/24/1992,
F. Baker, J. Watt <draft-ietf-trunkmib-ds1e1mib-02.txt>

"Definitions of Managed Objects for the DS3/E3 Interface Type", 11/24/1992,
T. Cox, K. Tesink <draft-ietf-trunkmib-ds3e3mib-02.txt>

2.3.4 Ethernet MIB (ethermib)

Charter

Chair(s):

Frank Kastenholz, `kasten@ftp.com`

Mailing Lists:

General Discussion: `enet_mib@ftp.com`
To Subscribe: `enet_mib-request@ftp.com`
Archive: `not available`

Description of Working Group:

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

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

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

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

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

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

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

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

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

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

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

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

Goals and Milestones:

TBD Draft Variable Status Rationale document.

TBD Develop Implementation Report.

Internet-Drafts:

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

Request For Comments:

RFC 1369 "Implementation Notes and Experience for The Internet Ethernet MIB"

2.3.5 FDDI MIB (fddimib)

Charter

Chair(s):

Jeffrey Case, `case@cs.utk.edu`

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done "Final" initial draft of required get/set variables.

Done Initial implementations of required get/set variables.

Done Revised "final" draft of required get/set variables.

Done Adoption of draft of required get/set variables.

Mar 1992 Submit the FDDI MIB to the IESG for consideration as a Proposed or Draft Standard depending on the magnitude of changes to RFC 1285.

Done Hold a meeting at the November IETF Plenary.

Dec 1992 Post an Internet-Draft aligned with current the current ANSI document factoring in implementation experience with RFC 1285.

Request For Comments:

RFC 1285 "FDDI Management Information Base"

CURRENT MEETING REPORT

Reported by Jeff Case/SNMP Research

Minutes of the FDDI MIB Working Group (FDDIMIB)

The FDDI MIB Working Group met on November 19, 1992 in Washington, D.C. This being the second meeting of the Group after a lengthy hiatus, the Chair invested the time to introduce the goals and for participants to introduce themselves to one another. The Agenda for the meeting was reviewed and adopted. The roster was circulated.

A brief report was received on the standardization status of SMT 7.2 within ANSI. Finalization is expected soon.

The current draft was introduced, including the algorithms used for object inclusion and exclusion and the algorithms used for transforming objects between GDMO and the SNMP/SMI concise MIB format.

The identified issues were discussed in turn:

- PATHConfiguration

  The Group agreed that the inclusion of the PATHConfiguration information objects is desirable, given the relative costs and benefits. The Group discussed two proposals that have been advanced, one from Anil and one from Sal. The editor was directed to incorporate the PATHConfiguration information in the next draft based on Sal's suggested table format.

- FddiTime

  The Group agreed to express FddiTime as a positive integer in 1 nanosecond units. The editor was directed to include a complete, appropriate, example to eliminate the confusion surrounding two's complement.

- Timestamp

  The Group agreed to rescale Timestamp to express it in 1 millisecond units.

- Correction of Equations and Relationships

  This is no longer a problem thanks to the assistance of Floyd Ross, who was able to supply machine readable text that could be successfully imported by the document's editor.

- SMT <-> SNMP Index Mappings

  The Group was pleased with the current draft which includes sparse mappings of the index values to match SMT.

The editor was directed to prepare a new draft to reflect these decisions and to have it posted to the internet-drafts directories.

Discussion of event notifications was deferred until after the MIB design work is completed.

Attendees

Caralyn Brown	cbrown@wellfleet.com
Jeff Case	case@cs.utk.edu
Wo Chang	wchang@nist.gov
Chris Chiotasso	chris@andr.ub.com
Wayne Foco	foco@ralvm6.vnet.ibm.com
Ed Heiner	eah@pau.synnet.com
Merike Kaeo	merike@alw.nih.gov
Mark Lewis	mlewis@telebit.com
Evan McGinnis	bem@3com.com
Daniel Myers	dan@nsd.3com.com
Luc Pariseau	pariseau@quiver.enet.dec.com
James Reeves	jreeves@synoptics.com
Guenter Roeck	roeck@conware.de
Dan Romascanu	dan@lannet.com
Assaf Rubissa	asaf@fibhaifa.com
Ahmet Tuncay	atuncay@synoptics.com
Jeff Young	young@alw.nih.gov

2.3.6 Host Resources MIB (hostmib)

Charter

Chair(s):

Steven Waldbusser, `waldbusser@andrew.cmu.edu`

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

All definitions produced by the Group will be consistent with the SNMP network management framework and all other internet-standard MIBs for SNMP. Wherever possible, the definitions produced will make use of or align with relevant work in progress with chartered working groups of the IETF. Also, wherever possible, the Working Group will take into consideration pre-existing, stable work produced by other, accredited standards bodies.

Goals and Milestones:

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

Internet-Drafts:

"Host Resources MIB", 10/07/1992, Pete Grillo, Steven Waldbusser <draft-ietf-hostmib-resources-00.txt>

INTERIM MEETING REPORT

Reported by Steven Waldbusser/CMU

Minutes of the Host Resources MIB Working Group (HOSTMIB)

Minutes of the Host Resources MIB Working Group held October 12-13, 1992 at Carnegie Mellon University.

Agenda

- Introduction (Waldbusser)
  - Local Arrangements
  - Attendance
  - Agenda
- Overview of Changes (Grillo, Waldbusser)
- Questions about current draft (Discussion)
- Proposals
- Summary

Overview of Changes from Last Draft

- New DateAndTime syntax (more changes in the works)
- Processor is now a device
- Storage allocation units
- New device types
- Nuke device location (as per mailing list)
- Add device status
- Add device errors
- Nuke video table (as per mailing list)
- Rearrange disk/partition/file system
- Full/Partial backup date
- Printer table
- SW type (OS, device driver,application)

Editor's Note (md): A detailed account of changes as well as resolutions to proposed topics is available via ftp under hostmib-minutes-92oct.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Attendees

Amatzia Ben-Artzi	amatzia@netmanage.com
Stephen Bush	sfb@ncoast.org
Raymond Edgerton	bellmf.edgerton@uwm.edu

Michael Erlinger	mike@jarthur.claremont.edu
Pete Grillo	pl0143@mail.psi.net
Todd Kulick	TK24@andrew.cmu.edu
Ed Reeder	ereeder@vnet.ibm.com
Jon Saperia	saperia@lkg.dec.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Walter Wong	wcw@cmu.edu

CURRENT MEETING REPORT

Reported by Steve Waldbusser/CMU

Minutes of the Host Resources MIB Working Group (HOSTMIB)

Agenda

- Administration
- Identify New Issues
- Resolve Issues
  - Queues
  - Networked printers
  - In hrPrinterPrintPages, what is a page?
- Plan/Conclusion

New Issues

- swRunCPU
- Parent ID
- User ID
- Terminal ID
- International character set
- Remote file systems (size, use, etc.)
- Chassis MIB (as it relates to HostMIB)
- Processor speed
- Remote FS

Editor's Note (md): A detailed account of the discussions and resolution of issues is available via ftp under hostmib-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Where Do We Go From Here?

The Group reached consensus that all appropriate substantive issues had been addressed and that the MIB should be submitted to the IESG for publication as an RFC at the Proposed Standard level. In particular, this consensus was that the appropriate amount of effort had been placed in identifying issues, and that each of these issues had been resolved. Any further issues that might be identified after the meeting would be addressed when the MIB is to go to Draft Standard. Bug fixes and clarifications during the final editing process would be agreed upon on the mailing list.

The plan for the document is that Pete Grillo and Steve Waldbusser will make changes based on the results of the meeting and make that document available via FTP and a notification on the mailing list. A week or two of discussion would occur on issues of bug fixes and clarifications. This would result in editing instructions that would be applied to the MIB and the MIB would then be submitted to the internet-drafts directory and to the IESG.

Attendees

Bill Bowman	bill@hpprsd.mayfield.hp.com
Stephen Bush	sfb@ncoast.org
Chris Chiotasso	chris@andr.ub.com
James Davin	davin@bellcore.com
Michael Davison	davison@fibercom.com
Manuel Diaz	diaz@davidsys.com
Jon Dreyer	Jon.Dreyer@east.sun.com
Raymond Edgerton	bellmf.edgerton@uwm.edu
Daniel Fauvarque	dfauvarq@france.sun.com
Pete Grillo	pl0143@mail.psi.net
Gerd Holzhauer	holzhauer1@applelink.apple.com
Kevin Jackson	kmj@concord.com
Ronald Jacoby	rj@sgi.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Dwain Kinghorn	microsoft!dwaink@uunet.uu.net
Andrew Knutsen	andrewk@sco.com
David Lin	lind@janus-ccm.zenith.com
William McKenzie	mckenzie@ralvma.vnet.ibm.com
Patrick Mullaney	mullaney@ctron.com
David Orelowitz	david@c1st.microcom.com
Gaige Paulsen	gaige@intercon.com
Ed Reeder	ereeder@vnet.ibm.com
Bradley Rhoades	bdrhoades@mail.mmmg.com
Dan Romascanu	dan@lannet.com
Assaf Rubissa	asaf@fibhaifa.com
Jon Saperia	saperia@lkg.dec.com
Michael Sapich	sapich@conware.de
Timon Sloane	timon@rahul.net
Dean Throop	throop@dg-rtp.dg.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Steven Waldbusser	waldbusser@andrew.cmu.edu

2.3.7 IEEE 802.3 Hub MIB (hubmib)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

The Working Group is further chartered at its discretion to define an SNMP MIB for management of IEEE 802.3 Medium Access Units (MAUs). An 802.3 Medium Attachment Unit (MAU) attaches a repeater port or Ethernet-like interface to the local network medium. The scope of this work may include several types of MAU units: 10BASE5 (thick coax), 10BASE2 (thin coax), 10BASE-T (twisted pair), FOIRL and 10BASE-F (fiber optic). Managed objects defined as part of the MAU MIB task may, for example, represent such information as MAU type, link status, and jabbering indications.

Goals and Milestones:

Done	Distribute first draft of documents and discuss via E-mail.
Done	Working Group meeting as part of IETF to review documents.
Done	Distribute updated documents for more E-mail discussion.
Done	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Done	Documents available with specified changes incorporated.
Done	Submit the Repeater MIB to the IESG for consideration as a Proposed Standard.
Nov 1992	Post the Media Access Unit MIB Definition as an Internet-Draft.
Apr 1993	Submit the Media Access Unit MIB to the IESG for consideration as a Proposed Standard.

Request For Comments:

RFC 1368 "Definitions of Managed Objects for IEEE 802.3 Repeater Devices"

CURRENT MEETING REPORT

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

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

Attendees

David Arneson	arneson@ctron.com
Andy Bierman	abierman@synoptics.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Chris Chiotasso	chris@andr.ub.com
Juan Cruz	juan@dss.com
Dave Cullerot	cullerot@ctron.com
Manuel Diaz	diaz@davidsys.com
David Engel	david@ods.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Maria Greene	mngreene@eng.xyplex.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
Robin Iddon	robini@cix.compulink.co.uk
Kevin Jackson	kmj@concord.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
David Lin	lind@janus-ccm.zenith.com
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
John Medicke	medicke@ralvm11.vnet.ibm.com
David Minnich	dwm@fibercom.com
Daniel Myers	dan@nsd.3com.com
Thomas Pusateri	pusateri@cs.duke.edu
Richard Ramos	ramos@mtunm.att.com
Sam Roberts	sroberts@farallon.com
Dan Romascanu	dan@lannet.com
Chris Rozman	chrisr@usr.com
Assaf Rubissa	asaf@fibhaifa.com
Michael Sapich	sapich@conware.de
Michael Scanlon	scanlon@interlan.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@rahul.net
Gerry White	gerry@lancity.com

Peter Wilson	peter_wilson@3com.com
Randall Worzella	worzella@ralvm29.unet.ibm.com
Kiho Yum	kxy@nsd.3com.com

2.3.8 Internet Accounting (acct)

Charter

Chair(s):

Cyndi Mills, cmills@nnsc.nsf.net
Gregory Ruth, gruth@bbn.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

Done	Policy models examined.
Done	Internet Accounting Background Working Draft written.
Done	Collection Protocols Working Papers written.
Done	Internet Accounting Background final draft submitted to the IESG for consideration as an Informational RFC.
Done	Collection protocol recommendation.
Done	Architecture submission as Internet-Draft.
Done	Post the Accounting Meter MIB as an Internet-Draft.
Jan 1993	Architecture document submitted to the IESG for consideration as a Proposed Standard.
Jan 1993	Submit the Accounting Meter MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

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

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

Request For Comments:

RFC 1272 "Internet Accounting: Background"

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

Minutes of the Internet Accounting Working Group (ACCT)

The Internet Accounting Working Group met Friday morning, November 20th to hear the interim results of work done by Nevil Brownlee of the University of Auckland and review proposed changes to the draft Accounting MIB to improve performance characteristics. Since progress towards a Standard now depends on the implementation of a working MIB, the Working Group has agreed to become dormant pending initial implementation. The Internet Accounting Architecture will be revised so as not to prohibit such changes and will be submitted for publication as an RFC.

Discussions and announcements will continue on the Working Group mailing list.

The following slides summarize Nevil Brownlee's presentation to the Working Group.

Attendees

James Alfieri	jdal@troy.cc.bellcore.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Henry Clark	henryc@oar.net
Ken Hayward	crm57d@bnr.ca
Robin Iddon	robini@cix.compulink.co.uk
Kevin Jackson	kmj@concord.com
Merike Kaeo	merike@alw.nih.gov
Cynthia Mills	cmills@nnsc.nsf.net
Gregory Ruth	gruth@bbn.com
Joo Young Song	jysong@ring.kotel.co.kr
Roy Spitzer	roy.spitzer@sprint.com
Morton Taragin	vsmorty@weizmann.weizmann.ac.il

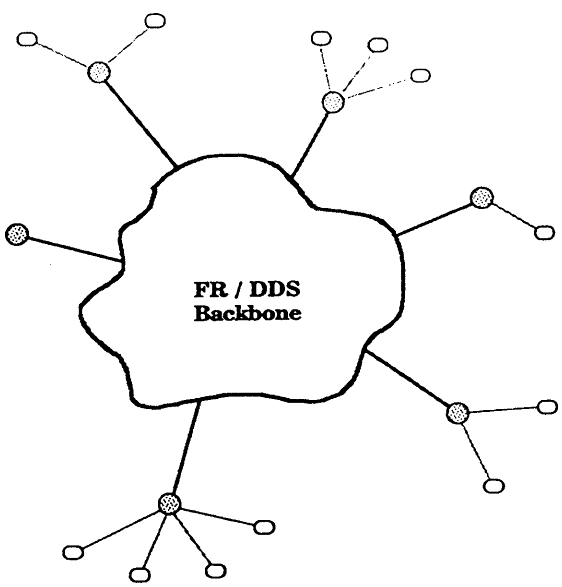

Tuia Network

Proximal Sites   Distal Sites

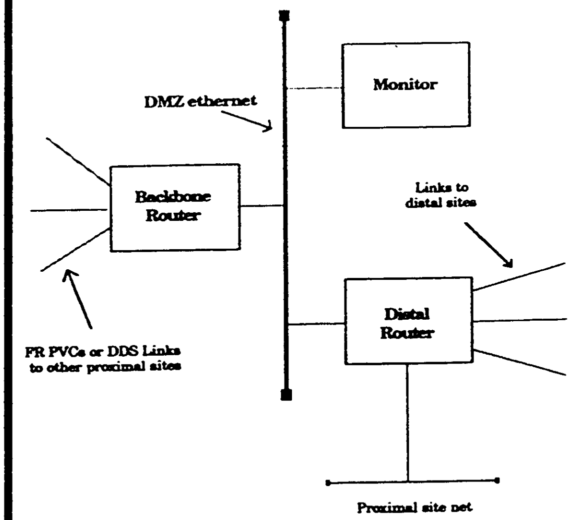

Proximal Site

Required Traffic Matrix

	World	P1 .. Pn	D1 .. Dn
World	*		
P1		*	
..			
Pn		*	
D1			*
..			
Dn			*

- Matrix is asymmetric
- Divided into regions

Collecting Counts

- Packet and byte counts
- Up and down directions
- 'Active' variable
- Saved counts
- Checkpoint => copy all counts, set 'active'

Manager/Collector

- Checkpoint at regular intervals
- Counters never reset
- Program outline:

```
for (;;) {
   make checkpoint;
   get new flows;
   get saved counts;
   pause;
   }
```

Deployment

- Meters running at Auckland, Waikato and Victoria universities
- Each controlled by a separate collector/manager at Auckland
- Need to separate manager and collectors

Meter Outline

```
for (;;) {
  get packet;
  extract variables for low
     and high addresses;

  search flow set;
  if (not found)
     try to create new flow;
  count;

  housekeeping:
     snmp request;
     keyboard;
     display;
  }
```

Defining Flows

* A flow has two addresses, ‘low’ and ‘high’
- Each address has components as follows:

  address_type
  address_value
  address_mask
  address_tally

- The two addresses make up the flow's ‘key’

Flow Directions

- 'High' address is >= 'low' address for flows
- Packet addresses are compared by a compare function (without using masks)
- Packet addr1 >= addr2: count in 'up' counters
- Otherwise count in 'down' counters

Address Masks

* When searching the flow set, addresses are ANDed with masks

- Only contiguous masks seem useful
- Address_masks used for searching flow set
- Tally_masks used in deciding whether to create new flows

Address types

IP: 4 bytes
Network | host
mask behaves like a
subnet mask

IPX: 4 bytes
Usually written in hex
Tuia uses IP numbers

DECnet: 3 bytes
1 byte area number
2 bytes host id

Storing the Flows

- Flows stored in forest of height-balanced binary search trees
- Each tree is initialised from a set of specified flows; these are the meter's 'rules.'
- Each rule has a key and a tree number
- The rules are downloaded by the manager

Searching the Flows

- Why a forest of trees?
- To handle flows which are subsets of other flows.
  e.g. IP, 130.216.3.0
  and 130.216.0.0
- Binary search will match either of the above for a subnet 3 packet.
- Handle this by putting most specific flow in tree 0, next in tree 1, ..

Creating New Flows (1)

- A tree is searched usng the address_masks.
- If the search fails, it is retried using the tally_masks.
- If the tally search fails, we move on to the next tree.
- Otherwise we create a new flow, using the packet addresses for values, and the masks

Creating New Flows (2)

Example: to make flows for all Class C subnets of 130.216 to any other Class C subnet.

Key: IP, 130.216.0.0,
255.255.255.0, 255.255.0.0
192.0.0.0,
255.255.255.0, 192.0.0.0

- We don't expect the value to mask (no subnet 0 packets)
- Tally masks allow match for any flow as specified above.
- e.g. 130.216.3.1 ->192.156.62.3
  creates flow
  130.216.3 -> 192.156.62

Search Efficiency

- Very good for a flow which is already in tree 0.
- Poor for flows in other trees, since must search each of them twice!
- Better to use a single digital search tree, for all the flow.

* Initial flows should be in a separate tree.

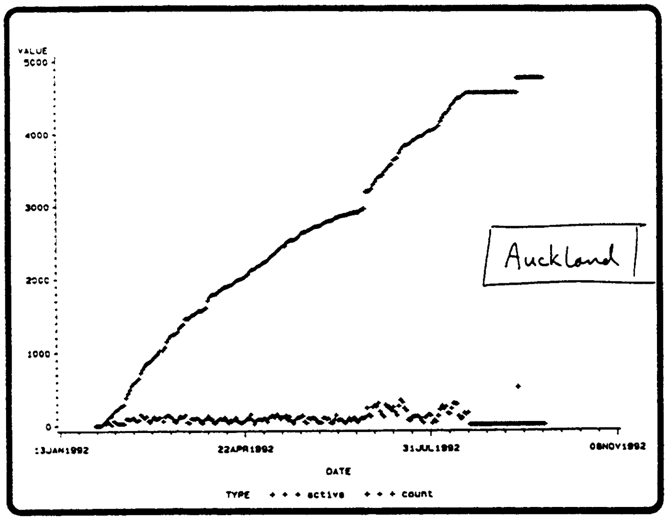
VALUE
5000
4000
3000
2000
1000
0
13JAN1992
22APR1992
31JUL1992
08NOV1992
DATE
TYPE active count
Auckland

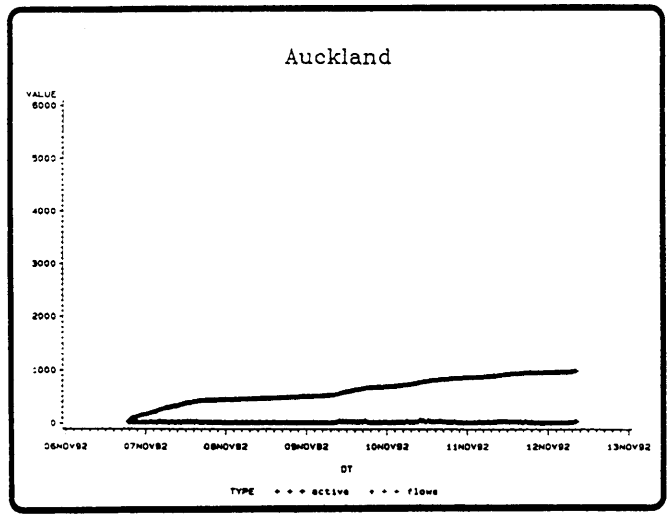
Auckland
VALUE
6000
5000
4000
3000
2000
1000
0
06NOV92
07NOV92
08NOV92
09NOV92
10NOV92
11NOV92
12NOV92
13NOV92
DT
TYPE active flows

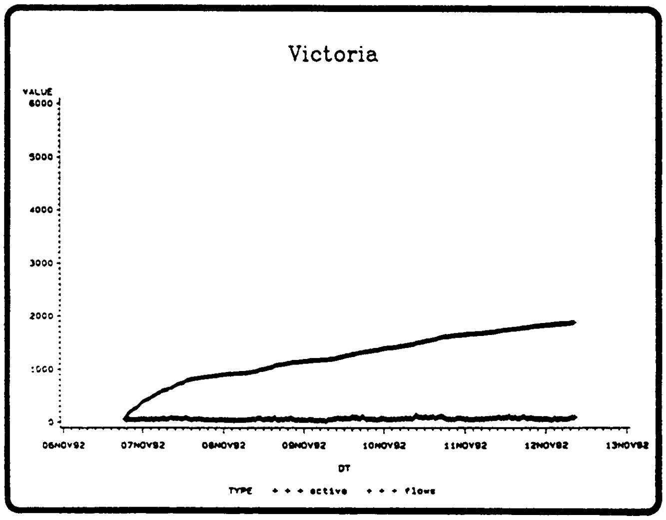
Victoria
VALUE
6000
5000
4000
3000
2000
1000
0
06NOV92
07NOV92
08NOV92
09NOV92
10NOV92
11NOV92
12NOV92
13NOV92
DT
TYPE active flows

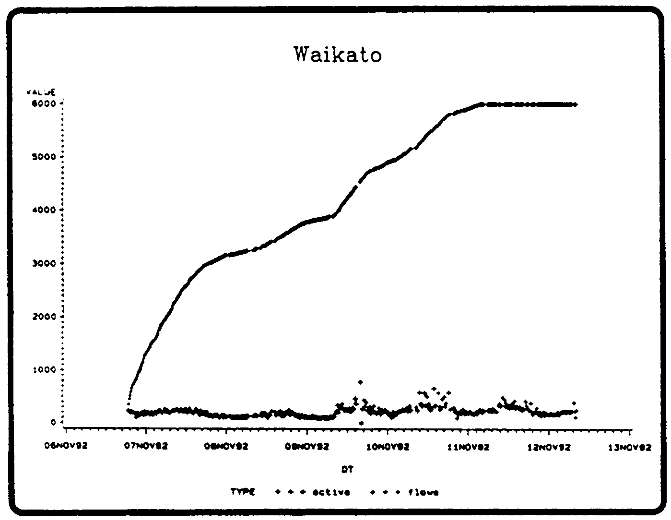
Waikato
VALUE
6000
5000
4000
3000
2000
1000
0
06NOV92
07NOV92
08NOV92
09NOV92
10NOV92
11NOV92
12NOV92
13NOV92
DT
TYPE active flows

Implementation (1)

- On a 12 MHz 286 the meter can handle about 400 packets/second.
- On a 25 MHz 386 it can handle about 1200.

* Implemented with Borland Turbo C, Waterloo TCP and CMU SNMP.

- Ethernet interface via CRYNWR packet driver.

Implementation (2)

- Manager/collector based on AARNET sample programs and CMU SNMP.
- Opaque object for recovering 'next 20 active flows.'
- Manager has routines to:
  - set/get a rule
  - make a checkpoint
  - get key for a flow
  - get next 20 active flows

Problems

- No way to re-use inactive flows. Manager re-starts trees instead.
- Need to improve search efficiency (hence packet- handling ability).
- Need to define better manager/collector interaction.

* Need a more general way to specify required flows.

Summary

- Prototype system provides useful traffic measurements.
- Performance adequate – could be better.
- Many MIB variables still to be implemented.
- Lots of insight gained from the implementation project.

2.3.9 SNMP Version 2 (snmpv2)

Charter

Chair(s):

Bob Stewart, `rlstewart@eng.xyplex.com`

Mailing Lists:

General Discussion: `snmp2@thumper.bellcore.com`
To Subscribe: `snmp2-request@thumper.bellcore.com`
Archive: `pub/davin/snmp2-archive@thumper.bellcore.com`

Description of Working Group:

This Working Group is chartered to consider technical contributions to the SNMP evolution process and to produce a single recommendation as to which contributions (or combinations or modifications thereof) should define the next generation SNMP network management framework.

The announced deadline for technical contributions to the SNMP evolution process is September 10, 1992. Any individual interested in contributing to this process should prepare and submit his/her contribution according to the requirements for detail, completeness, copyright, and format set forth in the original announcement. This Working Group is under no obligation to consider contributions that do not meet these basic requirements or contributions that are not submitted by the contribution deadline.

This Working Group has the option of (a) rejecting any or all contributions as the basis for positive evolution, (b) accepting any or all contributions as candidates for standardization, or (c) modifying or combining any or all contributions to produce consensus proposals for standardization.

The product of the Working Group will be a single recommendation to the IESG identifying those submitted specifications (or modifications thereof), if any, whose standardization as part of the SNMP framework is agreed to be warranted and desirable. The Working Group will not be chartered to produce tutorial, explanatory, advisory, or informational documents of any kind.

In its deliberations, the Working Group will take special cognizance of architectural principles on which the historic success of SNMP has rested:

(1) The SNMP framework minimizes the overall cost of a manageable network by minimizing the cost and complexity of those management system components that are most numerous.

(2) The SNMP framework fosters ubiquity of deployment by admitting the widest possible range of implementation strategies.

(3) The SNMP framework fosters operational robustness by realizing management system function as closely as possible to centers of responsible authority.

(4) The SNMP framework fosters operational robustness by locating control of resources consumed by the management activity (e.g., bandwidth, processing) as closely as possible to centers of responsible authority.

Moreover, the deliberations of the Working Group will take special cognizance of at least two aspects of evolutionary logistics:

(1) A single transition from existing SNMP technology to the next stage of SNMP evolution is highly desirable; multi-stage or protracted transitions are less desirable.

(2) Minimizing the number of distinct management technologies concurrently deployed in the Internet is highly desirable.

Consistent with the community desire for timely, deliberate progress, the Working Group may be disbanded at the time of the IETF plenary meeting in the Spring of 1993 regardless of whether or not it has produced the single recommendation required by its Charter.

This Working Group is not chartered to consider security aspects of the SNMP framework, as these are addressed as a matter of course by an existing IETF working group.

Goals and Milestones:

Done	Post as an Internet-Draft the technical proposals to the SNMP Evolution Working Group.
Done	Closing date for technical proposals.
Done	First Working Group meeting.
Done	Working Group meeting in Washington DC IETF Plenary.
Mar 1993	Working Group meeting at the IETF Plenary.
Mar 1993	Submit a proposal to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2)", 07/21/1992, J. Case, K. McCloghrie, M. Rose <draft-ietf-snmpv2-tm-05.txt>

"Coexistence between version 1 and version 2 of the Internet Network Management Framework", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-coex-05.txt>

"Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-tc-05.txt>

"Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-smi-05.txt>

"Introduction to version 2 of the Internet Network Management Framework", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-intro-05.txt>

"Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-proto-05.txt>

"Manager to Manager Management Information Base", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser

<draft-ietf-snmpv2-m2m-05.txt>

"Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)", 07/21/1992, J. Case, K. McCloghrie, M. Rose, S. Waldbusser <draft-ietf-snmpv2-mib-05.txt>

"Conformance Statements for version 2 of the Simple Network Management Protocol (SNMPv2)", 11/24/1992, J. Case, K. McCloghrie, M. Rose <draft-ietf-snmpv2-conf-02.txt>

INTERIM MEETING REPORT

Reported by James Davin/Bellcore

Minutes of the SNMP Version 2 Working Group (SNMPV2)

The SNMPV2 Working Group met October 5-6, 1992 in Knoxville, Tennessee. The Chair, Bob Stewart, called the meeting to order at 9:05 AM and circulated the attendance roster.

Agenda

- Introductions and Housekeeping
- Goals and Process
  - Credo
  - Organization
  - Stepwise Refinement to SNMP
- Easy Questions
- Proposals
- Summary

Introductions and Housekeeping

All present introduced themselves. The schedule for lunch and breaks was established. Changes to the Agenda were entertained. Local arrangements for reading email were explained.

Goals and Process

Bob presented some slides outlining his vision of where the Group was going and how it would get there. Under the rubric "Goals and Process," Bob introduced three topics:

1. Credo: As a "credo" for our collective work, Bob quoted a recent email statement by Dave Perkins as an illustration of the spirit he hoped everyone would bring to the discussion:

   > "to assist with creating a positive and long lasting solution for the community. This goal comes before any personal or company goals which I set aside when I communicate via EMAIL and attend IETF functions."

2. Organization: Bob noted that the Working Group was a chartered IETF Working Group. James "Chuck" Davin was appointed to take Minutes for this meeting. Bob stated that the Working Group would make decisions by discussion and consensus, both in meetings and via email.

   Marshall Rose was appointed editor of the Working Group documents. Bob noted that the Group would rely on Marshall to make appropriate changes without detailed instructions, except in those cases where "le mot juste" was required to capture the consensus properly. It was agreed that Marshall would clearly indicate all changes to the Working Group documents by change bars. A question was raised about whether the change bars should indicate differences from the originally posted documents or the most recent document versions. This question was deferred, because, at the moment, the latest version is the original posting.

It was noted that the Working Group Minutes would be available on-line in the usual IETF repositories.

3. Stepwise Refinement to SNMP: Bob explained what he meant by "Stepwise Refinement of SNMP" by presenting a slide with the following points:
   - Assume that SNMP is basically sound.
   - Widespread implementation.
   - Current level of technology, cooperation, understanding.
   - Choose improvements.
   - Maintain first principles.
   - High benefit-to-cost ratio.

   Bob identified the SMP proposal as the baseline documents from which the Working Group would proceed. He noted that there were eight documents, and four implementations; these latter are to be regarded as supporting the Working Group and building confidence in its baseline; the implementations will not in any way constrain the decisions or directions of the Group.

   At this point, Marshall said that all four of the SMP proponents look forward to making implementation changes based on the work of the Group.

   Bob next noted the need to coordinate with the SNMP Security Working Group. He noted the pledge of timely cooperation by the relevant Area Directors at the Cambridge IETF meeting. He noted that the liaison function is neatly realized insofar as Keith McCloghrie is both one of the SMP proponents and the co-Chair of the SNMP Security Working Group. Bob concluded by saying that, although the Group would not delve deeply into security issues, it could not and would not ignore them completely.

   Bob identified the "deliverables" of the Group as a set of Internet-Drafts, revised according to the judgement of the Group, together with a recommendation to the IESG that these documents (possibly together with revised documents produced by the SNMP Security Working Group) define the next generation SNMP framework. Assuming our ultimate agreement, the recommendation of the Group would be for Proposed Standard status for these documents.

   The schedule goal of the Group would be to finish up and, consistent with its Charter, to "drop dead" in the Spring of 1993 (shortly following the IETF plenary meeting, March 28 - April 2, 1993, in Columbus). A discussion of the schedule goal ensued:

   - Marshall and Jeff Case emphasized the need for quick progress.

- Dave emphasized that haste should in no way compromise the openness of the process.
- Chuck stressed that haste should not compromise the quality or thoroughness of the solution, because it is unlikely that revision of the standard framework will be undertaken again soon.

The Group agreed that its schedule and pace must be governed by all of these considerations. Recognizing considerable consensus, current and from the Cambridge BOF, that the work should be completed in December, the Group deferred accepting that as possible for the end of the second day.

Easy Questions

The focus of the Group turned to what Bob had identified as "Easy Questions." In this part of the meeting, Bob encouraged people to raise what they regarded as "easy issues" about the proposed framework. Those that could be quickly resolved, would be dispatched in real time. Those that proved more complicated would be noted for later consideration by the Group.

- Tracy Cox raised the question of whether or not the row-set-and-create mechanisms currently specified would be mandatory. Jeff suggested that it should be mandatory for new MIBs. Tracy sketched some scenarios in which the specified mechanism was undesirable owing to time delays between the processing of a SET request and the actual effecting of the requested alteration. The Group agreed that this point was not simple and warranted further discussion. Tracy accepted an action item to present more detail and analysis of the relevant scenarios and propose a solution.

- Satish Joshi asked whether or not the SMUX should be part of the standard framework. Marshall said that the SMUX is not part of the framework, but elements in the current proposal (the "or" table in the SMP MIB) permit the use of SMUX or SMUX-like mechanisms.

- Chuck expressed a general concern about uncertain conformance requirements and raised the particular question of whether or not use of the AGENT-CAPABILITIES macro would be required of conformant implementations. Chuck proposed that the specification language be clarified to either make the macro a requirement or to omit it from the standard framework (as is now the case). Keith said that use of the CAPABILITIES macro would be required because of its relationship to the "or" table in the proposed MIB. Marshall argued that use of the macro should not be required because it was not relevant to all of the constituencies of the proposal: it includes vendor tools, user tools, and Working Group tools. Chuck said that the different requirements in each of these three contexts should be written down unequivocally. Jon Saperia said that he favored requiring the AGENT-CAPABILITIES macro mandatory. After some discussion, it was proposed that the word "should" be applied to this issue. Chuck said that he found the use of "should" acceptable only if no other parts of the framework depended for their function or their unambiguous definition on the presence or use of this macro. The Working Group agreed to using "should" provided that this condition was met.

- Dave suggested that a list of "required reading" be prepared to help give everyone a common context for discussion. The Group agreed, and Dave accepted an action item. *Editor's Note (md): A copy of this list is available via ftp under snmpv2-minutes-92oct.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

- Dave also proposed that the new standards documents include a glossary of key terms. It was suggested that Marshall undertake this task and include the glossary in the introductory document. The Group agreed that Marshall would consider the effort involved and report back after lunch.

- Dave suggested that the Group prepare a detailed analysis of how well the baseline proposals addressed the concerns raised at the Atlanta IETF session on perceived deficiencies in SNMP. Jeff said that the basis of the current proposals was a list of problems he had maintained since 1988 that included the IETF session, a previous INTEROP BOF, and some additional items as well. Marshall said that preparing such an analysis would be too much effort. Jeff elaborated, saying that each item on the list was evaluated according to several criteria (e.g., compatibility with installed base, performance, impact in existing MIB object access methods).

- Peter Wilson raised the question of party proliferation. After brief discussion, this was identified as an issue for the SNMP Security Working Group, and further discussion of this topic was deferred for later in this meeting.

- Dave suggested that the Group consider a revision of the MIB-2 interfaces table. The consensus of the Group was that this was not in the scope of its Charter as it could be handled in the normal course of IETF business. The Group agreed to a recommendation that this work be pursued soon after the SNMP evolution work is completed.

- Dave raised a question about the definition of sysObjectId. It is ambiguous, but is also used by SNMP 2. Steve Waldbusser said that sysObjectId should identify the combination of software and hardware that makes up the managed system. Jeff agreed with Steve, and described various strategies used by OEM software vendors to address this question. Marshall said that the actual definition of sysObjectId is not ambiguous, but that the example text that follows it is bad.

- SMP assumes that sysObjectId names a protocol/MIB implementation but (not necessarily) the type of box (e.g., a bridge, router, etc.). Are we comfortable with this assumption? Do we want to legislate rules for assigning sysObjectId?

  Proposal: either fix the "or" table so that it doesn't refer to the (arguably ambiguous) sysObjectId or else define a new MIB table that tells what MIB objects are supported. Action (Dave): prepare a proposal if needed. If the Group agrees that the interpretation of sysObjectId that is implicit in the baseline proposals is correct, then this consensus must be documented in the standard.

After lunch, Bob suggested that the Group change its discussion mode to focus on brief discussions that would either result in quick resolution of topics or place those topics on a "deferred issues list" for later discussion. *sditor's Note (md): The deferred issues list is available via ftp under snmpv2-minutes-92oct.txt.*

There was a discussion of how Working Group consensus should be achieved, whether email or face-to-face meetings would dominate. Bob explained that neither would dominate. He would attempt to assure progress by posing straw conclusions and calls for consensus by requesting strong objections, but ultimately the Group would be governed by the soft principles that have been traditional in the IETF.

Proposals

The remainder of the day was spent in considering various proposals for amendment of the baseline documents.

Bob presented the list of pending proposals collected from the mailing list:

- Reliable Traps - Chuck Wegrzyn
- Party Proliferation - Pete Wilson
- Remove Counter64 Time Limit - Pete Wilson
- NAME Clause - Pete Wilson
- OID Optimization - Ilan Raab
- Redefinition of "Manager" - Bob Stewart
- Date and Time Textual Convention - Jon Saperia

He asked the Group for others and added:

- Miscellaneous changes from Jeff Case.
- Miscellaneous changes from Chuck Davin.
- Miscellaneous SMI changes from Dave Perkins.
- Ideas on Get Bulk OID compression from Satish Joshi.
- Two ideas from Robert Snyder.
  - Identifying MIB objects with constant values.
  - Contents of Set Responses.

Editor's Note (md): A detailed account of the discussions pertaining to each of the Proposals listed below is available via ftp under snmpv2-minutes-92oct.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

- Get Bulk OID Compression: Satish spoke about his ideas on Bulk Retrieval. He suggested compressing the OIDs in the varbind list of responses to Bulk Retrieval requests. He observed that the OIDs in the 2nd and subsequent repetitions in a response could be abbreviated. Compression could occur in the context of the original request or in the context of the preceding varbind.

- NAME Clause: Peter Wilson led a discussion for about 30 minutes on the addition of a NAME clause to the OBJECT TYPE macro. The Group concluded that the NAME clause should not be introduced because one could get the same effect by well-chosen object descriptors. However, it was also agreed that this sort of information might be included in macros exclusively for management stations. Dave Perkins accepted an action item to explore the feasibility of such a notation.

- Remove Counter64 Time Limit: Peter then offered a proposal that the time limit associated with the use of the 64-bit counter type be excised from the baseline documents. The consensus of the Group was to leave the restriction as it is.

- New Textual Convention: the Group took up Jon Saperia's proposal for a new Textual Convention for expressing dates. The Group spent some time tweaking the details of this proposal. Jon accepted an action item to post the agreed, amended proposal to the mailing list.

- New MIB Object Clause: Robert Snyder proposed a new MIB object clause that identifies an object as having a constant value: a manager need only retrieve it once. Marshall asked what macro it should go in. Chuck suggested that this information was really more of a manager aid than an essential property of a MIB object. Robert Snyder accepted an action item to go examine some MIBs and report back on these questions and on the number of cases in which this idea would yield actual benefits.

- Miscellaneous Changes from Jeff Case: Unless otherwise noted, all changes proposed by Jeff were accepted by the Group. *Editor's Note (md): A detailed listing of these changes is available via ftp under snmpv2-minutes-92oct.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

- Redefinition of "Manager": Bob Stewart offered a proposal that would clarify the definition of "manager" and "agent" in the framework:

  A "manager" is any active network management component that observes or controls one or more network devices, whether locally through implementation-specific interfaces or remotely via SNMP, with or without a human interface. Such a manager may use any subset of the SNMP manager functions. Definition of "agent" is unchanged.

- Miscellaneous SMI Changes: Dave Perkins led a discussion on a list of proposed changes to the SMI that he prepared. Dave actually submitted two separate lists of issues/suggested changes. One list covered the textual conventions SMP document.

It had 10 numbered points. The other list covered the SMI SMP document. It had 50 numbered points. In the two days at Knoxville, Dave was allocated approximately 2-3 hours to go over the lists. Only eight items from the SMI list were covered. The meeting attendees were given a paper copy of the lists. An electronic copy is available from the archive at thumper.bellcore.com. Dave plans to update the lists and submit them for consideration at a future meeting of the Working Group.

Summary

At this point, the Chair began the identification of residual issues and discussion of future schedule and meetings.

Bob said that, in order to encourage people to air proposals early, he would promise on the mailing list that proposals posted by Monday, November 9th would be assured time for discussion at the November meeting, while others might only be considered schedule permitting. The Group generally approved of this plan.

Bob emphasized the need for doing "homework" before the next meeting. An interim meeting date was set for December 14th in Atlanta. This date will only be used if the Group needs it.

Marshall said that new Internet-Draft documents reflecting the discussions at this meeting would be posted by Thursday.

The Group agreed that its work should be completed in December. If work can not be completed at the Washington IETF, the Group will hold a meeting at Georgia Tech in Atlanta. The suggested date for this meeting was December 14th.

Bob proposed that the deferred discussion on party proliferation be referred entirely to the SNMP Security Working Group. Keith accepted and the Group did not object.

The meeting closed with a brief review of the DISPLAY-HINT discussion in particular and the more general question of whether the Group should focus on technology that is primarily an aid to managers or leave that for future work. Chuck raised the question of whether display hint should be broken into two parts:

1. Representation on the wire.
2. Display format hint for a user interface.

The consensus was that the current text would stand for now pending any future proposals on this question.

Editor's Note (md): A list of changes to the SNMP Version 2 Documents is available via ftp under snmpv2-minutes-92oct.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Attendees

Steve Alexander	stevea@i88.isc.com
Uri Blumenthal	uri@watson.ibm.com
Jeff Case	case@cs.utk.edu
Tracy Cox	tacox@sabre.bellcore.com
James Davin	davin@bellcore.com
Michael Davison	davison@fibercom.com
Taso Devetzis	devetzis@bellcore.com
Gary Haney	hny@ornl.gov
Matthew Hecht	mhecht@cs.utk.edu
Susan Hicks	hny@ornl.gov
Satish Joshi	sjoshi@synoptics.com
Mark Kepke	mak@cnd.hp.com
Kenneth Key	key@cs.utk.edu
Michael Kornegay	mlk@bir.com
Deirdre Kostick	dck2@sabre.bellcore.com
Cheryl Krupczak	cheryl@cc.gatech.edu
Robert Lushbaugh	lus@ornl.gov
Keith McCloghrie	kzm@hls.com
David Minnich	dwm@fibercom.com
David Perkins	dperkins@synoptics.com
Shawn Routhier	sar@epilogue.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jon Saperia	saperia@tcpjon.ogo.dec.com
Robert Snyder	snyder@cisco.com
Bob Stewart	rlstewart@eng.xyplex.com
Maurice Turcotte	dnmrt@interlan.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Bert Wijnen	wijnen@vnet.ibm.com
Peter Will	will@isi.edu
Steven Wong	wong@took.enet.dec.com
Chris Young	cyoung@ctron.com
Kiho Yum	kxy@nsd.3com.com

CURRENT MEETING REPORT

Reported by Marshall Rose/DBC

Minutes of the SNMP Version 2 Working Group (SNMPV2)

The Agenda was reviewed and approved. In the majority of cases there was significant protracted discussion. In the interests of brevity, the decisions reached by the Working Group are summarized below.

Outstanding Procedural Issues were Discussed

- Deadline to Finish: Although a meeting slot has been identified for December, the Chair wanted to try to conclude business this week as several slots were scheduled for the Working Group. There was strong consensus that an additional meeting should be avoided if at all possible.

- No New Proposals: There was consensus that only "bug fixes" and "show stoppers" would be addressed after the conclusion of this meeting. The one exception is the row-creation and associated proposals, see III.10 below.

- Deadlock Shelf: There was consensus that deadlock shelf would remain in place for proposals for which consensus could not be reached. From time to time, these items will be taken off the shelf to see if there is a new consensus. *Editor's Note (md): Additional details on the Deadlock Shelf are available via ftp under snmpv2-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

- More Implementation Experience: There was consensus that no additional implementation requirements would be placed on the documents prior to the Working Group completing its work.

New Proposals

In order to facilitate the discussion, each presenter was required to first demonstrate a problem, before presenting a solution.

- Tracy Cox demonstrated that delayed operations (e.g., due to slow proxy) was a problem. Discussion of solutions was tabled until after the SNMP Security Working Group meeting later that evening. At that meeting, two proposals were suggested. As such, this issue has been moved to the SNMP Security Working Group.

- Dave Arneson suggested that efficient retrieval of tabular objects was a problem. There was consensus that, in bandwidth-limited environments, retrieval should be more efficient. However, there was no consensus that this problem was specific to tables.

- Anil Rijsinghani was absent, but a colleague demonstrated that auto-discovery of SNMP agents was a problem. There was consensus that the proposal was on the right track, but that this work could proceed independently from the SNMPv2 effort.

- Dave Perkins suggested that retrieval of non-rectangular tables was a problem and there was sufficient interest to look at the solution. However, there was consensus that there wasn't enough of a problem to warrant the solution.

- Dave Perkins and Sam Roberts presented their SMI issues. A few of these were postponed to an off-line editing meeting. *Editor's Note (md): A more detailed accounting of this issue is available via ftp under snmpv2-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

- Anil Rijsinghani was absent, but a Jon Saperia discussed a need for an unsigned 64-bit type, however the Group could not achieve consensus on any adequate choke rule. Eventually, there was consensus that despite some usefulness, such a type would not be added.

- Marshall Rose described a problem in the definition of the TEXTUAL-CONVENTION macro along with a solution. Textual conventions are now written as:

  ```
  <name> ::= TEXTUAL-CONVENTION
  <clauses ...>
  SYNTAX <syntax>
  ```

  This is necessary due to macro definition restrictions in ASN.1

- Jeff Case suggested that the limitations on enumerated values in INTEGERs was causing problems when translating MIBs written by other groups. There was consensus that the limitations should be removed with a recommendation that newly defined objects follow the old rules.

- Bill Norton presented the row-creation portion of the multi-part proposal by Guenther Schreiner, et. al. Discussion lasted for over two hours.

  Group consensus was that Create/Delete operators were not the solution to row creation, but there is a problem with complexity and multiple ways to use RowStatus. Jeff Case took an action to reconsider this problem. The Chair set a deadline of December 4th for final resolution on this issue and consideration of the other proposals that came with this one.

  In comparing the row-creation proposal to the RowStatus mechanism, it was agreed that the row-creation proposal did not solve the general problem of row creation, as:

  1. Sometimes multiple PDU exchanges were necessary in order to create a row, e.g., either because of a resource negotiation process between the agent and

manager, or because there might be too much data to fit in a single creation request.

2. The response from the creation PDU added varbinds in order to indicate what mandatory columns are missing. However, this could make the request too big to send back.

3. The creation request is not idempotent due to potential packet duplication and loss from the underlying transport service (i.e., the request gets duplicated, the first succeeds, but the response is lost, the second fails, and its response is returned.)

It was also observed that with the RowStatus mechanism, creation could be done in a single exchange, if the DEFVAL clause was active and the manager did a set to active. However, it was agreed that this text should be made more clear. An action was taken by Steve Waldbusser.

After much discussion, there was consensus that the real problem was that the community had three requirements:

1. A single, consistent way to do row-creation.
2. Some row-creations take more than 1 exchange.
3. Some agent writers wish to implement a simple table in such a way so that row creation must be done in a single exchange.

An action was taken by Jeff Case on behalf of the four SMP authors to see if some solution could be found which had these properties:

1. Avoided the tooBig problem.
2. Dealt with the discovery problem of missing columns and defvals.
3. Avoided stateful behavior.

Jeff Case was careful to stress that this issue had been looked at in great detail by the SMP authors prior to the publication of the SMP specification, and he was doubtful that a solution could be found.

Timetable

There was strong consensus that the row-creation issue and other associated, unresolved proposals. would be given until Friday, December 4th to achieve resolution.

There was complete consensus that the final deadline for comments on the nine SNMPv2 documents would be Friday, December 11th unless the SNMP Security effort raised new issues, then the documents would be sent forward to the IESG with a recommendation for advancement to the standards-track from the Working Group.

Finally, it was observed that the SNMPv2 documents could not go forward without the revisions the three SNMP Security documents. As such, it was suggested that the membership of the SNMPv2 Working Group now focus its energies on the issues before the SNMP Security Working Group.

Documents

Revised versions of the SNMPv2 documents were submitted to the Internet-Drafts area. In addition, "unofficial" copies are available via anonymous ftp:

```
host: ftp.ics.uci.edu
area: mrose/snmpv2/
files: *.txt
```

The documents can also be retrieved via e-mail:

```
mailbox: archive-server@ftp.ics.uci.edu
body: MIMESEND mrose/mh-mime/snmpv2
```

These documents will be removed once the actual Internet-Drafts are announced.

Attendees

Elizabeth Adams	adamse@attmail.com
Steve Alexander	stevea@i88.isc.com
David Arneson	arneson@ctron.com
Toshiya Asaba	asaba@wide.sfc.keio.ac.jp
Fred Baker	fbaker@acc.com
Jim Barnes	barnes@xylogics.com
Brian Bataille	bataillebc@afotec.af.mil
Andy Bierman	abierman@synoptics.com
Fred Bohle	fab@interlink.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Stephen Bush	sfb@ncoast.org
Jeff Case	case@cs.utk.edu
John Chang	changj@ralvm6.vnet.ibm.com
Szusin Chen	szusin.chen@eng.sun.com
Robert Ching	rching@nat.com
Chris Chiotasso	chris@andr.ub.com
Bobby Clay	clay@eagle.msfc.nasa.gov
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com

Juan Cruz	juan@dss.com
Dave Cullerot	cullerot@ctron.com
Cathy Cunningham	cmc@microcom.com
James Davin	davin@bellcore.com
Michael Davis	mad@spirit.clearpoint.com
Michael Davison	davison@fibercom.com
Cynthia Della Torre	cindy@gateway.mitre.org
Manuel Diaz	diaz@davidsys.com
Jon Dreyer	Jon.Dreyer@east.sun.com
Jacques Dugast	dugast@issy.cnet.fr
Donald Eastlake	dee@ranger.enet.dec.com
David Engel	david@ods.com
Michael Erlinger	mike@jarthur.claremont.edu
Roger Fajman	raf@cu.nih.gov
Daniel Fauvarque	dfauvarq@france.sun.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Shari Galitzer	shari@mitre.org
Shawn Gallagher	gallagher@quiver.enet.dec.com
Richard Graveman	rfg@ctt.bellcore.com
Maria Greene	mngreene@eng.xyplex.com
Michel Guittet	guittet1@applelink.apple.com
Robert Gutierrez	gutierre@nsipo.nasa.gov
William Haggerty	haggerty@ctron.com
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Ed Heiner	eah@pau.synnet.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
David Husak	dave@synnet.com
Robin Iddon	robini@cix.compulink.co.uk
Kevin Jackson	kmj@concord.com
Ole Jacobsen	ole@interop.com
Ronald Jacoby	rj@sgi.com
Frank Kastenholz	kasten@ftp.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Jong Yeol Kim	kimjy@ring.kotel.co.kr
Andrew Knutsen	andrewk@sco.com
Michael Kornegay	mlk@bir.com
Deirdre Kostick	dck2@sabre.bellcore.com
Michael Laufer	mlaufer@bbn.com
Mark Lewis	mlewis@telebit.com
David Lin	lind@janus-ccm.zenith.com
David Lindemulder	dcl@mtung.att.com
Benjamin Lisowski	Ben.Lisowski@sprint.sprint.com

David Liu	dliu@bnr.ca
John Lunny	jlunny@twg.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
William McKenzie	mckenzie@ralvma.vnet.ibm.com
Donna McMaster	mcmaster@synoptics.com
John Medicke	medicke@ralvm11.vnet.ibm.com
Douglas Miller	dmm@telebit.com
David Minnich	dwm@fibercom.com
Mohammad Mirhakkak	mmirhakk@mitre.org
Rohit Mital	rm@protools.com
George Mouradian	gvm@arch3.att.com
Patrick Mullaney	mullaney@ctron.com
Daniel Myers	dan@nsd.3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Hien Nguyen	h.nguyen@sprintintl@sprint.com
Mo Nikain	mo@bss.com
Tom Nisbet	nisbet@tt.com
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
David Perkins	dperkins@synoptics.com
Carl Powell	cpowell@bbn.com
Ilan Raab	iraab@synoptics.com
Richard Ramos	ramos@mtunm.att.com
Venkat Rangan	venkat@geoduck.matrix.com
Louise Reingold	l.reingold@sprint.sprint.com
Sam Roberts	sroberts@farallon.com
Kary Robertson	kr@concord.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Shawn Routhier	sar@epilogue.com
Chris Rozman	chrisr@usr.com
Assaf Rubissa	asaf@fibhaifa.com
Jon Saperia	saperia@tcpjon.ogo.dec.com
Michael Sapich	sapich@conware.de
Michael Scanlon	scanlon@interlan.com
Sam Schaen	schaen@mitre.org
John Seligson	johns@ultra.com
Paul Serice	serice@cos.com
Chris Shaw	cshaw@banyan.com
Timon Sloane	timon@rahul.net
Robert Snyder	snyder@cisco.com
Joo Young Song	jysong@ring.kotel.co.kr
Roy Spitzer	roy.spitzer@sprint.com

Einar Stefferud	stef@nma.com
John Stephens	john@cayman.com
Bob Stewart	rlstewart@eng.xyplex.com
Kaj Tesink	kaj@cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Ahmet Tuncay	atuncay@synoptics.com
Warren Vik	wmv@i88.isc.com
Ioannis Viniotis	candice@ececho.ncsu.edu
Steven Waldbusser	waldbusser@andrew.cmu.edu
Alice Wang	alice.wang@eng.sun.com
James Watt	james@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Gerry White	gerry@lancity.com
Peter Wilson	peter_wilson@3com.com
Steven Wong	wong@took.enet.dec.com
Randall Worzella	worzella@ralvm29.unet.ibm.com
Daniel Woycke	woycke@smiley.mitre.org
Honda Wu	honda@nat.com
Jeff Yarnell	jeffya@protools.com
Kiho Yum	kxy@nsd.3com.com

2.3.10 Token Ring Remote Monitoring (trmon)

Charter

Chair(s):

Michael Erlinger, mike@jarthur.claremont.edu

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Discussion and agreement on models and terminology. Comparison of RMON architecture and Token Ring requirements. Assign author and editor responsibilities.
Done	Working Group meeting at San Diego IETF.
Mar 1992	Post Internet-Draft of the Token Ring Monitoring MIB.
Done	Working Group meeting at Cambridge IETF.
Nov 1992	Submit the Token Ring MIB to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Mike Erlinger/Harvey Mudd College

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

Token Ring

The Group met for two sessions on Thursday. The major Working Group activity was a group by group discussion of the working draft for the Token Ring RMON. The end result was numerous changes for the author. The goal is to incorporate those changes by December 15th. To be followed by a call for consensus as to requesting that Token Ring RMON become a Proposed Standard with RFC status.

There were a couple of cantankerous issues, in particular the order table, for which consensus was reached. It was also stressed that further discussion should take place on the mailing list.

Attendees

David Arneson	arneson@ctron.com
Andy Bierman	abierman@synoptics.com
John Chang	changj@ralvm6.vnet.ibm.com
Manuel Diaz	diaz@davidsys.com
Michael Erlinger	mike@jarthur.claremont.edu
Maria Greene	mngreene@eng.xyplex.com
Paul Griffiths	griff@chang.austin.ibm.com
Daniel Hansen	dan@ngc.com
John Hopprich	hopprich@davidsys.com
Jeff Hughes	jeff@col.hp.com
Robin Iddon	robini@cix.compulink.co.uk
Kevin Jackson	kmj@concord.com
Mark Kepke	mak@cnd.hp.com
Keith Klamm	klamm@ods.com
Kenrick Kutzler	kkutzler@synoptics.com
Carl Madison	carl@startek.com
John Medicke	medicke@ralvm11.vnet.ibm.com
Rohit Mital	rm@protools.com
Patrick Mullaney	mullaney@ctron.com
Daniel Myers	dan@nsd.3com.com
Joe Nguyen	jnguyen@synoptics.com
Tom Nisbet	nisbet@tt.com
Venkat Rangan	venkat@geoduck.matrix.com
Sri Reddy	srireddy@synoptics.com
Dan Romascanu	dan@lannet.com

Avraham Rosenbach	armon@armon.hellnet.org
Assaf Rubissa	asaf@fibhaifa.com
Michael Scanlon	scanlon@interlan.com
Anil Singhal	anil@frontier.com
Richard Sweatt	rsweatt@synoptics.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Ian Wilson	ianw@spider.co.uk
Jeff Yarnell	jeffya@protools.com
Kiho Yum	kxy@nsd.3com.com

2.3.11 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: pub/ups-mib/mail-archive@ucs.utk.edu

Description of Working Group:

This Working Group will produce a document that defines MIB objects for use in monitoring and (possibly) control of both high-end and low-end UPSs and related systems (e.g., power distribution systems or power conditioning systems). Related devices may be addressed in this effort to the extent that the primary focus on UPSs is not compromised.

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

At its discretion, the Working Group may fulfill its Charter by the development of distinct MIB definitions for UPS systems of differing capabilities, but the number of MIB definitions produced by the Working Group will not exceed two.

At its discretion, the Working Group may produce an additional document defining traps that support the management of UPSs.

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

Because the structure of UPS implementations varies widely, the working group shall take special care that its definitions reflect a generic and consistent architectural model of UPS management rather than the structure of particular UPS implementations.

Goals and Milestones:

Done	Hold Interim Working Group meeting to review draft.
Nov 1992	Post initial draft MIB to Internet-Drafts.
Mar 1993	Meet at March IETF meeting to reach closure on MIB document.
Apr 1993	Submit the UPS MIB to the IESG for consideration as a Proposed Standard.

INTERIM MEETING REPORT

Reported by Jeff Case/UTenn

Minutes of the Uninterruptible Power Supply Working Group (UPSMIB)

The UPSMIB Working Group held its first meeting on October 16, 1992 in Dallas, Texas. This meeting was scheduled, as discussed at the formulative BOF, to coincide with the NETWORLD trade show held earlier that same week. This scheduling was at the request of the Working Group in order to maximize productivity while minimizing impact on travel budgets. The meeting began with introductions, circulation of the attendance sheet, and other administrative matters.

The second portion of the meeting was devoted to data collection. The attendees were divided into data providers and data collectors. Approximately one data provider was allocated per vendor represented. Approximately one data collector was allocated per MIB group. Data was collected for the cross product of each represented product and each proposed MIB object class. Each was rated using the following key:

UI:	Unimplementable in this product.
SM:	Static (constant) determined at time of manufacture.
SC:	Static (constant) determined at time of configuration or via a set request.
SD:	Always returns a static constant default value (such as unknown or noneOfTheBelow).
DR:	Can be implemented dynamic read/only, i.e., to reflect current values of instrumentation in the running system.
DW:	Can be implemented dynamic read/write.
NC:	No comment – no data provided.

This data was collected to help assess the implementability of each object. There were, as anticipated, several side benefits from the effort.

- First, it provided opportunities for Working Group members to work together on a shared goal – it got the members talking to each other in a non-confrontational setting. As a result, members got to know one another better.

- Second, it caused members to consider each MIB variable closely and uncovered considerable vagueness in the semantic content of many of the description clauses.

- Third, it identified meta-issues such as the difference between precision and accuracy. The matrix will be posted to the mailing list: ups-mib@cs.utk.edu (ups-mib-request@cs.utk.edu for administrative add/change requests).

The third phase of the meeting was devoted to discussion of the results of the survey, group by group. The editor was directed to make changes in the objects in the next draft as a result of these discussions. Not all groups were discussed owing to the limited time available in a one day meeting.

The final phase of the meeting was discussion of the future. The Group elected to not meet at the 25th IETF plenary due to a time conflict with the COMDEX show. A meeting at the COMDEX show was considered briefly but the Chair is unable to leave the plenary in order to meet with the Working Group in another city while the plenary is being held in Washington. The mailing list will be used to continue discussions and to decide on the next meeting date(s).

Attendees

Mark Allgeier	
John Bell	garringer@compuserve.com
Tom Brennan	brennan@exide.com
Jeff Case	case@cs.utk.edu
Yu Chin	76500.3160@compuserve.com
Roger Draper	rdraper@cerf.net
Bill Elliot	
Theodore Greene	
William Humphreys	
Rich Johnson	
Steve Loboyko	72360.2436@compuserve.com
Jess Marinez	
Lawren Markle	
Andrew McCartney	
Gary Mook	
Bob Nerz	
Rod Pullen	s.shepard@deltec.com
Doug Rademacher	
James Rigney	76500.3160@compuserve.com
Adam Stolinski	stolinsk@cerf.net
Grek Wilterdink	
Pete Yoest	
Brian Young	

2.3.12 X.25 Management Information Base (x25mib)

Charter

Chair(s):

Dean Throop, throop@dg-rtp.dg.com

Mailing Lists:

General Discussion: x25mib@dg-rtp.dg.com
To Subscribe: x25mib-request@dg-rtp.dg.com
Archive: dg-rtp.dg.com:x25mib/Current.Mail

Description of Working Group:

This Working Group will produce a set of three documents that describe the Management Information Base for X.25. The first document will specify the objects for the X.25 Link Layer. The second document will specify the objects for the X.25 Packet Layer. The third document will specify the objects for managing IP over X.25. The Working Group need not consider the Physical Layer because the "Definition of Managed Objects for RS-232-like Hardware Devices" already defines sufficient objects for the Physical Layer of a traditional X.25 stack. Any changes needed at the Physical Layer will be addressed as part of that activity.

The X.25 object definitions will be based on ISO documents 7776 and 8208 however nothing should preclude their use on other similar or interoperable protocols (i.e., implementations based on CCITT specifications).

The objects in the Link and Packet Layer documents, along with the RS-232-like document, should work together to define the objects necessary to manage a traditional X.25 stack. These objects will be independent of any client using the X.25 service. Both of these documents assume the interface table as defined in MIB-II contains entries for the Link and Packet Layer interfaces. Thus these documents will define tables of media specific objects which will have a one to one mapping with interfaces of ifType ddn-x25, rfc877-x25, or lapb. The objects for the IP to X.25 convergence functions will be defined analogously with the ipNetToMedia objects in MIB II.

The Working Group will endeavor to make each layer independent from other layers. The Link Layer will be independent of any Packet Layer protocol above it and should be capable of managing an ISO 7776 (or similar) Link Layer provider serving any client. Likewise the X.25 Packet Layer objects should be independent of the Link Layer below it and should be capable of managing an ISO 8208 (or similar) Packet Layer serving any client.

The Working Group will also produce a third document specifying the objects for managing IP traffic over X.25. These objects will reside in their own table but will be associated with the X.25 interfaces used by IP. These objects will not address policy decisions or other implementation specific operations associated with X.25 connection management decisions except as explicitly described in existing standards. These objects will manage the packet flow between IP and the X.25 Packet Layer specifically including observation of packet routing and diagnosis of error conditions. Progress on the Link and Packet Layer documents will not depend on progress of the IP over X.25 document. The IP over X.25 document will proceed on a time available basis after work on the Link and Packet Layer documents and as such the Link and Packet Layers may be completed before the IP over X.25 work.

All documents produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions (such as Concise MIB format). To the extent feasible, the object definitions will be consistent with other network management definitions. In particular ISO/IEC CD 10733 will be considered when defining the objects for the X.25 Packet Layer.

Goals and Milestones:

Done	Working Group meeting as part of IETF to review documents.
Done	Distribute first draft of documents and discuss via E-mail.
Done	Distribute updated documents for more E-mail discussion.
Done	Submit the LAPB MIB to the IESG for consideration as a Proposed Standard.
Done	Submit the X.25 Packet Layer MIB to the IESG for consideration as a Proposed Standard.
Nov 1992	Submit the Multiprotocol over X.25 MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"SNMP MIB extension for MultiProtocol Interconnect over X.25", 10/07/1991, Dean Throop <draft-ietf-x25mib-ipox25mib-04.txt>

Request For Comments:

RFC 1381	"SNMP MIB Extension for X.25 LAPB"
RFC 1382	"SNMP MIB Extension for the X.25 Packet Layer"

2.4 OSI Integration Area

Director(s):

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

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

The following Working Groups and/or BOFS in the OSI area met at the Washington IETF:

FTPFTAM	FTP-FTAM Gateway BOF
MHSDS	MHS-DS
NOOP	Network OSI Operations
OSIDS	OSI Directory Services
WHOIS	Shared Whois Project BOF
X400OPS	X.400 Operations
THINOSI	Xwindows over OSI and Skinny stack BOF

The MIME-MHS Working Group, dealing with mapping MIME into X.400(88) and back, did not meet in Washington. The Group finished the three drafts, and will submit them on the standards track.

FTP-FTAM Gateway BOF (FTPFTAM)

The FTP-FTAM Gateway Internet-Draft was previously discussed in the now defunct OSI-General Working Group. Josh Mindell and Robert Slaski gave a brief presentation of the status of the work since the July 1991 IETF, and described the changes introduced into the recently posted Internet Draft. Much of the work introduced is not radically new, but is not reflected in the current implementations. Steve Hardcastle-Kille indicated that the ISODE Consortium would be willing to consider implementation to upgrade the existing ISODE gateway if consortia members request it (and $ up).

The Working Group discussed quite frankly, the difficulties of sustaining interest in this project, which is like a classic chicken-egg situation. Absent an RFC to cite in procurement requests, it has been difficult to foster additional implementation efforts. The BOF requested that the OSI Area Directors inquire as to the possibility of progressing the Internet-Draft, which has been implemented, to Proposed Standard. It is expected that during the review and development period following the recommendation to Proposed, at least the two currently known implementations will be made to conform and interoperate against the draft.

MHS-DS Working Group (MHSDS)

At its meeting in Washington, the MHSDS Working Group accomplished the following:

- Approved an updated Charter which adds coordination of a pilot project to the scope of the Working Group.

- Wrote a formal statement of purpose for the pilot project, and established concrete goals, a time-frame, criteria for measuring success, participants, and a coordinator for it.

- Reviewed four of the Group's nine documents-in-progress, recommended two of them for advancement as Proposed Standards, and made good progress on its principal routing document.

Network OSI Operations Working Group (NOOP)

NOOP talked about the revision of RFC1139 and also the Tools RFC draft Both of these need some revision. Both need some specific text about MUST and SHOULD, etc. The Tools RFC is going to have the MIB information removed until there is a routing table MIB. Then the document will be modified to point to the routing table MIB. After the documents are revised, we will put them up as Internet-Drafts and try to move them on to Proposed Standard.

Some folks are going to work on getting a group together to make a routing table/forwarding table MIB. (Dave Piscitello is heading this effort).

Sue showed us the latest survey of OSI in the Internet. Some comments were made as to changes to the format of the survey to make it easier to fill out and understand. Sue is going to modify the survey and send it out to the Group. The survey results are availabe on merit.edu.

The second session of NOOP was a tutorial for folks a little less familiar with OSI and deployment issues. After the tutorial we discussed a particular network's topology and how it might be broken up into areas and domains.

OSI Directory Services Working Group (OSIDS)

The Working Group discussed several Internet-Drafts:

- Strategic Deployment of Directory Services on the Internet. No comments, will be published as Informational RFC.

- DUA Metrics. No comments, will be published as Informational RFC DSA metrics hold until tested it against an implementation.

- LDAP (Lightweight Access Protocol). This and associated syntax document will be submitted as Proposed Standard RFCs.

The Group discussed the RFC 1373, on portable DUAs and decided that the document is confusing and should not have been published as such.

Several drafts on representing network information and other non-personal information in the Directory were discussed. These drafts were deemed interesting, and the Working Group will start working on these.

Finally the Group discussed the Charter. It was concluded that most of the goals from the original Charter have been achieved. An inventory was made amongst the members on whether they thought the Group should close down, or whether there were new items in the directory area that needed work. The inventory showed that there is certainly interest to continue a Directory Services group, but with a slightly shifted focus, towards solving operational mid-term problems in the areas of datamanagement, provision of integrated DUAs, Database coupling interfaces, security and legal issues.

The Working Group Chair and Area Directors will draft a new Charter.

It was noted that absence of any representative of the ongoing pilots on X.500 is very unfortunate.

X.400 Operations Working Group (X400OPS)

The Working Group started off with a new co-Chair, Tony Genovese, taking over from Rob Hagens. Twenty-Nine participants from eight countries attended the meeting. The Working Group discussed various Internet-Drafts:

- Operational requirements for X.400 Management domains in the GO-MHS Community. Minor comments; will be published as Informational RFC and RTR.

- Using the Internet DNS to maintain RFC1327 mapping tables and X.400 routing information. This will be split into two documents. Progressed to prototype early 1993.

- Routing coordination for X.400...... As usual lots of comments. Routing is always a hot issue :-). Will now be advanced early 1993 to prototype.

- Assertion of C=US; A=Internet lively discussions on this document. Lots of opposition especially from outside of the US. A special design team was formed on this issue Chaired by Kevin Jordan and Allan Cargille. The US-RAC name registration and behaviour guidelines were presented under this item.

- Mapping between X.400 and Mail-11. No more comments on the document. Will be submitted as prototype RFC.

- X.400 use of extended character sets. No comments, will be published as an Informational RFC.

- X.400 postmaster convention will be discussed via E-mail and then put on standards track.

Xwindows over OSI and Skinny Stack BOF (THINOIS)

The THINOSI BOF was the second BOF on this subject, with fourteen participants. The conclusions were to propose a working group with three objectives:

1. Promote the deployment and testing of X-windows over OSI implementations and their generalization to be carrier of any byte- stream over ACSE and the OSI 7-layer. (Simple byte-stream, not equivalent to full TCP function).

2. Develop an RFC that defines the skinny bits for the generalized byte stream carrier: The protocol that the OSI standards require, but respecified without regard to which standard requires it.

3. Develop an RFC of skinny bits for some subset of Directory Access Protocol.

Items two and especially three are feasibility proofs to see if such a document can be produced and be usefull. Implementation in parallel is anticipated.

CURRENT MEETING REPORT

Reported by Joshua Mindel/Open Networks

Minutes of the FTP-FTAM Gateway BOF (FTPFTAM)

The purpose of the BOF was to discuss the status of the Internet-Draft FTP-FTAM Gateway Specification and the usefulness of transitioning the document to a more permanent status. Joshua Mindel initiated the discussion with an overview of the document's purpose, evolution, and content.

The purpose of the document is to provide a formal specification for a gateway that maps, in real-time, between the FTP and FTAM protocols. The specification will be of use to both implementors and customers. As an IETF-endorsed standard, the specification will permit implementors to justify the cost of building FTP-FTAM gateways. Robert Cooney, representing both the Defense Information Systems Agency (DISA) and the Navy Open Systems Environment Testbed, expressed the dilemma currently faced by customers that need to procure FTP-FTAM gateways - there is no specification available to reference. In the electronic mail world there is Request for Comments 1148 (RFC1148), Mapping between X.400 and RFC 822, to serve as a reference for electronic mail gateways.

The FTP-FTAM Gateway specification is based on the original NIST FTP-FTAM Gateway design, published by M.A. Wallace et al in 1986, as well as on lessons learned by the authors in fielding the ISODE FTP-FTAM Gateway on the MILNET in 1989. The following relevant historical events were also noted:

- MITRE implemented FTP-FTAM Gateway based on NIST design (1987)
- MITRE implementation was incorporated into ISODE
- FTP-FTAM Gateway (ISODE 6.0) was fielded on MILNET under the DoD GOSIP Gateway Initiative (1989)

The FTP-FTAM Gateway specification contributes to the advancement of the FTP-FTAM Gateway concept in the following ways:

- Clarifies and enhances the FTP and FTAM mappings documented by NIST.
- Enhances the user interaction capability provided by the ISODE implementation of the FTP-FTAM Gateway.
- Provides guidelines for fielding FTP-FTAM Gateways on the Internet(s10hs12V).
- Serves as a formal specification for the FTP-FTAM Gateway suitable for implementors to use in building additional FTP-FTAM Gateways.
- Serves as a formal specification for organizations desiring to procure FTP-FTAM Gateways.

The Internet-Draft was reviewed by the IETF in July 1991. In addition, Digital Equipment Corporation has developed an FTP-FTAM Gateway using the draft specification. Based on this implementation, DEC provided valuable input to the specification.

All BOF attendees agreed that the document is mature, represents stable technology, and has undergone sufficient peer review. The discussion then turned to the IETF status of the specification. In time, the usefulness of the specification will be proven by the number and quality of implementations based on it. Daniel Fauvarque, of Sun-France, indicated that his organization may have an interest in implementing an FTP-FTAM Gateway, but at this time, there is no published standard. He indicated that advancement of the specification to Proposed Standard would provide sufficient status to permit Sun to use the specification at the appropriate time.

Dave Piscitello, the OSI Area Director in attendance, agreed to advance the specification by submitting a request to the IESG to consider the Internet-Draft FTP-FTAM Gateway Specification as a Proposed Standard.

Copies of the FTP-FTAM Gateway specification were distributed at the BOF.

Attendees

Daniel Fauvarque	dfauvarq@france.sun.com
Steve Hardcastle-Kille	s.kille@isode.com
Joshua Mindel	mindel@netwrx1.nw1.com
David Piscitello	dave@sabre.bellcore.com
Robert Slaski	slaski@netwrx1.nw1.com

CURRENT MEETING REPORT

Reported by Sheri Repucci/Merit

Minutes of the Shared Whois Project BOF (WHOIS)

A presentation of the current state of the Shared Whois Project was presented, all core processes of the project are complete with some work left on the automation process. The next stage of this project includes converting the X.500 IP schema to that currently being proposed by Glenn Mansfield, Thomas Johannsen, and Mark Knopper. Furthermore, the GSI-NIC has recently brought up Quipu and will be attempting to represent their data in the various schemas and eventually to load their entire database of information into X.500.

Future reports on this project will be presented to the IETF community via the OSIDS Working Group as there will not be a separate working group formed for the shared Whois service.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Lou Berger	lberger@bbn.com
Stefan Fassbender	stf@easi.net
Tim Howes	tim@umich.edu.
Mark Knopper	mak@merit.edu
Andrew Partan	asp@uunet.uu.net
Sheri Repucci	smr@merit.edu
Mark Smith	mcs@umich.edu
Marten Terpstra	marten@ripe.net
Yung-Chao Yu	yy@qsun.att.com

CURRENT MEETING REPORT

Reported by Peter Furniss/PFC

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

Following the previous BOF meeting, a draft working group Charter had been submitted to the IESG by Peter Furniss. However, the aim of the Charter had not been sufficiently clear, and the BOF met again to clarify what was wanted and what it was appropriate to do in the IETF arena, if anything.

Peter Furniss suggested that the (or just his) overall objective was to show that the OSI upper-layer protocols were, or could be, lightweight. The documents are certainly heavy, and the OSI model is liable to lead to implementations that are heavyweight. A fully general-purpose implementation will be large, but an implementation designed for a particular purpose need not be. This was the essence of the "skinny stack" approach, which could also be summarised as an implementation of the protocols but not of the OSI documents. In the skinny stack:

- The OSI layers are merged.
- Pre-coded octet sequences are used for sending, where possible.
- In received protocol, only the values needed are looked for.

Additional principles are that only protocols conformant to the OSI standards are sent, and **any** conformant protocol can be received. Consequently a skinny implementation can interwork with a 'full' (non- skinny) implementation **that is supporting the same application**. It is implicit in the skinny approach that there is some kind of specialisation.

The possibility of light-weight implementations had contributed to the choice of mapping to OSI specified for the X Windows System protocol in an EWOS [European Workshop on Open Systems] Technical Guide (ETG 13) and in the draft ANSI standard dpANS X3.196 part IV. These define use of the full 7-layers of OSI, sending the separately defined X byte stream (as would be sent over TCP) over Presentation, with connection establishment using ACSE (Association Control Service Element).

It was pointed out by Keith Sklower that it would be perfectly feasible to carry X directly on OSI Transport, without the addition of Session, Presentation and ACSE. Possibly some additional specification would be needed to provide the equivalent of TCP graceful close. From the following discussion:

- Possibly no work would be needed for graceful close (it can be treated as a local matter).

- The whole point of the skinny approach was that the cost of the additional layers was minimal.

- The additional layers made X into a "normal" OSI application - it could use whatever support facilities became available for such.

- The most appropriate mapping rather depended on the anticipated environment - there were those who wanted to use X in an all-OSI environment.

A pilot implementation of the EWOS ETG13 mapping, using skinny techniques, was available at the University of London Computer Centre. There were versions for different interfaces to OSI Transport service, not all available yet. Brien Wheeler had an independent implementation using the ISODE upper-layers.

Peter suggested there were two possible directions to take the skinny approach from X - "wider" and "higher". "Wider" would be to extend it to support other TCP-using application protocols - this could be just to other "simple byte stream" protocols, or to provide equivalence of all TCP features, or to specific (standardized) applications. "Higher" would be to include the skinny implementation of OSI protocols - Directory Access Protocol, ROSE, CMIP, Transaction Processing.

The BOF then considered what the worthwhile future activities for a working group in this area were. The possibilities were:

1. Promote the deployment of X/osi, including interworking experiments.
2. Extend the skinny stack as an alternative carrier for other TCP-using protocols.
3. Produce specifications of skinny stack for some OSI application protocols.

The questions for two were how far to take the extension, and what exactly, if anything, needed to be done within the IETF. Specification of a profile for "migrant applications" is being progressed in the OSI Implementors Workshop (OIW). The possibility of defining the use of the Berkeley socket API for access to skinny stack OSI was considered - this had been the basis of the previous draft Charter, which had met problems. It was perceived that what was needed was a re-specification of the OSI protocols in simpler terms - the definition of the "skinny bits", the octet sequences that must be sent and received to conform to the protocol specifications. The re-specification would not be concerned with which (OSI) document required the particular bits, but just what they were. This could be limited to the octet sequences required for X, but it would be a minimal addition to extend this for other simple byte-stream protocols. It would not be extended to cover the full equivalence to TCP, nor for specific standardised protocols. Most of the details of this have already been worked out in developing the ULCC/Furniss X/osi pilot. The specification would also be usable as the supporting layers for OSI application protocols that only use the kernel and duplex session functional units and a single presentation context (apart from that for ACSE) – however, for these some other component of the system will be handling the ASN.1 encoding/decoding of the application protocol.

The development of implementations using this specification and their deployment would be encouraged in the usual way. The existing X/osi implementations are essentially using this specification.

Regarding Item three, various candidate application protocols were discussed, but the obvious example was the Directory Access Protocol. Again a specification of the "skinny bits" would be the best way to facilitate implementation. This would be an effective test of the skinny approach – it might not be possible to produce a useful, concise specification, or an efficient and reasonably small implementation. The level of functionality would be a deciding factor - an increasing scale would be:

- Look up P-address given application-entity title.
- Look up O/R name.
- Provide equivalent function to LDAP (lightweight directory access protocol).
- Everything in DAP.

If a lightweight DAP implementation is possible it will have the virtue of being able to interwork with a standard DSA, without requiring intervening converters or special DSAs.

Peter Furniss agreed to produce a draft Charter on these lines. The development tasks would be the "skinny bits" for simple byte-stream applications and the "skinny bits" for DAP.

A mailing list for the Working Group has now been set up: thinosi@ulcc.ac.uk with thinosi-request@ulcc.ac.uk as the place to send requests to join.

Attendees

Richard Colella	colella@osi.ncsl.nist.gov
John Dale	jdale@cos.com
Richard desJardins	desjardi@boa.gsfc.nasa.gov
Peter Furniss	p.furniss@ulcc.ac.uk
Steve Hardcastle-Kille	s.kille@isode.com
Susan Hares	skh@merit.edu
Triet Lu	triet@cseic.saic.com
David Piscitello	dave@sabre.bellcore.com
James Quigley	jim_quigley%Y0@hp6600.desk.hp.com
Keith Sklower	sklower@cs.berkeley.edu
Brien Wheeler	blw@mitre.org
Cathy Wittbrodt	cjw@nersc.gov

2.4.1 MHS-DS (mhsds)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

Mar 1992 Define a set of service requirements for MHS use of Directory. This should include: support for routing; support for security services; support for user agent capabilities; support for distribution lists. The extent to which existing standards can meet these requirements.

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

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

Internet-Drafts:

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

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

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

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

"Use of the Directory to support mapping between X.400 and RFC 822 Addresses", 04/09/1992, S. Hardcastle-Kille

<draft-ietf-mhsds-supmapping-02.txt, .ps>

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

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

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

CURRENT MEETING REPORT

Reported by Brien Wheeler/MITRE

Minutes of the MHS-DS Working Group (MHSDS)

Agenda

- Minutes of Boston Meeting
- Action Item Review
- Revision of Charter
- MHS-DS Pilot Project Planning
  - Project Coordinator
  - Project Participants
  - Software
- Document Review
- Adjournment

Minutes of Boston Meeting

The Minutes of the Boston meeting were approved as written.

Action Item Review

- Update of Charter (Kevin Jordan) - Done.
- Produce MHS-DS Overview Document (Kevin Jordan) - Not done.
- Revision of MHS-DS Documents (Steve Hardcastle-Kille) - Done.
- Produce Document on Other MHS-DS Issues (Jim Romaguera) - Not done.
- Production of Pseudocode for Routing Document (Harald Tveit Alvestrand) - Not done.

Revision of Charter

Kevin's revised Charter was discussed and approved.

Harald noted that some people have exhibited confusion as to whether or not user support is within the scope of MHS-DS. His opinion was that we should explicitly state in the Charter that that is not MHS-DS' function. Steve felt that there may be cases, particularly in the upcoming pilot, where MHS-DS may need to provide some user support.

MHS-DS Pilot Project Planning

Kevin stated that to make this pilot successful, we need people who have the resources and motivation to take key coordination roles. A list for indicating interest in participating in the pilot was circulated. Steve indicated that the PARADISE project has a deliverable in piloting the use of directory, so that they may be prepared to manage the MHS-DS pilot. David Goodman would be a good point of contact for pursuing this. Kevin pointed out that initial coordination of the pilot may not demand much time and effort, and that it can probably be accomplished via email, so he volunteered to be the coordinator.

Others who volunteered to participate in the pilot project include:

- Harald Tveit Alvestrand
- Allan Cargille
- John Dale
- Arlene Getchell
- John Hawthorne
- Kevin Jordan
- Sylvain Langlois
- Mary LaRoche
- Steve Hardcastle-Kille
- Jim Romaguera
- Panos Tsigaridas
- Karen Petraska-Veum
- Peter Yee
- Yung Yu

The question was raised as to whether Quipu 8.0 supports all the new attribute syntaxes defined within the MHS-DS attributes. Additionally, will the final public release of PP use these new attributes and algorithms?

Steve informed the Group that the situation is "unfortunate but a little complex." ISODE 8.0 was distributed in July, with a schema aligned to the then-current Internet-Drafts. This may cause certain problems in pilot use, but most likely only in complicated messaging situations. An alpha release of PP was made in August that was aligned to the July version of the MHS-DS documents. A beta release also aligned to the July documents will be both integrated into the ISODE Consortium beta release, and be released by X-Tel as a test under the auspices of JNT. It is expected that this release will evolve into the final public distribution of PP. The ISODE Consortium will be focusing on other aspects of the system, while X-Tel will concentrate on the use of directory aspects. X-Tel expects to make this release by the end of the year.

Harald raised the question of the relationship between local tables and directory information. Steve replied that there is a minor problem in the PP implementation in that there is not a strict one-to-one relationship between information in the directory and table information, so that people have to configure PP differently if it is using the directory.

Harald raised the question of the precedence of local routing information and directory information. Steve replied that currently no intelligence exists within the software to facilitate this, but reminded everyone that this software is of recent origin, and it will be evolving.

The question of putting the COSINE tables into the directory arose. This led to a discussion of automated tools for directory configuration. The conclusion is that at this time, no one can commit to developing these tools now, but some sites may be able to fund development projects later. In particular, the German research network, DFN, is planning to fund such development beginning early in 1993 (contact Panos Tsigaridas). Meanwhile, MHS-DS is encouraging sites to put their own configuration information into the directory. Specific participants will be responsible for adding the necessary infrastructure entries (i.e., ADMD and PRMD entries) to the directory in order to allow other participants to add their routing and configuration information independently.

After lunch, Kevin gave a presentation on "LONG BUD," the new MHS-DS pilot project (Lightweight Open Naming with Global Bearing Using Directory). LONG BUD has the following goals:

- Wider connectivity and greater direct connectivity between pilot MTAs.
- Provision of address mapping information as well as routing information.
- Reduction in complexity of participating in X.400 pilots.
- Support for X.400 communities without X.500 access (i.e., ADMDs and the COSINE MHS WEPS). This will require the development of some tools.
- Elimination of the need to distribute routing tables manually among pilot participants

LONG BUD will be international in scope, with participants already identified in DE, FR, NO, UK, and US. Other possible near-term participants include CH, IT, and FI. Further resources that may be able to help advance this pilot include the DFN (GMB) MHS-DS project, the ISODE Consortium, and the JNT UK project.

This pilot will run at least through the July 1993 IETF, with directory population of participant information being a defined goal for the March 1993 IETF meeting. Success can be declared when two main goals are met: tables are no longer used by directory-capable MTAs, and the directory is the principal source for routing information.

Certain pieces of functionality have been identified as critical to the success of LONG BUD: programs to extract routing information from the directory for table-based MTA's, an email interface to the directory (DFN) for end-users, and the ISODE Quipu and directory-capable PP software.

The distribution list for LONG BUD will be the MHS-DS discussion list, mhs-ds@mercury.udev.cdc.com (anonymous ftp archive server quixote.css.cdc.com), and Kevin Jordan (kej@mercury.udev.cdc.com) will be the Pilot Coordinator.

Action Item	Jim Romaguera will populate the Swiss portion of the DIT with ADMD names.

Action Item	Kevin Jordan will contact Wengyik Yeong about populating ADMDs under c=US.
Action Item	Harald Tveit Alvestrand will contact the appropriate person about populating the ADMDs under c=NO.
Action Item	Sylvain Langlois will contact the appropriate person about populating the ADMDs under c=FR.
Action Item	Panos Tsigaridas will contact the appropriate person about populating the ADMDs under c=DE.

These people will be points of contact for management of the PRMD namespace within their respective countries. Peter Yee will be the point of contact for registering PRMDs under ADMD=TELEMAIL, c=US.

Document Review

- Representing Tables and Subtrees in the Directory

  This document will be progressed with minor editorial changes as an Experimental Protocol.

- Representing the O/R Address Hierarchy in the Directory Information Tree.

  This document spurred a discussion of the problem of alias proliferation due to alternate names and underspecified O/R addresses. It was agreed that this issues bears on the routing document, and not the mapping document. This document will be progressed with minor editorial changes as an Experimental Protocol.

- Use of the Directory to Support Mapping Between X.400 and RFC 822 Addresses.

  This document will be progressed with minor editorial changes to a Experimental Protocol.

- A Simple Profile of MHS Use of the Directory.

  Although no changes were made to this document, its progression will be postponed so that it may be published with the routing document it profiles.

- MHS Use of Directory to Support MHS Routing.

  This document went through some technical revision at the meeting, and hence will remain an Internet-Draft until at least the March 1993 MHS-DS meeting.

Action Item: Kevin Jordan will send mail to Erik Huizer when the indicated documents are ready for progression.

Adjournment

Next meeting: The Spring 1993 IETF will be held in Columbus, Ohio, March 28th - April 2nd.

Attendees

Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
George Chang	gkc@ctt.bellcore.com
John Dale	jdale@cos.com
Daniel Fauvarque	dfauvarq@france.sun.com
Raphael Freiwirth	5242391@mcimail.com
Ella Gardner	epg@gateway.mitre.org
Tony Genovese	genovese@es.net
Arlene Getchell	getchell@es.net
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	kej@udev.cdc.com
Marko Kaittola	marko.kaittola@funet.fi
Mary La Roche	maryl@cos.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
John Myers	jgm+@cmu.edu
Karen Petraska-Veum	karen@nsisrv.gsfc.nasa.gov
Jim Romaguera	romaguera@cosine-mhs.switch.ch
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Brien Wheeler	blw@mitre.org
Peter Yee	yee@atlas.arc.nasa.gov

2.4.2 MIME-MHS Interworking (mimemhs)

Charter

Chair(s):

Steve Thompson, sjt@gateway.ssw.com

Mailing Lists:

General Discussion: mime-mhs@surfnet.nl
To Subscribe: mime-mhs-request@surfnet.nl
Archive:

Description of Working Group:

MIME, (Multipurpose Internet Mail Extensions) is currently a Proposed Standard. MIME redefines the format of message bodies to allow multi-part textual and non-textual message bodies to be represented and exchanged without loss of information. With the introduction of MIME as a Proposed Standard it is now possible to define mappings between RFC-822 content-types and X.400 body parts. The MIME-MHS Interworking Working Group is chartered to develop these mappings, providing an emphasis on both interworking between Internet and MHS mail environments and also on tunneling through these environments. These mappings will be made in the context of an RFC-1148bis environment.

Goals and Milestones:

Done Post an Internet-Draft describing MIME-MHS Interworking.

Done Post an Internet-Draft describing the "core" set of Registered conversions for bodyparts.

Jul 1992 Submit a completed document to the IESG describing MIME-MHS Interworking as a Proposed Standard.

Jul 1992 Submit the "core" bodyparts document to the IESG as a Proposed Standard.

Internet-Drafts:

"Mapping between X.400 and RFC-822 Message Bodies", 07/01/1992, H. Alvestrand, S. Hardcastle-Kille, R. Miles, M. Rose, S. Thompson <draft-ietf-mimemhs-mapping-01.txt>

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

"HARPOON: Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages", 09/28/1992, H. Alvestrand, J. Romaguera, K. Jordan <draft-ietf-mimemhs-harpoon-00.txt>

2.4.3 Network OSI Operations (noop)

Charter

Chair(s):

Susan Hares, skh@merit.edu
Cathy Wittbrodt, cjw@barrnet.net

Mailing Lists:

General Discussion: noop@merit.edu
To Subscribe: noop-request@merit.edu
Archive: merit.edu:pub/noop-archive

Description of Working Group:

The Working Group is chartered to work on issues related to the deployment of CLNP in the Internet. The first area of this Group's work has been the learning necessary to start deploying OSI in internet networks. This phase includes planning for OSI deployment by creating routing plans for regional networks and education on using OSI routing protocols.

This first area of the Group's work will be on-going as we continue to deploy OSI in the Internet. This step has lead to people deploying OSI for Pilot projects and demonstrations of OSI.

The second step of deploying OSI will be the transition of OSI from a pilot service to a production service. During this phase we will work on specifying the network debugging tools and test beds. We will need to track the level of OSI support in the Internet. We will need to provide documentation for new users of OSI on the Internet.

Goals and Milestones:

Ongoing	Provide a forum to discuss OSI routing plans by email or in group discussions.
Jan 1992	Post as an Internet-Draft, a tutorial for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.
Apr 1992	Post as an Internet-Draft, a requirements document specifying what OSI network tools are needed on every host and router.
Jul 1992	Post as an Internet-Draft, a collection of regional Routing and Addressing plans.
Done	Post as an Internet-Draft, a list of OSI Network Utilities available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.
Jul 1992	Post as an Internet-Draft, a description of OSI network layer debugging methods.

Done Post as an Internet-Draft, a list of OSI Network Layer NOC tools available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.

Jul 1992 Submit to the IESG for Proposed Standard, a requirements document specifying what network tools are needed on every OSI host and router.

Aug 1992 Submit to the IESG as an Informational RFC, a description of OSI network layer debugging methods.

Internet-Drafts:

"An Echo Function for ISO 8473", 11/10/1992, R. Hagens <draft-ietf-noop-echo-00.txt>

"Tools RFC", 11/10/1992, S. Hares, C. Wittbrodt <draft-ietf-noop-tools-00.txt>

CURRENT MEETING REPORT

Reported by Sue Hares/Merit

Minutes of the Network OSI Operations Working Group (NOOP)

Agenda

- RFC 1139
- Tools Document
- OSI Survey
- Technical Aspects of TUBA
- Tutorial on OSI
- Discussion of OSI Routing Plans

RFC 1139

The Group talked about the revision of RFC1139 and also the Tools Internet-Draft. Both of these need some revision. Both need some specific text before it should be put in the Internet-Draft, and in the ISO Documents (Richard Collela).

Tools RFC

Sue Hares wanted to move TOOLS RFC to a Proposed Standard. The Group had no problems over most of the section. Bill Manning suggested we try MIB as IS. Dino Farinacci from cisco wanting to see a routing MIB replace the current text for IS-IS, IDRP. The Area Director (Dave Piscitello) will help formulate a routing MIB working group since there seems to be interest. Dino will try to find someone from cisco to work on it. It was also indicated that some help from DEC would be useful.

The Tools RFC is going to have the MIB information removed until there is a routing table MIB. Then the document will be modified to point to the routing table MIB. Only sections 5.3.1.1 and 5.3.1.2 will be cut from document.

No one was opposed to the CMIP and GDMO sections being left in the document. The GDMO for the OSI specific management will be gone over by GDMO experts. Ross Callon said he could get Sue Hares in touch with a GDMO expert.

Second NOOP Session

The second session of NOOP was a tutorial for folks a little less familiar with OSI and deployment issues. After the tutorial we discussed a particular network's topology and how it might be broken up into areas and domains.

2.4.4 OSI Directory Services (osids)

Charter

Chair(s):

Steve Hardcastle-Kille, s.kille@isode.com

Mailing Lists:

General Discussion: ietf-osi-ds@cs.ucl.ac.uk
To Subscribe: ietf-osi-ds-request@cs.ucl.ac.uk
Archive:

Description of Working Group:

The OSI-DS Group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this Group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

Ongoing Maintain a Schema for the OSI Directory on the Internet.

Ongoing Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

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

Done Study the relationship of the OSI Directory to the Domain Name Service.

Internet-Drafts:

"Using the OSI Directory to Achieve User Friendly Naming", 11/26/1990, S. Kille <draft-ietf-osids-friendlynaming-04.txt, .ps>

"Naming Guidelines for Directory Pilots", 03/21/1991, P. Barker, S.E. Hardcastle-Kille <draft-ietf-osids-dirpilots-05.txt, .ps>

"A String Representation of Distinguished Names", 01/30/1992, S. E. Hardcastle-Kille <draft-ietf-osids-distnames-04.txt, .ps>

"Lightweight Directory Access Protocol", 04/17/1992, Wengyik Yeong, Tim Howes, Steve Hardcastle-Kille <draft-ietf-osids-lightdirect-02.txt>

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

"DUA Metrics", 09/23/1992, Paul Barker <draft-ietf-osids-dua-metrics-00.txt>

"DSA Metrics", 09/23/1992, P. Barker, S. Hardcastle-Kille <draft-ietf-osids-dsa-metrics-00.txt>

"A strategic Plan for deploying an Internet Directory Service", 10/14/1992, S. Hardcastle-Kille, E. Huizer, V. Cerf, R. Hobby, S. Kent, J. Postel <draft-ietf-osids-plan-directory-00.txt, .ps>

Request For Comments:

RFC 1275 "Replication Requirements to provide an Internet Directory using X.500"

RFC 1276 "Replication and Distributed Operations extensions to provide an Internet Directory using X.500"

RFC 1277 "Encoding Network Addresses to Support Operation Over Non-OSI Lower Layers"

RFC 1278 "A String Encoding of Presentation Address"

RFC 1279 "X.500 and Domains"

CURRENT MEETING REPORT

Minutes of the OSI Directory Services Working Group (OSIDS)

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

Attendees

Ed Albrigo	ealbrigo@cos.com
Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
Jules Aronson	aronson@nlm.nih.gov
George Chang	gkc@ctt.bellcore.com
James Conklin	jbc@bitnic.educom.edu
John Dale	jdale@cos.com
Letha Dugas	4371362@mcimail.com
William Edison	
Daniel Fauvarque	dfauvarq@france.sun.com
Catherine Foulston	cathyf@rice.edu
Ned Freed	ned@innosoft.com
Peter Furniss	p.furniss@ulcc.ac.uk
Ella Gardner	epg@gateway.mitre.org
Tony Genovese	genovese@es.net
Arlene Getchell	getchell@es.net
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Marco Hernandez	marco@mh-slip.educom.edu
Tim Howes	tim@umich.edu.
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	kej@udev.cdc.com
Marko Kaittola	marko.kaittola@funet.fi
Mark Knopper	mak@merit.edu
Mark Kosters	markk@nic.ddn.mil
John Kunze	jak@violet.berkeley.edu
Mary La Roche	maryl@cos.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Edward Levinson	levinson@pica.army.mil
John Myers	jgm+@cmu.edu
Chris Newman	chrisn+@cmu.edu
Rakesh Patel	patel@noc.rutgers.edu
Karen Petraska-Veum	karen@nsisrv.gsfc.nasa.gov
Sheri Repucci	smr@merit.edu
Jim Romaguera	romaguera@cosine-mhs.switch.ch

Marshall Rose	mrose@dbc.mtview.ca.us
Alan Roszkiewicz	alan@sprint.com
Srinivas Sataluri	sri@qsun.att.com
Richard Schmalgemeier	rgs@merit.edu
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgras@bitnic.educom.edu
Simon Spero	simon_spero@unc.edu
Catherine Summers	cts@cos.com
Fumio Teraoka	tera@csl.sony.co.jp
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Chris Weider	clw@merit.edu
Brien Wheeler	blw@mitre.org
Russ Wright	wright@lbl.gov
Peter Yee	yee@atlas.arc.nasa.gov
Yung-Chao Yu	yy@qsun.att.com

2.4.5 Office Document Architecture (oda)

Charter

Chair(s):

Peter Kirstein, P.Kirstein@cs.ucl.ac.uk

Mailing Lists:

General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk
Archive:

Description of Working Group:

The ODA Working Group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g., spreadsheets. The Working Group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liaison with the SMTP and X.400 Working Groups.

This Working Group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.
Done	Inaugural meeting.
Done	Produce a paper stating what ODA standards or profiles still need completing.
Done	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Jun 1992	Provide first feedback on the ODA Pilot.

2.4.6 SNMP over a Multi-protocol Internet (mpsnmp)

Charter

Chair(s):

Theodore Brunner, tob@thumper.bellcore.com

Mailing Lists:

General Discussion: snmp-foo@thumper.bellcore.com
To Subscribe: snmp-foo-request@thumper.bellcore.com
Archive: thumper.bellcore.com:pub/snmp-foo/archive

Description of Working Group:

Within the SNMP management framework, the philosophy is to place the burden of management processing on managers, not on agents. As the Internet evolves to accommodate multiple protocol suites, there may be SNMP agents in the Internet that do not support the recommended method of exchanging SNMP messages using UDP/IP. In these instances, the proper model for managing a multiprotocol internet should be that agents must only be required to support one method of exchanging SNMP messages (i.e., encapsulation of SNMP messages in *one* of the protocol suites of the multi-protocol internet), and the managers support as many encapsulation methods as needed (potentially, all) to communicate with all resources it manages.

The SNMP over a Multi-protocol Internet Working Group is chartered to identify and provide solutions for communication between SNMP agents and managers in those configurations where the recommended method of exchanging SNMP messages using UDP/IP cannot be used; i.e., where a managed resource supports a single protocol suite that protocol is not UDP/IP but another protocol suite of the multi-protocol internet (for example, OSI, AppleTalk, or XNS/IPX).

Questions to be considered include: What are the appropriate protocol suites to consider? What is the appropriate method of encapsulating SNMP? What are the addressing considerations for SNMP messages What new MIB Modules are required? What (positive) effect can SNMP-based management have on resource-sharing among multiple protocols?

Goals and Milestones:

Done	Post an Internet-Draft describing operation of SNMP over OSI.
Done	Post an Internet-Draft describing operation of SNMP over IPX.
Done	Post an Internet-Draft describing operation of SNMP over Appletalk.

Done Submit a document describing the operation of SNMP over OSI as a Proposed Standard.

Done Submit a document describing the operation of SNMP over IPX as a Proposed Standard.

Done Submit a document describing the operation of SNMP over Appletalk as a Proposed Standard.

Internet-Drafts:

"SNMP over AppleTalk", 12/23/1991, G. Minshall, M. Ritter <draft-ietf-mpsnmp-appletalk-02.txt>

"SNMP over OSI", 04/10/1992, Marshall Rose <draft-ietf-mpsnmp-overosi-03.txt>

"SNMP over IPX", 06/23/1992, Steve Bostock <draft-ietf-mpsnmp-overipx-01.txt>

2.4.7 X.400 Operations (x400ops)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: ietf-osi-x400ops@pilot.cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@pilot.cs.wisc.edu
Archive:

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this Group is to insure interoperability between Internet X.400 management domains and the existing Internet mail service. The specific task of this Group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

Done	Initial meeting, produce internal outline.
Done	Working draft, circulate to interested people.
Jul 1991	Internet-Draft available.
Dec 1991	Document ready for publication.

Internet-Drafts:

"Routing coordination for X.400 MHS services within a multi protocol / multi network environment", 03/03/1992, U. Eppenberger <draft-ietf-x400ops-mhs-service-03.txt>

"Mapping between X.400(1984/1988) and Mail-11 (DECnet mail)", 03/03/1992, Claudio Allocchio <draft-ietf-x400ops-mapsmail-02.txt>

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

"X.400 use of extended character sets", 06/18/1992, Harald Alvestrand <draft-ietf-x400ops-charactersets-01.txt>

"Postmaster Convention for X.400 Operations", 11/23/1992, C. A. Cargille <draft-ietf-x400ops-postmaster-00.txt>

"Assertion of C=US; A=<tbd>", 12/11/1992, E. Stefferud <draft-ietf-x400ops-admd-00.txt>

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

CURRENT MEETING REPORT

Reported by Alf Hansen/SINTEF

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

Welcome and Administration

Harald Tveit Alvestrand volunteered to take Minutes. Tony Genovese, ESNet, has volunteered to be the new co-Chair of the X400OPS Group now that Rob Hagens is no longer able to participate.

The Agenda was accepted, with the provision that some items under point three should be shuffled around a bit. They are presented below in the order in which they were discussed.

Action List from Boston

Most of the actions were marked "done" in the Agenda, and received no comment. Those that were discussed were:

- Tony Genovese - continue to work on a WEP which is accessible over public X.25.

  Tony reports: This item was tied in with ADMD connectivity and Corporation for Open Systems (COS), and was discussed at a meeting the evening before. Most probable is that an ADMD offers the connectivity, after ISOC based negotiations.

  The issue of an X.25 WEP in the US part of the GO-MHS community is still open. Nobody is willing to do it without further study of the financial implications and funding arrangements; UNINETT will continue to provide X.25 to TCP relay for the time being.

- Harald Alvestrand - update document on extended character sets and release as an Internet-Draft. [done on Thursday before meeting]

- Allan Cargille - write draft document about postmaster addresses and release as an Internet-Draft. [done - distributed at the meeting]

- Claudio Allochio - produce new document explaining how the X.400 DNS tables should be used and distribute to X400OPS list. Claudio promises that it will be ready for the next meeting.

Erik Huizer distributed copies of Rare Technical Report (RTR) versions of 1327 and 1328, and explained some of the procedures governing the RTR/RFC relationship. Marko Kaittola presented his ideas on mapping table updates using E-mail to carry them. Ideas presented, generally positive reception.

Review of Documents

Editor's Note (md): A detailed account of the Document reviews is available via ftp under x400ops-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

Liaison with Other Bodies

U.S. MHS MD Subcommittee: Ella Gardner of MITRE, Chair, MHS MD Subcommittee presented the output of MHS MD, which has completed the following documents for the U.S.:

- Behavior Guidelines for Voluntary Participation within the US National X.400 MTS.
- Registration Procedures for the United States Joint Registration Authority (US-JRA).
- Operating Guidelines for Registrars of MHS Management Domain Names Used within the US.

Ella Gardner gave an orientation about the status of naming authorities in the US.

A call asking for someone to operate the US national ADMD/PRMD register has gone out. Submissions are expected before December 1st with the hope that it will be operational in the First Quarter of 1993.

Any Other Business and Plan for Next Meeting

Next meeting: The Spring 1993 IETF will be held in Columbus, Ohio, March 28th - April 2nd.

Summary of Actions and Decisions

Claudio Allocchio	Produce new document explaining how the X.400 DNS tables should be used and distribute to X400OPS list.
	Will organize a pilot to get operational experience from the use of OPS-2.
	Part 1 (Mapping) of the document OPS-2 will be ready before January 1st; deadline for comments is January 15th; submission as an Experimental Protocol is expected by January 30th.
	Timescales for Part 2 (Routing) of OPS-2, now called OPS-2b, depend on the timescale for the routing document; expected date

	of an updated version is 2 weeks after the final version of the routing document (OPS-3).
	OPS-4 will be submitted as an Experimental Protocol within three weeks from now.
A. Cargille	Write the GO-MHS Community document together. [J. Romaguera]
	OPS-9 to be resolved by December 11th on the list; it will either be accepted, or the reference will be dropped from OPS-1.
Jim Romaguera	Compile a list of outstanding issues regarding OPS-3 and post it to the list. The list should be closed on December 4th; the final decisions should be taken by December 18th.
	The ADMD evaluation (OPS-8) should be refined and published as an Informational RFC (and RTR). Target date: December 15th.
Alf Hansen	Make the editorial changes, publish OPS-1 again as an Internet-Draft, and ask to have it published as an Informational RFC. Time limit: December 15th for the revised version, January 1st is the closing date for any comments before RFC submission. Will also be published as an RTR.
Harald Alvestrand	OPS-5 will be submitted for Informational RFC - Last Call will be sent to the Working Group on December 1st, with closing date December 14th. (The dates will be adjusted after consulting with the RARE TC and WG-MSG).
Urs Eppenberger	Add the text to paragraph 4.5 in OPS-3: "Secondary WEPs may require a testing period".
	A version number for the format will be added to all the documents defined in OPS-3.
	The START date should be mandatory, not optional in OPS-3.
List	OPS-6 to be discussed on the list; will be moved forward as an Experimental Protocol before December 18th.
Design Team	There will be formed a Design Team to define the necessary documents for U.S. operation of the Internet/GO-MHS community. This Design Team will report back to the main group when its conclusions are ready for review.

Attendees

Claudio Allocchio	Claudio.Allocchio@elettra.trieste.it
Harald Alvestrand	Harald.Alvestrand@delab.sintef.no
C. Allan Cargille	cargille@cs.wisc.edu
George Chang	gkc@ctt.bellcore.com
John Dale	jdale@cos.com
Daniel Fauvarque	dfauvarq@france.sun.com
Ned Freed	ned@innosoft.com
Ella Gardner	epg@gateway.mitre.org
Tony Genovese	genovese@es.net
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@isode.com
John Hawthorne	johnh@tigger.rl.af.mil
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	kej@udev.cdc.com
Marko Kaittola	marko.kaittola@funet.fi
Mary La Roche	maryl@cos.com
Sylvain Langlois	Sylvain.Langlois@der.edf.fr
Edward Levinson	levinson@pica.army.mil
Triet Lu	triet@cseic.saic.com
Bob Lynch	lynch@dsteg.dec.com
Karen Petraska-Veum	karen@nsisrv.gsfc.nasa.gov
Jim Romaguera	romaguera@cosine-mhs.switch.ch
John Sherburne	john.sherburne@sprintintl.sprint.com
Einar Stefferud	stef@nma.com
Panos-Gavriil Tsigaridas	Tsigaridas@fokus.berlin.gmd.dbp.de
Brien Wheeler	blw@mitre.org
Russ Wright	wright@lbl.gov

2.5 Operational Requirements Area

Director(s):

- Phill Gross: pgross@nis.ans.net
- Bernhard Stockman: boss@ebone.net

Area Summary reported by Bernhard Stockman/SUNET

BGP Deployment and Application Working Group (BGPDEPL)

The BGPDEPL Working Group met for one session during this IETF chaired by Matt Mathis.

Of the approximately 5000 networks which are currently reachable, almost 3000 are being announced with EGP2. This situation is pretty bleak. There is only a short time available to test and deploy BGP-4, and operators who have not yet deployed BGP-3, will face additional difficulties phasing out EGP-2. It is desirable for all network operators to have BGP experience as soon as possible.

Proposed deployment schedule for BGP-4/CIDR:

1/93-3/93	BGP-4 interoperability testing with the ANS testing facility (See below - all vendors are invited to participate).
3/93	BGP-4 capable code deployed in the production NSFnet backbone. Begin testing by propagating some test CIDR networks through the backbone.
6/93	Start aggregating **production** networks in the backbone.
12/93	Completely phase out all EGP-2 on NSFnet DMZs.

There was some discussion about how this interacts with the new network assignment rules and schedule specified in RFC1366 and RFC1367. The first aggregation of production networks (scheduled for June 1993) will be a flag day for any site requiring full routing tables and not running BGP-4.

Jordan Becker (ANS) estimated that full route aggregation will reduce the current routing tables by about 30%, because the old address assignment policies tended to allocate addresses in blocks anyhow.

cisco's next scheduled code freeze is February 1993, so even if bug-free BGP-4 code exists today, the earliest it will appear in General Availability products is October 1993. All customers who need BGP-4 before then must run pre-GA code.

Benchmarking Methodology Working Group (BMWG)

The BMWG Group met during one session chaired by Scott Bradner. The Group discussed various test frame formats to be used in conjunction with earlier described network device testing methods as described in RFC 1242.

Some router vendors have announced inadequate performance metrics with no consistent way defined for measuring of router performance.

In some network devices, packet forwarding has priority above other functions which could result in loss of learning tree for bridges and loss of routing information for routers when the device is loaded.

The Working Group discussed Performance impacts of filter lists. Various sizes of filter lists have been tested. Some vendors use hash-search where there is no significant difference in performance between various sizes of filter lists. When linear search is used the amount of list entries is proportional to the performance impact.

Finally the Working Group discussed the performance impact of network management. It was noted that some network products do not update the SNMP MIB database as often as the hardware updates its counters. There may thus be a discrepancy between what actually is going on and how this is reflected in the MIB database.

Network Status Report (NETSTAT) and Network Joint Management Working Group (NJM)

- **Mark Knopper, Merit, Jordan Becker, ANS - NSFnet:** Transition T1 -> T3 ongoing. In October, 18.9 billion packets carried on T3 while T1 steadily decreasing Number of nets is 7,354 whereof 2,566 is foreign networks. OSI traffic 600,000 - 1,000,000 packets per month during March to October 1992, August and September close to zero though. T3 not yet ready to forward native CLNP which will be carried encapsulated in IP. Of NSFnet/ANSnet configured networks nearly six thousands are actively announced. Around 90 percent of the networks are using T3 as primary.
  - The T3 backbone implementation. Dummy AS support for load splitting. Up to 5 high speed interfaces per router with 20 kpps in and out per interface and total of 50 kpps per router. Max performance is 22 Mbps in each direction at 270 byte packetsize. One way router hop delay = 0.165 msec which gives cross country router delay (8 hops) of 1.35 msec and a total cross country delay of 35 msec. A ping version using NTP for microsecond resolution is used. The dismantling of T1 backbone lines starts 12 Feb.
  - T3 Network Status. Announced corrections in peer behavior. Engineering changes in internal routing to minimize delay through T3 net. Map with delay numbers will be available on-line. Deployment underway of encapsulated CLNP across T3, to enable decommissioning of T1 very soon. Announced deployment of BGP4 in spring. Invited vendors and operators to use ANS testnet.

Support for CIDR is planned to start January 93.

- DECNET IV support
  - Multiprotol routers
  - Elimination of upgrade of tail circuits
  - Multinet DECNET in TCP encapsulation support
  - Throughput on T1 over 200 Kbps
  - Gradual transition (if any) to DECNET V
  - Native DECNET not considered due to severe loss of performance depending on DECNET resend algorithm.

- **Milo Medin - NASA Science Internet:** Awaiting US Department of Commerce clearance for connection to Russia. NASA portion of DoE/NASA ATM will initially include Langley, Lewis, Goddard, Ames, JPL.

- **Bernhard Stockman - EBONE:** Deploying security access scheme in EBONE routers - combination of kerberos and TACACS. Plan for link from Stockholm to Bonn.

- **Bob Collet - Sprint:** Sprint operates three logically distinct IP networks - domestic US, Atlantic-Europe-Mideast, and Pacific. Exclusively cisco routers showed new maps with new perspectives.

- **Rich Fisher - GSFC:** Satellite data collection and redistribution to distant research and processing centers.

- **Tony Hain - ESnet:** ATM project sites Livermore, LBL, LANL, Fermi, Oak Ridge, SuperCollider All local loops will be fiber.

- **Mark Knopper - ERNET:** Networking in India is funded by the Indian Government and the United Nations plans to "upgrade" to VSAT connections domestically to overcome shortcomings in domestic infrastructure.

Operational Requirement Area Directorate (ORAD)

The Group discussed requirements of ORAD and its members. ORAD is expected to guide other working groups and review documents with special attention to operational needs. Current Operations Area working group Chairs could be part of ORAD, but this is not implicitly required. To make ORAD have broad coverage it will be necessary to invite operators who have not traditionally participated in IETF meetings.

The meeting concluded that ORAD should not start off too big but initially concentrate on document review and presentation of issues to working groups.

Finally, the Group discussed various operational aspects of the ongoing audio and video multicast from IETFs. MBONE routers shall be positioned as high as possible in the topology. An ORAD operations recommendation was discussed. A variety of actions to improve the current MBONE implementation were identified. Tests shall happen before IETFs, which include announcements of tunneling and requests to be made further in advance of conferences, and a strict cut-off date after which there will be no more tunnels.

Operational Statistics Working Group (OPSTAT)

Before this meeting the Internet-Draft on a model for operational statistics had been submitted as an Informational RFC. This time the Working Group restarted the work on the client/server based protocol for retrieval of statistical data. Most of the simple commands were kept as is while the more complex parts were significantly modified. Some discussion centered around where the selection processing should be done. For example, should the conditionals be processed on the server or client? Great economies could be realized by processing the conditional on the server versus downloading all data to the client and processing it there. Some discussion revolved around the SQL-ness of the select command. There was a consensus not to make it more complex than it already is. As the storage format in the above mentioned RFC has changed since the client/server specification was initially drafted it was necessary to change some part of the client/server command language to reflect this. Finally the Goals and Milestones section of the OPSTAT Charter was reviewed and updated.

User Connectivity Problems Working Group (UCP)

The Group had previously defined a data structure that would enable Trouble Ticket hand-offs between NOCs. Paul Zawada had written an ASN.1-like description of the fields in this data structure.

Kaj Tesink drafted a document describing how some hand-off fields could be represented in electronic mail messages. The Group discussed this and agreed that the document needs to be revised to reflect more of the previously-defined hand-off fields. The goal is to allow trouble tickets to be mailed between NOCs both with and without internal trouble ticket systems. The format should be simple enough to enable humans to enter the data and yet regular enough to permit parsing. Paul and Kaj will work on this and get it out as an Internet-Draft. At that time, several groups agreed to experiment with the exchange format and to create a template to facilitate manual participation.

The UCP Internet-Draft on a Trouble Ticket Tracking System, originally written by Matt Mathis, had been discussed and revised heavily by the Group and it has now expired. Dan Long has volunteered to draft a new version which reflects the current consensus of the Group. This will also be published as an Internet-Draft.

The Group also discussed the current status of various publicly available internal Trouble Ticket systems.

CURRENT MEETING REPORT

Reported by Gene Hastings/PSC

Minutes of the Operational Area Directorate (ORAD)

Recruiting of ORAD Members

What is expected of recruits?

- Provide guidance as to what needs attention (what work groups need to be formed?) (Example: mbone coordination)

- Provide guidance to working groups in other areas, e.g. BGP Deployment and IPv7. For example Network Management and SAAG explicitly assign people to working groups.

- Document review. Particularly early on, i.e., Internet-Draftss, etc. As things are going in the POISED Working Group, it looks like the direction is for the IAB to delegate more of its activity and responsibility to the IESG which will increase the need for area advisory groups like ORAD.

  Two kinds of review:
  - All kind of Operations working group documents.
  - Selected review of other area groups (like ROAD stuff, etc.).

Discussion

There is a need for an explicitly nominated ORAD membership as distinct from the open ORAD meetings at the IETFs. This closed group will be responsible for the above listed topics. It is not necessary that those who are part of ORAD personally review documents but they shall see to it that such a review is made.

Current Operations Area working group Chairs could be part of ORAD but this is not implicitly required. It is necessary that individuals that have the interest and time to undertake the ORAD responsibilities.

There is a need for a method of flagging documents for ORAD review. If enough ORAD members thinks it needs an ORAD review, one member is assigned the responsibility to see that this happen.

To ensure broad coverage in ORAD, it will be necessary to invite operators that traditionally do not participate at IETFs.

Those interested in ORAD participation include:

Tony Bates	University of London
Nevil Brownlee	University of Auckland
Henry Clark	OARnet
Michael Conn	MCI
John Curran	NEARnet
Phill Gross	Advanced Network and Services, Inc.
Daniel Karrenberg	RIPE NCC
Peter Lothberg	EBONE
Bill Manning	SESQUINET
Bernhard Stockman	SUNET
Evan Wetstone	SESQUINET
Christopher Wheeler	University of Washington

Proposed Charter

1. What is the ORAD (and what is it not).
2. Forum for Operations groups.
3. Development of methods and practices.
4. Guidance and review.
5. Operations information and education.

The need for a backbone requirements document was discussed. There is value in having documents outline needs, services, and interoperation, but if it is too proscriptive, they may fail to accommodate all economic or organizational models.

Discussion around MBONE Coordination

There is a need to increase multicast performance in today routers

Matt Mathis volunteered to track MBONE contacts for the subversive purpose of collapsing connections to the highest level possible. The right thing to do is prevent the mbone from being heavily used until mrouted is fixed. If the operators were to turn it off, however, there would be a grass roots mbone appearing over which we would have NO control.

ORAD should issue a statement of recommendations on mbone utilization, requirements and operation. In the meantime, can we get ebugging tools, can we get multicast support from vendors?

Architectural weakness: Twenty-Five speakers at once fills a T1. This would create a situation of denial of service.

Mrouted needs more knobs. Must be able to do route pruning.

Action items for mbone:

- Major mrouted work
- Get together and list some bullets to take to Steve Deering and Steve Casner.
- Remove redundant tunnels.
- Public versions released as receive-only?
- Restrict to one audio and one video until more experiences.
- More tests and freeze of topology.

Tests shall happen before IETFs which includes announcements of tunneling and requests to be made further in advance of conferences. Strict cut-off date after which no more tunnels

Others actions:

- Need for more efficient diagnostic tools.
- mrouted related work.
  - Put throttling in the tunnels.
  - Treatment for misconfiguration (view others' configurations).
  - Pruning of the tree (no more than 12?).
  - Encaps, not LSRR.
  - Experiment with one-way path.
  - Encourage codings which conserve bandwidth.
  - Experiment outside of IETF meetings.

The Working Group needed to flesh these out with representatives from Merit, PSC, NEARnet.

Attendees

Tony Bates	t.bates@nosc.ja.net
Rebecca Bostwick	bostwick@es.net
J. Nevil Brownlee	nevil@aukuni.ac.uz
Henry Clark	henryc@oar.net
Michael Conn	4387451@mcimail.com
John Curran	jcurran@bbn.com
Hans Eriksson	hans@sics.se
Dennis Ferguson	dennis@ans.net
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Peter Ford	peter@goshawk.lanl.gov
Phillip Gross	pgross@nis.ans.net
Robert Gutierrez	gutierre@nsipo.nasa.gov
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Daniel Karrenberg	daniel@ripe.net
Mark Knopper	mak@merit.edu

Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Bill Manning	bmanning@sesqui.net
Dennis Morris	morrisd@imo-uvax.disa.mil
David O'Leary	doleary@cisco.com
Andrew Partan	asp@uunet.uu.net
Marsha Perrott	mlp+@andrew.cmu.edu
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

2.5.1 BGP Deployment and Application (bgpdepl)

Charter

Chair(s):

Jessica Yu, `jyy@merit.edu`

Mailing Lists:

General Discussion: `bgpd@merit.edu`
To Subscribe: `bgpd-request@merit.edu`
Archive: `/pub/bgpd-archive on merit.edu`

Description of Working Group:

The major purpose of this Group is to coordinate BGP deployment and application in the current Internet.

It intends to create a forum for BGP users to share BGP deployment experiences and also provide a channel for users to communicate with router vendors who implemented or who are implementing BGP. It also intends to discuss BGP policy application and coordinate policy implementation in the current internet routing environment which includes defining the usage of policy, defining a mechanism to share policy information, etc.

Goals and Milestones:

Ongoing	Facilitate the deployment of BGP as widely as possible.
TBD	Define the issues and the needs of policy routing in the current Internet architecture. Discuss how BGP policy routing capability applies to Internet policy routing needs. A document may be generated on this topic.
Dec 1992	Post as an Internet-Draft, a report of BGP deployment status.
Mar 1993	Post an Internet-Draft, defining a mechanism to share policy information between Administrative Domains.

CURRENT MEETING REPORT

Reported by Matt Mathis/PSC

Minutes of the BGP Deployment Working Group (BGPDEPL)

The immediate agenda of the BGP Working Group is to expedite BGP-4 deployment in the Internet. At this moment none of the vendors have BGP-4 code running. Network operators are all waiting on their vendors, and our task is to encourage the vendors to complete BGP-4 as soon as possible.

The AS path tree collected from Pittsburgh was discussed. *Editor's Note (md): A copy of the AS path tree is available via ftp under bgpdepl-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.* Of the approximately 5000 networks which are currently reachable, almost 3000 are being originated with EGP2. No single AS originating more than 100 networks is using BGP at all, including all other US Backbones and large regionals. NSFnet, Ebone, CAnet, a few mid-sized regionals, and ANS stub customers account for all deployed BGP in the Internet today. Four mid-level operators (representing relatively large networks) indicated that they had tested BGP-3, and had reverted to EGP2 because their vendor's implementations to be missing some features required in their environment. (Note that there are several BGP based transit networks including ANS, AlterNet, Sprint and the EBone, but these originate relatively few networks.)

This situation is pretty bleak. There is only a short time available to test and deploy BGP-4, and operators who have not yet deployed BGP-3 will face additional difficulties phasing out EGP2. It is desirable for all network operators to have BGP experience as soon as possible.

Implementor reports:

ANS	Currently has a "routed" implementation of BGP-1 and BGP-2 running in the backbone. Will be converting to "gated" in January or early February (This will deprecate BGP-1). They expect to begin testing BGP-4 w/CIDR in the production backbone during the First Quarter of 1993.
cisco	Currently support BGP-2 and BGP-3. Chose not to announce plans or timetable for BGP-4.
Wellfleet	Plans to have BGP-3 in summer of 1993 and BGP-4 sometime after. EMPB - PTT Telecom, The Netherlands). Will have BGP-3 Soon. Uncertain about schedule for BGP-4.
IBM	(Networking Systems at RTP) - BGP-3 has been announced for 6611 router in 2Q93.
3Com	Hopes to be testing BGP-4 by March, for possible release by August.

BBN	BGP-3 is currently deployed in T20's. They are working on BGP-4. Proteon was not represented, but several of their customers (who were present) have been promised BGP-4 in April 93.

Claudio Topolcic, from CNRI, presented RFC1366 and RFC1367, the proposed addressing plan and schedule to deploy CIDR. This served to initiate a general discussion about timetables.

The consensus was that the following schedule is reasonable:

1/93-3/93	BGP-4 interoperability testing with the ANS testing facility (See below - all vendors are invited to participate).
3/93	BGP-4 capable code deployed in the production NSFnet backbone. Begin testing by propagating some non-production/non-ABC networks through the backbone.
6/93	Start aggregating *production* networks in the backbone.
12/93	Completely phase out all EGP2 on NSFnet DMZs.

There was some discussion about how this interacts with the new network assignment rules and schedule specified in RFC1366 and RFC1367. It was noted that eight of the operators present had already obtained contiguous blocks of 256 class C addresses in advance of the formal plan. This has two implications: Since the NIC does not track assignments which are slated for subsequent reassignment, current growth figures are exaggerated by at at least 2000 networks and probably more.

Clusters of Class C networks will start appearing before predicted by RFC1367. This may slightly steepen the growth of the routing tables but most of the operators were not concerned. Several pointed out that most mature technology companies already had old network numbers. At the regional level most of the sites requesting new numbers were new to IP networking, and even if they had plans for hundreds or thousands of IP hosts, they currently only needed to announce one or two class C networks out of their entire block.

The first aggregation of production networks (scheduled for June 1993) will be a flag day for any site requiring full routing tables and not running BGP-4. There are only a few sites which may find themselves in this situation. All others can use default to route to their favorite backbone/mid-level provider. The sites which require full routing tables are either enforcing AUPs or implementing optimal routing to multiple backbones or mid-levels. In the latter case, not having BGP-4 will result in non-optimal routes for the aggregated networks only. This is not viewed as a show stopper. However, if some organization has a mission requirement to enforce an AUP but can not deploy BGP-4 by June 1993, there will be pressure to slip the schedule. This must be balanced against the cost to everyone else of carrying ever growing non-CIDR routing tables.

We also considered the scenario where BGP-4 isn't deployed on time for some reason (e.g., finding a serious bug in the specification.) Our conclusion was that the above discussion also applies, as long as the NSFnet backbone can still hold the routing table for the entire Internet. All other providers can (in principle) use the default to reach the NSFnet when they can no longer hold the entire routing table themselves. As a corollary, the Internet will survive the RO+AD crisis, if and only if, at least one network provider implements a solution. If no solutions are successful, the network provider of last resort will be blamed, even if the problem has become intractable.

Jordan Becker (ANS) estimated that full route aggregation will reduce the current routing tables by about 30%, because the old address assignment policies tended to allocate addresses in blocks anyhow.

Jordan repeated his invitation for all router vendors to make use of the ANS test facility for BGP interoperability testing. Vendors can make arrangements to bring equipment to the test facility, or to do remote testing through a "wide areax" BGP session across the internet. Contact him for further information (becker@ans.net).

Dennis Ferguson (ANS/gated implementor) announced that he had avoided implementing the supernet expansion code, which would map BGP-4 supernets into a block of BGP-3 class C networks. He was concerned that this is difficult to do correctly and unlikely to be useful for very long. After some discussion it was observed that the eight blocks of 256 class C networks represented in the room would be announced as 8 BGP-4 routes as soon as prudent. These would in turn map to 2048 class C BGP-3 routes, which are likely to immediately overwhelm any BGP-3 router. Ergo the supernet expansion code would probably be useful for less time than it would take to debug it. The BGP Deployment Working Group concurs that implementing the supernet expansion code would not be cost effective.

Tony Li of cisco made an interesting point about software release schedules. Releases are typically on a six month cycle with an eight month lead time between code freeze and "first customer ship" (aka: beginning of General Availability for some new feature). Cisco's next scheduled code freeze is February 1993, so even if bug-free BGP-4 code exists today the earliest it will appear in General Availability products is October 1993. All customers who need BGP-4 before then must run pre-GA code. Other vendors were in general agreement. Some of the operators were upset - one was operating under a prohibition against using non-GA code.

Attendees

Nagaraj Arunkumar	nak@3com.com
Tony Bates	t.bates@nosc.ja.net
Jordan Becker	becker@ans.net
Jeffrey Burgan	jeff@nsipo.nasa.gov
Henry Clark	henryc@oar.net

Stefan Fassbender	stf@easi.net
Dennis Ferguson	dennis@ans.net
Peter Ford	peter@goshawk.lanl.gov
Vince Fuller	vaf@stanford.edu
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Eugene Hastings	hastings@psc.edu
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@server.af.mil
Akira Kato	kato@wide.sfc.keio.ac.jp
Dave Katz	dkatz@cisco.com
John Krawczyk	jkrawczy@wellfleet.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Daniel Long	long@nic.near.net
Peter Lothberg	roll@stupi.se
Matt Mathis	mathis@a.psc.edu
David Meyer	meyer@oregon.uoregon.edu
Dennis Morris	morrisd@imo-uvax.disa.mil
Peder Norgaard	pcn@tbit.dk
David O'Leary	doleary@cisco.com
Yakov Rekhter	yakov@watson.ibm.com
Erik Sherk	sherk@sura.net
Bernhard Stockman	boss@ebone.net
Claudio Topolcic	topolcic@cnri.reston.va.us
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Chris Wheeler	cwheeler@cac.washington.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

2.5.2 Benchmarking Methodology (bmwg)

Charter

Chair(s):

Scott Bradner, sob@harvard.edu

Mailing Lists:

General Discussion: bmwg@harvard.edu
To Subscribe: bmwg-request@harvard.edu
Archive:

Description of Working Group:

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

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

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

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the Group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

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

Done The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.

Done Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Request For Comments:

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices"

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

Minutes of the Benchmarking Working Group (BMWG)

The Benchmarking Methodology Working Group met on Wednesday afternoon during the Washington IETF meeting. The topic of discussion was mostly the draft of the frame formats memo that is to accompany the methodology memo. The basic document was approved and some suggestions were made for improvement. Both the methodology memo and the frame formats memo will be submitted as Internet-Drafts after the requested editing.

A note for the record: The NIC has assigned a block of Class C IP network addresses to be used in the methodology memo. The intent of the assignment is to ensure that, even if a tester is accidentally attached to the Internet, minimal service interruption will occur since the addresses will never be "known" to the backbone. The addresses assigned are the Class C networks from 198.18.1.0 through 198.19.254.0.

Attendees

Lou Berger	lberger@bbn.com
Rebecca Bostwick	bostwick@es.net
Scott Bradner	sob@harvard.edu
Kay Chang	chang@chang.austin.ibm.com
Barbara Denny	denny@erg.sri.com
Kevin Dubray	kdubray@wellfleet.com
Jacques Dugast	dugast@issy.cnet.fr
Olafur Gudmundsson	ogud@cs.umd.edu
Jack Hahn	hahn@sura.net
Merike Kaeo	merike@alw.nih.gov
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Paul Sangster	sangster@ans.net
Cris Shuldiner	cws@ftp.com
Bernhard Stockman	boss@ebone.net
Von Welch	vwelch@ncsa.uiuc.edu

2.5.3 Network Joint Management (njm)

<u>Charter</u>

Chair(s):

Gene Hastings, hastings@psc.edu

Mailing Lists:

General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu
Archive:

Description of Working Group:

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the Internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the Agenda of this Group will include the IAB and its task forces, and groups within FARNET. In particular FARNET has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the Group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing.

Goals and Milestones:

None specified

CURRENT MEETING REPORT

Reported by Gene Hastings/PSC

Minutes of the Network Joint Management Working Group (NJM)

Announcements: Due to the number of presenters, there was an extension of Network Status Reports into part of the NJM session. Presentation slides for the NETSTAT portion can be found in Section 3 of the Proceedings.

There was positive feedback after both sessions for having combined presentations and discussions spurred by them. Consequently the sessions will be more tightly coupled in the future.

Mark Knopper, Jordan Becker - Merit/ANS T3 Network Status

A number of refinements and enhancements are planned, including RS960 FDDI deployment, and routing software changes such that an ENSS will not announce 140.222 if it is isolated from the backbone.

Changes in backbone internal configuration are planned to minimize coast-to- coast delay. Jordan said he will put the PostScript version of the delay map on line. (also included in Proceedings)

Change tonight (November 17th): CLNP will become encapsulated, [in one of the T1 PSPs] and transferred over IP, instead of being switched in a native stack. This is part of the migration plan to move all remaining traffic from the T1 net to the T3. It will initially remain on T1 net.

General notes and announcements:

- EASInet at CERN will have an ENSS, connected to NY.
- Traffic Source/Destination pair statistics are sampled, with a frequency of 1 in 50.
- Am map showing the T3 backbone with the T1 backup net is available online. Question: Do you have priority queuing for management traffic? - Not yet; AIX 3.2 will have priority queuing for routing traffic.
- The MTU in the backbone is set to 1500 on most interfaces. With deployment of new interfaces, many are being changed to 4000. ENSS FDDI can do MTU discovery. The deployment strategy is designed to avoid fragmentation on Ethernet interfaces, at the potential cost of inefficient FDDI use in the short term.
- The dismantling of the T1 backbone circuits begins December 2nd. Within CNSS PoPs, there are DSU upgrades planned, granting multipoint capability, and redundancy within PoP.

- Merit/ANS have issued an invitation to users, vendors and developers to come to the ANS test net to test interoperability of new BGP support.

Milo Medin - NASA Science Internet

An upgraded, year-round link to McMurdo base in Antarctica will be through ARC2. It was a seasonal 56kb, and has been upgraded to full time 384kb via IntelSat IV Lessons learned from NSI operations: Don't encapsulate (DECNOT) on 56kb lines.

NSI Currently does DECNOT on TCP. Doing it over UDP was a lose.

DEC NSP is "broken" the way TCP used to be w.r.t. retransmission. With DECNOT, TCP does retransmission, and NSP does not have to.

NSI network management is migrating to MSU (but not MCC) ("Polycenter 2000")

Future efforts:

- Awaiting Commerce Dept. OK for "Export of Internet" to Russia to allow network links to same and other parts of the Former Soviet Union.
- NASA participation with DoE ATM will have connections at Langley, Lewis, Goddard, Ames, JPL - AGS+.

Bernhard Stockman, EBONE

There is now a security access scheme in EBONE routers (Amalgam of Kerberos and TACACS). There will probably be a 256kb link from Bonn to Stockholm.

Network Joint Management - November 19, 1992

Bob Collet - Sprint

Sprint operates three internets, with Internet connectivity.

1. SprintLink (Domestic U.S.)
2. ICMnet-2 (Atlantic)
3. ICMnet-3 (Pacific)

Sprint uses cisco except in some private nets it uses Wellfleet. There are nine domestic customers, and nine more being provisioned. Sprint will be demonstrating T3 P-P at COMNET '93 in February. Most of the routers are owned by Sprint. There is some customer owned equipment, but only at customer premises. Sprint has management responsibility and sole configuration control of customer router.

Rich Fisher - GSFC

Project group at Goddard Space Flight Center does data acquisition from satellites, and redistributes it to numerous terrestrial labs. RAC :== Remote Analysis Computer (All VAX) All links provided by PSCN (NASA "phone co.")

Tony Hain - ESnet

[There are no accompanying slides for this presentation.] The Department of Energy's (DoE) original high speed RFP was two years ago, and began review in February 1992. An award was made to Sprint, with TRW and cisco as subcontractors. A protest was filed so the procurement is on hold. The General Accounting Office's (GAO) 90 day timer expires December 23rd. The GAO will say nothing before then...Therefore there will be a three month delay in deployment.

Initially access will be provided at T3, eventually a mix, up to OC-12. All local loops will be fiber. It is NOT a managed router service. The DoE and NASA will accept raw ATM. The router will be a cisco AGS+ with a CSC-4 processor and HSSI interface for T3. The router will connect via Digital Link CSUs. A new router will be needed to go beyond T3 rates. The net will do IP and CLNP. Planners are trying to figure out DECNET Phase IV, and negotiating with DEC to make Phase V genuinely be CLNP (as opposed to being only close). They are still discussing which IGP to use. (contemplating OSPF and IBGP).

Six DoE sites are planned as part of the initial project:

1. Lawrence Livermore National Laboratory
2. Lawrence Berkeley Laboratory
3. Los Alamos National Laboratory
4. FermiLab
5. Oak Ridge
6. Superconducting Supercollider, Waxahatchie, TX.

Continued plans call for the eventual connection of all ESnet sites, but sites that are not currently connected to ESnet will not necessarily be connected directly. If they were already being considered for a T1, they may get a connection.

In addition, there are five NASA planned as well:

1. Langley Air Force Base
2. NASA Lewis Research Center
3. Goddard Space Flight Center
4. Ames Research Center
5. Jet Propulsion Lab

Mark Knopper asks: What is the plan and timeline for testing vs. production?

- Service will be brought up off-line.
- Several sites will be cut into ESnet.
- Procurement includes an off-line testbed.

Since NASA and DoE will be sharing a fabric, they will eventually move their peering to this net (vs. at FIXes) There will be several logical subnets on the fabric:

- DoE internal
- NASA internal
- DoE-NASA "Phantom DMZ" for peering.

SDSC has gotten funding from NSF to participate. Will initially be part of the DoE subnet. Tony will have more technical information in January (encapsulation, etc.).

Mark Knopper: ERNET.(India Research And Education Net).

ERNET folks visited several U.S. operators and Internet researchers, including Merit. ERNET is funded by the Indian Government and the United Nations.

Phase III (refer to slides) uses VSAT for domestic networking due to inadequacy of domestic telecomm IP and CLNP over X.25. 128kb up, and 512kb down.

Original connection was Alternet UUCP, dialup uucp, analog leased uucp, analog leased SLIP (and improved UUCP performance, as TCP is better at line utilization) usually had > 20MB of mail in their queue.

Use mostly cisco MGS, some CGS. Currently has 1 class B network, nationwide...

NJM Discussion

- Route aggregation, BGP4 deployment can you do it?
- Does there exist a CIDR traceroute? does anyone know how to interpret one?
- Dennis Ferguson says that as a transition plan, it is easier for the midlevel to do route aggregation than the backbone. He suggests that such clients advertise both explicit routes and an aggregated route. The Backbone will install the explicit routes, and announce only the aggregated one.
- Continued from BGP Deployment Working Group. (See also BGPDEPL Minutes): Ref. Claudio Topolcic's timeline. Must examine transition to CIDR - Tools? If operators do not do BGP4 by (date), they will need to accept default only from backbone. What is that date? If Merit's deployment plan works, it may be December 31, 1993.
- ROAD transition issues. Who really needs to do BGP4? Many "stub" or tail nets will not need to soon. What is the operational impact of dinosaurs? Will they really die out? Can we afford to sustain them?

Attendees

Henry Clark	henryc@oar.net
Robert Collet	rcollet@icm1.icp.net
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Dennis Ferguson	dennis@ans.net
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Vince Fuller	vaf@stanford.edu
Tony Hain	alh@es.net
Eugene Hastings	hastings@psc.edu
Mark Knopper	mak@merit.edu
Kim Long	klong@sura.net
Matt Mathis	mathis@a.psc.edu
David O'Leary	doleary@cisco.com
Andrew Partan	asp@uunet.uu.net
Marsha Perrott	mlp+@andrew.cmu.edu
Tom Sandoski	tom@concert.net
Bernhard Stockman	boss@ebone.net
Claudio Topolcic	topolcic@cnri.reston.va.us
William Warner	warner@ohio.gov
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov
Paul Zawada	Zawada@ncsa.uiuc.edu

SprintLink Routing Services

Public Data Internet

and

International Connections Manager for the NSFNET

Robert Collet
Program Manager
rcollet@icm1.icp.net

NSFNET International Connections Manager Service

Cooperative Agreement for International Connections Management (ICM) between the National Science Foundation (NSF) and Sprint

Provides International Routing Services for the NSFNET

Service provisioned via two router-based networks:

ICMnet Atlantic

ICMnet Pacific

ICMnet service provides transit between NSFNET and foreign national networks

ICMnet service also provides some transit among some foreign networks

ICMnets integrated with Sprint's national SprintLink Public Data Internet

ICMnet supports both NSF Appropriate Use Policy (AUP), i.e., Research and Education (R&E) and Commercial (CO) requirements.

NSFNET International Connections Manager Service (continued)

ICMnet Atlantic:

EBONE - London (12/92)
EBONE - Paris
EBONE - Stockholm
UNINET-ZA (South Africa)
Kuwait (12/92)
United Arab Emirates (12/92)
Turkey (1/93)
CRnet (Costa Rica) (1/93)
ECUAnet (Ecuador) (1/93)
SprintLink
Alternet
SURAnet
More in 1993

ICMnet Pacific:

SINET (Japan)
MIMOS (Malaysia) (11/92)
SprintLink
More in 1993

SprintLink Public Data Internet

Public Internet Access Service

LAN to LAN routing

Service operational 7/92

Private Line Access: 9.6/19.2kbps analog
56/64 kbps - T1

Pricing is distance insensitive

Flat monthly rate (function of access line speed)

Commercial (CO) Internet Access

Research and Education (R&E) Internet access

Interconnections with Sprint Frame Relay and SprintNet X.25 networks

Planned Features: Dial-up via SprintNet
Dial-up via 800
Network Information Service Center (NISC)

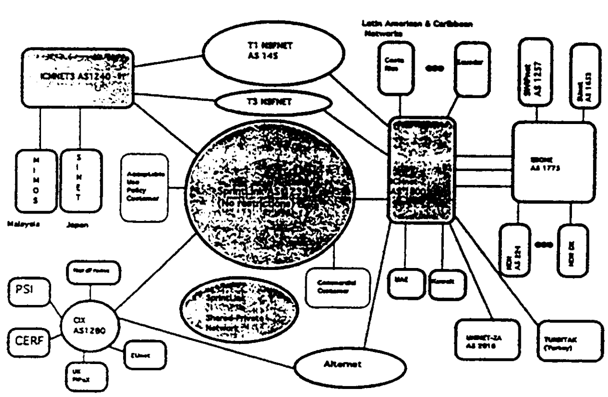

Figure 1. SprintLink/ICMnet Peering. SprintLink obtains connectivity to the AUP Internet via Sprint ICMnets East & West. (Shaded areas are Sprint networks)

October 16, 1992

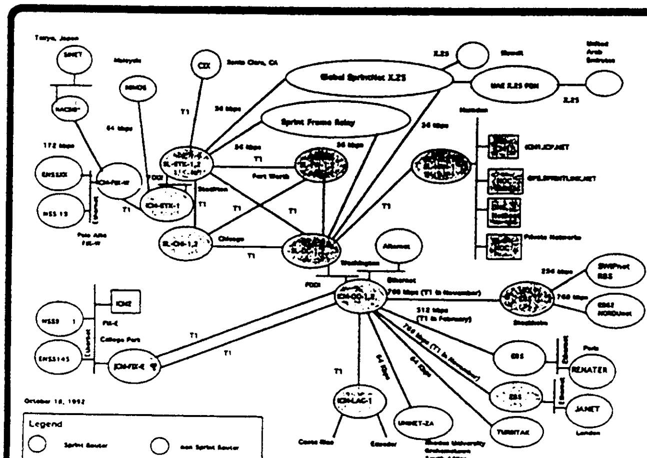

Figure 2. SprintLink/ICMnet Topology

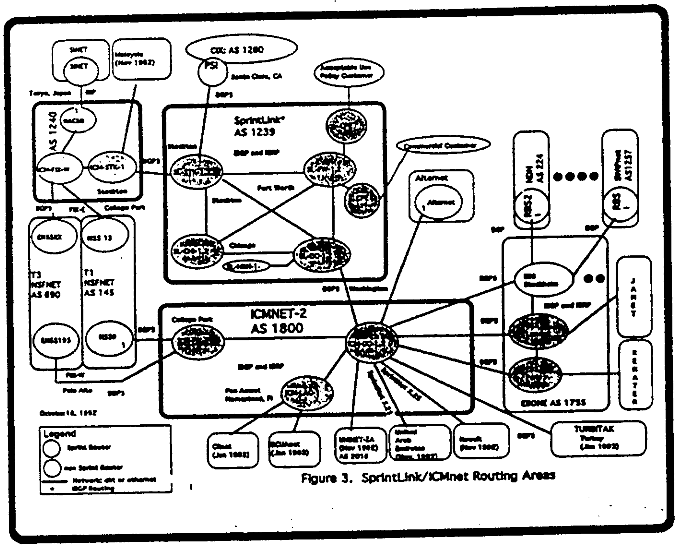

Figure 3. SprintLink/ICMnet Routing Areas

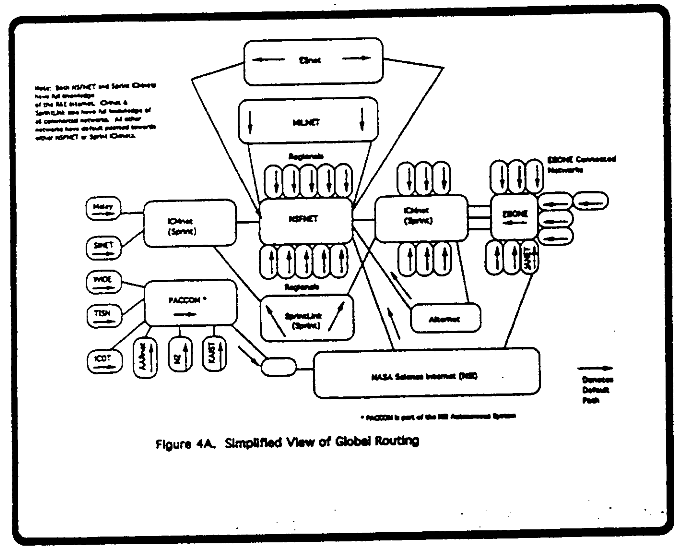

Figure 4A. Simplified View of Global Routing

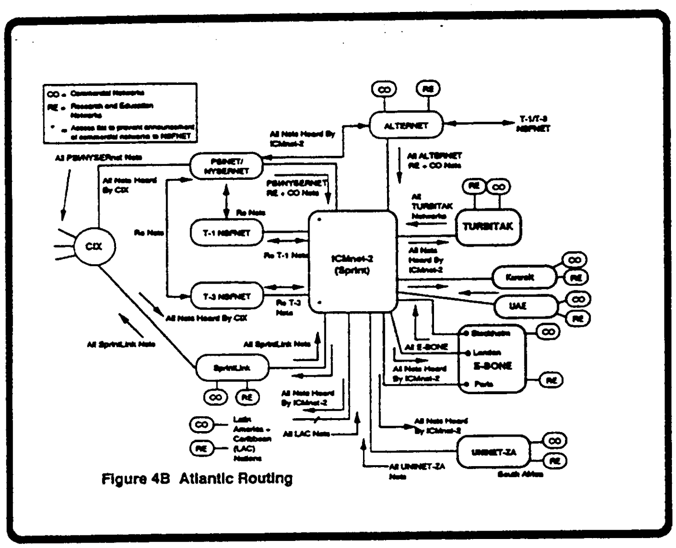

Figure 4B Atlantic Routing

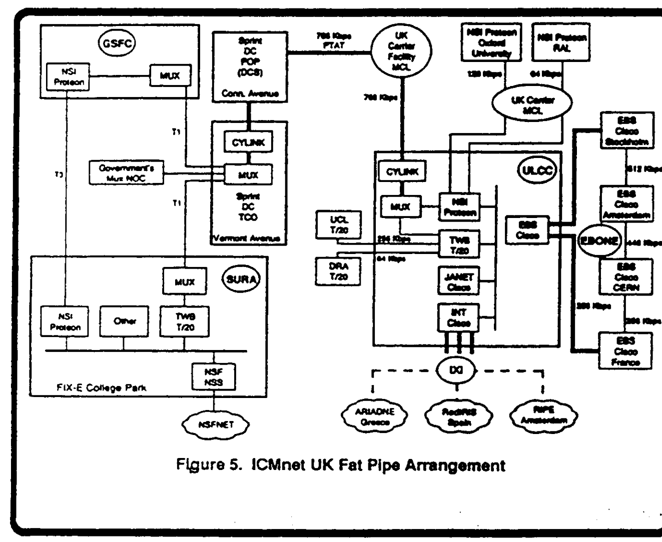

Figure 5. ICMnet UK Fat Pipe Arrangement

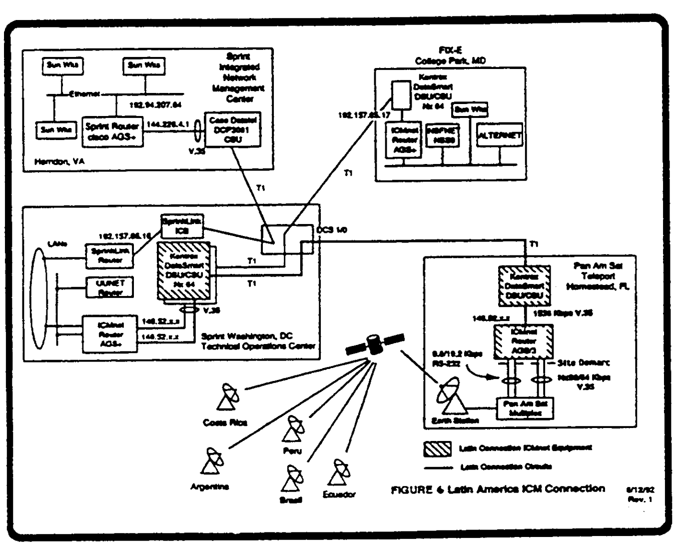

FIGURE 6 Latin America ICM Connection

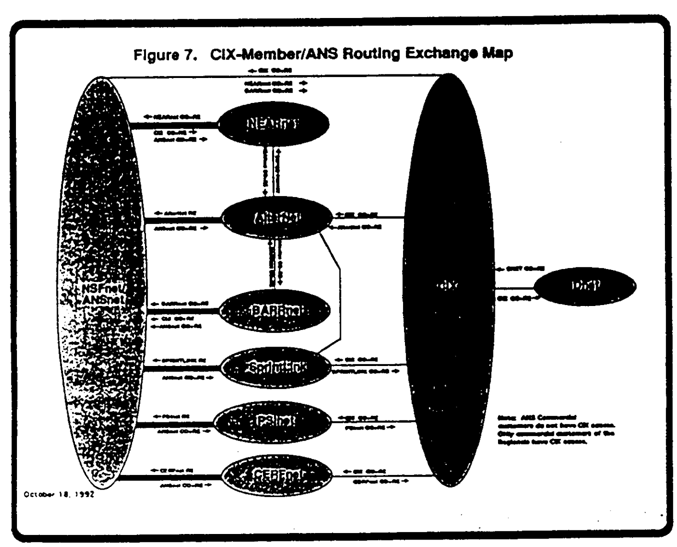

Figure 7. CIX-Member/ANS Routing Exchange Map

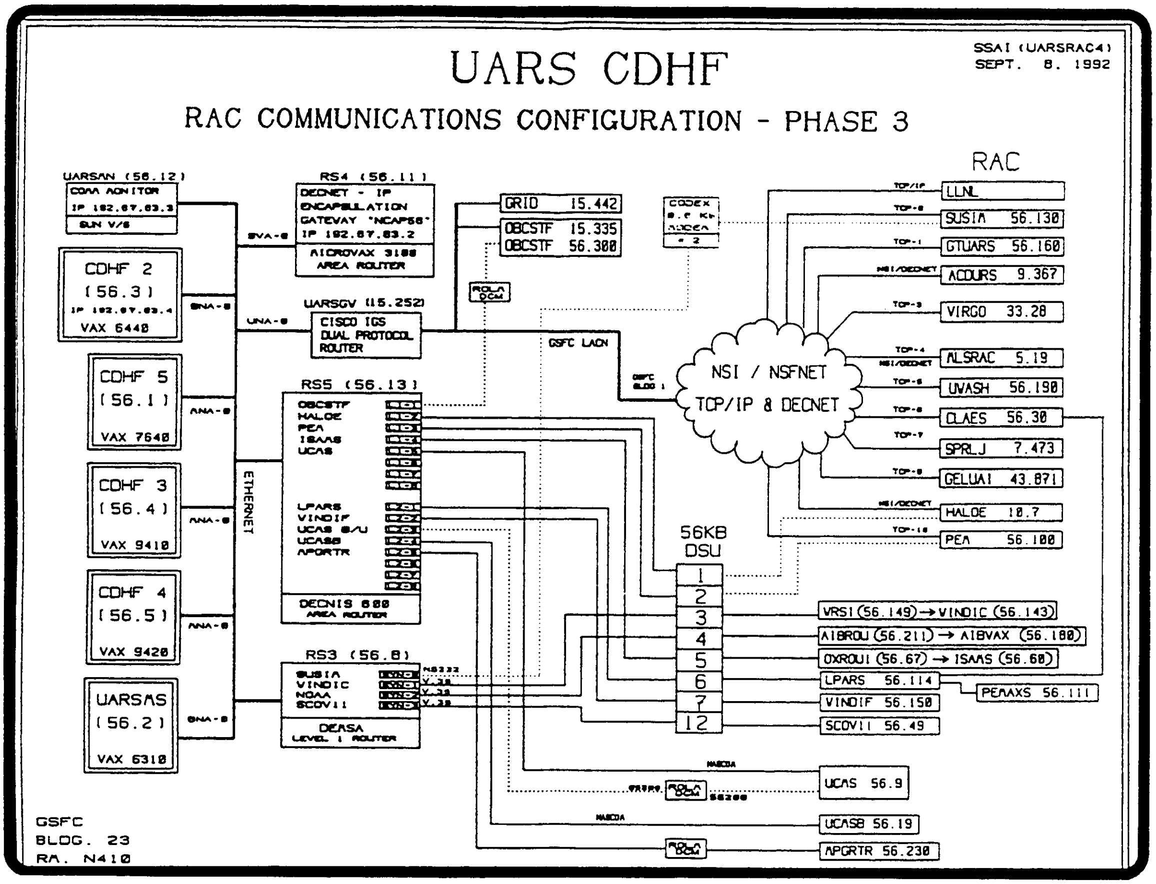
UARS CDHF
RAC COMMUNICATIONS CONFIGURATION - PHASE 3
SSAI (UARSRAC4)
SEPT. 8. 1992
RAC
UARSAN (56.12)
RS4 (56.11)
CDHF 2
(56.3)
VAX 6440
UARSGV (15.252)
CISCO IGS
DUAL PROTOCOL
ROUTER
CDHF 5
(56.1)
VAX 7640
RS5 (56.13)
CDHF 3
(56.4)
VAX 9410
ETHERNET
CDHF 4
(56.5)
VAX 9420
DECNIS 600
RS3 (56.8)
UARSMS
(56.2)
VAX 6310
GSFC
BLDG. 23
RM. N410
GRID 15.442
OBCSTF 15.335
OBCSTF 56.300
NSI / NSFNET
TCP/IP & DECNET
56KB
DSU
LLNL
SUSIM 56.130
GTUARS 56.160
ACOURS 9.367
VIRGO 33.28
ALSRAC 5.19
UWASH 56.190
CLAES 56.30
SPRLJ 7.473
GELUAI 43.871
HALOE 10.7
PEM 56.100
VRSI (56.149) → VINDIC (56.143)
AIBROU (56.211) → AIBVAX (56.180)
OXROUI (56.67) → ISAMS (56.60)
LPARS 56.114
PEMAXS 56.111
VINDIF 56.150
SCOVII 56.49
UCAS 56.9
UCASB 56.19
APGRTR 56.230

ERNET: Education and Research NETwork Project

- Goals
  - To establish a Computer Network for Academic and R&D Institutions
  - Technology development and demonstration
  - Manpower training
- Funded by
  - United Nations Development Program (UNDP) : $6 million
  - Govt. of India, through Dept. of Electronics : $3 million
- Initial participants for technology development and demonstration
  - The 5 IITs (Bombay, Delhi, Kanpur, Kharagpur, Madras)
  - IISc, Bangalore
  - National Centre for Software Technology (NCST), Bombay
  - Dept. of Electronics, Delhi

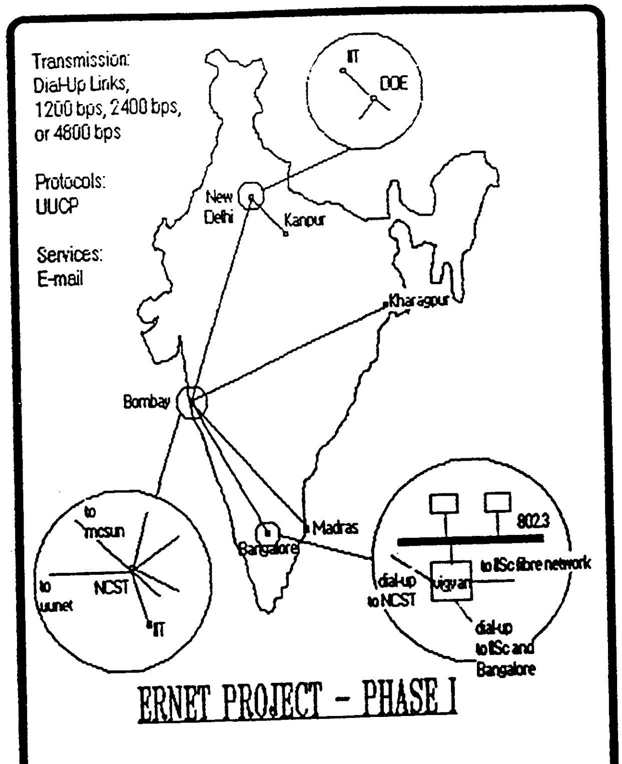

ERNET PROJECT - PHASE I

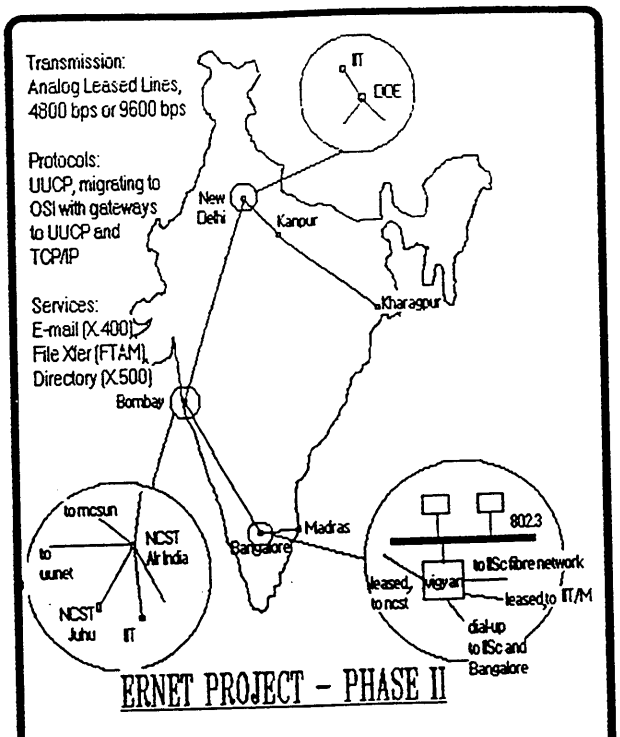

ERNET PROJECT - PHASE II

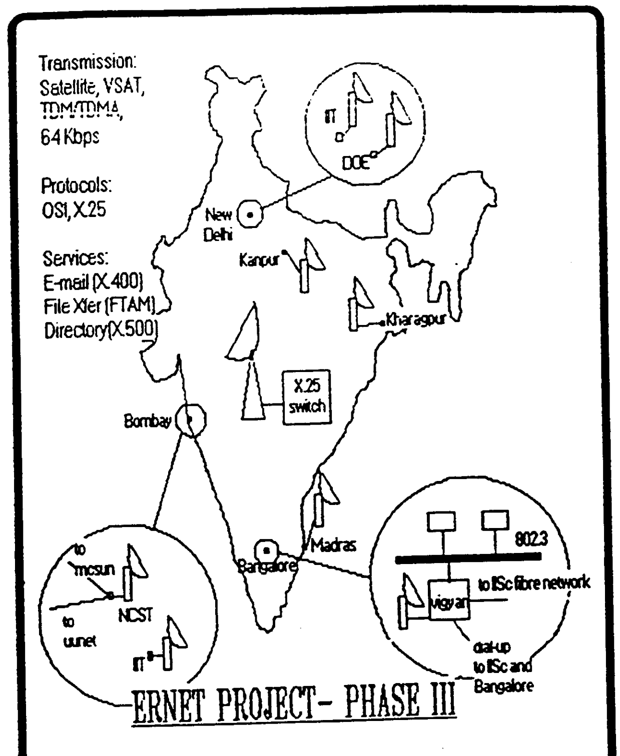

ERNET PROJECT- PHASE III

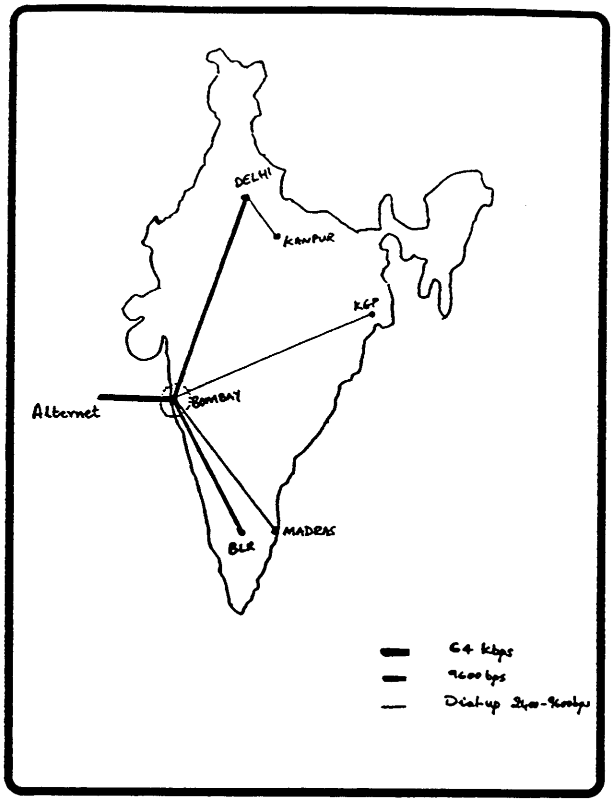

USER COMMUNITY

- 150 ORGANIZATIONS
- 30,000 USERS
- ACADEMIC, R&D

TRAFFIC

- 120MB/day, Largely e-mail

SERVICES

- E-MAIL
  - o INTERNET
  - o X.400 Mail
  - o UUCP Mail
- DIRECTORY SERVICES
  - o X.500
  - o Whois
- INFOSERVERS
  - o General Information Dissemination
- USENET NEWS
- INTERNET SERVICES
  - o telnet
  - o ftp
  - o r*

OTHER OSI SERVICES

- FTAM

MULTIMEDIA MAIL

- nsb's MIME implementation being tested out between ERNET sites.

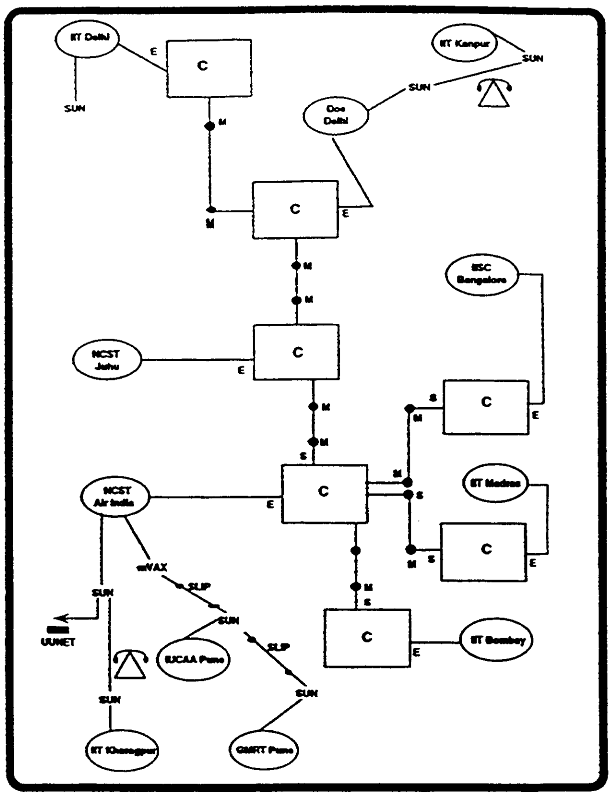

PROTOCOL STACK

TCP	TP4
IP	CLNS
HDLC	

SERVER CONFIGURATION

RFC 822 (Send Mail) NTP, DNS, UUCP	X400, X500 FTAM	
TCP	RFC1006	TP4
IP		CLNS
HDLC		

THE ERNET SATWAN SEGMENT

Why Satellite ?

- Provides easy accesspoints into the network from remote areas
- Ease of connection to the network
- Relatively low cost (VSAT technology)
- Flexible bandwidth allocation
- High Media Reliability (the need of the hour)
- High degree of scalability

Additional Services

- Library Information Services (INFLIBNET)
- Access to Multimedia Information Repositories, country wide
- Large volume data transfers
- Access to computing resources across the country

Unique

- First large scale deployment of VSAT tech for IP networking
- First attempt at experimenting with TOS based routing

Brace up for

- Operational Problems with X.25
  - maximise performance
  - Routing table overflows
- Single vendor package (HNS)
  - Not much room for creativity !?

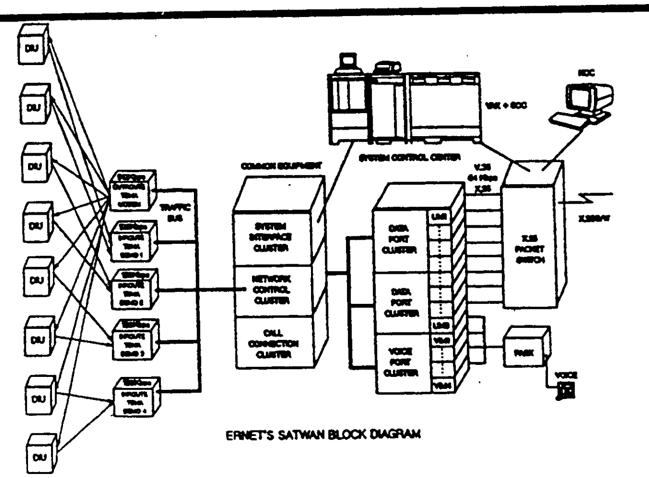

ERNET'S SATWAN BLOCK DIAGRAM

ERNET-SATWAN - 0

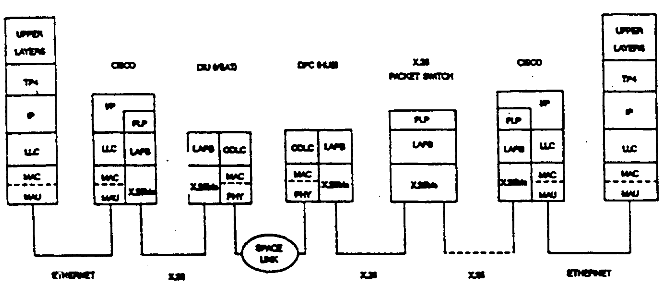

TYPICAL VSAT – PACKET SWITCH CONNECTION PROTOCOL SUITE

ERNET- SATWAN - 0

DATA NETWORKS IN INDIA

INET

- Dept of Telecommunications (DoT)
- Public X.25 Data Network
- Dial-in / Leased line connections 1200-9600 bps access to PADS
- 9600 bps switching between PADS
- Membership - 50 organisations
- Services
  - Remote database access
  - International database access (EASYNET, DIALOG, DOWJONES)
    Gateway Packet Switched Services - GPSS
  - Interconnection of organisation wide networks allowed
  - Interconnection of private networks allowed

SOFTNET

- Department of Electronics (DoE)
- TCP/IP national / international connectivity for commercial traffic
- Leased / dial-in / INET access to customers
- Links to USA, Europe and the Far East
- Operations expected to start by Jan 1993

INDONET

- Computer Maintainence Corporation - CMC (DoE)
- Closed user group of IBM users
- Leased 2400 access to 6 major cities
- Dial-in access only (2400 bps)

NICNET

- The Planning Commission
- Satellite based, connects upto 450 District HQ
- Exclusively for governmental use

Management Structure

Three layered management structure

PPC

APCU

IPMC

PPC - Policy Planning Committee
APCU - Apex Programme Coordination Committee
IPMC - Institute Project Monitoring Committee

IPMC

- Chairperson -> Director / Nominee (on-site Project Coordinator)
- Functions
  - o Review Work Plans at the site
  - o Oversee Project Implementation
  - o Review and Report Project Progress
  - o Other routine managerial tasks (even network monitoring, these days !)

APCU

- Chairperson -> National Project Coordinator
- Members
  - o Project Coordinators from participating sites
  - o Rep from the Telecom Commission
  - o Deputy Financial Advisor to the Govt. of India / Ministry of Finance
  - o Project Director

– Takes decisions on the operation of the project

Management Structure ....contd

PPC

- Chairperson -> Seceratary, DoE
- Members
  - Representatives from
    - The Planning Commission
    - Dept. of Economic Affairs
    - Ministry of Human Resource Development
    - The Telecom Commission
    - Department of Science and Technology (DST)
    - UNDP & Office of Project Services (UNDP-OPS)
- Is the project worthwhile ? Quantitative measures derived for impact on the exchequer

The Expert Panel

Dr. Klaus Ullman, DFN, Germany
Dr. Bob Cooper, JANET
Dr. Peter Kirstein, Head - Dept. of CS, UCL, LONDON
Dr. Vinton Cerf, CNRI, USA

Reviews

Monthly -> Onsite by the Project Coordinator
Biannual -> Sitewise by the Director and at the APCU
Annual -> By the Director, the expert panel and the UNDP

Immediate Action Plan

>> Consolidate Internet and OSI Services
(achieve by 15 Jun 1993)

- Set up formal NOCs at the backbone sites
  - o Formalise seperate Ops, Engg and Dev groups at each site
  - o Form a User Support Team and encourage user group feedback by way of user group meetings.
- Smoothen out the traffic pattern on the current international link
  - o Set up ftp archive servers on ERNET
  - o Cache / store info from the X500 DS database at PARADISE
- Provide more access points to ERNET both nationally and internationally
  - o National -> INET, SOFTNET (?)
  - o International -> Additional link to the US, a link to Europe and the Far East (Back-up routes and load balancing)
- Increase ERNET's visibility on the Internet
  - o Provide Internet access points to other South / South-East Asian countries
  - o Offer a WEP service for such countries
  - o Provide a the "Last Mile" to completing the global Internet circle
  - o Participate actively at forums like CCIRN, IEPG, INET and the IETF

Immediate Action Plan ....contd

>> Develop, Demonstrate and Propagate OSI across Asia. Contribute in the growth of OSI awareness and create a demand for OSI products

(achieve in 14 - 18 months - subject to revision)

- OSINET - ERNET's implementation of the OSI suite
  - o Complete interworking & conformance tests
  - o Check performance
  - o Package & Document
- Set up a Protocol Conformance Test Facility
- Implement OSI applications

ERNET - The forthcoming role

* Set up an NCC to coordinate Internetworking activity in the region and in India

* Implement, engineer and maintain a country wide high speed backbone for data networking

  - Funds ?
    Membership fees
    Sponsors

  - The obvious need for a coherent infrastructure to support n/w operations and management brings forth the requirement for a financially autonomous organisational setup to the fore !

* Network Expansion

  - encourage cooperative networking
  - help build regional networks and interconnect them to the ERNET backbone

The success of the ERNET project could be a good case study for both funding organisations as well as countries who intend to set up academic networks in their countries. ERNET could provide help in planning and technical assistance in implementation.

Three Cheers to all ERNET'ians for making ERNET happen

2.5.4 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, boss@ebone.net
Phillip Gross, pgross@nis.ans.net

Mailing Lists:

General Discussion: oswg-l@wugate.wustl.edu
To Subscribe: oswg-l-request@wugate.wustl.edu
Archive:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The Working Group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Done	Agreement on a model.
Done	Survey for most useful and popular metrics.
Done	Survey for most useful and popular presentation formats.
Dec 1990	Identify similar efforts being performed by other groups.
Done	Define a common minimal set of metrics.
Mar 1991	Propose a MIB for metrics not already there.
Done	Define a common storage format to facilitate data sharing.
Done	Define common presentation formats to make data comparable.
Mar 1991	Develop outline, and make writing assignments for paper (Opstat1) documenting March 1991 milestones.
May 1991	Complete paper Opstat1.

May 1991	Possible mid-term meeting to review Opstat1.
May 1991	Submit Opstat1 as Internet-Draft.
Jul 1991	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Jul 1991	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Jul 1991	Propose old tools to be retrofitted.
Jul 1991	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools.
Sep 1991	Complete paper Opstat2.
Sep 1991	Possible mid-term meeting to review Opstat2.
Sep 1991	Submit Opstat2 as Internet-Draft.
Dec 1991	Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Henry Clark/OARnet

Minutes of the Operational Statistics Working Group (OPSTAT)

Agenda

- Review of Client-Server Specification.
- Resource Utilization Criteria.
- Milestones/Goals Review.
- Statistical MIB.

Client-Server Protocol

Bernhard reviewed the client-server paper sent to the mailing list several weeks ago. *Editor's Note (md): A detailed summary of the commands within the protocol is available via ftp under opstat-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.*

Resource Utilization Criteria

The question has arisen before of the issues surrounding link utilizations and when a link should be upgraded and how to determine "fair" usage of a link by multiple organizations.

In terms of monitoring link usage, some networks query routers very frequently (as often as every 60 seconds) to detect peaks. Others try to track IP src/dest address pairs to track traffic flows. Some networks attempt to monitor usage at various points around their network to capture traffic flows. Some mention was made of an accounting mib such that traffic usage patterns could be withdrawn automatically via MIB queries. Some queries were to be made to the Internet Accounting Working Group to determine the relevance of their work to this topic.

There was an extended discussion on when to "upgrade" a link. When is it full? Should a link always run at max utilization in order to get maximum $$ from a link? Some mention was made of looking at the peak values, the duration of the peak values, the number and distribution of the peaks, and attempting to correspond the peak values to other events on the router such as errors and packet drops.

Bernhard felt that this topic should be moved from the Operational Statistics Working Group to another working group within the Operational Area. This was to be taken up at the ORAD meeting later in the week.

Goals & Milestones/Statistical MIB

The Common Storage Format RFC was submitted to the RFC editor in early November 1992. The initial re-examination of the client-server protocol has been completed. After some lengthy discussion, we moved the completion of the client-server Internet-Draft to the July 1993 IETF (with continuing discussions on the mailing list and at the March 1993 IETF in Columbus) and the completion of the Statistical MIB Internet-Draft at the March 1993 IETF with a first draft ready at the November 1993 IETF. Included in the discussions of dates was a discussion of the future SMP stuff and the get-bulk retrieval mechanisms for retrieval of data via the MIB which are to be examined in the future.

Attendees

Tony Bates	t.bates@nosc.ja.net
Rebecca Bostwick	bostwick@es.net
J. Nevil Brownlee	nevil@aukuni.ac.uz
Henry Clark	henryc@oar.net
Michael Conn	4387451@mcimail.com
John Curran	jcurran@bbn.com
Hans Eriksson	hans@sics.se
Dennis Ferguson	dennis@ans.net
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Peter Ford	peter@goshawk.lanl.gov
Phillip Gross	pgross@nis.ans.net
Robert Gutierrez	gutierre@nsipo.nasa.gov
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Daniel Karrenberg	daniel@ripe.net
Mark Knopper	mak@merit.edu
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Bill Manning	bmanning@sesqui.net
Dennis Morris	morrisd@imo-uvax.disa.mil
David O'Leary	doleary@cisco.com
Andrew Partan	asp@uunet.uu.net
Marsha Perrott	mlp+@andrew.cmu.edu
Bernhard Stockman	boss@ebone.net
Marten Terpstra	marten@ripe.net
Evan Wetstone	evan@rice.edu
Chris Wheeler	cwheeler@cac.washington.edu
Paul Zawada	Zawada@ncsa.uiuc.edu

2.5.5 User Connectivity (ucp)

Charter

Chair(s):

Dan Long, `long@nic.near.net`

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

Request For Comments:

RFC 1297	"NOC Internal Integrated Trouble Ticket System Functional Specification Wishlist ("NOC TT REQUIREMENTS")"

CURRENT MEETING REPORT

Reported by Dan Long/BBN

Minutes of the User Connectivity Working Group (UCP)

The User Connectivity Problems Working Group met for one session this time.

The Group had previously defined a data structure that would enable Trouble Ticket handoffs between NOCs. Paul Zawada had written an ASN.1-like description of the fields in this data structure.

Kaj Tesink drafted a document describing how some handoff fields could be represented in electronic mail messages. The Group discussed this and agreed that the document needs to be revised to reflect more of the previously-defined handoff fields. The goal is to allow trouble tickets to be mailed between NOCs both with and without internal trouble ticket systems. The format should be simple enough to enable humans to enter the data and yet regular enough to permit parsing. Paul and Kaj will work on this and get it out as an Internet-Draft. At that time, several groups agreed to experiment with the exchange format and to create a template to facilitate manual participation.

The UCP Internet-Draft on a Trouble Ticket Tracking System, originally written by Matt Mathis, had been discussed and revised heavily by the Group and it has now expired. Dan Long has volunteered to draft a new version which reflects the current consensus of the Group. This will also be published as an Internet-Draft.

The Group also discussed the current status of various publicly-available internal Trouble Ticket systems:

- CONCERT Trouble Ticket System, ftp.concert.net:dist/tickets/* (uses publicly-available database package).
- Help Desk Management System, ftp.delmarva.com:pub/hdms/HDMS.tar.Z (requires Unify database package).
- JVNC's NETLOG Trouble Ticket System, ftp.jvnc.net:pub/netlog-tt.tar.Z (requires no database package).
- NEARnet Trouble Ticket System, nic.near.net: pub/nearnet-ticket-system-v1.3.tar (requires Informix database package).
- UCP plans to have new drafts to discuss before the next IETF.

Attendees

Vikas Aggarwal	`vikas@jvnc.net`
David Conklin	`conklin@jvnc.net`

Tom Easterday	tom@cic.net
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Elise Gerich	epg@merit.edu
Eugene Hastings	hastings@psc.edu
Michael Laufer	mlaufer@bbn.com
Tony Li	tli@cisco.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Daniel Long	long@nic.near.net
Kim Long	klong@sura.net
Tom Sandoski	tom@concert.net
Kaj Tesink	kaj@cc.bellcore.com
Jack Waters	waters@sura.net
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov
Yung-Chao Yu	yy@qsun.att.com
Paul Zawada	Zawada@ncsa.uiuc.edu

2.6 Routing Area

Director(s):

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun

Border Gateway Working Group (BGP

The BGP Working Group agreed to make minor editorial changes to the BGP-4 draft and reissue it with the proposed changes. The BGP-4 MIB will be amended to remove potential conflicts with BGP-3 MIB.

The Group met with SIP/IPAE Working Group members to discuss how BGP-4 could be modified to accommodate SIP/IPAE inter-domain routing. The BGP Working Group recommended to the SIP/IPAE Group that it is preferable to not change BGP-4 to support SIP/IPAE routing, rather it is preferable to explore the possibility of using IDRP to carry SIP/IPAE routing information.

Inter-Domain Policy Routing Working Group (IDPR)

The IDPR Working Group met for two sessions on Monday. The Chair gave a status report, describing the GATED development effort, DNS modifications, IDPR MIB, and the scheduled Internet pilot installation. Following the status report, the IDPR Working Group focussed on the following issues:

1. How to get a software development team assembled which would enable the GATED effort to continue, independent of funding from individual sponsors.

2. What IDPR issues to pursue during the next year. For example, hierarchical addressing and host specification of source policy will likely be requested in the near future.

3. Conduct a survey of the policy needs of the Internet community, including the needs of transit service providers and the needs of users, to determine how IDPR meets these project needs.

4. Multicasting in the policy routing environment. The Group has a strawman proposal for this and intends to produce an Internet-Draft this spring, describing the proposal.

IP Over Large Public Data Networks (IPLPDN)

The Working Group discussed a revision to RFC 1294 "Multiprotocol over Frame Relay". The Group agreed to the following changes:

- Specification of the padding of encapsulated protocols to a two byte address.
- Discussion of remote bridging and clarification of bridge PDU encapsulation.
- Inclusion of connection oriented protocols in an appendix.

After an email review, the Draft will be submitted to the IESG as a Draft Standard.

The Group discussed RFC 1356 on Multiprotocol over X.25 and agreed to collect information on implementation experience and to then submit it the IESG as a Draft Standard.

The "Directed ARP" draft was reviewed again by the Group and their previous plan to recommend this as a Proposed Standard was reversed. The Working Group will recommend to the authors that it be published as an Experimental Protocol.

Work on IP over Circuit ISDN progressed with the following decisions:

- X.25 was approved as an additional encapsulation protocol.
- Multi-link transport was added. This will also be discussed with the PPP Extensions Working Group.

In addition, the Working Group will write an informational RFC describing the IEEE 802.6i Draft Standard on remote bridging over SMDS.

Open Shortest Path First (OSPF)

The Working Group started a last call for four documents:

1. Updated OSPF V2 specification (backward-compatible bug fixes to RFC 1247)
2. Updated OSPF MIB
3. OSPF NSSA area option
4. OSPF Trap MIB

The Working Group expects to submit these in the next month or two for publication as RFCs. There was continued discussion of a proposed user/implementation guide for OSPF over Frame relay networks (no protocol changes are actually required). The Group came to general agreement on a strategy for gracefully dealing with OSPF database overflow.

Mobile IP Working Group (MOBILEIP)

The Mobile IP Working Group met twice. The first session consisted of presentations by Pierre Dupont on current packet radio offerings from Motorola, by Tatsuya Ohnishi on a new mobile IP proposal from Matsushita, by Fumio Teraoka on recent changes to the Sony mobile IP proposal, and by Charlie Perkins of IBM on a more general architecture and terminology for mobility support in connectionless network protocols.

The second session consisted of a talk by John Ioannidis on security and authentication extensions to the Columbia mobile IP proposal, and discussion of a number of new topics, (including: agreement to submit all mobile IP proposals as Internet-Drafts, identification of multicast issues that arise with mobile hosts, the value of location-independent "EIDs" for supporting mobility, and the possible specification of a general "destination has moved" ICMP message to reduce the incidence of "triangular routing" in some mobile IP schemes).

Multicast OSPF Working Group (MOSPF)

The MOSPF Working Group decided to submit the Multicast OSPF Extensions draft as a Proposed Standard. They also decided to submit as a Proposed Standard a short document specifying how to map IP multicast addresses to a Token Ring Functional Address. The Group discussed possible designs of a multicast traceroute facility. Greg Minshall and Steve Deering agreed to write a concrete proposal.

Source Demand Routing Protocol BOF (SDRP)

The SDRP BOF decided to form a Working Group. A Charter for the Group was proposed, discussed and accepted. Work priorities were established and a schedule of work was generated. SDRP's capabilities and future goals were discussed. The mailing list will be established as sdrp-request@caldera.usc.edu.

Virtual Circuit (VC) Routing BOF (VCROUT)

The Virtual Circuit Routing BOF met to discuss adapting internet routing protocols for use in virtual circuit networks. The scope of the vc-routing work was covered including:

- Virtual Circuit-routing architecture
- Protocols to implement the architecture
- Class of service
- Multicast
- Private logical networks
- Internal switch addressing
- Switch to network address (ie IP) translation.
- VC setup protocol

The Group concluded that the Virtual Circuit routing work should focus on using existing protocols and existing standards work to the extent possible.

The Group reviewed and discussed several items from a previous meeting that was held at NET during INTEROP week. These included the following:

Scale	The IGP should scale up to 1000 switches.
Switch Addressing	The switch ID being a 32-bit address that can either be an IP address or be mapped into an IP address.
IGP	OSPF/RMP was discussed as a possibility for the IGP.

Inter-Domain Routing	IDRP may be a good choice for this solution.
Signaling	Use Q.93B (ATM)/Q.933 (FR) as the base for virtual circuit management with potentially a few modifications. The Group will attempt to define an inter-switch call set up protocol that will include an option to do inter-domain and intra-domain source route call setup.
Class Of Service	Discussed the potential of three classes of service: VC (PVC and SVC), Datagram, "Soft VC" (similar to VC but would not have the strict quality guarantees provided by VC.)
Address Resolution	There were a few issues considered that fall under the heading of address resolution including "network level" address to switch/port address resolution.

The Group decided that there was sufficient interest to form an IETF working group to continue this work.

CURRENT MEETING REPORT

Minutes of the New Internet Routing and Addressing Architecture BOF (NIMROD)

Report not submitted. Attendance was not taken.

Nimrod Overview

- Nimrod is 3 things so far:
  - Some fundamental principles of routing in very large networks
  - A broad-brush architecture using these principles, and general large system design rules
  - A deployment plan for this architecture
- Nimrod will be:
  - A set of protocols and algorithms (based on IDPR field experience) to implement the new architecture

Namespaces

- Address:
  - a structured name for a network attachment point
  - the structure is used by the routing
- DNS names:
  - a structured human usable name for a host, etc
  - the structure is used to facilitate the distribution and lookup

A new routing architecture should add a new name space to the packet layer: the "endpoint indentifier" (EID).

- An EID is:
  - Approximately, one end of a TCP connection
  - Can be hosts, mobile processes, etc.
  - A "fate-sharing" region
  - An "end-end" entity

Tradeoffs in Large Scale Routing

- In any large network, you have to discard detailed routing information, or incur the cost of massive amounts of potentially unneeded detail.
- Discarded routing information means non-optimal routes, i.e. costs for traffic taking extra hops, or failure to find a policy acceptable route which does exist but is not shown.
- The challenge of routing in very large networks is managing this choice of costs.

Addresses and Topology

- A hierarchical addressing structure is isomorphic to a topology.
- The hierarchy is chosen to minimize the sum of the previous two costs, for a given topology.
- As the topology changes over time, the cost (i.e. the sum of the two costs of large scale routing) is no longer minimal.
- In the long run, you have to change addresses. This should be dynamic, and semi-automatic.

Architectural Principles for Nimrod

- Maximize the lifetime of the design by minimizing the mandatory, system wide part of the architecture
- Maximize the lifetime of the design by making the design as flexible as possible
- Maximize the robustness of the design

Nimrod Fundamentals

- Link state (i.e. map distribution):
  - Allows maximally flexible policy routing
  - Interacts well with resource allocation systems
- Source routed (usually via flow setup)
  - Routing algorithm is *not* part of the spec
  - Allows policy routing (allows complete source control over packet forwarding path)
  - Interacts well with flexible abstraction mechanisms
  - Interacts well with resource allocation systems
  - Hop by hop forwarding requires *global consistency*

So far, identical to IDPR.

Nimrod Fundamentals Continued

- Flexible abstraction techniques (hierarchical)
  - Compression + Thinning ⟹ Abstraction
  - Abstraction is needed for practicality, but causes non-optimal routes
  - Nimrod allows users to pick their own balance between these two costs
  - Abstraction algorithms (incoming and outgoing) are not part of the spec

Reflections on the Architecture

- Note what is *not* part of the architecture:
  - Routing algorithm
  - Abstraction algorithm
  - LAC algorithm
- Key advantages:
  - Limits amount of work to be done.
  - Allows future improvement and easy deployment of new algorithms.
  - Reduces scope for errors and bad design by minimizing global mechanisms.
- What *is* part of the architecture:
  - Method of representing topology.
  - Method of distributing topology information.
  - Method of setting up flows.

Internet Evolution

The Internet needs several things:

- A new routing and addressing architecture. A routing and addressing architecture includes:
  - a system of naming networks, interfaces, etc
  - a way of exchanging information as to where these named things are
  - a way of computing routes between sets of endpoints
  - a way of causing traffic to take these routes
- A security architecture.
- A resource allocation architecture.
- Accounting.

We have a good idea how to do the former; the latter are still in progress.

Routing Architectures

A *packet format* is *not* a *routing architecture*.

- Nimrod is a routing architecture.
  - IDPR is a field trial of most of the Nimrod principles
  - A proposal has been submitted to design and implement Nimrod protocols
- IDRP is a routing architecture/protocol.
- Unified is a routing architecture/protocol.
- TUBA is primarily a *packet format*.
- IPAE is a *packet format*.
- PIP is a *packet format*.

Do We Need a New Packet Format Now?

Van Jacobson's Thesis: A new packet format now is a bad idea.

- Why?
  - Too many missing pieces (resource allocation, etc)
  - Substrate changing to rapidly (ATM, etc)

  Can we avoid a new packet format for a while? Don't we need a new packet format to deploy new addresses?
- No.
  - No new packet format is needed to deploy a new routing architecture.
- How?
  - The existing "address" fields in the IP header can be reinterpreted as EID fields.
  - This allows deployment of a new routing architecture without doing a new packet format at the same time.

New Packet Format?

Deploying a new packet format is a major task, as is deploying a new routing architecture.

Chiappa's Conjecture:

- To do both at once, in a larger operational network, is *folly*.

Conclusion:

- Deploy the new routing and addressing architecture as an adjunct to the existing IP layer, by reinterpreting the existing source and destination fields in the packet.
- After that is done, *then* we do a new packet format, once we understand resource allocation, etc.

CURRENT MEETING REPORT

Reported by Deborah Estrin/USC and Tony Li/cisco

Minutes of the Source Demand Routing Protocol BOF (SDRP)

The SDRP BOf began wit the introducion of the Charter and was followed by discussion on tasks and priorities.

Charter

The SDRP Working Group is chartered to specify, standardize and promote the use of SDRP (Source Demand Routing Protocol) as an interdomain routing protocol capability in conjunction with IDRP and BGP interdomain routing protocols. The purpose of SDRP (pronounced "esdrip", to go with "eyedrip", which is how IDRP is pronounced) is to support source-initiated selection of interdomain routes, to complement the intermediate node selection provided by BGP/IDRP.

Tasks and Issues

- Packet forwarding and control message format and protocol for IP. A draft specification is currently an Internet-Draft. It is under revision and the authors eagerly await more comments on the current or future drafts.

- Configuration and Usage. This is the highest priority after the draft specification in order to demonstrate how SDRP can be used to achieve concrete objectives. Work should begin on this document immediately.

- Internal SDRP. Several attendees commented that they would like to have SDRP functionality within domains and for that reason an item, intra-domain, or "Internal SDRP" (ISDRP) was added to the list.

- Extensions to BGP and IDRP. As mentioned in the Internet-Draft there are a few extensions to BGP/IDRP needed to support SDRP. These must be detailed and documented.

- MIBs. As needed for any protocol.

- Deployment. The Group will develop both an experimentation and deployment plan.

- Information Distribution. The initial versions of SDRP will use only information from IDRP/BGP and configured information as the basis for constructing source routes. The Group will develop mechanisms for distribution/acquisition of information to allow SDRP to construct a richer set of routes.

The Group also outlined a proposed timeline for Documents and Prototypes:

- Standards
  1. SDRP v1 Forwarding and Control Specification - January 1993
  2. MIB - March 1993
  3. Usage and Configuration - March 1993
  4. Internal SDRP - March 1993
  5. BGP/IDRP Extensions - June 1993
  6. Information Distribution - June 1993

- Prototypes
  1. SDRP v1 for IP v4 - March 1993

- Experimental Documents
  1. Draft Specification for Multicast - November 1993
  2. Draft Specification for Route Setup - November 1993

Attendees

David Arneson	arneson@ctron.com
Jules Aronson	aronson@nlm.nih.gov
Nagaraj Arunkumar	nak@3com.com
Fred Baker	fbaker@acc.com
Jim Barnes	barnes@xylogics.com
William Barns	barns@gateway.mitre.org
David Bolen	db31@ans.net
Scott Brim	Scott_Brim@cornell.edu
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ken Carlberg	Carlberg@cseic.saic.com
Charles Carvalho	charles@acc.com
Robert Ching	rching@nat.com
Richard Colella	colella@osi.ncsl.nist.gov
Dave Cullerot	cullerot@ctron.com
Michael Davis	mad@spirit.clearpoint.com
Steve Deering	deering@parc.xerox.com
Tim Dixon	dixon@rare.nl
Kurt Dobbins	dobbins@ctron.com
Jacques Dugast	dugast@issy.cnet.fr
Deborah Estrin	estrin@usc.edu
William Fink	bill@wizard.gsfc.nasa.gov
Shoji Fukutomi	fuku@furukawa.co.jp
Vince Fuller	vaf@stanford.edu
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com

Ken Hayward	crm57d@bnr.ca
Dwight Jamieson	djamies@bnr.ca
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
Merike Kaeo	merike@alw.nih.gov
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Peter Lothberg	roll@stupi.se
Gary Malkin	gmalkin@xylogics.com
Tracy Mallory	tracym@3com.com
Donald Merritt	don@brl.mil
David Meyer	meyer@oregon.uoregon.edu
Robert Moose	rmoose@gateway.mitre.org
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Laura Pate	pate@gateway.mitre.org
Yakov Rekhter	yakov@watson.ibm.com
April Richstein	abm@tycho.ncsc.mil
John Scudder	jgs@merit.edu
Frank Solensky	solensky@andr.ub.com
Michael St. Johns	stjohns@darpa.mil
Terrance Sullivan	terrys@newbridge.com
Morton Taragin	vsmorty@weizmann.weizmann.ac.il
Kannan Varadhan	kannan@oar.net
Warren Vik	wmv@i88.isc.com
Curtis Villamizar	curtis@ans.net
Luanne Waul	luanne@wwtc.timeplex.com
Cathy Wittbrodt	cjw@nersc.gov
Stephen Wolff	steve@nsf.gov
Robert Woodburn	woody@sparta.com
Honda Wu	honda@nat.com
Jeff Young	young@alw.nih.gov

CURRENT MEETING REPORT

Reported by Marco E. Sosa/Bellcore

Minutes of the Virtual Circuit Routing BOF (VCROUT)

The meeting was organized as a BOF to determine if there was enough interest in forming an IETF working group to develop a standard routing algorithm for choosing paths for virtual connection set-up in a Frame Relay or ATM network.

Rob Coltun gave a presentation covering the issues discussed at a previous meeting. Audience feedback/discussion on the scope of issues addressed was encouraged and received. (So much so that the Group didn't even get through the summary of the previous meeting's outcomes.) The talk included:

- Scope of the proposed working group.
- Initial set of issues/plan of attack discussed at previous meeting.

Juha Heinanen presented his proposal for integrated routing using NSAP addressing.

After much discussion the Group decided that there was interest in forming an IETF working group.

A mailing list has already been set up at vc-routing@gated.cornell.edu. Please use vc-routing-request@gated.cornell.edu to join.

The following topics were addressed:

- Rob and Marco Sosa (Bellcore) will co-Chair the Working Group and edit the resulting specification (standards-track RFC?).

- Rob and Marco will co-author an (informational) Internet-Draft, before the next IETF, which will outline the envisioned scope and architecture for which the resulting specification will apply, including what we think are the major issues to resolve. This will incorporate any input received from the mailing list.

- The resulting specification should be based on existing standards/implementation agreements. OSPF/RMP is one of the starting points for our proposal (RMP is an OSPF-based protocol used in multi-switch SMDS networks).

- The work of this Group should not duplicate the work that is being done in other standard groups/forums.

- The initial focus of the Group should be within a single routing domain (Autonomous System) with up to 1000 switches. Inter-domain routing issues will also be looked into, possibly based on IDRP.

- The Group will specify the routing protocol used between switches in the routing domain. A switch may in fact be distributed, running proprietary protocols internally.

- What address resolution is performed by attached end-systems (and how), is a hot topic that cannot be decided by this Group. However, it may be appropriate to specify a means of resolving the addresses received within a call set-up message into a network internal switch/ port identifier to be used by the routing protocol.

- Q.933 (FR)/Q.93B (ATM) signaling will be the base for virtual circuit management work. The Group will look into extending these as needed to define an inter-switch call set-up protocol (algorithm and format) that will include an option for inter-domain and inter-domain (loose) source route call set-up.

- Alternative methods for the format of the switch identifiers were discussed. Using a private 32-bit address (which could be an IP address) to ID a switch or using variable-length public NSAPs to identify end stations attached to a switch and routing based on these were the two main options.

- Assignment of link metrics was discussed. There seemed to be agreement that metrics should be settable through administration. There is also a proposal to define default metrics to provide a consistent metric all switches could handle (as metrics may change dynamically). The proposal is to look into using the SMDS RMP default metrics as the base. Modifications are needed to base the metric on "unreserved effective BW" rather than capacity. Effective BW allows for over-reservation of links and leaves open possibility of a better traffic descriptor than peak rate or CIR. Also, "damping" procedures will be needed to make sure excessive LSAs are not sent (don't want new LSA every time a new connection is reserved over a link).

- Three different classes of service were discussed, VC, Datagram, and "Soft VC" (like VC but without strict QOS guarantees, possibly allowing simpler call rerouting).

The co-Chairs will create a draft Charter resulting from the two meetings.

Attendees

Cynthia Bagwell	cbagwell@gateway.mitre.org
Fred Baker	fbaker@acc.com
Ken Benstead	kbenstead@coral.com
Lou Berger	lberger@bbn.com
Shiraz Bhanji	bhanji@gateway.mitre.org
Edo Biagioni	esb@fore.com
Ken Carlberg	Carlberg@cseic.saic.com
Kay Chang	chang@chang.austin.ibm.com
Dilip Chatwani	dilip@synoptics.com

Dean Cheng	dean@sun2.retix.com
Chi Chong	cchong@synoptics.com
George Clapp	clapp@ameris.center.il.ameritech.com
Robert Cole	rgc@qsun.att.com
Osmund de Souza	osmund.desouza@att.com
Art Dertke	dertke@gateway.mitre.org
Jacques Dugast	dugast@issy.cnet.fr
Robert Enger	enger@reston.ans.net
Mike Goguen	goguen@src.dec.com
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Ken Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
David Husak	dave@synnet.com
David Jacobson	dnjake@vnet.ibm.com
Merike Kaeo	merike@alw.nih.gov
George Kajos	kajos@coral.com
Fong-Ching Liaw	fong@eng.sun.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Robin Littlefield	rlittlef@wellfleet.com
Andrew Malis	malis@bbn.com
Tracy Mallory	tracym@3com.com
Robert Moose	rmoose@gateway.mitre.org
Julianne Myers	jmyers@network.com
Erik Nordmark	nordmark@eng.sun.com
Bala Rajagopalan	braja@qsun.att.com
Jim Scott	scott@kali.enet.dec.com
Frank Solensky	solensky@andr.ub.com
Marco Sosa	mxs@sabre.bellcore.com
Brad Steinka	brad@microcom.com
Terrance Sullivan	terrys@newbridge.com
John Tavs	tavs@vnet.ibm.com
Dono van-Mierop	dono_van_mierop@3mail.3com.com
James Watt	james@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Guy Wells	guy2@uswest.com
Ian Wilson	ianw@spider.co.uk
Liang Wu	ltw99@bellcore.com

Agenda

- Scope Of Work
- Recap Of Last Meeting
- SVC & Datagrams
  - -- Address Resolution
  - -- Datagram Techniques

Scope Of Work

- Make Use Of Existing Standards
- Private And Public Networks
- Recommend Routing Architecture
  - -- Routing Protocols
  - -- Class Of Service
  - -- Multicast

- -- Multiple Logical Nets
- -- PVC/SVC/Datagram
- -- Internal Addressing
- -- Address Resolution Scheme For SVC

- Recommend VC Management Protocol
  - -- Interaction With Routing
  - -- NNI
  - -- UNI

Recap Of Last Meeting

- Scale <= 1000 Switches
- Node Address = 32 Bits
  - -- Can Be IP Address Or IP Address Mapped
  - -- Can Use IP Tools (Telnet/FTP/SNMP)
- IGP (OSPF/RMP)
- Inter-Domain Routing (IDRP?)

- Signaling
  - -- Q.93B/Q.933
  - -- Hop-By-Hop Or Source Route Call Setup
  - -- Inter-Domain Source Route?
- Class Of Service
  - -- VC/Soft VC/Datagram
  - -- Metric Is Based On Unused Bandwidth With Damping Factor

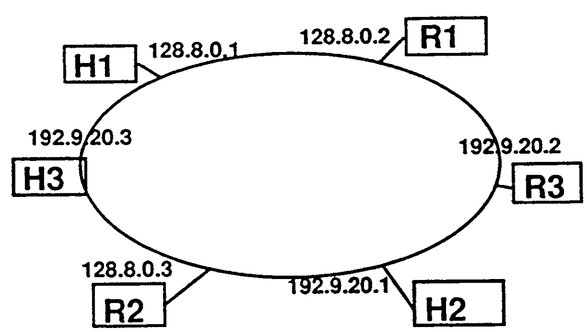

Address Binding Problem

- Need To Eliminate Notion Of Direct Subnet Attachment
- Techniques
  - -- Arp Broadcast
  - -- Inverse Arp Servers
  - -- BGP Route Servers
  - -- Shortcut Routing

Datagram Techiques

- Datgram VC Server(s)
- Multicast
- Fast Select
- Embeded Router

2.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):

Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:

General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu
Archive:

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Complete development of Version 2 of the Border Gateway Protocol (BGP).
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
Done	Develop a MIB for BGP Version 3.
Done	Work with the Security Area to enhance the provision for security in BGP.
Done	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.
Done	Post an Internet-Draft specifying multicast extensions to BGP.
Done	Post the specfication of BGP 4 as an Internet-Draft.
Done	Post an Internet-Draft specifying a MIB for BGP Version 4.
Jan 1993	Submit the multicast extensions to BGP to the IESG as a Proposed Standard.
Jan 1993	Submit the specification for BGP Version 4 to the IESG for consideration as a Proposed Standard.
Jan 1993	Submit the BGP Version 4 MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"IP Multicast Communications Using BGP", 08/26/1991, Scott Brim, Yakov Rekhter <draft-ietf-bgp-multicast-02.txt>

"A Border Gateway Protocol 4 (BGP-4)", 05/05/1992, Y. Rekhter, T. Li <draft-ietf-bgp-bgp4-04.txt>

"Definitions of Managed Objects for the Border Gateway Protocol (Version 4)", 09/01/1992, S. Willis, J. Burruss, J. Chu <draft-ietf-bgp-mibv4-01.txt>

"BGP4 OSPF Interaction", 09/15/1992, K. Varadhan <draft-ietf-bgp-bgp4ospf-interact-00.txt>

"Application of the Border Gateway Protocol in the Internet", 09/28/1992, Y. Rekhter, P. Gross <draft-ietf-bgp-application-01.txt>

Request For Comments:

RFC 1105 "Border Gateway Protocol BGP"

RFC 1163 "A Border Gateway Protocol (BGP)"

RFC 1164 "Application of the Border Gateway Protocol in the Internet"

RFC 1265 "BGP Protocol Analysis"

RFC 1266 "Experience with the BGP Protocol"

RFC 1267 "A Border Gateway Protocol 3 (BGP-3)"

RFC 1268 "Application of the Border Gateway Protocol in the Internet"

RFC 1269 "Definitions of Managed Objects for the Border Gateway Protocol (Version 3)"

RFC 1364 "BGP OSPF Interaction"

RFC 1397 "Default Route Advertisement In BGP2 And BGP3 Versions Of The Border Gateway Protocol"

RFC 1403 "BGP OSPF Interaction"

CURRENT MEETING REPORT

Reported by John Tavs/IBM

Minutes of the Border Gateway Protocol Working Group (BGP)

The BGP Working Group met twice on Monday, November 16, 1992, and discussed the following issues:

BGP-4 Editorial Change

The Working Group approved minor editorial change to the BGP-4 protocol specifications suggested by Yakov Rekhter. The revised text of the Internet-Draft will be posted shortly.

BGP MIB

It was pointed out that there are differences between MIB for BGP-3 and MIB for BGP-4. The Working Group decided to get a new OID for the bgpPathAttrEntry object. It is expected that BGP-3 MIB will eventually be deprecated.

BGP-4 Implementation

Dennis Ferguson (ANS) discussed his implementation of BGP-4 in GATED. It is expected that ANSNet will be running BGP-4 1Q93. A copy of the presentation should be available upon request from Dennis.

BGP and IPAE/SIP

The Working Group met jointly with several key members of the IPAE/SIP Working Group to discuss a possibility of supporting IPAE/SIP inter-domain routing with BGP. The Group suggested that inter-domain routing with IPAE/SIP should not be supported by BGP, but rather should be accommodated by IDRP. The rationale for this recommendation was based on the urgency of deploying BGP-4, unwillingness to change BGP-4 protocol specifications, and the overall consensus that IDRP is more suitable (than BGP) as an inter-domain routing protocol for IPAE/SIP.

Attendees

Roland Acra	acra@cisco.com
Nagaraj Arunkumar	nak@3com.com
Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
William Barns	barns@gateway.mitre.org
David Bolen	db31@ans.net
Scott Brim	Scott_Brim@cornell.edu
Jeffrey Burgan	jeff@nsipo.nasa.gov

Richard Colella	colella@osi.ncsl.nist.gov
Peter Ford	peter@goshawk.lanl.gov
Shoji Fukutomi	fuku@furukawa.co.jp
Vince Fuller	vaf@stanford.edu
Thomas Hacker	hacker@citi.umich.edu
Kevin Jackson	kmj@concord.com
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@server.af.mil
Dan Jordt	danj@nwnet.net
Akira Kato	kato@wide.sfc.keio.ac.jp
John Krawczyk	jkrawczy@wellfleet.com
Tony Li	tli@cisco.com
Peter Lothberg	roll@stupi.se
David Meyer	meyer@oregon.uoregon.edu
Robert Moose	rmoose@gateway.mitre.org
Dennis Morris	morrisd@imo-uvax.disa.mil
John Moy	jmoy@proteon.com
Julianne Myers	jmyers@network.com
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Peder Norgaard	pcn@tbit.dk
Laura Pate	pate@gateway.mitre.org
Michael Patton	map@bbn.com
April Richstein	abm@tycho.ncsc.mil
John Tavs	tavs@vnet.ibm.com
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Iain Wacey	cat@pluto.dss.com
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov

2.6.2 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: iplpdn@nri.reston.va.us
To Subscribe: iplpdn-request@nri.reston.va.us
Archive: /ietf.mail.archives/iplpdn.mail.archive

Description of Working Group:

The IP over Large Public Data Networks Working Group 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 internets implemented SMDS and on other PDNs.
TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.
Done	Establish priorities and dates of completion for documents.

Request For Comments:

RFC 1293	"Inverse Address Resolution Protocol"
RFC 1294	"Multiprotocol Interconnect over Frame Relay"
RFC 1315	"Management Information Base for Frame Relay DTEs"
RFC 1356	"Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode"

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

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

RFC 1294

The revised version of RFC 1294, "Multi-protocol over Frame Relay", was discussed and approved with the following further changes:

- Specification of padding of encapsulated protocols to a two byte address boundary assuming either two or four octet DLCI's.

- Discussion of Remote Bridging and clarification of Bridge PDU (BPDU) encapsulation.

- Inclusion of connection-oriented protocols in an appendix.

After review of the document via email, the Group agreed to submit this draft to the IESG as a Draft Standard.

RFC 1356

The Group discussed RFC 1356 on Multi-protocol over X.25 and agreed to gather documentation of interoperable implementations and then to submit the RFC to the IESG to advance the standard to Draft Standard status.

"Directed ARP"

The "Directed ARP" Internet-Draft was released by the Working Group at the previous IETF meeting for approval as an RFC on the standards track. The Group reconsidered and revised the request to approval as an Experimental RFC.

IP over Circuit ISDN

Work on IP over Circuit ISDN progressed with the following decisions:

- X.25 was approved as an additional encapsulation protocol.
- "Multi-link Transport", in which multiple B channels may be aggregated into single higher bandwidth channel, was approved as a feature. The Group agreed to work with the PPP Extensions Working Group (PPPEXT) on this topic.

Parameter Negotiation

The Group agreed to use the existing PPP specification for negotiation through SNAP encapsulation.

The PPPEXT Working Group agreed to take on the task of supporting a "compound packet" format, in which multiple PPP control packets are transmitted in a single data-link layer protocol.

IP over SMDS and CCITT I.364

The Group agreed to the release of an Informational RFC describing the IEEE802.6i draft standard on the support of remote bridging over SMDS.

Contributions were requested on the topic of address resolution

Attendees

George Abe	abe@infonet.com
Roland Acra	acra@cisco.com
Floyd Backes	backes@dsmail.enet.dec.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
Ken Benstead	kbenstead@coral.com
Daniel Brennan	dmb@teleoscom.com
Caralyn Brown	cbrown@wellfleet.com
Steve Bucey	sabucey@ns.pacbell.com
Brian Carpenter	brian@dxcern.cern.ch
Charles Carvalho	charles@acc.com
Kay Chang	chang@chang.austin.ibm.com
Dean Cheng	dean@sun2.retix.com
Chi Chong	cchong@synoptics.com
George Clapp	clapp@ameris.center.il.ameritech.com
Michael Collins	collinsms@es.net
Tracy Cox	tacox@sabre.bellcore.com
Bruce Davie	bsd@bellcore.com
Greg Dobrich	gdobrich@nalusda.gov
Jacques Dugast	dugast@issy.cnet.fr
Tom Easterday	tom@cic.net
James Forster	forster@cisco.com
Cliff Frost	cliff@cmsa.berkeley.edu
Shoji Fukutomi	fuku@furukawa.co.jp
Paul Griffiths	griff@chang.austin.ibm.com
Thomas Hacker	hacker@citi.umich.edu
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Ken Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
George Kajos	kajos@coral.com
Michael Laufer	mlaufer@bbn.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com

Andrew Malis	malis@bbn.com
Tracy Mallory	tracym@3com.com
Jun Matsukata	jm@eng.isas.ac.jp
Julianne Myers	jmyers@network.com
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Mo Nikain	mo@bss.com
Brad Parker	brad@fcr.com
David Piscitello	dave@sabre.bellcore.com
Jim Rees	jim.rees@umich.edu
Tony Richards	
Benny Rodrig	4373580@mcimail.com
Guenter Roeck	roeck@conware.de
Timothy Salo	tjs@msc.edu
Allyson Showalter	allyson@nsipo.arc.nasa.gov
William Simpson	Bill.Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Lansing Sloan	ljsloan@llnl.gov
Marco Sosa	mxs@sabre.bellcore.com
Brad Steinka	brad@microcom.com
Terrance Sullivan	terrys@newbridge.com
Iain Wacey	cat@pluto.dss.com
William Warner	warner@ohio.gov
James Watt	james@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Guy Wells	guy2@uswest.com
Peter Will	will@isi.edu
Kirk Williams	kirk@sbctri.sbc.com
Ian Wilson	ianw@spider.co.uk
Chin Yuan	cxyuan@pacbell.com

2.6.3 IP Routing for Wireless/Mobile Hosts (mobileip)

Charter

Chair(s):
Steve Deering, deering@parc.xerox.com

Mailing Lists:
General Discussion: mobile-ip@parc.xerox.com
To Subscribe: mobile-ip-request@parc.xerox.com
Archive: pub/mobile-ip/mail-archive@parcftp.xerox.com

Description of Working Group:

The Mobile IP Working Group is chartered to develop or adopt architectures and protocols to support mobility within the Internet. In the near-term, protocols for supporting transparent host "roaming" among different subnetworks and different media (e.g., LANs, dial-up links, and wireless communication channels) shall be developed and entered into the Internet Standards track. The work is expected to consist mainly of new and/or revised protocols at the (inter) network layer, but may also include proposed modifications to higher-layer protocols (e.g., transport or directory). However, it shall be a requirement that the proposed solutions allow mobile hosts to interoperate with existing Internet systems.

Longer term, the Group may address, to the extent not covered by the mobile host solutions, other types of internet mobility, such as mobile subnets (e.g., a local network within a vehicle), or mobile clusters of subnets (e.g., a collection of hosts, routers, and subnets within a large vehicle, like a ship or spacecraft, or a collection of wireless, mobile routers that provide a dynamically changing internet topology).

Goals and Milestones:

Done	Review and approve the Charter, making any changes deemed necessary.
Nov 1992	Post an Internet-Draft documenting the Mobile Hosts protocol.
Mar 1993	Review the Charter of the Mobile IP Working Group for additional work required to facilitate non-host mobility.
Mar 1993	Submit the Mobile Host Protocol to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Minutes of the IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

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

Attendees

George Abe	abe@infonet.com
Cengiz Alaettinoglu	ca@cs.umd.edu
Tony Bates	t.bates@nosc.ja.net
Pravin Bhagwat	pravin@cs.umd.edu
Shiraz Bhanji	bhanji@gateway.mitre.org
Rebecca Bostwick	bostwick@es.net
Robert Braden	braden@isi.edu
Gregory Bruell	gob@wellfleet.com
Matthew Busche	mtb@anchor.ho.att.com
Ken Carlberg	Carlberg@cseic.saic.com
Wo Chang	wchang@nist.gov
Dilip Chatwani	dilip@synoptics.com
Richard Cogger	rhx@cornell.cit.bitnet
Cathy Cunningham	cmc@microcom.com
Steve Deering	deering@parc.xerox.com
Barbara Denny	denny@erg.sri.com
Pierre Dupont	dupont@mdd.comm.mot.com
Eric Fleischman	ericf@act.boeing.com
Karen Frisa	karen.frisa@andrew.cmu.edu
James Galvin	galvin@tis.com
Robert Gilligan	Bob.Gilligan@eng.sun.com
Heather Gray	heather@zk3.dec.com
Paul Griffiths	griff@chang.austin.ibm.com
William Haggerty	haggerty@ctron.com
Neil Haller	nmh@thumper.bellcore.com
Patrick Hanel	hanel@yoyodyne.dco.ntc.nokia.com
Jonathan Hsu	brenda@penril.com
Christian Huitema	christian.huitema@sophia.inria.fr
John Ioannidis	ji@cs.columbia.edu
Ronald Jacoby	rj@sgi.com
Dwight Jamieson	djamies@bnr.ca
Phil Karn	karn@qualcomm.com
Paulina Knibbe	knibbe@cisco.com
Christopher Kolb	kolb@psi.com
John Larson	jlarson@parc.xerox.com

David LeRoy	dleroy@mitchell.cit.cornell.edu
David Lin	lind@janus-ccm.zenith.com
John Linn	linn@erlang.enet.dec.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Joshua Littlefield	josh@cayman.com
Dan Magorian	magorian@ni.umd.edu
Brian Marsh	marsh@mitl.com
Greg Minshall	minshall@wc.novell.com
Andy Nicholson	droid@cray.com
Erik Nordmark	nordmark@eng.sun.com
William Nowicki	nowicki@legato.com
Tatsuya Ohnishi	ohnishi@isl.mei.co.jp
Charles Perkins	perk@watson.ibm.com
Drew Perkins	ddp@andrew.cmu.edu
Mike Petry	petry@ni.umd.edu
Bala Rajagopalan	braja@qsun.att.com
Yakov Rekhter	yakov@watson.ibm.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel_rodrigues@att.com
Allan Rubens	acr@merit.edu
Henry Sanders	henrysa@microsoft.com
William Simpson	Bill.Simpson@um.cc.umich.edu
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@andr.ub.com
Michael St. Johns	stjohns@darpa.mil
Martha Steenstrup	msteenst@bbn.com
Brad Steinka	brad@microcom.com
Terrance Sullivan	terrys@newbridge.com
John Tavs	tavs@vnet.ibm.com
Fumio Teraoka	tera@csl.sony.co.jp
Marten Terpstra	marten@ripe.net
Thierry Turletti	turletti@sophia.inria.fr
John Veizades	veizades@apple.com
Daniel Wilson	dvw@bellcore.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Robert Woodburn	woody@sparta.com

2.6.4 ISIS for IP Internets (isis)

Charter

Chair(s):

Ross Callon, callon@bigfut.lkg.dec.com

Mailing Lists:

General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu
Archive:

Description of Working Group:

The IETF ISIS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

Done	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.
Jan 1993	Post a revision of the IS-IS as an Internet-Draft.
Mar 1993	Submit the revised IS-IS to the IESG as a Draft Standard.
Mar 1993	Submit the IS-IS MIB to the IESG as a Proposed Standard.

Internet-Drafts:

"Integrated IS-IS Management Information Base", 11/05/1991, Chris Gunner <draft-ietf-isis-mib-01.txt>

"Use of OSI IS-IS for Routing in TCP/IP and Multi-Protocol Environments", 01/11/1993, R. Callon <draft-ietf-isis-tcpip-00.txt, .ps>

Request For Comments:

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

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:

"An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-idpr-architecture-05.txt, .ps>

"Inter-Domain Policy Routing Protocol Specification: Version 1", 03/05/1991, M. Steenstrup <draft-ietf-idpr-specv1-02.txt, .ps>

"IDPR as a Proposed Standard", 04/28/1992, M. Steenstrup <draft-ietf-idpr-summary-00.txt, .ps>

Request For Comments:

RFC 1126	"Goals and functional requirements for inter-autonomous system routing"

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

Minutes of the Inter-Domain Policy Routing Working Group (IDPR)

At the November 1992 IETF meeting, the IDPR Working Group met for two consecutive sessions during the afternoon of Monday the 16th. The first session was a working meeting, while the second session was conducted as an overview for newcomers. The Group organized the first session as follows:

1. General Status Report

   - The IESG and IAB accepted the IDPR architecture and protocol documents as Proposed Standards in August 1992.

   - SRI is expecting to implement a large part of the IDPR MIB.

   - Rob Austein has designed the the DNS changes (address to domain identifier mapping queries and responses) required for IDPR.

   - The Group is seeking eager volunteers to produce an independent implementation of IDPR.

2. Gated version of IDPR

   Woody Woodburn of Sparta led the gated implementation effort, with additional participation by BBN. SRI is presently using the gated version of IDPR as the basis for policy routing in a network for one of their clients. Currently, SRI and BBN are taking responsibility for the IDPR gated software. The Group will eventually turn over the gated version of IDPR to Cornell, but before doing so, the Group needs to ensure that the software:

   - Conforms to the protocol specification.
   - Has clear and complete documentation.
   - Has been tuned to provide good performance.

   The Group welcomes all those interested in working on the IDPR gated software or in developing their own IDPR implementations. Please send a message to idpr-wg@bbn.com, if you're interested in working on IDPR software development.

3. Planned Internet Pilot Installation

   The target date is February 1993. The installation will initially include three backbone domains (NSFnet, NSInet, and TWBnet) and four source domains. The Group will exercise both source and transit policies. This will give transit service providers a chance to observe IDPR in action. The results of the pilot installation, including ease of use and management, general performance, and any problems encountered, will be published as an Internet-Draft.

4. Policy Survey

   The policies initially available with IDPR were extrapolated from a survey of federal agencies conducted several years ago. As IDPR moves from the testbed to the Internet, the Group should reevaluate the policy support provided. The Group intend to conduct a systematic survey of users and transit service providers to determine what types of source and transit policies are most desired. Results of this survey will be folded back into the policy offerings within IDPR. Anyone interested in helping to conduct the survey, please respond to the idpr-wg mailing list.

5. Multicast IDPR

   To provide multicast support in an internetwork in which policy is important, one cannot leave the forwarding decisions to intermediate routers. Rather multicast distribution should be defined by the source, just as it is for unicast distribution. To provide multicast support within IDPR, the Group plans to make the following modifications to IDPR:

   - All multicast groups, of which hosts within a domain are members, will be distributed as part of the existing routing information messages for the domain. This information will be used by a source to generate a multicast tree to other members of a multicast group.

   - The path identifier will carry a special multicast bit indicating that it is a multicast packet. All paths in a multicast tree will carry the same path identifier.

   - One or more path setup packets will be used to set up the multicast tree in sections or all at once. Each intermediate policy gateway in a path must keep track of all of the destination domains in the multicast tree that are reachable through the subtree of which it is the root.

   - The source will be notified through a teardown message when all hosts within a domain leave a the multicast group. The teardown will only affect the portion of the tree set up to that domain. A source should be able to initiate teardown to selected destinations or to all destinations within a multicast tree.

   - Intra-domain multicast, when available, will be used in conjunction with IDPR multicast.

In early 1993, the Group will distribute an Internet-Draft describing the initial version of multicast routing for IDPR.

Attendees

Cengiz Alaettinoglu	ca@cs.umd.edu
Ken Carlberg	Carlberg@cseic.saic.com
Dilip Chatwani	dilip@synoptics.com
Osmund de Souza	osmund.desouza@att.com
Barbara Denny	denny@erg.sri.com
Paul Griffiths	griff@chang.austin.ibm.com
John Hedderman	jjh@ans.net
Jonathan Hsu	brenda@penril.com
Dwight Jamieson	djamies@bnr.ca
Fong-Ching Liaw	fong@eng.sun.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Peder Norgaard	pcn@tbit.dk
John Scudder	jgs@merit.edu
William Simpson	Bill.Simpson@um.cc.umich.edu
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@andr.ub.com
Martha Steenstrup	msteenst@bbn.com
Robert Woodburn	woody@sparta.com

2.6.6 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):

Steve Deering, deering@parc.xerox.com

Mailing Lists:

General Discussion: mospf@comet.cit.cornell.edu
To Subscribe: mospf-request@comet.cit.cornell.edu
Archive:

Description of Working Group:

This Working Group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done	Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.
Done	Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.
Done	We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using 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.
Feb 1992	Submit the MOSPF Specification to the IESG as a Proposed Standard.

Internet-Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-02.txt, .ps>

CURRENT MEETING REPORT

Minutes of the Multicast Extensions to OSPF Working Group (MOSPF)

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

Attendees

Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Ken Benstead	kbenstead@coral.com
Scott Brim	Scott_Brim@cornell.edu
Matthew Busche	mtb@anchor.ho.att.com
Stephen Casner	casner@isi.edu
Robert Ching	rching@nat.com
Steve Deering	deering@parc.xerox.com
Dennis Ferguson	dennis@ans.net
Shoji Fukutomi	fuku@furukawa.co.jp
Jonathan Hsu	brenda@penril.com
Charley Kline	cvk@uiuc.edu
John Krawczyk	jkrawczy@wellfleet.com
David LeRoy	dleroy@mitchell.cit.cornell.edu
Tony Li	tli@cisco.com
Robin Littlefield	rlittlef@wellfleet.com
Kent Malave	kent@bach.austin.ibm.com
David Marlow	dmarlow@relay.nswc.navy.mil
Greg Minshall	minshall@wc.novell.com
Robert Moose	rmoose@gateway.mitre.org
John Moy	jmoy@proteon.com
Julianne Myers	jmyers@network.com
Erik Nordmark	nordmark@eng.sun.com
Laura Pate	pate@gateway.mitre.org
Jim Perchik	perchik@athena.mit.edu
Benny Rodrig	4373580@mcimail.com
John Tavs	tavs@vnet.ibm.com
Luanne Waul	luanne@wwtc.timeplex.com
Douglas Williams	dougw@ralvmg.vnet.ibm.com

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 to define backward compatible extensions and/or network adaptation layers to enable this protocol to be used in the TCP/IP internets. If any ISO standardization efforts overlap this area of work, it is intended that the ISO work will supersede the standards proposed by this Group.

1) IDRP for IP over IP document (standards track)

This document contains the appropriate adaptations of the IDRP protocol definition that enables it to be used as a protocol for exchange of "inter-autonomous system information" among routers to support forwarding of IP packets across multiple autonomous systems.

2) IDRP MIB document (standards track)

This document contains the MIB Definitions for IDRP. These MIB Definitions are done in two parts; IDRP General MIB, and IDRP for IP MIB. An appendix is planned; IDRP For IP GDMO

3) IDRP - OSPF Interactions (standards track)

This document will specify the interactions between IDRP and OSPF. This document will be based on a combination of BGP-OSPF interactions document and IDRP - ISIS interaction document.

4) IDRP for IP Usage document (standards track)

Most of the IDRP for IP Usage will reference the CIDR (Supernetting document) Internet Draft. Any additional terms or protocol definitions needed for IDRP for IP will also be specified here.

Goals and Milestones:

Done	IDRP for IP submitted for Internet-Draft.

Jun 1992 IDRP MIB document submitted for Internet-Draft.

Jun 1992 IDRP - OSPF Interactions document submitted for Internet-Draft.

Jun 1992 IDRP Usage document submitted for Internet-Draft.

Nov 1992 IDRP for IP submitted to the IESG for Proposed Standard.

Nov 1992 IDRP Usage document submitted to the IESG for Proposed Standard.

Nov 1992 IDPR MIB Submitted to the IESG for Proposed Standard.

Nov 1992 IDRP - OSPF Interactions document submitted to the IESG for Proposed Standard.

CURRENT MEETING REPORT

Reported by Sue Hares/Merit

Minutes of the OSI IDRP for IP over IP Working Group (IPIDRP)

John Krawczyk responded to Sue Hares' request for volunteers to work with her on MIB issues.

A merger of the IDRP-IP and BGP4 Working Group meetings was proposed. There was consensus that this would be a good idea. It was strongly emphasized by a number of those present that BGP is not going to suddenly disappear, simply that it would not be evolved past BGP4 (and that further interdomain protocol work would go into IDRP instead). No one is being forced to transition to IDRP!

Dave Katz gave a status report on the ISO process for IDRP. In summary:

- October 21st DIS ballot closes April 21, 1993. The DIS text is available from the FTP archive on merit.edu (/pub/iso/idrp.ps). U.S. ballot comments are pending an ANSI meeting in January.
- Specific issues expected to be addressed in the U.S. ballot comments are specification of protocol-specific data formats (this is somewhat CLNP-centric in the current DIS) and some state machine bashing.
- Barring substantial technical objections to the DIS an IS text will be issued at the September SC6 meeting in Korea.

Dave Katz gave a quick IDRP tutorial.

Discussion of IDRP-(other) Interactions.

In general, it seems that the IDRP interactions with other protocols are virtually the same as for BGP4, and it may even be appropriate to adopt entire BGP4 interaction documents (Sue will look into this).

IDRP-EGP Interaction

A concern was raised regarding exploding large aggregate routes into EGP. This isn't a bug per se, but text should be co-opted from the CIDR document to warn against doing this. A possible fix would be to allow for the configuration of filters in the IDRP speaker to constrain propagation of exploded information.

IDRP-BGP4

Bill Manning suggested that the Group "clean up the verbiage." Tony Li suggested that the document be revised to follow the BGP4-BGP3 document. Since there is little difference between the two, this can be a short document. The Group may specify that ARP should be used to get SNPA of next-hop router.

IDRP-BGPx

In general, in an integrated IP/OSI IDRP environment, there needs to be a way of mapping ASes back and forth into RDIs, and RDIs should be assigned with this in mind. The question arose of what to do with RDIs which don't map cleanly into an AS, e.g., RDCs. Options for this:

- Don't propagate non-translatable RDIs.
- Translate the RDI anyway (into what?) and set the external-info flag.
- Continue and syslog an error.

Another option (beyond the scope of BGP) would be to hand-configure a list of RDC memberships into each BGP-speaker which is a member of an RDC.

The usage document needs to clarify this issue. Dave Katz volunteered to work on this.

IDRP-OSPF

As for BGP4-OSPF. There was a great deal of debate on the subject of using multiple exits from an RD to reach another RD, i.e., load-sharing. Perhaps the IDRP-ISIS document needs to refine the mechanisms for this? IDRP for IPAE/SIP was discussed.

Miscellaneous Items

Sue presented a list of don't/do implements.

A glossary was proposed for the IDRP-IP document. The example used was "NLRI," which is never described in detail with respect to IP addresses.

MIB and Policy. Again, we can reference the analogous BGP4 document.

Additions to first Minutes from Dennis Ferguson

Issues to be resolved for BGP4 - IDRP:

1. What happens to unrecognized optional transitive attributes, both from BGP4->IDRP and from IDRP->BGP4. (Editor's Note (sh): The optional transitive attributes should be passed as unrecognized optional transitive attributes. The bgp4-idrp document needs to note this.)

2. Clarification of Dual IDRP functions

   If domain's external IP connectivity is congruent to it's external CLNP connectivity, then the domain may have a single RDI for IP and CLNP and may run an integrated IDRP. In all other cases, the domain has to have two different RDIs and must run Ships in the Night protocol. Ships in the Night protocol implies that one instances of the IDRP protocols runs over IP and a second instance runs over CLNP.

Attendees

Lee Alley	iadlra@asuvm.inre.asu.edu
Nagaraj Arunkumar	nak@3com.com
William Barns	barns@gateway.mitre.org
David Bolen	db31@ans.net
Vickie Brown	brown@osi540sn.gsfc.nasa.gov
Ross Callon	callon@bigfut.lkg.dec.com
Brian Carpenter	brian@dxcern.cern.ch
Richard Colella	colella@osi.ncsl.nist.gov
Michael Collins	collinsms@es.net
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@ans.net
Thomas Hacker	hacker@citi.umich.edu
Joel Halpern	jmh@network.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Robert Hinden	hinden@eng.sun.com
Kathleen Huber	khuber@bbn.com
David Jacobson	dnjake@vnet.ibm.com
Matthew Jonson	jonson@server.af.mil
Dave Katz	dkatz@cisco.com
Zbigniew Kielczewski	zbig@eicon.qc.ca
Charley Kline	cvk@uiuc.edu
John Krawczyk	jkrawczy@wellfleet.com
Padma Krishnaswamy	kri@sabre.bellcore.com
Tony Li	tli@cisco.com
Tracy Mallory	tracym@3com.com
Bill Manning	bmanning@sesqui.net
David Meyer	meyer@oregon.uoregon.edu
Robert Moose	rmoose@gateway.mitre.org
Dennis Morris	morrisd@imo-uvax.disa.mil
John Moy	jmoy@proteon.com
Andrew Partan	asp@uunet.uu.net
Laura Pate	pate@gateway.mitre.org
April Richstein	abm@tycho.ncsc.mil
Dallas Scott	scott@fluky.mitre.org
John Scudder	jgs@merit.edu
Paul Serice	serice@cos.com
Frank Solensky	solensky@andr.ub.com

Catherine Summers	cts@cos.com
John Tavs	tavs@vnet.ibm.com
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Iain Wacey	cat@pluto.dss.com
Luanne Waul	luanne@wwtc.timeplex.com
Cathy Wittbrodt	cjw@nersc.gov

IDRP

Interdomain Routing Protocol (IDRP) for CLNP and IP

Dave Katz
cisco Systems, Inc.

(email: dkatz@cisco.com)

Dave Katz, cisco Systems

IDRP

What is a Routing Domain?

The "cloud" in which an intradomain routing protocol (i.e., an IGP) is run

Intradomain routing routes within a routing domain (RD)

Interdomain routing routes between RDs

Dave Katz, cisco Systems

IDRP

A Routing Domain

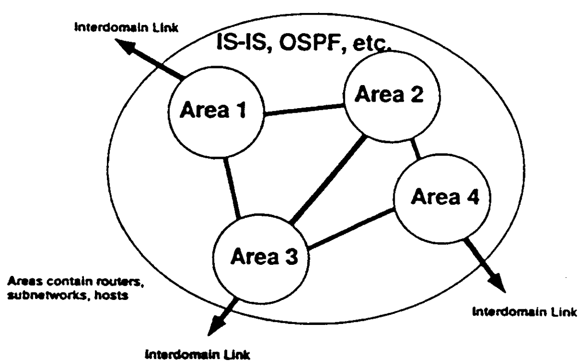

Dave Katz, cisco Systems

IDRP

Interdomain Routing

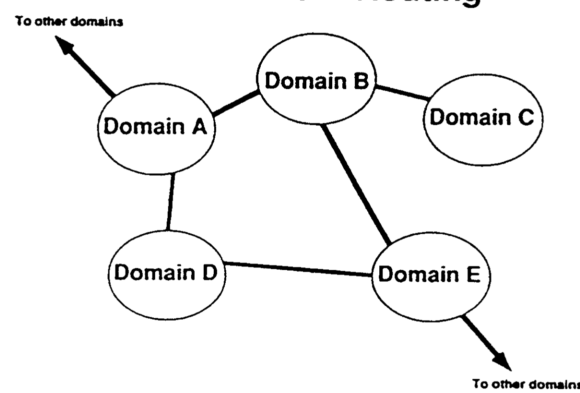

Dave Katz, cisco Systems

IDRP

Characteristics of Intradomain Routing

Single administrative structure

High degree of trust

Tightly coupled

"Best" path = shortest path

Routes between hosts, subnetworks, areas

Scales to thousands of hosts

Dave Katz, cisco Systems

IDRP

Characteristics of Interdomain Routing

Spans multiple administrations

Moderate level of trust

Loosely coupled

"Best" path = whatever policy dictates

Routes between routing domains, clusters of routing domains

Must scale to route to the entire universe

Dave Katz, cisco Systems

IDRP

Interdomain Routing Protocol (IDRP)

Scalable interdomain routing for OSI, IP, etc.
Originally derived from BGP3 in 1990
Now a Draft International Standard
Full International Standard expected in mid 1993
Several implementations underway

Dave Katz, cisco Systems

IDRP

IDRP Mechanisms

Runs over connectionless network layer (CLNP, IP, ...)
Private transport protocol (NL protocol independent)
Only changes are propagated
Digital signature in all packets for integrity, authentication

Runs in Border Intermediate Systems (BISs) at RD edge
Internal IDRP used to form single "Distributed BIS"
Decoupled from Intradomain routing

Routes carry a path of RD Indentifiers (RDIs)
Policy based on route selection/filtering, not in protocol

Dave Katz, cisco Systems

IDRP

Border Intermediate Systems

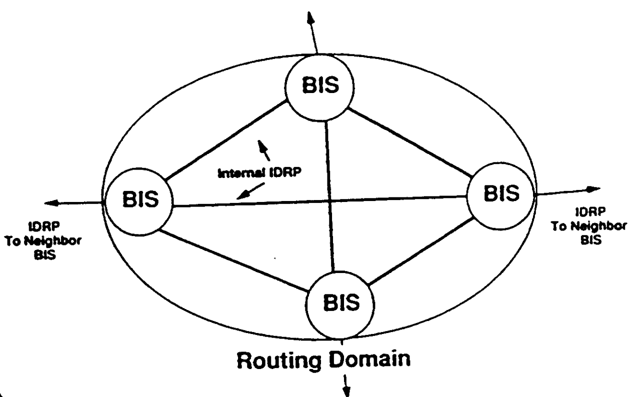

Dave Katz, cisco Systems

IDRP

Route Propagation

Reject routes with own RDI in path (loop suppression)
Add local RDI to path
Assign degree of preference to each route (policy)
Pick routes with highest preference value
Choose neighbors to which selected routes will be propagated (based on policy)

Dave Katz, cisco Systems

IDRP

Route Propagation

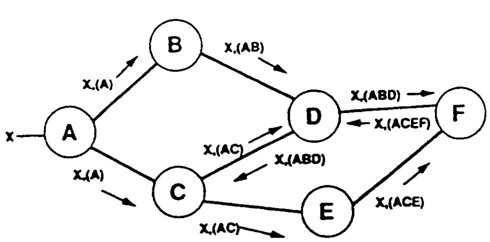

Rejected routes not shown
(for clarity)

Dave Katz, cisco Systems

IDRP

Scalability

What gets big? List of destinations
Reduce by aggregation:
- Use shorter prefixes to represent larger set of destinations (requires hierarchical address assignment)
- Merge path information (form RD Sets)

Dave Katz, cisco Systems

Route Aggregation

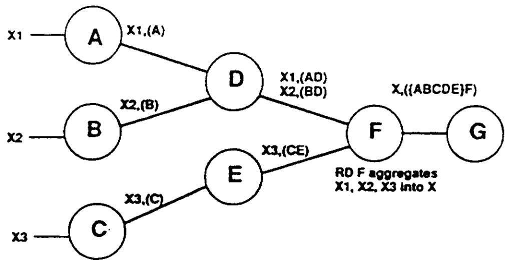

Scalability

- Route aggregation results in large paths (set union of all RDs in the universe)
- Still not sufficiently scalable
- What to do?

Routing Domain Confederations

- Routing domain confederation (RDC) is a grouping of RDs (a "superdomain")
- From the outside, an RDC looks like a single RD
- Path list of RDs within an RDC becomes a single RDI

Routing Domain Confederations

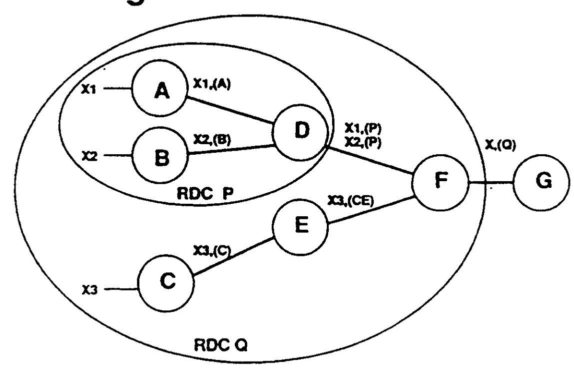

Routing Domain Confederations

- Loose coupling of RDs
- Minimal administrative agreements necessary
- No coupling of RD policies, except:
  - Path cannot exit/reenter RDC
- Enforced by protocol (automatic)

Multiprotocol Support

- IDRP mechanisms are protocol-independent
- Protocol-dependent stuff is opaque to IDRP:
  - Reachability information (addresses)
  - Next Hop address
- This information qualified by protocol ID
- Can support all protocols identified by NLPID, LSAP, or Ethertype

IDRP

Integrated or Ships in the Night?

Yes.

Dave Katz, cisco Systems

IDRP

Integrated Multiprotocol IDRP

Routing information for multiple protocols passed over single IDRP connection

RD boundaries are congruent for all protocols

IDRP itself can run over any of the protocols being routed (or even another one...)

Fits together nicely with integrated intradomain routing (e.g., Integrated IS-IS), but not necessary

Dave Katz, cisco Systems

IDRP

Ships in the Night Model

Routing independent between protocols

RD boundaries need not be congruent

One IDRP connection per protocol

Connections demultiplexed by NL protocol

RDIs must be different

Dave Katz, cisco Systems

IDRP

IDRP for IP

First application of IDRP outside of OSI

Specification is within the purview of the IETF

No changes to the base IDRP specification necessary

Mailing list:

idrp-for-ip-request@merit.edu

Dave Katz, cisco Systems

IDRP

Conclusion

IDRP supports sufficient policy control for the vast majority of user requirements (anything else requires substantially different technology)

IDRP has the capability to scale to arbitrarily large internetworks (if they are deployed intelligently)

IDRP can provide this functionality while supporting multiple protocols--don't continually reinvent the wheel

IDRP is a counterargument to the conventional wisdom concerning the International Standards process

Text available via anonymous FTP:

merit.edu:pub/iso/dis10747.ps.Z

Dave Katz, cisco Systems

2.6.8 Open Shortest Path First IGP (ospf)

Charter

Chair(s):

Mike Petry, `petry@ni.umd.edu`
John Moy, `jmoy@proteon.com`

Mailing Lists:

General Discussion: `ospfigp@trantor.umd.edu`
To Subscribe: `ospfigp-request@trantor.umd.edu`
Archive:

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

Done	Design the routing protocol, and write its specification.
Done	Develop multiple implementations, and test against each other.
Done	Obtain performance data for the protocol.
Done	Make changes to the specification (if necessary) and publish the protocol as a Draft Standard RFC.
TBD	Gather operational experience with the OSPF protocol and submit the document as a Standard.

Internet-Drafts:

"OSPF Version 2 Traps", 07/23/1991, Rob Coltun <draft-ietf-ospf-trapmib-01.txt>

"The OSPF NSSA Option", 10/13/1992, R. Coltun, V. Fuller <draft-ietf-ospf-nssa-option-00.txt>

"OSPF Version 2 Management Information Base", 11/03/1992, F. Baker, R. Coltun <draft-ietf-ospf-mib-00.txt>

"OSPF Version 2", 11/11/1992, J. Moy <draft-ietf-ospf-version2-00.txt, .ps>

Request For Comments:

RFC 1131	"OSPF specification"
RFC 1245	"OSPF Protocol Analysis"
RFC 1246	"Experience with the OSPF Protocol"
RFC 1247	"OSPF Version 2"
RFC 1248	"OSPF Version 2 Management Information Base"
RFC 1252	"OSPF Version 2 Management Information Base"
RFC 1253	"OSPF Version 2 Management Information Base"

CURRENT MEETING REPORT

Reported by John Moy/Proteon

Minutes of the Open Shortest Path First IGP Working Group (OSPF)

The meeting began with some administrative details. For those that had missed the announcement, it was mentioned that RFCs 1370/1371 had been published, officially making OSPF the recommended IGP for the TCP/IP Internet. Next, an informal poll was taken concerning attendance at next year's Amsterdam IETF. Based on that poll, the OSPF Working Group will probably not meet at that IETF. Lastly, there was some discussion of maintenance of the OSPF mailing list (trantor.umd.edu needs a forwarding record for ospf-request).

The last call for comments was started on four documents. These documents have been stable for some time, and the Group hopes to have them issued as RFCs before the next IETF (the latest versions had been issued as Internet-Drafts before the meeting):

- The OSPF V2 specification.
- The OSPF MIB.
- The NSSA option.
- The OSPF Trap MIB.

A number of issues were brought up by Robert Ching, on behalf of the OSPF Forum. Most were just requests for technical clarification (addressed in the meeting, but omitted from these notes in the interest of brevity). There was desire for an OSPF/RIP transition document (any volunteers?). Also, there was some confusion over the way OSPF represented serial lines. As John Moy explained, they are represented in router-LSAs as a direct connection to a neighbor, with each neighboring router advertising the other's serial line address as a (stub) host route. This encourages pings to a serial line address to actually traverse the serial line. However, two other representations are also possible: each neighboring router advertising its own address as a host route, or each neighboring router advertising a stub route to a subnet that has been allocated to the serial line. It was pointed out that the latter representation had the problem that traffic addressed to a non- existent host on the serial line had a tendency to loop until its TTL expired.

Osmund deSouza outlined a proposed usage document for OSPF over Frame relay. Requiring no protocol changes, this document would allow a Frame relay network to be configured as an arbitrary collection of NBMA networks, numbered and unnumbered serial lines.

Tom Pusateri presented his document on running IP multicast over 802.5 networks. A functional address (03-00-00-20-00-00) has been allocated, and Tom's document mandates that the token-ring address be configurable as either the all-ones broadcast MAC address (current practice), the new functional address or a group address (for possible future definition when token ring controller support is available). This document should soon be published as an RFC.

John Moy led a discussion of his proposal for how to deal with OSPF database overflow. It was decided to:

1. Exempt default routes from the limit calculation.

2. Automatically regenerate routes that have been earlier flushed due to database overflow (this regeneration will be done after some random interval between 1 minute and a configurable upper bound, with an option to completely disable the regeneration) and,

3. Set the LSA limit (which must be the same through all routers) through SNMP. Hopefully the Group will have a document describing this in detail next meeting.

At the end of the meeting, mention was made of two possible new work items and possible subjects for the next meeting.

1. A scheme to use the OSPF tag field and a new LSA type to replace IBGP and,
2. A new authentication type using something like MD5.

Attendees

Fred Baker	fbaker@acc.com
Ken Benstead	kbenstead@coral.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Dean Cheng	dean@sun2.retix.com
Robert Ching	rching@nat.com
Rob Coltun	rcoltun@ni.umd.edu
Osmund de Souza	osmund.desouza@att.com
Vince Fuller	vaf@stanford.edu
William Haggerty	haggerty@ctron.com
Jonathan Hsu	brenda@penril.com
Akira Kato	kato@wide.sfc.keio.ac.jp
David LeRoy	dleroy@mitchell.cit.cornell.edu
Tony Li	tli@cisco.com
Olli-Pekka Lintula	olli-pekka.lintula@ntc.nokia.com
Robin Littlefield	rlittlef@wellfleet.com
Kent Malave	kent@bach.austin.ibm.com
Jun Matsukata	jm@eng.isas.ac.jp
David Meyer	meyer@oregon.uoregon.edu
Douglas Miller	dmm@telebit.com
John Moy	jmoy@proteon.com
Julianne Myers	jmyers@network.com
Laura Pate	pate@gateway.mitre.org
Thomas Pusateri	pusateri@cs.duke.edu

Manoel Rodrigues	manoel_rodrigues@att.com
Paul Serice	serice@cos.com
Erik Sherk	sherk@sura.net
Roy Spitzer	roy.spitzer@sprint.com
John Tavs	tavs@vnet.ibm.com
Paul Traina	pst@cisco.com
Iain Wacey	cat@pluto.dss.com
James Watt	james@newbridge.com
Luanne Waul	luanne@wwtc.timeplex.com
Douglas Williams	dougw@ralvmg.vnet.ibm.com
Linda Winkler	lwinkler@anl.gov

OSPF
WG
Agenda

- Administrative
  - RFCs 1370/1371
  - Amsterdam
  - mailing list ospf@gated.cornell.edu
- Document review
  - OSPF spec
  - OSPF MIB
  - NSSAs
  - Trap MIB
- OSPF over Frame relay
- Database overflow
- 802.5 multicast
- New items
  - IBGP alternative
  - Better authentication
- more OSPF MIB

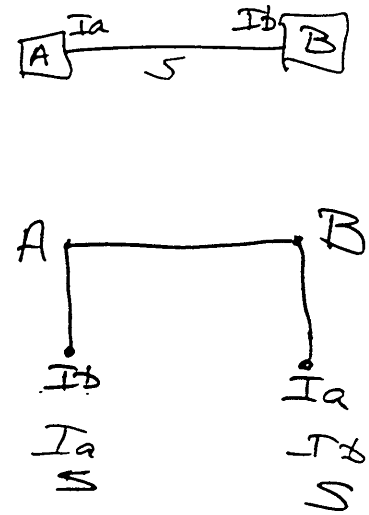

OSPF
Over
Frame-Relay

* FR networks defined as nbma
  - full mesh connectivity not always possible
* Define/implement a more flexible model.
  - nbma
  - point-to-point links
  - combinations
* Guidelines for implementation/use
  - informational RFC
  - no protocol changes

1/4

Alternatives

* nbma recommended for small networks
  - use current model
* Point-to-point PVCs configured as numbered links
* Point-to-point PVCs configured as un-numbered links

2/4

Numbered Links

* Multiple IP addresses on a FR interface
* Need variable length subnetting for efficiency
* DLCI↔IP address binding
  - Inverse ARP
  - ARP
  - Configure
* OSPF interface↔IP address binding

3/4

Un-Numbered Links

* Don't need multiple IP addresses
* OSPF needs a handle
  - ospfAddressLessIf → ifIndex
* Define P2P DLCI as distinct (OSPF?) interface.
  - create distinct ifEntry
  - ifSpecific→frDlcmiEntry
* Creates DLCI- OSPF interface binding

4/4

3 possible mappings for IP multicast to Token-Ring MAC address

1. All rings broadcast address (currently in use)
2. Assigned Token-Ring Functional Address
3. Assigned Hardware Group Addresses

   01-00-5e-xx-xx-xx

- all entities on a physical ring must agree
- Group Address preferred but not possible at this time
- Functional Address recommended
- Broadcast required for backward compatibility

Database overflow
choices

3, 5, 7

1) For which types of advertisements?

2) Reoriginate after some interval? random 0-10 min

3) How to set limits config, minimum 1 minute

- [illegible]

*Manually → default=0 no limit

[illegible]

4) Always Keep defaults

2.6.9 RIP Version II (ripv2)

Charter

Chair(s):
Gary Malkin, gmalkin@xylogics.com

Mailing Lists:
General Discussion: ietf-rip@xylogics.com
To Subscribe: ietf-rip-request@xylogics.com
Archive: xylogics.com:gmalkin/rip/rip-arc

Description of Working Group:

The RIP Version 2 Working Group is chartered to expand the RIP protocol, as defined in RFC 1058. The expansion will include the addition of subnet masks to the routing entries. The expansion may also include authentication, AS numbers, next hop address, MTU, or linkspeed. Since all routing protocols are required to have a MIB, one will be defined. The primary issue is the maintainance of backwards compatibility, which must be preserved.

The purpose of improving RIP is to make a simple, widely available protocol more useful. It is not intended that RIP-II be used in places where OSPF would be far better suited.

Goals and Milestones:

TBD	Given successful implementation experience, advancement of RIP-II to Draft Standard. Submission of MIB into the standards track.
TBD	Final meeting to achieve closure on any pending issues.
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.

Internet Drafts:

"RIP Version 2 Carrying Additional Information", 08/14/1991, Gary Malkin <draft-malkin-rip-05.txt>

"RIP Version 2 MIB Extension", 04/09/1992, Gary Malkin, Fred Baker <draft-ietf-ripv2-mibext-03.txt>

"RIP Version 2 Protocol Analysis", 08/14/1992, G. Malkin <draft-ietf-ripv2-analysis-00.txt>

2.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker/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 work for the SAAG consists of a set of formal work items. These work items correspond to four types of activities.

1. Working groups within the IETF Security Area. These are marked as "Security."

2. Working groups in allied organizations that function as part of the IETF Security Area. These are marked either "PSRG" for the Privacy and Security Research Group, or "TSIG" for the Trusted Systems Interoperability Group.

3. Security relevant developments within working groups in areas other than Security. These are marked according to the relevant area, viz., Applications, Internet Services, Management, OSI, Operations, Routing, Standards, or User Services.

4. Internal SAAG work items. These are topics which do not merit the creation of a formal working group but which do need some level of attention. These are assigned to a SAAG member and followed for one or more SAAG meetings. These are marked as "SAAG".

The following is the status of each of the currently open work items under the Security Area.

Authorization and Access Control (AAC)

To develop distributed authorization mechanisms and propose a standard access control application programmer's interface.

A proposed working group Charter submitted prior to the meeting was discussed. No changes to the text were proposed at the meeting. After a couple weeks to provide opportunity for suggestions through the mailing list, the Charter will be submitted to the IESG for approval and the creation of a working Group. There were differences in opinion with respect to how quickly the Group should proceed on standardizing distributed authorization mechanisms with many feeling it is still premature and others feeling we need to move quickly. All agreed that work on a local access control API should proceed. An early strawman for an access control API was distributed and discussed. A more detailed proposal will be presented prior to the March IETF.

Internet Protocol Security Protocol (IPSEC)

To develop mechanisms to protect client protocols of IP by providing support for authentication, integrity, access control, and confidentiality security services.

A BOF met for the second time at the Washington, DC IETF. A Charter has been drafted to be submitted to the IESG for approval to create a working group to continue work in this area. *Editor's Note (md): This Group has subsequently been approved as an official Working Group and a Charter has been included in these Proceedings.*

Commercial Internet Protocol Security Option (CIPSO)

To define an IP security option that can be used to pass security information within and between security domains of the commercial, U.S. civilian and non-U.S. communities.

CIPSO meets principally under the auspices of the Trusted Systems Interoperability Group. The following Internet-Draft is available: draft-ietf-cipso-ipsecurity-01.txt

There is a revised draft circulating within the CIPSO Working Group that needs to be posted as an Internet-Draft. The name of the protocol has been changed to Common IP Security Option from Commercial IP Security Option.

Common Authentication Technology (CAT)

To provide security services to a range of IETF protocol callers in a manner which insulates those callers from the specifics of underlying cryptographic security mechanisms, enabling modular separation between protocol and security implementation activities.

The following Internet-Drafts are available.

- The DASS architecture: draft-ietf-cat-dass-00.txt, .ps
- The GSS-API base specification: draft-ietf-cat-genericsec-02.txt, .ps
- A specification for Kerberos Version 5, proposed for consideration as a CAT mechanism: draft-ietf-cat-kerberos-01.txt, .ps
- A companion document defining C language bindings: draft-ietf-cat-secservice-01.txt

The integration of Kerberos and DASS has stalled due to DECs stalling of the DASS effort.

The CAT Working Group met at the Washington IETF. The first discussion topic was the status of the Working Group's current Internet-Drafts: the Security Area Director will proceed with advancement of the GSS-API and GSS-API C Bindings drafts; it was also agreed that the Kerberos V5 Draft, with pending minor changes, should be recommended for advancement. Ted Ts'o gave a presentation suggesting that a number of candidate GSS-API customer protocols could be conveniently served with an overlay which transfers GSS-API tokens across a stream established by the caller and temporarily handed to the overlay; its implications were discussed and follow-on work is planned. Further work is also planned on integration of a non-disclosing password mechanism (despite recognized global naming limitation) under the GSS-API framework, and a clarification and verification of token tagging conventions. The session concluded with a short discussion of FTP security issues, led by Sam Sjogren.

Network Access Server Requirements (NASREQ)

To specify the requirements, especially security requirements, of network access servers.

A requirements document was available for review during the meeting. There was a fair amount of discussion about what was being protected. No consensus was reached.

Privacy Enhanced Mail (PEM)

To specify the inclusion of security services in Internet text-based mail messages, specifically message origin authentication, message integrity, message confidentiality, and optionally non-repudiation.

The successor to RFCs 1113, 1114, 1115, respectively, are:

- draft-ietf-pem-msgproc-02.txt
- draft-ietf-pem-keymgmt-01.txt
- draft-ietf-pem-algorithms-02.txt

The "FORMS" document (not previously published as an RFC) is available as an Internet-Draft: draft-ietf-pem-forms-01.txt

The Working Group decided to remove the optional use of DES-MAC due to the recently discovered vulnerabilities that result from improper use. With this change, the documents are ready to be submitted to the IESG for consideration as Proposed Standards.

The remainder of the time was consumed by a discussion of the integration of PEM and MIME. Internet-Draft, draft-ietf-pem-mime-00.txt, is available with the proposal that was discussed at the meeting. Discussion will continue on the mailing list and at the next meeting.

SNMP Security (SNMPSEC)

To revise the SNMP Security Protocol RFCs 1351, 1352, and 1353 according to the SNMP Version 2 work and implementation experience.

The Working Group's Charter has been revised to include tracking the SNMP Version 2 work so changes and enhancements can be included in the next version of the documents. The Working Group met twice at the DC IETF leaving two agenda items to be resolved on the mailing list. SAAG members are encouraged to review the revised documents sooner rather than later.

TCP Client Identity Protocol (IDENT)

To review a proposed update to RFC931, a specification for an authentication server. The current draft was approved by the Working Group and will be submitted to the IESG for consideration for publication.

Work currently being carried out in other Areas include:

Automated Internet Mailing List Services - Applications Area

To select and integrate security services into a list management system. In particular, authentication and privacy are identified priorities.

This work item was officially closed at the Washington, DC IETF in November 1992. No work was ever completed since the Working Group never got started.

Internet Message Extensions - Applications Area

To extend "822 E-Mail" to handle multi-media mail, including a data integrity service suitable for detecting accidental modifications to messages in transit.

It was decided not to provide an integrity service at this time. Ideally, PEM should subsume this requirement. This work item was officially closed at this meeting.

Network News Transport Protocol - Applications Area

To select and integrate security services into the specification of NNTP.

When it is ready, Elliot Lear will submit the draft transport document to the IESG for publication. The NNTP Group will then go dormant, though a reader Working Group may be started at the next IETF.

Ted Ts'o made a presentation of the concept of a CAT protocol subroutine that was enthusiastically received. The need for support for a non-disclosing password mechanism was re-iterated. This work item was officially closed at this meeting.

Network Printing Protocol - Applications Area

To select and integrate security services into this new Internet protocol.

A document was reviewed in March of 1992 but the Working Group appears to have disbanded. This work item was officially closed at the Washington, DC IETF, November 1992.

OSI Directory Services - Applications Area

To include the use of strong authentication in the various directory pilot projects.

This work item is stalled due to the lack of interest and resources among the participants in the Working Group to proceed.

Telnet Security Services - Applications Area

To enhance the TELNET protocol to include security services.

The following documents are available as Internet-Drafts:

- "Telnet Authentication Option" draft-ietf-telnet-authentication-04.txt
- "Telnet Authentication: Kerberos Version 4" draft-ietf-telnet-authker-v5-01.txt
- "Telnet Authentication: SPX" draft-ietf-telnet-authspx-00.txt
- "Telnet Encryption Option" draft-ietf-telnet-encryption-01.txt

The proposed encryption option is vulnerable to active attacks. It has been decided to merge it with the authentication option so it can be done securely. As a result, the authentication documents will be submitted to be published as Experimental Protocols. When the merge of the encryption option is complete, the document will be proposed for publication on the standards track.

Point-to-Point Protocol Authentication - Internet Area

To assist the Point-to-Point Protocol (PPP) Extensions Working Group in the enhancement of PPP to include an authentication service.

The PPP Authenticaion Protocol has been published as RFC 1334. This work item was officially closed during the Washington, DC IETF in November 1992.

Router Requirements - Internet Area

To identify and resolve security issues as needed by the Working Group.

The Router Requirements document was deleted and is stalled for reasons beyond our control.

Internet Accounting - Network Management Area

To select and integrate security services with the process of generating and collecting accounting data within the Internet. In particular, integrity and confidentiality are identified priorities.

The Working Group will be submitting their work on Internet Accounting to the IESG for consideration as a Proposed Standard soon. The security of the accounting mechanisms is not met directly by the accounting system, but instead is provided by SNMP, over which accounting data is transferred. As such, the security of the accounting mechanisms depend on the security of SNMP. The accounting documents specify which security services are needed, or may be needed, for the transfer of accounting data. This work item was officially closed at the Washington, DC IETF in November 1992.

Trusted Network File System - Transport and Services Area

To select and integrate security services into NFS.

The TNFS Working Group meets principally under the auspices of the Trusted Systems Interoperability Group. The following internet drafts are available: draft-ietf-tnfs-spec-00.txt, .ps

The current specification meets short-term needs but is vulnerable to covert attacks. It is currently targeted to "secure network" customers. Authentication and securing the communications channel are future work items.

Internet Anonymous FTP Archives - User Services Area

To establish guidelines for providing anonymous FTP services on the Internet.

The IAFA Working Group will be submitting several documents for consideration for publication by the IESG. Of the documents to be submitted, the guidelines for running an FTP site include security related information about how to set up an FTP site, the anonymous FTP account, and protections for files and directories. This document has been reviewed by several SAAG members and adequately covers the issues we identified. This work item was officially closed at the Washington, DC IETF in November 1992.

Authentication Requirements

To identify a set of authentication requirements to help guide the work of protocols requesting authentication services.

This work has been subsumed by various other groups including CAT, authorization and access control, and network access service requirements. It was officially closed at the Washington, DC IETF in November 1992.

Export Control Issues

To prepare an informational document that summarizes the export control issues. Appropriate points of contact within the various government agencies will also be included. The document will focus on U.S. policy, mentioning others where information is readily available.

Vint Cerf has begun writing a document. It is on hold pending changes to export rules to Eastern Europe.

A presentation was made at the Washington, DC IETF. An overview of the rules was presented with some time remaining for questions and answers.

Mobile IP Security

To address the security issues of using a mobile IP protocol.

A design for Mobile IP was presented and discussed; it was agreed that mobile IP does not add additional threats but exacerbates current problems. If there existed an IP security option security would not be an issue for this Working Group. A revised document will be made available before the next meeting.

Network Database Privacy

To consider the issues of database privacy and accuracy.

The NISI Working Group documents have some words about database privacy. This needs to be reviewed before the next meeting to determine if it is sufficient.

PEM and MIME Integration

To develop a specification that integrates PEM and MIME.

A draft specification has been published as an Internet-Draft: draft-ietf-pem-mime-00.txt

The discussions have been taking place under the auspices of the PEM Working Group. They will continue on the mailing list and at the next meeting.

Random Number Generation Issues

To identify a critical set of issues about random number generation that are important to the Internet.

Jeff Schiller will provide a summary of the salient sections of X9.17, plus a few things it does not say, by December 15th.

Routing Security Plan

To develop a plan for addressing the security requirements of routing.

No progress to report.

CURRENT MEETING REPORT

Reported by Clifford Neuman/USC

Minutes of the Authorization and Access Control BOF (AAC)

Agenda

1. Discuss strawman of Charter (to be distributed at meeting).
   (a) Discuss possible goals of the Working Group.
      - Common mechanism for specifying local access control information.
         - Discuss strawman.
      - Improving interoperability for distributed authorization mechanisms.
      - Others

   (b) Evaluate each in terms of whether sufficient experience exists, or whether the Group would be premature in its efforts.
      - Discuss approaches and alternatives.
      - Assign work item to prepare strawman.

   (c) Select achievable goals and prepare a timetable.

   (d) Agree on mechanics for preparation and approval of Charter.

2. Any other business.

Overview

The second meeting of the BOF on Authorization and Access Control met at the November IETF. The purpose of the BOF was to organize a Working Group to address authorization and access control issues for the Internet. The discussion was centered primarily around two issues:

1. Development of a Charter and Milestones for the Working Group.
2. Initial work to develop an application program interface (API) supporting authorization.

Though not discussed in depth at this meeting, the Group is also concerned with mechanisms for distributed authorization on the Internet.

Charter and Goals of the Group

A draft Charter for the Working Group was distributed at the meeting. The Charter had been sent to the mailing list a day earlier and was made available by FTP for remote participants at the meeting. The first goal of the Working Group will be to develop a common mechanism for specifying access control information that will work well with distributed authentication mechanisms that are becoming available. The Working Group will also examine evolving mechanisms and architectures for authorization in distributed systems and to establish criteria that enable interworking of confidence and trust across systems and to encourage the evolution of (or develop ourselves) credential formats that more readily allow support for or translation across multiple mechanisms.

A timetable for these deliverables was discussed. There seemed to be agreement that we should move rapidly toward developing a common mechanism for specifying access control information. Clifford Neuman will submit a draft API for discussion prior to the March IETF. By the July IETF we hope to have examples of its use for selected applications, and the goal is to submit the specification of the API to the IESG by next November.

There was considerably less agreement on the timetable for work on distributed authorization mechanisms. The original timetable was less specific for work on distributed authorization mechanisms, initially exploring the area, trying not to constrain evolving implementations until more experience is gained. Several attendees, in particular Steve Crocker and Bill Simpson, felt that we should develop our own protocol and credential formats before incompatible mechanisms arise. There was no resolution on this issue, and it will be discussed further at the next meeting. Piers McMahon will submit specifications for DCE authorization, particularly with respect to proposed enhancements from SESAME for DCE. Clifford Neuman will submit additional information on authorization through restricted proxies.

As part of the discussion, a question was raised about whether the output of the Group would be a protocol. Our work on the API will not result in a protocol, instead it will yield a common mechanism for making authorization decisions based on authentication information obtained through other protocols (application protocols, and authentication and authorization protocols). The work on distributed authorization mechanisms, however, would result in a protocol or at least a common credential format to be used by other protocols. Even before distributed authorization mechanisms are in place, the API together with existing authentication protocols (e.g., CAT), would allow the retrieval and evaluation of fine-grained access control information allowing access by specific principals not previously registered (in terms of having an account) on a server.

During discussion of the API, Piers McMahon suggested that the scope of the API should support the specification of delegated principal identifiers, though the mechanism for delegation would be the subject of subsequent work. Richard Graveman suggested that the mechanism should support the specification of groups. Mechanisms to certify membership in groups would be the subject of the distributed authorization work, but the specification of required group membership does belong in any access control list mechanism, and this should be part of the API.

Steve Lunt pointed out that the naming of principals is an important issue that must be addressed by the API. Steve Crocker pointed out that the need for a common naming mechanism is a problem that the IAB is aware of, but that we shouldn't expect such a mechanism to be in place soon, we must support the multiple existing mechanisms for now. John Linn pointed out that the GSSAPI exports names tagged with a type and provides a function to compare two names for equality, and that that mechanism may be sufficient for our needs.

Piers McMahon asked about the scope of our mechanism. Is it to be Internet specific, or is it to extend beyond the Internet? The answer was that the Group would like it to be universal, but to the extent that making it so adds complexity or hinders progress it should be restricted to the Internet. In any event, the Group will look at mechanisms and APIs developed in other contexts, including DCE and Posix.

Comments on the Charter should be sent to ietf-aac@isi.edu. After a couple weeks for discussion, the Charter will be submitted for approval by Steve Crocker (the Security Area Director) and the IAB.

Authentication Requirements

After discussion of the Charter, Neil Haller spoke about an Internet-Draft he and Randall Atkinson submitted on Internet Authentication Requirements. The draft discusses authentication requirements and guidelines for different applications. The mechanisms covered include simple password mechanisms, non-disclosing passwords, Kerberos, DASS, and CAT.

It is not clear which working group is best for discussion of this document. It was felt that in general this work item fits best under the Common Authentication Technology (CAT) Working Group and John Linn indicated his willingness to take it on as a separate work item for the CAT Group. Some issues, in particular how authentication requirements interact with authorization mechanisms used by particular applications (the login application was presented as an example) should be considered in this (AAC) Group.

Neil Haller did not receive many comments when the Internet-Draft was first submitted to the INET-AUTH list. He will resubmit it to the CAT list in hope that CAT will provide the input required.

Authorization and Access Control API

The next topic of discussion was an API for access control. A strawman outline was distributed and made available to remote participants. The strawman called for a function:

answer = check_acl(id,(object/acl/multiple_acl),operation)

and each input and output was discussed. The first item of discussion was the ID structure.

In the strawman:

> id = user and/or group identification from distributed authentication mechanisms and future authorization mechanisms. We should support the passing of multiple identifiers to support user and group and to support Access Control List (ACL) entries naming compound principals (i.e., two principals must be present to perform an operation).

John Linn suggested that perhaps there should be separate inputs for the clients identity and other authorization credentials. It was felt that this was a bad idea. It was resolved that there should be a single identifier if at all possible, that this identifier might be a GSSAPI security context, and that the security context might need to be extended to include addition information as required.

> object/acl/multiple_acl = a reference or identifier for the object to be accessed, or a reference to a specific access control list associated with the object. Multiple acl's might be necessary for example if an acl is associated with both a directory and an object within the directory.

The topic of discussion here was whether one names an object whose ACL is to be checked, or pass the ACL itself. In either case there is an abstraction violation. In one case the application must manage ACLs so that they may be passed to the API. In the other case, the code implementing the API many require knowledge about the application.

Steve Lunt suggested that each ACL should be named, and that the application would decide how to map the object into the name of the appropriate ACL. The name of the ACL might be simply the name of the object. If a system wide authorization database is shared by more than one application, it would be important to make sure that no name conflicts arise.

> operation = a list of those operations to be performed, or more precisely a list of those rights needed to perform the requested operation.

The issue here was whether the operation should be passed as input and checked by the API returning a yes/no answer, or whether the API should return the operations allowed and let the application decide. The resolution is that we should support two calls, one that returns the rights and one that checks them returning yes or no.

Sam Sjorgen suggested support for VMS/Tops-20 style enabling and disabling of capabilities during the checking of rights. Unfortunately, it is not clear how such a capability would work in a distributed environment. In particular, the rights that are enabled are simply those passed to the server, and checked by the API. Disabled rights would not be visible to the process checking for access.

answer = A yes/no response indicating whether the operation is allowed, and optionally a list of restrictions to be applied by the application. Applications that don't require or can't interpret restrictions in a response would not have an authorization database that provides them. Thus if you don't need this functionality, your ACL mechanism doesn't need to support it. If your ACL mechanism does return a restriction that the application can't understand the response will be treated as not authorized.

Discussion on this topic centered around the use of restrictions. Does the use of restrictions place too great a burden on the application to understand what they mean? Some restrictions, for example time of day, are relatively common and could be interpreted by the code implementing the API, but some are inherently application specific and could not be interpreted by the code implementing the API.

Bill Simpson raised the network access server as an example of an application that could use the API. He wants a mechanism that they can put in their boxes. Restrictions for the network access server might be an address mask restricting where a user can connect. John Linn asked what the objects are that are being protected. The answer is network addresses to which one can connect. Clifford Neuman pointed out that the restrictions allow one to specify ACLs for fewer objects. For the same fine-grained control without restrictions, one would specify an ACL for each address (or at least each subnet). With restrictions, one has a single ACL with an application restriction that provides finer grained control. Whether an application choose to use restrictions is a design decision, we should not make the decision for them.

Piers McMahon asked whether the Group had considered existing APIs for access control. Posix was looked at, but it is not suited to distributed principals not previously registered as users on a system. Piers asked if the OSF API had been considered for access control. The answer was no, since the Group was not aware of it. Conditioned on obtaining OSF approval to do so, Piers will submit a copy of the OSF access control API to the list.

To Proceed

- Comments on the Charter should be sent to ietf-aac@isi.edu.
- The Charter will be submitted for approval in the next few weeks.
- The ACL API will be refined. Discussion will take place on the ietf-aac mailing list.
- Addition information on authorization in DCE, ECMA, and using restricted proxies will be submitted to the list by Piers McMahon and Clifford Neuman.

Attendees

Vickie Brown	brown@osi540sn.gsfc.nasa.gov
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov

Barbara Fraser	byf@cert.org
Shari Galitzer	shari@mitre.org
Richard Graveman	rfg@ctt.bellcore.com
Thomas Hacker	hacker@citi.umich.edu
Neil Haller	nmh@thumper.bellcore.com
Ken Hirata	khirata@emulex.com
David Katinsky	dmk@rutgers.edu
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Clifford Neuman	bcn@isi.edu
Tu Nguyen	Nguyen1T@cc.ims.disa.mil
Rakesh Patel	patel@noc.rutgers.edu
Tim Seaver	tas@concert.net
William Simpson	Bill.Simpson@um.cc.umich.edu
Sam Sjogren	sjogren@tgv.com
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov

2.7.1 Commercial Internet Protocol Security Option (cipso)

Charter

Chair(s):

Ron Sharp, `rls@neptune.att.com`

Mailing Lists:

General Discussion: `cipso@wdl1.wdl.loral.com`
To Subscribe: `cipso-request@wdl1.wdl.loral.com`
Archive: `archive-server@wdl1.wdl.loral.com`

Description of Working Group:

The Commercial Internet Protocol Security Option Working Group is chartered to define an IP security option that can be used to pass security information within and between security domains. This new security option will be modular in design to provide developers with a single software environment which can support multiple security domains.

The CIPSO protocol will support a large number of security domains. New security domains will be registered with the Internet Assigned Numbers Authority (IANA) and will be available with minimal difficulty to all parties.

There is currently in progress another IP security option referred to as IPSO (RFC 1108). IPSO is designed to support the security labels used by the U.S. Department of Defense. CIPSO will be designed to provide labeling for the commercial, U.S. civilian and non-U.S. communities.

The Trusted Systems Interoperability Group (TSIG) has developed a document which defines a structure for the proposed CIPSO option. The Working Group will use this document as a foundation for developing an IETF CIPSO specification.

Goals and Milestones:

Ongoing	Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard.
Done	Review and approve the Charter for the IETF CIPSO Working Group. Review revised TSIG CIPSO Specification.
Done	Review outstanding comments/issues from mailing list. Continue work on specification and prepare it for submission as an Internet-Draft by the end of May.
Jul 1991	Review outstanding comments/issues from mailing list. The specification will be submitted to the IESG for consideration as a Proposed Standard.

Mar 1992 Submit specification to the IESG for consideration as a Draft Standard. There must be at least two interoperable implementations by this time.

Internet-Drafts:

"COMMERCIAL IP SECURITY OPTION (CIPSO 2.2)", 12/03/1991, Trusted Sys Interop. Group (TSIG)

<draft-ietf-cipso-ipsecurity-01.txt>

2.7.2 Common Authentication Technology (cat)

Charter

Chair(s):

John Linn, `linn@erlang.enet.dec.com`

Mailing Lists:

General Discussion: `cat-ietf@mit.edu`
To Subscribe: `cat-ietf-request@mit.edu`
Archive: `/cat-ietf/archive@bitsy.mit.edu`

Description of Working Group:

The goal of the Common Authentication Technology Working Group is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security-oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accommodating evolution of underlying technologies.

In support of these goals, the Working Group will pursue several interrelated tasks. We will work towards agreement on a common service interface allowing callers to invoke security services, and towards agreement on a common authentication token format, incorporating means to identify the mechanism type in conjunction with which authentication data elements should be interpreted. The CAT Working Group will also work towards agreements on suitable underlying mechanisms to implement security functions; two candidate architectures (Kerberos V5, based on secret-key technology and contributed by MIT, and X.509-based public-key Distributed Authentication Services being prepared for contribution by DEC) are under current consideration. The CAT Working Group will consult with other IETF working groups responsible for candidate caller protocols, pursuing and supporting design refinements as appropriate.

Goals and Milestones:

Ongoing	Progress Internet-Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms.

Done Preliminary BOF session at IETF meeting, discussions with Telnet and Network Printing Working Groups.

Done Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet-Draft process.

Done First IETF meeting as full Working Group: review Charter distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.

Oct 1991 Update mechanism-independent Internet-Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.

Nov 1991 Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.

Dec 1991 Submit service interface specification to RFC standards track.

Internet-Drafts:

"Generic Security Service Application Program Interface", 06/12/1991, John Linn <draft-ietf-cat-genericsec-03.txt, .ps>

"The Kerberos Network Authentication Service (V5)", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-01.txt, .ps>

"Generic Security Service API : C-bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-01.txt>

"Distributed Authentication Security Service", 11/04/1991, Charles Kaufman <draft-ietf-cat-dass-01.txt, .ps>

CURRENT MEETING REPORT

Reported by John Linn/DEC

Minutes of the Common Authentication Technology Working Group (CAT)

The CAT Working Group met for one session at the November 1992 IETF. Primary discussion topics were:

- Status of documents.
- Future work items and issues.
- Token representation and integration.
- FTP security.

Status of Documents

Steve Crocker stated his belief that the GSS-API and associated C bindings Internet-Drafts were ready for advancement to Proposed Standard RFCs, and that he would recommend this action shortly. The Kerberos V5 Internet-Draft is pending certain local and specific edits, and is to be included in the same advancement recommendation. Despite the fact that Kerberos is the only CAT technology visibly under active development and support at this time, it was still viewed as a desirable goal that applications use CAT/GSS-API rather than mechanism-specific interfaces so as to support future portability.

Future Work Items and Issues

Ted Ts'o led a discussion suggesting future work items and issues for CAT. He divided client applications for authentication services into four groups:

1. Datagram protocols, generally (e.g., SNMP) not viewed as a good fit for CAT, though better suitability to other connectionless protocols was considered as a possibility.

2. Store-and-forward protocols, also not viewed as a good fit for CAT.

3. Stream protocols (e.g., Telnet, rpc, lpr, NNTP), considered by Ted as the best-fitting candidates for CAT usage.

4. Multiple-stream protocols (e.g., FTP), with suitability not evaluated.

His list of thoughts on future work included [with editor's annotations in square brackets]:

- A need for better [more complete and fully-tested] GSS-API clients and mechanism implementations, e.g., to implement rlogin with less reliance on local or mechanism-specific routines in addition to GSS-API.

- Development of an easy-to-use layer overlaid atop GSS-API to embody token-passing, analogous to Kerberos's krb_sendauth.

Ted's list of issues and near-term action recommendations was as follows:

- Negotiation of mechanisms: recognized as important in the eventual term, but not needed in the near-term, given a presumption that GSS-API callers (in an environment with only a limited number of mechanisms in use) would not be burdened by a requirement to pass in explicit mechanism type specifiers.

- Strength ranking of mechanisms: as with negotiation, not needed at first.

- Naming: generalized translation between types not needed, but a canonical flat ASCII representation was desirable for ACLs, etc. Internationalization was recognized as an issue here, with a character set selection tag being requested. The long-requested and as-long-frustrated desire for a unifying Internet naming framework also arose as an issue here.

- Infrastructure requirements: given that many sites don't want to pay the prices attendant to use of Kerberos, SPX, or similar cryptographic mechanisms, peer-peer key/password exchange and non-disclosing password systems should be considered as CAT mechanisms. An issue here is the fact that such lower-function mechanisms don't generally authenticate principals in terms of global names; use of an interface facility [e.g., a name type tag] to distinguish local from global names is a partial approach to the issue. Many lower-function mechanisms do not yield session keys for per-message protection as a result of authentication, but mechanisms with this characteristic are accommodated with existing interface indicators.

Availability of a Kerberos V4 GSS-API implementation would be convenient; while some activities had been undertaken to this end in previous years, no complete implementation compatible with current specifications is known to exist. The GSS-API modules within the Kerberos V5 implementation (as of the re- cent Beta 2 release) have been unit tested, and code exists to support all calls, but have not been linked and tested with a sample client. Ted Ts'o indicated that he would like to coordinate with anyone interested in performing this testing, but that he cannot himself provide the resources needed to develop or carry out the tests in the near future; Steve Lunt expressed interest in this activity.

Token Representation and Integration

The present Kerberos V5 GSS-API implementation includes tagging facilities on its tokens, but (unsurprisingly, given the order of events) the tags do not include the object identifier recently assigned to Kerberos by the IANA and to be included in the upcoming revision to the Kerberos specification. As a goal, it was agreed desirable that applications into which authentication is being newly integrated should use OID-identified mechanism tags. It was noted that use of ASN.1 in tagging should be constrained (and, in the GSS-API appendix's recommendation, is constrained to X.509-DER) so that the use of a fully general ASN.1 parser is not required; further clarification on the encoding conventions and their processing requirements was requested.

Ted Ts'o suggested development of a generic "plug-in" authentication protocol layered on GSS-API, to be embedded within applications which are built over stream-oriented communications. NNTP was specifically cited as an example; Telnet (given the fact that it acts itself as a sophisticated stream manager and is oriented to transfer of data elements within options) was not considered as a customer for this proposed technology and would be more appropriately served by calling CAT directly. In the "plug-inx" approach, a stream would be established by the application and then handed over for use by the authentication protocol while authentication tokens were exchanged. Subsequent to token exchange, within which mechanism negotiation could also be incorporated, the stream could (optionally) either be handed back to the application or the application's communications could be encapsulated and thereby protected by the "plug-in" protocol. The ability to reinitiate the "plug-in" protocol on an already-authenticated stream, thereby accomplishing reauthentication, was requested in discussion and considered to be supportable.

The format of CAT tokens was not perceived as a particularly hard issue from the viewpoint of caller protocols; the prospect the token exchanges in the course of carrying out GSS-API continuation scenarios raises qualitatively different complexity to callers, which use of the "plug-in" could simplify. It was observed that existing mechanisms involve exchange of no more than two tokens, one from an initiator to a target and a second returned from the target to the initiator, and that perhaps the most likely scenario in which need for longer exchanges might arise would be design of a "negotiated" mechanism in which authentication elements were preceded by tokens transferred in order to establish a mechanism shared between peers.

FTP Security

At the end of the meeting, Sam Sjogren led a brief discussion on security for FTP, a topic for which he has established a discussion group. Interest exists in Kerberized FTP in order to eliminate transmission of cleartext passwords across networks. The FTP specification states that FTP's control connection "follows Telnet protocol", but is silent about use of Telnet options on the control connection and it was believed that at least most FTP implementations would not accept Telnet options on an FTP control port. The FTP specification also states that data elements in FTP commands are usually to be interpreted by humans, but informal communication with Jon Postel suggests that he would not oppose the inclusion of encoded data intended for machine interpretation (e.g., cryptographic authentication tokens) so long as the data elements' contents were properly specified. It was suggested that authentication information for an FTP control connection could be represented either through use of the Telnet authentication option (if Telnet options are found to be supported or easily supportable within FTP) or by direct calls to CAT and textual encoding of CAT tokens.

In addition to security on FTP's control connection, there was also interest in protecting the data connection, most efficiently in a block mode. Any such protection would need to be compatible with the variety of transfer modes supported within FTP.

Actions

Ted Ts'o plans further work on documenting the stream-oriented "plug-in" overlay.

Neil Haller plans further work on integrating a lower-function authentication mechanism, probably to be based on the S/key technology, under the GSS-API.

John Linn plans further work on documenting token encoding conventions and their attendant requirements.

Attendees

David Conklin	conklin@jvnc.net
Stephen Crocker	crocker@tis.com
Cathy Cunningham	cmc@microcom.com
Art Dertke	dertke@gateway.mitre.org
William Edison	
Richard Graveman	rfg@ctt.bellcore.com
Neil Haller	nmh@thumper.bellcore.com
Ken Hirata	khirata@emulex.com
Russell Housley	Housley.McLean_CSD@Xerox.Com
Frank Kastenholz	kasten@ftp.com
David Katinsky	dmk@rutgers.edu
John Kunze	jak@violet.berkeley.edu
Paul Lambert	paul_lambert@email.mot.com
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Mohammad Mirhakkak	mmirhakk@mitre.org
Clifford Neuman	bcn@isi.edu
Hilarie Orman	ho@cs.arizona.edu
Paul Sangster	sangster@ans.net
Sam Schaen	schaen@mitre.org
Sam Sjogren	sjogren@tgv.com
Tang Tang	tt@virginia.edu
John Vollbrecht	jrv@merit.edu
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Daniel Woycke	woycke@smiley.mitre.org

2.7.3 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:

Mar 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 Key Management Protocol to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT

Minutes of the Internet Protocol Security Protocol Working Group (IPSEC)

Report not submitted. Attendance was not taken. Please refer to the Security Area Report for a brief summary.

2.7.4 Network Access Server Requirements (nasreq)

Charter

Chair(s):

Allan Rubens, acr@merit.edu
John Vollbrecht, jrv@merit.edu

Mailing Lists:

General Discussion: auth-acct@merit.edu
To Subscribe: auth-acct-request@merit.edu
Archive:

Description of Working Group:

The Network Access Server Requirements Working Group has as its primary goal, to identify functions and services that should be present in IP Network Access Servers (NAS's) and to specify the standards that provide for these functions and services. The term "Network Access Server" is used instead of the more conventional term "Terminal Server" as it more accurately describes the functions of interest to this Group. A "Network Access Server" is a device that provides for the attachment of both traditional "dumb terminals" and terminal emulators as well as workstations, PC's or routers utilizing a serial line framing protocol such as PPP or SLIP. A NAS is viewed as a device that sits on the boundary of an IP network, providing serial line points of attachment to the network. A NAS is not necessarily a separate physical entity; for example, a host system supporting serial line attachments is viewed as providing NAS functionality and should abide by NAS requirements.

This Group will adopt (or define, if need be) a set of standard protocols to meet the needs of organizations providing network access. The immediate needs to be addressed by the Group are in the areas of authentication, authorization, and accounting (AAA). In general, this Group will select a set of existing standards as requirements for a NAS. If necessary, the Group will identify areas of need where internet standards don't already exist and new standardization efforts may be required.

Initially the Group will independently investigate the two cases of character and frame oriented access to the NAS. This investigation will be aimed at determining what work is being done, or needs to be done, in this and other working groups in order to be able to define the set of NAS requirements. While the ultimate goal of this Group is to produce a NAS Requirements document, it may be necessary to define standards as well. This initial investigation will help determine what the goals of this Group need to be. The Group will also work with appropriate Working Groups to define required NAS standards that fall into the areas of these other groups.

Goals and Milestones:

Done	NAS Requirements Document posted as an Internet-Draft.
Nov 1992	Post an Internet-Draft on Character oriented Authentication, Authorization, and Accounting(AAA).
Nov 1992	Post an Internet-Draft on frame oriented AAA requirements.
Nov 1993	Submit the NAS Requirements document to the IESG as a Proposed Standard.

Internet-Drafts:

"Network Access Server Proposed Requirements Document", 10/01/1992, J. Vollbrecht, A. Rubens, G. McGregor, L. Blunk, Richard Conto <draft-ietf-nasreq-nasrequirements-00.txt>

CURRENT MEETING REPORT

Reported by John Vollbrecht/Merit

Minutes of the Network Access Server Requirements Working Group (NASREQ)

The NASREQ Working Group met on Wednesday afternoon at the Washington IETF meeting. Allan Rubens chaired the meeting. Allan announced that John Vollbrecht would be acting as co-Chair for the Group. Note that the mail Group for this has a new alias <nas-req@merit.edu> in addition to the old <auth-acct@merit.edu>

Agenda

- Discuss the NAS Proposed Requirements Document (Internet-Draft).
- Go over Jesse Walker's comments on the Draft.
- Plan next steps.

Discussion of NAS Proposed Requirements Document

Copies of the Draft NAS requirements were available at the session. John Vollbrecht talked through the main points. A major change of focus between the Draft and the previous Draft is that the current Draft considers the NAS to be a router which supports temporary connections to a net rather than as a terminal server which also supports framed access.

Terminal support in a NAS (if available) is provided by a Character Stream client (e.g., Telnet) that converts the character stream to framed output. The output of the Character Stream client is then input to the router.

The major thrust of the NASREQ Working Group is to define support requirements for systems providing temporary connections to a network. The main requirements were seen to be:

1. (Mutual) authentication of NAS and user.

2. Per user configuration of ports on the NAS and/or per user authorization of user for network access.

3. Per user session record keeping.

Some discussion of the NAS model took place. Jeff Schiller asked if the ability to have character stream terminal sessions authenticate without sending passwords in the clear was being considered by this Working Group. The response was that so far this had been outside the area being considered, but perhaps could be included if standards for this are developed.

There was some question of the need to authenticate for access to the network at all. Presumably hosts and servers can demand authentication if they need to know who is using their system, and to monitor and control scarce resources (modems and phone lines). The response was that the NAS would authenticate in order to know who is using it. It is a special server that provides access to the network. Network providers use it to give their clients access to their network. A NAS may use the same or different authentication methods (and servers) as a file or print server.

A good deal of discussion of authorization and per user configuration took place. The issue of whether the NAS would screen access to other services on the network was discussed. The concept in the Draft document is that the NAS only controls NAS functions, and other hosts need to screen themselves. If one views what is required in NAS authorization as per user port configuration, then the concept becomes clearer. A user connects, gets authenticated, then has its port set up according to the user's preconfigured requirements.

The authentication and authorization must be supported by (possibly) remote servers. A set of NAS's would be able to be authorized by a set of authorization servers. Bill Simpson asked if this was aimed at ultimately supporting a situation where a user could connect to a NAS on one network and get authenticated and authorized by another (connected) network. Indeed this is one of the goals.

Some discussion of two approaches to multiple domain authorization took place. The first is a hierarchical approach where each NAS goes to a specific server (or backup). The server then talks with other servers if necessary to get authorization.

The second is to have the NAS contact different authentication and authorization servers itself. This might be driven by having the user identify the server for the NAS to use as part of the connection sequence. This could be useful where a number of sites have independent authentication and configuration/authorization server. Both methods should be investigated.

Authentication Issues

The NAS provides access to the network to which the authentication server is connected. The user and authentication server must communicate before the user is formally connected to the NAS. This requirement means that the NAS must provide a capability at connection time for this communication to happen. Two possible approaches were discussed.

1. The user-NAS dialog includes Pasword Authentication Protocol (PAP) or some other sequence that provides an id/password combination. The NAS would then take the password/id and go to an authentication server (e.g., Kerberos) on behalf of the user. It was pointed out that authentication will need to be done before the NAS knows the IP address of the user. This is because the IP address assignment may be based on the user id. It may be necessary to use a "temporary" ip address during the authentication phase.

2. The user-NAS dialog would use Challenge Handshake Authentication Protocol (CHAP). In this case the NAS does not receive the id/password, so the most it seems is possible would be to have it act as a CHAP forwarder. The NAS would forward messages between the user and a remote CHAP server.

In both these cases some additional issues need to be worked out. In the first case a question is whether the response from the authentication server will reach the user. The user would presumably get a "ticket" which it then passes to the NAS to request access. Alternatively the NAS could act as proxy for the user, which might be better in general since the user doesn't then need to support Kerberos or whatever authentication protocol is used.

In the second case the question is how does the NAS get informed of the result of the CHAP exchange if all it does is act as a forwarding agent. Clearly it will need to interpret some of the exchange as well as forward so that it will know if the authentication succeeded. It may be that the remote CHAP server will need to have extensions to the protocol defined to allow it to communicate with the NAS as well as the user.

Authorization/per User Configuration Issues

Per user configuration requires that the port to which a user is attached be configured from that user's predefined setup. For the general case the port could be configured with route filters, an IP or other protocol address, static routes, routing protocols supported, and anything else that is needed to configure a router port.

Authorization is implied by the configuration. Route filters act as restrictions, static routes are specific authorizations. In the discussion of how this fits in with the model of the Authorization and Access Control Group, Cliff Neuman suggested that this could be considered as a single Access Control List (ACL) for network access with restrictions (filters) to provide finer grain control. The alternative would be to have an ACL for each address or network reachable via the NAS - a potentially very large number for a NAS connected to the Internet.

Accounting Issues

It would be good to use the work done by the Internet Accounting (ACCT) Working Group as the basis for what is required in a NAS. A couple of issues need to be sorted out to be sure this is workable.

The ACCT Group does not seem to have a well described way to handle multiple sessions on the same port. This may be possible, but it needs to be worked out.

The method for collecting information being proposed by the ACCT Working Group is to use SNMP to query for information stored. The definition of MIB for multiple sessions per port needs to be clarified. It would seem reasonable to consider alternatives to SNMP (like an rpc) for passing accounting information to the remote collector.

Review of Jesse Walker's Comments

Allan went over a number of issues raised in Jesse's comments. The Group thought it was important to air these issues at this meeting to get input from others before making changes to the Requirements Draft. Also, since many of the issues raised deal with the question of focus for this Working Group, it was beneficial to raise these issues in order to solicit input from the Group on directions that should be taken. Jesse's message containing the issues and a response generated by John are available in the auth-acctg archives. The resolution of the raised issues will be incorporated in the next Requirements Draft, to the best of our ability. A few of the major points of contention are described briefly next.

One of the issues raised was that security ultimately hinges upon secure loading and configuration of the NAS itself and this is not an issue being addressed as yet. There was no consensus as to what to do about this problem. It is definitely not within the bounds of this Working Group to solve this problem, but we should incorporate any solution as a NAS requirement.

As far as Kerberos being sufficient to handle security, it may help but it doesn't completely solve the problem. As discussed above, it may be used with PAP, but it doesn't seem to be useful with CHAP.

The issue of how long an authentication is valid was said to be a matter of policy, not an issue of concern to this Group. This is probably true, but it brings up a related matter - the issue of how to deal with inactive PPP sessions tying up NAS resources. This needs further discussion.

The issue of a "reliable" transport mechanism for the collection of accounting information was brought up. It was explained that "reliable" was not intended to mean absolutely fail-safe, rather it meant that a best-effort mechanism was needed so that accounting/auditing information was not frequently lost. The NAS document will be modified to make this clear.

Date and timestamping of accounting needs to be optional as it requires a clock synchronization mechanism. Again, this is not really the case because the Group is only talking about times corresponding within "reasonable" limits. The document will be changed to clarify this.

The topic of Account limits was also discussed. One thing that was clear from this discussion was that this shouldn't just be written off as being beyond the scope of the Group or of being a policy matter - at least not without further discussion.

Plans for Future Action

We discussed a number of possible next steps. Allan and John agreed to clean up the NAS document and resubmit it as an Internet-Draft. The changes will reflect discussion of the document and of Walker's specific comments.

We would like to have more detailed requirements for how the NAS will do authentication. The PAP/Kerberos and CHAP/CHAP cases both should be defined in more detail. A number of people expressed some interest in this but no specific plan was made to do something.

The configuration/authorization issues need further work. A specific proposal for how to manage per user configuration is needed. No specific plan for this was initiated.

Finally, there needs to be some work with the ACCT Working Group to define how specific requirements for the NAS will fit. Some of the ACCT Group members showed a lot of interest in working on this, and John and Allan will follow up with them to come up with a proposal for inclusion in the NAS Requirements document.

Attendees

Jim Barnes	barnes@xylogics.com
Daniel Brennan	dmb@teleoscom.com
Vickie Brown	brown@osi540sn.gsfc.nasa.gov
J. Nevil Brownlee	nevil@aukuni.ac.uz
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Barbara Fraser	byf@cert.org
James Galvin	galvin@tis.com
Ken Hirata	khirata@emulex.com
Nat Howard	nrh@bellcore.com
Miriam Kadansky	mckadansky@eng.xyplex.com
John Linn	linn@erlang.enet.dec.com
Joshua Littlefield	josh@cayman.com
Steven Lunt	lunt@bellcore.com
Mohammad Mirhakkak	mmirhakk@mitre.org
Clifford Neuman	bcn@isi.edu
Hilarie Orman	ho@cs.arizona.edu
Brad Parker	brad@fcr.com
Drew Perkins	ddp@andrew.cmu.edu
Joseph Ramus	ramus@nersc.gov
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
Tim Seaver	tas@concert.net
William Simpson	Bill.Simpson@um.cc.umich.edu
Brad Steinka	brad@microcom.com
John Vollbrecht	jrv@merit.edu

2.7.5 Privacy-Enhanced Electronic Mail (pem)

Charter

Chair(s):
Stephen Kent, kent@bbn.com

Mailing Lists:
General Discussion: pem-dev@tis.com
To Subscribe: pem-dev-request@tis.com
Archive: pem-dev-request@tis.com

Description of Working Group:

PEM is the outgrowth of work by the Privacy and Security Research Group (PSRG) of the IRTF. At the heart of PEM is a set of procedures for transforming RFC 822 messages in such a fashion as to provide integrity, data origin authenticity, and optionally, confidentiality. PEM may be employed with either symmetric or asymmetric cryptographic key distribution mechanisms. Because the asymmetric (public-key) mechanisms are better suited to the large scale, heterogeneously administered environment characteristic of the Internet, to date only those mechanisms have been standardized. The standard form adopted by PEM is largely a profile of the CCITT X.509 (Directory Authentication Framework) recommendation.

PEM is defined by a series of documents. The first in the series defines the message processing procedures. The second defines the public-key certification system adopted for use with PEM. The third provides definitions and identifiers for various algorithms used by PEM. The fourth defines message formats and conventions for user registration, Certificate Revocation List (CRL) distribution, etc. (The first three of these were previously issued as RFCs 1113, 1114 and 1115. All documents have been revised and are being issed first as Internet Drafts.)

Goals and Milestones:

Done	Submit first, third, and fourth documents as Internet-Drafts.
Ongoing	Revise Proposed Standards and submit to IESG for consideration as Draft Standard, and repeat for consideration as Internet Standard.
Done	Submit second document as Internet-Draft.
Done	First IETF Working Group meeting to review Internet-Drafts.
Done	Submit revised Internet-Drafts based on comments received during Working Group meeting, from pem-dev mailing list, etc.

Nov 1991 Submit Internet-Drafts to IESG for consideration as Proposed Standards.

Internet-Drafts:

"Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", 03/26/1991, John Linn <draft-ietf-pem-msgproc-02.txt>

"Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", 07/17/1991, Steve Kent <draft-ietf-pem-keymgmt-01.txt>

"Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", 08/22/1991, David Balenson <draft-ietf-pem-algorithms-02.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services", 09/01/1992, B. Kaliski <draft-ietf-pem-forms-01.txt>

"MIME-PEM Interaction", 11/23/1992, S. Crocker, N. Freed, M. Rose <draft-ietf-pem-mime-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"

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

Minutes of the Privacy-Enhanced Mail Working Group (PEM)

A review of document status was provided by Steve Crocker, the Security Area Director. Three of the four documents are ready for progression, and the fourth (RFC 1115bis) needs to be edited to make it clear that additional algorithm suites will be published via new RFCs, but this is viewed as a minor edit and thus should not hold up progression of the documents. Steve indicated that the documents will be recommended for progression very soon, perhaps at the Friday IESG meeting.

RFC 1115bis also needs to be revised to remove use of Data Encryption Standard (DES) MAC as a message integrity code (MIC). Recent work has indicated that use of DES MAC is unsuitable with either symmetric or asymmetric key management algorithms, even in the limited contexts already defined in 1115bis. Only one party who might object to this removal of DES MAC was identified and he will be promptly notified of the planned change. here too, the change is considered minor as it involves removal of what is viewed as an option which was not expected to see much, if any, use.

The Working Group received a hardcopy handout of an Internet-Draft written by Steve Crocker, Ned Freed, and Marshall Rose. The Internet-Draft proposes an approach to integrating MIME and PEM.

Ned Freed presented the following approach to the Working Group and discussion ensued:

- There was agreement that the current processing description for submission should not be proscriptive, i.e., alternative user interface options for invoking PEM for a MIME message are permitted. Thus section 5.1 needs to be revised to avoid any implications that the pre-submission processing is a description of a user interface requirement. The goal here is to convey what needs to be done, but not to imply a required user interface form.

- It was suggested that additional formats could be defined in MIME to transport certificates, exclusive of their transport in the PEM header. This is not in conflict with the current proposal, but was generally regarded as a very useful addition.

- There was a discussion of what 5.1.2 in the Internet-Draft implies. The intent was that step would transform any input into MIME canonical form. Discussion explored the use of the new (as of last IETF meeting) PEM header field "Content-Domain" to represent the canonicalization performed by PEM. This field was intended to allow other than vanilla 822 canonicalization to be performed on the input to PEM, in an effort to avoid redundant encoding steps.

- It was suggested that the PEM MIC-ONLY option is not required in the MIME environment as MIME will employ a transfer encoding that preserves the PEM message. Thus MIC-CLEAR could be used in lieu of MIC-ONLY, avoiding a redundant encoding step. However, MIC-CLEAR does pose real danger when "helpful" mail relays are involved, i.e., if MIME is not available at all recipients, even if some recipients do have (non-MIME) PEM. It also is suggested that inclusion of a redundant, cleartext body part is a means of accommodating the recipients for whom MIC-CLEAR was developed. Thus this issue is unresolved.

- It is not clear that the canonical encoding options now used in MIME preserve the reversibility required for signature preservation in forwarded messages. This is a cause for some concern and requires further examination.

- There was debate over whether the preferred approach here is to define a new PEM Content-Domain for use with MIME, allowing any (8-bit) input and avoiding possibly redundant base64 encoding, or to use only the existing PEM 822 Content-Domain and impose the base64 encoding in all cases.

- The question was raised as to whether the PEM header needs to indicate Content-Domain MIME when the PEM header is already within a MIME message, or is it redundant? the issue was not resolved and requires further study.

- It is suggested by several attendees that the Content-Annotation proposed in section 6, needs to be dropped or improved. It does not provide enough information to preserve all of the security information that PEM provides. There is considerable feeling that there is a difference between what is displayed to the user as part of message reading, vs. what is retained when the message is stored. The message may be stored in enciphered form, in signed only form, or without any cryptographic (PEM) protection. It was argued that the labeling of a stored message which was previously protected by PEM is strictly a local matter and thus should not be part of the MIME header (nor part of the MIME-PEM specification).

- There was agreement to continue this discussion on the PEM-DEV mailing list and at the next IETF meeting. The authors of the Internet-Draft, which was the focal point of this discussion, agreed to work on a successor version, taking into account the various issues raised and discussed during this meeting. The PEM and MIME Working Group Chairs agreed to request that future PEM and MIME Working Group meetings during IETF be explicitly scheduled to not conflict with one another.

Attendees

David Balenson	balenson@tis.com
Fred Bohle	fab@interlink.com
David Conklin	conklin@jvnc.net

James Conklin	jbc@bitnic.educom.edu
Chuck Cranor	chuck@maria.wustl.edu
Stephen Crocker	crocker@tis.com
Art Dertke	dertke@gateway.mitre.org
Steve Dusse	spock@rsa.com
William Edison	
Barbara Fraser	byf@cert.org
Shari Galitzer	shari@mitre.org
James Galvin	galvin@tis.com
Richard Graveman	rfg@ctt.bellcore.com
Terry Gray	gray@cac.washington.edu
Neil Haller	nmh@thumper.bellcore.com
Alton Hoover	hoover@ans.net
Christian Huitema	christian.huitema@sophia.inria.fr
Phil Karn	karn@qualcomm.com
Stephen Kent	kent@bbn.com
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Louis Mamakos	louie@ni.umd.edu
Mohammad Mirhakkak	mmirhakk@mitre.org
Clifford Neuman	bcn@isi.edu
Hilarie Orman	ho@cs.arizona.edu
Joseph Ramus	ramus@nersc.gov
Alan Roszkiewicz	alan@sprint.com
Paul Sangster	sangster@ans.net
Sam Schaen	schaen@mitre.org
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
Tang Tang	tt@virginia.edu
Theodore Ts'o	tytso@mit.edu
Gregory Vaudreuil	gvaudre@cnri.reston.va.us
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Moira West	mjw@cert.org
Daniel Woycke	woycke@smiley.mitre.org
Peter Yee	yee@atlas.arc.nasa.gov

2.7.6 SNMP Security (snmpsec)

<u>Charter</u>

Chair(s):

James Galvin, galvin@tis.com
Keith McCloghrie, kzm@hls.com

Mailing Lists:

General Discussion: snmp-sec-dev@tis.com
To Subscribe: snmp-sec-dev-request@tis.com
Archive: snmp-sec-dev-request@tis.com

Description of Working Group:

Enhancements to the SNMP network management framework are being contemplated within the SNMP Version 2 Working Group of the IETF. The SNMP Security Working Group is chartered to consider changes to RFCs 1351, 1352, 1353 that may be required either for consistency with this SNMP evolution effort or to reflect implementation experience with the current specifications.

Goals and Milestones:

Done	Publish Internet-Draft specifications.
Done	Submit specification to IESG for consideration as a Proposed Standard.
Done	At the November IETF meeting, review and discuss feedback from implementation experience of the present specifications and requirements from the evolution of the SNMP Framework.
Done	Publish updated SNMP Security documents as Internet-Drafts.
Feb 1993	Submit the SNMP Security Documents to the IESG for consideration as a Draft Standard.

Internet-Drafts:

"Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)", 12/23/1992, K. McCloghrie, J. Davin, J. Galvin <draft-ietf-snmpsec-partyv2-01.txt>

"Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)", 12/23/1992, J. Davin, J. Galvin, K. McCloghrie <draft-ietf-snmpsec-adminv2-01.txt>

"Security Protocols for version 2 of the Simple Network Management Protocol (SNMPv2)", 12/23/1992, J. Galvin, K. McCloghrie, J. Davin <draft-ietf-snmpsec-secv2-01.txt>

"Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2)", 12/23/1992, J. Case, K. McCloghrie, M. Rose <draft-ietf-snmpsec-tmv2-00.txt>

"Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)", 12/23/1992, J. Case, K. McCloghrie, M. Rose <draft-ietf-snmpsec-mibv2-00.txt>

"Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)", 01/14/1993, J. Case, K. McCloghrie, M. Rose

<draft-ietf-snmpsec-protov2-00.txt>

Request For Comments:

RFC 1351 "SNMP Administrative Model"

RFC 1352 "SNMP Security Protocols"

RFC 1353 "Definitions of Managed Objects for Administration of SNMP Parties"

CURRENT MEETING REPORT

Reported by James Galvin/TIS

Minutes of the SNMP Security Working Group (SNMPSEC)

The meeting was called to order at 4:00 p.m. on November 17th. The Agenda was reviewed and accepted as presented. It was mentioned that the Charter was revised to reflect the need to align with SNMP Version 2; there were no comments on the Charter as distributed to the mailing list. Finally, it was noted that the Security Area has been alerted to the aggressive schedule we have planned and has been requested to provide a security review of the soon to be revised documents as soon as possible.

There were quite a number of editorial changes as well as technical clarifications which were suggested and approved for all three documents.

Editor's Note (md): A detailed listing of these changes is available via ftp under snmpsec-minutes-92nov.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

There were three outstanding issues after the meetings closed for the week:

1. Compliance Levels for Party MIBs with the new SNMPv2 compliance macros; it is necessary that compliance levels be specified for the party MIB; a proposal to specify three such levels was made

2. Party creation without the Data Encryption Standard (DES); a new proposal was distributed, and presented at the meeting, detailing mechanisms by which parties could be "cloned" from other parties, both to allow party creation without DES and to allow systems with limited Now Volatile RAM (NVRAM) to support temporary parties which would be re-created after each reboot.

3. Party Proliferation; "proliferation of parties" had been passed to this Working Group from the SNMPv2 Working Group as a problem to be solved; two different presentations were made on this issue; one detailed mechanisms to reduce the amount of NVRAM required to support SNMP Security through a reduction in the number of parties and a recasting of the tables in the Party MIB; the other presentation, with the same aims, suggested that spatial and temporal semantics should be embodied in the Management Information model, which therefore argued against some of the MIB recasting of the first proposal.

These three issues were left unresolved due to a lack of time and due to their interdependence (e.g., the number and types of compliance levels depend on whether DES is required for party creation).

Attendees

Steve Alexander	stevea@i88.isc.com
David Arneson	arneson@ctron.com
Jim Barnes	barnes@xylogics.com
Andy Bierman	abierman@synoptics.com
Fred Bohle	fab@interlink.com
Vickie Brown	brown@osi540sn.gsfc.nasa.gov
J. Nevil Brownlee	nevil@aukuni.ac.uz
Theodore Brunner	tob@thumper.bellcore.com
Matthew Busche	mtb@anchor.ho.att.com
Stephen Bush	sfb@ncoast.org
Dilip Chatwani	dilip@synoptics.com
Szusin Chen	szusin.chen@eng.sun.com
Robert Ching	rching@nat.com
Chris Chiotasso	chris@andr.ub.com
Bobby Clay	clay@eagle.msfc.nasa.gov
Tracy Cox	tacox@sabre.bellcore.com
Juan Cruz	juan@dss.com
Dave Cullerot	cullerot@ctron.com
Cathy Cunningham	cmc@microcom.com
James Davin	davin@bellcore.com
Cynthia Della Torre	cindy@gateway.mitre.org
Art Dertke	dertke@gateway.mitre.org
Manuel Diaz	diaz@davidsys.com
William Edison	
David Engel	david@ods.com
Raphael Freiwirth	5242391@mcimail.com
James Galvin	galvin@tis.com
Richard Graveman	rfg@ctt.bellcore.com
Maria Greene	mngreene@eng.xyplex.com
Pete Grillo	pl0143@mail.psi.net
Michel Guittet	guittet1@applelink.apple.com
Ed Heiner	eah@pau.synnet.com
Ken Hirata	khirata@emulex.com
Gerd Holzhauer	holzhauer1@applelink.apple.com
John Hopprich	hopprich@davidsys.com
Kevin Jackson	kmj@concord.com
Bob Jeckell	rrj@3com.com
Mark Kepke	mak@cnd.hp.com
John Kimmins	kimm@ctt.bellcore.com
Keith Klamm	klamm@ods.com
Andrew Knutsen	andrewk@sco.com
Michael Kornegay	mlk@bir.com
Deirdre Kostick	dck2@sabre.bellcore.com

David Lin	lind@janus-ccm.zenith.com
David Lindemulder	dcl@mtung.att.com
John Linn	linn@erlang.enet.dec.com
Benjamin Lisowski	Ben.Lisowski@sprint.sprint.com
John Lunny	jlunny@twg.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Douglas Miller	dmm@telebit.com
David Minnich	dwm@fibercom.com
Mohammad Mirhakkak	mmirhakk@mitre.org
George Mouradian	gvm@arch3.att.com
Patrick Mullaney	mullaney@ctron.com
Daniel Myers	dan@nsd.3com.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Bill Norton	wbn@merit.edu
Steven Onishi	sonishi@wellfleet.com
David Perkins	dperkins@synoptics.com
Ilan Raab	iraab@synoptics.com
Richard Ramos	ramos@mtunm.att.com
Venkat Rangan	venkat@geoduck.matrix.com
Louise Reingold	l.reingold@sprint.sprint.com
Sam Roberts	sroberts@farallon.com
Kary Robertson	kr@concord.com
Dan Romascanu	dan@lannet.com
Avraham Rosenbach	armon@armon.hellnet.org
Shawn Routhier	sar@epilogue.com
Chris Rozman	chrisr@usr.com
Assaf Rubissa	asaf@fibhaifa.com
Jon Saperia	saperia@tcpjon.ogo.dec.com
Michael Scanlon	scanlon@interlan.com
Sam Schaen	schaen@mitre.org
Jim Scott	scott@kali.enet.dec.com
John Seligson	johns@ultra.com
Timon Sloane	timon@rahul.net
Robert Snyder	snyder@cisco.com
Roy Spitzer	roy.spitzer@sprint.com
Bob Stewart	rlstewart@eng.xyplex.com
Kaj Tesink	kaj@cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Warren Vik	wmv@i88.isc.com
Ioannis Viniotis	candice@ececho.ncsu.edu
Steven Waldbusser	waldbusser@andrew.cmu.edu
Alice Wang	alice.wang@eng.sun.com

Gerry White	gerry@lancity.com
Peter Wilson	peter_wilson@3com.com
Steven Wong	wong@took.enet.dec.com
Randall Worzella	worzella@ralvm29.unet.ibm.com
Daniel Woycke	woycke@smiley.mitre.org
Kiho Yum	kxy@nsd.3com.com

2.7.7 TCP Client Identity Protocol (ident)

Charter

Chair(s):

Mike St. Johns, stjohns@darpa.mil

Mailing Lists:

General Discussion: ident@nri.reston.va.us
To Subscribe: ident-request@nri.reston.va.us
Archive: nri.reston.va.us:ietf.mailing.lists/ident.mailing.list

Description of Working Group:

The TCP Client Identity Protocol Working Group is chartered to define a protocol for returning the identity of the user initiating a TCP connection. When a client on host A initiates a TCP connection to host B, host B may query a server on host A to determine the identity of the client on host A. The primary purpose of this protocol is to record the identity of requesters initiating a connection.

This work is a clarification and standardization of the Experimental Protocol currently published as RFC 931.

Goals and Milestones:

Done Review implementations, and resolve outstanding issues in preparation for Draft Standard.

Done Post an Internet-Draft of the revised RFC 931 Identity Server Protocol.

Done Submit the Identity Server Protocol to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

"Ident MIB", 04/13/1992, Michael St. Johns, Marshall Rose <draft-ietf-ident-mib-03.txt>

"Identification Server", 06/02/1992, Mike StJohns <draft-ietf-ident-idserver-03.txt>

CURRENT MEETING REPORT

Minutes of the TCP Client Identity Protocol Working Group (IDENT)

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

Attendees

David Conklin	conklin@jvnc.net
Chuck Cranor	chuck@maria.wustl.edu
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Barbara Fraser	byf@cert.org
Cliff Frost	cliff@cmsa.berkeley.edu
James Galvin	galvin@tis.com
Frank Kastenholz	kasten@ftp.com
John Myers	jgm+@cmu.edu
Joseph Ramus	ramus@nersc.gov
Michael St. Johns	stjohns@darpa.mil
Chuck Warlick	warlick@theophilis.nsfc.nasa.gov
Walter Wimer	walter.wimer@andrew.cmu.edu

2.8 Transport and Services Area

Director(s):

- David Borman: dab@cray.com

Area Report not submitted.

CURRENT MEETING REPORT

Minutes of the DNSS II BOF (dns2)

Report not submitted.

Attendees

Roland Acra	acra@cisco.com
Robert Austein	sra@epilogue.com
Bryan Beecher	bryan@umich.edu
Robert Braden	braden@isi.edu
Hans Eriksson	hans@sics.se
Roger Fajman	raf@cu.nih.gov
Peter Ford	peter@goshawk.lanl.gov
Barbara Fraser	byf@cert.org
Shari Galitzer	shari@mitre.org
Steve Hardcastle-Kille	s.kille@isode.com
Erik Huizer	huizer@surfnet.nl
Scott Kaplan	scott@ftp.com
Daniel Karrenberg	daniel@ripe.net
Andrew Knutsen	andrewk@sco.com
Mark Laubach	laubach@hpl.hp.com
Louis Mamakos	louie@ni.umd.edu
Bill Manning	bmanning@sesqui.net
Milo Medin	medin@nsipo.nasa.gov
Clifford Neuman	bcn@isi.edu
Rakesh Patel	patel@noc.rutgers.edu
Michael Patton	map@bbn.com
Drew Perkins	perkins+@cmu.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Jon Postel	postel@isi.edu
Tim Seaver	tas@concert.net
Mark Smith	mcs@umich.edu
Marten Terpstra	marten@ripe.net
Warren Vik	wmv@i88.isc.com
Moira West	mjw@cert.org
Evan Wetstone	evan@rice.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
Liang Wu	ltw99@bellcore.com

2.8.1 Audio/Video Transport (avt)

Charter

Chair(s):

Stephen Casner, casner@isi.edu

Mailing Lists:

General Discussion: rem-conf@es.net
To Subscribe: rem-conf-request@es.net
Archive: nic.es.net:[anonymous.ietf.rem-conf]/av-transport-archiv

Description of Working Group:

The Audio/Video Transport 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. Our first step will be to solicit contributions of potential protocols from projects that have already developed packet audio and video. From these contributions we will distill the appropriate protocol features.

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.

Mar 1993 Submit to the IESG the Audio/Video Transport protocol for publication as an Experimental Protocol.

Internet-Drafts:

"Issues in Designing a Transport Protocol for Audio and Video Conferences and other Multiparticipant Real-Time Applications", 10/27/1992, H. Schulzrinne <draft-ietf-avt-issues-00.txt, .ps>

"A Transport Protocol for Real-Time Applications", 12/16/1992, H. Schulzrinne <draft-ietf-avt-rtp-00.txt>

"Media Encodings", 12/16/1992, H. Schulzrinne <draft-ietf-avt-encodings-00.txt>

"Sample Profile for the Use of RTP for Audio and Video Conferences with Minimal Control", 12/16/1992, H. Schulzrinne <draft-ietf-avt-profile-00.txt>

CURRENT MEETING REPORT

Reported by Steve Casner/USC-ISI

Minutes of Audio/Video Transport Working Group (AVT)

The AVT Working Group met for three sessions. In the first two, the Group reviewed the draft specification for the Real-time Transport Protocol (RTP); the third session was an "implementors agreement" session focusing on software video encoding.

Presentations of Draft RTP Specification

The Group are indebted to Henning Schulzrinne for his efforts in writing a summary of the discussion from the Working Group meeting at the July IETF, and subsequently developing that into a concise RTP specification and a separate rationale document comparing the design tradeoffs considered by the Working Group. The Group began with a presentation by Henning on the draft protocol specification. In brief, RTP supports the following functions:

- Transfer of media data.
- Demultiplexing of multiple flows.
- Content identification.
- Synchronization and sequencing.
- Options for simple control functions such as identification of participants.

RTP consists primarily of protocol header for real-time data packets. In the typical case, the RTP header is just 8 octets long and composed of the following fields (this includes some changes since the meeting):

- Protocol version (2 bits, value 1)
- Flow identifier (6 bits)
- Option present bit
- Synchronization bit (marks end of synchronization unit)
- Content type index (6 bits)
- Packet sequence number (16 bits)
- Timestamp, middle 32 bits of NTP-format timestamp

The slides are not included here, but full details on the protocol are available in the Internet-Drafts just released (see section 4).

Discussion of the Specification Seeking "rough consensus".

There were many issues discussed, but no roadblocks were identified. Some items simply required additional explanation in the text. All items were resolved sufficiently for the editor to produce the next draft. The following items are expanded in the text below:

- Framing of data units
- End-of-synchronization-unit flag
- Conference announcement protocol as a separate document
- Timestamp mechanisms
- Encoding/flow descriptors
- Backchannel information, including QoS measurement
- Profiles and mapping to port numbers

Framing is required when using RTP over a stream-oriented protocol layer, but we discussed here that it is also needed to allow multiple data units (e.g., from different media) in one packet. To allow alignment and to avoid length constraints, the frame length field was increased to 32 bits.

There was no objection to the change of the header flag from start-to end-of-synchronization-unit. This gives a few advantages with only a slight addition in complexity.

In the protocol draft sent out just before the meeting, a Conference Announcement Protocol (CAP) was added. CAP is intended as one near-term method of simple conference control until more sophisticated control protocols are developed. However, this protocol was deemed by some to be outside the scope of the Working Group, and in any case the Working Group agreed it should be specified separately from the RTP. It was agreed also that no specific references to audio or video encoding should be made in the RTP specification because it should be usable for other applications as well.

Unlike the previous two meetings, there was relatively little discussion of timestamp formats. The Group has settled on a real-time timestamp, rather than a timestamp based on the media sample clock, to allow the timestamp to be independent of the content type and to aid inter-media synchronization. However, implementation experience is needed to validate this choice. The Group discussed the need to clarify the wording in the specification to say that globally synchronized time is not required if it is not available (and inter-media synchronization is not required); also to specify that timestamps within a synchronization unit should be derived from media timing.

The topic receiving the most discussion was the encoding/flow (EF) field and the EF description (EFDESC). The idea was that the value of the EF field would be used as an index into a table both the flow (or sequence state space) and the encoding (renamed content) which is opaque to the RTP layer. Since the meeting, this combined-function field has been found difficult to implement, and it has been separated into two fields by sacrificing the "option length" field and replacing it with just an "option present" bit. This requires parsing of all the options to determine where the data starts, but that may not be a disadvantage if all options must be processed before the data anyway.

Another topic receiving substantial discussion was the need to provide a backchannel from receivers to the sender. The draft contained a "quality of service measurement" option that could be multicast by receivers with or without their own data, but there may also be a need to unicast encoding control information back to the sender for error control or flow control. There is a need to identify to which flow from the sender the backchannel information pertains.

A new idea was that RTP may be used in various ways for different applications, and that the Group must define and indicate those modes of use. The term "profile" is taken for this purpose. A profile might indicate that one or more options are always present in a specified order, effectively increasing the fixed size of the header. The profile would also specify how content types are defined (statically in the profile, or dynamically through some higher-level control protocol). It is expected that use of RTP with a particular profile may be identified by a registered port number for IP multicast service. Since unicast service may require dynamically assigned port numbers, the profile will have to be identified (perhaps by the registered port number) in the control protocol that communicates the dynamic port numbers between the endpoints. More work is needed on this topic.

It was suggested that a model for "entity addressing" is needed covering both the multicast and unicast cases. This touches on the use of IP multicast addresses, port numbers, flow identifiers, and identification of multiple sources within one host. Should the model of a flow be unidirectional or bidirectional? These questions were not answered.

In addition to the topics listed above, the Group discussed the need to address security measures (authentication, confidentiality, integrity) before this protocol draft can become an RFC. However, the Group did not define those measures yet.

"Implementors Agreement" Session

The real-time transport protocol should be independent of the media encoding algorithms and formats that belong to the next higher layer. However, several members of the Working Group are developing packages for software video compression, so we devoted the third Working Group session to an "implementors agreement" discussion to promote convergence and interoperation among these packages.

The Group heard presentations by Thierry Turletti on the INRIA "IVS" system implementing software H.261 encoding; by Richard Cogger from Cornell on the CUSeeMe package for Macintosh; and a short description was given remotely by Ron Frederick at Xerox PARC on the "nv" package. Paul Milazzo, who has previously made a presentation on the BBN "DVC" system, and Bob Clements of BBN, also participated in the discussion remotely over the packet audio channel. Oliver Jones from PictureTel made a presentation on coding standards applicable to this effort.

It was found that there was much in common among these systems, and several of the implementors agreed to work together toward convergence and interoperation. A first step is for a description of each of these systems to be posted to the mailing list rem-conf@es.net (some information has already been posted). There was some discussion of defining an API for the software compression algorithms so they could be plugged into application frameworks on different platforms. However, Paul Milazzo pointed out that it may be necessary to interleave compression operations into the acquisition process to reduce processing time, so it may infeasible or at least premature to define an API between the two steps.

The Group also determined there were no conflicts between the draft RTP protocol and the requirements of these packages. The Group will need to define an enumeration of these experimental encodings to allow systems to process multiple formats.

Further Working Group Activities

Subsequent to this meeting, an updated set of Internet-Drafts on RTP was issued on December 18th to incorporate the changes discussed at the meeting. These are:

- draft-ietf-avt-rtp-00.txt
- draft-ietf-avt-encoding-00.txt
- draft-ietf-avt-profile-00.txt
- draft-ietf-avt-issues-00.ps, .txt

The first draft is the specification of the real-time transport protocol itself. The second and third drafts define a set of media encodings and a sample profile for use of those encodings to implement audio and video multiparticipant conferences with minimal control. The last draft is an updated discussion of the issues and decisions involved in the design of the protocol.

Before these drafts are issued as RFCs, it is important that sufficient implementation and operational experience be obtained to validate or revise the protocol. The goal should be to implement the protocol for both audio and video, experiment with it and have implementations ready for use to multicast the next IETF meeting in March. Assuming success in this process, the drafts should then be submitted to become RFCs after review at the March meeting.

.ttendees

Vikas Aggarwal	vikas@jvnc.net
Brian Bataille	bataillebc@afotec.af.mil
Lou Berger	lberger@bbn.com
Dean Blackketter	deanb@apple.com
Rita Brennan	brennan@apple.com
Stephen Casner	casner@isi.edu
Kay Chang	chang@chang.austin.ibm.com
Wo Chang	wchang@nist.gov
Richard Cogger	rhx@cornell.cit.bitnet
Robert Cole	rgc@qsun.att.com
Hans Eriksson	hans@sics.se
Jerry Friesen	jafries@sandia.llnl.gov

James Geddes	wk05020@worldlink.com
Robert Gilligan	Bob.Gilligan@eng.sun.com
Mark Green	markg@apple.com
Thomas Hacker	hacker@citi.umich.edu
Don Hoffman	don.hoffman@eng.sun.com
Christian Huitema	christian.huitema@sophia.inria.fr
Oliver Jones	oj@pictel.com
Phil Karn	karn@qualcomm.com
Charley Kline	cvk@uiuc.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Christopher Kolb	kolb@psi.com
Fong-Ching Liaw	fong@eng.sun.com
Louis Mamakos	louie@ni.umd.edu
Donald Merritt	don@brl.mil
Greg Minshall	minshall@wc.novell.com
Mitra	mitra@pandora.sf.ca.us
Kathleen Nichols	nichols@apple.com
Ari Ollikainen	ari@es.net
Michael Patton	map@bbn.com
Jim Perchik	perchik@athena.mit.edu
Mike Petry	petry@ni.umd.edu
Joe Ragland	jrr@concert.net
Bala Rajagopalan	braja@qsun.att.com
Allan Rubens	acr@merit.edu
Tom Sandoski	tom@concert.net
Eve Schooler	schooler@isi.edu
Dallas Scott	scott@fluky.mitre.org
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@andr.ub.com
Joo Young Song	jysong@ring.kotel.co.kr
Terrance Sullivan	terrys@newbridge.com
Tang Tang	tt@virginia.edu
Morton Taragin	vsmorty@weizmann.weizmann.ac.il
Sally Tarquinio	sallyt@gateway.mitre.org
Claudio Topolcic	topolcic@cnri.reston.va.us
Thierry Turletti	turletti@sophia.inria.fr
Zheng Wang	z.wang@cs.ucl.ac.uk
Von Welch	vwelch@ncsa.uiuc.edu
Peter Will	will@isi.edu
Kirk Williams	kirk@sbctri.sbc.com
Jeff Young	young@alw.nih.gov
Paul Zawada	Zawada@ncsa.uiuc.edu

2.8.2 Distributed File Systems (dfs)

Charter

Chair(s):
Peter Honeyman, honey@citi.umich.edu

Mailing Lists:
General Discussion: dfs-wg@citi.umich.edu
To Subscribe: dfs-wg-request@citi.umich.edu
Archive:

Description of Working Group:

Trans- and inter-continental distributed file systems are upon us. The consequences to the Internet of distributed file system protocol design and implementation decisions are sufficiently dire that we need to investigate whether the protocols being deployed are really suitable for use on the Internet. There's some evidence that the opposite is true, e.g., some distributed file systems protocols don't checksum their data, don't use reasonable MTUs, don't offer credible authentication or authorization services, don't attempt to avoid congestion, etc. Accordingly, a Working Group on DFS has been formed by the IETF. The Working Group will attempt to define guidelines for ways that distributed file systems should make use of the network, and to consider whether any existing distributed file systems are appropriate candidates for Internet standardization. The Working Group will also take a look at the various file system protocols to see whether they make data more vulnerable. This is a problem that is especially severe for Internet users, and a place where the IETF may wish to exert some influence, both on vendor offerings and user expectations.

Goals and Milestones:

May 1990 Generate an RFC with guidelines that define appropriate behavior of distributed file systems in an internet environment.

2.8.3 Domain Name System (dns)

Charter

Chair(s):

Rob Austin, sra@epilogue.com

Mailing Lists:

General Discussion: dns-wg@nsl.dec.com
To Subscribe: dns-wg-request@nsl.dec.com
Archive:

Description of Working Group:

The DNS Working Group is concerned with the operation of name servers on the Internet. We do not operate name servers but serve as a focal point for the people who do operate them. We are also concerned with the Domain Name System itself. Changes to the existing RFC's, for example, are discussed by the Working Group. If changes to the RFC's or additional DNS related RFC's are deemed necessary the Working Group will propose them and will prepare the associated documents.

Because we intend to serve as the focal point for people operating name servers, one of our projects will be to assist anyone bringing up a name server by publishing a collection of useful hints, tips and operational experience learned by the people already running name servers.

The DNS Working Group will also take an active role in the dissemination of solutions to problems and bugs encountered while running various name server implementations. We will also provide guidance to anyone writing a new name server implementation, whenever possible.

Goals and Milestones:

TBD Adding DNS variables to the MIB.

TBD Hints, tips, and operations guide for DNS software.

TBD Implementation catalog for DNS software.

TBD Discussion of adding load balancing capability to the DNS.

TBD Discussion of adding a Responsible Person Record.

TBD Discussion of adding network naming capability to the DNS.

Internet-Drafts:

"DNS MIB Extensions", 03/05/1992, R. Austein, J. Saperia <draft-ietf-dns-mibext-05.txt, .ps>

CURRENT MEETING REPORT

Reported by Bill Manning/RICE

Minutes of the Domain Name System Working Group (DNS)

Agenda

- Review proposed changes to RFC1348
- Shuffle and load averaging RR support
- Large zones problems & policy
- IDPR
- DNS values set from SNMP

RFC1348 Changes

The Draft clears up the reverse lookup method and adds the inclusion of Network Entity Titles. NSAP prefix storage appears to be unworkable without a seperate RR type. If this feature is desired, then a new RR may be added. The label NSAP-PTR needs to be corrected, since it is not a pointer. Bill will post the revised Draft shortly after the IETF.

IDPR

They need a RR to lable an administrative domain or AD. The plan is to use a 32bit number. Rob will follow up and make sure this happens.

Shuffle and Load

Mimi described a problem that needs to query for load average of a logical cluster of machines. Mel and Walter both came back with two distinct ways to try and address this problem. Mimi will write up a draft that better covers the problem space.

Large Zones

Mike pointed out the problem with the com and edu zones. The Group should recommend to the new NIC that following rule of thumb be used for new assignments in the EDU zone:

- xxxx.edu - 4 year schools that issue masters/phd degrees or equivalent.
- xxxx.state.country - 2 year schools or those that issue associate or bachelor degrees.
- xxxx.k12.state.country - pre-collage

The COM domain is still a problem. The Group might recommend the second level be assigned an industry type. The problem would also be helped by distributing the servers off root platforms.

MIB: IESG to decide on Proposed Standard this week. The discussion regarding setting values via SNMP will be discussed on the list. Scott will lead the charge here.

Attendees

Roland Acra	acra@cisco.com
Robert Austein	sra@epilogue.com
Bryan Beecher	bryan@umich.edu
David Conklin	conklin@jvnc.net
Chuck Cranor	chuck@maria.wustl.edu
Elise Gerich	epg@merit.edu
Heather Gray	heather@zk3.dec.com
Scott Kaplan	scott@ftp.com
David Katinsky	dmk@rutgers.edu
John Larson	jlarson@parc.xerox.com
Daniel Long	long@nic.near.net
E. Paul Love	loveep@sdsc.edu
Kent Malave	kent@bach.austin.ibm.com
Louis Mamakos	louie@ni.umd.edu
Bill Manning	bmanning@sesqui.net
Clifford Neuman	bcn@isi.edu
Andrew Partan	asp@uunet.uu.net
Mel Pleasant	pleasant@hardees.rutgers.edu
Cris Shuldiner	cws@ftp.com
Michael St. Johns	stjohns@darpa.mil
Tang Tang	tt@virginia.edu
Walter Wimer	walter.wimer@andrew.cmu.edu
Miriam Zohar	zohar@watson.ibm.com

2.8.4 Service Location Protocol (svrloc)

Charter

Chair(s):
John Veizades, `veizades@apple.com`
Scott Kaplan, `scott@ftp.com`

Mailing Lists:
General Discussion: `srv-location@apple.com`
To Subscribe: `srv-location-request@apple.com`
Archive: `pub/srv-location/svr-loc-archive`

Description of Working Group:

The Service Location Working Group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based services location protocols. Protocols that would be designed by this Group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the services location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this Group to understand the operation of services location and then determine the correct action in their view whether it be to use current protocols to suggest a services location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

Done	Open discussion and determine if a Working Group should be formed.
Done	Continue discussion trying to refine the problem statement and possible resolutions.
Jul 1991	Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the Service Location Protocol Working Group (SVRLOC)

The Service Location Working Group met to discuss the latest proposal on service location by Scott Kaplan. The document was discussed and comments were made relating to:

- Scalablity - This proposal has to work from the dentist office case to some upward bound. How is that upward bound found and how does the protocol make the transition at this boundary case? This is particularly of issue with the initial dictionary query.

- Security - How does the protocol pass around security information for authorization control.

- Multilingual - The protocol can be designed to support a variety of languages in one administrative domain by the use of a translation service. The issue of data encoding came up and ASCII english will be used for transmission of character information used in text matching and the translation service will handle multilingual issues. This should be addressed in the next draft of the document.

The Group thought that the proposal was in general on the right track. John Veizades and Scott Kaplan will continue to evolve the proposal with intent to have an Internet-Draft by the beginning of next year.

Bill Nowiciki presented work on RAP by Legato. Interesting insights observed are that protocols like this provide one user interface to services per platform solving the cross product problem of services to platforms.

Scott Kaplan is now the co-Chair of this Working Group taking over from Leo. There will be a new Charter coming out shortly reflecting this and an accurate schedule for the Working Group.

Attendees

Richard Basch	basch@mit.edu
Naomi Courter	naomi@concert.net
Karen Frisa	karen.frisa@andrew.cmu.edu
Nat Howard	nrh@bellcore.com
Ronald Jacoby	rj@sgi.com
Scott Kaplan	scott@ftp.com
Andrew Knutsen	andrewk@sco.com
Michael Laufer	mlaufer@bbn.com

Edward Levinson	levinson@pica.army.mil
Brian Marsh	marsh@mitl.com
John Myers	jgm+@cmu.edu
William Nowicki	nowicki@legato.com
Tatsuya Ohnishi	ohnishi@isl.mei.co.jp
Gaige Paulsen	gaige@intercon.com
Charles Perkins	perk@watson.ibm.com
Bradley Rhoades	bdrhoades@mail.mmmg.com
Mike Ritter	mwritter@applelink.apple.com
Richard Schmalgemeier	rgs@merit.edu
Cris Shuldiner	cws@ftp.com
Brad Steinka	brad@microcom.com
Jim Thompson	jim@tadpole.com
John Veizades	veizades@apple.com

2.8.5 TCP Large Windows (tcplw)

Charter

Chair(s):

David Borman, dab@cray.com

Mailing Lists:

General Discussion: tcplw@cray.com
To Subscribe: tcplw-request@cray.com
Archive:

Description of Working Group:

The TCP Large Windows Working Group is chartered to produce a specification for the use of TCP on high delay, high bandwidth paths. To this end, this Working Group recommended RFC 1072 "TCP extensions for long-delay paths" and RFC 1185 "TCP Extension for High-Speed Paths" be published jointly as a Proposed Standard. Deficiencies in the technical details of the documents were identified by the End-to-End Research Group of the IRTF. Rather than progress the standard with known deficiencies, the IESG tasked the End-to-End Research Group to fix and merge these two documents into a single protocol specification document. This review was done on the eze-interest@isi.edu mailing list.

The TCP Large Windows Working Group is being resurrected for a one time meeting, to review and if appropriate, approve this new document.

Goals and Milestones:

Done	Review the TCP Extended Window Size proposal from the IRSG End to End Research Group and if acceptable, recommend it for standards status.

Request For Comments:

RFC 1323	"TCP Extensions for High Performance"

2.8.6 Trusted Network File Systems (tnfs)

Charter

Chair(s):
Fred Glover, `fglover@zk3.dec.com`

Mailing Lists:
General Discussion: `tnfs@wdl1.wdl.loral.com`
To Subscribe: `tnfs-request@wdl1.wdl.loral.com`
Archive: `archive-server@wdl1.wdl.loral.com`

Description of Working Group:

The Trusted Network File System Working Group is chartered to define protocol extensions to the Network File System (NFS) Version 2 protocol which support network file access in a Multilevel Secure (MLS) Internet environment. MLS functionality includes Mandatory Access Control (MAC), Discretionary Access Control (DAC), authentication, auditing, documentation, and other items as identified in the Trusted Computer System Evaluation Criteria (TCSEC) and Compartmented Mode Workstation (CMW) documents.

The primary objective of this Working Group is to specify extensions to the NFS V2 protocol which support network file access between MLS systems. It is intended that these extensions should introduce only a minimal impact on the existing NFS V2 environment, and that unmodified NFS V2 clients and servers will continue to be fully supported.

Transferring information between MLS systems requires exchanging additional security information along with the file data. The general approach to be used in extending the NFS V2 protocol is to transport additional user context in the form of an extended NFS UNIX style credential between a Trusted NFS (TNFS) client and server, and to map that context into the appropriate server security policies which address file access. In addition, file security attributes are to be returned with each TNFS procedure call. Otherwise, the NFS V2 protocol remains essentially unchanged.

The Trusted System Interoperability Group (TSIG) has already developed a specification which defines a set of MLS extensions for NFS V2, and has also planned for the future integration of Kerberos as the authentication mechanism. The TNFS Working Group should be able to use the TSIG Trusted NFS document as a foundation, and to complete the IETF TNFS specification within the next 3-6 months.

Goals and Milestones:

Mar 1991 Verify the interoperability of TNFS implementations at the 1992 NFS Connectathon.

Done Review and approve the TNFS Working Group Charter, review revised TSIG TNFS Specification, and publish a proposed standard following the July meeting.

Jul 1991 Review revised TSIG TNFS Specification.

Oct 1991 Review outstanding comments/issues from mailing list.

Oct 1991 Make any final revisions to TNFS document based on comments, issues, and interoperability testing.

Nov 1991 Publish a Proposed Standard following the July meeting.

Mar 1992 Request IESG to make the revised document a Draft Standard.

Internet-Drafts:

"A Specification of Trusted NFS (TNFS) Protocol Extensions", 07/23/1991,
Fred Glover <draft-ietf-tnfs-spec-01.txt>

2.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Seven working groups and one BOF (Birds of a Feather) met in Washington, D.C..

Directory Information Services (pilot) Infrastructure Working Group (DISI)

The DISI Working Group provides a forum to define user requirements in X.500. It is a combined effort of the User Services Area and the OSI Integration Area.

The three papers (Advanced Usages Guide, the Pilot Project Catalog, and the revision of RFC 1292/FYI 11) are all nicely on track. Drafts of the first two of those papers have been submitted, and the template for responses to the RFC 1292 revision has been finalized.

DISI was originally chartered as a Working Group for X.500 documentation. However, with the introduction of new directory services protocols in the Internet (whois++), DISI decided that there was a role for it in working on specifications for integrating the directory services, and living up to the Directory Information Services Infrastructure title of the Group. Consequently, DISI has decided to revise the Charter to reflect the new goals, and Tim Howes of the University of Michigan has agreed to be co-Chair of the revamped DISI in order to bring a sense of balance between X.500 and whois++.

Internet Anonymous FTP Archives Working Group (IAFA).

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.

The Internet Anonymous FTP Archives Working Group met for what is expected to be the last time. In a discussion of the current User and System Administration draft documents it was decided that a final editing pass was required and that they could then be submitted for Internet-Draft status. The general consensus was that IAFA had completed its mandate and that additional work in information distribution on the network would be better performed in new and existing working groups specifically tasked for the purpose.

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.

Three topics were discussed:

1. Development of an FAQ (Frequently Asked Questions) list for K-12 questions.
2. Connectivity models for school-or school-district-size organizations.
3. Domain name conventions for K-12 organizations.

Results of the meeting were: a temporary group was developed to work with Jennifer Sellers and April Marine to post a first FAQ list; there was agreement that the initial list of candidate connectivity models would be fleshed out and offered for discussion and amplification; and John Clement agreed to work with Ellen Hoffman to post a preliminary list of current K-12 domain names and to request input of further existing domain names.

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. There were three items on the Agenda.

1. Status of nic-profiles
2. Net-help
3. Future of NISI.

Nic-profiles. Nic-profiles has been made an action item. The near-term goal is to have as many NICs as possible in our db by the March IETF. Three people have volunteered to work with Pat Smith on pro-active outreach to the obtain the nic information. For the record, a mail server is now in place to add/modify/retrieve information. Send to x500test@merit.edu and put "help" in the subject field to receive further information.

Net-help. Two short summary papers are going to be written. One will summarize the "content" side, what will be on the screen, etc. Much of this was discussed at the Boston meeting. The other short paper will describe the envisioned plan for implementation. The Group plans on working with representatives from another area (probably Applications) to handle this area.

The future of NISI. Again, there was a short period of discussion regarding the future of NISI. The general tone is that NISI is needed and will continue as a viable working group.

NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

The "Son of NOCTools" Working Group is updating and revising their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

NOCTool2 has submitted an Internet-Draft at this IETF for review and forward to the RFC Editor for publication. This Working Group has now come to closure.

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 Working Group session began with a brief overview of the shorter bibliography and with a few minor revisions. The document is currently available within an internet archive file which is available at four sites across the country. The Group sought volunteers to make this information available at the international level because we are conscious of the fact that the current information is U.S. centric. The Group also decided that as a Group we would explore the mechanisms of document delivery with the goal of facilitating the transfer of information to the users home host. It was further ascertained that some of the original material within the RFC 1175 was no longer current and other information such as conference proceedings of an applicable nature would be more appropriate at this point. The Working roup will meet again at the next IETF in Columbus, Ohio and will ascertain our future endeavors at that time.

Internet User Glossary Working Group (USERGLOS)

The USERGLOS Working Group is chartered to create an Internet specific glossary of networking terms and acronyms for the Internet community.

USERGLOS met for the final time. Some final editorial adjustments were made to the Internet-Draft. There were also a couple of entries removed and a couple added. Send the edited Internet-Draft to the RFC Editor in December. This Working Group has now come to closure.

User Services Working Group (USWG)

The USWG provides a regular forum for people interested in all user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet.

At this IETF session announcements included working groups coming to closure (DISI (will be revised with a new Charter and co-Chair), IAFA, NOCTOOL2, and USERGLOS) and new working groups starting up (WNILS, NIR, URI, IIIR). The four new working groups will be a joint effort of the Applications and User Services areas. *Editor's Note (md): Please refer to the Applications Area Report for a brief summary of these four working groups.* New FYI RFC publications, current user services related Internet-Drafts postings include:

Internet-Drafts:

- draft-malkin-newcomers-guide-00.txt
- draft-ietf-noctool2-debug-tcpip-00.txt
- draft-ietf-userdoc2-fyi-novice-01.txt
- draft-ietf-userglos-glossary-00.txt

FYI RFCs:

- FYI 16 "Connecting to the Internet: What Connecting Institutions Should Anticipate", (Also RFC 1359), August 1992.

- FYI 15 "Privacy and Accuracy Issues in Network Information Center Databases", (Also RFC 1355), August 1992.

Jill Foster provided an update on RARE ISUS activities. Joyce Reynolds provided reports on RIPE activities, and on NETF activities.

Steve Coya led a discussion on the Internet-Draft, "draft-malkin-newcomers-guide-00.txt" and requested comments from the USWG. Peter Deutsch led a discussion on a USWG project in development called, "Internet Quick and Dirty". It is intended to be a short document on descriptions of each network service with pointers on where to obtain additional information. Gary Malkin requested new volunteers to help update the FAQ for New Internet Users and the FAQ for Experienced Internet Users. FYI16/RFC1359 was discussed, primarily focusing on expanding and updating the document.

Training Materials BOF (TRAINMAT)

One BOF was held in the User Services area regarding a working group formation on Training Materials. Ideas and thoughts on forming a working group on training materials were discussed, including having the endeavor be a joint effort between RARE and the IETF. The BOF attendees decided that a working group should be formed. Jill Foster and Ellen Hoffman will be the co-Chairs of this new Group.

CURRENT MEETING REPORT

Minutes of the Training Materials BOF (TRAINMAT)

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

Attendees

Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Henry Clark	henryc@oar.net
Naomi Courter	naomi@concert.net
Jill Foster	jill.foster@newcastle.ac.uk
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Lenore Jackson	jackson@nsipo.nasa.gov
Laura Kelleher	lak@merit.edu
Janet Marcisak	jlm@ftp.com
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Jennifer Sellers	sellers@nsipo.arc.nasa.gov
Patricia Smith	psmith@merit.edu
Simon Spero	simon_spero@unc.edu
Tang Tang	tt@virginia.edu
Moira West	mjw@cert.org

2.9.1 Directory Information Services Infrastructure (disi)

Charter

Chair(s):
Chris Weider, clw@merit.edu

Mailing Lists:
General Discussion: disi@merit.edu
To Subscribe: disi-request@merit.edu
Archive: pub/disi-archive@merit.edu

Description of Working Group:

The Directory Information Services (pilot) Infrastructure Working Group is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

The DISI Working Group is an offshoot of the OSI Directory Services Group, and, accordingly, is a combined effort of the OSI Integration Area and User Services Area of the IETF. The current OSIDS Working Group was chartered to smooth out technical differences in information storage schema and difficulties in the interoperability and coherence of various X.500 implementations. The DISI Group is concerned solely with expanding the Directory Services infrastructure. As DISI will be providing infrastructure with an eye towards truly operational status, DISI will need to form liaisons with COSINE, Paradyse, and perhaps the RARE WG3.

As a final document, the DISI Working Group shall write a Charter for a new working group concerned with user services, integration, maintenance, and operations of Directory Services, the Internet Directory User Services Group.

Goals and Milestones:

Done	Submit an Internet-Draft on 'Catalog of available X.500 Implementations'
Done	Submit to the IESG the 'Catalog of available X.500 Implementations' as an informational document.
Done	Submit an Internet-Draft on 'Executive Introduction to X.500'

Done	Submit to the IESG the 'Executive Introduction to X.500' as an informational document.
Done	Submit an Internet-Draft on 'A Technical Overview of Directory ervices and X.500'.
Done	Submit to the IESG the 'Technical Overview of Directory Services and X.500' as an informational document.
Done	First IETF Meeting: review and approve the Charter making any changes necessary. Examine needs and resources for the documentation to be produced, using as a first draft a document produced by Chris Weider, Merit, which will be brought to the IETF. Assign writing assignments. Further work will be done electronically.
Done	Submit as an Internet-Draft the 'Advanced Usages' paper.
Jul 1992	Submit as an Internet-Draft the 'How to get registered' paper.
Nov 1992	Submit to the IESG the 'How to get registered' paper as an informational document.
Nov 1992	Submit to the IESG the 'Advanced Usages' paper as an informational document.
Nov 1992	Submit as an Internet-Draft the 'Pilot Projects Catalog' paper.
Nov 1992	Submit as an Internet-Draft the 'Where do I belong in the Directory' paper.
Mar 1993	Submit to the IESG the 'Pilot Projects Catalog' as an informational document.
Mar 1993	Submit to the IESG the 'Where do I belong in the Directory' paper as an informational document.
Mar 1993	Submit as an Internet-Draft the 'Guide to setting up a DSA'.
Jul 1993	Submit to the IESG the 'Guide to setting up a DSA' as an informational document.

Internet-Drafts:

"A Survey of Advanced Usages of X.500", 10/07/1992, Chris Weider, Russ Wright, Elizabeth Feinler <draft-ietf-disi-x500-survey-00.txt>

Request For Comments:

RFC 1292	"A Catalog of Available X.500 Implementations"
RFC 1308	"Executive Introduction to Directory Services Using the X.500 Protocol"
RFC 1309	"Technical Overview of Directory Services Using the X.500 Protocol"

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Directory Information Services Infrastructure Working Group (DISI)

Agenda

- Review of Old Minutes
- Advanced Usages Document
- Pilot Project Catalog
- Revision of RFC 1292
- New Papers if Needed

Minutes

There were no changes or additions to the old Minutes.

Advanced Usages Document

Russ Wright and Chris mentioned that drafts had (have) been circulated, and that a final document should be ready for the Columbus, IETF.

Pilot Project Catalog

Tim Howes spoke for April Marine, who produced the document, and distributed a draft just before the Washington IETF. Chris mentioned that April might not be able to see the document through as SRI is eliminating its NIC services; Tim graciously agreed to adopt the document if April has to move on. Most people had not read the paper yet, so comments should go to the list.

Revision of RFC 1292

Arlene Getchell and Srinivas Sataluri spoke about the expanded survey that they will be sending out to solicit updates for RFC 1292. There was some discussion of the survey, but there were no substantive changes made. A first draft of the new FYI should be ready for Columbus.

New Papers if Needed

There was a general consensus that there were no more papers necessary. Chris Weider mentioned that since there were now several directory services protocols (whois++ in addition to X.500) that the focus of the Group could profitably be shifted to working on integration between the existing (and future) directory services to allow the creation of a unified directory service. Erik Huizer felt that there was a need for guides for directory services providers and administrators, and a need to work on data integrity and maintainability. Consequently, Chris Weider and Tim Howes agreed to create a Charter for a new Working Group called Integrated Directory Services, which they would both co-Chair as representatives of the different protocols. This new Group will take over all the documents currently under construction by DISI (especially RFC 1292, which as an FYI must be kept up to date), and will work on these new issues also. A Charter will be released for this new Group by the end of 1992; once the Charter has been approved, DISI will be dissolved.

Attendees

Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Naomi Courter	naomi@concert.net
Arlene Getchell	getchell@es.net
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Tim Howes	tim@umich.edu.
Erik Huizer	huizer@surfnet.nl
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Barbara Jennings	bjjenni@sandia.gov
Neil Katin	neil.katin@eng.sun.com
Edward Krol	e-krol@uiuc.edu
Janet Marcisak	jlm@ftp.com
Bradley Rhoades	bdrhoades@mail.mmmg.com
Srinivas Sataluri	sri@qsun.att.com
Jennifer Sellers	sellers@nsipo.arc.nasa.gov
Mark Smith	mcs@umich.edu
Larry Snodgrass	snodgras@bitnic.educom.edu
Simon Spero	simon_spero@unc.edu
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov
Yung-Chao Yu	yy@qsun.att.com

2.9.2 Integration of Internet Information Resources (iiir)

Charter

Chair(s):
Chris Weider, clw@merit.edu

Mailing Lists:
General Discussion: iiir@merit.edu
To Subscribe: iiir-request@merit.edu
Archive: merit.edu:pub/iiir-archive

Description of Working Group:

The Integration of Internet Information Resources Working Group (IIIR) is chartered to facilitate interoperability between Internet Information Services, and to develop, specify, and align protocols designed to integrate the plethora of Internet information services (WAIS, ARCHIE, Prospero, etc.) into a single "virtually unified information service" (VUIS). Such protocols would include (but are not limited to) update protocols for distributed servers, a 'query routing protocol' to pass queries between existing services, protocols for gateways between existing and future services, and standard exchange formats (perhaps based on Z39.50) for cross-listing specific information.

Also, where necessary, IIIR will create technical documentation for protocols used for information services in the Internet.

Goals and Milestones:

Ongoing	Track emerging Internet information services in order to specify technical requirements for their integration into the VUIS.
Ongoing	Liaise with other groups working on deployment and integration of Internet information services: e.g., The Coalition for Networked Information, RARE Working Group 3, etc.
Nov 1992	Post an Internet-Draft defining gateway protocols.
Mar 1993	Submit final version of gateway protocol design to the IESG as a Proposed Standard.
Mar 1993	Post an Internet-Draft defining common exchange formats.
Jun 1993	Post an Internet-Draft defining update protocols for distributed servers.
Jun 1993	Post an Internet-Draft defining a Query Routing Protocol.
Jun 1993	Submit final version of common exchange format to the IESG as a Proposed Standard.

Nov 1993 Submit final version of update protocol to the IESG as a Proposed Standard.

Mar 1994 Submit final version of Query Routing Protocol to the IESG as a Proposed Standard.

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Integration of Internet Information Resources Working Group (IIIR)

This was the first meeting of the IIIR Working Group. There are a growing number of "information services" popping up on the Internet, all of which have their own clients, servers, paradigms, etc. It is the Chair's opinion, that it is time to start trying to tie them together to provide a "virtually unified" information service to the end user.

After a brief overview of the Charter, the Group heard from Tim, Peter, and Jim, each of whom have experience in integrating various information services together. Each of them was asked to speak for a bit on what they had done to tie resources together, and then to share their views on what needed to be done next. The Group then discussed the Charter.

Most people felt that it was not terribly useful for the Working Group to get into the business of specifying gateway protocols, and that it was too early to attempt some of the other milestones, such as common exchange formats and query routing protocols. However, after an hour of wrangling back and forth, agreement was reach, that it would be useful at this time to start work on two documents, one to provide a vision of the future for Internet information services in 1995, and the second to build a taxonomy of currently available information service tools.

The first document would allow the Group to define a shared vision of information services, to which the Group could start integration of existing resources. The second document would allow the Group to determine which conceptual tools it currently had to start building its future vision, and determine which tools needed to be developed to allow integration of all the information resources.

These two papers will be assigned on the mailing list, so that everyone who wishes to contribute may. The Charter will also be revised to reflect these new goals, and will be published to the list by the end of 1992.

Attendees

George Brett	George.Brett@cnidr.org
J. Nevil Brownlee	nevil@aukuni.ac.uz
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@bunyip.com
Alan Emtage	bajan@bunyip.com
Jill Foster	jill.foster@newcastle.ac.uk

Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
Thomas Hacker	hacker@citi.umich.edu
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
J. Paul Holbrook	holbrook@cic.net
Nat Howard	nrh@bellcore.com
Tim Howes	tim@umich.edu.
Erik Huizer	huizer@surfnet.nl
Dan Magorian	magorian@ni.umd.edu
Janet Marcisak	jlm@ftp.com
Mitra	mitra@pandora.sf.ca.us
Charlotte Mooers	mooers@nnsc.nsf.net
Clifford Neuman	bcn@isi.edu
Jim Perchik	perchik@athena.mit.edu
Pete Percival	percival@indiana.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Richard Rodgers	rodgers@nlm.nih.gov
Richard Schmalgemeier	rgs@merit.edu
Jane Smith	jds@jazz.concert.net
Mark Smith	mcs@umich.edu
Simon Spero	simon_spero@unc.edu
Chris Weider	clw@merit.edu
Yung-Chao Yu	yy@qsun.att.com

2.9.3 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: pub/iafa-archive@archive.cc.mcgill.ca

Description of Working Group:

The Internet Anonymous FTP Archives Working Group is chartered to define a set of recommended standard procedures for the access and administration of anonymous ftp archive sites on the Internet. Such a set of procedures will provide a framework for:

(a) Allowing the inexperienced Internet user the ability to more easily navigate the hundreds of publically accessible archive sites.

(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.

CURRENT MEETING REPORT

Reported by Alan Emtage/Bunyip and Peter Deutsch/Bunyip

Minutes of the Internet Anonymous FTP Archives Working Group (IAFA)

The Minutes of the Boston meeting were accepted as read.

A draft of April Marine's IAFA User's Document was reviewed. There was some discussion about the need for coordinating changes to this document if IAFA were to close, which led in turn to a discussion about whether such additional work should be done through IAFA or returned to the User Services Working Group. It was decided that it was not desirable to continue IAFA if its principal work on the IAFA documents was completed. it was agreed that if it was decided by the end of the meeting that IAFA's primary work on the Site Administrators' documents is done, then future editing of the User Document should be done in the User Services Working Group at the next IETF.

A final discussion was made on the two part Site Administrator's Draft Documents. A number of minor changes were proposed, including the need to spell out that such template items as Mailing List support free form text, that server types can include LISTSERV or MAILBASE and that mailing lists can be OPEN, CLOSED or PRIVATE.

It was also agreed that we should include a specific citation of the privacy and accuracy policies RFC in these documents.

It was agreed that we should have an "X-field" to allow users to add experimental fields. These could be ignored by automatic indexing software or treated as they saw fit. It was also suggested that a "See-Also" field be included.

It was agreed that the Group could beef up the "reasonable practices" field in the documents to include more specific recommendations about what a new site administrator should know about running a site. Examples suggested include:

- Use descriptive filenames on newly created files.
- Don't arbitrarily change filenames copied from other sites.
- Update README files regularly, to ensure users get current (and valid) information.

Alan Emtage will coordinate making these last recommended changes in the final draft and circulate to the list as soon as possible. Once this is done, the two drafts will be submitted to the RFC editor for advancement to draft status.

Some discussion was held about coordinating the implementation of IAFA templates across the Internet. Volunteers were sought to help coordinate an "implementation phase" in which sites were encouraged to install IAFA templates describing their anonymous FTP archive's contents.

Several possible avenues for promoting the use of the IAFA templates were suggested, including approaching the moderators of appropriate source-distribution Usenet newsgroups, plus the moderators of prominent archive sites, to persuade them to request the use of IAFA templates for future submissions. In addition, it was agreed that those involved in the implementation phase should coordinate closely with those working on other projects or information delivery schemes, including Gopher, TopNode and others.

Several people volunteered to help with a cataloging effort and once the final editing copy goes to draft status a mailing list will be set up at bunyip.com to coordinate the authoring and caching of completed templates for the most popular packages already on the net. Volunteers to work on this effort will be sought both through the IAFA Working Group and through such avenues as comp.archives.admin and other archive-related newsgroups.

The following have volunteered their sites to act as repositories for completed templates:

- Paul Holbrook (holbrook@cic.net)
- Laura Kelleher (lak@merit.edu)
- Mitra (mitra@pandora.sf.ca.us)
- Simon Spero (ses@sunsite.unc.edu)

New Business

The possibility for liaison work with other groups was discussed briefly (including the new IIIR, NIR, URI Working Groups at the IETF and the library community through such efforts as TopNode, etc.) but it was agreed that IAFA has already fulfilled its Charter and should leave such work to other working groups. Given the successful completion of the Site Administration documents, it was voted to close the Working Group after this meeting. The mailing list and archive will continue while the documents advance through the draft stage and the Co-Chairs will help coordinate the implementation recommendations outlined above.

Attendees

Brian Bataille	bataillebc@afotec.af.mil
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
David Conklin	conklin@jvnc.net
Naomi Courter	naomi@concert.net
Peter Deutsch	peterd@bunyip.com
Alan Emtage	bajan@bunyip.com
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu

Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
J. Paul Holbrook	holbrook@cic.net
Lenore Jackson	jackson@nsipo.nasa.gov
Laura Kelleher	lak@merit.edu
Edward Krol	e-krol@uiuc.edu
John Kunze	jak@violet.berkeley.edu
Edward Levinson	levinson@pica.army.mil
E. Paul Love	loveep@sdsc.edu
Gary Malkin	gmalkin@xylogics.com
Janet Marcisak	jlm@ftp.com
Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Charlotte Mooers	mooers@nnsc.nsf.net
Mark Needleman	mhn@stubbs.ucop.edu
Pete Percival	percival@indiana.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Richard Rodgers	rodgers@nlm.nih.gov
Anthony Rutkowski	amr@sprint.com
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Joo Young Song	jysong@ring.kotel.co.kr
Simon Spero	simon_spero@unc.edu
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Evan Wetstone	evan@rice.edu
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.9.4 Internet School Networking (isn)

Charter

Chair(s):

John Clement, clement@educom.edu
Arthur St. George, stgeorge@bootes.unm.edu
Connie Stout, cstout@tenet.edu

Mailing Lists:

General Discussion: cosndisc@bitnic.educom.edu
To Subscribe: listserv@bitnic.educom.edu (Sub: cosndisc)
Archive: listserv@bitnic.educom.edu

Description of Working Group:

The Internet School Networking Working Group is chartered to facilitate the connection of the United States' K-12 (Kindergarten-12th Grade) schools, public and private, to the Internet, and school networking in general.

It is critically important that national networking for K-12 education proceed along established lines of protocol, using existing network structures. The Working Group's first priority will be to establish guidelines for specialized user interfaces. K-12 networking will also require other support services, such as directories, online and hotline help, specialized training programs and collaborative projects with instructional and curriculum groups, disciplinary groups and postsecondary institutions.

While the initial focus is school networking in the U.S., the Working Group will coordinate its efforts with similar activities in other countries and regions of the world.

Goals and Milestones:

Done	Meet for the first time at IETF and establish approval of Charter. Examine the status of projects in process when Working Group was created. Begin work on list of deliverables.
Jan 1992	Release X.500 "K-12 People Directory" version in collaboration with Merit. Develop plans and milestones for K-12 Resources Directory.
Mar 1992	First draft of information packet document for computing directors to assist them in connecting K-12 schools. First draft of user interface guideline statement.
May 1992	Release X.500 K-12 Resource Directory version in collaboration with Merit. Present final draft guideline statement.

CURRENT MEETING REPORT

Minutes of the Internet School Networking Working Group (ISN)

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

Attendees

J. Nevil Brownlee	nevil@aukuni.ac.uz
Steve Bucey	sabucey@ns.pacbell.com
Robert Ching	rching@nat.com
John Clement	clement@educom.edu
Naomi Courter	naomi@concert.net
Sallie Fellows	sallie%ed@psc.plymouth.edu
Barbara Fraser	byf@cert.org
Joan Gargano	jcgargano@ucdavis.edu
Jack Hahn	hahn@sura.net
Eugene Hastings	hastings@psc.edu
J. Paul Holbrook	holbrook@cic.net
Lenore Jackson	jackson@nsipo.nasa.gov
Barbara Jennings	bjjenni@sandia.gov
Mark Kosters	markk@nic.ddn.mil
Daniel Long	long@nic.near.net
Andrew Maffei	arm@aqua.whoi.edu
Dan Magorian	magorian@ni.umd.edu
Cynthia Mills	cmills@nnsc.nsf.net
Mitra	mitra@pandora.sf.ca.us
Kate O'Mara	kate@acfcluster.nyu.edu
Frank Odasz	franko@bigsky.dillon.mt.us
Marsha Perrott	mlp+@andrew.cmu.edu
Joseph Ramus	ramus@nersc.gov
Joyce K. Reynolds	jkrey@isi.edu
Richard Rodgers	rodgers@nlm.nih.gov
Anthony Rutkowski	amr@sprint.com
Richard Schmalgemeier	rgs@merit.edu
Jennifer Sellers	sellers@nsipo.arc.nasa.gov
Joo Young Song	jysong@ring.kotel.co.kr
Michael St. Johns	stjohns@darpa.mil
Tang Tang	tt@virginia.edu
John Vollbrecht	jrv@merit.edu
Moira West	mjw@cert.org
Gerry White	gerry@lancity.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov
John Zalubski	zalubski@nic.ddn.mil

2.9.5 Internet User Glossary (userglos)

Charter

Chair(s):

Tracy LaQuey Parker, tracy@utexas.edu
Gary Malkin, gmalkin@xylogics.com

Mailing Lists:

General Discussion: usergloss@xylogics.com
To Subscribe: usergloss-request@xylogics.com
Archive: xylogics.com:gmalkin/usergloss/usergloss-arc

Description of Working Group:

The Internet User Glossary Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
Done	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet-Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.
Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the Charter as necessary. Discussion of Userglos Working Group Chair nominations submitted by

Internet Drafts:

"Internet Users' Glossary", 10/01/1992, G. Malkin, T. Parker <draft-ietf-userglos-glossary-00.txt>

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the Internet User Glossary Working Group (USERGLOS)

Agenda

- Review the Charter and previous meetings Minutes.
- Review comments sent by those who won't be able to attend.
- Page-by-page review.
- Close Working Group.

The Internet User Glossary Working Group met to take care of some final editorial adjustments to the Internet-Draft. A couple of entries were removed and a couple were added. The edited Internet-Draft will be sent to the RFC Editor in December.

Most editorial changes were accepted as read. Some entries were replaced by definitions from "authoritative" sources. Overall, the format and quality of the document were accepted by all those who made comments.

Several people volunteered to submit new and replacement entries. Gary Malkin will incorporate those changes into the final version which will be submitted to the RFC Editor in December 1992.

Everyone agreed that the Working Group accomplished the goals set forth in its Charter and should therefore disband.

Attendees

Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Alisa Hata	hata@cac.washington.edu
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Edward Krol	e-krol@uiuc.edu
Tracy LaQuey Parker	tracy@utexas.edu
Gary Malkin	gmalkin@xylogics.com
Janet Marcisak	jlm@ftp.com
Charlotte Mooers	mooers@nnsc.nsf.net
Marsha Perrott	mlp+@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.9.6 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@reston.ans.net
Darren Kinley, kinley@crim.ca

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu
Archive:

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.

- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.

- Identify additional new or improved tools that have become apparent since the last compilation of the reference document.

- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.

- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Done	Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.
Done	Draft of catalog will be prepared, draft to be reviewed and modified. Initiate IETF Internet-Draft review process by submission of a "Son of NOCTools" catalog draft to IESG Secretary.
Done	Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI RFC for publication.

Internet-Drafts:

"FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices", 11/11/1992, R. Enger, J. Reynolds <draft-ietf-noctool2-debug-tcpip-00.txt>

2.9.7 Network Information Services Infrastructure (nisi)

Charter

Chair(s):

April Marine, april@nisc.sri.com
Pat Smith, psmith@merit.edu

Mailing Lists:

General Discussion: nisi@merit.edu
To Subscribe: nisi-request@merit.edu
Archive:

Description of Working Group:

The NISI Working Group will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. The work will begin with existing NIC functions and services and should build upon work already being done within the Internet community. It should address areas such as common information formats, methods of access, user interface, and issues relating to security and privacy of Internet databases.

Goals and Milestones:

Done	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.

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 Pat Smith/Merit

Minutes of the Network Information Services Infrastructure Working Group (NISI)

Agenda

1. nic-profiles
2. net-help, immediate and future
3. Future of NISI.

Thirty-Three people attended the NISI session. April Marine (co-Chair) was unable to attend.

Nic-profiles

History: the nic-profiles concept was conceived at the St. Louis meeting in March 1991. The idea behind nic-profiles is to obtain and maintain key pieces of information about NICs around the world that will be easily accessible to all the NICs.

A mail server is in place to add/modify/retrieve nic-profiles information. Send to x500test@merit.edu and put "help" in the subj: field to receive further information.

This project is now an action item at the top of the NISI priority list. The near-term goal is to have information on as many NICs as possible by the March IETF.

Suggestions were made that we make every effort to utilize any collections of information that may already exist and that we try to approach the "right" people when attempting to obtain this information so that the request doesn't get tossed in the black hole.

Three volunteers, Charlotte Mooers (NNSC/BBN), Marsha Perrott (PREPnet) and Bill Yurcik (NASA) will work with Pat Smith (Merit) on proactive outreach to discover as many NICs as possible, group them in some logical way, and then proceed with making the contacts to interest them in filling out a template and submitting it to nic-profiles electronically.

Net-help

Current plans. Extensive discussion took place regarding how we are going to keep net-help as simple as possible. A sampling of questions/comments:

- Where will the information coming from?

- The contents of the help screen should be managed locally. Each site should be able to control where they send their users.

- How "smart" will net help be?

- What about dial-ins where the IP is dynamically assigned?

After discussion of these and other points, Ed Kroll volunteered to write an "implementation requirements" brief document as a starting point. Ed will send his draft out to the list. The plan is that the Group will hack on it at the next meeting and, hopefully, come up with agreed-upon written requirements which can then be used when working with the applications folks and working toward a feasible implementation plan.

The other very brief paper will summarize the nethelp plan as it stands now, including the screen that was agreed on at the July '92 meeting. Jack Hahn (SURAnet) volunteered to write this one.

There was some further discussion on how nethelp will be enhanced in the future.

Future of NISI

Again there was a short of period of discussion regarding the future of NISI. The Group consensus is that NISI is needed and will continue as a viable Working Group.

Attendees

Steve Bucey	sabucey@ns.pacbell.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
James Conklin	jbc@bitnic.educom.edu
Naomi Courter	naomi@concert.net
Tim Dixon	dixon@rare.nl
Greg Dobrich	gdobrich@nalusda.gov
Roger Fajman	raf@cu.nih.gov
Jill Foster	jill.foster@newcastle.ac.uk
Jack Hahn	hahn@sura.net
Deborah Hamilton	debbie@qsun.att.com
Alf Hansen	Alf.Hansen@delab.sintef.no
Alisa Hata	hata@cac.washington.edu
Masaki Hirabaru	hi@nic.ad.jp
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Laura Kelleher	lak@merit.edu
Edward Krol	e-krol@uiuc.edu
Gary Malkin	gmalkin@xylogics.com
Janet Marcisak	jlm@ftp.com
Bob McCollum	bobm@nic.ddn.mil

Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Charlotte Mooers	mooers@nnsc.nsf.net
Pete Percival	percival@indiana.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Joseph Ramus	ramus@nersc.gov
Joyce K. Reynolds	jkrey@isi.edu
Richard Rodgers	rodgers@nlm.nih.gov
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Scott Williamson	scottw@nic.ddn.mil
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.9.8 Networked Information Retrieval (nir)

Charter

Chair(s):

Jill Foster, jill.foster@newcastle.ac.uk
George Brett, George.Brett@cnidr.org

Mailing Lists:

General Discussion: nir@cc.mcgill.ca
To Subscribe: nir-request@cc.mcgill.ca
Archive: pub/mailing-lists/nir:archives.cc.mcgill.ca

Description of Working Group:

As the network has grown, along with it there has been an increase in the number of software tools and applications to navigate the network and make use of the many, varied resources which are part of the network. Within the past year and a half we have seen a wide spread adoption of tools such as the ARCHIE servers, the Wide Area Information Servers (WAIS), the Internet Gopher, and the WorldWide Web (WWW). In addition to the acceptance of these tools there are also diverse efforts to enhance and customize these tools to meet the needs of particular network communities.

There are many organizations and associations that have recently begun to focus on the proliferating resources and tools for networked information retrieval (NIR). The Networked Information Retrieval Group will be a cooperative effort of three major players in the field of NIR: IETF, RARE, and the Coalition for Networked Information (CNI) specifically tasked to collect and disseminate information about the tools and to discuss and encourage cooperative development of current and future tools.

The NIR Working Group intends to increase the useful base of information about networked information retrieval (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 NIRtools, 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.

CURRENT MEETING REPORT

Minutes of the Networked Information Retrieval Working Group (NIR)

Report not submitted. This Group is jointly chartered under the User Services/Applications Areas. Please refer to the Applications Area Report for a brief summary.

Attendees

Vikas Aggarwal	vikas@jvnc.net
George Brett	George.Brett@cnidr.org
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Peter Deutsch	peterd@bunyip.com
Tim Dixon	dixon@rare.nl
Alan Emtage	bajan@bunyip.com
Roger Fajman	raf@cu.nih.gov
Richard Fisher	rfisher@cdhf1.gsfc.nasa.gov
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Laura Kelleher	lak@merit.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Edward Krol	e-krol@uiuc.edu
David LeRoy	dleroy@mitchell.cit.cornell.edu
Edward Levinson	levinson@pica.army.mil
Gary Malkin	gmalkin@xylogics.com
Janet Marcisak	jlm@ftp.com
Michael Mealling	michael@fantasy.gatech.edu
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Hilarie Orman	ho@cs.arizona.edu
Jim Perchik	perchik@athena.mit.edu
Pete Percival	percival@indiana.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Alan Roszkiewicz	alan@sprint.com
Srinivas Sataluri	sri@qsun.att.com
Jane Smith	jds@jazz.concert.net
Larry Snodgrass	snodgras@bitnic.educom.edu
Simon Spero	simon_spero@unc.edu
Claudio Topolcic	topolcic@cnri.reston.va.us
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Janet Vratny	janet@apple.com

Chris Weider	`clw@merit.edu`
Yung-Chao Yu	`yy@qsun.att.com`
William Yurcik	`yurcik@dftnic.gsfc.nasa.gov`

2.9.9 Uniform Resource Identifiers (uri)

Charter

Chair(s):

Jim Fullton, jim_fullton@unc.edu
Alan Emtage, bajan@bunyip.com

Mailing Lists:

General Discussion: uri@bunyip.com
To Subscribe: uri-request@bunyip.com
Archive: pub/uri-archive@archives.cc.mcgill.ca

Description of Working Group:

The Uniform Resource Identifiers Archives Working Group is chartered to define a set of standards for the encoding of system independent Resource Location and Identification information for the use of Internet information services.

This Working Group is expected to produce a set of documents that will specify standardized representations of Uniform Resource Locators (URLs) which specify a standardized method for encoding location and access information across multiple information systems. Such standards are expected to build upon the document discussed at the UDI BOF session held during the 24th IETF meeting in Boston, Unique Resource Serial Numbers (URSNs) which specify a standardized method for encoding unique resource identification information for Internet resources and Uniform Resource Identifiers (URIs), which specify a standardized method for encoding combined resource identification and location information systems to be used for resource discovery and access systems in an Internet environment.

Such a set of standards will provide a framework that: allows the Internet user to specify the location and access information for files and other resources on the Internet, allows users and network-based tools to uniquely identify specific resources on the Internet, and allows the creation and operation of resource discovery and access systems for the Internet. The security of such resource discovery services will also be considered to be an integral part of the work of this Group.

Goals and Milestones:

Done	Review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended documents. Review the first draft of a proposal for Uniform Resource Locators already available.
Mar 1993	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.

CURRENT MEETING REPORT

Reported by Alan Emtage/Bunyip

Minutes of the Uniform Resource Identifiers Working Group (URI)

The Agenda for the first meeting of the URI Working Group was approved. The Charter for the Group was reviewed and approved. It was noted that the "Goals and Milestones" may need to be changed in the future depending on the progress in this very new area.

Peter Deutsch/Bunyip who was initially named to co-Chair the Group resigned from the position in order to follow a more activist role and avoid any potential conflict of interest. Jim Fullton/CNIDR was installed as new co-Chair. Before stepping down, however, Peter took the opportunity to make a few personal observations and commitments:

- Peter has offered to co-author an overview paper along with Chris Weider. This paper would propose a possible architecture to the Group describing the use and the form of the various Uniform Resource objects such as URI's (Uniform Resource Identifiers), URL's (Uniform Resource Locators) and URSN (Uniform Resource Serial Numbers) and how the would interoperate.
- Peter gave a basic overview of his ideas about what the UR objects looked liked. By his definitions:
  - A URL identifies a particular object on the network and is composed of a named scheme (e.g., FTP, WAIS, Gopher) and information specific to that scheme. It was noted that this idea already exists in a similar form in the World Wide Web (WWW) system, and has been codified in a paper by Tim Berners-Lee/CERN.
  - A URSN can be broken down into a "virtual user" and an actual serial number. Related topics were the issue of the "producer" of an network object and the "owner"; some possible schemes for implementation of the virtual user (whois++ handle, X.500) ; and what the serial number would looked like (possibly an MD5 checksum and other methods).

  It was decided in the interests of time that further discussions should be carried out on the mailing list.

The paper currently titled "Universal Resource Locators" by Tim Berners-Lee was reviewed and the following comments were made:

- The use of the term "protocol" in the document is ambiguous given the context of the IETF and should be replaced or more specifically defined.
- The use of the term "name" was considered to be unclear and again should be clarified. It was suggested that it be removed and another term used in its place.

- The document should be written as a "standalone" unit. However, the objects described therein should be viewed as part of a larger architecture and an explicit description of their purpose should be added. It was suggested that the document could be further generalized from a _perceived_ WWW bias.

- The question of the "partial form" of the URL brought heated discussion between two factions: one which wanted the removal of the form altogether and one which suggested their continued existence with restrictions. Some consensus developed around the idea that partial forms could be used internally for individual information systems but should not be used when exchanged externally. It was decided that further discussion should occur on the mailing list.

- Consensus was reached that the document should specifically state URLs are to be considered transient and should not be used in static objects (hardcopy documents, etc.). Their use as references should be specifically discouraged. Such references were considered to be in the domain of the URSN, whatever they ultimately look like.

- The paper should describe the general scheme being proposed without reference to particular systems (other than as examples). All detailed descriptions of individual systems should be put in an appendix. It was decided that the most likely repository for the individual definitions would ultimately be the Internet Assigned Numbers Authority (IANA) but that the original document may propose the definitions for a basic range of services (such as FTP).

- It was suggested by Thomas Hacker/UMich that they use the OSF DCE DFS (Open Software Foundation Distributed Computing Environment Distributed File System).

- Mitra proposed a "fragment specifier" scheme to be incorporated into the URL document. It was decided that detailed discussion of this was best left to the mailing list.

- Other points were:
  - Some of the text and examples did not agree.
  - The use of percentage signs should be reviewed on the mailing list.
  - Use of blank characters was again questioned.

  All were referred back to the mailing list for further discussion.

A discussion about URI's followed. The questions that were raised were:

1. Given the current definitions what **exactly** does URI mean?
   - Alan Emtage suggested that they may be defined as URI = URL + URSN + "Uniform Resource Representator" (URR) since the current definitions of URL and URSN do not give sufficient information for a user/client to determine if in fact the information available is useful and that such things as filename

extensions are not a reliable method of determining content format (and in the case of processes is meaningless). However he declined to be committed on what exactly these URR's would look like.

- It was suggested that the concept of the "URI" may be defunct now since it as been decomposed into several constituent parts.

2. The proposal that John Kunze/UCBerkeley had made on the mailing list previously was briefly discussed and it was suggested that he and Clifford Lynch/UC co-author an alternate document to that produced by Peter Deutsch and Chris Weider, more from the perspective of the library community. John's proposal for access lists, descriptive fields, functional types and a "UR Citation" were suggested as being better handled in detail on the mailing list.

3. In addition to the document describing the general UR system, Peter Deutsch and Chris Weider have agreed to co-author a paper proposing the structure of URSN's.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@bunyip.com
Alan Emtage	bajan@bunyip.com
Jill Foster	jill.foster@newcastle.ac.uk
Joan Gargano	jcgargano@ucdavis.edu
Thomas Hacker	hacker@citi.umich.edu
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
J. Paul Holbrook	holbrook@cic.net
Ole Jacobsen	ole@interop.com
Edward Krol	e-krol@uiuc.edu
John Kunze	jak@violet.berkeley.edu
Clifford Lynch	calur@uccmvsa.ucop.edu
Janet Marcisak	jlm@ftp.com
Michael Mealling	michael@fantasy.gatech.edu
Mitra	mitra@pandora.sf.ca.us
Charlotte Mooers	mooers@nnsc.nsf.net
Mark Needleman	mhn@stubbs.ucop.edu
Kate O'Mara	kate@acfcluster.nyu.edu
Pete Percival	percival@indiana.edu
Joyce K. Reynolds	jkrey@isi.edu
Bradley Rhoades	bdrhoades@mail.mmmg.com
Richard Rodgers	rodgers@nlm.nih.gov

Jennifer Sellers	sellers@nsipo.arc.nasa.gov
Jane Smith	jds@jazz.concert.net
Simon Spero	simon_spero@unc.edu
Craig Summerhill	craig@cni.org
Claudio Topolcic	topolcic@cnri.reston.va.us
Janet Vratny	janet@apple.com
Chris Weider	clw@merit.edu
Moira West	mjw@cert.org
Yung-Chao Yu	yy@qsun.att.com

2.9.10 User Documents Revisions (userdoc2)

Charter

Chair(s):

Ellen Hoffman, ellen_hoffman@um.cc.umich.edu
Lenore Jackson, jackson@nsipo.nasa.gov

Mailing Lists:

General Discussion: user-doc@nnsc.nsf.net
To Subscribe: user-doc-request@nnsc.nsf.net
Archive:

Description of Working Group:

The focus of the USER-DOC2 Working Group is on identifying and locating documentation about the Internet. A major activity is the revision of an existing bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "How to use the Internet" (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 USER-DOC2 Working Group will:

(1) Identify and categorize useful documents, reference materials, training tools, and other publications about the Internet, particularly those available on-line.

(2) Publish on-line and hard copies of the bibliography(s) produced and other reference material on documentation as needs are identified.

(3) Develop and implement procedures to maintain and update the bibliography and investigate methods to provide the information in an on-line format.

(4) As a part of the update process, identify new materials for inclusion into the active bibliography and identify additional needs which are required for locating documentation and other publications.

(5) Review procedures for periodic review of the bibliography by the User Services Working Group.

(6) Examine methods for delivering documentation and work with providers to improve the availability of basic Internet documentation.

Goals and Milestones:

Done	Identify new "sources of information" (e.g., individuals, mailing lists, bulletins, etc.) Review existing document and obtain comments from others in USWG about needed revisions at the San Diego IETF.

Done	Publish Internet-Draft of the short bibliography for novice users.
Done	Submit the revised FYI document to the IESG for publication as an RFC.
Jan 1993	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.

Internet-Drafts:

"FYI on Introducing the Internet–A Short Bibliography of Introductory Internetworking Readings for the Network Novice", 09/25/1992, Ellen Hoffman, Lenore Jackson <draft-ietf-userdoc2-fyi-novice-01.txt>

CURRENT MEETING REPORT

Minutes of the User Document Revisions Working Group (USERDOC2)

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

Attendees

Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Roger Fajman	raf@cu.nih.gov
Jill Foster	jill.foster@newcastle.ac.uk
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Lenore Jackson	jackson@nsipo.nasa.gov
Laura Kelleher	lak@merit.edu
Edward Krol	e-krol@uiuc.edu
Janet Marcisak	jlm@ftp.com
Marsha Perrott	mlp+@andrew.cmu.edu
Patricia Smith	psmith@merit.edu
Janet Vratny	janet@apple.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

2.9.11 User Services (uswg)

Charter

Chair(s):

Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:

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

Description of Working Group:

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

(1) Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should:

- Clearly address user assistance needs; - Produce an end-result (e.g., a document, a program plan, etc.); - Have a reasonably clear approach to achieving the end-result (with an estimated time for completion); - Not duplicate existing or previous efforts.

(2) Create working groups or other focus groups to carry out projects deemed worthy of pursuing.

(3) Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

Ongoing This is an oversight group with continuing responsibilities.

Request For Comments:

RFC 1150 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. notes"

RFC 1177 "FYI on Questions and Answers - Answers to Commonly Asked "New Internet User" Questions"

RFC 1206 "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions"

RFC 1207 "Answers to Commonly asked "Experienced Internet User" Questions"

RFC 1325 "FYI on Questions and Answers Answers to Commonly asked "New Internet User" Questions"

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI

Minutes of the User Services Working Group (USWG)

At this IETF session announcements included working groups coming to closure (DISI (will be revised with a new Charter and co-Chair), IAFA, NOCTOOL2, and User-Gloss) and new working groups starting up (WNILS, NIR, URL, IIIR). The four new working groups will be a joint effort of the Applications and User Services areas. New FYI RFC publications, current user services related Internet-Drafts postings include:

Internet-Drafts

- draft-malkin-newcomers-guide-00.txt
- draft-ietf-noctool2-debug-tcpip-00.txt
- draft-ietf-userdoc2-fyi-novice-01.txt
- draft-ietf-userglos-glossary-00.txt

FYI RFCs

- FYI 16 "Connecting to the Internet: What Connecting Institutions Should Anticipate", (Also RFC 1359), August 1992.

- FYI 15 "Privacy and Accuracy Issues in Network Information Center Databases", (Also RFC 1355), August 1992.

Jill Foster provided an update on RARE ISUS activities. Joyce Reynolds provided reports on RIPE activities, and on NETF activities.

Steve Coya led a discussion on the Internet-Draft, "draft-malkin-newcomers-guide-00.txt" and requested comments from the USWG. This document is specifically for the new IETF attendee. It has been estimated that as many as 30% of the IETF attendees is now comprised of first time participants. The User Services area is assisting the IETF Secretariat in this documentation.

Peter Deutsch led a discussion on a USWG project in development called, "Internet Quick and Dirty". It is intended to be a short document on descriptions of each network service with pointers on where to obtain additional information. A draft was posted to the USWG mailing list. It is currently eight pages and broken into sections, including a basic introduction and useful collections of data. But, is it too long??

Gary Malkin requested new volunteers to help update the FAQ (Frequently Asked Questionis) for New Internet Users and the FAQ for Experienced Internet Users.

FYI16/RFC1359 was discussed, primarily focusing on expanding and updating the document. General comments were that this document is too U.S. centric, focusing specifically at the University level. The USWG intends to work with Martyne Hallgren and ACM SIGUCCS to take apart this document, restart, and expand to make it as global as possible.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Naomi Courter	naomi@concert.net
Stephen Coya	scoya@cnri.reston.va.us
Michael Davis	mad@spirit.clearpoint.com
Peter Deutsch	peterd@bunyip.com
Susan Estrada	estradas@cerf.net
Sallie Fellows	sallie%ed@psc.plymouth.edu
Jill Foster	jill.foster@newcastle.ac.uk
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
James Geddes	wk05020@worldlink.com
Jack Hahn	hahn@sura.net
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Masaki Hirabaru	hi@nic.ad.jp
Ellen Hoffman	ellen_hoffman@um.cc.umich.edu
Laura Kelleher	lak@merit.edu
Mark Kosters	markk@nic.ddn.mil
Edward Krol	e-krol@uiuc.edu
Tracy LaQuey Parker	tracy@utexas.edu
Gary Malkin	gmalkin@xylogics.com
Janet Marcisak	jlm@ftp.com
Charlotte Mooers	mooers@nnsc.nsf.net
Joyce K. Reynolds	jkrey@isi.edu
Jane Smith	jds@jazz.concert.net
Patricia Smith	psmith@merit.edu
Simon Spero	simon_spero@unc.edu
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Janet Vratny	janet@apple.com
Peter Will	will@isi.edu
Scott Williamson	scottw@nic.ddn.mil
Yung-Chao Yu	yy@qsun.att.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov
John Zalubski	zalubski@nic.ddn.mil

2.9.12 Whois and Network Information Lookup Service (wnils)

Charter

Chair(s):

Joan Gargano, jcgargano@ucdavis.edu

Mailing Lists:

General Discussion: ietf-wnils@ucdavis.edu
To Subscribe: ietf-wnils-request@ucdavis.edu
Archive: pub/ietf-wnils-archive@ucdavis.edu

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.
Dec 1992	Submit the Whois and Network Information Lookup Service Recommendations document to the IESG as an Informational document.
Dec 1992	Post a revised WHOIS protocols specification as an Internet-Draft.
Dec 1992	Submit the revised WHOIS protocol documents to the IESG as Draft Standards.

Internet-Drafts:

"Architecture of the Whois++ Index Service", 11/23/1992, C. Weider, J. Fullton, S. Spero <draft-ietf-wnils-whois-00.txt>

CURRENT MEETING REPORT

Reported by Joan Gargano/UCDavis

Minutes of the Whois and Network Information Lookup Service Working Group (WNILS)

Introduction and Overview

Joan Gargano spoke briefly about the impetus for starting the Working Group and development of the Whois++ application. This was followed by presentations by the developers.

Architecture of Whois++

Peter Deutsch presented an overview of the Whois++ architecture. Peter described the prototype system, the currently supported query syntax and areas for improvement.

Centroids

Chris Weider provided an overview of the mechanism used to build a distributed directory service called "centroids". Simon Spero continued with a discussion of the underlying theory of centroids which has been used by other groups studying information retrieval issues.

Sample Server

Jim Fullton presented a brief overview of the Whois++ server developed at CNIDR. Jim describe his system which was built using standard Unix utilities. Response times of 15 seconds to query a database of 30,000 records was reported, however Jim felt significant improvements could be achieved with software designed and coded to optimize the system. Servers for testing are available at CNIDR and UC Davis. It is anticipated about ten new servers will be in place within a couple of weeks.

Questions

The floor was opened for questions and general discussion of the protocol.

Future

The Whois++ developers solicited input about desired features. The following work was recommended.

- A feature to allow servers to pass information about servers that poll them to optimize searching.
- A description of printer output format.

- There needs to be a clearinghouse for templates. It was recommended this be performed by CNIDR.
- Privacy and security features.
- A method for handling replication of services.
- Further discussion about usernames and attributes.
- A method for using synonyms with queries.
- A method to abort searches and continue searches.
- Provide the ability to store and pass images.
- A way to limit the attributes returned by a search.
- A method for tagging attributes.

Attendees

Jules Aronson	aronson@nlm.nih.gov
Anthony Ballardie	A.Ballardie@cs.ucl.ac.uk
Steve Bucey	sabucey@ns.pacbell.com
Peter DiCamillo	Peter_DiCamillo@brown.edu
Letha Dugas	4371362@mcimail.com
Cliff Frost	cliff@cmsa.berkeley.edu
Jim Fullton	jim_fullton@unc.edu
Joan Gargano	jcgargano@ucdavis.edu
Masayoshi Gohara	mg@sinet.ad.jp
Deborah Hamilton	debbie@qsun.att.com
Alisa Hata	hata@cac.washington.edu
Russ Hobby	rdhobby@ucdavis.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
John Kunze	jak@violet.berkeley.edu
Edward Levinson	levinson@pica.army.mil
Clifford Lynch	calur@uccmvsa.ucop.edu
Mitra	mitra@pandora.sf.ca.us
Charlotte Mooers	mooers@nnsc.nsf.net
Clifford Neuman	bcn@isi.edu
Pete Percival	percival@indiana.edu
Karen Petraska-Veum	karen@nsisrv.gsfc.nasa.gov
Sheri Repucci	smr@merit.edu
Joyce K. Reynolds	jkrey@isi.edu

Richard Rodgers	rodgers@nlm.nih.gov
Patricia Smith	psmith@merit.edu
Larry Snodgrass	snodgras@bitnic.educom.edu
Janet Vratny	janet@apple.com
Yung-Chao Yu	yy@qsun.att.com
William Yurcik	yurcik@dftnic.gsfc.nasa.gov

Chapter 3

Network Status Briefings

3.1 EBONE

Bernhard Stockman/SUNET

EBONE STATUS REPORT

EAT AND EOT HAVE MEET SEVERAL TIMES DURING SPRING AND SUMMER

FOUR SPECIFICATIONS READY

- IMPLEMENTATION PLAN
- EBS TECHNICAL DESCRIPTION
- ROUTING PLAN
- OPERATION PLAN

FAST PROGRESS OF IMPLEMENTATION
LAST LINE IN PLACE SEPTEMBER 23

ROUTING AT LEADING EDGE OF TODAY TECHNOLOGY

KTH, STOCKHOLM EBONE NOC

MOVEMENT OF FRANCE EBS TO PARIS

BONN AND WIEN PROPOSED TO BE NEW EBONE BORDER SYSTEMS

UPGRADE OF ROUTERS NECESSARY

SECURITY ACCESS SCHEME INSTALLED IN THE ROUTERS

EBONE CLNS SPECIFICATION PRODUCED BY RARE/COSINE CLNS PROJECT

INFORMATION DOCUMENT NEEDED
EBS-RBS TUTORIAL DOCUMENT BEING WORKED ON

SOME STATISTICS PRODUCED

EBONE

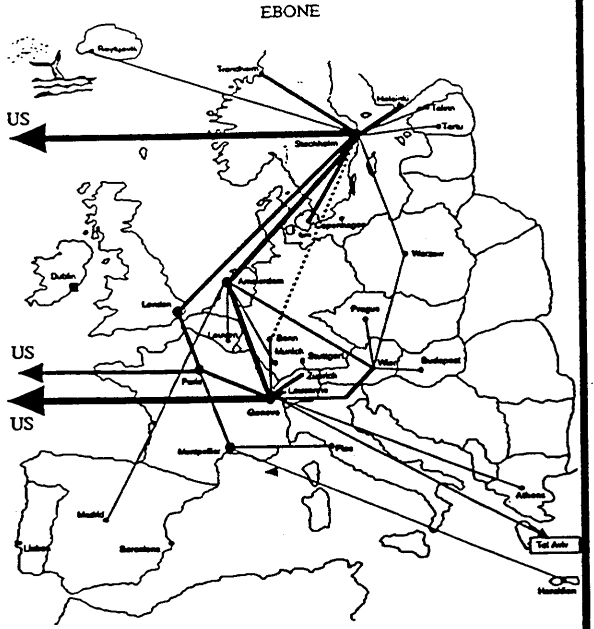

EBONE TODAY

STOCKHOLM - LONDON 256 KBPS

LONDON - MONTPELLIER 256 KBPS

LONDON - PARIS 256 KBPS ON ORDER

STOCKHOLM - AMSTERDAM 512 KBS

AMSTERDAM - GENEVA 512 KBPS

GENEVA - MONTPELLIER 256 KBS

GENEVA - PARIS 256 KBPS

EBONE - IXI GATEWAYS IN AMSTERDAM AND LONDON

EBONE NOC IN STOCKHOLM

ONGOING REFINEMENTS OF ROUTING PLAN

AMSTERDAM	SURFNET (Netherlands) REDIRIS (Spain) EUNET (Europe) IXI (Europe) ACONET (Austria) RCCN (Portugal) ECRC (Germany) KUL (Belgium)
GENEVA	SWITCH (Switzerland) CERN (Cern) EARN (Europe) ARIADNET (Greece) ACONET (Austria) ILAN (Israel) EASINET (Europe)
LONDON	JANET (United Kingdom) PIPEX (United Kingdom) UKNET (United Kingdom) HEANET (Ireland)
MONTPELLIER	RENATER (France) FORTH (Greece)
STOCKHOLM	NORDUNET (Nordic) DATANET (Finland) SWIPNET (Sweden) TIPNET (Sweden)

EBONE DEVELOPMENT

INTEGRATION OF CLNS PILOT

EXPANSION TO CENTRAL AND EAST EUROPE

INCREASE OF BANDWIDTH IN CORE LINES UREGENTLY NEEDED

FORMAL CONNECTIVITY AGREEMENTS FOR 1993 AND ONWARDS

GENERAL PURPOSE TRANSIT CAPACITY INSTALLED 1993 VIA GLOBAL INTERNET EXCHANGE (GIX)

NEW EBS'S PROPOSED IN BONN AND WIEN

EBONE STRATEGICAL COMMITTEE FORMED TO PROPOSE ORGANISATIONAL AND FINANCIAL STRUCTURES FOR EBONE AFTER 1993. FIRST REPORT IN FEBRUARY 1993.

THE GLOBAL INTERNET EXCHANGE

GIX

THE NEED FOR A TOP LEVEL POLICY-FREE EXCHANGE

HIERARCHICAL MODEL WITH LESS RESTRICTIONS AT TOP

CONNECTIVITY BETWEEN NETWORKS AT SAME LEVEL

ORDERED CONNECTIVITY BETWEEN NETWORKS AT ADJACENT LEVELS

METROPOLITAN FIBER SYSTEM (MFS) INSTALLATION NOW IN WASHINGTON DC.

IEPG MEETING NOVEMBER 15 ON GIX IMPLEMENTATION

3.2 NSI

Reported by Milo Medin/NASA Ames

NASA Science Internet Update

Internet Engineering Task Force Meeting
11/17/92
Washington, DC

Milo S. Medin
medin@nsipo.nasa.gov

Vital statistics

- 114 Proteon routers (45 support DECNET IV + IP)
- 2 Cisco AGS+ routers
- Handful of DEC DECNET IV only routers (being transitioned)
- Multiple T1 backbone (FTS-2000 circuits)
- T3 link between GSFC and FIX-E
- System runs as single OSPF backbone area
- CLNP support in parts of the network
- Managed and supported centrally from Ames Research Center

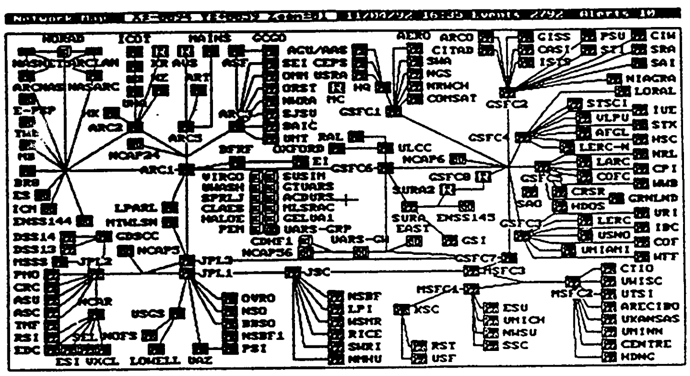

DECNET IV support

- Replacement of Digital hardware with standard multiprotocol routers
- Elimination or upgrade of tail circuits
- MultiNet DECNET in TCP encapsulation support
  - 4 VAXsystem 3000 encapsulation hosts deployed
  - Throughput on T1 path over 200 Kbps
  - MultiNet software provided to remote sites
  - Saves cost of leased line, and more robust connectivity provided
- Gradual if any transition to Phase V
- DECNET V DTS and DNS hosts in place

Network Operations

- NOC located at Ames Research Center in dedicated facility
- 24 hr/day x 7 day/week (w/ holidays) operations coverage
- 7 dedicated operations staff (no beepers for ops)
- 1-800-424-9920 hotline number, w/ international toll free access
- Out of band access to routers
- Migration to Remedy Trouble Ticketing system (ARS)
- Goal is end to end coverage, even when other nets are involved

Network Management

- Main support from Proteon Overview system
- Being transitioned to DEC MSU product (not MCCI)
  - Supports native DECNET IV management protocol (NICE)
  - SNMP, NICE, etc data imported into Ingres database
  - Displays driven off of database information
- Link utilization statistics produced by local utility, using SNMP and NICE support

NASA Science Internet Project Office
Ames Research Center
Information and Communications Systems Division

Milo Medin	NASA Science Internet Update	11/16/92

OCONUS efforts

- T1 link to Fairbanks, AK (fiber via NPC)
- 512 Kbps link to the UK (shared by NSF for infrastructure)
- 256 Kbps link to CNES, France
- 768 Kbps links (w/ESA) to ESTEC (Netherlands) and ESOC (Germany)
- 56 Kbps link (VSAT) to CTIO, La Serena, Chile
- 56 Kbps link to Sonderstrom, Greenland
- 384 Kbps link to McMurdo base, Antarctica (year-round)

page 6

NASA Science Internet Project Office
Ames Research Center
Information and Communications Systems Division

Milo Medin	NASA Science Internet Update	11/16/92

OCONUS efforts

- PACCOM

  T1 link to U. of Hawaii

  512Kbps to AARNet, Australia

  64 Kbps link to U. of Waikato, New Zealand

  Multiple links to Japan

  56 Kbps to KAIST, Korea

  56 Kbps to Chinese University, Hong Kong

page 7

NASA Science Internet Project Office
Ames Research Center
Information and Communications Systems Division

Milo Medin	NASA Science Internet Update	11/16/92

Future efforts

- New links overseas, including Russia
- 800 v.32bis dialup support (w/SLIP, PPP, etc...)
- Working with INMARSAT for messaging relay and dial-up access
- ATM integration as part of NREN effort at five sites in FY 93
- Upgrade of FIX-W and FIX-E to FDDI

page 8

Chapter 4

IETF Protocol Presentations

4.1 PIP

Presented by Paul Tsuchiya/Bellcore

Bio:*Mr. Tsuchiya is currently a Member Technical Staff in Bellcore's Information Networking Research Laboratory. He is also currently a PhD student at the University College London. He has been researching routing and addressing problems for seven years. He is the author of several RFCs and Internet Drafts, covering, among other things, the topics of subnet number assignment, shortcut routing, and network address translation.*

PIP (P Internet Protocol) using EIP (Extended Internet Protocol)

Paul F. Tsuchiya

Bellcore

PIP STATUS

- Not mature enough to be chosen at this time
  - Needs solid year of design/implementation/testing
- Is it worth waiting for?

PIP STATUS (cont)

- Done so far.......
  - Pip Forwarding Spec (Tsuchiya)
  - Pip Objects Spec (Tsuchiya)
  - Pip ID Spec (Tsuchiya)
  - EIP Spec (Zheng Wang)
  - EIPIP Spec (Wang, Tsuchiya)
  - Transition Comparison (Wang)
  - IESG Criteria Evaluation
  - Simple implementations of Pip host and router
- Still much to be done

PIP PHILOSOPHY

- Very general
  - All known routing/address paradigms
- But relatively efficient
- Easy to evolve
  - Internet will never stand still

PIP OVERVIEW

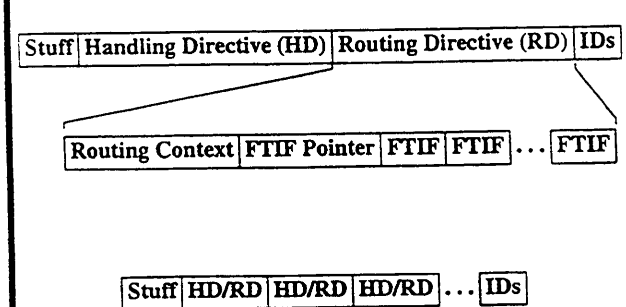

ROUTING DIRECTIVE: BASIC MECHANISM

- Use Routing Context value to pick forwarding table
- Use active FTIF to index into forwarding table
- Forward packet

FTIF CHAIN EXAMPLE: Inter-domain Hierarchical Address *(C to Z)*

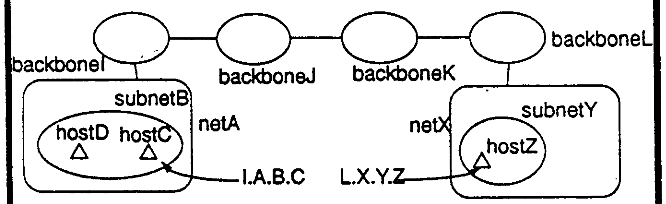

FTIF Pointer	Source Address				Destination Address			
5	C	B	A	I	L	X	Y	Z

FTIF CHAIN EXAMPLE: Policy Route (C to Z)

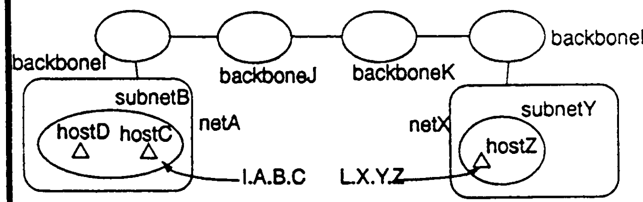

					Policy Route					
4	C	B	A	I	J	K	L	X	Y	Z
	Source						Destination			

FTIF CHAIN EXAMPLE: Intra-domain Hierarchical Address (C to D)

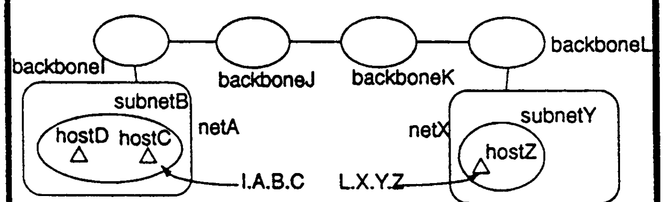

2	C	D

ROUTING CONTEXT

- Treated as flat field by router
- In reality composed of multiple fields
  - QOS
  - Address type (unicast, multicast, family)
  - Hierarchical level
  - etc.
- Each router determines how RC is formatted
  - RC modified be each router to conform to next-hop router's format

ROUTING CONTEXT (cont)

- RC modification also a powerful tagging mechanism
  - For instance, tag packet to take alternate path
- Neighbor Pip systems exchange RC (and HD) format information
  - Globally understood "Pip Objects" are mapped into RC (and HD) contents
  - Useful for evolution

HANDLING DIRECTIVE

- Used only to determine handling of packet (not routing)
  - Queueing priority
  - Drop priority
  - Congestion experienced bit
  - Flow Descriptor
- HD mechanism similar to RC
  - Contents and format of HD "field" determined locally
  - Contents and format dynamically exchanged by neighbor Pip systems
- Therefore also used for tagging

ID FIELDS

- Ignored by routers
- Used as flat fields by hosts to identify source and destination of packet
- But have hierarchical structure none-the-less
  - For ease of administrative assignment
  - Perhaps for other purposes such as reverse DNS lookup

TUNNELING

- HD and RD can be stacked
- Hosts and routers understand tunneling
- Useful for....
  - Evolution
  - Firewalling
  - Etc.

INTER-DOMAIN ADDRESSING ISSUE

- Geographic addressing forces topological restrictions (or doesn't scale)
- Provider-based addressing forces changes in subscriber's addresses
- Pip allows for provider-based addressing without requiring changes in subscriber's host or router addresses
  - Only directory service and border routers need know inter-domain address part
- Intra-domain Routing Directive doesn't include inter-domain address part
- Border routers fill in inter-domain address part as packet exits border

PIP TRANSITION

- Option 1: New Pip header
  - Transition similar to CLNP
- Option 2: IPAE scheme
  - Transition similar to IPAE/Sip
- Proposed method is Extended IP scheme (EIP)
  - Pip over EIP is called EIPip

EIPIP

- "Pip" part of header is placed in what looks like an IP Option field
- To non-EIPip system, packet looks like IP header with unknown options field
- While IP addresses still unique (transition period)
  - No changes to ARP/RARP, ICMP, intra-domain routers, FTP, TCP/UDP Checksum
- After transition period
  - No changes to ARP/RARP, ICMP, TCP/UDP Checksum
- Have already sent EIPIP packets to hosts all over the internet

PIP DEMO-WARE (ACTUALLY EIPIP)

- Acknowledge Zheng Wang, Rob Coltun, Mike Deaddio, Bob Smart
  - All wrote code for demo (though not all of their code is in use)
- Very simple
  - Basically "good faith" effort for this meeting
- EIPIP Host
  - Takes IP address from DNS, forms EIPIP header with network numbers in level 1 FTIFs and IP address in level 0 FTIFs
- EIPIP Router
  - Forwards EIPIP Host output

PIP EVOLUTION: HOSTS

- Hosts have a "dumb" mode
  - HD/RD spoonfed to host by smart box
  - Therefore, when ROAD paradigm changes, old Pip hosts don't need to understand new paradigm
- "Host Version" field in Host Part of Pip header
  - Tells router what algorithms host has implemented (icmp-type messages, etc.)

PIP EVOLUTION: ROUTERS

- Sometimes can change ROAD paradigm without changing routers
  - For instance, hierarchical addressing to policy routing
- Tunneling to get through old-style routers
- Use of dynamic "Pip Object" assignment to ease introduction of new HD/RC
  - Even if router does not "understand" new Pip Object, can still pass it on untouched

4.2 Simple Internet Protocol

Presented by Steve Deering/Xerox PARC

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, for mobility, and for multi-media services.*

SIP — A Simple Internet Protocol

Steve Deering
Xerox PARC

November 1992

A New Version of IP

Some lessons learned from IPv4:

- globally-unique addresses, hierarchically-structured for efficient routing, work well!
  (IPv4 address is just a little too short and a little too flat)

  —> SIP uses 64-bit, hierarchical addresses

- performance can be improved by ignoring or avoiding some inessential header fields

  —> SIP omits some fields, restructures others for improved performance

- bound on IP option length constrains some usage, e.g., source routing, security tagging

  —> SIP relaxes limits on optional data

The SIP Header

← 32 bits →

Vers.	Reserved		
Payload Length		Payload Type	Hop Limit
Source Address			
Destination Address			

Version — IP version 6

Payload Length — length of following data

Payload Type — type of data, e.g., TCP

Hop Limit — like TTL with no notion of time

Source & Destination Addresses...

Fields Omitted from IPv4

← 32 bits →

Vers	HLen	Prec	TOS	Total Length
Identification			flags	Fragment Offset
Time to Live		Protocol		Header Checksum
Source Address				
Destination Address				

Header Length — SIP header is fixed-length
(SIP options have own headers)

Precedence & Type-of-Service
— not used in SIP (may be subsumed by new "Flow ID" field, TBD)

Ident, Flags & Offset — moved to SIP option

Checksum — unnecessary (transport pseudo-header checksum prevents misdelivery)

SIP Addresses

metro-based unicast addresses:

C	country + metro ID	site ID	intra-site part

- geographic hierarchy for scalable routing; needs provider coordination in each metro
- provider not identified in address => can keep addresses when changing provider
- can easily encode all countries + metro IDs in 14 bits, using CIDR-like assignment

SIP Addresses (cont.)

provider-based unicast addresses:

C	country + provider ID	subscriber ID	intra-subscriber part

- similar to current guidelines for NSAP addresses and remaining IPv4 addresses
- for use with uncoordinated providers
- also used for source routing via specific providers, for policy reasons

multicast addresses:

C	mc prefix + scope	group ID

SIP Routing

same routing architecture as IPv4 & CLNP:

- multi-level hierarchy; heterogeneous routing protocols (OSPF, IS-IS, RIP2, BGP, IDRP)
- hop-by-hop routing; "longest-prefix-match" route lookup
- source route option for source-directed routing (incl. routing through specific domains, as well as through specific routers)

compatible with the "Unified" architecture of Estrin, Li, and Rekhter

SIP Source Route

Reserved			
Num Addrs	Next Addr	Payload Type	Reserved
Address[0]			
Address[1]			
⋮			

- identified by a specific Payload Type in SIP header; in turn identifies following Payload Type
- more efficient than IPv4 — examined only by the identified routers, not any intermediate ones
- allows up to 255 addresses

Packet Size & Fragmentation

SIP requires each link to have MTU ≥ 576 (cf. 68 for IPv4)

any higher-layer protocol that wants to send > 576 must do Path MTU Discovery

SIP has a fragmentation option for IPv4 compatibility, and for "tunneling":

Identification			
0 0 M	Fragment Offset	Payload Type	Reserved

more efficient than IPv4 — examined only by the destination system, not by any forwarding systems

Changes to Related Protocols

ICMP

- pseudo-header added to checksum
- error messages return larger piece of erroneous packet
- Redirect & Router Advertisement modified to carry larger addresses

IGMP

- modified to carry larger addresses

transport -layer

- changes to service interface (addr. size, option format, omitted fields)
- must protect self from SIP header corruption
- must protect self from old packets
- must perform Path MTU Discovery to send packets longer than 576

link-layer

- must provide own frag./reasm. if MTU < 576

IPv4 —> SIP Migration

- IPAE WG has developed a multi-step scheme for transitioning to SIP, using encapsulation of SIP in IPv4, initially
- IPv4 addresses can be used as low-order 32 bits of SIP addresses, while still unique
- SIP designed for straightforward translation between IPv4 and SIP packets:
  - — IPv4 compatibility bit in SIP address to indicate need for translation
  - — kept IPv4 fragmentation algorithm
  - — some IPv4 features lost in translation, e.g., Precedence, rarely-used options

Future Directions

- use Reserved field in SIP header to carry a "Flow ID", for real-time / special-QOS traffic
- add ICMP "Destination Has Moved" msg, for mobile hosts, rehomed domains
- add Trace Route option, for unicast and multicast path tracing
- new Host-to-Router protocol (ES-IS-like) for ARP, black-hole detection, router disc., auto-config and reconfig, and "beaconing"
- relaxation of strict subnet model, and need for each interface to have its own address
- authentication and integrity-assurance mechanisms, perhaps using SP3/4

4.3 IPAE

Presented by Bob Hinden/Sun

Bio:*Bob Hinden is the Manager of Internet Engineering at Sun Microsystems. He has been involved in the Internet community since 1980 and has been the IESG Routing Area Director since 1989. He is currently involved in work in internet routing and addressing, and the issues relating to internetworking using Asynchronous Transfer Mode (ATM).*

IP ADDRESS ENCAPSULATION

or

HOW TO TRANSITION TO A NEW IP

Robert Hinden

Sun Microsystems

November 16, 1992

CREDITS

- Dave Crocker / The Branch Office
  - Chair of IPAE working group
- Steve Deering / Xerox PARC
  - Inventor of new IP described here
- The IPAE and SIP Working Groups

INTRODUCTION

- Focus & Approach
- What is IPAE?
- Transition Plan
- Benefits and Differences
- Conclusions and Status

IPAE FOCUS

- Solve Internet Routing and Addressing Problems
  - Routing explosion (short term)
  - IP network address exhaustion (medium term)
- Change as Few Things as Possible
  - Protocol modules
  - Router and host devices
  - People and operations
- Build a Strong Base for the Future
- Best Cost / Benefit Ratio

HISTORY

- March 1992
  - *New Scheme for Internet Routing and Addressing*
- June 1992
  - *A Proposal for IP Address Encapsulation* (IPAE)
- September 1992
  - *Simple Internet Protocol*
- October 1992
  - Decision to Focus IPAE on a Transition Mechanism for a New Internet Protocol
  - Work with SIP w.g. to Refine SIP to Ease Transition Issues

APPROACH

- Protect the Installed Base
  - Solve only critical problems
- IP Works Well
  - User base is comfortable
  - Core operations extremely stable
- Transition to New IP
  - Provide Forward and backward compatibility
- Concentrate on Transition Issues

HOW IT WORKS

- New IP Protocol (SIP) with Same Semantics as Current IP
  - Makes translation between *IP <=> SIP* possible
- Addressing Support for Transition
  - 32-bit IP Addresses embedded in SIP address
  - Compatibility bit in address
- Use Current IP to Carry New IP during Transition Period

Sun Microsystems November 16, 1992

HOW TO GET THERE

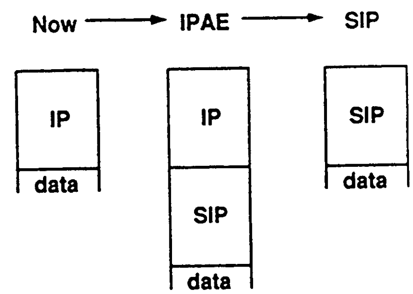

Sun Microsystems November 16, 1992

IPAE TRANSITION FEATURES

- First Deployment in Border Routers
- Uses Current IP Infrastructure to carry New IP during Transition
  - Doesn't require new infrastructure at start
  - Could run in this mode forever
- Very Flexible Transition Steps and Timing
  - Individual sites deploy at own schedule
- Benefits Accrued as Deployed
  - Don't have to wait until everyone else deploys
  - Hosts don't have to deploy first

Sun Microsystems November 16, 1992

MAPPING FEATURES

- Modes of Operation

IP <-> IPAE	IP <-> IP
IPAE <-> SIP	IPAE <-> IPAE
IP <-> SIP	SIP <-> SIP

- Fragmentation

  IP Fragment <-> SIP Fragment
- Control Fields

  TTL <-> HOP LIMIT

  LENGTH <-> LENGTH

  PROTOCOL ID <-> PAYLOAD
- Addresses

  IP Address <-> SIP Address

Sun Microsystems November 16, 1992

IPAE PACKET FORMAT

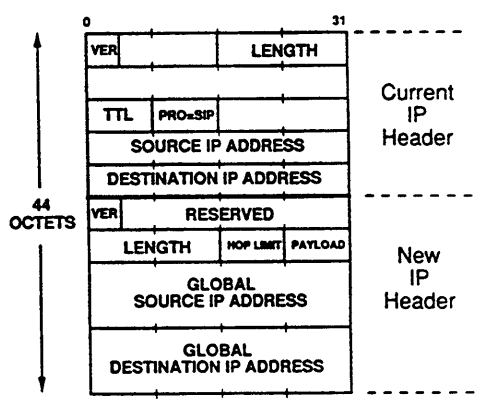

Sun Microsystems November 16, 1992

ROUTING MODEL

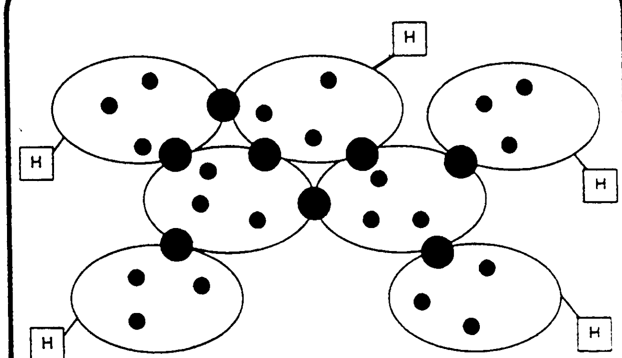

Sun Microsystems November 16, 1992

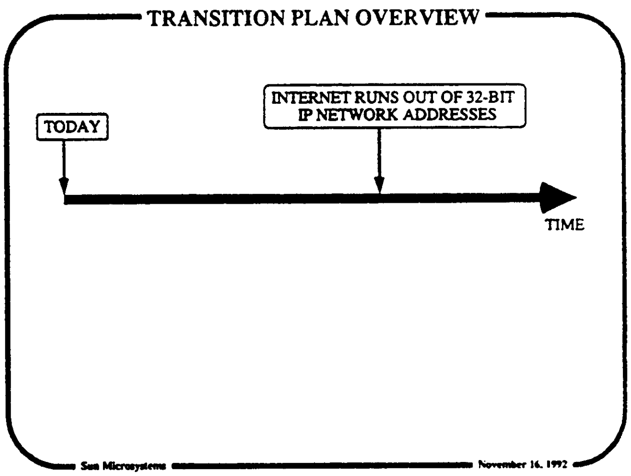
TRANSITION PLAN OVERVIEW
TODAY
INTERNET RUNS OUT OF 32-BIT IP NETWORK ADDRESSES
TIME
Sun Microsystems
November 16, 1992

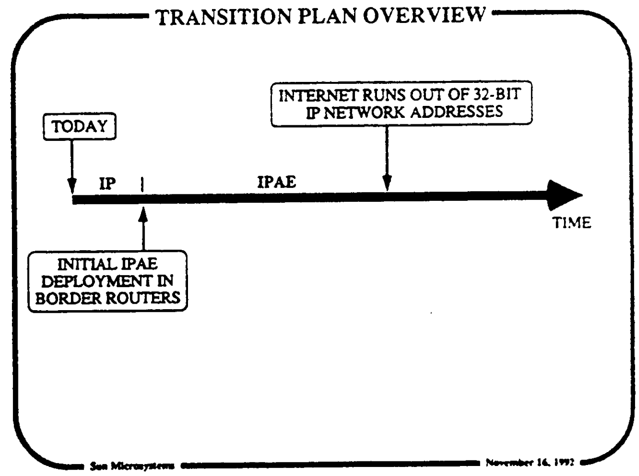
TRANSITION PLAN OVERVIEW
TODAY
INTERNET RUNS OUT OF 32-BIT IP NETWORK ADDRESSES
IP
IPAE
TIME
INITIAL IPAE DEPLOYMENT IN BORDER ROUTERS
Sun Microsystems
November 16, 1992

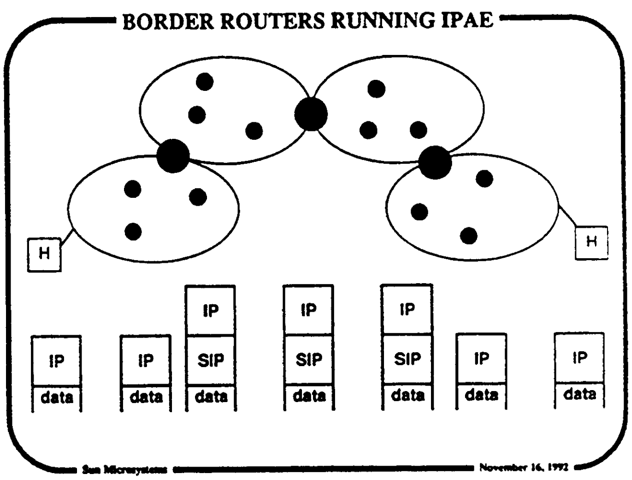
BORDER ROUTERS RUNNING IPAE
H
H
IP
data
IP
data
IP
SIP
data
IP
SIP
data
IP
SIP
data
IP
data
IP
data
Sun Microsystems
November 16, 1992

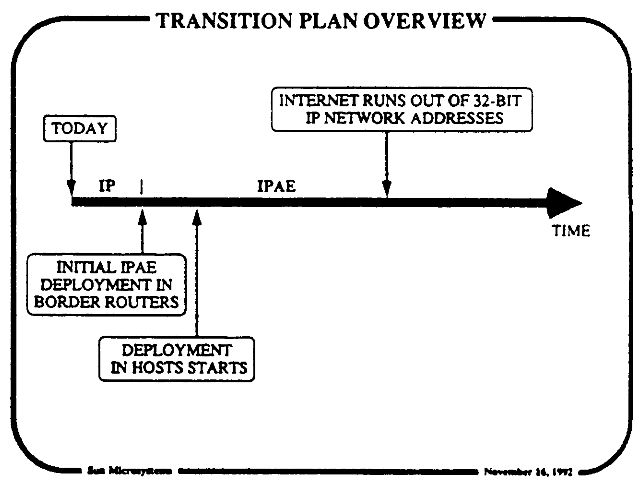
TRANSITION PLAN OVERVIEW
TODAY
INTERNET RUNS OUT OF 32-BIT IP NETWORK ADDRESSES
IP
IPAE
TIME
INITIAL IPAE DEPLOYMENT IN BORDER ROUTERS
DEPLOYMENT IN HOSTS STARTS
Sun Microsystems
November 16, 1992

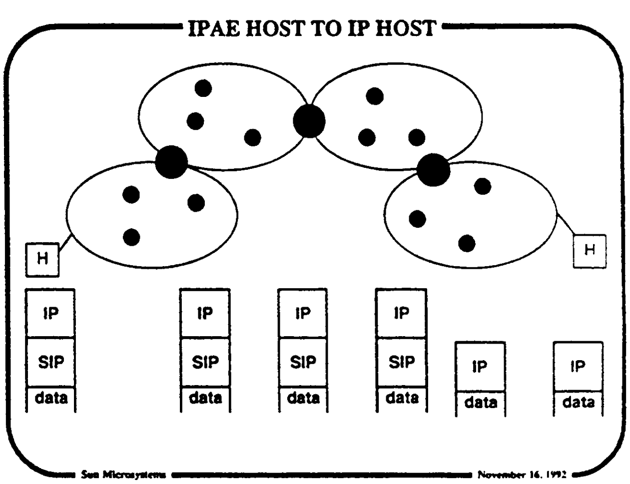
IPAE HOST TO IP HOST
H
H
IP
SIP
data
IP
SIP
data
IP
SIP
data
IP
SIP
data
IP
data
IP
data
Sun Microsystems
November 16, 1992

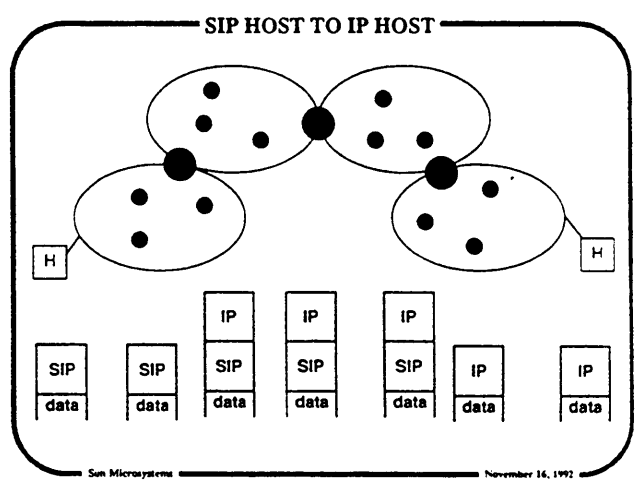
SIP HOST TO IP HOST
H
H
SIP
data
SIP
data
IP
SIP
data
IP
SIP
data
IP
SIP
data
IP
data
IP
data
Sun Microsystems
November 16, 1992

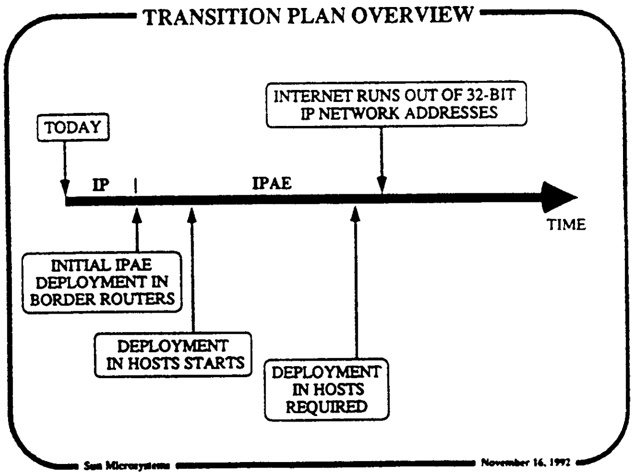

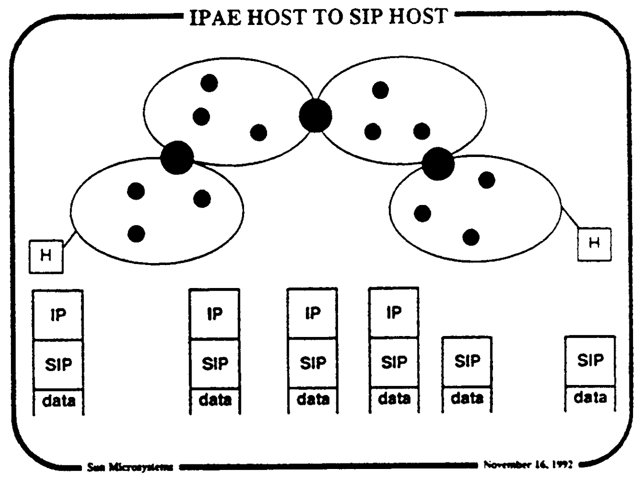

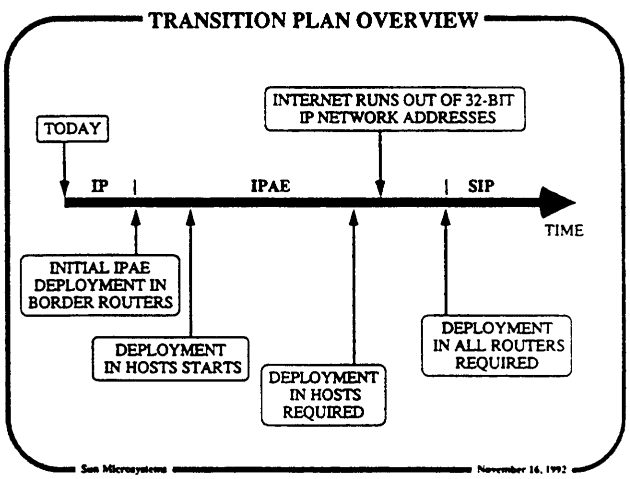

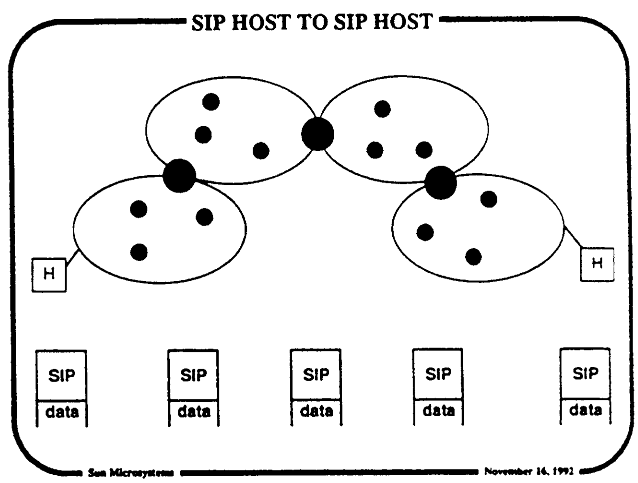

BENEFITS

- Routing Table Explosion Problem Solved in Short Term
- IP Address Exhaustion Problem Solved in Medium Term
- New IP Routing Infrastructure can Support both
  - New IP routing
  - IP routing

  Real Unified Routing

Sun Microsystems — November 16, 1992

BENEFITS (CONTINUED)

- Protocols
  - Current IP layer protocols retained as long as desired
  - Multicast retained
  - IP-over-<MEDIA> retained
  - Compatible with current IP mobility work

Sun Microsystems — November 16, 1992

BENEFITS (CONTINUED)

- Devices
  - Most interior and exterior routers not required to change until late in transition
  - Hosts do not change current 32-Bit IP Addresses
  - Hosts not required to implement IPAE until second transition step
  - IP hosts can continue to use IP inside site forever

Sun Microsystems November 16, 1992

BENEFITS (CONTINUED)

- People and Operations
  - Existing formats and terminology retained
  - Operational tools continue to function
  - Current investment in training, procedures, documentation retained

Sun Microsystems November 16, 1992

HOW THIS IS DIFFERENT FROM TUBA

- Deals with Routing Table Size and Routing Computation Problems When Initially Deployed
- Retains Current IP Layer infrastructure
- No Changes to Transport Protocols Required for IPAE
- Contains Protocol Identifier in Header
- No Host Changes Required at Initial Step

Sun Microsystems November 16, 1992

HOW THIS IS DIFFERENT FROM TUBA (CONTINUED)

- People and Operations
  - Flexible transition
  - Global communication between IP hosts and New IP hosts until IP runs out of networks addresses
  - Local communication between IP hosts and New IP Hosts always
  - Extends current IP technology base
  - No issues of protocol ownership

Sun Microsystems November 16, 1992

HOW THIS IS DIFFERENT FROM TUBA (CONTINUED)

- All of the Benefits of TUBA but at a Much Lower Cost
  + Provides a base for new features
- High Probability that IPAE can be Deployed in Time to Save Internet

Sun Microsystems November 16, 1992

CURRENT STATUS

- Four Implementations Underway
  - Sun, Silicon Graphics, Proteon, DEC
  - First datagrams exchanged!
- Internet Drafts Published
  - IPv7 Criteria Analysis for IP Address Encapsulation and the Simple Internet Protocol
  - Simple Internet Protocol
  - IP Address Encapsulation (IPAE): A Mechanism for Introducing an New IP
- To Join Working Group
  - ip-encaps-request@sunroof.eng.sun.com

Sun Microsystems November 16, 1992

Chapter 5

Technical Presentations

5.1 Report form the POISED Working Group

Presented by Steve Crocker/TIS

Bio:*Steve Crocker has been Vice President of Trusted Information Systems since 1986. He is responsible for research and development in network security, integration of cryptography and trusted systems, privacy enchanced mail, and program verification. His current professional activities include, Area Director for Security in the Internet Engineering Task Force and Treasurer, of the IEEE Technical Committee on Security and Privacy. Steve's past positions include: Director of Computer Science Laboratory at The Aerospace Corporation, El Segundo, Senior Research Associate, USC Information Sciences Institute and R&D Program Manager, DARPA/ISTO.*

Bio:*Carl Malamud is the author of "Exploring the Internet" (Prentice Hall, 1992) and other professional reference books.*

The POISED Working Group was chartered to examine the procedures of the IAB, IESG and IETF and propose changes in the way we select people to fill key positions and the way our various groups interact with each other. The relationship between the IAB/IESG/IETF and the new Internet Society is also of concern.

Specific proposals have been generated, including significant changes in the organizational structure. This talk will report on the main ideas and serve as a baseline for the Working Group meeting later in the day. Depending on how rapidly consensus is achieved, the changes contemplated by this Working Group may be adopted quickly, slowly or never.

Process Organization for Internet Standards (and Development)

POISED

Steve Crocker
POISED WG Chair

Process Organization for Internet Standards (POISED)

Outline

- Who am I? Why am I here?
- POISED activities at this IETF
- Online resources
- Draft proposals
- Sense of the community
- Internet Society Trustees Nominations (Craig Partridge)

Process Organization for Internet Standards (POISED)

Who am I? Why am I here?

- IPv7
- POISED WG
- Real problems
  - Scaling
  - Delay
  - Communication difficulty
  - Surprise
  - Focus
  - Accountability
  - Burn out

Process Organization for Internet Standards (POISED)

POISED Activities

- 15-20 MBytes Mail
  - poised-request@cnri.reston.va.us
- Internet Draft
- 3 Meetings
  - This Presentation
  - Working Group: Tues. 4-6
  - Plenary, Thursday Night
- Presentation to ISoc trustees, December 10

Process Organization for Internet Standards (POISED)

Draft Proposals

- Unwritten status quo
- ITTF draft (Crocker/Malamud)
- Davin constitution
- Plan fragments
  - Form new society
  - Humming

Process Organization for Internet Standards (POISED)

Internet Technical Task Force (ITTF)

- Working groups and design teams
- WGs remain main focus
  - open forums
- multiple types
  - engineering
  - + research (advanced development?)
  - + architectural
- Design teams
  - Self-selecting group

Process Organization for Internet Standards (POISED)

ITTF Leadership

- Technical Board
  - Area chairs
  - ITTF Chair
  + Architect
- Editor
- Process Board
- Internet Society

Process Organization for Internet Standards (POISED)

Accountability and Selection

+ ITTF Technical Board members
  - Nomination WG
  - Each design team proposes specific selection
  - WG consensus? ITTF consensus?
  - Forward one or more names
  - Process Board advice and consent
+ Process Board members
  - Same as above plus
  - ISoc Trustees advice and consent
+ Recall

Process Organization for Internet Standards (POISED)

Process and Procedure

- Accountability and selection
- Objective Criteria
  - Technical -- working code
  - Procedural
- Open Meetings
- Hearings

Process Organization for Internet Standards (POISED)

Transition Plan

- Is there IETF consensus?
- Internet Society trustees
- RFC publication of charter
- Transition
  - Process Board and ITTF chair selection
  - Interim operation of existing IAB/IESG
  - Selection of Technical Board
  - Flag day (undetectable, seamless)

Process Organization for Internet Standards (POISED)

Other proposals

- Status quo (almost)
  - Leave IAB, IESG, IETF, IRTF intact
  - need some sort of selection process
- Davin's constitution
- Humming

5.2 Connection-Oriented Internet Protocols

Presented by Chuck Cranor/Ph.D. Candidate - Washington University

Bio: *Chuck Cranor received his undergraduate degree in electrical engineering from the University of Delaware in 1989. He received his masters in computer science in 1992 from Washington University in St. Louis, Missouri. He is currently starting work on his doctoral degree under the advisorship of Dr. Gurudatta M. Parulkar.*

Recently a number of research groups have proposed connection-oriented access protocols that can provide a variable grade of service with performance guarantees on top of diverse networks (e.g., MCHIP, ST-II, FLOW). These connection-oriented internet protocols (COIPs), while similar in many aspects, have different performance trade-offs.

A COIP-Kernel which can be used as a toolkit to implement the proposed COIPs has been created. COIP-K forms the core of a COIPprotocol and includes the minimum functionality necessary for a wide range of multicast connection-oriented protocols. It includes appropriate provisions to interface with protocol-specific functional modules to form an instance of a COIP protocol. COIP-K features module interchange and incremental software support. This talk briefly presents COIP-K's design and implementation.

An Implementation Model For

Connection-Oriented Internet Protocols

(COIP-K)

Charles D. Cranor
(314)-935-4203
chuck@maria.wustl.edu

Guru Parulkar
(314)-935-4621
guru@flora.wustl.edu

Washington University, St Louis MO.

This work is supported in part by Bellcore, BNR, DEC, Italtel SIT, NEC, NTT, Synoptics, and NSF.

Existing Internet Model

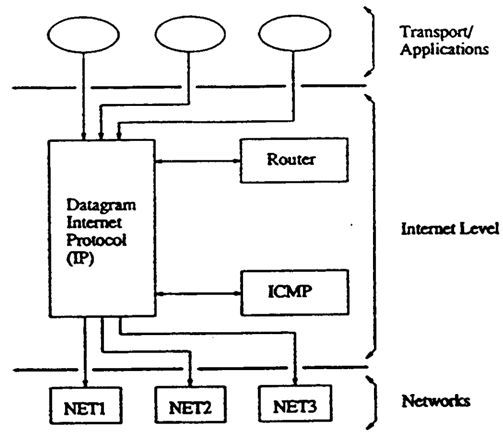

Does not allow variable grades of service with performance guarantees (e.g. rfc1363 FLOWs)

2

Congram-Oriented Internetworking

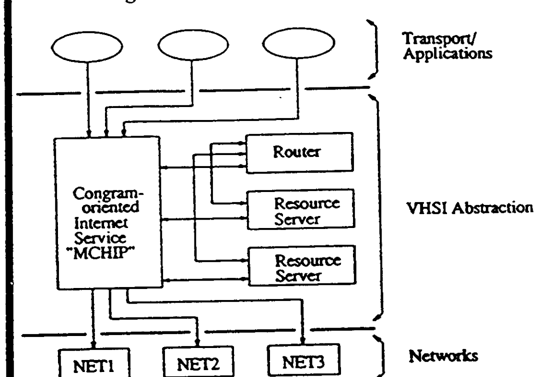

VHSI goal: variable grade of service with performance guarantees across an internet of heterogeneous networks

Provide homogeneous networking abstraction to higher level protocols on top of a set of heterogeneous networks

Congram service: provides datagram and connection oriented

Resource server for datagram networks

Proposed COIPs: MCHIP, ST, FLOW

3

Congram Phases

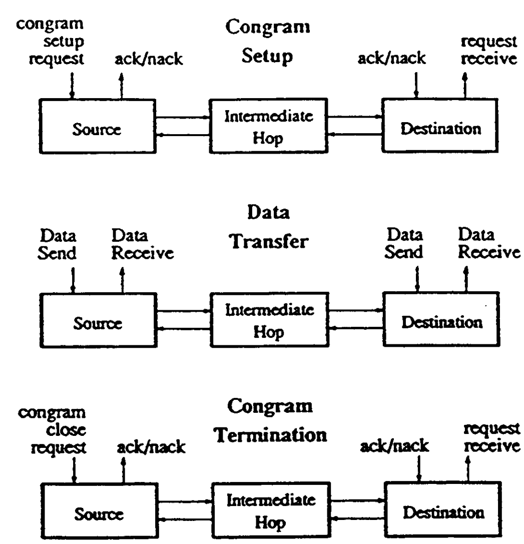

4

COIP Protocol Differences

Connection management:

ST: one-to-many
FLOW: point-to-point
MCHIP: many-to-many

ST allows addition and deletion from MTP connection

Resource management:

FLOW: average bandwidth, average interval, virtual clock

ST: peak bandwidth

MCHIP: peak bandwidth, average bandwidth, peak to average ratio, delay, loss...

Packet formats

Motivation

Observations:

- Proposed COIPs have similarities and differences
- It's important to pursue, compare, and contrast these protocols
- Independent implementation of protocols considered unwise
  - duplicate work
  - complexity of Unix kernel

COIP-K proposed:

- Forms the core of a COIP protocol
- When combined with a set of functional modules, it creates an instance of a COIP protocol

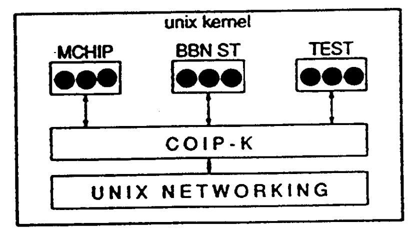

Implementation Requirements

1. COIP-K must be implemented in the Unix kernel
2. COIP-K should allow implementation of various COIP protocols by module interchange
3. COIP-K should refrain from modifying the user-level socket interface
4. COIP-K should have efficient per-packet processing
5. COIP-K should support multipoint connections

Our COIP-K Implementation

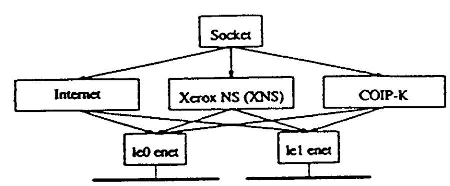

COIP-K is its own communication domain.

Each protocol under the COIP-K domain will have it's own protosw structure.

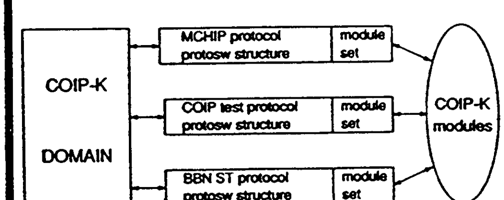

Socket Layer COIP-K Functions

COIP-K Client

```
s = socket(PF_COIP, SOCK_STREAM, 0)
e = setsockopt(s, 1, CODE, &p, sizeof(p)))
e = connect(s, cin, cinlen)
```

COIP-K Server

```
s = socket(PF_COIP, SOCK_STREAM, 0)
e = setsockopt(s, 1, CODE, &p, sizeof(p)))
e = bind(s, &cin, sizeof(cin))
e = listen(s, 5)
s_new = accept(s, &cn, sizeof(cn))
```

After Establishment

```
err = read(s, buf, buflen)
err = write(s, buf, buflen)
close(s)
```

9

Protocol Layer Data Structures (PCB)

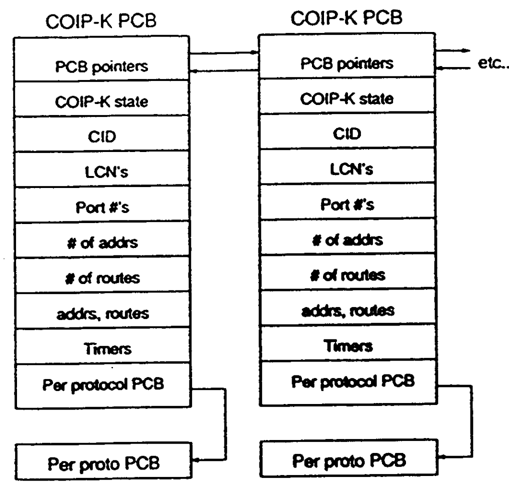

PCBs stored in linked list of mbufs

MTP connections require extra mbufs for addrs/routes

10

Addresses in the COIP-K PCB

Since COIP-K supports multipoint connections, it must be able to store a variable number of addresses and routes in the PCB.

Exception: point to point connections store all addressing and routing info in the cinpcb struct.

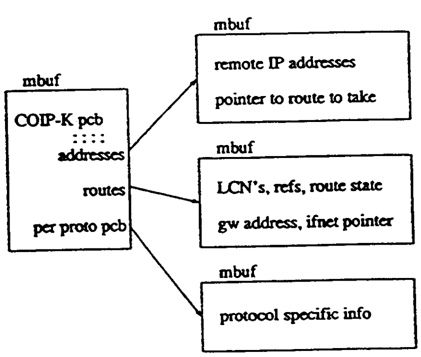

11

Protocol Layer Functions

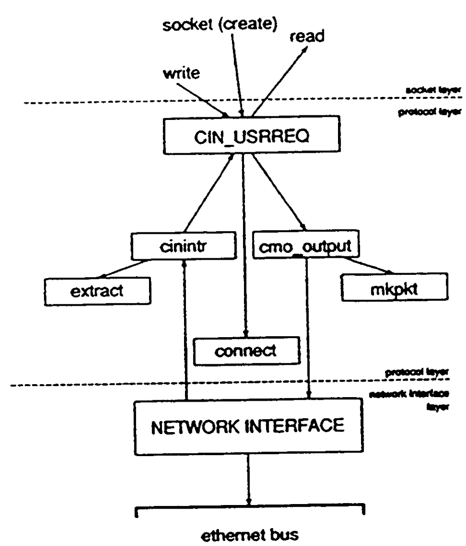

12

COIP-K Modules

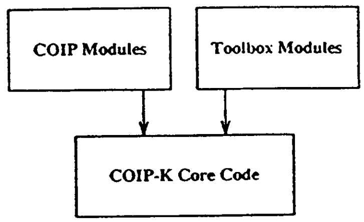

There are two types of COIP-K modules: required and optional

Required:

Extract	Packet type
Make packet	PCB lookup

Optional:

Init	Fast / Slow timers
Attach	PCB set up
Detach	Output
Data input	Localized
Control input	Reject
Disconnect	Performance packet input
Connect	Set performance requirement

13

COIP-K Feasibility and Viability

The objectives of COIP-K have been successfully achieved as outlined below:

- COIP-K Test Protocol (CTP)
  - Implemented using COIP-K
  - A subset of the MCHIP protocol
  - Used for testing of COIP-K
- COIP-K module interchange
  - Critical in realizing and comparing different COIPs
  - Example: CTP and CTP2
- COIP-K performance results
  - Characterize per-packet processing
  - Quantify the cost of the COIP-K concept
- COIP-K tests and demonstrations
  - Test and verify several capabilities of COIP-K
  - Show that applications can be easily ported
  - COIP-K can be used to create useful apps.

14

COIP-K Module Interchange

Recall:

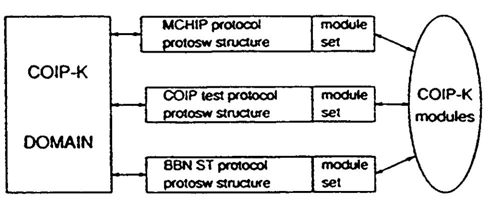

modules plug in to get a COIP protocol

a module set is formed from protocol specific and toolbox modules

When a module set is put in COIP-K the result is a COIP protocol

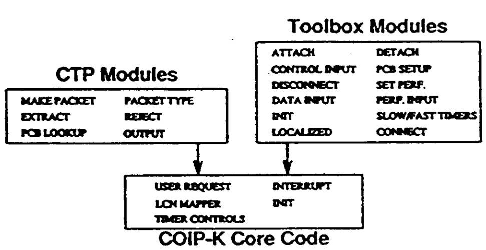

15

COIP-K Module Sets

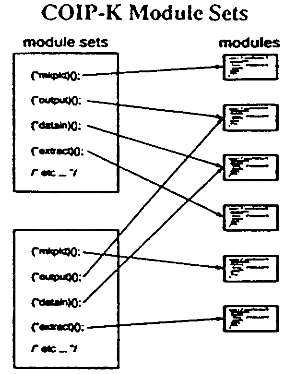

Notes:

1. All modules are part of a "module pool"
2. Modules can be shared between module sets
3. To create a new instantiation: copy COIP's module set
4. To override a default module: add new module to pool and change appropriate pointer
5. Program chooses module set with socket() sys call

16

Module Interchange Example: CTP2

To demonstrate the usefulness of COIP-K module interchange, a new COIP protocol based on CTP called CTP2 has been created.

Two ways to do multipoint connections:

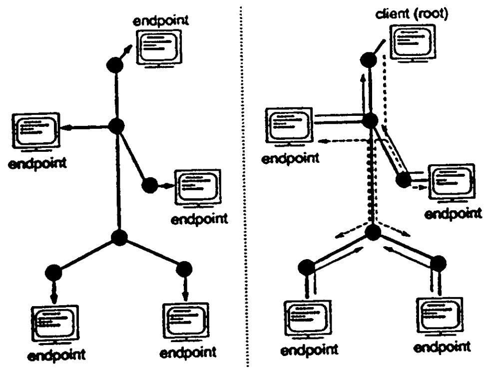

CTP: many-to-many

CTP2: one-to-many

Changes for CTP2: new data input module
modify toolbox output module

CTP2 implementation time: ~1 hour

17

COIP-K Performance

The cost of COIP-K:

What is the cost of using COIP-K as compared to direct protocol implementation...?

Cost: the added overhead of a module call.

CALL	INSTRUCTIONS
`fcall()`	1
`(*ptr1->fptr)()`	4
`(*ptr1->ptr2->fptr)()`	5

Cost: data access via accessor functions (1 function call)

Throughput (on same local sun):

UDP = 7.57 Mbps

COIP-k/CTP = 7.48 Mbps

Processing time/delay:

UDP = 0.146 msec (client), 0.073 msec (server)

CTP = 0.115 msec (client), 0.048 msec (server)

Theoretical max throughput: 71Mbps

18

Demos

Simple demo: send a packet and close

File transfer

File transfer through "gateway"

telnet/telnetd - port of 4.3bsd telnet/telnetd programs

Simple MTP demo

MTP chat program

MTP script program

19

Conclusions

We have realized the vision of COIP-K

In order to realize this vision, we set and met our 5 implementation requirements...

1. COIP-K is implemented in the Unix kernel
2. COIP-K allows implementation of various COIP protocols easily with module interchange:
   - partitioning of COIP protocols
   - ease of interchange of modules
   - variable levels of kernel support

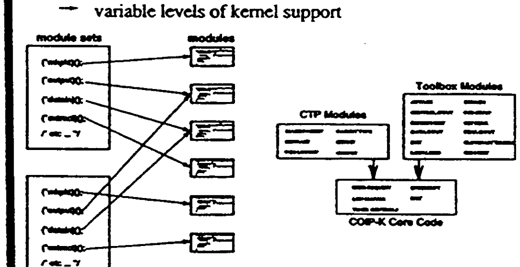

20

Conclusions

We have met all 5 of our implementation requirements:

1. Unix kernel
2. Module interchange
3. COIP-K has retained the user-level socket interface, thus allowing easy porting of applications such as telnet/telnetd.
4. COIP-K has efficient per-packet processing (comparable to UDP and much better than TCP)
5. COIP-K supports multipoint connections, which are useful for multiparticipant collaborations

COIP-k runs on SunOS 4.0.3c/4.1.1/4.1.2 sun4/sun4c arch

ftp: dworkin.wustl.edu [128.252.169.2], dir: /dist/coip-k

21

5.3 Source Demand Routing Protocol Specification (Version 1)

Presented by Deborah Estrin/USC, Tony Li/cisco and Yakov Rekhter/IBM

Bio:*Deborah Estrin is an Associate Professor of Computer Science at the University of Southern California where she teaches and does research in the area of Computer Networks. Deborah Chairs the Autonomous Networks Research Group and is a member of the IRTF. She was one of the designers of IDPR and now works on: Unified routing, adaptive routing, reservation setup for integrated services networks, and large scale simulation/emulation.*

Bio:*Tony Li received his B.S. from Harvey Mudd College and Ph.D. from the University of Southern California. He currently slings bits for cisco Systems, Inc., specializing in IP exterior routing protocols and fixing nasty bugs.*

Yakov Rekhter received his M.S. in Physics from St. Petersburg (Leningrad) University, Russia (USSR), his M.S. in Computer Science from New York University, and his Ph.D. in Computer Science from Polytechnic University. Yakov is a manager of High Performance Networking group at the T.J. Watson Research Center, IBM Corporation. He is a Chair of the BGP Working Group of the IETF and an active participant in the ANSI X3S3.3 committee. He is one of the principal designers and implementors of the NSFNET Backbone routing architecture and protocols.

The purpose of SDRP is to support source-initiated selection of inter-domain routes to complement the intermediate-node route selection provided by BFP ([1], [2], [3]) or IDRP ([4]). This document refers to such source-initiated routes as "SDRP routes".

The protocol makes minimal assumption about the distribution and acquisition of routing information needed to construct the SDRP routes. These minimal assumptions are believed to be sufficient for the existing Internet. Future versions of the protocol will extend capabilities in this area and others in a largely backward-compatible manner.

This version of the protocol sends all packets with the complete SDRP route in the SDRP header. Future versions will address route setup and other enhancements and optimizations.

Source Domain Routing Protocol (SDRP): A Component of Unified Approach to Inter-Domain Routing

Deborah Estrin (USC)
Tony Li (cisco)
Yakov Rekhter (IBM)

- The problem
- Overview of Unified Approach
- SDRP
  - Design constraints
  - Some details
  - Status

The Problem

To provide efficient inter-domain routing in future global internetworks.

Protocols must scale to networks that are increasingly large and complex.

- Large: requires good complexity characteristics (logarithmic is good; N^2 is unacceptable)
  - forwarding table space.
  - overhead of routing information distribution, storage and processing.
  - convergence of nodes' routing decisions.
- Complex: functional requirements
  - type of service and policy-sensitive routing.
  - facilitate autonomy of routing decisions.
  - require little global coordination.

The Pain of Policy

There is no ONE best route to a destination.

- Multiple types of service (TOS)
- Multiple carriers, charging schemes, user classifications, etc.
- Restrictions imposed by transit domains govern who/what may traverse resources.
- Source domain policies guide which allowable paths are selected/preferred.

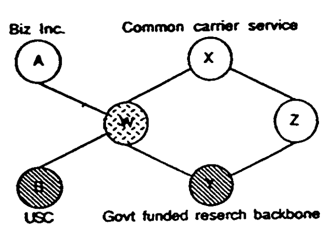

Scale and Heterogeneity: Contradictory Goals

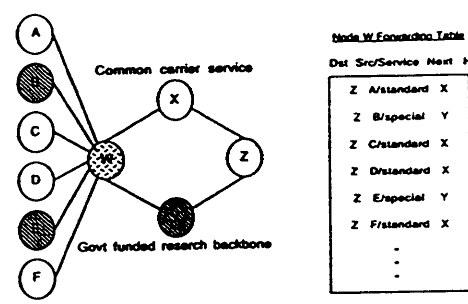

Previous protocols were engineered to perform well for particular internetwork models; emphasized scaling OR heterogeneous control.

- Scaling mandates information reduction.
- Heterogeneous control mandates need for detailed information.

Real world is both large and heterogeneous.

Conclusion: Contradictory goals require complementary mechanisms.

Unified Approach

Node Routing and Source-Demand Routing.

- Avoid over-engineering individual protocols
- Design as integrated whole; avoid doubling overhead/mechanism.

Complementary components deal efficiently with both very large and complex internetwork model.

- Generic routes (widely used) supported by Node routing that emphasizes **aggregation**.
- Special routes (sparsely used) supported by Source Demand Routing that emphasizes **flexibility** and **extensibility**, minimizes global system impact of special requirements.
- Adapt over time to changing traffic/demand: if a routing requirement becomes widely used it becomes more efficient to anticipate the demand and produce a generic route.

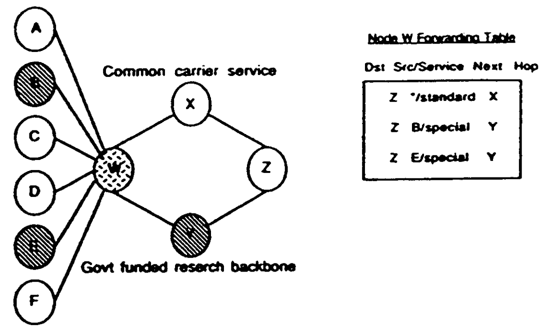

Node Routing Overview

Path Vector routing mechanism as in BGP, IDRP

- Information exchanged includes a record of domain-level path that the information has traversed.
- Path information used to avoid routing loops.
- Reduces dependence on globally-known information and globally-coordinated decision criteria.
    - ° Heterogeneous route selection criteria.
    - ° Partial and selectively distributed routing information.
    - ° Flexible, selective aggregation: TOS/policy, topology, reachability info may be aggregated differently by different domains.

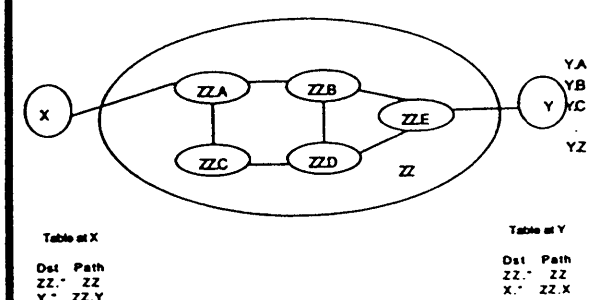

Source Demand Routing Protocol SDRP

Traditional LS protocols distribute state information globally, within a hierarchical level (OSPF, IS-IS, IDPR).

We exploit sparse utilization and NR as information infrastructure.

"Lazy Link-State" to reduce the overhead required to distribute, and compute over, detailed global routing information.

Decouple route construction and route validation.

- Slowly-changing topology and policy map.
- Exploit NR information.
- Dynamic (status) information acquired on-demand, in response to active probes or failure messages.
- Additional mechanisms TBD for acquiring selective portions of map on demand (e.g., subscriptions, search, etc.)

Design Constraints

Isolate the cost and mechanism to those that want it; a form of autonomy...

Minimize overhead, particularly for those domains that are not requesting or imposing the special requirements.

- No global flooding
- Ability to specify SDR for PART of the route, use NR for the rest.
- Support "loose" domain-level source route.

"Lite"

- No mandatory setup; efficient DG support.
- No HBH-reliable setup.
- Incremental deployment: may start with very simple route construction support and get more sophisticated as demand grows.

SDRP Packet Forwarding and Control Msgs

SDRP route is sequence of domains; may be strict or loose.

Encapsulation used to send packets over SDRP routes.

- Packet from originating source host is encapsulated in an SDRP packet; SDRP header carries source route and other control info.
- SDRP packet carried across domains as data portion of IP packet with SDRP protocol number.

Delivery header | SDRP header | Payload

SDRP data-packet header carries: Source route, Pointer to next domain in source route, Probe and Strict/Loose route flags, BR hop count, Source route and payload protocol types, route identifier.

SDRP control-packet header also carries: Notification codes, Target BR.

Originating SDRP Packets

BR receives IP datagram originated by host within domain.

BR uses information in packet and local criteria to determine whether the datagram should be forwarded along a particular SDRP route (criteria are local matter; see future usage document.)

BR constructs SDRP packet:

Original packet into payload.

Selected route as source route.

Next hop pointer set to 0,

TTL copied into hop count.

If loose source route: add NLRI information to avoid loops.

SDRP packet encapsulated in IP delivery header:

DF, TOS, Security fields copied from payload.

Source address in IP delivery header is of originating BR.

Forwarding SDRP Packets

If Domain pointed to by next hop pointer is domain of the current BR, increment pointer and route packet to this next domain.

D-FIB contains routes to domains, based on BGP/IDRP or configuration (Domain, next-hop-BR)

If Strict source route, next domain pointed to must be adjacent.

Look up Next domain in D-FIB to find next SDRP hop (use NLRI information also if loose source route).

Place IP address of next SDRP hop in destination field of IP delivery header.

When end of source route reached, headers stripped off and payload packet delivered to destination using normal FIB information.

Error messages returned if no D-FIB or FIB entries found.

-Control messages

Sent by last BR in SDRP route in response to probe bit set by originator; used to determine route feasibility.

Sent by any transit SDRP BR to report errors back to originating BR.

Control information also inferred from ICMP messsage carrying Source Route Identifier of SDRP route as last 4 bytes of ICMP message. Used to identify problem SDRP route.

Control messaged used for MTU discovery: set DF and Probe bit, or infer from ICMP dst unreachable and local information.

Constructing SDRP routes

Initially use BGP/IDRP routes plus dynamic information from SDRP control messages.

Construct prioritized list of SDRP routes to support dynamic rerouting upon failure.

TBD (To Be Designed)

- Develop route construction and selection algorithms/heuristics/rules.
- Develop protocol for obtaining (relatively-static) maps that will be more complete.
- Develop more sophisticated dynamic (on demand) update techniques; but no global flooding.

Summary

Contradictory goals require complementary mechanisms.

HBH/PV selected for widely-used generic routes because of flexible and efficient aggregation properties.

Source routes constructed and applied on-demand to support special routes.

SDRP currently being developed to support special routes.

Join SDRP BOF this afternoon.

Email add request to
sdrp-request@caldera.usc.edu

→ Existing Doc describes V1. of pkt
Forwarding + Control Msg Protocol

→ To appear : usage doc
: Route info protocol

5.4 Export Controls on Cryptographic Software

Presented by Gerard Rainville/NSA

EXPORT CONTROL
OF
CRYPTOGRAPHIC PRODUCTS

Twenty Fifth IETF

Thursday, November 19, 1992

Washington, D.C.

Gerard A. Rainville, Jr.

National Security Agency

EXPORT REGIMES

DEPARTMENT OF STATE	DEPARTMENT OF COMMERCE
ITAR	EAR
USML	CCL (MCTL)
DOD Role Mandatory	DOD Advice Optional

COCOM

GENTLEMAN'S AGREEMENT

IL, IML

OTHER

CRYPTOGRAPHY

ANCIENT RULE: IF IT'S CRYPTO, IT'S NOT EXPORTED!

OLD RULE: IF IT'S CRYPTO, IT'S UNDER ITAR CONTROL.

CURRENT RULE: IF IT'S CRYPTO, IT'S BORN UNDER ITAR CONTROL, BUT MAY MIGRATE TO EAR CONTROL.

FUTURE RULE: ???????????????????????????????

INTEGRITY and ACCESS CONTROL	U.S. and Canada	U.S. subs / F.I.	All other
DES	No restriction.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.
Other symmetric key systems	No restriction.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.
RSA and other asymmetric key systems	No restriction.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.	General Commerce Department license except for Eastern Bloc. Individual licenses for Eastern Bloc usually granted.

KEY MANAGEMENT	U.S. and Canada	U.S. subs / F.I.	All other
DES	No restriction. A warning label that the product requires a license for export is recommended..	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.
Other symmetric key systems	No restriction. A warning label that the product requires a license for export is recommended.	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.
RSA and other asymmetric key systems	No restriction. A label warning that the product requires a license for export is recommended.	A State Department license is required and is generally granted provided the modulus does not exceed 512bits.	A State Department license is required and is generally granted provided the modulus does not exceed 512bits.

ENCRYPTION	U.S. and Canada	U.S. subs / F.I.	All other
DES	No restriction. A warning label that the product requires a license for export is recommended..	State Department export license is required and will generally be granted.	State Department export license is required and generally will NOT be granted.
Other symmetric key systems	No restriction. A warning label that the product requires a license for export is recommended.	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.	Examined on a case-by-case basis. For some products, a Commerce Department license may be available.
RSA and other asymmetric key systems	No restriction. A label warning that the product requires a license for export is recommended.	State Department export license is required and generally will NOT be granted.	State Department export license is required and generally will NOT be granted.

POC PHONE NUMBERS

STATE:

GENERAL INFORMATION: (703) 875-6644

COMMODITY JURISDICTION: (703) 875-5655

LICENSE STATUS: (703) 875-6652

COMMERCE:

JOE YOUNG: (202) 377-0708

EXPORT TECHNICAL ADVISERS: (301) 688-7834

EXPORT PROBLEMS UNIQUE TO SOFTWARE

VAPORWARE

- What is it that is out there?
- Who has what?

CULTURE

- Shareware.
- Unclear legal landscape.

MARKET

- Gizmo-mania.
- Rapid product cycle.
- Mass market.

RESPONSE TO SOFTWARE CHALLENGE

Separate encryption from other uses of cryptography.

Separate privacy from confidentiality.

Mass market easy out.

SPA negotiations.

5.5 IPX Protocol Suite

Presented by Paul Turner/Novell

Bio:*Paul Turner has worked in Novell's Systems Engineering Division for four years. During that time, he has written several technical papers covering the NetWare IPX protocol architecture and regularly appears at industry trade shows and conferences to discuss the subject. Paul received a degree in Electrical Engineering from the University of South Florida in 1987.*

Novell's Internetwork Packet Exchange (IPX) protocol suite enjoys the largest installed base of any internetwork protocol, however, the details of its operation remain a mystery for many. This presentation provides a detailed explanation of IPX protocol stack and its operation. The information presented includes protocol frame structures and algorithms. The protocols discussed include: IPX, Sequenced Packet Exchange (SPX), Routing Information Protocol (RIP), Service Advertising Protocol (SAP), NetWare Core Protocol (NCP), and NetWare Link Services Protocol (NLSP).

NOVELL

The IPX Protocol Suite

Paul Turner

NOVELL

History

- IPX suite derived from XNS spec in early 1980's
- Protocols added or enhanced to better support file, print, gateway services
- Installed base in 1992: over 8 million nodes (Forrester)

3 © Novell, Inc
SED-SIO-212I

IPX Protocols Stacked Against OSI Model

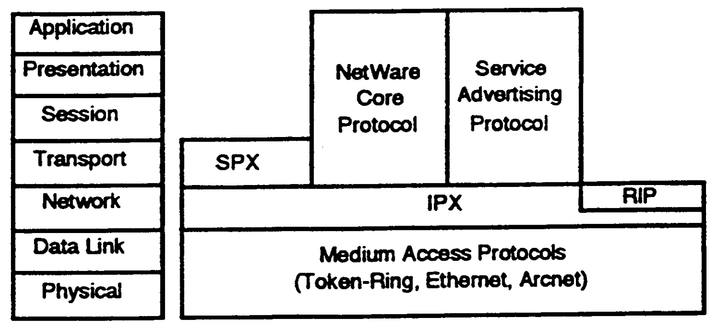

4 ©Novell, Inc
SED-SIO-212I

Internetwork Packet Exchange (IPX)

- TCP/IP counterparts: IP and UDP
- Same structure as XNS IDP
- Datagram/connectionless protocol
- Supports multiple media and frame types
- Defines internetwork addressing scheme

5 ©Novell, Inc
SED-SIO-212I

Structure of an IPX Packet

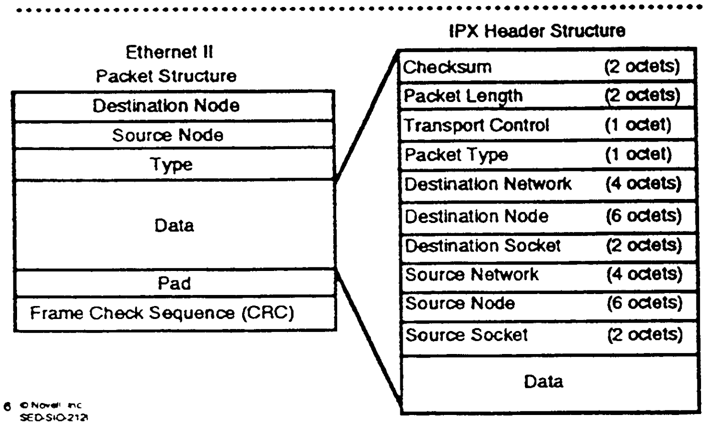

6 ©Novell, Inc
SED-SIO-212I

IPX Addressing

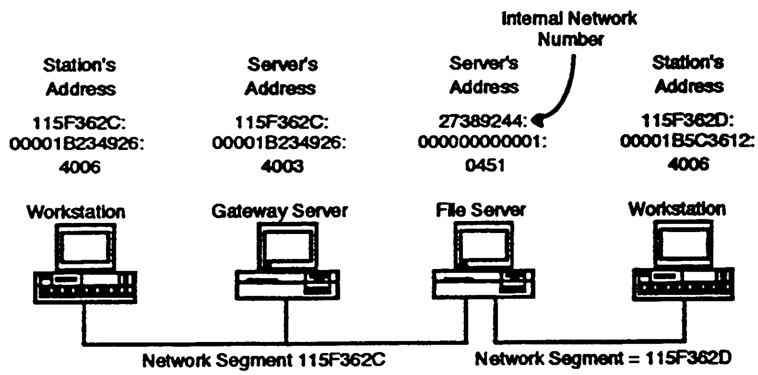

7 © Novell, Inc.
SED-

IPX Routing Methodology

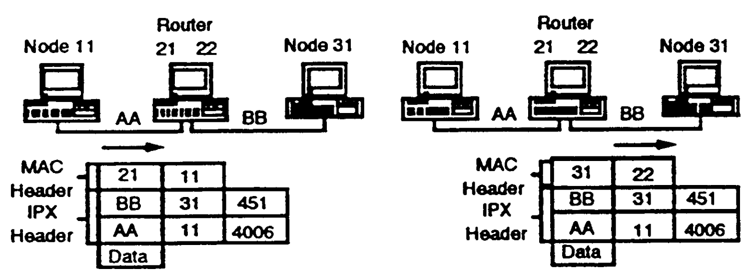

8 © Novell, Inc.
SED-SIO-2121

Routing Information Protocol (RIP)

- TCP/IP counterpart: RIP
- Used for transfer of internetwork routing information
- Distance-vector protocol
- Differs slightly from XNS RIP: Time delay field added
- IPX routers use RIP to exchange route information
- IPX clients use RIP to find fastest route to destination networks

9 © Novell, Inc.
SED-SIO-2123

RIP Packet Structure

RIP Packet Structure

Operation	(2 octets)
Network Number (1)	(4 octets)
Number of Hops (1)	(2 octets)
Number of Ticks (1)	(2 octets)
:	
Network Number (n)	(4 octets)
Number of Hops (n)	(2 octets)
Number of Ticks (n)	(2 octets)

IPX Header Structure

Checksum
Packet Length
Transport Control
Packet Type
Destination Network
Destination Node
Destination Socket
Source Network
Source Node
Source Socket
Data

10

Broadcasts When Router is Brought Up

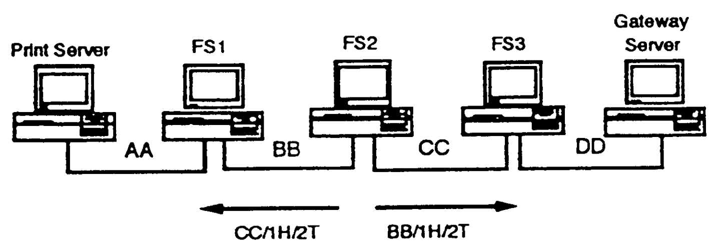

H = Hop = A pass through a router
T = Tick = 1/18 second

11 © Novell, Inc
SED-SIO-2121

Request for Information by Router

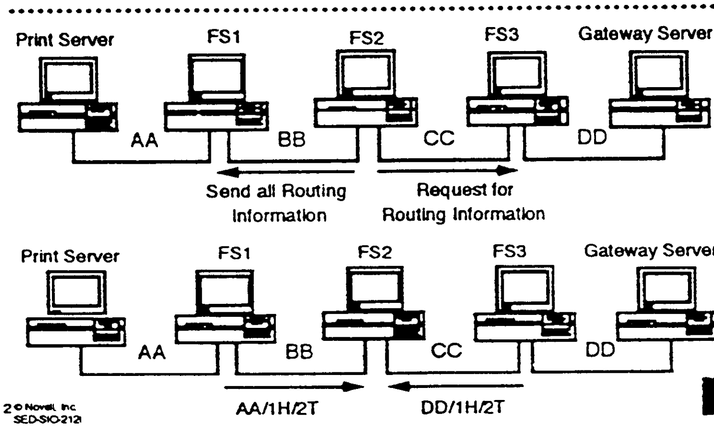

12 © Novell, Inc
SED-SIO-2121

Routing Information Broadcasts

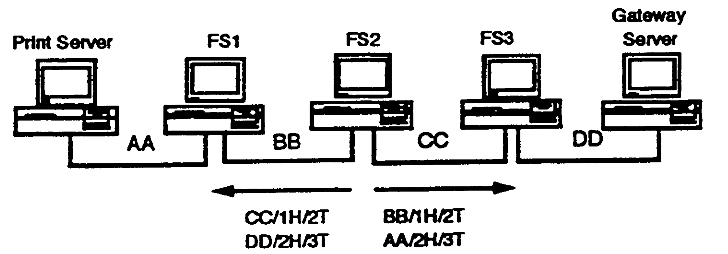

H = Hop = a pass through a router
T = Tick = 1/18 second

Service Advertising Protocol (SAP)

- TCP/IP counterpart: DNS Not Really!
- Servers use SAP to advertise their services, name, and address
- Routers use SAP to exchange server information
- Workstations use SAP to locate servers

SAP Packet Structure

SAP Packet Structure

Operation	(2 octets)
Service Type(1)	(2 octets)
Server Name (1)	(48 octets)
Network Address (1)	(4 octets)
Node Address (1)	(6 octets)
Socket Address (1)	(2 octets)
Hops to Server (1)	(2 octets)
:	
Service Type (n)	(2 octets)
Server Name (n)	(48 octets)
Network Address (n)	(4 octets)
Node Address (n)	(6 octets)
Socket Address (n)	(2 octets)
Hops to Server (n)	(2 octets)

IPX Header Structure

Checksum
Packet Length
Transport Control
Packet Type
Destination Network
Destination Node
Destination Socket
Source Network
Source Node
Source Socket
Data

Server Information Broadcasts

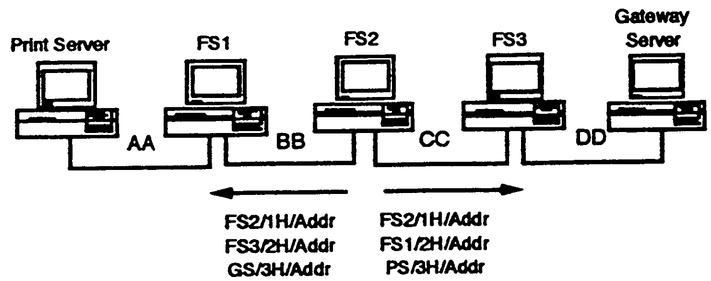

H = Hop
Addr = Full Internetwork Address

Frequency of Broadcasts

- Every 60 seconds
  - Token-Ring
  - Ethernet
  - Arcnet
- Every 60 seconds or only when information changes
  - T1, Fractional T1, 56Kbps
  - X.25
  - Async

Connecting to a File Server

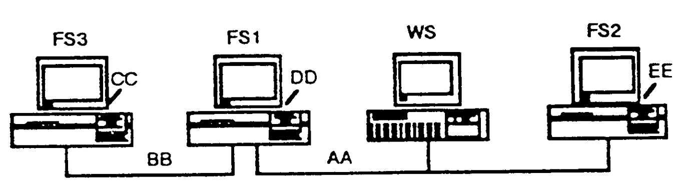

Initial Connection Sequence

Call	Source	Dest.	Protocol
1. Get Nearest Server	Client	Broadcast	SAP
2. Give Nearest Server	Router	Client	SAP
3. Get Fastest Route	Client	Broadcast	RIP
4. Give Fastest Route	Router	Client	RIP

19 ©Novell, Inc
SED-SIO-2121

Initial Connection Sequence

(continued)

Call	Source	Dest.	Protocol
5. Create Connection	Client	FS1	NCP
6. Requested Processed; Connection # Assigned	FS1	Client	NCP
7. Propose Buffer Size	Client	FS1	NCP
8. Return Common Buffer Size	FS1	Client	NCP

20

Get Nearest File Server Request

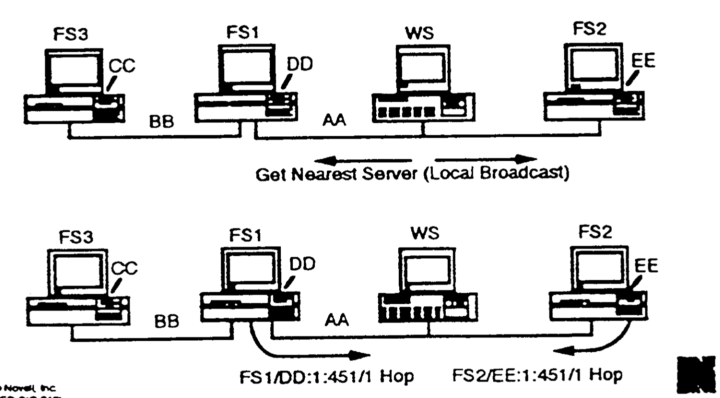

21 ©Novell, Inc
SED-SIO-2121

Get Routing Information Request

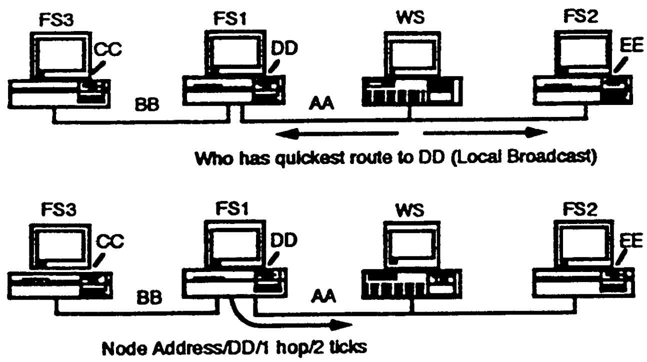

22 © Novell, Inc
SED-SIO-2121

Create Connection

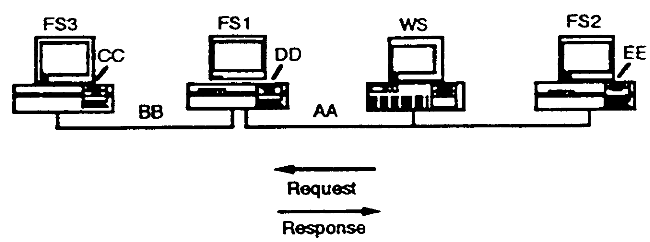

23 © Novell, Inc
SED-SIO-2121

Query of Bindery

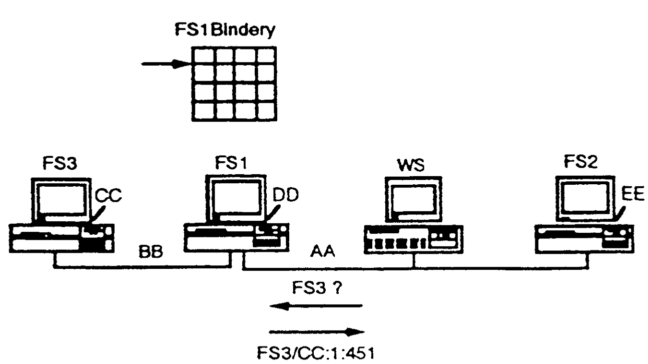

24 © Novell, Inc
SED-SIO-2121

Get Routing Information Request

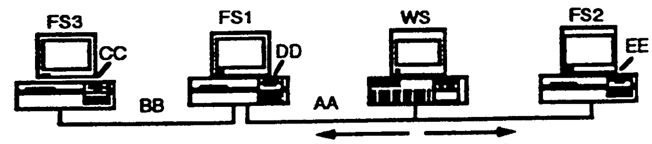

Who has quickest route to CC (Local Broadcast)

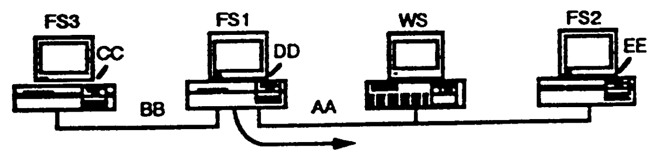

Node Address/CC/2 hops/3 ticks

25 ©Novell, Inc
SED-SIO-2121

Sequenced Packet Exchange (SPX)

- TCP/IP Counterpart: TCP
- Same packet structure as XNS SPP
- Transport Layer protocol

26 ©Novell, Inc
SED-SIO-2121

SPX Packet Structure

SPX Header Structure

Connection Control	(1 octet)
Datastream Type	(1 octet)
Source Connection ID	(2 octets)
Dest. Connection ID	(2 octets)
Sequence Number	(2 octets)
Acknowledge Number	(2 octets)
Allocation Number	(2 octets)
Data	

IPX Header Structure

Checksum
Packet Length
Transport Control
Packet Type
Destination Network
Destination Node
Destination Socket
Source Network
Source Node
Source Socket
Data

27

NetWare Core Protocol (NCP)

- TCP/IP counterparts: NFS, RPC
- Service requests and responses
- Transport layer services (connection and sequence numbers)
- Burst mode: provides adaptive sliding window

Structure of an NCP Packet

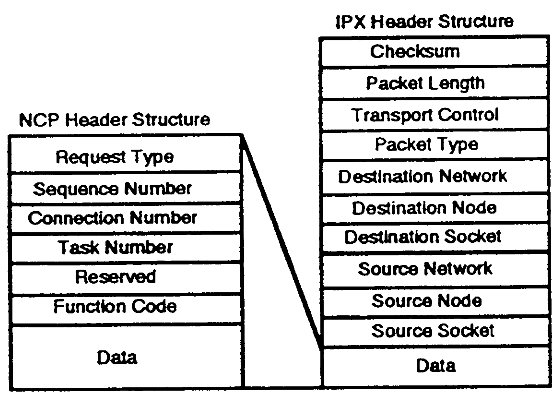

Burst Mode Protocol

Features

- Protocol utilized for file reads and writes
- Uses maximum size packets
- Error recovery retransmits only dropped packets
- Flow control is adaptive and self tuning

Burst Mode Packet Structure

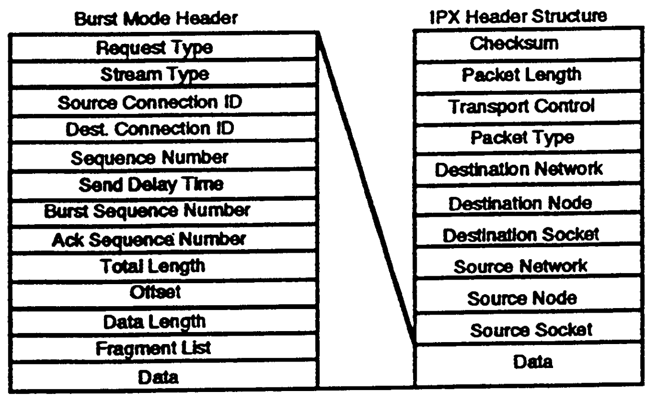

31

IPX WAN

RFC 1362

- Simple IPX protocol to facilitate WAN connectivity
- Works over PPP, X.25, Frame Relay, ISDN without media awareness
- Not limited to the above media
- Provides mechanism to calculate routing delay metrics
- Allows arbitration for common network number
- Provides mechanism to learn remote routers network number and name
- Extensible by third parties to new WAN Protocols and options

32

Future Directions

- SPX II
- NetWare Link Services Protocol
- NetWare Directory Services

33
SED-310-212

SPX II

- Sliding window implementation
- Tests and adapts for maximum frame size on internetwork links

NetWare Link Services Protocol

- Link State Protocol for IPX
- Based on IS-IS (SIN implementation)
- Works over IPX; independent of media access protocol (Ethernet, Token Ring, etc.)
- Backward compatible with RIP and SAP
- Configurable parameters

NetWare Directory Services

- Architecture derived from X.500
- Allows for hierarchical naming
- Allows for partitioning and replication of name database
- Provides for time synchronization

5.6 Internet Accounting

Presented by Cyndi Mills/BBN

Bio: *Cyndi Mills co-Chairs the Internet Engineering Task Force's Internet Accounting Working Group with Dr. Gregory Ruth. Cyndi has been involved with network and internetwork communications architecture issues, including OSI transition planning and heterogeneous protocol development and testing since 1978, first in Europe and then in the U.S. She currently manages the National Science Foundation (NSF) Network Services Center for BBN.*

Nov 1992

Internet Accounting

Cyndi Mills
Internet Engineering Task Force
Technical Presentation
19-Nov-92

C.Mills — BBN Systems and Technologies 1

Nov 1992

Overview

- Introduction
  - Let me count the ways...
- Accounting Architecture and Protocols
  - Nuts and bolts
  - Using the tools
- Where next?

C.Mills — BBN Systems and Technologies 2

Nov 1992

Goals of the Internet Accounting Working Group

- Provide an Internet accounting framework
  (RFC 1272, Internet Accounting Background)
- Provide a near term architecture for Internet accounting
  (Internet-Draft Internet Accounting Architecture)
- Provide a draft MIB for experimentation
  (Internet-Draft Internet Accounting: Meter Services Experimental MIB)

Why near term?

C.Mills — BBN Systems and Technologies 3

Nov 1992

Dopirak's First Law

If you live in the future, you'll die in the present.

C.Mills — BBN Systems and Technologies 4

Nov 1992

Benefits of Network Accounting

Accounting -
1. the art or system of keeping and analyzing financial records
2. an explanation of one's behavior
(merriam-webster)

Financial
- Allocate costs based on usage
- Generate revenue

Feedback
- Understand user behavior
- Motivate efficient use of resources
- Measure compliance with network policies

C.Mills — BBN Systems and Technologies 5

Nov 1992

Cost of Network Accounting

- Network overhead
  - maintaining complete counts (statistical sampling techniques used for network management are often incompatible with this goal)
  - performance and traffic
- Administrative overhead
  - report generation
  - report/bill distribution
  - user services
- Maintenance of traffic flow confidentiality
  - access control
  - encryption

C.Mills — BBN Systems and Technologies 6

Nov 1992

Conventional Telephone System

- Multi-vendor community requiring coordination of policy (tariff) and accounting information
  - need agreement on
    - billable entities
    - types of counts
  - need to exchange information

BBN Systems and Technologies 7

Nov 1992

Packets Aren't Calls: Connection-oriented vs Connectionless

Telephone system
- connection oriented (dedicated to a single subscriber or reserved)
- homogeneous type of service (single voice or data channel)

IP Internet
- datagram oriented (shared unpredictably by arbitrarily many subscribers)
- diverse types of service (video conference, e-mail)

Units metered per connection
- by subscriber
- elapsed time (start & end)
- constant quality of service

Units metered per datagram
- by subscriber
- per packet
- variable type of service

pipes vs. packets

BBN Systems and Technologies 8

Nov 1992

Concerns about Accounting on the Internet

"Fear and loathing in the research community"

- **Internet as Infrastructure**

  Free dissemination of information is beneficial to the community. Bulletin boards, e-mail, public databases and servers are currently available for "free".

- **Internet tariffs**

  Infrastructure economics are "different".

  User network behavior is modified through tariffs (cost-avoidance) and quality of service (frustration).

  Researchers fear that individual administrations will effectively inhibit users from rendering community services to avoid costs.

These factors lead to a preference for billing by pipe over billing by packet.

BBN Systems and Technologies 9

Nov 1992

BUT!

In some situations, usage based accounting is desirable:

- **Scarce (expensive) resources**

  Resources may be limited by cost, policy, or delays in aquisition. Usage-based accounting can influence resource hogs (e.g. through quotas or financial means) to fairly share scarce resources.

- **Competitive pressure**

  Local-area, low-volume traffic *is* cheaper to carry than high-volume global traffic. Subscribers with less expensive traffic patterns want to share in the cost savings.

- **Regulation**

  Gain experience with various cost recovery models including usage-based models to establish appropriate norms *before* regulation. (It is easier to recommend in advance than to change after the fact.)

In other words ...

BBN Systems and Technologies 10

Nov 1992

Dopirak's Second Law

Sleazebags never rest.

BBN Systems and Technologies 11

Nov 1992

Accounting Model

Usage Reporting, Billing, Network Management, and Policy Research are applications which operate on an accounting application data store.

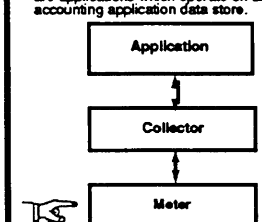

- **Accounting** *applications*: manage the accounting database, correlate information from separate metering systems, direct monitoring and control of meters.
- **Collection and Network Management** *protocols*: provide transport between applications and meters.
- **Meters perform the aggregation and attribution of** *counts* (packets, bytes).

BBN Systems and Technologies 12

Nov 1992

Minimize Overhead by Metering at the IP Level

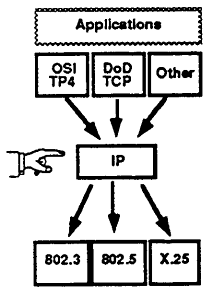

Protocol "Hourglass"

- **Readily Accessible Information**
  - **Hosts:** Application, Transport, IP
  - **IP Switches:** Internet Only
- **Minimize implementation and maintenance cost by minimizing the number of**
  - protocols
  - vendors
  - administrations
  - physical units

affected.

Nov 1992

Use Information Available in the IP Header

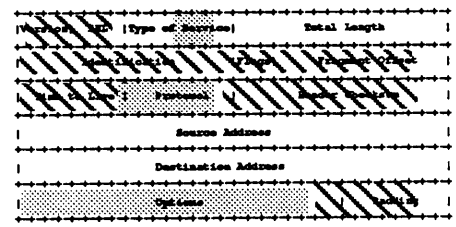

Nov 1992

View the Internet as a Loose Hierarchy for Internet Accounting

Problem

- Counting every packet for every user in every network at every router is too fine a granularity--too much overhead. (This is the same problem that routing protocols have!)

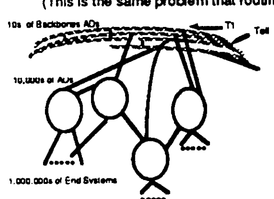

Approaches to Reducing Overhead

- Hold each network responsible / accountable for its own traffic.
- Perform accounting only at entry and/or exit gateways.
- Use coarser granularity for measuring transit traffic.
- Use multiple granularities to differentiate exception traffic.
- Use simple collection mechanisms real-time (if reliable enough) and post-process into accounting records.

Nov 1992

Selecting the "Right" Granularity

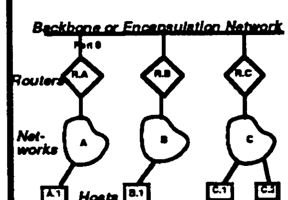

Granularity	Sample collector (subscriber)
Coarse: IF(Port 0 R.A)	*backbone (regional)*
Intermediate: IF(Port 0 R.A) to Adj(R.C)	*regional (regional)*
Mixed: NetAdr(A.1) toAdj R.C)	*local (host/regional)*
Fine: NetAdr(A.)1 to NetAdr(C.1)	*any (host/any)*

Nov 1992

Internet Accounting Process

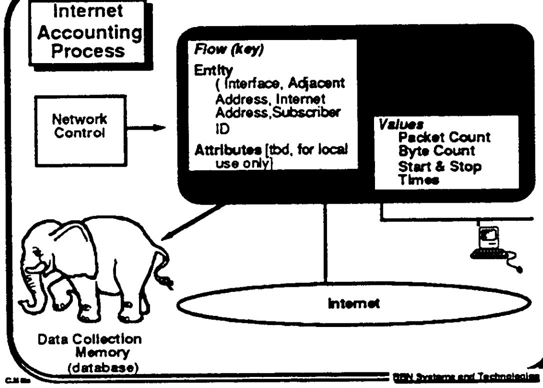

Nov 1992

Dopirak's Third Law

Never look in a dog's mouth.

Meter Control

Network management minimizes loss of data:

- starts/stops accounting at meter
- controls reporting interval and flow timeout (lifetime)
- regulates granularity of connection
- prioritizes groups of flows
- determines primary and alternate accounting data recipients

Collection Protocol

Collection protocol issues:

- **reliability**

  reliable protocol preferred to unacknowledged datagram service
- **security**

  data integrity, confidentiality, authentication, denial of service, non-repudiation desirable
- **atomicity of indexed entries**

  accounting record (complex key and values)
- **bulk data transfer**

  accounting produces large tables of data to be saved periodically
- **wildcarding objects to be retrieved**

  e.g. fetch all accounting records that have changed since time x.

SNMP currently used anyway.

Local storage of data at meter and transfer via ftp is also an option.

Internet Accounting Records

Each Record consists of a sequence of *flows*

Flows aggregate internet accounting information, with one *"bucket"* or set of counters per flow.

Information which spans multiple flows includes

- **Groups**

  An eight-bit identifier which allows flows to be prioritized and grouped for selective retrieval or discard
- **Scaling Factors**

  Where high volumes of traffic are anticipated, packets and/or byte counts may be shifted by a power of two to avoid counts larger than the protocol can handle, e.g. a count of 1 may represent (1*64 packets or 1*1024 packets, etc.)
- **Statistical Sampling Frequencies**

  In high-bandwidth applications it may not be feasible to count every single packet. In such cases, the average sampling frequency must be known, and the sampling algorithm sufficiently random and fair to satisfy the bean counters.

Flow Record

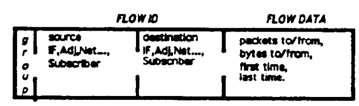

- **Flow Identifier**

  Protocol addresses and options *identify* traffic types and provide the information necessary to determine the accountable entity (the subscriber)
- **Values**

  Packet counts (actually PDUs or fragments), byte counts, start and stop times, *quantify* flows

Flow IDs

- **Source and Destination Addresses/Identifiers**

  Interface Address (port, physical layer)

  Adjacent Address (neighbor, Layer N-1)

  Network Address (IP address or NSAP address or ...)*

  (may extend to Transport Address, Session Address, Application Address, but these extensions are not explicitly defined by working group)

  Subscriber ID (optional)
- **Subscriber ID**

  covers dial-up/mobile subscribers, third-party billing, and the creation of arbitrary flows which count multiple address combinations into a single bucket
- ***Attributes***

  type of service, options, etc.
  (recommended for local use only, not currently defined by working group)

* *Need an address type RFC to standardize list and assign numbers*

Rule Table

The rule table is a list of comparisons which are performed on each packet to determine the correct flow bucket.

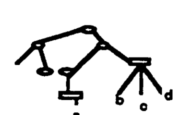

Rule tables must be kept as simple as possible to avoid overhead, particularly in routers.

A Nodal Rule describes

- **The address on which to operate** source, destination, layer, type
- **A mask for comparison** e.g. 010010***1
- **Action to take if comparison true (fall through to next rule if false)**
  - goto another rule

A Leaf Rule describes the flow bucket with masks.

e.g. 010010***1 to individuate (b,c,d)
or 0100100001 to aggregate (a)
or empty to ignore (don't count)

Nov 1992

Rule Table (cont.)

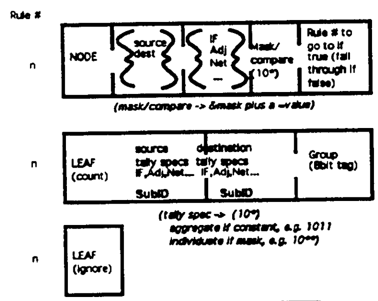

BBN Systems and Technologies 25

Nov 1992

Implementation Tip: Using Trees to Avoid Processing the Rule Table

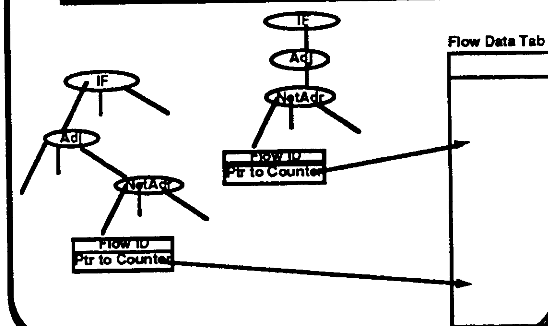

BBN Systems and Technologies 26

Nov 1992

Dopirak's Fourth Law

The chainsaw won't help you find the tree.

BBN Systems and Technologies 27

Nov 1992

Future Issues

- **Finer Granularities**
  - Quality of Service
  - Host participation in identifying end-users
  - Upper-layer protocols (application-level accounting)
- **Quota Systems**
  - Limiting costs and resources
  - Real-time interactions with access control
- **Additional Features**
  - Reserved Bandwidth
  - Inter-Domain Policy Routing
  - Distance
- **Standardized File Formats**
  - Exchanging accumulated and/or reduced usage data between accounting applications

BBN Systems and Technologies 28

Nov 1992

How You Can Help!

- If we're changing IP anyway, think about subscriber identification.
- A reliable secure network management/collection protocol that handles complex atomic objects and transfers tables efficiently wouldn't hurt either.
- An RFC assigning numbers which identify protocols would be useful.

BBN Systems and Technologies 29

Nov 1992

Summary

1. If you live in the future you'll die in the present.
2. Sleazebags never rest.
3. Never look in a dog's mouth.
4. The chainsaw won't help you find the tree.

BBN Systems and Technologies 30

Chapter 6

Process for Organization of Internet Standards

6.0.1 Process for Organization of Internet Standards (poised)

Charter

Chair(s):

Steve Crocker, crocker@tis.com

Mailing Lists:

General Discussion: poised@nri.reston.va.us
To Subscribe: poised-request@nri.reston.va.us
Archive: nri.reston.va.us:~/poised/current

Description of Working Group:

The goal of this Working Group is to examine the Internet standards process and the responsibilities of the IAB, with attention to the relationship between the IAB and IETF/IESG.

The need for this Working Group was suggested during discussions at the July 1992 IETF. This led to a request from the Internet Society president to form such a Working Group.

The Working Group will consider the following matters:

1. Procedures for making appointments to the Internet Architecture Board.

2. Procedures for resolving disagreements among IETF, IESG and IAB in matters pertaining to the Internet Standards.

3. Methods for assuring that for any particular Internet Standard, procedures have been followed satisfactorily by all parties so that everyone with an interest has had a fair opportunity to be heard.

The Working Group will begin with a review of the procedures for making IAB appointments as documented in RFC 1358 and a review of the standards-making process documented in RFC 1310.

The Working Group has a goal of issuing a final report in time for IESG consideration and publication as an RFC before the ISOC Board Trustee's meeting in December 1992. Given the compressed timescale, the Working Group will conduct most of its deliberations by electronic mail on the POISED Working Group mailing list. There will also be a preliminary report and discussions at the November 1992 IETF meeting in Washington, DC.

This will be a normal IETF Working Group, i.e., the mailing list and all discussions will be completely open.

Goals and Milestones:

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

Done	Gather initial set of issues and write a preliminary report.
Oct 1992	Post as an Internet-Draft the initial recommendations to the ISOC Board.
Done	Open discussion and presentation of the work of the POISED Working Group at Washington D.C. IETF meeting.
Dec 1992	Submit the recommendations document to the IESG for posting as an Informational RFC. This document will be subsequently transmitted to the ISOC Board.

CURRENT MEETING REPORT

Reported by Steve Crocker/TIS

Minutes of the Process for Organization of Internet Standards Working Group (POISED)

This report provides a summary of the POISED Working Group, starting from the events leading to the formation of the Working Group to the end of 1992. Necessarily, this synopsis represents my own perception, particularly for the "prehistory" period. Quite a few people hold strong views about both the overall sequence and specific events. My intent here is to convey as neutral a point of view as possible.

Background and Formation of POISED Working Group

The POISED Working Group resulted from two sequences of activity, both intimately related to the growth of the Internet. During 1991, there was great concern that the IP address space was being depleted and that the routing tables were growing too large. Some change in the IP addressing and routing mechanisms seemed inevitable, and it became urgent to explore and choose what those changes should be. The ROAD Working Group was formed to study the issues and recommend changes. The ROAD Group returned with a specific recommendation for the short-term, but did not reach a conclusion on a long term-plan.

The IESG then formulated a plan of action for further exploration of the issues and forwarded these recommendations to the IAB. In June 1992, after the INET '92 meeting in Kobe, Japan, the IAB met and considered the IESG's recommendations. After considering the IESG's recommendations, the IAB felt that additional ideas were also important, particularly some of the addressing ideas in the CLNP protocol. The IAB communicated its concerns, and there was immediate controversy along two dimensions. One dimension was technical: What is the best course for evolving the IP protocol? How important or useful are the ideas in the OSI protocol stack? The other dimension was political: Who makes decisions within the Internet community? Who chooses who makes these decisions?

As often happens during periods of conflict, communication suffered among the several parties. The June communication from the IAB was understood by many as an IAB decision or, equivalently, a sense of the decisions the IAB would make in the future. In contrast, many if not all on the IAB felt that they were trying to open up the discussion and their memos were intended as advice and not decisions. From my perspective, this form of miscommunication was partly due to the extended size of the Internet technical community. When the community was much smaller, the IAB was in close contact with the day to day workings of the technical groups. With the creation of the IESG and Area Directorates, there are now two or three layers between a working group and the IAB.

These matters came to a head during the IETF meeting in July in Cambridge, Massachusetts. It was made clear that the consideration of changes to the IP protocol remained open. Work on that topic has proceeded and is reported in the appropriate forums.

However, it became clear that it was necessary to examine the decision process and the procedures for populating the IESG and IAB. With respect to the procedures for selecting IAB and IESG members, the procedures that were in place derived from the creation of the Internet Society (ISOC) and the ISOC's sponsorship of the IAB. These procedures had been developed during the early part of 1992 and had been adopted by the ISOC during its meeting in Kobe in June. Hence, as fast as the ISOC was building the framework for supporting the Internet community, the community was questioning its structure and processes.

Following the IETF meeting, Vint Cerf, Internet Society president, called for the formation of a working group to examine the processes and particularly the selection process During August, the Working Group was formed, I was asked to chair it, and a Charter for the Working Group was formulated. (The acronym is due to Erik Huizer and originally stood for The Process for Organization of Internet Standards and Development. It was shortened to fit into the space available on paper and in the IETF Secretariat's database.)

Deliberations: August through mid-November

The formation of the POISED Working Group provided a forum for discussion of process issues. An estimated 20 MB of messages filled up disks all over the world. Much of this discussion was fragmented or focused on narrow issues. The salient point that emerged was the need for a well defined process for selecting leaders with explicit community representation in the selection process. There was also substantial discussion of the role of the IAB – to what extent should it make decisions and to what extent should it provide technical guidance? – and the relationship between the IAB and IESG.

After several weeks of discussion, Carl Malamud and I attempted to capture the main elements of the discussion by presenting a specific proposal for the reorganization of the entire structure. The main elements of the proposal were:

- Retention of the Working Group and area structure now in place within the IETF.
- Replacement of the IAB and IESG by two boards, one devoted to technical management and one devoted to oversight of the process.
- Well defined terms for members of both boards.
- Selection by committees with input from the community.

This proposal was technically radical in the sense that it proposed new structures to replace existing structures instead of proposing changes within the existing system. The proposal focused all further discussion and set the stage for the fall IETF in mid-November in Washington D.C.

November IETF Meeting

By virtue of the intensity of interest throughout the community, the POISED Working Group was one of the focal points of the IETF meeting. The schedule included a plenary session Tuesday morning to present the current state of the POISED Working Group discussions, a formal POISED Working Group session Tuesday afternoon and an open IESG meeting Thursday evening devoted to the POISED issues. The formal schedule was only the tip of the iceberg; numerous meetings took place over breakfast, lunch and dinner, in the halls and off in the corners. The more active participants probably had a dozen or more separate meetings on this over the three most active days, Tuesday, Wednesday and Thursday.

Amidst all this frenetic activity, remarkable progress occurred at two key points. At the Tuesday afternoon POISED Working Group meeting, Lyman Chapin, IAB Chair, Phill Gross, IETF Chair and IAB member, and other IAB members proposed changes within the existing IAB/IESG structure which converged with the process management elements of the Malamud-Crocker proposal. The key point was that all processing of standards actions, including the final decision to advance a specification along the standards track, would be made by the IESG. This change in the process shortens the decision cycle and brings it a step closer to the Working Group. Convergence on this key point obviated a radical proposal and signaled the building of a consensus on how the standards process should evolve. Over the next two days attention then turned to the selection process.

As indicated above, there was a strong feeling in the community that the IAB and IESG members should be selected with the consensus of the community. A natural mechanism for doing this is through formal voting. However, a formal voting process requires formal delineation of who's enfranchised. One of the strengths of the IETF is there isn't any formal membership requirement, nor is there a tradition of decision through votes. Decisions are generally reached by consensus with mediation by leaders when necessary.

Various formulas were considered, and the one that emerged was that IAB and IESG members would be selected by a nomination and recruiting committee. The committee is to consist of seven members from the community, with non-voting representatives from the IAB and IESG and a non-voting Chair provided by the ISOC. The seven members are to be volunteers, with selection by lot if there are more than seven volunteers. The only requirement for volunteers is they must have attended two IETF meetings. This requirement is designed to ensure the nomination committee has some familiarity with the Internet community and the standards process.

IAB and IESG members are to serve two years. Half of each body is to have terms starting in odd years, and half is to have terms starting in even years. Selections to the IESG have to be ratified by the IAB, and selections to the IAB had to be ratified by the ISOC. In the event that the nomination committee is unable to reach a consensus on a single candidate for each position, it may forward multiple nominations to the ratifying body, and the ratifying body will select the candidate.

In addition to this selection process, a recall mechanism was outlined using a similar scheme. The ISOC is to supply an ombudsman who will field complaints after all oversight processes have been exhausted. If the ombudsman is unable to resolve a complaint after a cooling off period, a recall committee, selected at random among volunteering community members, will consider the matter. A two thirds vote by the committee is necessary to remove someone.

This proposal was formulated and circulated during Wednesday and Thursday and presented at the IESG Open Plenary. In contrast to the extraordinarily contentious open IESG meeting in Cambridge, this meeting was characterized by a strenuous effort by numerous people, representing diverse points of view, to reach consensus on this proposal, and the meeting ended with a distinct decision to proceed on this basis. Given the strong consensus that emerged at that meeting, the Group decided to implement the selection process by the next IETF meeting, with the new IESG and IAB members to begin their terms at the termination of the IETF meeting in March.

On Friday, the IAB and IESG met jointly to determine what to do next. Both Groups agreed to implement the change in processing standards actions quickly and cooperatively and to identify the positions which are open for selection. Within a couple of weeks, the IAB finished processing the standards actions in its queue, and IESG began to handle standards actions on its own.

December ISOC Meeting

The Internet Society Trustees met December 10th and 11th at CNRI in Reston, VA. The process and organization of the IAB and IETF was one of their major concerns. A session of the Trustees at 3:00 p.m. EST, December 10th, was broadcast via the Internet. It was not clear how many people listened, but Geoff Huston, Internet Trustee, was spliced in separately from Australia.

At this session, I presented the POISED Working Group results deliberations and asked on behalf of the IETF that the Trustees approve the selection process described above. For the long run, a new charter is needed. Given the very compressed schedule for these activities, there has not been time to draft and refine a new charter, so the Trustees were asked to approve the general direction of the reorganization of the IAB and IETF and give temporary approval to the selection process in order to permit the first round of selections to proceed.

The Trustees expressed strong approval for the work of the POISED Working Group and general approval for the direction of the effort. One area of concern for the Trustees is the legal liability of the Internet Society regarding decisions the IESG might make in the future. The Trustees made it quite clear that they are not inclined to micromanage the IETF process, but they do feel compelled to understand the legal issues and help construct a charter which is consistent with their responsibilities as Trustees.

The session adjourned with agreement to proceed on the current course and for the IETF, IAB and ISOC Trustees to work together to draft the appropriate charter.

Future Activities

Both the IESG and IAB have selected the positions which must be filled through the new selection process. As I write this, Vint Cerf has been working to find a Chair for the nominations committee, and the process should move forward during January and February. Communications on the details of the nomination process will be published on the IETF mailing list and possibly other forums. As described above, the selection process should be complete well before the next IETF meeting, and preferably by the end of February.

The other open Agenda item is the draft of a new Charter for the IAB and IETF and adoption of the Charter by the ISOC. This is the next order of business for the POISED Working Group.

Attendees

Jeff Case	case@cs.utk.edu
Vinton Cerf	vcerf@cnri.reston.va.us
A. Lyman Chapin	lyman@bbn.com
Stephen Crocker	crocker@tis.com
James Davin	davin@bellcore.com
Frank Kastenholz	kasten@ftp.com
Carl Malamud	carl@malamud.com
Craig Partridge	craig@bbn.com
Michael Roberts	mmr@isoc.org

Appendix A

Attendees

George Abe
Infonet Services Corporation
Phone: 310-335-2867
Fax: 310-335-2876
abe@infonet.com

Ahmed Abutaleb
Sprint International
Phone: 703-689-7172
Fax: 703-689-5633
a.abutaleb@sprint.com

Roland Acra
cisco Systems, Europe
Phone: +33 1 60 92 20 00
Fax: +33 1 69 28 83 26
acra@cisco.com

Elizabeth Adams
NM Forum
Phone: 908-766-1544
Fax: 908-766-5741
adamse@attmail.com

Vikas Aggarwal
JvNCnet
Phone: 609-258-2403
Fax: 609-258-2424
vikas@jvnc.net

Robert Aiken
Lawrence Livermore National Laboratory
Phone: 510-422-4017
Fax: 510-422-0435
aiken@es.net

Cengiz Alaettinoglu
University of Maryland
Phone: 301-405-2723
Fax: 301-405-6707
ca@cs.umd.edu

Ed Albrigo
Corporation for Open Systems International
Phone: 703-205-2723
Fax: 703-846-8590
ealbrigo@cos.com

Steve Alexander
INTERACTIVE Systems Corporation
Phone: 708-505-9100
Fax: 708-505-9133
stevea@i88.isc.com

James Alfieri
Bell Communications Research
Phone: 908-758-5460
Fax: 908-758-4331
jdal@troy.cc.bellcore.com

James Allard
Microsoft Corporation
Phone: 206-882-8080
Fax: 206-936-7329
jallard@microsoft.com

Lee Alley
Arizona State University
Phone: 602-965-6398
Fax: 602-965-1608
iadlra@asuvm.inre.asu.edu

Claudio Allocchio
National Institute for Nuclear Physics - Italy
Phone: +39 40 3758523
Fax: +39 40 226338
Claudio.Allocchio@elettra.trieste.it

Guy Almes
Advanced Network and Services, Inc.
Phone: 914-789-5333
Fax: 914-789-5310
almes@ans.net

Philip Almquist
Consultant
Phone: 415-752-2427
almquist@jessica.stanford.edu

Harald Alvestrand
SINTEF DELAB
Phone: +47 7 59 70 94
Fax: +47 7 53 25 86
Harald.Alvestrand@delab.sintef.no

Joseph Apple
Sprint Government Systems Division
Phone: 703-904-2353
Fax: 703-904-2351
japple@lcm1.icp.net

David Arneson
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-424-8121
arneson@ctron.com

Jules Aronson
National Library of Medicine
Phone: 301-496-9300
Fax: 301-496-0673
aronson@nlm.nih.gov

Nagaraj Arunkumar
3Com Corporation
Phone: 408-764-5104
Fax: 408-764-5003
nak@3com.com

Toshiya Asaba
WIDE Project
Phone: +81 473 50 1658
Fax: +81 473 50 1658
asaba@wide.sfc.keio.ac.jp

Randall Atkinson
Naval Research Laboratory
Phone: 202-767-1111
Fax: 202-404-8590
atkinson@itd.nrl.navy.mil

Robert Austein
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sra@epilogue.com

Floyd Backes
Digital Equipment Corporation
Phone: 508-486-7076
backes@dsmail.enet.dec.com

Cynthia Bagwell
The MITRE Corporation
Phone: 703-883-5438
Fax: 703-883-7142
cbagwell@gateway.mitre.org

Fred Baker
Advanced Computer Communications
Phone: 805-685-4455
Fax: 805-685-4465
fbaker@acc.com

David Balenson
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
balenson@tis.com

Anthony Ballardie
University College London
Phone: +44 71 387 7050
A.Ballardie@cs.ucl.ac.uk

David Ballowe
Compatible Systems Corporation
Phone: 303-444-9532
Fax: 303-444-9595
ballowe@compatible.com

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

Richard Basch
Massachusetts Institute of Technology
Phone: 617-253-0100
Fax: 617-258-8736
basch@mit.edu

Brian Bataille
United States Air Force
Phone: 505-846-2828
Fax: 505-846-2585
bataillebc@afotec.af.mil

Tony Bates
University of London Computer Center
Phone: +44 71 405 8400
Fax: +44 71 242 1845
t.bates@nosc.ja.net

Jordan Becker
Advanced Network and Services, Inc.
Phone: 914-789-5308
Fax: 914-789-5310
becker@ans.net

Bryan Beecher
University of Michigan
Phone: 313-747-4050
bryan@umich.edu

Tom Benkart
ACC Systems
Phone: 410-290-8100
Fax: 410-290-8106
teb@saturn.acc.com

Ken Benstead
Coral Network Corporation
Phone: 508-366-3600
Fax: 508-870-1777
kbenstead@coral.com

Lou Berger
Consultant
Phone: 301-530-9790
lberger@bbn.com

Pravin Bhagwat
IBM Corporation
Phone: 914-784-6866
pravin@cs.umd.edu

Shiraz Bhanji
The MITRE Corporation
Phone: 703-883-6527
Fax: 703-883-7142
bhanji@gateway.mitre.org

Edo Biagioni
Fore Systems, Inc.
Phone: 412-967-4040
Fax: 412-967-4044
esb@fore.com

Andy Bierman
SynOptics Communications, Inc.
Phone: 408-764-1673
Fax: 408-988-5525
abierman@synoptics.com

Dean Blackketter
Apple Computer, Inc.
Phone: 408-974-4213
Fax: 408-974-0781
deanb@apple.com

Fred Bohle
Interlink Computer Sciences
Phone: 301-317-6600
Fax: 301-317-6608
fab@interlink.com

Frederick Boland
National Institute of
Standards and Technology
Phone: 301-975-3608
Fax: 301-975-2128
boland@ecl.nist.gov

David Bolen
Advanced Network and
Services, Inc.
Phone: 914-789-5327
Fax: 914-789-5310
db31@ans.net

Nathaniel Borenstein
Bell Communications Research
Phone: 201-829-4270
Fax: 201-829-5963
nsb@thumper.bellcore.com

David Borman
Cray Research, Inc.
Phone: 612-683-5571
Fax: 612-683-5599
dab@cray.com

Rebecca Bostwick
Lawrence Livermore National
Laboratory
Phone: 510-422-4017
Fax: 510-422-1482
bostwick@es.net

Bill Bowman
Hewlett-Packard
Phone: 415-691-5678
Fax: 415-691-5956
bill@hpprsd.mayfield.hp.com

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

Daniel Brennan
Teleos Communications, Inc.
Phone: 908-544-6432
dmb@teleoscom.com

Rita Brennan
Apple Computer, Inc.
Phone: 408-862-8063
Fax: 408-974-2691
brennan@apple.com

George Brett
MCNC - CNIDR
Phone: 919-248-1886
Fax: 919-248-1405
George.Brett@cnidr.org

Scott Brim
Cornell University
Phone: 607-255-9392
Fax: 607-255-5771
Scott_Brim@cornell.edu

Caralyn Brown
Wellfleet Communications, Inc.
Phone: 617-275-2400
Fax: 617-275-5001
cbrown@wellfleet.com

Jack Brown
United States Army
Computer Engineer Center
Phone: 602-533-2834
Fax: 602-533-5676
jbrown@huachuca-emh8.army.mil

Vickie Brown
Computer Sciences Corporation
Phone: 301-937-0760
Fax: 301-937-0826
brown@osi540sn.gsfc.nasa.gov

J. Nevil Brownlee
University of Auckland
Phone: +64 9 373 7599
Fax: +64 9 373 7425
nevil@aukuni.ac.uz

Gregory Bruell
Wellfleet Communications, Inc.
Phone: 617-275-2400
Fax: 617-275-5001
gob@wellfleet.com

Theodore Brunner
Bell Communications Research
Phone: 201-829-4578
Fax: 201-829-2504
tob@thumper.bellcore.com

Steve Bucey
Pacific Bell
Phone: 916-972-2066
Fax: 916-483-0877
sabucey@ns.pacbell.com

Philip Budne
Shiva Corporation
Phone: 617-252-6356
Fax: 617-252-6852
phil@shiva.com

Jeffrey Burgan
NASA Ames Research Center
Phone: 415-604-5705
Fax: 415-604-0063
jeff@nsipo.nasa.gov

Matthew Busche
AT&T Bell Laboratories
Phone: 908-949-4136
Fax: 908-949-0707
mtb@anchor.ho.att.com

Stephen Bush
General Electric Information Services
Phone: 216-362-5718
Fax: 216-362-5606
sfb@ncoast.org

Ross Callon
Digital Equipment Corporation
Phone: 508-486-5009
Fax: 508-486-5279
callon@bigfut.lkg.dec.com

C. Allan Cargille
University of Wisconsin
Phone: 608-262-5084
Fax: 608-262-9777
cargille@cs.wisc.edu

Ken Carlberg
SAIC - CSEIC
Phone: 703-734-9000
Carlberg@cseic.saic.com

Brian Carpenter
CERN
Phone: +41 22 767 4967
Fax: +41 22 767 7155
brian@dxcern.cern.ch

Charles Carvalho
Advanced Computer Communications
Phone: 805-685-4455
Fax: 805-685-4465
charles@acc.com

Jeff Case
University of Tennessee
Phone: 615-573-1434
Fax: 615-573-9197
case@cs.utk.edu

Stephen Casner
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
casner@isi.edu

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: 908-336-2830
gkc@ctt.bellcore.com

John Chang
IBM Corporation
Phone: 919-254-4393
Fax: 919-254-5483
changj@ralvm6.vnet.ibm.com

Kay Chang
IBM Corporation
Phone: 512-838-3542
Fax: 512-838-3509
chang@chang.austin.ibm.com

Wo Chang
National Institute of
Standards and Technology
Phone: 301-975-3439
Fax: 301-963-9137
wchang@nist.gov

Yee-Hsiang Chang
MCNC
Phone: 919-248-1169
Fax: 919-248-1405
yhc@concert.net

A. Lyman Chapin
Bolt Beranek and Newman
Phone: 617-873-3133
Fax: 617-873-4086
lyman@bbn.com

Brett Chappell
NSWC-DD
Phone: 703-663-1566
Fax: 703-663-8673
bchappe@relay.nswc.navy.mil

Dilip Chatwani
SynOptics Communications, Inc.
Phone: 408-764-1624
dilip@synoptics.com

Szusin Chen
Sun Microsystems, Inc.
Phone: 415-336-2424
Fax: 415-336-3153
szusin.chen@eng.sun.com

Dean Cheng
Retix
Phone: 310-828-3400
Fax: 310-829-5278
dean@sun2.retix.com

J. Noel Chiappa
Phone: 804-898-8183
Fax: 804-898-7663
jnc@lcs.mit.edu

Robert Ching
Network Application Technology
Phone: 408-370-4270
Fax: 408-370-4222
rching@nat.com

Chris Chiotasso
Ungermann-Bass, Inc.
Phone: 508-687-4445
chris@andr.ub.com

Chi Chong
SynOptics Communications, Inc.
Phone: 408-764-1466
Fax: 408-988-5525
cchong@synoptics.com

Jim Christy
Apple Computer, Inc.
Phone: 408-974-3080
christy@apple.com

Jodi-Ann Chu
University of Hawaii
Phone: 808-956-2400
Fax: 808-956-2412
jodi@uhunix.uhcc.hawaii.edu

George Clapp
Ameritech Advanced Data Services
Phone: 708-248-6507
Fax: 708-248-6038
clapp@ameris.center.il.ameritech.com

Cynthia Clark
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
cclark@cnri.reston.va.us

Henry Clark
OARnet
Phone: 614-292-6483
Fax: 614-292-7168
henryc@oar.net

Bobby Clay
Boeing
Phone: 205-544-9035
clay@eagle.msfc.nasa.gov

John Clement
EDUCOM
Phone: 202-466-6296
Fax: 202-872-4318
clement@educom.edu

Tim Clifford
Sprint
Phone: 202-416-6211
Fax: 202-833-8621
tcliff@icml.icm.net

Richard Cogger
Cornell University
Phone: 607-255-7566
Fax: 607-255-5771
rhx@cornell.cit.bitnet

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@osi.ncsl.nist.gov

Robert Collet
Sprint Government Systems Division
Phone: 703-904-2353
Fax: 703-904-2115
rcollet@icm1.icp.net

Michael Collins
Lawrence Livermore National Laboratory
Phone: 510-422-4018
Fax: 510-422-0435
collinsms@es.net

Rob Coltun
Consultant
Phone: 301-340-9416
rcoltun@ni.umd.edu

David Conklin
JvNCnet
Phone: 609-258-2414
Fax: 609-258-2424
conklin@jvnc.net

James Conklin
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
jbc@bitnic.educom.edu

Michael Conn
MCI
Phone: 214-918-5710
Fax: 214-918-6024
4387451@mcimail.com

John Cook
Chipcom Corporation
Phone: 508-490-5654
Fax: 508-460-8952
cook@chipcom.com

R. Bryan Cook
IBM Corporation
Phone: 914-435-7914
Fax: 914-432-9413
rbcook@tdcsys2.vnet.ibm.com

Naomi Courter
MCNC - CONCERT Network
Phone: 919-248-1807
Fax: 919-248-1405
naomi@concert.net

Tracy Cox
Bell Communications Research
Phone: 908-758-2107
Fax: 908-758-4177
tacox@sabre.bellcore.com

Stephen Coya
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
scoya@cnri.reston.va.us

Wesley Craig
University of Michigan
Phone: 313-764-2278
Fax: 313-764-5140
wes@umich.edu

Chuck Cranor
Washington University
Phone: 314-935-4203
chuck@maria.wustl.edu

David Crocker
The Branch Office
Phone: 408-246-8253
Fax: 408-249-6205
dcrocker@mordor.stanford.edu

Stephen Crocker
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
crocker@tis.com

Juan Cruz
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-2688
juan@dss.com

Dave Cullerot
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-335-1739
cullerot@ctron.com

Cathy Cunningham
Microcom, Inc.
Phone: 617-551-1131
Fax: 617-551-1015
cmc@microcom.com

John Curran
NEARnet
Phone: 617-873-4398
Fax: 617-873-5620
jcurran@bbn.com

John Dale
Corporation for Open Systems
Phone: 703-205-2742
Fax: 703-846-8590
jdale@cos.com

Nabil Damouny
NEC America, Inc.
Phone: 408-433-2956
Fax: 408-433-2950
nabil@tdd.sj.nec.com

Bruce Davie
Bell Communications Research
Phone: 201-829-4838
Fax: 201-829-2504
bsd@bellcore.com

Megan Davies
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
mdavies@cnri.reston.va.us

James Davin
Bell Communications Research
Phone: 201-829-4798
Fax: 201-829-2504
davin@bellcore.com

Michael Davis
Clearpoint, Inc.
Phone: 508-435-7452
Fax: 508-435-7504
mad@spirit.clearpoint.com

Michael Davison
FiberCom, Inc.
Phone: 703-342-6700
Fax: 703-342-5961
davison@fibercom.com

Osmund de Souza
AT&T Bell Laboratories
Phone: 908-949-1393
Fax: 908-949-4673
osmund.desouza@att.com

Michael DeAddio
Bell Communications Research
Phone: 201-829-5024
Fax: 201-829-5204
deaddio@thumper.bellcore.com

Steve Deering
Xerox Corporation
Phone: 415-812-4839
Fax: 415-812-4471
deering@parc.xerox.com

Cynthia Della Torre
The MITRE Corporation
Phone: 703-883-1914
Fax: 703-883-7142
cindy@gateway.mitre.org

Barbara Denny
SRI International
Phone: 415-859-2812
Fax: 415-859-4812
denny@erg.sri.com

Art Dertke
The MITRE Corporation
Phone: 703-883-7155
Fax: 703-883-1397
dertke@gateway.mitre.org

Richard desJardins
NASA
Phone: 301-286-8554
Fax: 301-286-5269
desjardi@boa.gsfc.nasa.gov

Peter Deutsch
Bunyip Information Systems
Phone: 514-398-3709
Fax: 514-398-3709
peterd@bunyip.com

Manuel Diaz
David Systems, Inc.
Phone: 408-720-6865
Fax: 408-720-1337
diaz@davidsys.com

Peter DiCamillo
Brown University
Phone: 401-863-7582
Fax: 401-863-7329
Peter_DiCamillo@brown.edu

Tim Dixon
RARE
Phone: +31 20 639 1131
Fax: +31 20 639 3289
dixon@rare.nl

Kurt Dobbins
Cabletron Systems, Inc.
Phone: 603-337-5049
Fax: 603-424-8121
dobbins@ctron.com

Greg Dobrich
National Agricultural Library
Phone: 301-504-6813
Fax: 301-504-7473
gdobrich@nalusda.gov

Jack Drescher
MCNC - CONCERT Network
Phone: 919-248-1412
Fax: 919-248-1405
drescher@concert.net

Jon Dreyer
SunSelect
Phone: 508-442-0385
Fax: 508-442-0011
Jon.Dreyer@east.sun.com

Ralph Droms
Bucknell University
Phone: 717-524-1145
Fax: 717-524-3760
droms@bucknell.edu

David Dubois
Pacer Software, Inc.
Phone: 508-898-3300
Fax: 508-366-1356
dad@pacersoft.com

Kevin Dubray
Wellfleet Communications, Inc.
Phone: 617-280-2362
Fax: 617-280-2300
kdubray@wellfleet.com

Letha Dugas
MCI
Phone: 202-416-5685
Fax: 202-416-5858
4371362@mcimail.com

Jacques Dugast
France Telecom
Phone: +33 1 45 29 43 74
Fax: +33 1 45 29 65 17
dugast@issy.cnet.fr

Thomas Dunkenberger
Sprint Government Systems Division
Phone: 202-416-6269
Fax: 202-833-8621
tdunk@icm1.icp.net

Pierre Dupont
Motorola, Inc.
Phone: 604-241-6243
Fax: 604-241-6243
dupont@mdd.comm.mot.com

Steve Dusse
RSA Data Security, Inc.
Phone: 415-595-7703
Fax: 415-595-4126
spock@rsa.com

Tom Easterday
CICNet, Inc.
Phone: 313-998-6285
Fax: 313-998-6105
tom@cic.net

Donald Eastlake
Digital Equipment Corporation
Phone: 508-486-2358
Fax: 508-486-2311
dee@ranger.enet.dec.com

Raymond Edgerton
Bell Atlantic Business Systems
Phone: 215-296-6159
Fax: 215-296-0460
bellmf.edgerton@uwm.edu

William Edison
IBM Corporation
Phone: 301-240-6899

Alan Emtage
Bunyip Information Systems
Phone: 514-398-8117
Fax: 514-398-6876
bajan@bunyip.com

David Engel
Optical Data Systems, Inc.
Phone: 214-234-6400
Fax: 214-234-4059
david@ods.com

Robert Enger
Advanced Network and Services, Inc.
Phone: 703-758-7722
Fax: 703-758-7717
enger@reston.ans.net

Hans Eriksson
Swedish Institute of Computer Science
Phone: +46 8 752 1527
Fax: +46 8 751 7230
hans@sics.se

Michael Erlinger
Harvey Mudd College
Phone: 714-693-2617
Fax: 714-621-8465
mike@jarthur.claremont.edu

Susan Estrada
CERFnet
Phone: 619-455-3902
Fax: 619-455-3990
estradas@cerf.net

Deborah Estrin
University of Southern California
Phone: 213-740-4524
Fax: 213-740-7285
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: 301-402-1246
Fax: 301-480-6245
raf@cu.nih.gov

Dino Farinacci
cisco Systems
Phone: 415-688-4696
Fax: 415-688-8282
dino@cisco.com

Stefan Fassbender
GMD
Phone: +49 2241 14 3180
Fax: +49 2241 14 3002
stf@easi.net

Daniel Fauvarque
Sun Microsystems, Inc.
Phone: +33 76 41 42 26
Fax: +33 76 41 42 41
dfauvarq@france.sun.com

Steven Fellini
National Institutes of Health
Phone: 301-496-4823
steve@helix.nih.gov

Sallie Fellows
National Forum on Education Statistics
Phone: 603-271-2778
Fax: 603-271-1953
sallie%ed@psc.plymouth.edu

Dennis Ferguson
Advanced Network and Services, Inc.
Phone: 914-789-5367
Fax: 914-789-5310
dennis@ans.net

William Fink
NASA Goddard Space Flight Center
Phone: 301-286-9423
Fax: 301-286-5152
bill@wizard.gsfc.nasa.gov

Richard Fisher
NASA Goddard Space Flight Center
rfisher@cdhf1.gsfc.nasa.gov

Eric Fleischman
Boeing Computer Services
Phone: 206-957-5334
Fax: 206-957-5050
ericf@act.boeing.com

Wayne Foco
IBM Corporation
Phone: 919-254-4422
Fax: 919-254-5483
foco@ralvm6.vnet.ibm.com

Peter Ford
Los Alamos National Laboratory
Phone: 505-665-0058
Fax: 505-665-7793
peter@goshawk.lanl.gov

Richard Ford
Falcon Microsystems
Phone: 301-386-6440
Fax: 301-341-0187
72510.553@compuserve.com

James Forster
cisco Systems
Phone: 415-688-7854
Fax: 415-688-4575
forster@cisco.com

Jill Foster
University of Newcastle-Upon-Tyne
Phone: +44 91 222 8250
Fax: +44 91 222 8232
jill.foster@newcastle.ac.uk

Catherine Foulston
Rice University
Phone: 713-285-5233
Fax: 713-527-6099
cathyf@rice.edu

Barbara Fraser
CERT Coordination Center
Phone: 412-268-5010
Fax: 412-268-6989
byf@cert.org

Ned Freed
Innosoft International, Inc.
Phone: 714-624-7907
Fax: 714-621-5319
ned@innosoft.com

Raphael Freiwirth
RCI, Inc.
Phone: 508-453-7857
Fax: 508-934-9066
5242391@mcimail.com

Jerry Friesen
Sandia National Laboratories
Phone: 510-294-3144
Fax: 510-294-1225
jafries@sandia.llnl.gov

Karen Frisa
Carnegie Mellon University
Phone: 412-268-8576
Fax: 412-268-4987
karen.frisa@andrew.cmu.edu

Cliff Frost
University of California, Berkeley
Phone: 510-642-5360
Fax: 510-643-5385
cliff@cmsa.berkeley.edu

Shoji Fukutomi
Furukawa Electoric
Phone: +81 463 23 1045
Fax: +81 463 23 0416
fuku@furukawa.co.jp

Vince Fuller
BARRNet
Phone: 415-723-6860
Fax: 415-723-0010
vaf@stanford.edu

Jim Fullton
University of North Carolina
Phone: 919-962-9107
Fax: 919-962-5604
jim_fullton@unc.edu

Peter Furniss
Peter Furniss Consultants
Phone: +44 81 313 1833
Fax: +44 81 313 1833
p.furniss@ulcc.ac.uk

Shari Galitzer
The MITRE Corporation
Phone: 703-883-6229
Fax: 703-883-5200
shari@mitre.org

Shawn Gallagher
Digital Equipment Corporation
Phone: 508-952-3519
gallagher@quiver.enet.dec.com

James Galvin
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
galvin@tis.com

Ella Gardner
The MITRE Corporation
Phone: 703-883-5826
Fax: 703-883-7142
epg@gateway.mitre.org

Joan Gargano
University of California, Davis
Phone: 916-752-2591
Fax: 916-752-9154
jcgargano@ucdavis.edu

James Geddes
InternetWare, Inc.
Phone: 703-560-7926
Fax: 703-560-4364
wk05020@worldlink.com

Tony Genovese
Lawrence Livermore National Laboratory
Phone: 415-423-2471
Fax: 415-422-0435
genovese@es.net

Elise Gerich
University of Michigan
Phone: 313-936-2120
Fax: 313-747-3745
epg@merit.edu

Arlene Getchell
Lawrence Livermore National Laboratory
Phone: 510-423-6349
Fax: 510-422-0435
getchell@es.net

Robert Gilligan
Sun Microsystems, Inc.
Phone: 415-336-1012
Fax: 415-968-4298
Bob.Gilligan@eng.sun.com

Joseph Godsil
National Center for Supercomputing Applications
Phone: 217-244-7715
Fax: 217-244-1987
jgodsil@ncsa.uiuc.edu

Mike Goguen
Digital Equipment Corporation
Phone: 415-853-2268
Fax: 415-324-4873
goguen@src.dec.com

Masayoshi Gohara
National Center for Science
Information Systems
Phone: +81 3 3942 6948
Fax: +81 3 3942 9398
mg@sinet.ad.jp

Richard Graveman
Bell Communications Research
Phone: 908-699-4611
Fax: 908-336-2943
rfg@ctt.bellcore.com

Heather Gray
Digital Equipment Corporation
Phone: 603-881-1079
heather@zk3.dec.com

Terry Gray
University of Washington
Phone: 206-543-3880
Fax: 206-543-3909
gray@cac.washington.edu

Mark Green
Apple Computer, Inc.
Phone: 510-204-0934
Fax: 510-649-4008
markg@apple.com

Maria Greene
Xyplex, Inc.
Phone: 508-264-9900
Fax: 508-264-9930
mngreene@eng.xyplex.com

Paul Griffiths
IBM Corporation
Phone: 512-838-3544
Fax: 512-838-3509
griff@chang.austin.ibm.com

Pete Grillo
Network Innovations
Phone: 503-526-9766
Fax: 503-526-9766
pl0143@mail.psi.net

Phillip Gross
Advanced Network and
Services, Inc.
Phone: 914-789-5341
Fax: 914-789-5310
pgross@nis.ans.net

Olafur Gudmundsson
University of Maryland
Phone: 301-405-2767
Fax: 301-405-6707
ogud@cs.umd.edu

Michel Guittet
Apple Computer Europe, Inc.
Phone: +33 1 49 01 49 47
Fax: +33 1 40 90 09 54
guittet1@applelink.apple.com

Robert Gutierrez
NASA Science Internet
Phone: 415-604-3655
Fax: 415-604-0063
gutierre@nsipo.nasa.gov

Thomas Hacker
University of Michigan
Phone: 313-763-6188
Fax: 313-763-4434
hacker@citi.umich.edu

Robert Hagens
Advanced Network and
Services, Inc.
Phone: 703-758-7700
Fax: 703-758-7717
hagens@ans.net

William Haggerty
Cabletron Systems, Inc.
Phone: 603-337-5050
Fax: 603-424-8121
haggerty@ctron.com

Jack Hahn
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
hahn@sura.net

Tony Hain
Lawrence Livermore National Laboratory
Phone: 510-422-4200
Fax: 510-422-1482
alh@es.net

Neil Haller
Bell Communications Research
Phone: 201-829-4478
Fax: 201-829-2504
nmh@thumper.bellcore.com

Joel Halpern
Network Systems Corporation
Phone: 612-424-1606
Fax: 612-424-1736
jmh@network.com

Deborah Hamilton
AT&T Bell Laboratories
Phone: 908-949-9459
Fax: 908-949-1726
debbie@qsun.att.com

Patrick Hanel
Nokia Telecommunications
Phone: +358 0 50603729
Fax: +358 0 5061652
hanel@yoyodyne.dco.ntc.nokia.com

Craig Haney
Sprint
Phone: 703-904-2084
Fax: 703-904-2251
craig@icp.net

Alf Hansen
SINTEF DELAB
Phone: +47 7 592982
Fax: +47 7 532586
Alf.Hansen@delab.sintef.no

Daniel Hansen
Network General Corporation
Phone: 201-616-0986
Fax: 201-616-8311
dan@ngc.com

Steve Hardcastle-Kille
ISODE Consortium
Phone: +44 71 223 4062
Fax: +44 71 223 3846
s.kille@isode.com

Susan Hares
Merit Network, Inc.
Phone: 313-936-2095
Fax: 313-747-3745
skh@merit.edu

Eugene Hastings
Pittsburgh Supercomputing Center
Phone: 412-268-4960
Fax: 412-268-5832
hastings@psc.edu

Alisa Hata
University of Washington
Phone: 206-543-7732
Fax: 206-543-3909
hata@cac.washington.edu

John Hawthorne
Rome Research Corporation
Phone: 315-330-7928
Fax: 315-330-2396
johnh@tigger.rl.af.mil

Ken Hayward
Bell Northern Research
Phone: 613-763-4042
Fax: 613-763-2626
crm57d@bnr.ca

Frank Heath
Rockwell CMC
Phone: 714-833-6864
Fax: 714-833-6863
heath@cmc.com

John Hedderman
Advanced Network and
Services, Inc.
Phone: 914-789-5321
Fax: 914-789-5310
jjh@ans.net

Juha Heinanen
Telecom Finland
Phone: +358 49 500958
Fax: +358 31 2432211
juha.heinanen@datanet.tele.fi

Ed Heiner
Synernetics, Inc.
Phone: 508-670-9009
Fax: 508-670-9015
eah@pau.synnet.com

Marco Hernandez
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
marco@mh-slip.educom.edu

Robert Hinden
Sun Microsystems, Inc.
Phone: 415-336-2082
Fax: 415-336-6015
hinden@eng.sun.com

Masaki Hirabaru
University of Tokyo
Phone: +81 3 38122111
Fax: +81 3 38147279
hi@nic.ad.jp

Ken Hirata
Emulex Corporation
Phone: 714-662-5600
Fax: 714-754-0731
khirata@emulex.com

Russ Hobby
University of California,
Davis
Phone: 916-752-0236
Fax: 916-752-9154
rdhobby@ucdavis.edu

Don Hoffman
Sun Microsystems, Inc.
Phone: 415-336-4339
Fax: 415-965-4903
don.hoffman@eng.sun.com

Ellen Hoffman
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
ellen_hoffman@um.cc.umich.edu

J. Paul Holbrook
CICNet, Inc.
Phone: 313-998-7680
Fax: 313-998-6105
holbrook@cic.net

G. Ken Holman
Microstar Software, Ltd.
Phone: 613-727-5696
Fax: 613-727-9491
holman@tmn.com

Gerd Holzhauer
Apple Computer Europe, Inc.
Phone: +33 1 49 01 49 47
Fax: +33 1 40 90 09 54
holzhauer1@applelink.apple.com

Alton Hoover
Advanced Network and
Services, Inc.
Phone: 914-789-5325
Fax: 914-789-5310
hoover@ans.net

John Hopprich
David Systems, Inc.
Phone: 408-720-7459
Fax: 408-720-1337
hopprich@davidsys.com

Russell Housley
Xerox Special Information Systems
Phone: 703-790-3767
Fax: 703-790-3764
Housley.McLean_CSD@Xerox.Com

Nat Howard
Bell Communications Research
Phone: 201-829-5196
Fax: 201-829-2645
nrh@bellcore.com

Tim Howes
University of Michigan
Phone: 313-747-4454
Fax: 313-764-5140
tim@umich.edu.

Jonathan Hsu
Penril DataComm Networks, Inc.
Phone: 301-921-8200
brenda@penril.com

Kathleen Huber
Bolt Beranek and Newman
Phone: 617-873-2520
Fax: 617-873-4086
khuber@bbn.com

David Hughes
Old Colorado City
Communications
Phone: 719-636-2040
dave@oldcolo.com

Jeff Hughes
Hewlett-Packard
Phone: 719-531-4777
Fax: 719-531-4505
jeff@col.hp.com

Christian Huitema
INRIA
Phone: +33 93 65 77 15
Fax: +33 93 65 77 66
christian.huitema@sophia.inria.fr

Erik Huizer
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
huizer@surfnet.nl

Craig Hunt
National Institute of
Standards and Technology
Phone: 301-975-3827
Fax: 301-963-9137
chunt@nist.gov

David Husak
Synernetics, Inc.
Phone: 508-670-9009
Fax: 508-670-9015
dave@synnet.com

Robin Iddon
AXON Networks, Inc.
Phone: +44 31 654 0409
Fax: +44 31 654 0409
robini@cix.compulink.co.uk

John Ioannidis
Columbia University
Phone: 212-939-7029
Fax: 212-666-0140
ji@cs.columbia.edu

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
kmj@concord.com

Lenore Jackson
NASA Science Internet
Phone: 415-604-0455
jackson@nsipo.nasa.gov

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

Ronald Jacoby
Silicon Graphics, Inc.
Phone: 415-390-1371
Fax: 415-969-2314
rj@sgi.com

Dwight Jamieson
Bell Northern Research
Phone: 613-763-8761
Fax: 613-763-2626
djamies@bnr.ca

Bob Jeckell
3Com Corporation
Phone: 408-764-6049
Fax: 408-764-5003
rrj@3com.com

Barbara Jennings
Sandia National Laboratories
Phone: 505-845-8554
Fax: 505-844-2067
bjjenni@sandia.gov

George Johnston
Massachusetts Institute of Technology
Phone: 617-253-8652
Fax: 617-253-0807
glj@nerus.pfc.mit.edu

Oliver Jones
PictureTel Corporation
Phone: 508-977-8396
Fax: 508-532-6893
oj@pictel.com

Matthew Jonson
United States Air Force
Phone: 205-416-4075
Fax: 205-416-4738
jonson@server.af.mil

Kevin Jordan
Control Data Systems, Inc.
Phone: 612-482-6835
Fax: 612-482-4455
kej@udev.cdc.com

Dan Jordt
NorthWestNet
Phone: 206-562-3000
Fax: 206-562-4822
danj@nwnet.net

Miriam Kadansky
Xyplex, Inc.
Phone: 508-264-9900
Fax: 508-264-9930
mckadansky@eng.xyplex.com

Merike Kaeo
National Institutes of Health
Phone: 301-402-1548
Fax: 301-402-2867
merike@alw.nih.gov

Marko Kaittola
FUNET
Phone: +358 31 162 933
Fax: +358 31 162 913
marko.kaittola@funet.fi

George Kajos
Coral Network Corporation
Phone: 508-366-3600
Fax: 508-870-1777
kajos@coral.com

Scott Kaplan
FTP Software, Inc.
Phone: 415-543-9001
Fax: 415-543-9002
scott@ftp.com

Phil Karn
Qualcomm, Inc.
Phone: 619-597-5501
Fax: 619-455-0571
karn@qualcomm.com

Daniel Karrenberg
RIPE NCC
Phone: +31 20 592 5102
Fax: +31 20 592 5090
daniel@ripe.net

Frank Kastenholz
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
kasten@ftp.com

Neil Katin
SunSoft
Phone: 415-336-3718
Fax: 415-336-6776
neil.katin@eng.sun.com

David Katinsky
Rutgers University
Phone: 908-932-4783
Fax: 908-932-2968
dmk@rutgers.edu

Akira Kato
Keio University
Phone: +81 466 49 1100
Fax: +81 466 49 1101
kato@wide.sfc.keio.ac.jp

Dave Katz
cisco Systems
Phone: 415-668-8284
Fax: 415-326-1989
dkatz@cisco.com

Laura Kelleher
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3185
lak@merit.edu

Stephen Kent
Bolt Beranek and Newman
Phone: 617-873-3988
Fax: 617-873-4086
kent@bbn.com

Mark Kepke
Hewlett-Packard
Phone: 303-229-4201
Fax: 303-229-3526
mak@cnd.hp.com

Zbigniew Kielczewski
Eicon Technology, Inc.
Phone: 514-631-2592
Fax: 514-631-3092
zbig@eicon.qc.ca

Jong Yeol Kim
Korea Telecom Research Center
Phone: +82 2 526 5077
Fax: +82 2 526 5570
kimjy@ring.kotel.co.kr

John Kimmins
Bell Communications Research
Phone: 908-699-4026
Fax: 908-336-2943
kimm@ctt.bellcore.com

Dwain Kinghorn
Microsoft Corporation
Phone: 206-936-8499
Fax: 206-883-8101
microsoft!dwaink@uunet.uu.net

Keith Klamm
Optical Data Systems, Inc.
Phone: 214-234-6400
Fax: 214-234-4059
klamm@ods.com

John Klensin
INFOODS Secretariat
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

Paulina Knibbe
cisco Systems
Phone: 415-688-7814
Fax: 415-688-7666
knibbe@cisco.com

Mark Knopper
Merit Network, Inc.
Phone: 313-763-6061
Fax: 313-747-3745
mak@merit.edu

Jim Knowles
Sterling Software
Federal Systems
Phone: 415-604-3056
Fax: 415-604-7300
jknowles@binky.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
stev@ftp.com

Andrew Knutsen
Santa Cruz Operation
Phone: 408-427-7538
Fax: 408-458-0811
andrewk@sco.com

Christopher Kolb
Performance Systems
International, Inc.
Phone: 703-620-6651
Fax: 703-620-4586
kolb@psi.com

Michael Kornegay
VisiSoft/Technology Dynamics
Phone: 404-874-0428
Fax: 404-874-6412
mlk@bir.com

Mark Kosters
Network Solutions, Inc.
Phone: 703-802-8484
Fax: 703-802-8376
markk@nic.ddn.mil

Deirdre Kostick
Bell Communications Research
Phone: 908-758-2642
Fax: 908-758-4177
dck2@sabre.bellcore.com

Grisha Kotlyar
Farallon Computing, Inc.
Phone: 510-596-9204
Fax: 510-596-9024
grisha@farallon.com

John Krawczyk
Wellfleet Communications, Inc.
Phone: 617-280-2311
Fax: 617-275-5001
jkrawczy@wellfleet.com

Padma Krishnaswamy
Bell Communications Research
Phone: 908-758-2362
Fax: 908-758-4177
kri@sabre.bellcore.com

Edward Krol
University of Illinois
Urbana
Phone: 217-333-7886
Fax: 217-244-7089
e-krol@uiuc.edu

John Kunze
University of California,
Berkeley
Phone: 510-642-1530
Fax: 510-643-5385
jak@violet.berkeley.edu

Kenrick Kutzler
SynOptics Communications, Inc.
Phone: 408-764-1256
Fax: 408-988-5525
kkutzler@synoptics.com

Mary La Roche
Corporation for Open Systems
International
Phone: 703-205-2741
Fax: 703-846-8590
maryl@cos.com

Paul Lambert
Motorola, Inc.
Phone: 602-441-3646
Fax: 602-441-8377
paul_lambert@email.mot.com

Albert Landberg
National Institute of
Standards and Technology
Phone: 301-975-2245
Fax: 301-590-0932
landberg@mml.ncsl.nist.gov

Sylvain Langlois
Electricite de France
Phone: +33 147654402
Fax: +33 147653523
Sylvain.Langlois@der.edf.fr

Tracy LaQuey Parker
University of Texas at Austin
Phone: 512-471-2488
Fax: 512-471-2449
tracy@utexas.edu

John Larson
Xerox Corporation
Phone: 408-662-9755
Fax: 408-662-9756
jlarson@parc.xerox.com

Mark Laubach
Hewlett-Packard
Phone: 415-857-3513
Fax: 415-813-3152
laubach@hpl.hp.com

Anthony Lauck
Digital Equipment Corporation
Phone: 508-486-7644
Fax: 508-486-5279
lauck@tl.enet.dec.com

Michael Laufer
Bolt Beranek and Newman
Phone: 410-290-5008
Fax: 410-290-5015
mlaufer@bbn.com

Eliot Lear
Silicon Graphics, Inc.
Phone: 415-960-1980
Fax: 415-961-9584
lear@sgi.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

David LeRoy
IBM Corporation
Phone: 919-254-1670
dleroy@mitchell.cit.cornell.edu

Edward Levinson
Accurate Information Systems
Phone: 201-316-0055
Fax: 201-316-0474
levinson@pica.army.mil

Mark Lewis
Telebit Corporation
Phone: 408-745-3232
Fax: 408-745-3810
mlewis@telebit.com

Tony Li
cisco Systems
Phone: 415-688-8186
Fax: 415-688-4575
tli@cisco.com

Fong-Ching Liaw
Sun Microsystems, Inc.
Phone: 415-336-5415
Fax: 415-336-3156
fong@eng.sun.com

David Lin
Zenith Electronics Corporation
Phone: 708-391-8912
Fax: 708-391-8919
lind@janus-ccm.zenith.com

David Lindemulder
NCR
Phone: 908-576-2276
Fax: 908-576-4613
dcl@mtung.att.com

John Linn
Phone: 508-486-7384
Fax: 508-486-5279
linn@erlang.enet.dec.com

Olli-Pekka Lintula
Nokia Telecommunications
Phone: +358 0 506031
Fax: +358 0 5061654
olli-pekka.lintula@ntc.nokia.com

Benjamin Lisowski
Sprint
Phone: 703-689-6531
Fax: 703-689-6144
Ben.Lisowski@sprint.sprint.com

Joshua Littlefield
Cayman Systems, Inc.
Phone: 617-494-1999
Fax: 617-494-5167
josh@cayman.com

Robin Littlefield
Wellfleet Communications, Inc.
Phone: 617-280-2312
Fax: 617-280-2300
rlittlef@wellfleet.com

David Liu
Northern Telecom
Phone: 919-991-8440
Fax: 919-991-7608
dliu@bnr.ca

Daniel Long
Bolt Beranek and Newman
Phone: 617-873-2766
Fax: 617-873-5620
long@nic.near.net

Kim Long
SURAnet
Phone: 302-982-3214
Fax: 302-982-4605
klong@sura.net

Peter Lothberg
Ebone
Phone: +46 8 7906514
roll@stupi.se

E. Paul Love
San Diego Supercomputer Center
Phone: 619-534-5043
Fax: 619-534-5152
loveep@sdsc.edu

Triet Lu
SAIC
Phone: 703-749-5419
Fax: 703-749-5997
triet@cseic.saic.com

John Lunny
The Wollongong Group, Inc.
Phone: 703-847-4554
Fax: 703-847-4520
jlunny@twg.com

Steven Lunt
Bell Communications Research
Phone: 908-699-4244
Fax: 908-336-2943
lunt@bellcore.com

Bryan Lyles
Xerox Corporation
Phone: 415-812-4466
Fax: 415-812-4471
lyles@parc.xerox.com

Bob Lynch
Digital Equipment Corporation
Phone: 603-894-3288
lynch@dsteg.dec.com

Clifford Lynch
University of California
Phone: 510-987-0522
Fax: 510-839-3573
calur@uccmvsa.ucop.edu

Tom Lyon
Sun Microsystems, Inc.
Phone: 415-336-2942
pugs@eng.sun.com

Carl Madison
Star-Tek, Inc.
Phone: 508-393-9393
Fax: 508-393-6934
carl@startek.com

Andrew Maffei
Woods Hole
Oceanographic Institution
Phone: 508-457-2000
Fax: 508-457-2174
arm@aqua.whoi.edu

Dan Magorian
University of Maryland
Phone: 301-405-3004
Fax: 301-314-9220
magorian@ni.umd.edu

Carl Malamud
carl@malamud.com

Kent Malave
IBM Corporation
Phone: 512-838-3486
Fax: 512-838-3509
kent@bach.austin.ibm.com

Andrew Malis
Bolt Beranek and Newman
Phone: 617-873-3419
Fax: 617-873-4086
malis@bbn.com

Gary Malkin
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-3159
gmalkin@xylogics.com

Tracy Mallory
3Com Corporation
Phone: 408-764-6518
Fax: 408-764-5003
tracym@3com.com

Louis Mamakos
University of Maryland
Phone: 301-405-7836
Fax: 301-405-7416
louie@ni.umd.edu

Allison Mankin
The MITRE Corporation
Phone: 703-883-7907
Fax: 703-883-7142
mankin@gateway.mitre.org

Bill Manning
SESQUINET
Phone: 713-285-5415
Fax: 713-527-6099
bmanning@sesqui.net

Janet Marcisak
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
jlm@ftp.com

David Marlow
Naval Surface Warfare Center
Dahlgren Division
Phone: 703-663-1675
dmarlow@relay.nswc.navy.mil

Brian Marsh
MITL
Phone: 609-497-4600
Fax: 609-497-4013
marsh@mitl.com

George Marshall
ADAPTIVE Corporation
Phone: 415-366-9500
Fax: 415-364-0123
george@adaptive.com

Matt Mathis
Pittsburgh Supercomputing
Center
Phone: 412-268-4960
Fax: 412-268-5832
mathis@a.psc.edu

Jun Matsukata
Institute of Space and
Astronautical Science
Phone: +81 427 51 3911
Fax: +81 427 59 4252
jm@eng.isas.ac.jp

Cynthia Matthews
Corporation for National
Research Initiatives
Phone: 703-648-9888
Fax: 703-620-0913
cmatthew@cnri.reston.va.us

Keith McCloghrie
Hughes LAN Systems
Phone: 415-966-7934
Fax: 415-960-3738
kzm@hls.com

Bob McCollum
Network Solutions, Inc.
Phone: 703-802-8476
Fax: 703-802-8376
bobm@nic.ddn.mil

Emily McCoy
The MITRE Corporation
Phone: 703-883-5544
Fax: 703-883-7142
mccoy@gateway.mitre.org

Evan McGinnis
3Com Corporation
Phone: 408-764-6064
Fax: 408-764-6134
bem@3com.com

William McKenzie
IBM Corporation
Phone: 919-254-5705
mckenzie@ralvma.vnet.ibm.com

Donna McMaster
SynOptics Communications, Inc.
Phone: 408-764-1206
Fax: 408-764-1299
mcmaster@synoptics.com

Michael Mealling
Georgia Institute of Technology
Phone: 404-894-4669
Fax: 404-894-6157
michael@fantasy.gatech.edu

John Medicke
IBM Corporation
Phone: 301-240-3472
Fax: 301-240-3228
medicke@ralvm11.vnet.ibm.com

Milo Medin
NASA Ames Research Center
Phone: 415-604-6440
Fax: 415-604-3622
medin@nsipo.nasa.gov

Greg Merrell
Digital Equipment Corporation
Phone: 415-691-4658
Fax: 415-691-4589
merrell@caldec.enet.dec.com

Donald Merritt
Army Research Laboratory
Phone: 410-278-6808
Fax: 410-278-5075
don@brl.mil

David Meyer
University of Oregon
Phone: 503-346-1745
Fax: 503-346-4397
meyer@oregon.uoregon.edu

Douglas Miller
Telebit Corporation
Phone: 408-745-3124
Fax: 408-745-3810
dmm@telebit.com

Cynthia Mills
Bolt Beranek and Newman
Phone: 617-873-4143
Fax: 617-873-5620
cmills@nnsc.nsf.net

Joshua Mindel
Open Networks, Inc.
Phone: 703-648-0013
Fax: 703-648-0016
mindel@netwrx1.nw1.com

David Minnich
FiberCom, Inc.
Phone: 703-342-6700
Fax: 703-342-5961
dwm@fibercom.com

Greg Minshall
Novell, Inc.
Phone: 510-975-4507
Fax: 510-947-1238
minshall@wc.novell.com

Mohammad Mirhakkak
The MITRE Corporation
Phone: 703-883-7820
mmirhakk@mitre.org

Rohit Mital
ProTools, Inc.
Phone: 503-645-5400
Fax: 503-645-3577
rm@protools.com

Mitra
Pandora Systems
Phone: 415-488-0944
mitra@pandora.sf.ca.us

Dave Monachello
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
dave@pluto.dss.com

Doug Montgomery
National Institute of
Standards and Technology
Phone: 301-975-3630
Fax: 301-590-0932
dougm@osi.ncsl.nist.gov

Charlotte Mooers
NSF Network Service Center
Phone: 617-873-3400
Fax: 617-873-5620
mooers@nnsc.nsf.net

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

Dennis Morris
Defense Information Systems
Agency
Phone: 703-487-3329
Fax: 703-487-3190
morrisd@imo-uvax.disa.mil

George Mouradian
AT&T Bell Laboratories
Phone: 908-949-7671
Fax: 908-949-8569
gvm@arch3.att.com

John Moy
Proteon, Inc.
Phone: 508-898-2800
Fax: 508-898-3176
jmoy@proteon.com

Patrick Mullaney
Cabletron Systems, Inc.
Phone: 603-332-9400
Fax: 603-424-8121
mullaney@ctron.com

Jun Murai
Keio University
Phone: +81 466475111
Fax: +81 466491101
jun@wide.ad.jp

Daniel Myers
3Com Corporation
Phone: 408-764-6865
Fax: 408-764-6134
dan@nsd.3com.com

John Myers
Carnegie Mellon University
Phone: 412-268-2314
Fax: 412-268-5249
jgm+@cmu.edu

Julianne Myers
Network Systems Corporation
Phone: 612-424-4888
Fax: 612-424-2853
jmyers@network.com

Sheryl Namoglu
Soft-Switch, Inc.
Phone: 215-640-7506
Fax: 215-640-7550
sfn@softsw.ssw.com

Rina Nathaniel
RAD Network Devices
Phone: +972 3 6458556
Fax: +972 3 5447368
rina!rnd!rndi@uunet.uu.net

Mark Needleman
University of California
Phone: 510-987-0530
Fax: 510-763-9673
mhn@stubbs.ucop.edu

Clifford Neuman
Information Sciences Institute
Phone: 310-822-1511
bcn@isi.edu

Michael Newell
Sterling Software
Phone: 202-434-8954
mnewell@nsipo.nasa.gov

Chris Newman
Carnegie Mellon University
Phone: 412-268-6403
Fax: 412-268-4987
chrisn+@cmu.edu

Scott Newman
Digital Equipment Corporation
Phone: 603-884-6123
Fax: 603-884-0829
newman@broke.enet.dec.com

Hien Nguyen
Sprint International
Phone: 703-689-6908
h.nguyen@sprintintl@sprint.com

Joe Nguyen
SynOptics Communications, Inc.
Phone: 408-764-1613
Fax: 408-988-8548
jnguyen@synoptics.com

Tu Nguyen
Defense Information Systems Agency
Phone: 703-487-3037
Fax: 703-487-3190
Nguyen1T@cc.ims.disa.mil

Kathleen Nichols
Apple Computer, Inc.
Phone: 408-974-1136
Fax: 408-974-9793
nichols@apple.com

Andy Nicholson
Cray Research, Inc.
Phone: 612-683-5473
Fax: 612-683-5599
droid@cray.com

Mo Nikain
BellSouth
Phone: 404-332-2224
Fax: 404-880-9011
mo@bss.com

Tom Nisbet
Telecommunications Techniques Corporation
Phone: 301-353-1550
Fax: 301-353-1536
nisbet@tt.com

Erik Nordmark
SunSoft
Phone: 415-962-8441
Fax: 415-336-6015
nordmark@eng.sun.com

Peder Norgaard
PTT Telecom, The Netherlands
Phone: +45 86 28 81 76
Fax: +45 86 28 81 86
pcn@tbit.dk

Bill Norton
Merit Network, Inc.
Phone: 313-936-2656
Fax: 313-747-3745
wbn@merit.edu

William Nowicki
Legato Systems, Inc.
Phone: 415-329-7856
Fax: 415-329-8898
nowicki@legato.com

Michael O'Brien
Sprint
Phone: 913-624-6955
Fax: 913-624-5670
m.obrien@sprintmail.com

David O'Leary
cisco Systems
Phone: 714-261-5681
Fax: 714-261-5684
doleary@cisco.com

Sean O'Malley
University of Arizona
Phone: 602-621-6613
Fax: 602-621-4246
sean@cs.utexas.edu

Kate O'Mara
National Computer Systems
Phone: 703-276-1440
kate@acfcluster.nyu.edu

Frank Odasz
Western Montana College
Phone: 406-683-7338
Fax: 406-683-7493
franko@bigsky.dillon.mt.us

Tatsuya Ohnishi
Matsushita Electronic
Industrial Company, Ltd.
Phone: +81 6 906 2431
Fax: +81 6 906 5547
ohnishi@isl.mei.co.jp

Ari Ollikainen
Energy Sciences Network
Phone: 510-423-5962
Fax: 510-423-8744
ari@es.net

Steven Onishi
Wellfleet Communications, Inc.
Phone: 617-280-2316
Fax: 617-275-5001
sonishi@wellfleet.com

Alan Oppenheimer
Apple Computer, Inc.
Phone: 408-974-4744
Fax: 408-862-1159
oppenheime1@applelink.apple.com

David Orelowitz
Microcom, Inc.
Phone: 914-377-2703
Fax: 914-968-7100
david@c1st.microcom.com

Hilarie Orman
University of Arizona
Phone: 602-621-4891
Fax: 602-621-4246
ho@cs.arizona.edu

Joe Pagan
Department of Defense
Phone: 301-688-7504
jrp@afterlife.ncsc.mil

Luc Pariseau
Digital Equipment Corporation
Phone: 508-952-3563
Fax: 508-952-3023
pariseau@quiver.enet.dec.com

Brad Parker
FCR, Inc.
Phone: 617-424-6552
Fax: 617-424-6552
brad@fcr.com

Andrew Partan
UUNET Technologies, Inc.
Phone: 703-204-8000
Fax: 703-204-8001
asp@uunet.uu.net

Craig Partridge
Bolt Beranek and Newman
Phone: 415-326-4541
craig@bbn.com

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
patel@noc.rutgers.edu

Michael Patton
Bolt Beranek and Newman
map@bbn.com

Gaige Paulsen
InterCon Ststems Corporation
Phone: 703-709-5509
Fax: 703-709-5555
gaige@intercon.com

Jim Perchik
IBM High Performance Computing and Communications
Phone: 617-253-1967
Fax: 617-258-8831
perchik@athena.mit.edu

Pete Percival
Indiana University
Phone: 812-855-9146
Fax: 812-855-8299
percival@indiana.edu

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
Perkins Software
Phone: 412-325-1785
Fax: 412-325-1344
perkins+@cmu.edu

Marsha Perrott
PREPnet
Phone: 412-268-7870
Fax: 412-268-7875
mlp+@andrew.cmu.edu

Roy Perry
US West Advanced Technologies
Phone: 303-541-6245
Fax: 303-541-6276
rperry@advtech.uswest.com

Karen Petraska-Veum
Hughes/STX
Phone: 301-286-3348
Fax: 301-286-5152
karen@nsisrv.gsfc.nasa.gov

Mike Petry
University of Maryland
Phone: 301-405-2994
Fax: 301-405-7566
petry@ni.umd.edu

David Piscitello
Bell Communications Research
Phone: 908-758-2286
Fax: 908-758-4177
dave@eve.bellcore.com

Mel Pleasant
Rutgers University
Phone: 908-932-2023
Fax: 908-932-2968
pleasant@hardees.rutgers.edu

Jon Postel
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
postel@isi.edu

Carl Powell
Bolt Beranek and Newman
Phone: 410-290-5010
Fax: 410-290-5015
cpowell@bbn.com

Thomas Pusateri
Thomas & Willow Consulting
Phone: 919-676-9941
pusateri@cs.duke.edu

James Quigley
Hewlett-Packard
Phone: 408-447-3505
Fax: 408-447-3660
jim_quigley%YO@hp6600.desk.hp.com

Ilan Raab
SynOptics Communications, Inc.
Phone: 408-764-1542
Fax: 408-988-5525
iraab@synoptics.com

Joe Ragland
MCNC - CONCERT Network
Phone: 919-248-1404
Fax: 919-248-1405
jrr@concert.net

Bala Rajagopalan
AT&T Bell Laboratories
Phone: 908-949-8017
Fax: 908-949-1726
braja@qsun.att.com

Richard Ramos
AT&T Bell Laboratories
Phone: 908-576-6267
Fax: 908-576-4613
ramos@mtunm.att.com

Joseph Ramus
Lawrence Livermore National Laboratory
Phone: 510-423-8917
Fax: 510-423-8744
ramus@nersc.gov

Christopher Ranch
Novell, Inc.
Phone: 408-473-8667
Fax: 408-946-9761
cranch@novell.com

Venkat Rangan
Metrix Network Systems, Inc.
Phone: 603-888-7000
Fax: 603-891-2796
venkat@geoduck.matrix.com

Greg Ratta
Bell Communications Research
Phone: 908-758-5696
Fax: 908-758-4437
gregra@cc.bellcore.com

Sri Reddy
SynOptics Communications, Inc.
Phone: 408-764-1622
Fax: 408-988-6210
srireddy@synoptics.com

Ed Reeder
IBM Corporation
Phone: 919-543-2154
Fax: 919-543-0723
ereeder@vnet.ibm.com

Jim Rees
University of Michigan
Phone: 313-763-4174
Fax: 313-763-4434
jim.rees@umich.edu

James Reeves
SynOptics Communications, Inc.
Phone: 408-764-1440
Fax: 408-988-5525
jreeves@synoptics.com

Louise Reingold
Sprint
Phone: 703-689-5940
Fax: 703-689-6998
l.reingold@sprint.sprint.com

Yakov Rekhter
IBM Corporation
Phone: 914-945-3896
yakov@watson.ibm.com

Sheri Repucci
Merit Network, Inc.
Phone: 313-936-3335
smr@merit.edu

Joyce K. Reynolds
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
jkrey@isi.edu

Bradley Rhoades
Metal Matrix Composites
Phone: 612-736-2874
Fax: 612-736-0431
bdrhoades@mail.mmmg.com

Tony Richards
Sprint Government Systems Division
Phone: 202-416-6208
Fax: 202-833-8621

April Richstein
Department of Defense
Phone: 301-688-0849
abm@tycho.ncsc.mil

Mike Ritter
Apple Computer, Inc.
Phone: 408-862-8088
Fax: 408-862-1159
mwritter@applelink.apple.com

Michael Roberts
EDUCOM
Phone: 703-648-9888
Fax: 703-620-0913
mmr@isoc.org

Sam Roberts
Farallon Computing, Inc.
Phone: 510-814-5215
sroberts@farallon.com

Kary Robertson
Concord Communications, Inc.
Phone: 508-460-4646
Fax: 508-481-9772
kr@concord.com

Richard Rodgers
National Library of Medicine
Phone: 301-496-9300
rodgers@nlm.nih.gov

Benny Rodrig
RAD Network Devices
Phone: +972 3 6458552
Fax: +972 3 5447368
4373580@mcimail.com

Manoel Rodrigues
AT&T Bell Laboratories
Phone: 908-949-4655
Fax: 908-949-1726
manoel_rodrigues@att.com

Guenter Roeck
Conware Computer Consulting GmbH
Phone: +49 721 9495 0
Fax: +49 721 9495 146
roeck@conware.de

Jim Romaguera
NetConsult AG
Phone: +41 31 235750
Fax: +41 31 235792
romaguera@cosine-mhs.switch.ch

Dan Romascanu
Lannet Data Communications
Phone: +972 3 6458414
Fax: +972 3 5447146
dan@lannet.com

Marshall Rose
Dover Beach Consulting, Inc.
Phone: 415-968-1052
Fax: 415-968-2510
mrose@dbc.mtview.ca.us

Avraham Rosenbach
Armon Networking, Ltd.
Phone: +972 3 490702
Fax: +972 3 49701
armon@armon.hellnet.org

Alan Roszkiewicz
Sprint International
Phone: 703-689-5319
Fax: 703-689-5644
alan@sprint.com

Shawn Routhier
Epilogue Technology Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sar@epilogue.com

Chris Rozman
US Robotics, Inc.
Phone: 708-982-5025
Fax: 708-982-1348
chrisr@usr.com

Allan Rubens
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
acr@merit.edu

Assaf Rubissa
Fibronics, Ltd.
Phone: +972 4 313604
Fax: +972 4 550550
asaf@fibhaifa.com

Gregory Ruth
Bolt Beranek and Newman
Phone: 617-873-3150
Fax: 617-873-8202
gruth@bbn.com

Anthony Rutkowski
Sprint International
Phone: 703-689-5080
Fax: 703-689-6715
amr@sprint.com

Michael Safly
Boeing
Phone: 205-544-6445
Fax: 205-544-4618
saf@tank1.msfc.nasa.gov

Timothy Salo
Minnesota Supercomputer Center
Phone: 612-626-0347
Fax: 612-624-6550
tjs@msc.edu

Fred Sammartino
Sun Microsystems, Inc.
Phone: 415-336-2938
Fax: 415-336-3156
fred@eng.sun.com

Henry Sanders
Microsoft Corporation
Phone: 206-882-8080
henrysa@microsoft.com

Tom Sandoski
MCNC - CONCERT Network
Phone: 919-248-1104
Fax: 919-248-1405
tom@concert.net

Paul Sangster
Advanced Network and Services, Inc.
Phone: 703-758-7706
Fax: 703-758-7717
sangster@ans.net

Jon Saperia
Digital Equipment Corporation
Phone: 508-496-8333
Fax: 508-496-9929
saperia@lkg.dec.com

Michael Sapich
Conware Computer Consulting GmbH
Phone: +49 721 9495 0
Fax: +49 721 9495 146
sapich@conware.de

Ashok Saraf
Fujitsu America, Inc.
Phone: 214-997-7773
Fax: 214-497-6991

Srinivas Sataluri
AT&T Bell Laboratories
Phone: 908-949-7782
Fax: 908-949-1726
sri@qsun.att.com

Michael Scanlon
Racal-Datacom, Inc.
Phone: 508-263-9929
Fax: 508-263-8655
scanlon@interlan.com

Sam Schaen
The MITRE Corporation
Phone: 703-883-7297
Fax: 703-883-1397
schaen@mitre.org

Paul Scheffler
Teleos Communications, Inc.
Phone: 908-544-6442
Fax: 908-544-9890
pes@teleoscom.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

Eve Schooler
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
schooler@isi.edu

Martin Schulman
Sprint International
Phone: 703-689-7171
Fax: 703-689-5176
mas@loyola.edu

Henning Schulzrinne
AT&T Bell Laboratories
Phone: 908-582-2262
Fax: 908-582-5809
hgs@research.att.com

Dallas Scott
The MITRE Corporation
Phone: 703-883-6700
Fax: 703-883-1364
scott@fluky.mitre.org

Jim Scott
Digital Equipment Corporation
Phone: 508-493-5514
scott@kali.enet.dec.com

John Scudder
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
jgs@merit.edu

Tim Seaver
MCNC - CONCERT Network
Phone: 919-248-1973
Fax: 919-248-1105
tas@concert.net

John Seligson
Ultra Network Technologies
Phone: 408-922-0100
Fax: 408-433-9287
johns@ultra.com

Jennifer Sellers
NASA NREN
Phone: 202-434-8954
Fax: 202-434-4599
sellers@nsipo.arc.nasa.gov

Paul Serice
Corporation for Open Systems International
Phone: 703-205-2713
Fax: 703-846-8590
serice@cos.com

Chris Shaw
Banyan Systems, Inc.
Phone: 508-836-2898
Fax: 508-898-3604
cshaw@banyan.com

John Sherburne
Sprint International
Phone: 703-689-6297
Fax: 703-689-6066
john.sherburne@sprintintl.sprint.com

Erik Sherk
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
sherk@sura.net

Brian Shiflett
Sprint Government Systems Division
Phone: 703-904-2074
Fax: 703-904-2251
bshiflett@icm1.icp.net

Robert Shirey
The MITRE Corporation
Phone: 703-883-7210
Fax: 703-883-1397
shirey@mitre.org

Allyson Showalter
NASA Science Internet
Phone: 202-434-8954
Fax: 202-434-4599
allyson@nsipo.arc.nasa.gov

Cris Shuldiner
FTP Software, Inc.
Phone: 508-685-3600
Fax: 508-794-4484
cws@ftp.com

William Simpson
Computer Systems Consulting Services
Phone: 313-623-7979
Bill.Simpson@um.cc.umich.edu

Anil Singhal
Frontier Software Development
Phone: 508-851-8872
Fax: 508-851-6956
anil@frontier.com

Sam Sjogren
TGV, Inc.
Phone: 408-427-4366
Fax: 408-427-4365
sjogren@tgv.com

Dean Skidmore
IBM Corporation
Phone: 919-254-4806
Fax: 919-254-6430
deans@ralvmiz.vnet.ibm.com

Keith Sklower
University of California, Berkeley
Phone: 510-642-9587
sklower@cs.berkeley.edu

Robert Slaski
Open Networks, Inc.
Phone: 703-648-0013
Fax: 703-648-0016
slaski@netwrx1.nw1.com

Lansing Sloan
Lawrence Livermore National Laboratory
Phone: 510-422-4356
Fax: 510-423-8534
ljsloan@llnl.gov

Timon Sloane
Timon Sloane Consulting Services
Phone: 415-515-0741
Fax: 415-941-4679
timon@rahul.net

Jane Smith
MCNC - CNIDR
Phone: 919-248-9213
Fax: 919-248-1405
jds@jazz.concert.net

Mark Smith
University of Michigan
Phone: 313-764-2277
Fax: 313-764-5140
mcs@umich.edu

Patricia Smith
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
psmith@merit.edu

Larry Snodgrass
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
snodgras@bitnic.educom.edu

Robert Snyder
cisco Systems
Phone: 415-688-8245
Fax: 415-688-8282
snyder@cisco.com

Frank Solensky
Ungermann-Bass, Inc.
Phone: 508-691-4604
Fax: 508-687-4442
solensky@andr.ub.com

Joo Young Song
Korea Telecom Research Center
Phone: +82 2 526-5077
Fax: +82 2 526-5570
jysong@ring.kotel.co.kr

Marco Sosa
Bell Communications Research
Phone: 908-758-2651
Fax: 908-758-4177
mxs@sabre.bellcore.com

Simon Spero
University of North Carolina at Chapel Hill
Phone: 919-962-9107
Fax: 919-962-5604
simon_spero@unc.edu

Roy Spitzer
Sprint International
Phone: 703-689-6899
Fax: 703-689-5176
roy.spitzer@sprintintl.sprint.com

Michael St. Johns
DARPA/CSTO
Phone: 703-696-2271
Fax: 703-696-2202
stjohns@darpa.mil

Martha Steenstrup
Bolt Beranek and Newman
Phone: 617-873-3192
Fax: 617-873-3776
msteenst@bbn.com

Einar Stefferud
Network Management
Associates, Inc.
Phone: 714-842-3711
Fax: 714-848-2091
stef@nma.com

Scott Stein
Apple Computer, Inc.
Phone: 408-974-3124
Fax: 408-974-0781
scotts@apple.com

Brad Steinka
Microcom, Inc.
Phone: 617-551-1209
Fax: 617-551-1015
brad@microcom.com

John Stephens
Cayman Systems, Inc.
Phone: 617-494-1999
Fax: 617-494-5167
john@cayman.com

Bob Stewart
Xyplex, Inc.
Phone: 508-264-9900
Fax: 508-264-9930
rlstewart@eng.xyplex.com

Bernhard Stockman
SUNET
Phone: +46 8 7906519
Fax: +46 8 241179
boss@ebone.net

Terrance Sullivan
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
terrys@newbridge.com

Craig Summerhill
Coalition for
Networked Information
Phone: 202-232-2466
Fax: 202-462-7849
craig@cni.org

Catherine Summers
Corporation for Open Systems
Phone: 703-205-2801
Fax: 703-846-8590
cts@cos.com

Richard Sweatt
SynOptics Communications, Inc.
Phone: 408-764-1260
Fax: 408-764-1299
rsweatt@synoptics.com

Tang Tang
University of Virginia
Phone: 804-924-0632
Fax: 804-982-4715
tt@virginia.edu

Morton Taragin
The Weizmann Institute of
Science
vsmorty@weizmann.weizmann.ac.il

Sally Tarquinio
The MITRE Corporation
Phone: 703-883-7987
Fax: 703-883-7142
sallyt@gateway.mitre.org

John Tavs
IBM Corporation
Phone: 919-254-7610
Fax: 919-254-0151
tavs@vnet.ibm.com

Fumio Teraoka
Sony Computer Science
Laboratory, Inc.
Phone: +81 3 34484380
Fax: +81 3 34484273
tera@csl.sony.co.jp

Marten Terpstra
RIPE NCC
Phone: +31 20 592 5102
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

Richard Thomas
Bell Northern Research
Phone: 613-763-7025
Fax: 613-765-4920
rjthomas@bnr.ca

Geoffrey Thompson
SynOptics Communications, Inc.
Phone: 408-764-1339
Fax: 408-988-5525
thompson@synoptics.com

Jim Thompson
Tadpole Technology, Inc.
Phone: 512-219-2200
Fax: 512-219-2244
jim@tadpole.com

Dean Throop
Data General Corporation
Phone: 919-549-8421
Fax: 919-248-5942
throop@dg-rtp.dg.com

Claudio Topolcic
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
topolcic@cnri.reston.va.us

Paul Traina
cisco Systems
Phone: 415-688-8189
Fax: 415-688-4575
pst@cisco.com

Theodore Ts'o
Massachusetts Institute of Technology
Phone: 617-253-8091
Fax: 617-258-8736
tytso@mit.edu

Panos-Gavriil Tsigaridas
GMD-FOKUS, Germany
Phone: +49 30 25499232
Fax: +49 30 25499202
Tsigaridas@fokus.berlin.gmd.dbp.de

Paul Tsuchiya
Bell Communications Research
Phone: 201-829-4484
Fax: 201-829-2504
tsuchiya@thumper.bellcore.com

Ahmet Tuncay
SynOptics Communications, Inc.
Phone: 408-764-1279
Fax: 408-764-1299
atuncay@synoptics.com

Thierry Turletti
INRIA
Phone: +33 93 65 77 18
Fax: +33 93 65 77 66
turletti@sophia.inria.fr

Dono van-Mierop
3Com Corporation
Phone: 408-764-5930
Fax: 408-764-5722
dono_van_mierop@3mail.3com.com

Stuart Vance
TGV, Inc.
Phone: 408-427-4366
Fax: 408-427-4365
vance@tgv.com

Elizabeth Vanderbeck
IBM Corporation
Phone: 914-385-0385
beth@tdcsys2.vnet.ibm.com

Kannan Varadhan
OARnet
Phone: 614-292-4137
Fax: 614-292-7168
kannan@oar.net

Gregory Vaudreuil
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
gvaudre@cnri.reston.va.us

John Veizades
Apple Computer, Inc.
Phone: 408-974-2672
Fax: 408-862-1159
veizades@apple.com

Warren Vik
INTERACTIVE Systems
Corporation
Phone: 708-505-9100
Fax: 708-505-9133
wmv@i88.isc.com

Curtis Villamizar
Advanced Network and
Services, Inc.
Phone: 914-789-5300
Fax: 914-789-5310
curtis@ans.net

Ioannis Viniotis
North Carolina State
University
Phone: 919-515-5148
Fax: 919-515-5523
candice@ececho.ncsu.edu

John Vollbrecht
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
jrv@merit.edu

Janet Vratny
Apple Library
Phone: 408-974-2048
Fax: 408-725-8502
janet@apple.com

Iain Wacey
Datability, Inc.
Phone: 201-438-2400
Fax: 201-438-7767
cat@pluto.dss.com

A. Lee Wade
NASA Science Internet
Phone: 415-604-4789
Fax: 415-604-0063
wade@nsipo.nasa.gov

Steven Waldbusser
Carnegie Mellon University
Phone: 412-268-6628
Fax: 412-268-4987
waldbusser@andrew.cmu.edu

Alice Wang
Sun Microsystems, Inc.
Phone: 415-336-5394
Fax: 415-336-3156
alice.wang@eng.sun.com

Eva Wang
Penril DataComm Networks, Inc.
Phone: 301-921-8200
eva@penril.com

Zheng Wang
University of College London
Phone: +44 71 787 7050
Fax: +44 71 387 1397
z.wang@cs.ucl.ac.uk

Chuck Warlick
Boeing
Phone: 205-544-6270
Fax: 205-544-4618
warlick@theophilis.nsfc.nasa.gov

William Warner
Department of Administrative Services, State of Ohio
Phone: 614-466-6683
Fax: 614-752-6108
warner@ohio.gov

Jack Waters
SURAnet
Phone: 301-982-4600
Fax: 301-982-4605
waters@sura.net

James Watt
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
james@newbridge.com

Luanne Waul
Ascom Timeplex, Inc.
Phone: 310-443-5651
Fax: 310-443-5582
luanne@wwtc.timeplex.com

Chris Weider
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3185
clw@merit.edu

Terry Weigler
Corporation for National Research Initiatives
Phone: 703-620-8990
Fax: 703-620-8990
tweigler@cnri.reston.va.us

Abel Weinrib
Bell Communications Research
Phone: 201-829-4454
Fax: 201-829-5889
abel@bellcore.com

Von Welch
National Center for Supercomputing Applications
Phone: 217-244-3095
Fax: 217-244-1987
vwelch@ncsa.uiuc.edu

Guy Wells
US West Advanced Technologies
Phone: 303-541-6759
Fax: 303-541-6276
guy2@uswest.com

Jonathan Wenocur
Shiva Corporation
Phone: 617-252-6347
Fax: 617-252-6852
jhw@shiva.com

Moira West
CERT Coordination Center
Phone: 412-268-7090
Fax: 412-268-6989
mjw@cert.org

Evan Wetstone
SESQUINET
Phone: 713-285-5173
Fax: 713-527-6099
evan@rice.edu

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: 206-543-3909
cwheeler@cac.washington.edu

Gerry White
Applitek Corporation
Phone: 508-475-4050
Fax: 508-475-0550
gerry@lancity.com

Brenda Whitehurst
Penril DataComm Networks, Inc.
Phone: 301-921-8200
brenda@penril.com

Gary Wightman
Sprint
Phone: 816-854-2367
Fax: 816-854-2301

Peter Will
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
will@isi.edu

Douglas Williams
IBM Networking Systems
Phone: 919-254-4801
dougw@ralvmg.vnet.ibm.com

Kirk Williams
Southwestern Bell Technology Resources
Phone: 314-529-7586
Fax: 314-529-7674
kirk@sbctri.sbc.com

Scott Williamson
Network Solutions, Inc.
Phone: 703-802-8405
Fax: 703-802-8376
scottw@nic.ddn.mil

Daniel Wilson
Bell Communications Research
Phone: 201-829-4375
Fax: 201-829-1910
dvw@bellcore.com

Ian Wilson
Spider Systems, Ltd.
Phone: +44 31 554 9424
Fax: +44 31 554 0649
ianw@spider.co.uk

Peter Wilson
3Com Limited
Phone: +44 442 231000
Fax: +44 442 254701
peter_wilson@3com.com

Walter Wimer
Carnegie Mellon University
Phone: 412-268-6252
Fax: 412-268-4987
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
Phone: 708-252-7236
Fax: 708-252-5983
lwinkler@anl.gov

Cathy Wittbrodt
Lawrence Livermore National Laboratory
Phone: 510-422-4016
Fax: 510-422-0435
cjw@nersc.gov

Stephen Wolff
National Science Foundation
Phone: 202-357-9717
Fax: 202-357-0320
steve@nsf.gov

Steven Wong
Digital Equipment Corporation
Phone: 508-486-5344
Fax: 508-486-7417
wong@took.enet.dec.com

Robert Woodburn
Sparta, Inc.
Phone: 703-448-0210
Fax: 703-734-9000
woody@sparta.com

Randall Worzella
IBM Corporation
Phone: 919-254-1728
Fax: 919-254-6963
worzella@ralvm29.unet.ibm.com

Daniel Woycke
The MITRE Corporation
Phone: 703-883-1362
Fax: 703-883-1397
woycke@smiley.mitre.org

Russ Wright
Lawrence Berkeley Laboratory
Phone: 510-486-6965
Fax: 510-486-6363
wright@lbl.gov

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
ltw99@bellcore.com

Jeff Yarnell
ProTools, Inc.
Phone: 503-645-5400
Fax: 503-645-3577
jeffya@protools.com

Peter Yee
NASA Ames Research Center
Phone: 415-604-3812
Fax: 415-604-6999
yee@atlas.arc.nasa.gov

Jeff Young
National Institutes of Health
Phone: 301-496-2251
Fax: 301-402-2867
young@alw.nih.gov

Susan Young
Sprint Government Systems Division
Phone: 202-416-6250
Fax: 202-833-8621

Yung-Chao Yu
AT&T Bell Laboratories
Phone: 908-949-0555
Fax: 908-949-4673
yy@qsun.att.com

Chin Yuan
Pacific Bell
Phone: 510-867-5016
Fax: 510-867-1405
cxyuan@pacbell.com

Kiho Yum
3Com Corporation
Phone: 408-764-6400
Fax: 408-754-5003
kxy@nsd.3com.com

William Yurcik
NASA Science Internet
Phone: 301-286-1376
Fax: 301-286-7251
yurcik@dftnic.gsfc.nasa.gov

John Zalubski
Networking Information Center
Phone: 703-802-8462
Fax: 703-802-8376
zalubski@nic.ddn.mil

Paul Zawada
National Center for Supercomputing Applications
Phone: 217-244-4728
Fax: 217-244-7396
Zawada@ncsa.uiuc.edu

James Zmuda
Hughes Aircraft Company
Phone: 714-732-7240
Fax: 714-732-2613
zmuda@mls1.hac.com

Miriam Zohar
IBM Corporation
Phone: 914-945-2567
zohar@watson.ibm.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