ACKNOWLEDGEMENTS

The Twenty-Third Internet Engineering Task Force was held in San Diego, California during the week of March 16th and boasted an unprecedented 530 attendees. Access to the Internet was provided through the Local Host, San Diego Supercomputer Center, and a number of local vendors. E. Paul Love, Jr. arranged for the setup of the Terminal Room and was assisted in that endeavor by Hans-Werner Braun, Bilal Chinoy, Don Doering, Jay Dombrowhski, Kevin Fall, Richard Gallup, and several others. We wish to thank Paul and the others for the long hours and hard work which resulted in a terrific Terminal Room. Recognition should also be extended to the following organizations for their generous contributions:

Pacific Bell	T1 link
Helfrich Co. and Verilink	CSU/DSU pair for link
CERFnet	cisco Routers; Internet Link
Digital Equipment	Workstation plus X-Terms
Sun	Workstations
Hewlett-Packard	Workstations plus X-Terms

We would be remiss if we did not also acknowledge those individuals who participated in the various Technical Presentations throughout the week. The quality of their presentations significantly contributed to the overall success of the meeting. They are, Ashok Agrawala, Jordan Becker, Dave Borman, Bob Braden, Scott Bradner, Steve Casner, Peter Danzig, Steve Deering, Ralph Droms, Peter Ford, Phill Gross, Steve Hardcastle-Kille, Mark Knopper, Patrice Lyons, Mark McCahill, Andy Malis, Greg Vaudreuil, Scott Williamson and Steve Wolff.

A final thanks goes to Debra Legare and Cynthia Clark not only for their capable handling of the registration process, but also for their invaluable contributions to these Proceedings.

Megan Davies/CNRI

Contents

Chairman's Message 1

Agenda of the Twenty-Third IETF 7

1 IETF Overview 15
- 1.1 Future IETF Meeting Sites . . . 19
- 1.2 On Line IETF Information . . . 21
  - 1.2.1 The IETF Directory . . . 22
  - 1.2.2 The Internet-Drafts Directory . . . 23
- 1.3 Guidelines to Authors of Internet Drafts . . . 25

2 Steering Group Report 27
- 2.1 Standards Progress Report . . . 29
- 2.2 Minutes of the Open Plenary and IESG . . . 31
  - 2.2.1 IETF Growth . . . 31
  - 2.2.2 Routing and Addressing . . . 31

3 Area and Working Group Reports 35
- 3.1 Applications Area . . . 37
  - 3.1.1 Automated Internet Mailing List Services (list) . . . 53
  - 3.1.2 Distributed Scheduling Protocol (chronos) . . . 55
  - 3.1.3 Internet Mail Extensions (smtpext) . . . 57
  - 3.1.4 Internet Message Extensions (822ext) . . . 61
  - 3.1.5 Network Database (netdata) . . . 63
  - 3.1.6 Network Fax (netfax) . . . 65
  - 3.1.7 Network News Transport Protocol (nntp) . . . 69
  - 3.1.8 Network Printing Protocol (npp) . . . 73
  - 3.1.9 TELNET (telnet) . . . 75
- 3.2 Internet Area . . . 79
  - 3.2.1 Connection IP (cip) . . . 85
  - 3.2.2 Dynamic Host Configuration (dhc) . . . 87
  - 3.2.3 IP over AppleTalk (appleip) . . . 91
  - 3.2.4 IP over Asynchronous Transfer Mode (atm) . . . 101
  - 3.2.5 IP over FDDI (fddi) . . . 107
  - 3.2.6 Multi-Media Bridging (mmb) . . . 109
  - 3.2.7 Point-to-Point Protocol Extensions (pppext) . . . 111

3.2.8 Router Requirements (rreq) 119
3.3 Network Management Area 121
3.3.1 Chassis MIB (chassis) 147
3.3.2 DS1/DS3 MIB (trunkmib) 155
3.3.3 Bridge MIB (bridge) 159
3.3.4 Character MIB (charmib) 161
3.3.5 DECnet Phase IV MIB (decnetiv) 163
3.3.6 Ethernet MIB (ethermib) 165
3.3.7 IEEE 802.3 Hub MIB (hubmib) 171
3.3.8 Internet Accounting (acct) 179
3.3.9 Remote LAN Monitoring (rmonmib) 185
3.3.10 Token Ring Remote Monitoring (trmon) 191
3.3.11 X.25 Management Information Base (x25mib) 193
3.4 OSI Integration Area 197
3.4.1 MHS-DS (mhsds) 209
3.4.2 Network OSI Operations (noop) 215
3.4.3 OSI Directory Services (osids) 225
3.4.4 OSI General (osigen) 243
3.4.5 Office Document Architecture (oda) 245
3.4.6 SNMP over a Multi-protocol Internet (mpsnmp) 249
3.4.7 X.400 Operations (x400ops) 257
3.5 Operational Requirements Area 265
3.5.1 BGP Deployment and Application (bgpdepl) 277
3.5.2 Benchmarking Methodology (bmwg) 287
3.5.3 Network Joint Management (njm) 291
3.5.4 Operational Statistics (opstat) 293
3.5.5 User Connectivity (ucp) 299
3.6 Routing Area 303
3.6.1 Border Gateway Protocol (bgp) 321
3.6.2 IP over Large Public Data Networks (iplpdn) 335
3.6.3 ISIS for IP Internets (isis) 341
3.6.4 Inter-Domain Policy Routing (idpr) 345
3.6.5 Multicast Extensions to OSPF (mospf) 347
3.6.6 Open Shortest Path First IGP (ospf) 353
3.6.7 RIP Version II (ripv2) 363
3.7 Security Area 367
3.7.1 Commercial Internet Protocol Security Option (cipso) 371
3.7.2 Common Authentication Technology (cat) 373
3.7.3 Privacy-Enhanced Electronic Mail (pem) 383
3.7.4 SNMP Security (snmpsec) 387
3.8 Transport and Services Area 389
3.8.1 Audio/Video Transport (avt) 391

3.8.2 Distributed File Systems (dfs) 399
3.8.3 Domain Name System (dns) 401
3.8.4 Service Location Protocol (svrloc) 405
3.8.5 TCP Large Windows (tcplw) 409
3.8.6 Trusted Network File Systems (tnfs) 411
3.9 User Services Area 417
3.9.1 Directory Information Services Infrastructure (disi) 425
3.9.2 Internet Anonymous FTP Archives (iafa) 429
3.9.3 Internet School Networking (isn) 433
3.9.4 Internet User Glossary (userglos) 439
3.9.5 NOC-Tool Catalogue Revisions (noctool2) 445
3.9.6 Network Information Services Infrastructure (nisi) 449
3.9.7 User Services (uswg) 455

4 Network Status Briefings **459**
4.1 CERFnet 461
4.2 EBONE Report 463
4.3 ESnet Report 465
4.4 NSI Report 469
4.5 PACnet 471
4.6 RIPE Report 473

5 IETF Protocol Presentations **479**
5.1 Dynamic Host Configuration Protocol 481
5.2 TCP Large Windows 487
5.3 The Internet Gopher Protocol 495
5.4 IP Over X.25 503

6 Technical Presentations **505**
6.1 ROAD Presentation 507
6.2 Copyright Claims and Standards 517
6.3 First IETF Internet Audiocast 523
6.4 Traffic Characterization in Wide Area Networks 525
6.5 Multi-Media Mail Extensions 531
6.6 NSFnet Update 541
6.7 NEARnet Presentation 547
6.8 NIC Services 555

7 Distributed Papers **557**
7.1 Flow and Congestion Control in Networks 559

A Attendees **565**

Chair's Message

Our local hosts, E. Paul Love, Jr., and the San Diego Supercomputer Center (SDSC), did an outstanding job of supporting our meeting. Paul was our host for the Ninth IETF meeting in March 1988, and it was a special pleasure to return to San Diego for the Twenty-Third meeting.

The IETF has evolved beyond the point of having a simple "terminal room" - we now have a complete computer center and Local Area Network. The SDSC "IETF computer center" had over 25 workstations, X-window terminals, and conventional terminals. They provided multiple printers, a network monitor, numerous laptop ports, and even an appletalk network.

I had personal experience using the facility when Peter Ford and I put together the ROAD presentation, and later when I created foils for the Open IESG session. So I can say with conviction how grateful I am to Paul for all the work the SDSC folks did for the IETF.

The list of the hard working folks to whom we owe thanks is listed in the acknowledgements. Paul, thanks again.

Demographics of the San Diego Meeting

This meeting had the largest increase in attendees we've ever experienced. The last IETF meeting in November 1991 in Sante Fe had approximately 360 attendees, and this meeting had 529 attendees. This was the first IETF meeting at which we had more than 400 total attendees.

In San Diego, 46 Working Groups, 17 BOFs, and 2 Area Directorates convened in 94 separate meetings. In addition, the Internet Activities Board met for a full day and evening session. This represents the largest number of Working Group and BOF sessions at an IETF.

So, growth continues to be a concern for us – both as a logistics problem and as a potential impediment for continuing to produce high quality, timely protocols.

Although growth is an important concern, I do not feel we have a cause for alarm... yet. It turns out that we have had one other very large jump in attendance, and this was the last time we held a meeting in California. In July 1989, we met at Stanford University, hosted by BARRnet. That meeting had slightly more than 200 attendees, which represented a 100% increase over the previous meeting. We were very concerned at the time that IETF attendance might be on the verge of exploding, but then the next meeting dropped back down to about 130 attendees. Subsequent meetings showed a steady but manageable increase which we were able to absorb as the Secretariat matured.

Based on the single attendance spike at the Stanford meeting, the feeling is that the proximity of the meeting to the California computer industry may have been the main reason for the increased attendance at the Stanford, and perhaps the San Diego, meetings.

There has also been a technology influx over the last few meetings. In particular, the attendance at the newly created IP Over ATM and Packet Video Working Groups was very high, and this may also have played a role in the steep attendance jump at this meeting.

We will continue to monitor growth closely over the next few meetings to determine if any major changes are needed in the way we conduct our meetings.

Internationalization

One very positive outcome of our growth has been the increase in non-US attendees. At this meeting we had 43 non-US attendees from an impressive number of countries. Although this does not yet represent a big percentage of total attendance, it is none-the-less quite impressive. We had attendees from the following countries:

Canada	9	Korea	3
Denmark	1	Netherlands	3
Finland	1	New Zealand	1
France	2	Norway	3
Germany	3	Sweden	3
Israel	4	Switzerland	2
Italy	1	United Kingdon	7
		Total Non-US	43

By Region:

North America (non-US)	9
European	26
Pacific	4
Middle East	4

This is an area in which we hope to see continued growth as the Internet Society grows in size and importance.

Packet Audio Experiment

Thanks to the organizing efforts of Steve Casner (ISI) and Steve Deering (Xerox), and the behind-the-scenes efforts of Van Jacobson and others, we had a very exciting demonstration of the DARPA packet audio experiment in San Diego. These stalwart experimentors set up IP multicast tunnels through the NSFnet backbone, and broadcast the Plenary proceedings of the IETF to multiple sites across the Internet. Sweden, UK, and Australia all took part in this exercise. We even had a brief 2-way communication, in which several remote listeners spoke to the assembled Plenary.

The quality was not perfect. Some sites had much better reception than others. For some sites, the broadcast was apparently unintelligible at times. Still, for all its imperfections, this demonstration was an impressive promise of services to come.

Some of us speculated that this new technology might play an important role in helping to deal with our future growth. For example, if the Proceedings of the Plenary (and perhaps even certain Working Groups) could be made available as a reliable Internet service, especially if it provided a robust 2-way interaction, it might give prospective attendees an alternative way to participate, rather than flying to attend the meeting in person. The "Information Age" could truly be at hand!

One of the traditional strengths of the Internet community is that we use the technology we are developing to assist that very development. This exciting packet audio demonstration offers the promise of adding to that tradition. We hope to see more packet audio, and perhaps even packet video, experiments at the IETF in the future. More information on this particular demonstration can be found in a presentation by Deering and Casner later in these Proceedings.

ROAD Activities

At the last IETF meeting in Santa Fe (November 1991), we organized a Group to investigate the specific problems of:

- Class B exhaustion
- Routing table scaling
- IP address exhaustion

(See the Proceedings of the November 1991 IETF meeting. Also, see related reports in these Proceedings.)

The "Routing and Addressing" (ROAD) Group grew jointly out of some specific discussions in the BGP Working Group regarding address aggregation using address masks, and out of activities from the IAB Architecture Retreat in July 1991.

I had the pleasure, and indeed, the honor, of co-Chairing the ROAD Group with Peter Ford (LANL). Peter was the driving force in bringing the IAB and BGP Working Group efforts together into a single coherent effort.

The Group met several times between the Santa Fe and San Diego meetings, and explored several alternative solutions for each of the mentioned problems. In San Diego we scheduled several specific BOFs on the ROAD Group findings. We also scheduled several specific BOFs on ROAD related activities during the week. Perhaps most importantly, we began planning how to bring the ROAD results into the IETF in the form of specific working group activities.

On Monday morning, Peter and I briefed the IETF on the ROAD results to help set the stage for the other ROAD related activities during the rest of the week. On Thursday, I presented a summary of the activities, with a preliminary approach on what activities the IETF would undertake. The slides for both these presentations are in these Proceedings.

(Both the Monday morning slides, and the main IETF Agenda give pointers to the other ROAD-related activities during the week).

Perhaps most important is the slide of the names of all the folks who participated in the ROAD meetings. These folks deserve the real credit for an amazing amount of hard work between the two IETF meetings. I truly learned the meaning of the phrase "Road mode" (along with every other "Road" pun and metaphor you could think of!).

The work is by no means over. In fact, the challenge is just starting. We now have to translate the results of the ROAD meetings into concrete actions in the IETF. In Santa Fe, we assumed that we would convene a single "ROAD Working Group" in San Diego. In fact, my Chair's message in the Santa Fe Proceedings was written under this assumption.

However, as we drew nearer to the San Diego IETF, it became clear that we could spin up separate activities to attack the separate parts of the problems. It is a tribute to the ROAD Group that they explored the problems in a very compressed period in sufficient depth to make this possible.

We are still making the final determination on how to proceed in the IETF. Between now and the IETF meeting in Boston, the IESG will take the output from the ROAD Group and develop a recommendation to the IAB on how best to proceed in the IETF.

I look forward to the Boston meeting as the next step toward solving these important Internet problems. I can't wait!

Other Highlights

In San Diego, I was consumed by the ROAD activities. So I must confess that I may have slighted some of the other important activities at this meeting. For that I apologize.

However, one significant achievement was too impressive to escape even my diverted notice – 780 Mbps TCP/IP throughput from Cray to Cray over HIPPI. Wow! The software loopback measurements were even higher, but I think the machine-to-machine numbers are even more impressive.

Congratulations, Dave Borman. You and Van Jacobson are magicians!

New IESG Members

I am pleased to welcome Erik Huizer (SURFnet) and David Piscitello (Bellcore) to the IESG in the role as joint Area Directors for OSI Integration. Erik is our second non-US representative on the IESG (along with Bernhard Stockman of NORDUnet). Dave has been a long-time participant and former Vice-Chair of ANSI X3S3.3, the Group that standardizes network and transport layer protocols.

Unfortunately, this joining is occasioned by a matching departure. Ross Callon (DEC) and Rob Hagens (U. Wisc) previously shared the Area Director responsibilities for OSI Integration. Rob left the IESG last summer and Ross left last Fall after the Santa Fe meeting. The IESG is a significant time commitment, and I appreciate the many contributions Ross and Rob made during their tenure. We will miss their good counsel on the IESG, but we can take solace that both will continue to be active participants in the IETF.

Future Meetings (Don't forget Europe!)

Our next meeting will be held at the Hyatt Regency in Cambridge, Massachusettes, July 13-17, 1992. The Hosting organization is MIT and James Davin and Jeff Schiller are the local hosts.

Our planning for the first IETF meeting in Europe in 1993 continues to take shape. We hope to be able to provide firmer information on the European meeting at the July 1992 meeting.

Agenda of the Twenty-Third IETF

(March 16-20, 1992)

MONDAY, March 16, 1992

7:30-8:00 am	IETF Registration and Continental Breakfast	
8:00-9:30 am	Introductions and Technical Presentation	
	• Local Host Orientation (E. Paul Love, Jr./SDSC) • "Internet Routing and Addressing Considerations" (Peter Ford/LANL and Phill Gross/ANS)	
9:30-12:00 noon	Morning Sessions	
	APP	Network Fax WG (Mark Needleman/UC)
	APP	Network News Transport Protocol WG (Eliot Lear/Silicon Graphics)
	INT	IP over Asynchronous Transfer Mode WG (Bob Hinden/BBN)
	MGT	X.25 Management Information Base WG (Dean Throop/Data General)
	OPS	Network Status Reports (Gene Hastings/PSC)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	OSI	OSI Directory Services WG (Steve Hardcastle-Kille/UCL)
	SEC	Security Area Advisory Group (Stephen Crocker/TIS)
	USV	User Services WG (Joyce Reynolds/ISI)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	Host Resources MIBs BOF (Steve Waldbusser/CMU)
	BOF	Teleconferencing Architecture BOF (Jack Drescher/MCNC and Ari Ollikainen/LLNL)
	APP	Internet SMTP Extensions WG (John Klensin/MIT)
	APP	Network Printing Protocol WG (Glenn Trewitt/DEC)
	OPS	Operational Area Directorate (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
	OSI	OSI Directory Services WG (Steve Hardcastle-Kille/UCL)

	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	SEC	Common Authentication Technology WG (John Linn/DEC)
	USV	Network Information Services Infrastructure WG (April Marine/SRI and Pat Smith/Merit)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	BOF	Extensions to FTP BOF (Jordan Brown/Quarterdeck)
	BOF	IP over HIPPI (Andy Nicholson/Cray Research and John Renwick/Cray Research)
	BOF	Terminal Server Accounting and Authentication BOF (Allan Rubens/Merit)
	APP	Network News Transport Protocol WG (Eliot Lear/Silicon Graphics)
	MGT	DS1/DS3 MIB WG (Fred Baker/ACC and Tracy Cox/Bellcore)
	MGT	Remote LAN Monitoring WG (Mike Erlinger/Lexcel)
	OSI	OSI Directory Services WG (Steve Hardcastle-Kille/UCL)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	Multicast Extensions to OSPF WG (Steve Deering/Xerox PARC)

TUESDAY, March 17, 1992

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	• "Copyright Claims and Standards" (Patrice Lyons, Esquire)	
	• "NSFNET and the NREN" (Stephen Wolff/NSF)	
9:30-12:00 noon	Morning Sessions	
	BOF	SNMP Agent Description BOF (Marshall Rose/DBC and Dave Perkins/SynOptics)
	APP	Telnet WG (Steve Alexander/INTERACTIVE Systems)
	INT	IP over Asynchronous Transfer Mode WG (Bob Hinden/BBN)
	MGT	Internet Accounting WG (Cyndi Mills/BBN and Gregory Ruth/BBN)
	OPS	User Connectivity Problems WG (Dan Long/BBN)
	OSI	MHS-DS WG (Harald Alvestrand/SINTEF DELAB and Kevin Jordan/CDC)
	TSV	Audio/Video Transport WG (Stephen Casner/ISI)
	TSV	Service Location Protocol WG (John Veizades/Apple)
	USV	Internet User Glossary WG (Tracy LaQuey Parker/UTexas and Gary Malkin)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	Internet Routing and Addressing BOF (Peter Ford/LANL and Phill Gross/ANS)
	BOF	MIME to MHS Mapping BOF (Marshall Rose/DBC)
	MGT	Ethernet MIB WG (Frank Kastenholz/Clearpoint)
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	SEC	Common Authentication Technology WG (John Linn/DEC)

	TSV	Audio/Video Transport WG (Stephen Casner/ISI)
	TSV	Dynamic Host Configuration WG (Ralph Droms/Bucknell)
	USV	Internet School Networking WG (John Clement/EDUCOM, Connie Stout/TheNet and Art St. George/UNM)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	BOF	Multiplexing SNMP Agents BOF (Karl Auerbach/Sun Microsystems)
	BOF	Internet Routing and Addressing BOF (Peter Ford/LANL and Phill Gross/ANS)
	APP	Internet SMTP Extensions WG (John Klensin/MIT)
	INT	IP over Appletalk WG (John Veizades/Apple)
	OPS	Network Joint Management WG (Gene Hastings/PSC)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	Open Shortest Path First IGP WG (John Moy/Proteon)
	TSV	Audio/Video Transport WG (Stephen Casner/ISI)
	USV	Internet User Glossary WG (Tracy LaQuey Parker/UTexas and Gary Malkin)
7:00-10:00 pm	Evening Sessions	
	BOF	BGP Deployment and Application BOF (Jessica Yu/Merit)
	BOF	Dynamic Creation of Network Links (T3 Circuits) BOF (Andy Nicholson/Cray Research)
	BOF	Living Documents BOF (Peter Deutsch/McGill and Alan Emtage/McGill)
	MGT	Chassis MIB WG (Jeff Case/UTenn and Bob Stewart/Xyplex)

WEDNESDAY, March 18, 1992

8:30-9:00 am	Continental Breakfast	
9:00-12:00 noon	Morning Sessions	
	APP	Network News Transport Protocol WG (Eliot Lear/Silicon Graphics)
	INT	IP over Appletalk WG (John Veizades/Apple)
	INT	IP over Asynchronous Transfer Mode WG (Bob Hinden/BBN)
	MGT	IEEE 802.3 Hub MIB WG (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
	MGT	Internet Accounting WG (Cyndi Mills/BBN and Gregory Ruth/BBN)
	OPS	Benchmarking WG (Scott Bradner/Harvard)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB)
	RTG	RIP Version II WG (Gary Malkin)
	SEC	Privacy-Enhanced Electronic Mail WG (Steve Kent/BBN)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	CIDR Supernetting BOF* (Phill Gross/ANS and Peter Ford/LANL)
	BOF	IP Routing for Wireless/Mobile Hosts BOF (Steve Deering/Xerox)
	BOF	Teleconferencing Architecture BOF (Jack Drescher/MCNC and Ari Ollikainen/LLNL)
	INT	Point-to-Point Protocol Extensions WG (Brian Lloyd/Consultant)
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	OSI	SNMP over a Multi-protocol Internet WG (Ted Brunner/Bellcore)
	OSI	X.400 Operations WG (Alf Hansen/SINTEF DELAB)
	RTG	Border Gateway Protocol WG (Yakov Rekhter/IBM)*
	SEC	Privacy-Enhanced Electronic Mail WG (Steve Kent/BBN)

	TSV	Domain Name System WG (Mike Reilly/DEC)
	USV	Internet Anonymous FTP Archives WG (Peter Deutsch/McGill and Alan Emtage/McGill)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Technical Presentations	

- "IETF Internet Audiocast" (Steve Casner/ISI and Steve Deering/Xerox PARC)
- "Dynamic Host Configuration Protocol" (Ralph Droms/Bucknell)
- "Traffic Characterization in Wide Area Networks" (Peter Danzig/USC)
- "TCP Large Windows" (David Borman/Cray Research and Bob Braden/ISI)

7:00-10:00pm	Evening Session	
	BOF	IAB Standards Procedure BOF (Lyman Chapin/BBN)
	BOF	Routing Table Lookup Algorithm BOF (Paul Tsuchiya/Bellcore)
	BOF	WAIS and Directory Integration BOF (Steve Hardcastle-Kille/UCL)
	MGT	X.25 Management Information Base WG (Dean Throop/Data General)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	USV	NOC-Tool Catalogue Revisions WG (Robert Enger/ANS and Darren Kinley/RISQ)

*BGP and CIDR will meet jointly.

THURSDAY, March 19, 1992

8:30-9:00 am	Continental Breakfast	
9:00-12:00 noon	Morning Sessions	
	BOF	Traffic Collection, Measurement & Characterization BOF (Peter Danzig/USC, Hans-Werner Braun/SDSC and Kim Claffy/SDSC)
	INT	Point-to-Point Protocol Extensions WG (Brian Lloyd/Consultant)
	MGT	Remote LAN Monitoring WG (Mike Erlinger/Lexcel)
	OPS	User Connectivity Problems WG (Dan Long/BBN)
	OSI	Office Document Architecture WG (Peter Kirstein/UCL)
	RTG	Border Gateway Protocol WG (Yakov Rekhter/IBM)*
	RTG	ISIS for IP Internets WG (Ross Callon/DEC)*
	SEC	Security Area Advisory Group (Stephen Crocker/TIS)
	TSV	Dynamic Host Configuration WG (Ralph Droms/Bucknell)
	USV	Directory Information Services Infrastructure WG (Chris Weider/Merit)
Breaks	Coffee available throughout the morning.	
1:30-3:30 pm	Technical Presentations	
	• "The Internet Gopher Protocol" (Mark McCahill/UMinn) • "Send Time Flow Control Experiments" (Ashok Agrawala/UMD) • "Multi-Media Mail Extensions" (Greg Vaudreuil/CNRI)	
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Open Plenary and IESG	
	• "IP Over X.25" (Andy Malis/BBN) • Protocol Standards Actions	

*BGP and ISIS will meet jointly.

FRIDAY, March 20, 1992

8:30-9:00 am	Continental Breakfast	
9:00-12:00pm	Technical Presentations	
	• "NSFnet Update" (Mark Knopper/Merit and Jordan Becker/ANS)	
	• "NEARnet Presentation" (Scott Bradner/Harvard)	
	• "NIC Services" (Scott Williamson/Network Solutions)	
	• "ISODE Consortium" (Steve Hardcastle-Kille/UCL)	
1:30-3:30	Summary Reports	
	APP	Applications Area (Russ Hobby/UC Davis)
	INT	Internet Area (Noel Chiappa and Philip Almquist/Consultant)
	MGT	Network Management Area (Chuck Davin/MIT)
	OPS	Operations Area (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
	OSI	OSI Integration Area (Erik Huizer/SURFnet and David Piscitello/Bellcore)
	RTG	Routing Area (Bob Hinden/BBN)
	SEC	Security Area (Steve Crocker/TIS)
	TSV	Transport and Services Area (David Borman/Cray Research)
	USV	User Services Area (Joyce K. Reynolds/ISI)
3:30-4:00 pm	Concluding Remarks (Phill Gross/ANS)	

Key to Abbreviations

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

Chapter 1

IETF Overview

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

The IETF mission includes:

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

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

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

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

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

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

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

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

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

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

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

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

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

Mailing Lists

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

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

Information and logistics about upcoming meetings of the IETF are distributed on the general IETF mailing list. For general inquiries about the IETF, requests should be sent to `ietf-info@nri.reston.va.us`. An archive of mail sent to the IETF list is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 Future IETF Meeting Sites

Summer 1992

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

Fall 1992

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

Spring 1993

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

1.2 On Line IETF Information

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

Directory Locations

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

cd ietf
cd internet-drafts

Individual files can then be retrieved using the GET command:

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

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

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

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

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

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

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

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

Europe Address: nic.nordu.net (192.36.148.17)

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

1.2.1 The IETF Directory

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

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

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

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

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

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

1.2.2 The Internet-Drafts Directory

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

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

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

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

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

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

where <author> is the author's name.

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

1.3 Guidelines to Authors of Internet Drafts

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

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

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

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

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

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

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

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

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

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

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

Chapter 2

Steering Group Report

2.1 Standards Progress Report

Between the November meeting hosted by Los Alamos National Laboratories in Santa Fe and the November meeting hosted by the San Diego Supercomputer Center, there have been many IETF originating protocols and informational documents published as RFC's.

The IAB and IESG are working to better document the Standard Process and the protocols in the standard track.

RFC1280	IAB Official Protocol Standards (Informational)
RFC1310	The Internet Standards Process (Informational)
RFC1311	Introduction to the STD Notes (Informational)

Efforts to document widely implemented application protocols continued.

RFC1282	BSD RLogin (Informational)
RFC1288	The Finger User Information Protocol (Draft Standard)

Users and network operations have been the focus of new work.

RFC1290	A Catalogue of Available X.500 Implementations. (Informational)
RFC1302	Building a Network Information Services Infrastructure (Informational)
RFC1308	Executive Introduction to Directory Services using the X.500 Protocol (Informational)
RFC1309	Technical Overview of Directory Services (Informational)

Work continued to specify protocols required to run IP in the public data network environment.

RFC1293	Inverse Address Resolution Protocol (Proposed Standard)
RFC1294	Multiprotocol Interconnect over Frame Relay (Proposed Standard)

Efforts are continuing to instrument the multiprotocol Internet.

RFC1283	SNMP over OSI (Experimental)

RFC1298	SNMP over IPX (Informational)
RFC1284	Definitions of Managed Objects for the Ethernet-like Interface Types (Proposed Standard)
RFC1285	FDDI Management Information Base (Proposed Standard)
RFC1286	Definition of Managed Objects for Bridges (Proposed Standard)
RFC1289	DECnet Phase IV MIB Extensions (Proposed Standard)
RFC1303	A Convention for Describing SNMP-based Agents (Proposed Standard)

2.2 Minutes of the Open Plenary and IESG

Agenda:

- IETF Growth
- Routing and Addressing

2.2.1 IETF Growth

The San Diego IETF was the largest IETF to date. The number of working groups and, Internet Drafts, and RFCs produced is rising at a similar rate to the growth in attendance. Phill Gross briefed the IETF on the growth patterns and encouraged discussion on the impact of this growth.

2.2.2 Routing and Addressing

Routing and Addressing was the primary topic in the Open Plenary. The IAB chartered ROAD Group completed their work to define a set of options for evolution of Internet Routing and Addressing. Phill Gross presented to the Plenary a summary of the options and a possible set of actions to be taken to develop the technology to implement these ideas.

The IETF was generally supportive of the CIDR approach to near term address assignment and aggregation. While there was concern that engineering work had not begun earlier, a strong level of support was shown for making Routing and Addressing a primary focus of near term work.

There was a high level discussion on the possible use of CLNP as a replacement for IP. While well specified proposals for the use of CLNP have not been widely distributed, there was an examination at a high level of security and deployment considerations.

ROAD Results And The IETF

Review

- At the Santa Fe IETF (Nov 1991), a special effort was chartered to look at some key IP issues
- Activities stemming from the BGP WG and from the IAB Architecture Retreat were combined into the "Routing and Addressing (ROAD) Group"
- Goal was to explore possible solutions and report back to IETF in San Diego (March 1992)
- The ROAD co-chairs (Peter Ford and Phill Gross) reported to the IETF plenary in San Diego, and several specific follow-up activities took place during the week
- This is a summary of the initial thinking of how the ROAD results will be followed up in the IETF

Problems

- Class B address exhaustion
- Routing table explosion
- IP address exhaustion

Directions

- #1– Better way to deal with current addresses (e.g., hierarchical assignments for aggregation in routing)
  - (to help slow class B exhaustion and routing table explosion)
- #2– Bigger addresses for the Internet
  - (to stem IP address exhaustion)
- #3– Need focus on future directions for Internet routing and advanced features

Course of Action

- Separate #1 (dealing with IP addresses better) from #2 (bigger Internet addresses), and pursue each in the IETF in parallel
  - (i.e., to the extent technically feasible, do "not" allow #1 to be a gating item for #2)
- Pursue #3 (future routing) in IRTF
- Start #1, #2, and #3 now ("ALL" are vitally important!)
- Estimated delivery window:
  - for #1, ~ 12 - 18 months
  - for #2, ~ 3 - 5 years
- CIDR for #1., and make "VERY" high priority

Classless Inter-Domain Routing (CIDR)

- Addressing and routing plan
  - Hierarchical assignment, aggregation for routing
- Enhanced Inter-domain protocols
- Inter-Domain routing "Usage documents" for using addressing and routing plan with the enhanced inter-domain protocols, and for interacting with IGPs

In The IETF -- CIDR

- a. Publish CIDR document (overview and guidance)
- b. New WG for "IP Address Assignment Guidelines"
- c. BGP WG for any BGP extensions
- d. New WG for IDRP/IP
- e. "CIDR Deployment WG" (which could include CIDR/IGP interactions)

In The IETF -- Bigger Internet Addresses

- Need to flesh out various proposals with goal of picking a single direction "SOON"
  - CNAT and "Simple CLNP"
  - IP/IP Encapsulation
  - Others?
- IAB/IETF needs to ensure WG formation/review

Summary

- CIDR: Start "NOW", deploy in 12-18 months
  - Components of CIDR now proceeding in several WGs
  - Treat "Addressing Plan" and deployment planning as operational issues
- Bigger Internet addresses: Pick single solution "very soon", deploy within 3-5 years, BOFs and/or WGs by next IETF
- IESG will accept ROAD results and issue recommendations to the IAB by next IETF
- IESG and IAB must monitor all activities closely to ensure progress; IESG, ADs, WGs will report on ROAD-related activities at future IETFs

Acknowledgements

- All the details are not completely worked out, but the ROAD group has helped us take the first step on an important path for the Internet
- Grateful thanks are due to the folks who initiated the ROAD activity (the IAB and the BGP WG) and the ROAD participants themselves for their combined hard work and vision

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UCDavis

Area Overview

The general goal of the Area is to define the protocols to create an interoperable multimedia distributed computing environment for the Internet. A lot of attention at the San Diego meeting was directed toward the creation of workstation based conferencing. There were demonstrations of how video and audio could be carried between workstations over the Internet today. The Teleconferencing Architecture BOF discussed what is necessary to create a multi-workstation multimedia conference environment. With the greater international interest in the IETF there is a desire to be able to extend the "face-to-face" meetings over the network around the world.

Extensions to FTP BOF (FTPEXT)

The FTP Extensions BOF was to determine if there is sufficient interest to work on the definition of new extensions to FTP. There were many ideas were discussed, in particular the passing of file attributes with the transfer. Jordan Brown (Quarterdeck) agreed to Chair a new Working Group to continue the work.

Internet Mail Extensions Working Group (SMTPEXT)

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

Network FAX Working Group (NETFAX)

The Working Group finished work on the image format to be used for transporting FAX on the Internet. The Document has been submitted for approval as a Proposed Standard. This completes the chartered work and the Working Group will go dormant.

Network News Transport Protocol Working Group (NNTP)

The Working Group met three times in San Diego to walk through the existing NNTP v2 draft. After final editing changes the document will be submitted to be a Proposed Standard. The Working Group will go on to look at a possible News Reader Protocol.

Teleconferencing Architecture BOF (TELARCH)

The BOF explored how a Working Group could define an overall architecture to cover all aspects of remote conferencing. This included things such as session management and groupware. Separate working groups may be generated to work on specific protocols specified by the architecture. Jack Drescher (MCNC) and Ari Ollikainen (LLNL) will co-Chair the new Working Group.

TELNET Working Group (TELNET)

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

CURRENT MEETING REPORT

Reported by Jordan Brown/Quarterdeck

Minutes of the Extensions to FTP BOF (FTPEXT)

The mail traffic and discussion at the meeting basically involved people's wish lists for the protocol. Topics included:

- Passing "auxiliary" information about files - dates, etc.

  The goal would be to build an extensible mechanism allowing a client and server to pass "auxiliary" information about files. Extended versions of LIST, RETR, STOR, etc., would pass this information, and a new command would be added to change it. Matching client and server should be able to pass all of the information their native system supports; non-matching pairs would pass as much as they have in common. A major open issue is whether the data should be passed in binary as type-length-value or in some printable-ASCII form.

- Automatic authentication

  There are two basic ways in which authentication data might be passed at present - using FTP commands or, relaxing the specification a bit, using TELNET options on the control connection. It was suggested that authentication and encryption are big complex issues on their own, and that they be split off from the rest of the items on the wish lists.

- Encryption

  There was interest in encryption of both the data and the control channel. Encryption is tightly tied to authentication, and the two should probably be treated as a unit.

- On-the-fly compression

  Some servers already implement on-the-fly compression triggered by variations in the file name. The patent status of LZW is an issue which needs to be researched and resolved.

- Checkpointing/restart

  Some attendees sought official blessing for Rick Adams' stream mode restart capability (present in some BSD clients and servers). It was noted that it is unclear whether or not this mechanism truly works for NVT-ASCII mode transfers. It was clear that the restart marker for this mechanism should be measured in data-connection octets. Implementing such a restart mechanism might be tricky on systems where the local <-> network translation is not strictly invertible.

- Language selection for messages

  Seems pretty self-explanatory; obviously no system will support all languages, but support for multiple languages seems reasonable. This issue will come up in other contexts (SMTP, etc.,); perhaps there should be a more global framework.

- Message digest calculation

  The goal is to allow automatic mirroring of archives without having to transfer all of the data.

- Atomic store

  The disposition of the file resulting from a failing STOR is unspecified; a new command would require that the file be deleted if the transfer was not completed successfully.

- Various protocol cleanups/clarifications

  RFC 1127 lists several response code cleanups and clarifications. Experience with newer servers which make more extensive use of multiline responses indicates that not all clients can handle them. The syntax for multiline responses is apparently not clear enough; there has been confusion. Note that FTP multiline responses are more liberal than SMTP multiline responses.

- More sophisticated conversions - character set, application levels, etc.

  An extended version of NVT-ASCII mode would allow for transmission of non-USASCII characters; a mechanism would be needed to specify what character set is in use and what translations should be applied. This issue has already been addressed in Kermit and the lessons learned there should be applied. A still more sophisticated mechanism to automatically do application-level transformation (e.g., Microsoft Word to TeX) would certainly be useful, but is obviously a very complex topic.

- Should write both a specification and an "implementor's guide"

  It was mentioned that FTP has numerous common pitfalls, and an informational document pointing out these pitfalls and suggesting implementation schemes would help implementors and improve interoperability.

- Time conversion issues - time zones, DST, etc.

  Once FTP is passing around time information (file modification times, mostly), it becomes important to know what the times really mean, so that meaningful comparisons and conversions can be done. One obvious answer is to require that all times be expressed in GMT, but that is awkward for the large (and ever- increasing) number of machines which don't know what time zone they're in. One scheme for dealing with this would be to provide a command which gives the server's current time with respect to whatever TZ it finds convenient; the client can compare this with its current time to determine the offset to be applied to other times. This requires very loosely synchronized clocks - less than 15 minutes difference. It's not clear whether DST confuses this issue - a file stored under DST and later retrieved under ST might have its times mistranslated. (Portable computers make time issues a nightmare.)

There was consensus that a Working Group should be formed, and when a deafening silence resulted from a call for volunteers to Chair it, Jordan Brown agreed to. (Volunteers are still solicited!) Russ Hobby offered to host the mailing list and archives. The initial mailing list is the BOF attendees.

Mailing list:	ietf-ftpext@ucdavis.edu
Requests:	ietf-ftpext-request@ucdavis.edu
Archive:	ucdavis.edu: /archive/ftpext-archive

No date was set for the next meeting.

Attendees

Rick Adams	rick@uunet.uu.net
James Alfieri	jdal@troy.cc.bellcore.com
J. Allard	jallard@microsoft.com
Mark Baushke	mdb@nsd.3com.com
David Borman	dab@cray.com
Jordan Brown	jbrown@qdeck.com
Russ Hobby	rdhobby@ucdavis.edu
John Klensin	klensin@infoods.mit.edu
Henry Miller	henrym@sacusr.mp.usbr.gov
Keith Moore	moore@cs.utk.edu
Chris Myers	chris@wugate.wustl.edu
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Jon Postel	postel@isi.edu
Sam Sjogren	sjogren@tgv.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Veizades	veizades@apple.com
Kathleen Wilde	wilde@decvax.dec.com

CURRENT MEETING REPORT

Reported by Jack Drescher/MCNC and Ari Ollikainen/LLNL

Minutes of the Teleconferencing Architecture BOF (TELARCH)

Two BOF sessions were held on Teleconferencing Architecture as a follow-on to both the Santa Fe IETF BOF on Teleconferencing and the Packet Video Videoconferencing Workshop at MCNC, held in late 1991. The general objective of the BOFs was to move the process of forming a new Teleconferencing Architecture Working Group forward as much as possible. A proposed Draft Charter and set of Milestones was reviewed, relevant current work efforts were presented and discussed and a number of recommended action items were presented from the floor. A request for input to the future architecture content was made. Thirty-eight people attended all or part of the sessions.

The Monday session was introduced by showing consistency of transition from the Santa Fe BOF Minutes and the MCNC workshop objectives. The Proposed Draft of the Teleconferencing Architecture Working Group was presented in outline form, along with an initial set of milestones. A brief review of the MCNC Packet Video Project, for timing perspective, was presented. Eve Schooler, ISI, presented the Connection Manager and Connection Control Protocol work being done at ISI as part of the Multimedia Conferencing Project. Abel Weinrib, Bellcore, presented an additional (to Santa Fe BOF) presentation on the Touring Machine project at Bellcore. Both presentations were in response to the request for input to the overall teleconferencing architecture process. The objectives for the Wednesday session were reviewed.

The Wednesday session was begun with some observations and discussion about the Audio/Visual Transport Working Group session and the Intellectual Property Law plenary presentation of Tuesday. MCNC stated that they would name someone to be an active member of the AVT Working Group. It was also concluded that we all need to pay attention to Intellectual Property Law, that there was nothing unique about it to this Working Group, and that any information submitted to this Working Group and architecture produced by it should not be proprietary.

Paul Milazzo of BBN led a presentation/discussion of a number of questions, issues and definitions requiring answers and clarification to better define and bound the architecture and Working Group Charter.

A summary of key points, conclusions, decisions follows:

1. The completed architecture, targeted for 1st Quarter 1993, needs to include:
    - Supporting Protocols
    - Statement of Conferencing Applications Supported
    - Network Resources Required
2. Applications need to be modular and scalable. Models are needed.

3. "Groupware" needs to be better defined and expanded (as used in the draft Charter)
4. Set a new milestone, as near in as possible, for sharper definition of applications supported.
5. Identify separation of pieces that make up the total Teleconferencing Architecture as soon as possible. Note: MCNC will try to provide first pass structure proposal for list to review by May 1, 1992.
6. On the question of separating Connection Management from proposed Teleconferencing Architecture Working Group, the Group decided that it didn't know enough about all the pieces to do that at this time. Connection Management work should continue on its current course.

   Note: Agreement on (5) above should make questions like this easier to answer correctly.
7. Charter update process:
   - Update Charter, Pass by list for review.
   - Submit to Russ Hobby for input to Working Group .
   - Formalization process.

Proposals Made:

1. The Working Group mailing list, along with that of AVT Working Group, should be rem-conf@es.net.
2. Consider renaming proposed Teleconferencing Architecture Working Group to "Remote Conferencing" Working Group

Attendees

Joe Blackmon	blackmon@ncsa.uiuc.edu
John Burnett	jlb@adaptive.com
Stephen Casner	casner@isi.edu
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Cogger	rhx@cornell.cit.bitnet
David Crocker	dcrocker@mordor.stanford.edu
Jack Drescher	drescher@concert.net
Maria Gallagher	maria@nsipo.nasa.gov
Tony Hain	hain@es.net
Russ Hobby	rdhobby@ucdavis.edu
Alton Hoover	hoover@nis.ans.net
H. A. Kippenhan	kippenhan@fndcd.fnal.gov
Peter Kirstein	kirstein@cs.ucl.ac.uk
Holly Knight	holly@apple.com

Jim Knowles	jknowles@binky.arc.nasa.gov
Linda Leibengood	ldl@ans.net
Arthur Lin	yalin@srv.pacbell.com
Larry Masinter	masinter@parc.xerox.com
Paul Milazzo	milazzo@bbn.com
Ari Ollikainen	ari@es.net
Ayal Opher	aopher@synoptics.com
Geir Pedersen	geir.pedersen@usit.uio.no
Michael Roberts	roberts@educom.edu
Ron Roberts	roberts@jessica.stanford.edu
Kevin Rowett	kevinr@tandem.com
Allan Rubens	acr@merit.edu
Eve Schooler	schooler@isi.edu
Michael Smith	mes5674@hertz.njit.edu
Scott Stein	scotts@apple.com
R.B. Swenson	rswenson@igc.org
Claudio Topolcic	topolcic@nri.reston.va.us
Dono van-Mierop	dono_vanmierop@3mail.3com.com
John Veizades	veizades@apple.com
James Watt	jamesw@newbridge.com
Abel Weinrib	abel@bellcore.com
Jil Westcott	Westcott@BBN.COM
Jonathan Whaley	whaley@concert.net
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk

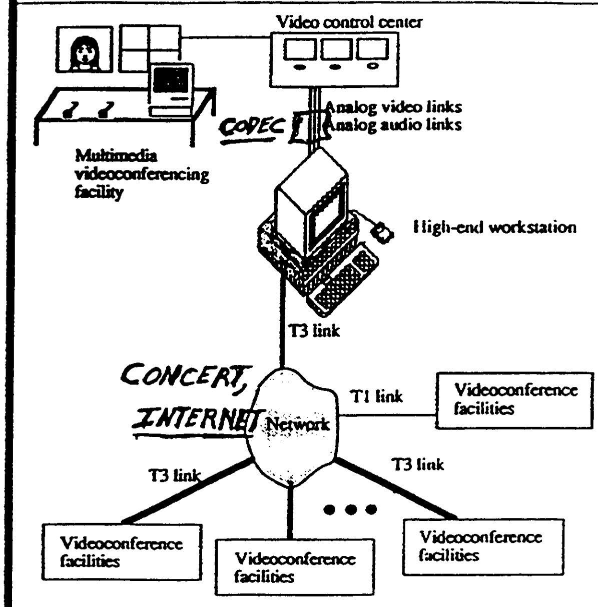

Figure 1. Network Interconnect for Multimedia Conferencing Using Packet Video over the INTERNET

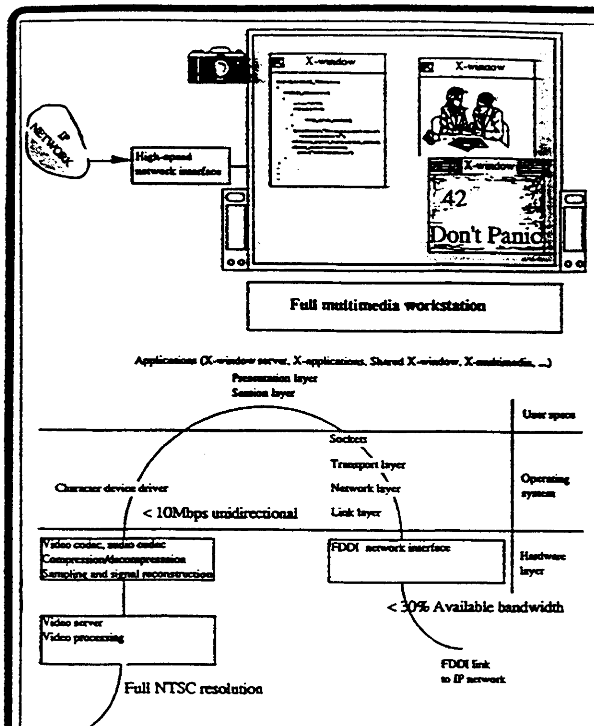

Figure 6. Phase V System

3.0 Schedule

The project is divided into five different phases that are distinguished by the associated tasks.

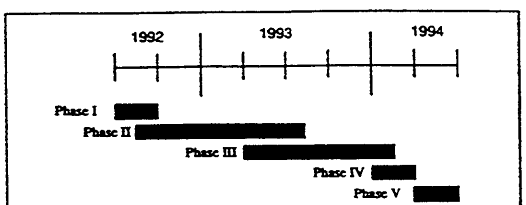

Figure 7. Schedule

Phase I:

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

Phase II:

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

3.0 Schedule

The project is divided into five different phases that are distinguished by the associated tasks.

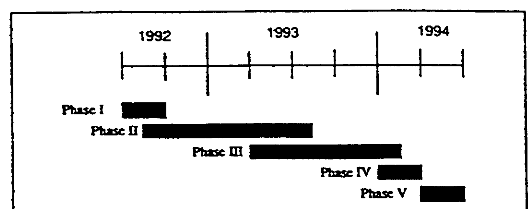

Figure 7. Schedule

Phase III:

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

Phase IV:

- Transport protocol comparison.
- Application software.

Phase V:

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

PACKET VIDEO CONSORTIUM

ORGANIZATIONS THAT HAVE EXPRESSED INTEREST

(C)	MCNC	Drescher
	NSF	Strawn
(C)	ADVANCED NETWORK AND SERVICES	Gross
	SUN MICROSYSTEMS	Austin
(C)	IBM	Bosco
	COMPRESSION LABS, INC.	Lookabaugh/Glancy
(C)	SEMICONDUCTOR RESEARCH CORPORATION	Lucic
(C)	N.C. STATE	Abbott/Nilsson
(C)	UNC-CHAPEL HILL	Jeffay/Smith
(C)	UNC-CHARLOTTE	Stahl

NOTE: ASSUMPTION IS THAT CONSORTIUM MEMBERSHIP IS OPEN-ENDED

NOTE: (C) - COMMITTED

An Architecture for Multimedia Connection Management

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

Layered Protocol Stack

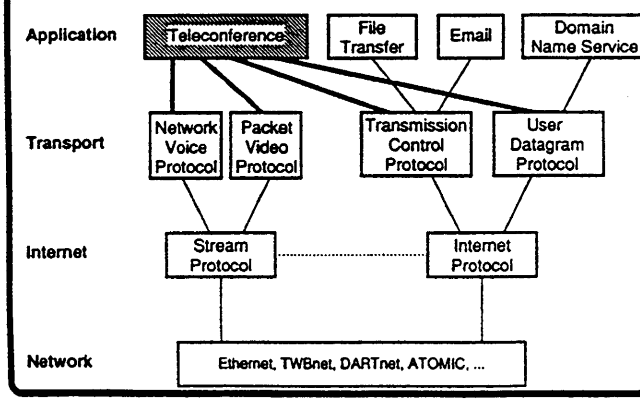

I. A Connection Management Architecture

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

    - *Conduit through which control info flows (locally and remotely)*

Coordinated Management of Separate Services

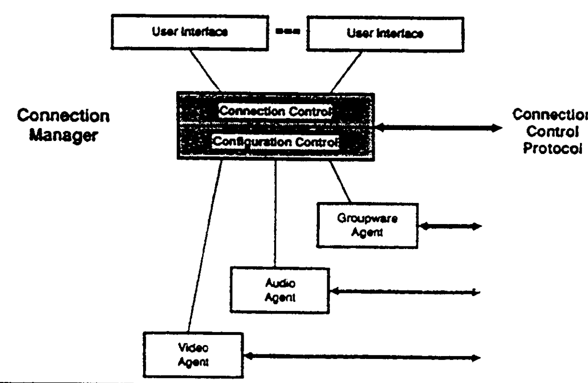

II. The Connection Control Protocol (CCP)

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

A Distributed, Peer-to-peer Model

- Peer connection managers reside throughout the Internet
  - Each acts as both client and server
- Conference orchestration entails:
  - The *initiator* is designated leader for duration of setup
  - Communication with peer CMs
  - Communication with local agents
- Four-phase connection establishment procedure

  1. Negotiate a common set of *capabilities*
  2. Request others' participation
  3. Initiate underlying voice, video and groupware data flows
  4. Propagate info among peers, then revert to having no special status

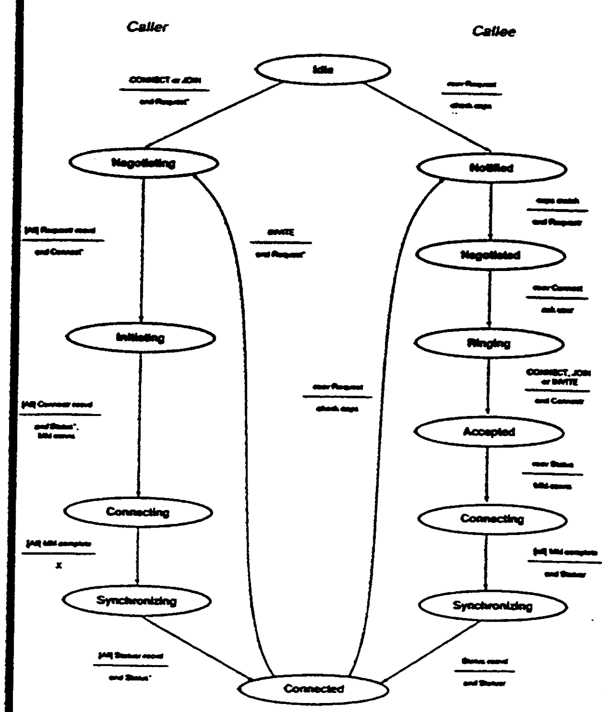

Other CCP Attributes

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

III. Configuration Management

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

Several mechanisms proposed to combat heterogeneity:

- Configuration language
- Distributed locator service
- Resource synthesizer

Dynamic Configuration Management of Shared Resources

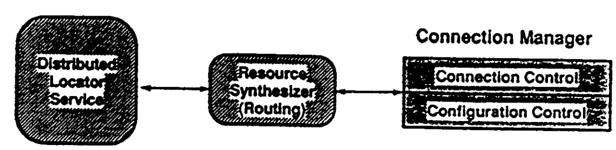

The Site Wide Operating Environment

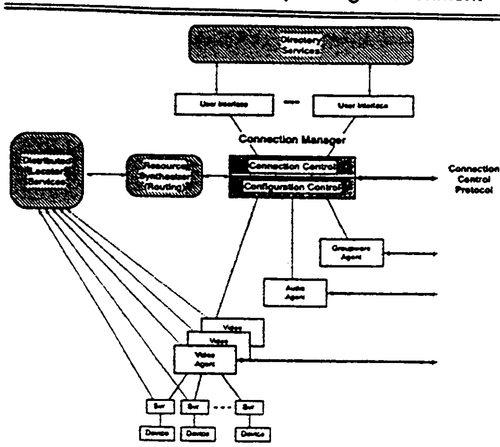

Distributed Program Complications

- Temporary connectivity outages
  - Multicast transaction service must report failures
  - CCP must repair partitioned state afterwards
  - Institute a range of policies toward communication errors
- Different routes taken by different media
  - Convey synchronization hints to synch proto
- A range of delays due to Internet, due to config variations
  - Incorporate adaptive timeouts

TOURING MACHINE:

Distributed Systems for Multimedia Communication

Abel Weinrib
abel@bellcore.com

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

Bellcore

Touring Machine Project Goals

Explore design and implementation issues of a large-scale multimedia communication system
- bridge together computer and communications systems research
- construct a system experiment to explore various approaches
- operate under "Real World" conditions

Design and support Application Programming Interface that makes available network's capabilities and core services
- API developed in collaboration with applications research projects Cruiser™ and Rendezvous™ (Rob Fish, Carlin Lowery, John Patterson)
- valuable input from EXPANSE (Steve Minzer) and Patch-Cords (Rich Clayton)
- Cruiser project is building applications using Touring Machine software infrastructure
  - use technology to overcome distance barrier
  - informal communication by "cruising" hallways

In contrast to many research projects in this area, we are primarily interested in developing the overall software architecture to support ubiquitous multimedia communications, rather than studying the applications themselves.

AW 12/14/91

Current Touring Machine Hardware

Analog voice and video
- existing inexpensive off-the-shelf hardware
- existing wiring

In every user's office
- station equipment
- workstation to act as a station controller

In each "central office"
- audio/video switches
- computer to control switches
- special purpose hardware
  - video and audio bridges and mixers
  - inputs to provide cues to users

In the wiring plant
- balanced audio signals through telephone wiring
- NTSC video signals through RG59 coax
- cross-connect system to allow easy reconfiguration

AW 12/14/91

Multimedia Telecommunications Service

One, of many, Touring Machine applications
- separate control of media
- control of multiple sessions
- suspend/resume
- add/delete parties
- change configuration

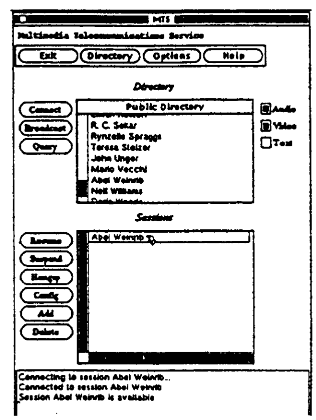

AW 12/14/91

Touring Machine Features

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

Separate control of media
- audio
- video
- data

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

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

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

AW 12/14/91

Touring Machine World View

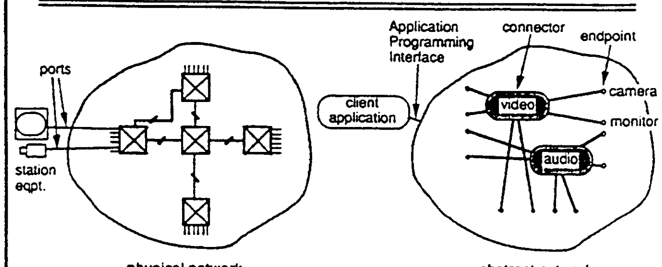

Ports **are network access channels connected to physical station (CPE) hardware**

Connectors **are associations of source and sink endpoints for transport in a given medium**

Endpoints **are "logical ports" defined by an application**

Endpoints are *assigned* to ports, and then *mapped* and *unmapped*

AW 12/14/91

API – Registration

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

Users initiate interactions with the "network" by starting up a client application that registers with Touring Machine:

```
(registerClient abel:MTS
  (registerEndpoint
     camera:video:source monitor:video:sink...)
  (setPrivacy group))
```

AW 11/8/91

API – Sessions

Session establishment and modification

The "session" separates *control* from *transport* for communicating applications
- negotiate and set privacy, billing, and other policies
- permits "publishing" a session, allowing others to locate and join it
- negotiate and define transport topology (as connectors)

User creates session to establish a three-way audio-video call:

```
(sessionCreate abel:MTS:1
  (setPrivacy all)
  (setPermission protected)
  (addClient abel:MTS vicmak:MTS wel:MTS)
  (addCon v1 video)
  (addSource v1 abel:MTS:camera vicmak:MTS:camera
     wel:MTS:camera)
  (addSink v1 abel:MTS:monitor vicmak:MTS:monitor
     wel:MTS:monitor)
  (addCon a1 audio)...)
```

AW 11/8/91

Connectors

a connector is an association of sources and sinks for transport

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

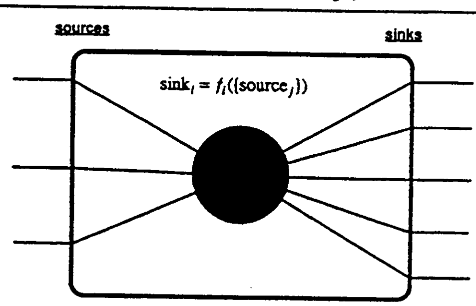

AW 1/16/92

Allocation of Connectors

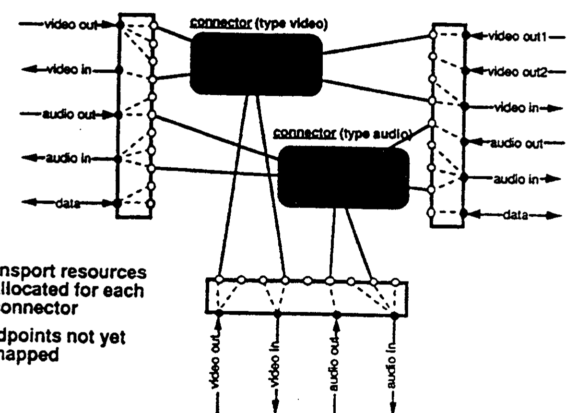

transport resources allocated for each connector

endpoints not yet mapped

AW 1/16/92

API – Network Access

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

Map and unmap give applications control over the network-access resources (channels connected to cameras, etc.) separate from the network resources allocated for transport

User "suspends" the video for a connection:

```
(endpointUnmap abel:MTS:1 abel:MTS:monitor abel:MTS:camera)
```

User moves video to port `vout2` *and resumes:*

```
(endpointAssign abel:MTS:1 abel:MTS:monitor,station7:vout2)
(endpointMap abel:MTS:1 abel:MTS:monitor)
```

AW 11/8/91

Mapping of Endpoints to Ports

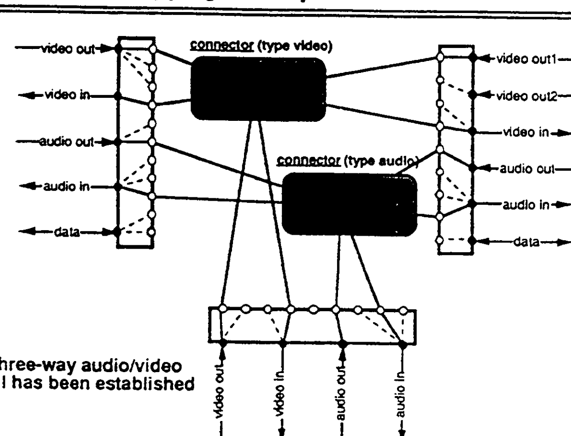

a three-way audio/video call has been established

AW 1/16/92

More Complicated Example

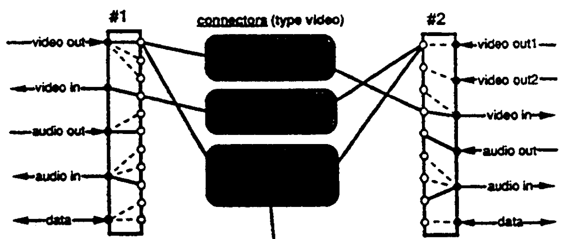

user #3 monitors a two-way conversation, but user #2 has suspended her video output

(source endpoints are shared between connectors)

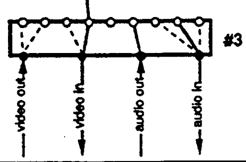

AW 1/16/92

API – Name Service and Messaging

Name server queries

Allows clients to access system information

User searches for all members of a session `CSS-seminar`***:***

```
(nsQuery ... session.name=CSS-seminar session.client)
```

User searches for all sessions of which user `vicmak` ***is a member:***

```
(nsQuery ... session.client=vicmak
  session.name, session.permission)
```

Inter-client message forwarding

Provides facility for clients to exchange control messages through Touring Machine

Coordinate appearance of RENDEZVOUS™ shared whiteboard application:

```
(messageSend abel:whiteboard vicmak:whiteboard "ready")
```

AW 11/8/91

Software Structure

Touring Machine is a distributed system constructed of several types of objects
- *not* a single monolithic program
- each object encapsulates a set of capabilities and/or system state
  - many changes are local to a single object
- layered design, from API to control of physical resources

Touring Machine objects
- can run on any UNIX™ machine
- communicate by exchanging messages
- behavior characterized by well defined interfaces (suite of messages)

AW 12/14/91

Anatomy of an Object

Touring Machine object
- implemented as a passive server
- behavior characterized by its message suite
- IPC uses the Connection Manager (Peter Bates)
  - built on UNIX™ sockets
  - objects listen for messages on well-defined ports
  - objects referenced by `machine:service`
    `rocket:tm.Station`
  - messages are LISP-like S-expressions
    `(sessionCreate 101 "POCS" "ms" ...)`
  - LISP function names bound to C functions and automatically called by the Connection Manager
- asynchronous
  - requests and replies are matched with "reply tokens"
  - timeout mechanism aborts actions that receive no reply

AW 12/14/91

Version 2.0 Architecture

Long-lived objects

Station Manager (optional)
- implements resource-sharing policies among clients
- provides common functionality for clients (e.g., call forwarding)

Station Object
- provides interface to Touring Machine
- manages station resources
- registers clients

Resource Manager
- manages physical resources
- maintains approximate state information ("hints")

Resource objects
- control physical resources

Name Server
- repository for static and dynamic system information

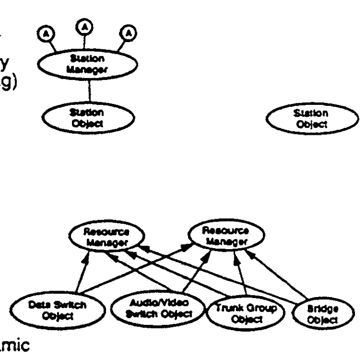

AW 9/25/91

Version 2.0 Architecture

Station Object creates Session Object

Session Object
- site for negotiation between Station Objects
- maintains logical state of session
- derives "call graph" specified as connectors

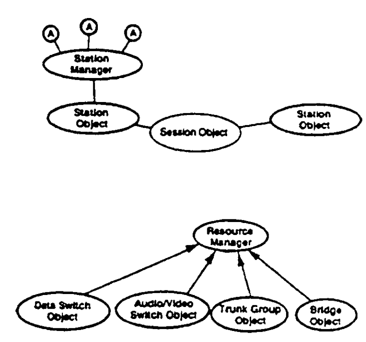

AW 9/25/91

Version 2.0 Architecture

Session Object connects to available Resource Manager

Resource Manager
- allocates physical resources to realize connectors
- "spawns" Transport Object for the session

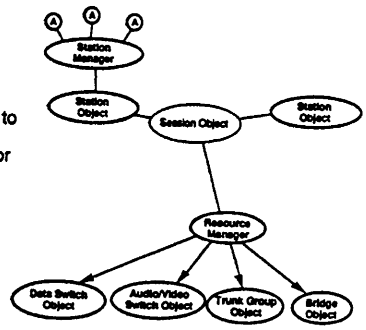

AW 9/25/91

Version 2.0 Architecture

Transport Object
- responds to Session Object
- maintains logical-to-physical mapping for session
- as *optimization*, processes some commands (e.g. map/unmap)

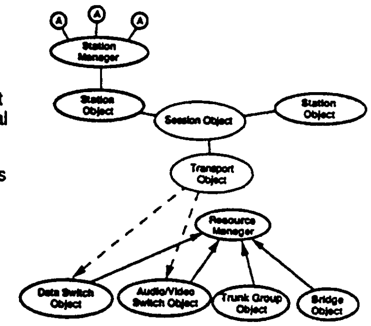

AW 9/25/91

Version 2.0 Architecture

Transport Object connects to any available Resource Manager for more complex operations
- change transport resources
- "tear down" connectors

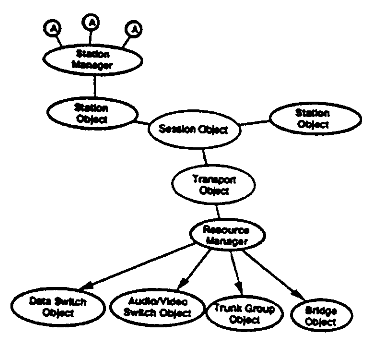

AW 9/25/91

Bellcore's Touring Machine Network

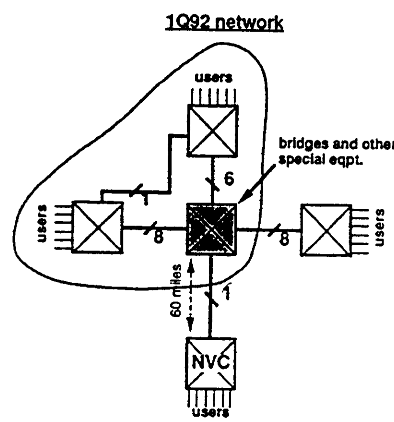

AW 12/14/91

1992 Research Directions

Applications Programming Interface
- extensions to API as applications are developed
  - general presentation control
- abstractions for hybrid analog/digital network

System software
- multiple (mistrustful) administrative domains
- fault management
- system instrumentation and observation
- on-line replacement of software modules
- privacy and security
- resource finding
- naming and addressing
- support for user mobility

All-digital network

Mechanisms for discovering and controlling interactions between different applications

Deployment opportunities

AW 12/14/91

3.1.1 Automated Internet Mailing List Services (list)

Charter

Chair(s):

David Lippke, lippke@utdallas.edu

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

Goals and Milestones:

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

3.1.2 Distributed Scheduling Protocol (chronos)

<u>Charter</u>

Chair(s):

Paul Linder, lindner@boombox.micro.umn.edu

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

Jan 1991	Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. Prototype implementations.
Feb 1991	Make document an Internet Draft. Continue revisions based on comments received over e-mail.
Mar 1991	Spring IETF meeting. Review final draft and if OK, give to IESG for publication as RFC. Begin implementations.
Jul 1991	Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

3.1.3 Internet Mail Extensions (smtpext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done	Review the Charter of the Group. Determine if changes to SMTP are necessary. Discuss the needs for backward compatability, and interoperability. This discussion will be held by email.
Aug 1991	Discuss the elimination of the 7 bit restrictions in SMTP, and the implications of removing this restriction in terms of interoperation.
Aug 1991	Discuss the issues involved with binary transmission. Determine whether a "binary" mode should be pursued, and whether the SMTP line length restriction should be eliminated.
Dec 1991	Write a document specifying the changes to SMTP agreed to by the Group. Post as an Internet Draft.
Mar 1992	Review and finalize the SMTP Extensions document.
Mar 1992	Submit the SMTP Extensions document as a Proposed Standard.

Internet Drafts:

"SMTP Extensions for Transport of Enhanced Text-Based Messages", 07/10/1991, John Klensin <draft-ietf-smtpext-8bittransport-04.txt>

CURRENT MEETING REPORT

Reported by John Klensin/MIT

Minutes of the Internet Mail Extensions Working Group (SMTPEXT)

During the San Diego meeting, the Working Group reviewed several issues that had been settled earlier, but for which it appeared that new technical issues had been identified. The Working Group concluded that there were, in fact, no significant new technical issues being raised, and no significant changes to the working document were made. The Working Group did succeed in tying up the remaining loose ends in the document, including identifying locations where additional explanatory text was needed and providing exact keywords, syntax, and definition for concepts agreed upon some time ago.

The present draft provides an extension model and compatible extensions to SMTP for mail transport of 8-bit characters. Using the same extension model, it provides some additional extensions to supplement SMTP and improve its efficiency, especially when very large files are being transmitted.

It is expected that a new Internet Draft, reflecting agreements made in San Diego, will be produced shortly after the IETF, reviewed quickly on the mailing list, and then submitted for processing as a Proposed Standard.

The meeting contained an extended discussion of the issues raised the previous day, including a review of whether the Working Group's work and the RFC-ZZZZ model fit well into a "transition plan"/"final target" model. Several other issues were revisited, including the question of whether we might be better off with a new protocol on a new port. The assertion was made that the present model either constituted two separate services over the same port (in which case it was muddled and wrong) or some attempt at hidden gateways (in which case there was fear of other problems). It was pointed out that the two port model made it very difficult to distinguish between no service on the "new" port and temporary unavailability. The advocates of the two-port model felt that this was a non-problem unless one wanted to mix services or have implied gateways. It was pointed out that no gateway was implied if the originating system was prepared to send a message in either 8-bit or 7-bit form, but preferred 8 if that was available. There was a brief religious argument as to whether or not this case was interesting.

The Working Group concluded that it did not want to revisit the "new port" issue.

A second heated discussion over the CPBL command with a contention that the command would increase the number of round trips pitted against a contention that it would give intelligent clients the ability to reduce the actual number of round trips. The Working Group decided to retain CPBL.

The issue of "SIZE" was reviewed, both with regard to the information to be returned by CPBL (redesignated as "LIMIT" after the meeting to reduce possible confusion) and with regard to the estimation and treatment of the value to be sent with the SIZE verb. The Working Group concluded that the capability limit should be reported in terms of two administrative values, the size that the implementation would try to provide in all cases and the size that would probably always be rejected. The editor was directed to provide some guidance in the document for estimating, in hosts with single-character newline conventions internally, the size to be transmitted. See the new version of the draft document for the results of these discussions.

Details of Discussions and Decisions

As discussed in the summary above, much of the two Working Group sessions at this meeting was devoted to a review of previously-settled issues, sometimes from a new point of view. Issues and proposals raised included:

- Syntax and semantics for the response to the CBPL command/inquiry. The Working Group decided that this should list all capabilities of the server, on a one-verb-per-line basis, using the existing syntax for multiline responses. This is one of the options considered in Santa Fe and tentatively approved. The handling of the "message size" portion of the response was as agreed on in Santa Fe, i.e., providing the administrative limits.

- Syntax and semantics of the SIZE command. The decisions made in Santa Fe were affirmed and refined.

  See the forthcoming version of the Internet Draft for additional details on the two features above.

- Possible separation of a "transition model" from the "protocol" as it would exist in deployed form, e.g., creating two clearly separated documents. The Working Group concluded that this was neither necessary nor appropriate.

- Possible replacement of the notion of capabilities inquiries with a clear "version" model with no optional features. The theory behind this was that the Internet has succeeded because of the small number of options (Telnet negotiation notwithstanding) and that few clients have really been designed to take advantage of optional features in any event. This discussion led to the insight that some confusion has been created by describing EMAL and related features in "negotiation" terms. They are really verification that particular expected capabilities are present.

  The next version of the Internet Draft will be modified to reflect the "verification" terminology and to remove "negotiation". This does not affect the operation of the proposed protocol.

  After some discussion, the Working Group concluded that a "version" model was inappropriate. Different members saw different problems, but the most serious included the fact that SMTP already contains optional features (e.g., the interactive mail commands SEND FROM, SOML FROM, and SAML FROM) and that introducing strict "all or nothing at this level" versioning would require, either requiring support for those verbs and concepts, or deprecating them. The Working Group was unwilling to consider doing either; some members considered such tampering with the requirement level for existing features to be outside the Working Group's scope. There were also concerns about excessively raising the level of effort required for a minimal implementation of the new features, thereby defeating the Working Group's goal of providing as smooth and easy a transition path to existing (nonconforming) 8bit-sending implementations as possible.

- There was interest expressed in "address streaming" and, in particular, return of sequence numbers in replies to RCPT TO commands.

  The Working Group concluded that this was not an extension it was willing to make to SMTP during the current round, especially since no concrete or specific proposal was on the table.

Attendees

Serge Aumont	aumont@cicb.fr
Robert Austein	sra@asylum.sf.ca.us
Mark Baushke	mdb@nsd.3com.com
Ronald Broersma	ron@nosc.mil
James Conklin	conklin@bitnic.educom.edu
David Crocker	dcrocker@mordor.stanford.edu
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Jisoo Geiter	geiter@gateway.mitre.org
Olafur Gudmundsson	ogud@cs.umd.edu
Max Hillson	hillson@koala.enet.dec.com
Tim Howes	Tim.Howes@umich.edu.
Neil Katin	katin@eng.sun.com
Tom Kessler	kessler@eng.sun.com
John Klensin	klensin@infoods.mit.edu
Jim Knowles	jknowles@binky.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Eliot Lear	lear@sgi.com
John Leong	john.leong@andrew.cmu.edu
Henry Miller	henrym@sacusr.mp.usbr.gov
Keith Moore	moore@cs.utk.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Mark Smith	mcs@umich.edu
Einar Stefferud	stef@nma.com
Theodore Tso	tytso@mit.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.4 Internet Message Extensions (822ext)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done	Review the Charter, and refine the Group's focus. Decide whether this is a worthwhile effort.
Done	Discuss, debate, and choose a framework for the solution. Assign writing assignments, and identify issues to be resolved.
Done	Review exiting writing, resolve outstanding issues, identify new work, and work toward a complete document.
Done	Post a first Internet Draft.
Nov 1991	Review and finalize the draft document.
Dec 1991	Submit the document as a Proposed Standard.

Internet Drafts:

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

"A User Agent Configuration Mechanism For Multimedia Mail Format Information", 06/18/1991, Nathaniel Borenstein <draft-ietf-borenstein-configmech-04.txt, .ps>

"Mnemonic Character Sets", 07/08/1991, Keld Simonsen <draft-ietf-822ext-mnemonics-03.txt>

"Character Mnemonics and Character Sets", 07/08/1991, Keld Simonsen <draft-ietf-822ext-charsets-04.txt>

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

3.1.5 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.

Aug 1991 Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting.

Dec 1991 Continue revisions based on comments received at meeting and e-mail. Start making document an Internet Draft.

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

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

Internet Drafts:

"Network Database Protocol", 06/26/1991, Daisy Shen <draft-ietf-netdata-netdata-02.txt>

"Network Database Implementation Information", 12/16/1991, Daisy Shen <draft-ietf-netdata-implement-01.txt>

3.1.6 Network Fax (netfax)

Charter

Chair(s):

Mark Needleman, mhn@stubbs.ucop.edu

Mailing Lists:

General Discussion: netfax@stubbs.ucop.edu
To Subscribe: netfax-request@stubbs.ucop.edu
Archive: /pub/netfax@stubbs.ucop.edu

Description of Working Group:

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

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

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

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

Goals and Milestones:

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

Done — Continue work on definition of issues and protocols. Work to be conducted on mailing list.

Aug 1991 — First draft of RFC to be completed. To be discussed at IETF meeting and revised as necessary.

Dec 1991 Continue revisions based on comments received and submit to IESG for publication as RFC.

Mar 1992 Overlapping with activities listed above may be implementations based on ideas and work done by the Working Group. If so revise RFC to include knowledge gained from such implementations.

Request For Comments:

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

CURRENT MEETING REPORT

Reported by Mark Needleman/U California

Minutes of the Network Fax Working Group (NETFAX)

The IETF Netfax Working Group met at the San Diego IETF meeting. The primary purpose of the meeting was to get a progress report and plan future work.

Mark Needleman reported that the file format document has been accepted by the IESG which in turn has recommended to the IAB that it be accepted as a Proposed Standard. There is every indication that this will happen.

Mark also reported on the effort to convene a meeting of all known groups doing applications involving networked fax, for the purpose of talking about converting to the RFC file format and planning interoperability testing. This meeting would be sponsored by the Coalition for Networked Information (CNI). Anyone knowing of organizations who are involved in this area and is encouraged to provide Mark with the organization(s) name. Mark promised to keep the mailing list informed on the status and outcome of this effort.

Mark also discussed the fact that after two years of discussion and work no consensus had developed on any of the other issues the Group had been working on beyond the file format, such as addressing and routing issues. He suggested that it did not seem likely that such consensus was going to develop any time soon. For that reason, it was decided that, the Group would quiesce for awhile and not meet again until something happened to create such a need. It was felt that perhaps the CNI work would act as a catalyst and that maybe in some of the interoperability testing undertaken by that group, some solutions to some of the unresolved issues such as routing and addressing would be created that could act as a basis for future work in the IETF Working Group.

The Group therefore decided not to meet again until there was a need for it, and to follow the efforts of CNI. Mark promised to keep the Working Group updated on those efforts by means of the Working Group mailing list.

Attendees

Cerafin Castillo	cec@emulex.com
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Pete Grillo	pl0143@mail.psi.net
Jeff Gumpf	jag@po.cwru.edu
Max Hillson	hillson@koala.enet.dec.com
Russ Hobby	rdhobby@ucdavis.edu
Ronald Jacoby	rj@sgi.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Michael Marcinkevicz	mdm@csu.net

Mark Needleman	mhn@stubbs.ucop.edu
Jon Postel	postel@isi.edu
Stephanie Price Marasciullo	price@cmc.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.7 Network News Transport Protocol (nntp)

Charter

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

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

Description of Working Group:

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

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

o User level and site designated authentication methods; o Binary transfer capability; o Minimization of line turnaround; and o Stronger article selection capability.

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

Goals and Milestones:

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

Internet Drafts:

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

CURRENT MEETING REPORT

Reported by Eliot Lear/Silicon Graphics

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

The IETF-NNTP Working Group met three times in San Diego to walk through the existing NNTP v2 draft and get it out the door. All changes to the NNTP document have been made, and after several formatting changes are made a new version will be put out for comments.

Agenda

- Discussion of Security Issues in the NNTP Architecture
- Use of formats in NNTP
- Document walk-through
- News Reader Issues
- Action Items

Authentication Issues

During the first meeting, we discussed the current security mechanism. It was believed that AUTHINFO was still not general enough for sites to implement certain types of authentication. The conclusion was to essentially hand over the TCP stream to a mutually agreed upon authentication system, and take it back when it is done.

Seven/Eight Bit Issues

After much wrangling on the topic of 7/8 bit, it was decided that the 7-bit restriction on NNTP should be removed. Commands must still be sent with the high order bit cleared, but data may contain octets with the high order bit set. In fact this is existing practice of the most common servers. The BINARY format has been removed until such time as someone can define a message format for it (see below). The IMAGE format has been renamed to PREARRANGED, to heighten the point that the option should only be used by consenting partners.

Document Walkthrough Highlights

There were a bunch of minor changes and clarifications. The error codes section has been reworded slightly. Specifically, certain response codes should be properly processed any time a response code is expected, either for debugging purposes (09x codes) or certain other error conditions that may occur at any time (e.g., new authentication required).

A new VERSION command has been added and required so that each side can determine the software being used by the other. The syntax of this command allows for comments at the end of either the command or the response. The expectation is that version information about particular implementations will be collected.

The Connect sequence has been clarified, and discussion has been moved from Section 3 to Section 2.

Continuation characters a'la SMTP and FTP have been allowed. This documents existing practice in most implementations. It was thought that this would be useful to communicate site specific connection information to the other side of a connection.

The LIST command has been restructured to return the same information it gave for version 1. This was done because the change brought up more problems than it solved, and the extensions were primarily for news readers.

The NEWGROUPS command is deprecated in favor of LIST.

The option command has been changed so that a possible return is UNIMPLEMENTED. This is for non-standard options that are asked of unwitting servers.

The BATCH option has been changed so that no articles greater than the agreed upon batch size may be transferred. To transfer larger articles the other side must first turn batch mode off.

The SIMPLE authentication mechanism has been reworked to fit the changed authentication model.

A new appendix is being added on implementation issues. Currently there are numerous implementations that exacerbate the worst features of the current protocol. For example, opening a connection, offering one article, not sending it, and closing the connection transfers no data. Data presented at this IETF indicates that this happens a lot.

News Reader Issues

We began discussing news reading issues in the last session, and came up with more questions than answers in that time frame. The following comments were made:

Are we talking about an SQL interface? Should predicates be specified in the protocol? Clearly the user needs some better way to prioritize what gets presented.

Should the protocol be more server event driven than NNTP? Currently the news reader software is responsible for ordering articles. But suppose an article with higher priority hits the server. How should this best be communicated to the user?

Should the protocol BE NNTP with yet more extensions?

Is an RPC interface the ideal? Should the client even have to know that it is going to another machine for its information? If not, are we giving up on .newsrcs?

What about other work in this area? The World-Wide Web (WWW), Archie, and WAIS people are all dealing with similar questions, not to mention every librarian.

Should the protocol be tied to netnews? Should the distinction between netnews and EMail be eliminated, as far as this protocol is concerned? What are the differences?

Also, should multiple remote sources be supported in the protocol? Should there be discovery?

It doesn't take much of an imagination to see that one could easily bloat a protocol. Further discussion on how to limit the scope of a reader protocol ensued.

Action Items

- Get a draft on NNTP out within two weeks.
- Get an implementation of the new version out within four weeks.
- Update Charter.

The former is all but done, the latter two are still being worked on.

Attendees

Jordan Brown	`jbrown@qdeck.com`
Philip Budne	`phil@shiva.com`
Erik Fair	`fair@apple.com`
Olafur Gudmundsson	`ogud@cs.umd.edu`
Ittai Hershman	`ittai@nis.ans.net`
Brian Kantor	`brian@ucsd.edu`
Neil Katin	`katin@eng.sun.com`
Eliot Lear	`lear@sgi.com`
John Leong	`john.leong@andrew.cmu.edu`
David Martin	`dem@fnal.fnal.gov`
Leo McLaughlin	`ljm@wco.ftp.com`
Mel Pleasant	`pleasant@rutgers.edu`
Walter Stickle	`wls@ftp.com`
Theodore Tso	`tytso@mit.edu`
David Waitzman	`djw@bbn.com`
Frank Welch	`frank_welch@ccm.hf.intel.com`

3.1.8 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.

CURRENT MEETING REPORT

Minutes of the Network Printing Protocol Working Group (NPP)

Report not submitted.

Attendees

Charles Bazaar	bazaar@emulex.com
Scott Bradner	sob@harvard.edu
David Bridgham	dab@epilogue.com
Christopher Bucci	bucci@dss.com
Philip Budne	phil@shiva.com
Cerafin Castillo	cec@emulex.com
Nat Howard	nrh@bellcore.com
James Jones	jones@regent.enet.dec.com
Stev Knowles	stev@ftp.com
Leo McLaughlin	ljm@wco.ftp.com
Donald Merritt	don@brl.mil
Paul Raison	raison@xylogics.com
Glenn Trewitt	trewitt@pa.dec.com
Steve Willens	steve@livingston.com

3.1.9 TELNET (telnet)

Charter

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

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

Description of Working Group:

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

- Re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol.
- Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically:
  - Environment variable passing
  - Authentication
  - Encryption
  - Compression
- Act as a clearing-house for all proposed RFCs that deal with the TELNET protocol.

Goals and Milestones:

Done	Write an environment option
Dec 1990	Write an authentication option
Dec 1990	Write an encryption option
Mar 1991	Rewrite RFC 854

Internet Drafts:

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

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

"Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-03.txt>

"Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-03.txt>

"Telnet Environment Option", 03/03/1992, D. Borman <draft-ietf-telnet-environment-02.txt>

"Telnet Remote Flow Control Option", 03/03/1992, C. Hedrick, D. Borman <draft-ietf-telnet-remflow-cntrl-00.txt>

"Telnet Authentication: Kerberos Version 4", 03/03/1992, D. Borman <draft-ietf-telnet-authker-v4-00.txt>

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

Request For Comments:

RFC 1116 "Telnet Linemode option"

RFC 1184 "Telnet Linemode Option"

CURRENT MEETING REPORT

Reported by Steve Alexander/INTERACTIVE Systems

Minutes of the TELNET Working Group (TELNET)

Agenda

- Authentication Option
- Environment Option
- Remote Flow Control Option
- Encryption Option

A discussion of authentication was held. It was decided that the authentication drafts should be put forth as experimental RFCs after some minor editorial changes. Steve Alexander will edit the drafts and send them on to be published. The changes include removing several authentication types from the list in the main document and specifying the Kerberos V4 Authenticator in Section 2 of the Kerberos IV draft.

Dave Borman has made a minor extension to remote flow control to allow any character to restart output. It was decided that Dave will send this document to the IESG on his own, since this has not been discussed by the Working Group in detail and the change is straightforward.

The Environment option was discussed in detail. Dave Borman explained the history behind the VAR and USERVAR separation. The usefulness of variables such as PRINTER was debated. Some Group members raised the issue of tying the environment option to the authentication option; it was decided that they did not need to be tied together. It was decided not to add ACCT/JOB fields to the authentication option. The Group has decided to put the environment option forward as a Proposed Standard after some clarification of VAR & USERVAR in Section V. Steve Alexander will make the changes and submit the document.

No serious discussion of encryption was held. Steve volunteered to begin looking at merging encryption with authentication as discussed in Santa Fe.

Action Items

Steve Alexander

- Revise authentication documents and submit as Experimental.
- Revise environment document and submit as Proposed.
- Look at merging encryption/authentication.

Dave Borman

- Submit remote flow control option as a Proposed Standard.

Attendees

Steve Alexander	stevea@i88.isc.com
Mark Baushke	mdb@nsd.3com.com
David Borman	dab@cray.com
David Carrel	carrel@munin.com
Neil Haller	nmh@bellcore.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Jeff Hayward	j.hayward@utexas.edu
Rajeev Kochhar	rxk@3com.com
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Glenn McGregor	ghm@merit.edu
Marjo Mercado	marjo@hpda.cup.hp.com
Henry Miller	henrym@sacusr.mp.usbr.gov
Paul Raison	raison@xylogics.com
Michael Reilly	reilly@nsl.dec.com
Jonathan Rodin	rodin@ftp.com
Jeffrey Schiller	jis@mit.edu
Sam Sjogren	sjogren@tgv.com
Eric Smith	eric@telebit.com
Chuck Townsend	townsend@ctron.com
Theodore Tso	tytso@mit.edu
Kathleen Wilde	wilde@decvax.dec.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.2 Internet Area

Director(s):

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

Area Summary reported by Greg Vaudreuil /CNRI

There were four Internet Area working groups and two Birds of a Feather sessions during this IETF meeting.

Dynamic Configuration of Network Links BOF (DCNL)

This BOF has been rather long lived. The intent was to receive feedback on a mechanism developed by Cray Research to dynamically set up T3 connections between supercomputer centers. This mechanism has been published as an Experimental RFC. This meeting was spent explaining the basic mechanisms.

Dynamic Host Configuration Working Group (DHC)

The Dynamic Host Configuration Working Group is nearly finished there long-livedeffort to allow automatic configuration of Internet Hosts. A presentation was made to the Plenary session and the protocol is expected to be submitted to the IESG as a Proposed Standard shortly.

IP over Appletalk Working Group (APPLEIP)

The Appletalk Working Group discussed many AppleTalk/Internet Interworking Issues. The Group is finishing a document describing SNMP over Appletalk for environments where UDP is not available. Work is beginning on a Appletalk over X.25, using the mechanisms being reworked for IP over X.25. Appletalk over PPP is in the final stages, and was discussed both in this Working Group and in the Point to Point Working Group. New work is beginning on a protocol for distributing routing information, currently called the AA protocol.

IP over ATM Working Group (ATM)

The IP over ATM Working Group sprung from a BOF held in Santa Fe. This Working Group session was very well attended with 97 people! The Group discussed issued of coordination their work with other standard bodies responsible for defining and profiling ATM. The Group discussed several encapsulation proposals, both the use of the currently defined adaptation layers, and a new proposal to put IP or other network protocols directly over ATM.

IP over HIPPI BOF (HIPPI)

This BOF was called to review a proposal for IP over HIPPI. An Internet Draft has been available. Some minor twiddling was suggested and a new document is expected.

Point to Point Protocol Extensions Working Group (PPPEXT)

The PPPEXT has submitted the bulk of their work to the IESG and IAB. At this meeting the Group focused on defining and narrowing the scope of a PPP MIB. Work is continuing on defining a LCP for IPX, a LCP for OSI, and a LCP for DECnet.

CURRENT MEETING REPORT

Reported by Andy Nicholson/Cray Research

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

The documents describing the work done at Cray Research in this area have been recently published as RFC's, 1306 and 1307.

At this meeting there was an almost total turnover of interested attendees, but there still seems to be a significant level of interest in the topic of Dynamic Control of Links. Andy Nicholson had prepared a suggested scope for work to be done, and was not prepared for a reintroduction to the material, which was necessary.

After the assembled group was brought up to date on the proceeding work, and some alternatives were discussed, there was agreement that Andy's proposal was on the right track. He had put this in the form of a problem statement:

"To develop a standard method of access and control for network links which are available but not normally active within the Internet."

Andy offered to publish all the previous Minutes and references to the existing documents on the mailing list. After everyone has had a chance to refer to this work and begin discussing it, the Group will be able to determine if there is reason to form a Working Group within the IETF.

Philip Almquist, one of the Internet Area Directors, was consulted on this. He pointed out that anything produced would be best documented as experimental, as there is not enough experience with this work in the Internet to suggest the adoption of a Proposed Standard.

Attendees

Richard Basch	basch@mit.edu
Mark Baushke	mdb@nsd.3com.com
Larry Blunk	ljb@merit.edu
David Borman	dab@cray.com
Robert Bosen	bbosen@world.std.com
David Carrel	carrel@munin.com
Barbara Fraser	byf@cert.sei.cmu.edu
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
James Galvin	galvin@tis.com
Terry Gray	gray@cac.washington.edu
Michael Grieves	mgrieves@chipcom.com
Neil Haller	nmh@bellcore.com
Alton Hoover	hoover@nis.ans.net
Thomas Kile	tgkile@brl.mil

John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Kent Malave	kent@chang.austin.ibm.com
Scott Marcus	smarcus@bbn.com
Richard McBride	
Glenn McGregor	ghm@merit.edu
Clifford Neuman	bcn@isi.edu
Andy Nicholson	droid@cray.com
Brad Passwaters	bjp@sura.net
P. Rajaram	rajaram@eng.sun.com
Michael Reilly	reilly@nsl.dec.com
Jim Reinstedler	jimr@ub.com
Gregory Ruth	gruth@bbn.com
Jeffrey Schiller	jis@mit.edu
Tom Simons	nes!simons@phx.mcd.mot.com
William Streilein	bstreile@wellfleet.com
Theodore Tso	tytso@mit.edu
Huyen Vu	vi@polaris.dca.mil
Preston Wilson	preston@i88.isc.com
Steven Wong	wong@took.enet.dec.com

CURRENT MEETING REPORT

Reported by Andy Nicholson/Cray Research

Minutes of the IP over HIPPI BOF (HIPPI)

The purpose of this meeting was to discuss the draft of IP over HIPPI by John Renwick and Andy Nicholson, currently available as an Internet Draft, "draft-renwick-hippilan-01.txt."

The meeting was well attended, including some of the members of the ANSI X3T9.3 Working Group on HIPPI. There was general agreement that the current document is very close to being ready to be promoted as a Proposed Standard.

There was one point of contention, whether it is reasonable to require that IP over HIPPI use the HIPPI LE frame header rather than simply using another HIPPI FP ULP Internet Draft. There was general agreement that the interoperability made possible by using LE and 802 format headers was worth the minor extra cost of the headers.

There were some other suggestions which will be incorporated into another draft which will be distributed in the first week of April.

- The scope should be more clearly stated.
- Relax the requirement that a destination be able to buffer 68 bursts before accepting a connection.
- Don't require that a host know its 48 bit ULA.
- Check the ARP optimization strategy against standards and don't require any optimizations which could propagate incorrect information.
- More clearly state the ARP example.
- Discuss switch address bootstrap issues against possible changes in HIPPI LE or HIPPI SC.
- More discussion of what the HIPPI standards are for people who have not read them.
- Update the HIPPI standards references to their latest numbering.

After this draft goes out for review, if there are no substantive changes or comments on the mailing list, we will offer it to the IESG and the IAB as a Proposed Standard. However, if there is still important discussion, we will hold a Working Group meeting at the Boston IETF.

Attendees

Philip Almquist	almquist@jessica.stanford.edu

Zorica Avramovic	zorica@sparta.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
Tom Benkart	teb@saturn.acc.com
Richard Binder	rbinder@nri.reston.va.us
Randy Butler	rbutler@ncsa.uiuc.edu
Dean Cheng	dean@sun2.retix.com
Richard Cogger	rhx@cornell.cit.bitnet
Bob Doolittle	rad@think.com
Julio Escobar	jescobar@bbn.com
Stefan Fassbender	stf@easi.net
David Getchell	dag@llnl.gov
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Paul Griffiths	griff@chang.austin.ibm.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Randy Hoebelheinrich	rho@rowan.lanl.gov
Holly Knight	holly@apple.com
John Lekashman	lekash@nas.nasa.gov
Tony Li	tli@cisco.com
Kurt Lidl	lidl@sura.net
Matt Mathis	mathis@psc.edu
Milo Medin	medin@nsipo.nasa.gov
Ed Menze	menze@cs.arizona.edu
Donald Morris	morris@ucar.edu
Andy Nicholson	droid@cray.com
L. Michael Sabo	lms@denver.ssds.com
Timothy Salo	tjs@msc.edu
John Seligson	johns@ultra.com
Lansing Sloan	ljsloan@llnl.gov
Frank Solensky	solensky@clearpoint.com
Richard Thomsen	rgt@lanl.gov
Don Tolmie	No net address
Jonathan Whaley	whaley@concert.net
Linda Winkler	lwinkler@anl.gov
Frances Yeh	fyeh@rafael.jpl.nasa.gov
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.2.1 Connection IP (cip)

Charter

Chair(s):
Claudio Topolcic, topolcic@nri.reston.va.us

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

Description of Working Group:

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

Goals and Milestones:

Done Produce a new specification of ST.

Done Define common hardware and software platform.

Done Implement hardware and software platform.

May 1991 Implement experimental modules and perform experiments.

May 1992 Produce a specification of a next generation connection oriented protocol.

Internet Drafts:

"Notes for Application Implementors on ST-II Socket API", 03/11/1992, Charles Lynn <draft-ietf-cip-apisocket-00.txt>

Request For Comments:

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

3.2.2 Dynamic Host Configuration (dhc)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

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

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

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

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

TBD — Write a bootp extensions document.

Internet Drafts:

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

"Dynamic Host Configuration Protocol", 07/09/1991, R. Droms <draft-ietf-dhc-protocol-02.txt, .ps>

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell

Minutes of the Dynamic Host Configuration Working Group (DHC)

The DHC Working Group held two meetings in San Diego, on Tuesday afternoon and Thursday morning. In addition, Ralph Droms gave a technical presentation to the IETF summarizing Dynamic Host Configuration Protocol (DHCP). In the Tuesday meeting, the Working Group discussed changes to DHCP, as reflected in a new version of the DHCP Internet Draft, dated March, 1992. Several minor changes were discussed, which will be integrated into a new version of the DHCP Internet Draft.

Editor's Note (md): A listing of the proposed changes is available via ftp under dhc-minutes-92mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions.

There were also several topics that sparked longer discussions. These topics, in general, remain unresolved:

- The relationship between BOOTP and DHCP should be clarified by renaming DHCP as BOOTP.
- Various levels of compliance with the DHCP/BOOTP specification should be indicated by "BOOTP Level 0," "BOOTP Level 1" and "BOOTP Level 2."

  Discussion: There is the potential for confusion on the part of naive users about the relationship between BOOTP and DHCP. For example, BOOTP clients will continue to work with DHCP servers; will DHCP clients work with BOOTP servers? The relay agents will still be known as "BOOTP relay agents"; will those work with DHCP clients and servers?

  Some client implementations of DHCP may not include all the functions specified in the DHCP Internet Draft. In particular, members of the Working Group have expressed concern that some DHCP clients may be unable to enforce the network address lease rules, and may always require allocation of a network address with an infinite lease.

  The Working Group proposes renaming DHCP to BOOTP, and subsuming the specification of the current BOOTP protocol into the DHCP/BOOTP document. The various type of DHCP/BOOTP service would be names:

  - BOOTP Level 0 – BOOTP as defined in RFC 951.
  - BOOTP Level 1 – DHCP/BOOTP with automatic network address allocation and only infinite leases.
  - BOOTP Level 2 – DHCP/BOOTP with fully dynamic (finite lease) network allocation.

- Does DHCP need to deliver more parameters in vendor extensions than will fit in a single DHCP message?
- If DHCP delivers more parameters than will fit in a single DHCP message, how should those additional parameters be delivered?

  Discussion: Some members of the Working Group contend that DHCP must, indeed, be prepared to deliver more parameters than will fit in a single DHCP message. One argument in favor of that view is that specific sites will want to deliver additional DHCP parameters, and, if the DHCP document does not explicitly specify the use of a particular mechanism, each site may choose a local (and potentially incompatible) mechanism. Other members of the Working Group feel that DHCP can deliver sufficient parameters in a single message to configure a host. Based on the currently defined vendor extensions, and assuming "reasonable" lengths for variable length vendor extensions, DHCP could send all parameters specified by vendor extensions in roughly 280 octets. Counting the "sname" and "file" fields, a single DHCP message can deliver 502 octets of vendor extensions. Thus, DHCP can, at present, deliver all necessary parameters to a host in a single DHCP message. The Working Group discussed a mechanism called a "bill of lading" to provide reliable delivery of vendor extensions in multiple DHCP messages, if DHCP is defined to support vendor extensions in excess of 502 octets. A bill of lading is a bit vector representing all 256 vendor extensions. The DHCP server will deliver a bill of lading to the host, indicating which vendor extensions the host must receive with a 1 in the corresponding bit in the bit vector. The host will then repeatedly ask for DHCP parameters, checking off specific vendor extensions as they arrive, until all specified vendor extensions have been delivered.

- Should a host use DHCP at every reboot to reacquire network parameters?
- If the host cannot contact a DHCP server at reboot, should a host be allowed to reuse previously acquired network parameters?

  1. Required use of DHCP: A DHCP host should likely be required to use DHCP whenever the local network parameters may have changed (e.g., system reboot, network failure and restart), as connected network may change w/o host's or user's knowledge (consider change inside wiring closet w/o user notification).
  2. Application of lease: does the lease apply to just the network address or does it apply to all network parameters? i.e., host is allowed to reuse stored network address (careful here; host may have bogus network address and not be able to contact server).
  3. Authority of DHCP server: should server be able to force host to take new network values; i.e., "remote manage" the host?

     Discussion: The conversation about this topic began with consideration of whether or not a DHCP host should be allowed to continue if it cannot contact a DHCP server at reboot. There were strong arguments for and against, – for: allows for partial network operation even if DHCP servers are inaccessible;

against: may allow misconfigured hosts to operate. Note that only network parameters can be misconfigured; DHCP guarantees that any network address will be assigned to only one DHCP host.

The Working Group then discussed the relation between the "lease" and network parameters; should the lease apply to all parameters or just to the network address. A proposal was floated for another bit-vector to describe those network parameters to which the lease applied. The idea here is to provide some degree of network management through the guarantee that DHCP hosts would periodically reacquire new (possibly changed) network parameters.

Attendees

Steve Alexander	stevea@i88.isc.com
J. Allard	jallard@microsoft.com
David Bridgham	dab@epilogue.com
Andrew Cherenson	arc@sgi.com
Ralph Droms	droms@bucknell.edu
Robert Gilligan	gilligan@eng.sun.com
Tom Hemp	themp@ralvm11.vnet.ibm.com
Nat Howard	nrh@bellcore.com
Hung Huang	huang@ics.uci.edu
Scott Kaplan	scott@wco.ftp.com
Holly Knight	holly@apple.com
Stev Knowles	stev@ftp.com
Rajeev Kochhar	rxk@3com.com
Leo McLaughlin	ljm@wco.ftp.com
Greg Minshall	minshall@wc.novell.com
Robert Morgan	morgan@jessica.stanford.edu
Rex Pugh	pugh@hprnd.rose.hp.com
Kevin Rowett	kevinr@tandem.com
Walter Stickle	wls@ftp.com
William Streilein	bstreile@wellfleet.com
Dave Thompson	daveth@ingate.microsoft.com
Chiaohuey Tung	tung@ashland.intel.com
Kathleen Wilde	wilde@decvax.dec.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Bahman Zargham	zargham_bahman@tandem.com

3.2.3 IP over AppleTalk (appleip)

<u>Charter</u>

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

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

Description of Working Group:

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

Goals and Milestones:

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

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

Internet Drafts:

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

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

"SNMP over AppleTalk", 12/23/1991, G. Minshall, M. Ritter <draft-ietf-appleip-snmp-appletalk-01.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

INTERIM MEETING REPORT

Reported by Mike Ritter/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)

January 15, 1992

PPP/AppleTalk

New drafts are available in the Internet-Drafts Directory, they will also be posted on AppleLink in the Networking Standards Folder. RFC 1171 & 1172 have been rewritten into one document, with additions for authentication and dial-back. Brad Parker has been writing the AppleTalk specifics and has changed the address negotiation so it has a better chance of working. The specifications should be reviewed and commented on by those interested in AppleTalk and PPP.

The following specifications are available in the Internet-Drafts Directory on nic.ddn.mil or nnsc.nsf.net

```
draft-ietf-pppext-appletalk-00.txt
draft-ietf-pppext-authentication-02.txt
draft-ietf-pppext-lcp-02.txt
```

MacIP

John Veizades and Tom Evans are negotiating over the final draft and will have it posted by mid-February.

MIBs

People expressed interest in the following MIBs:

- AppleTalk MIB+ transports and configuration
- Local Talk Repeater
- AT over PPP
- AURP
- ARAP
- MacIP
- AFP
- Printer Server
- DecNet over LT
- Mac System MIB - informational RFC

Anyone interested on working on any of these MIBs should send their names and a proposed Charter to the Apple-IP Mailing list. Peter Caswell and Garth Conboy from Pacer expressed intense interest in working on an AppleTalk Services MIB (Print, FileShare, etc.) Some others expressed an interest in the ARAP Server MIB. Apple is also interested in helping out on these, please contact Mike Ritter at MWRitter@applelink.apple.com and he will try to get the correct parties together.

Karen Frisa from CMU has been volunteered to work on the AppleTalk MIB+ (by Steve Waldbusser). There was general consensus that this MIB should cover all of Inside AppleTalk (except AFP) and that it should try to address the problem of router configuration.

Authentication for the AppleTalk MIB and routers was discussed. It was agreed that any serious attempts should implement the SNMP security protocols. Since these are still in the draft stages, vendors discussed what they are using today. Most routers implement a trusted IP address. For SNMP over AppleTalk it was noted that, due to dynamic node addressing, a trusted node address was unworkable, but that a trusted net number range was basically equivalent to a trusted IP host list.

There is a "connectathon" at Apple which will test inter-operation of AURP, ARAP, and SNMP over AppleTalk. Tools to exercise the implementations are available from Apple. (For SNMP test tools (both over IP and AppleTalk) contact Blee@applelink.apple.com.) Mike Ritter of Apple promised to try to get the source of the SNMP test tools released. InterConn offered demo versions of their management console product that does SNMP over DDP (and runs on a Macintosh). People who have SNMP Agents should talk to InterConn about incorporating their MIBs into the product.

AURP

Alan Oppenheimer reported that AURP router operational experience with several universities was being set up. For AURP test tools contact Oppenheime1@applelink.apple.com. He also has copies of a document for the routing protocol and update state diagram.

Alan has withdrawn the AURP draft from the IETF standards track. He is working on an AURP document with APDA that will be an Apple standard and recommended that an RFC should be written that describes how to run AURP on the Internet. The APDA document will also be published as an informational RFC.

The MIB description for AURP was sent out to the list - please send comments to Alan and the list.

AppleTalk Directorate

John Veizades promised to be an open and unbiased Working Group Chair and professed an ability to keep things confidential if requested. He asked two questions:

1. What would an AppleTalk Directorate be? and
2. What is the process and what concerns do people have?

Frank Slaughter volunteered his thoughts. He was disgruntled because he thought the Group was working on an open standard (AURP) and it turned out to be an Apple standard. He does not want to work on Apple standards, but does want to work on open standards. He said everyone was standardizing on TCP/IP for wide-area connectivity because the IETF process really works. He said the difference between vendor and IETF standards are that an IETF standard has an open forum available to work on it, try it out, and later it's declared officially done by the IAB. A vendor specific standard has the vendor product and the protocol done at the same time. He thought that to make AppleTalk really usable for wide area connectivity and large networks the IETF open standards process must be adhered to.

Greg Minshall, from Novell said that in the IETF this Group is unique in that it is very vendor specific. The only reason there is interest in making these protocols open is because of the number of Macintoshes there are in the market. He suggested that Macintosh-to-local routers as opposed to Router-to-Router would be a logical place to split the ownership of the protocol. He also reminded people that what is done in the Group relates to all Apple products. Don't get too idealistic about how working groups work and the problems that do or do not follow. Greg felt that it was wrong to attempt to do everything in this one Group and recommended that the Group split into separate working groups and have a specific Charter for each topic.

John Veizades pointed out that splitting working groups into other existing Directorates may cause problems because the Directors don't know anything about AppleTalk.

Jonathan Wenocur seconded the idea of splitting up into separate working groups and recommended that when an issue comes up it should be made clear where the future control of a protocol lies - will it be an IETF protocol or will it be an Apple standard?

Frank Slaughter wanted the Group to make this statement: "Things that are done by this Working Group are open standards and that is what we are working towards." There was general consensus that this was correct.

Alan Oppenheimer said that AURP 2.0 could be a standard protocol of the IETF, Apple can put it under the IETF's control.

Frank Slaughter said that he does not feel the need to have Apple's stamp of approval for future work.

John Veizades summarized that the Working Group has carved out a niche that it is willing to work on open standards for AppleTalk solutions (under IETF control) and doesn't require Apple's approval. Apple can be a participant just like any other vendor. There was general consensus that this was correct.

Greg Minshall put forth that the IAB doesn't want to take over Apple protocols that Apple wants to keep proprietary.

Bob Morgan asked if there was really a need for an AppleTalk Directorate that stays around to guide all of the working groups in this area.

Jonathan Wenocour said that a Directorate is important because of AppleTalk integration into IP areas (AURP, KIP, CAP, etc.) The IETF would feel safer with more control over AppleTalk, especially with the possibility of large AppleTalk internets tunneling through the Internet, but there is the potential for conflicts with Apple.

John Veizades pointed to the work being done for AppleTalk over PPP. He said that this Group can do important work without touching on what's inside AppleTalk or other Apple proprietary protocols. He said the Group is wildly enthusiastic about being able to do work in areas that are not Apple proprietary protocols and everyone agreed.

John Veizades recommended that someone should write up a Working Group Charter for the following areas if they were interested: AppleTalk interior (and exterior) routing protocols that scale better than the present one, network management and MIBs, AppleTalk over PPP, AppleTalk over a variety of link layers, wide-area naming, and configuration management and any other areas that people thought needed work on.

Greg Minshall posed the question: Does Apple have an obligation to cede a protocol or tell the Working Group that they are working in the same area? He answered it by saying that it was impossible to expect this. What working groups do are open standards, what Apple does can be Apple standards. In addition he recommended that the working groups have Charters and limit their discussions to the Agenda.

The next meeting will be held in March during the IETF Plenary in San Diego.

Attendees

Debbie Alsop	Alsop.D@applelink.apple.com
Philipe Boulanger	Access.priv@applelink.apple.com
Craig Brenner	Brenner2@applelink.apple.com
Rich Brown	Richard.E.Brown@DARTMOUTH.EDU
Jerome Calvo	transware@Applelink.Apple.Com
Peter Caswell	pfc@pacer.soft.com
Edward Ching	eching@wc.novell.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Garth Conboy	gc@pacer.soft.com
Joy Cordell	J_CORDELL@VMSNET.ENET.DEC.COM
Donald Eastlake	dee@ranger.enet.dec.com
Richard Ford	72510.553@compuserve.com
Karen Frisa	karen.frisa@andrew.cmu.edu
John Gawf	gawf@compatible.com
Bob Jeckell	robert_jeckell@nso.3com.com
Grisha Kotlyar	grisha@farallon.com
Fidelia Kuang	kuang@apple.com
Louise Laier	laierl@applelink.apple.com
Greg Merrell	merrell@greg.enet.dec.com
Greg Minshall	minshall@wc.novell.com

Robert Morgan	morgan@jessica.stanford.edu
Alan Oppenheimer	oppenheimer1@applelink.apple.com
J. Bradford Parker	brad@cayman.com
Christopher Ranch	cranch@novell.com
Michael Ritter	mwritter@applelink.apple.com
Randy Ryals	D2604@applelink.apple.com
John Schafer	john_schafer@um.cc.umich.edu
Mike Schumacher	ms@telebit.com
Frank Slaughter	fgs@shiva.com
Evan Solley	solley@applelink.apple.com
David S.A. Stine	dstine@cisco.com
Michael Swan	neon.sw@applelink.apple.com
Steve Sweeney	steves@farallon.com
Bob Van Andel	BVA@applelink.apple.com
John Veizades	veizades@apple.com
Jonathan Wenocur	jhw@shiva.com

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the IP over AppleTalk Working Group (APPLEIP)

The following subgroups made presentations at the Apple-IP Working Group meeting.

MacIP - Encapsulation of IP Datagrams in AppleTalk Networks

The MacIP Group decided to wait for a new proposal from Tom Evans at Webster Computer Australia to make any further determination on the status of this protocol.

SNMP over AppleTalk

There are currently at least five implementations of SNMP over AppleTalk. The document is final and will be moved to the proposed standard level. Some discussion of this protocol occurred in the OSI area meeting on SNMP over protocols other than IP.

AppleTalk over X.25

The specification of AppleTalk over X.25 will be worked on by Apple Computer. It is suggested that the specification for IP over X.25 be considered.

AppleTalk over PPP

The protocol proposal for AppleTalk over PPP was discussed and various minor changes were made. The document will be resubmitted.

AppleTalk MIB +

The proposal for the follow on AppleTalk MIB was discussed and various issues as to the ability to use this MIB for configuration were raised. There is no clear view that this MIB is any where near final and further discussion of contentious issues will be discussed on the mailing list.

AppleTalk Half Router MIB

The AppleTalk half router MIB is an extension to the MIB + proposal. These issues should be incorporated in the MIB+ document.

AppleTalk File Sharing MIB

The proposal for a AppleTalk File Sharing MIB was presented. The MIB will be posted to the mailing list and comment will be solicited from various AppleTalk File Sharing implementors.

AURP

There were no outstanding issues related to AURP. Field trials of the protocol are continuing.

AA Protocol

Distribution of router configuration information to interested routers. The AA protocol proposal is available for implementation.

There was considerable discussion on the breaking up of the IP over Appletalk Working Group into several working groups. Issues of how this would effect the ability of the various groups to fit within the framework of the IETF is at issue as well as how this Group would relate to the direction and control that Apple Computer would want to retain over the protocols. The representative from Apple will report to the Group as to the relationship between Apple and its developer and user community after these issues are discussed within Apple. The Working Group Chair will work with the IAB and IESG to better understand the relationship between this Group and its governing body. The disposition of these issues will be reported on the mailing list when more is known.

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Bill Bliss	billbl@microsoft.com
Gregory Bruell	gob@shiva.com
Philip Budne	phil@shiva.com
Peter Caswell	pfc@pacer.soft.com
Richard Cherry	rcherry@wc.novell.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Garth Conboy	gc@pacer.soft.com
Karl Fox	karl@morningstar.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Bob Jeckell	rrj@3com.com
Holly Knight	holly@apple.com
Fidelia Kuang	kuang@apple.com
Kenneth Lindahl	lindahl@violet.berkeley.edu
Joshua Littlefield	josh@cayman.com
John Lynn	lynn@jalps2.cit.cornell.edu
Kent Malave	kent@chang.austin.ibm.com
Greg Merrell	merrell@caldec.enet.dec.com
Greg Minshall	minshall@wc.novell.com

Robert Morgan	morgan@jessica.stanford.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Alan Oppenheimer	oppenheime1@applelink.apple.com
J. Bradford Parker	brad@cayman.com
Tushar Patel	tpatel@doelz.temeplex.com
Radia Perlman	perlman@radia.enet.dec.com
Stephanie Price Marasciullo	price@cmc.com
Rex Pugh	pugh@hprnd.rose.hp.com
Christopher Ranch	cranch@novell.com
Sam Roberts	sroberts@farallon.com
William Simpson	bsimpson@vela.acs.oakland.edu
Eric Smith	eric@telebit.com
Evan Solley	solley@farallon.com
Michael Swan	neon.sw@applelink.apple.com
Wayne Tackabury	wft@pacersoft.com
John Veizades	veizades@apple.com
Lee Wade	wade@nsipo.arc.nasa.gov
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com

3.2.4 IP over Asynchronous Transfer Mode (atm)

Charter

Chair(s):

Robert Hinden, hinden@eng.sun.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

Done First Meeting. Establish detailed goals and milestones for Working

Group.

Jan 1992 Circulate drafts of IP over ATM Specifications.

Mar 1992 Review approaches to running IP over ATM.

CURRENT MEETING REPORT

Reported by Bob Hinden/Sun

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

Over 100 people attended the IP over ATM Working Group meetings at the San Diego IETF Meeting. The Group met for three consecutive morning sessions. Topics on the Agenda to be discussed at the meetings included the following:

- Relationship with T1S1
- Encapsulation Proposals
- Addressing Approaches
- Mapping Network Layers to Data Link Layers
- ATM Signaling
- IP/ATM MIB
- Congestion Avoidance/Flow Control
- Type of Service and Quality of Service
- Multicast

The Working Group agreed to initiate forming an official liaison with the ANSI T1S1.5. The Chair will draft a letter that after IESG/IAB approval will be sent to the Chair of T1S1. The Group also discussed forming a relationship with the ATM Forum. The Chair was tasked to investigate this and report back to the Working Group.

Juha Heinanen's encapsulation proposal was discussed at great length. The proposal consisted of several approaches for multiprotocol encapsulation over ATM Adaptation Layer 5 (AAL5). After much detailed discussion the Group agreed that a revised version of the document would be produced that documents the reasonable approaches. This would include three approaches:

1. Virtual Circuit Multiplexing
2. NLPID/SNAP Based Multiplexing
3. LLC/SNAP Based Multiplexing

This document will be published initially as an Internet Draft and then as an Experimental RFC. After we obtain some experience with these approaches, the Group will select one or more for standardization. Another encapsulation approach discussed by the Working Group was to use the BISDN I.cls encapsulation used by SMDS. This will be documented in a separate document.

There was general discussion on addressing and signaling approaches. Several people presented their views and what is available in BISDN. There was also the beginnings of a discussion on what is required for ATM routing. All of these topics will be explored further at future meetings.

Attendees

J. Allard	jallard@microsoft.com
Zorica Avramovic	zorica@sparta.com
Cynthia Bagwell	cbagwell@gateway.mitre.org
John Ballard	jballard@microsoft.com
Tom Benkart	teb@saturn.acc.com
Richard Binder	rbinder@nri.reston.va.us
Joe Blackmon	blackmon@ncsa.uiuc.edu
Bill Bliss	billbl@microsoft.com
Robert Braden	braden@isi.edu
Scott Bradner	sob@harvard.edu
Caralyn Brown	cbrown@wellfleet.com
Theodore Brunner	tob@thumper.bellcore.com
Rick Bubenik	rick@arl.wustl.edu
John Burnett	jlb@adaptive.com
Randy Butler	rbutler@ncsa.uiuc.edu
Dilip Chatwani	dilip@synoptics.com
Dean Cheng	dean@sun2.retix.com
Bilal Chinoy	bac@sdsc.edu
Isaac Choi	ichoi@synoptics.com
George Clapp	clapp@ameris.center.il.ameritech.com
Richard Cogger	rhx@cornell.cit.bitnet
Robert Cole	robert.g.cole@qsun.att.com
Tracy Cox	tacox@sabre.bellcore.com
Bob Doolittle	rad@think.com
Guy Dubuc	guyd@newbridge.com
Bill Edwards	
Julio Escobar	jescobar@bbn.com
Bill Fardy	fardy@ctron.com
Dino Farinacci	dino@cisco.com
Mark Fox	m_fox@took.enet.dec.com
Cliff Frost	cliff@cmsa.berkeley.edu
Shawn Gallagher	gallagher@quiver.enet.dec.com
Eugene Geer	ewg@nvuxr.bellcore.com
David Getchell	dag@llnl.gov
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Terry Gray	gray@cac.washington.edu
Paul Griffiths	griff@chang.austin.ibm.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Walter Guilarte	70026.1715@compuserve.com
Joel Halpern	jmh@network.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Jeff Hayward	j.hayward@utexas.edu

Kenneth Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Darryl Henry	henryd@imo-uvax.dca.mil
Robert Hinden	hinden@eng.sun.com
Duke Hong	duke@ics.uci.edu
Hung Huang	huang@ics.uci.edu
Ronald Jacoby	rj@sgi.com
Ed Juskevicius	cnv29@bnr.ca
George Kajos	kajos@coral.com
Prakash Kashyap	prakash.kashyap@sun.com
Manu Kaycee	kaycee@ctron.com
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELEMAIL/C=US/@sprint.com
Darren Kinley	kinley@crim.ca
Holly Knight	holly@apple.com
Rajeev Kochhar	rajeev_kochhar@3com.com
Lawerence Lang	llang@cisco.com
Anthony Lauck	lauck@tl.enet.dec.com
Joseph Lawrence	jcl@sabre.bellcore.com
Dory Leifer	del@terminator.um.cc.umich
John Lekashman	lekash@nas.nasa.gov
John Leong	john.leong@andrew.cmu.edu
Mark Lewis	mlewis@telebit.com
Arthur Lin	yalin@srv.pacbell.com
Susan Lin	suelin@ibm.com
Tony Li	tli@cisco.com
Bryan Lyles	lyles@parc.xerox.com
Carl Madison	carl@startek.com
Andrew Malis	malis@bbn.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Ed Menze	menze@cs.arizona.edu
Donald Merritt	don@brl.mil
Henry Miller	henrym@sacusr.mp.usbr.gov
John Moy	jmoy@proteon.com
William Nichols	nichols@took.enet.dec.com
Chan-Ok O'Brien	chan@bilbo.jdssc.dca.mil
Ari Ollikainen	ari@es.net
Ayal Opher	aopher@synoptics.com
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Edison Paw	esp@3com.com
Radia Perlman	perlman@radia.enet.dec.com
Venkat Prasad	vsp@3com.com
Stephanie Price Marasciullo	price@cmc.com
Rex Pugh	pugh@hprnd.rose.hp.com

Scott Quinn	smq@aries.bellcore.com
Paul Raison	raison@xylogics.com
K. Ramakrishnan	rama@erlang.enet.dec.com
Benny Rodrig	4373580@mcimail.com
Allyn Romanow	allyn@eng.sun.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
Allan Rubens	acr@merit.edu
Anthony Rutkowski	amr@sprint.com
L. Michael Sabo	lms@denver.ssds.com
Timothy Salo	tjs@msc.edu
Miguel Sasson	sasson@xylogics.com
Eve Schooler	schooler@isi.edu
John Seligson	johns@ultra.com
Anil Singhal	rnadkarn@cs.ulowell.edu
Keith Sklower	sklower@cs.berkeley.edu
Michael Smith	mes5674@hertz.njit.edu
Michael St. Johns	stjohns@umd5.umd.edu
Martha Steenstrup	msteenst@bbn.com
Scott Stein	scotts@apple.com
Walter Stickle	wls@ftp.com
William Streilein	bstreile@wellfleet.com
Tatsuya Suda	suda@ics.uci.edu
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Glenn Trewitt	trewitt@pa.dec.com
Mike Truskowski	truskowski@cisco.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Chiaohuey Tung	tung@ashland.intel.com
Dono van-Mierop	dono_vanmierop@3mail.3com.com
Justin Walker	justin@apple.com
Yuan Wang	natadm!ycw@uunet.uu.net
James Watt	jamesw@newbridge.com
Jehu Westmark	westmark@dsi.scri.fsu.edu
Jonathan Whaley	whaley@concert.net
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Brian Wyld	brianw@spider.co.uk
Frances Yeh	fyeh@rafael.jpl.nasa.gov
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr
Chin Yuan	cxyuan@pacbell.com
Jessica Yu	jyy@merit.edu

3.2.5 IP over FDDI (fddi)

Charter

Chair(s):
Dave Katz, dkatz@merit.edu

Mailing Lists:
General Discussion: FDDI@merit.edu
To Subscribe: FDDI-request@merit.edu
Archive:

Description of Working Group:

The IP over FDDI Working Group is chartered to create Internet Standards for the use of the Internet Protocol and related protocols on the Fiber Distributed Data Interface (FDDI) medium. This protocol will provide support for the wide variety of FDDI configurations (e.g., dual MAC stations) in such a way as to not constrain their application, while maintaining the architectural philosophy of the Internet protocol suite. The Group will maintain liaison with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards. This Group is specifically not chartered to provide solutions to mixed media bridging problems.

Goals and Milestones:

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

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

Request For Comments:

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

3.2.6 Multi-Media Bridging (mmb)

Charter

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

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

Description of Working Group:

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

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

Goals and Milestones:

Done Finalize Charter of Group.

Aug 1991 Document mulit-media bridging technology and its affect on TCP/IP Internets.

Aug 1991 Document issues to be addressed by Working Group.

3.2.7 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:

"Definitions of Managed Objects for the Point-to-Point Protocol", 09/10/1990, Frank Kastenholz <draft-ietf-pppext-pppmib-02.txt>

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

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

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

"The Point-to-Point Protocol for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links", 07/01/1991, W A Simpson <draft-ietf-pppext-lcp-03.txt>

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

"Proposed AppleTalk Control Protocol (ATCP)", 07/08/1991, Brad Parker <draft-ietf-pppext-appletalk-01.txt>

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

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

<draft-ietf-pppext-authentication-03.txt>

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

Request For Comments:

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

CURRENT MEETING REPORT

Reported Brian Lloyd/Lloyd and Associates

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

IETF PPP

- AppleTalk
- LQM
- MIB
- IPX
- DECnet
- CLNP
- Physical Layer

Brian Lloyd distributed a memo titled "The PPP Internetwork Packet Exchange (IPX) Control Protocol" submitted by Novell. Karl Fox distributed a PPP Pocket Reference card from Morningstar Technologies.

There will be a delay in the issuance of an RFC for LCP, IPCP, and Authentication due to an oversite within the IAB. However, there are not going to be any changes to the drafts prior to them becoming RFCs, so it is safe to implement, and still be in compliance.

Questions regarding IP Address Negotiation. The implementor needs to support old format until PPP becomes a full standard. First check to see if the peer is using the old format. If so, negotiate IP addressing using the old algorithm. This procedure applies until PPP is a full standard. After that, support will not be provided for old algorithms for IP address negotiation. If you do IP you need to go ahead and do the new IP address negotiation scheme.

Each of LCP, IP Control, LQM, and Authentication have their own document.

Brian asked, "How many here are actively working to implement PPP?" Approximately twelve hands went up.

As for penetration in the market it was noted that BARRnet now runs PPP on links to member sites.

AppleTalk

Brad Parker of Cayman presented an updated draft of his AppleTalk over PPP document. There was some feedback from Bill Simpson and Chris Ranch. The document was forwarded to the IETF drafts mailing list. It received a good review from AppleTalk community, supports ARAP. Will support router to router half routing. The document will be placed on Merit.edu and Angband.stanford.edu in addition to the usual places.

Link Quality Monitoring (LQM)

The previous version of LQM was not widely implemented so major changes were deemed acceptable (this choice was made at the Santa Fe IETF meeting). As a result the general mechanism was redefined and should be able to determine if a synchronous link is up. Flow control monitoring is not recommended for asynchronous links. LQM is useful for high speed point-to-point links between router vendors. LQM will give continuous information on the state of the link. This is good for OSPF type link state relative protocols.

PPP MIB

Frank Kastenholz of Clearpoint updated the MIB for PPP. Discussion has been open on the mailing list. Frank presented an update. PPP is a complex protocol so the MIB grew to almost 200 variables. Frank says this MIB has to be trimmed down, but others are asking for more. This MIB doesn't even address AppleTalk, DECnet, or CLNP.

It was asked whether this MIB should cover NCPs.

Frank drew on the overhead. There were four columns: Protocol, Mandatory, Conditional Mandatory (if you are trying to control PPP instead of just monitor), and Optional. This graph allowed the members of the Working Group to assign each variable to a category.

One reason to have lots of MIB variables is the need to configure PPP in routers via SNMP (the router from NAT was used as an example since it is only manageable via SNMP). It was suggested that all configuration variables be in the optional column and get rid of the Conditional Mandatory column.

Discussion continued as to how necessary it might be to trim down the variables. It was determined that MIB variables present for debugging purposes be discarded. Brian requested that Frank Kastenholz, Bill Simpson, and Glenn McGregor meet to pare down the MIB prior to the next day's Working Group meeting.

IPX

Christopher Ranch of Novell took the floor to lead the discussion of IPX over PPP. The Novell NCP has no options and this is in conflict with what Shiva has proposed. Brian recommended that Novell and Shiva hammer out the differences and produce a single unifying document. The Working Group indicated that they wanted to see an address negotiation and a compression option added to Novell's proposal. Brian also asked Chris to consider adding negotiation of a routing protocol IF he thinks it would be useful.

DECnet

There appeared to be no progress in the area of DECnet over PPP. Further work on this subject is awaiting an implementation and/or a new draft document.

OSI/CLNP

Bill opened discussion with the remark that there is a well-written document for which there are no implementations. This is a no-option document that differentiated between the three different kinds of CLNP. This will be re-addressed when there is an implementation. Christopher Ranch will forward requests on this to the correct person at Novell who is beginning an implementation.

Bridging

Fred Baker is looking for implementation experiences to document. 3-COM has done bridging over PPP. Currently the document needs to:

1. Clarify the concept of a virtual ring, and
2. Tighten up the language.

32-BIT FCS

Bill Simpson stated the issues with 32-bit FCS. These being that DEC owns patents on a procedure of combining the 32bit and 16bit FCS into a 48bit FCS to be used while 32bit FCS is being negotiated. Noel Chiappa said that DEC will make a license to their process freely available. DEC will provide a general grant of right to use the technology and will provide a letter to the IETF stating so.

Action Item: Karl Fox of Morningstar Technologies, (a vendor company with an implementation), is going to take the task of getting the letter from DEC releasing the rights to the process to the world.

Physical Layer

Where/how to handle the physical layer information. The PPP mailing list concluded that the LCP layer is not the place. Bill Simpson stated that PPP is supposed to run over anything; in other words if you have two wires you should be able to run PPP. Brian Lloyd suggested the need for an implementation reference. There was agreement to this. Someone said this should be an informational RFC. Items to be covered included: PPP SYNC interface with an eye towards RS232, V35, V36, RS422/RS449; async implementations; switched circuits, i.e., Hayes compatible, X21, V25bis dialing; and definition of physical layer up/down determination; etc. Questions presented:

- How are we going to deal with ISDN? This is a topic of future discussion and work with the IPLPDN Working Group.
- Chat scripts and dealing with a login sequence on an async link. What is the other end going to send?

MIB Revisted

Frank took the floor to revisit the issue of MIB variables. Together with Bill Simpson, Glenn McGregor, and some input from Jeff Case they got the number of variables to just over a hundred. This is down from 196. They did not deal with every section, and some still need to be added for AppleTalk, and IPX. It will be necessary to know if and what will be monitored in IPX over PPP.

Changes: link extensions table is gone, FSM table(s) are gone, these were deemed to be debugging information. It was decided that it made more sense to return the link quality reports as a single aggregate MIB variable instead of permitting each field withing the LQR to be queried separately. Individual variables in the LQR are not very useful by themselves plus in order to make sense of the timely information it is necessary to see a complete "snapshot" in one operation.

On the philosophy of configurable parameters: the choice seems to be, a rich set of "knobs" or allowing the vendor to completely control the initial and desired state of the implementation. There was no hard-and-fast decision so it was left up to Frank to clean up what was decided and to post all changes to the MIBs to the mailing list in a few weeks where discussion will begin anew.

Attendees

Steve Alexander	stevea@i88.isc.com
Syed Ashraf	irfan@pluto.dss.com
James Barnes	barnes@xylogics.com
Larry Blunk	ljb@merit.edu
Robert Bosen	bbosen@world.std.com
Caralyn Brown	cbrown@wellfleet.com
John Denune	jdenune@saturn.sdsu.edu
Constance Fleenor Lloyd	cfleenor@lloyd.com
Karl Fox	karl@morningstar.com
Paul Griffiths	griff@chang.austin.ibm.com
Joel Halpern	jmh@network.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Frank Kastenholz	kasten@europa.clearpoint.com
Tom Kessler	kessler@eng.sun.com
Holly Knight	holly@apple.com
Roman Kochan	roman@doelz.timeplex.com
Mark Lewis	mlewis@telebit.com
Brian Lloyd	brian@lloyd.com
David Martin	dem@fnal.fnal.gov
Glenn McGregor	ghm@merit.edu
Scott Miller	smiller@qualcomm.com
John Moy	jmoy@proteon.com
Alan Oppenheimer	oppenheime1@applelink.apple.com
Mely Pacheco	mely@tdd.sj.nec.com

J. Bradford Parker	brad@cayman.com
Lars Poulsen	lars@cmc.com
Venkat Prasad	vsp@3com.com
Rex Pugh	pugh@hprnd.rose.hp.com
Christopher Ranch	cranch@novell.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
William Simpson	bsimpson@vela.acs.oakland.edu
Michael Smith	mes5674@hertz.njit.edu
Mark Smith	mcs@umich.edu
Pyda Srisuresh	suresh@3com.com
Bill Thomas	jthomas@novell.com
Maurice Turcotte	dnmrt@interlan.com
John Veizades	veizades@apple.com
Vernor Vinge	vinge@sdsu.edu
Huyen Vu	vi@polaris.dca.mil
Yuan Wang	natadm!ycw@uunet.uu.net
Preston Wilson	preston@i88.isc.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.2.8 Router Requirements (rreq)

Charter

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

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

Description of Working Group:

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

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

The purposes of this project include:

- Defining what an IP router does in sufficient detail that routers from different vendors are truly interoperable.
- Providing guidance to vendors, implementors, and purchasers of IP routers.

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

Goals and Milestones:

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

Internet Drafts:

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

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

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

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

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

"IP Forwarding Table MIB", 08/14/1991, Fred Baker <draft-ietf-rreq-forwarding-05.txt>

3.3 Network Management Area

Director(s):

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

Area Summary reported by James Davin/MIT

At the San Diego meeting of the IETF, seven working groups of the Network Management Area held one or more sessions throughout the week. Also, three Birds of a Feather sessions were held. Owing to a brief hiatus in pending MIB reviews, the SNMP Network Management Directorate did not meet.

Chassis MIB Working Group (CHASSIS)

The Chassis MIB Working Group met for the first time at the San Diego IETF meeting. 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 discussed the optional items of its charter. It decided to address the instrumentation of power supplies and other physical properties of a chassis box by attempting to align existing proprietary MIBs in these areas. The co-chair will prepare a distillation of such MIBs for consideration at the next meeting.

The Working Group also decided that development of an "aggregation MIB" for instrumenting aggregate properties of a collection of network elements was worthy of effort but that work in that direction should not be permitted to interfere with the highest priority goal of the Group — to wit, a MIB to represent the mapping of logical network devices onto the physical components of a chassis.

The Working Group heard two presentations about the possible shape of a chassis MIB. Draft documents will be prepared and discussed at the next meeting.

DS1/DS3 MIB Working Group (TRUNKMIB)

The TRUNKMIB Working Group is chartered to formulate any necessary revisions to the DS1 and DS3 transmission MIBs (RFC 1232 and RFC 1233) as these specifications are considered for elevation from Proposed Standard to Draft Standard status. The Group is considering those changes motivated by implementation experience and those motivated by the desire to align with relevant work within ANSI T1M1. The Working Group discussed what changes to the MIB might be desirable in order to support both traditional transmission interfaces as well as proxy representation of discrete CSU/DSU devices. The Working Group also discussed identified conformance groupings of the MIB objects that better reflect existing implementations. A number of issues remain outstanding, and it is hoped that these can be resolved by an email discussion of revised drafts.

Ethernet MIB Working Group (ETHERMIB)

The Ethernet MIB Working Group met to continue its work on conformance issues related to the MIB for Ethernet-like transmission media. Discussion took a giant step forward with the presentation of data on the implementation of objects defined in the current Ethernet MIB. Based on this data, the Group decided to reduce significantly the number of required MIB objects. The Group also worked on documenting the rationale for remaining objects.

Host Resources MIB BOF (HOSTMIB)

A Birds of a Feather session on SNMP instrumentation of host resources was conducted by Steve Waldbusser of CMU. An emphasis in this discussion was the need for commonality among existing proprietary MIBs that instrument various types of resources frequently associated with internet hosts (e.g., processing capacity, memory, disk space). The consensus of those assembled was that internet hosts include not only Unix machines but also various types of personal computers. A number of participants attributed some urgency to the need for commonality in this area, and many felt that forward progress was desirable. Organization of a working group effort with a practical, limited scope and timeframe is likely.

IEEE 802.3 Hub MIB Working Group (HUbmib)

This Working Group met to discuss the current draft of a SNMP MIB for 802.3 Repeater devices. The Group resolved outstanding issues about the treatment of counter totals for repeater implementations whose physical configuration may change dynamically. The penultimate revision of the repeater MIB will be circulated via electronic mail for review by the Working Group members. The Working Group also discussed strategies for developing SNMP MIB instrumentation for 802.3 medium access units based on recent IEEE definitions.

Internet Accounting Working Group (ACCT)

This Working Group met in two sessions during the San Diego IETF meeting during which the efforts of the Group were effectively concluded. The final revision of the Accounting Architecture document will be prepared and reviewed by Working Group members via electronic mail. This document is intended for publication as an Informational RFC. An SNMP MIB to support the accounting architecture was discussed at the meeting. A final revision of this MIB definition will be prepared to reflect comments made at the meeting and will be reviewed via electronic mail. The final MIB text is intended for publication as an Experimental RFC. With the successful completion of its charter, no further meetings of this Working Group are anticipated.

Multiplexing SNMP Agents BOF (MPLXMIB)

A Birds of a Feather session on composite SNMP agent implementations was conducted by Karl Auerbach of Sun Microsystems. The purpose of the session was to discuss the problem of SNMP agent implementations that may be composed of software units from different sources. A comprehensive list of technical issues and constraints was presented and discussed. At least five distinct approaches were identified as appropriate depending on the goals of the implementor, the implementation environment, or other circumstances:

- Strategies based on SNMP proxy using SNMP security mechanisms,
- SMUX (RFC 1227),
- DPI (RFC 1228),
- Application software conventions, and
- Operating system software technologies.

Remote LAN Monitoring Working Group (RMONMIB)

This was both the final session of the RMONMIB Working Group and the first "unofficial" session of the Token Ring Remote Monitoring Working Group. Its Charter is to extend the work begun in the Remote LAN Monitoring MIB (RFC 1271) to the domain of IEEE 802.5 Token Ring media. The Group spent much of its time discussing the first draft of a MIB specification to support token ring monitoring. Among the issues raised was the extent to which the defined instrumentation ought to be oriented towards promiscuous or non-promiscuous models of token ring monitoring. Discussion of a revised version of the specification will continue at the next meeting.

SNMP Agent Description BOF (SNMPAGEN)

This BOF on the utility and desirability of the notation described in RFC 1303, was conducted by Dave Perkins of SynOptics. This notation is designed to document those respects in which a particular SNMP agent implementation may deviate from standardized protocol and MIB specifications. Part of the session was devoted to the identification and discussion of possible problems or improvements to the notation itself. A recurring theme of the session was whether the effect of the proposed notation would be more to document agent "variations" or more to document agent "deficiencies." Opinion was mixed on this issue, and, accordingly, opinion was mixed on the desirability of this or like notations.

X.25 Management Information Base Working Group (X25MIB)

The X.25 MIB Working Group met twice during the San Diego IETF meeting. One session focused on issues of MIB form and structure; another session focused on the X.25 protocol functionality being instrumented. Based on this discussion, a penultimate revision of the LAPB and X.25 Packet Layer MIBs will be circulated on the mailing list for final review and closure. A remaining task for this Working Group is to rework the MIB instrumenting IP over X.25 convergence functions to embrace recent work in the Internet Services Area on multiprotocol use of X.25.

CURRENT MEETING REPORT

Reported by Steven Waldbusser/CMU

Minutes of the Host Resources MIB BOF (HOSTMIB)

Agenda

- Introduce Charter and Mailing list
- Discuss goals
- Identify resources
- Plan for future work
- Straw Proposal

The session opened with the introduction of the proposed Charter. The purpose of the BOF will be to explore interest in consolidating existing private vendor MIBs for SNMP that instrument computer hosts. Such instrumentation represents non-network specific information such as operational state and hardware and software configuration information. This BOF will attempt to outline a plan for future activity in this area.

Because the Internet comprises a wide variety of workstation types, discussion will not be confined to workstations that run Unix. The goal of such an activity would be to improve interoperability in the Internet environment by consolidating the large number of SNMP SMI based objects which have been and are being created by a number of vendors in this area. Some attempt will be made to identify those areas of existing workstation instrumentation in which the desire for commonality is most urgently felt. The mailing list for this Group is: hostmib@andrew.cmu.edu. Addition and removal request should be sent to hostmib-request@andrew.cmu.edu.

The Group discussed the goals for this effort. Several goals outlined in the Charter were reiterated:

- Goal 1: Improve interoperability by consolidating existing vendor MIBs.
- Goal 2: Any objects should be common across platforms (e.g., Unix, DOS, Mac).
- Goal 3: Identify the most urgent areas and solve them quickly. The desire for timeliness, was felt rather strongly by the Group.

A discussion then ensued to attempt to identify the scope of this process. It was noted that there are many problems to solve, some of which are very complex, and others which are more general in nature than instrumentation of hosts.

It was also unclear what comprises a host, but it was noted that focusing on "personal computers" or "workstations" was most important and makes it much easier to find commonality.

It was proposed that the most important instrumentation would cover:

- Operational state/performance
- Hardware configuration
- Software configuration

Furthermore, the focus would be on the following parts of a "host": CPU, memory, disk, devices, and software (OS, drivers, applications).

The Group agreed that a model should be prepared to support this work.

The Group then began identifying resources for this work. Many existing MIBs (and related specifications) were identified including:

- Character MIB
- Unix MIB (Marshall Rose)
- Unix MIB (UTK)
- Unix MIB (U. of Colorado)
- Apple System MIB
- Unix MIB (HP)
- Unix MIB (Dec)
- Microcom LanLord
- Racal Interlan
- Unix MIB (Posix)
- Netware 3rd party Lan inventory products

In addition, in order to judge if there was sufficient interest in moving forward with this work, a show of hands was conducted as to who would participate in a Host Resource MIB Working Group, if formed. A significant number of people (> 30) indicated that they would participate.

The Group achieved consensus on the following plans for the future:

- Charter a Working Group (Action item: Waldbusser)
- Define a Model
- Write a MIB

In the remaining time, a straw proposal was outlined by Pete Grillo. This proposal defined a model and a MIB that instrumented common objects for "Desktop" hosts.

Attendees

Miriam Amos Nihart	`miriam@ltning.enet.dec.com`
Karl Auerbach	`karl@eng.sun.com`
James Barnes	`barnes@xylogics.com`
Ken Boggs	`boggs@ralvm12.vnet.ibm.com`
Steve Bostock	`steveb@novell.com`

Bill Bowman	bill@hpprsd.mayfield.hp.com
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Carson Cheung	carson@bnr.ca
Dave Cullerot	cullerot@ctron.com
Dave Dalton	dave_dalton@ccm.hf.intel.com
James Davin	jrd@ptt.lcs.mit.edu
Kathleen Dodd	kathy@gateway.mitre.org
Bob Doolittle	rad@think.com
Raymond Edgerton	bellmf!edgerton@uwm.edu
Michael Erlinger	mike@jarthur.claremont.edu
Michael Erlinger	mike@lexcel.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Rob Graham	rob@protools.com
Pete Grillo	pl0143@mail.psi.net
Tom Hemp	themp@ralvm11.vnet.ibm.com
Jon Infante	jon@rv.icl.com
Ronald Jacoby	rj@sgi.com
Yigal Jacoby	jacoby@Tekelec.com
Scott Kaplan	scott@wco.ftp.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Kenneth Lindahl	lindahl@violet.berkeley.edu
David Lindemulder	dcl@mtung.att.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Rohit Mital	rm@protools.com
Lynn Monsanto	monsanto@sun.com
Robert Morgan	morgan@jessica.stanford.edu
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Chandy Nilakantan	csn@3com.com
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
Mike Ritter	mwritter@applelink.apple.com
Marshall Rose	mrose@dbc.mtview.ca.us
L. Michael Sabo	lms@denver.ssds.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
John Seligson	johns@ultra.com
Surinder Singh	singh@irv.icl.com
Anil Singhal	rnadkarn@cs.ulowell.edu
Mark Sleeper	mws@sparta.com

Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Wayne Tackabury	wft@pacersoft.com
Chris Thomas	chris_thomas@ccm.hf.intel.com
Dean Throop	throop@dg-rtp.dg.com
Chiaohuey Tung	tung@ashland.intel.com
Maurice Turcotte	dnmrt@interlan.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Timothy Walden	tmwalden@saturn.acc.com
Frank Welch	frank_welch@ccm.hf.intel.com
Bert Wijnen	wijnen@vnet.ibm.com
Steven Wong	wong@took.enet.dec.com
Henry Yip	natadm!henry@uunet.uu.net

CURRENT MEETING REPORT

Reported by Ed Alcoff/The Wollongong Group

Minutes of the Multiplexing SNMP Agents BOF (MPLXMIB)

The meeting focused primarily on two issues, Multiplexing SNMP Agents and a Straw Man proposal. The purpose of the BOF is to determine:

1. Whether SMUX or DPI are adequate.
2. Whether the proxy capabilities in secure SNMP are adequate.
3. If neither (1) nor (2) are adequate, why not? In other words, what sorts of functions do users think they need from a multiplexing SNMP agent?
4. What sort of solutions might exist?
   - Are the solutions limited to the Unix operating system?
   - Are the solutions generalized to cover Unix and other environments?

Goal: The goal of this BOF is to make an inquiry regarding the scope of the issue and the range of potential solutions.

Discussion on the Straw Man Proposal covered the following nine points.

1. New MIB sub-trees may be attached *and* detached at any time.
2. Sub-trees may not be nested. In other words, an attached sub-tree may not have dynamically attached lower level sub-trees.
3. The master agent must not be required to have a priori knowledge of what is in the subtrees. (In other words, that agent ought to be able to accept any arbitrary subtree.)
4. Get-next must work across subtree boundaries.
5. Get-requests and Get-next requests must be allowed to have varBind entries which refer to elements in multiple subtrees. In other words, a single get/getnext ought to be able to fetch stuff in multiple sub-trees.
6. A set-request ought to be able to contain varBind entries which refer to elements in multiple subtrees. And the "as if simultaneous" requirement must be preserved across subtrees. This one we might want to debate.
7. The multiplexing scheme must be robust despite sub-agent failures. (i.e., a request ought not hang forever if a sub-agent is non-responsive.)
8. The multiplexing scheme need only support sub-agents on the same machine as the primary agent.

9. Multi-agent/multi-protocol capability: The multiplexing protocol ought to be designed so that a given subordinate agent could support multiple superior agents. In addition, the protocol ought to be rich enough that those superior agents aren't just SNMP agents.

Discussion

1. No one strongly disputed the need for dynamic growth or contraction of the MIB tree.

2. It was pointed out that one may want to have different instances of a variable or group "owned" by a different sub-agents. For example, each row of a table may be provided by a separate sub-agent rather than having the entire table provided by a single sub-agent. This appears to be a useful point of view. It is, however, significantly at odds with previous thinking in which all instances of any given variable would be "owned" by a single sub-agent.

3. A model was proposed in which there would be multiple, independent sub-agents rather than a single master multiplexing SNMP agent. In this new model, each agent would emit a start-up trap announcing its service address. A management station would then address independent SNMP queries to the appropriate agent.

4. Problem with both SMUX & DPI was noted: Because of the uncoordinated activities of the various sub-agents, the correlation of sysUpTime with sub-agent derived information may be weak and may vary unpredictably. It may be difficult, if not impossible, to provide a useful correlation between time stamps (such as sysUpTime) and readings of management variables.

5. A concern was raised whether effective network management requires that a management station be able to issue an SNMP varBindList which has items spanning MIB subtrees owned by separate sub-agents. One participant asked whether this was a non-issue. In particular, does it really matter whether a query is split among agents (or sub-agents) by the SNMP manager station itself or by a master agent? In further discussion, it was suggested that since the agent is "closer" to the sub-agents it is in a better position to know how to best partition queries. Partitioning by the agent also has the advantage that get-next semantics can be preserved even where the next lexicographic item lies across a sub-agent boundary.

   Another participant commented that whenever a MIB is partitioned among sub-agents, it is necessary to replicate at least the System group in each partition. Thus it would be possible to retrieve timestamps which are correlated with the data values.

   A question was raised whether the sysUpTime value of the various partitions need to be synchronized. A well known pragmatist answered that on most hosts, it is quite easy to keep the various instances of sysUpTime fairly well synchronized. This left unanswered the question whether such synchronization is needed when the sub-agents reside on separate processors.

6. Looking to practice, do people actually issue queries which deal with logically separate

information bases (each of which presumably would be handled by a separate sub-agent)? On one hand would one ever realistically want to ask about routing tables and sendmail in the same SNMP query? Probably not. On the other hand, one could conceive of a query which tried to correlate network traffic load with changes in routing topology. And it is not unreasonable to believe that load measurement and routing topology would be maintained in separate sub-agents.

It was pointed out that many SNMP managers do not recognize the notion that a single managed device may contain multiple SNMP entities. Consequently, many managers today present such devices on the user interface as if they were multiple, separate devices.

7. It was asked to what extent existing multiplexed SNMP agents enforce the "as if simultaneous" atomic requirements of SNMP Set-Requests. It appears that a significant number of existing multiplexed agents do not make this guarantee. This has not, to date, appeared to have caused any operational difficulties. However, this may be the result of simplistic user interfaces which limit set-requests to one value, proprietary MIBs which are designed so as to avoid the need for atomic "sets", or to adolescent tools which do not yet push Set-Requests very hard.

8. There was discussion regarding the implementation burden on sub-agent writers. At a minimum, there was a desire to avoid encoding and decoding ASN.1/BER. Alternatives suggested were XDR or simplified BER. If the agent-sub-agent interface did not cross machine boundaries then one could even use internal, host specific data formats.

   Some people wanted to go further and isolate sub-agents from the issues of object naming, object instancing, and lexicographic ordering. It was hoped that the agent-to-sub-agent interface could be hidden behind a clean programming API. There was no consensus, however, whether this is feasible unless done in the context of a specific operating system.

9. DEC, HP, IBM, and Peer Networks quickly described their own methods of dealing with some or all of the issues.

10. Marshall Rose described a means using the Secure SNMP protocols and MIB to partition the management variable space among multiple sub-agents. Each "party" would be mapped to a separate sub-agent. It was pointed out that this is really a variation of the "completely disjoint agent" method #3, above.

Summary

There is strong interest in multiplexing SNMP agents. A number of multiplexing agents or extensible agents have been constructed. Various attendees have built multi-protocol agents and managers.

The requirements are not yet well understood. In particular, a significant number of attendees were of the opinion that it is not necessary to preserve full SNMP query semantics across sub-agent boundaries or that it is acceptable for an agent to fail to honor the "as it simultaneous" and atomic properties of SET requests.

Using the proxy facility of the forthcoming secure SNMP would easily and directly provide a means to divide MIBs among separate sub-agents. But it would require that a management station be aware of the MIB partions.

Attendees

Edward Alcoff	oldera@twg.com
Steve Alexander	stevea@i88.isc.com
Miriam Amos Nihart	miriam@ltning.enet.dec.com
Karl Auerbach	karl@empirical.com
John Ballard	jballard@microsoft.com
James Barnes	barnes@xylogics.com
Amatzia Ben-Artzi	amatzia@netmanage.com
Ken Boggs	boggs@ralvm12.vnet.ibm.com
Steve Bostock	steveb@novell.com
Bill Bowman	bill@hpprsd.mayfield.hp.com
David Bridgham	dab@epilogue.com
Theodore Brunner	tob@thumper.bellcore.com
Niels Brunsgaard	nob@dowtyns.dk
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Carson Cheung	carson@bnr.ca
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Steve Deering	deering@xerox.com
Kathleen Dodd	kathy@gateway.mitre.org
Billy Durham	durham@fishery.honeywell.com
Gary Ellis	garye@hpspd.spd.hp.com
David Engel	david@ods.com
Michael Erlinger	mike@lexcel.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Shawn Gallagher	gallagher@quiver.enet.dec.com
Kenneth Goodwin	goodwin@psc.edu
Michael Grieves	mgrieves@chipcom.com
Walter Guilarte	70026.1715@compuserve.com
Tom Hemp	themp@ralvm11.vnet.ibm.com
Ronald Jacoby	rj@sgi.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicom.qc.ca

Yoav Kluger	ykluger@fibhaifa.com
Rajeev Kochhar	rxk@3com.com
Deidre Kostick	dck2@sabre.bellcore.com
Walter Lazear	lazear@gateway.mitre.org
Scott Marcus	smarcus@bbn.com
Evan McGinnis	bem@3com.com
Lynn Monsanto	monsanto@sun.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Chandy Nilakantan	csn@3com.com
Chan-Ok O'Brien	chan@bilbo.jdssc.dca.mil
Ayal Opher	aopher@synoptics.com
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
Jim Reinstedler	jimr@ub.com
Sam Roberts	sroberts@farallon.com
Kary Robertson	kr@concord.com
Jonathan Rodin	rodin@ftp.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Surinder Singh	singh@irv.icl.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Dave Thompson	daveth@ingate.microsoft.com
Dean Throop	throop@dg-rtp.dg.com
Chuck Townsend	townsend@ctron.com
Maurice Turcotte	dnmrt@interlan.com
Dono van-Mierop	dono_vanmierop@3mail.3com.com
David Waitzman	djw@bbn.com
Timothy Walden	tmwalden@saturn.acc.com
Jehu Westmark	westmark@dsi.scri.fsu.edu
Bert Wijnen	wijnen@vnet.ibm.com
Preston Wilson	preston@i88.isc.com
Steven Wong	wong@took.enet.dec.com
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Frances Yeh	fyeh@rafael.jpl.nasa.gov
Henry Yip	natadm!henry@uunet.uu.net

CURRENT MEETING REPORT

Reported by David Perkins/SynOptics

Minutes of the SNMP Agent Description BOF (SNMPAGEN)

A Birds of a Feather session was held on Tuesday to present the contents of the recently published RFC1303 and to gather feedback from the attendees. Leading the presentation was David Perkins from SynOptics with assistance from Marshall Rose from Dover Beach Consulting, one of the joint authors of the RFC.

RFC 1303, an informational document, proposes a method to describe the MIB objects implemented by an SNMP agent using an ASN.1 macro. This concise and precise technique provides a crisp communication between agent implementors and users of an agent.

The set of overheads that were presented are included in the Proceedings. The overheads included a description of the proposal, uses of the proposal, problems with the proposal, and suggested changes.

There were close to eighty people in attendance with a good mix of agent technology vendors, managed device vendors (i.e., those who ship devices with agents), network management system vendors, and users of network management systems.

The feedback to the presentation was generally that the proposal in RFC1303 was timely and pretty much on the mark. However, most people found the proposal to have the problems as presented in the BOF. The problems ranked in order of importance were: 1) missing support for traps; 2) problems in applying to SMUX type agents; 2) incomplete definition of syntax refinements; 3) document problems with row creation and deletion; 4) fuzzy definition of MIB groups; and 5) potential for a large number of macro instances.

The attendees were urged to use the macro defined by the proposal to describe some agents to determine if there was a good fit and report the results back to the authors. There was some interest in forming a Working Group and taking the proposal on the standards track. This action was not the intent of the BOF. Further consultation with the authors, the Network Management Director, and interested parties was needed to determine a course of action.

Attendees

Edward Alcoff	oldera@twg.com
Miriam Amos Nihart	miriam@ltning.enet.dec.com
Michael Anello	mike@xlnt.com
Fred Baker	fbaker@emerald.acc.com
John Ballard	jballard@microsoft.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com

Bill Bowman	bill@hpprsd.mayfield.hp.com
David Bridgham	dab@epilogue.com
Niels Brunsgaard	nob@dowtyns.dk
Lida Carrier	lida@apple.com
Peter Caswell	pfc@pacer.soft.com
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
Richard Cherry	rcherry@wc.novell.com
Carson Cheung	carson@bnr.ca
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Kathleen Dodd	kathy@gateway.mitre.org
Gary Ellis	garye@hpspd.spd.hp.com
Michael Erlinger	mike@lexcel.com
Bill Fardy	fardy@ctron.com
Stefan Fassbender	stf@easi.net
Jeff Fried	jmf@relay.proteon.com
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Shawn Gallagher	gallagher@quiver.enet.dec.com
Kenneth Goodwin	goodwin@psc.edu
Michael Grieves	mgrieves@chipcom.com
Pete Grillo	pl0143@mail.psi.net
Tom Hemp	themp@ralvm11.vnet.ibm.com
Jeffrey Honig	jch@risci.cit.cornell.edu
William Jackson	jackson@manta.nosc.mil
Ronald Jacoby	rj@sgi.com
Bob Jeckell	rrj@3com.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Yoav Kluger	ykluger@fibhaifa.com
Rajeev Kochhar	rxk@3com.com
Deidre Kostick	dck2@sabre.bellcore.com
John Lynn	lynn@jalps2.cit.cornell.edu
Carl Madison	carl@startek.com
Kent Malave	kent@chang.austin.ibm.com
Richard McBride	
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Lynn Monsanto	monsanto@sun.com
Robert Morgan	morgan@jessica.stanford.edu
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
David Preston	djp@epilogue.com

Jim Reinstedler	jimr@ub.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Erik Sherk	sherk@sura.net
Surinder Singh	singh@irv.icl.com
Anil Singhal	rnadkarn@cs.ulowell.edu
Timon Sloane	peernet!timon@uunet.uu.net
Brad Steinka	brad@python.eng.microcom.com
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
Steve Thompson	steve@ods.com
Dean Throop	throop@dg-rtp.dg.com
Glenn Trewitt	trewitt@pa.dec.com
Maurice Turcotte	dnmrt@interlan.com
David Waitzman	djw@bbn.com
Jonathan Wenocur	jhw@shiva.com
Bert Wijnen	wijnen@vnet.ibm.com
Steven Wong	wong@took.enet.dec.com
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Brian Wyld	brianw@spider.co.uk

Agent Description BOF

A review of

RFC 1303: A Convention for Describing SNMP-based Agents

David T. Perkins, SynOptics
Marshall T. Rose, Dover Beach Consulting
17-mar-92

Goals of Proposal

Goals are:

- To describe with great precision the MIB objects implemented by an Agent; and
- Define a method for a management application to determine those implemented objects.

For use by:

- Agent implementors;
- Management applications;
- Managed device description literature; and
- Users of managed devices.

Current Agent Descriptions

Ad hoc - question and answer, incomplete

User: What MIB(s) does your device implement?
Vendor: We support MIB-2.
User: All groups, including EGP and TCP?
Vendor: Well, our device is an 802.3 repeater so we only support system, interface, at, ip, icmp, udp, and snmp groups.
User: Do you support writing to ifAdminStatus?
Vendor: Sorry, we don't.
User: Do you support creation of ipNetToMedia table rows?
Vendor: Sorry, we don't do that. But we are MIB-2 conformant.

Lots of work to find out an incorrect answer!

Where Descriptions are Used

Interactions between Groups – Need effective communications

- User asks for features ==> product marketing
- Product marketing decides on features ==> agent implementor
- Agent implementor describes limitations ==> NMS application implementor, Product marketing
- Product marketing describes features ==> product user
- user configures NMS ==> NMS applications use description

New Method for Agent Description

In RFC 1303:

A new ASN.1 Macro called MODULE-CONFORMANCE is defined.

For an Agent, it specifies implemented:

- MIB MODULES;
- Groups within a Module;
- Variant objects; and
- Tables that allow row creation.

A management application determines the agent's implemented MIB objects by inspecting the value of object sysObjectID.

New Method vs Old Method

Old:
- Ad hoc
- Imprecise
- Incomplete
- No standard format

New:
- Standard parsable format
- High precision
- Concise
- Complete (mostly)

Description Used By...

- Agent implementors
  - within an organization or 3rd party contractors
  - can be used to generate agent MIB data structures
- Testing organizations
  - within an organization or 3rd party
  - can be used to generate test scripts
- Vendors of managed devices
  - part of product description
- Management station application implementors
  - decisions based on optional and mandatory groups
  - checks for implementation of required objects
  - checks for limitations of key objects

Description Used By, continued...

- Operators of management stations
  - to know what an agent implements, so can choose application to use
  - determine if deviation is a bug or a limitation
- Purchasers of managed devices
  - what MIB objects and exceptions or extensions are in the managed device
  - easy basis for choice
- Managed device reviewers
  - easy to determine management features
  - easy to report and summarize comparisons

Management Station Use

- Computer readable agent descriptions are loaded into network management station (NMS) database - (which can be ASCII file/SQL database/proprietary binary).
- Computer readable MIB definitions areloaded into NMS database - (which can be ASCII file/SQL database/proprietary binary).
- Information about managed devices such as name, network address, security parties and keys, and agent description loaded into NMS database.
- Applications when run use NMS databases to determine if they can perform the desired task by comparing needed MIB objects with those in agent description.

Network Management Station Use

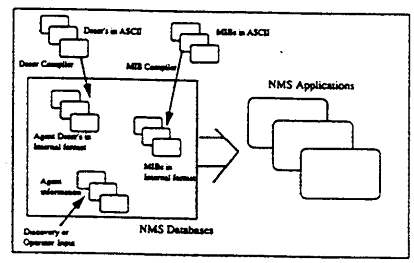

Network Management Station

Agent Description Format

- Uses ASN.1 Macro
- One or more macro instances per ASN.1 MODULE
- Macro definition given in RFC1303

Macro Syntax

Syntax of Macro in BNF format:

```
moduleConformance =
    descrName MODULE-CONFORMANCE
    "LAST-UPDATED" utcTime
    "PRODUCT-RELEASE" releaseDescr
    "DESCRIPTION" agentDescr
    [ SUPPORTS moduleName
        "INCLUDES" "{" groupName ["," groupName]... "}"
        [ variation]... ]...
    ::= oidValue
```

Where:

descrName - name of instance of Conformance macro
utcTime - quoted time in YYMMDDhhmmZ format
releaseDescr - quoted description of the product release
agentDescr - quoted description of the agent
moduleName - name of a MIB module
groupName - name of a group within a MIB module

Macro Syntax, continued...

```
variation =
    "VARIATION" objectName
    [ "SYNTAX" syntaxVariation ]
    [ "WRITE-SYNTAX" writeSyntaxVariation ]
    [ "ACCESS" accessVariation ]
    [ "CREATION-REQUIRES" "{" colName ["," colName]... "}" ]
    [ "DEFVAL" defvalVariation ]
    "DESCRIPTION" variationDescr
```

Where:

- objectName - name of an object within a MIB group
- syntaxVariation - refined syntax for the object
- writeSyntaxVariation - refined syntax for write operations
- accessVariation - refined access for an object
- colName - name of column required for row creation
- defvalVariation - refined DEFVAL value
- variationDescr - quoted description of the variation or refinement

Macro Example

```
exampleAgent MODULE-CONFORMANCE
   LAST-UPDATED        "9202061223Z"
   PRODUCT-RELEASE     "version 4.1"
   DESCRIPTION         "an example agent"

   SUPPORTS            RFC1213-MIB    -- MIB-II
      INCLUDES         { system, interfaces, at, ip, icmp,
                           udp, snmp }

      VARIATION        ifAdminStatus
         SYNTAX        INTEGER { up(1), down(2) }
         DESCRIPTION   "Unable to set test mode"

      VARIATION        ipInAddrErrors
         ACCESS        not-accessible
         DESCRIPTION   "Information not available"

   ::= { agentProfile 23 }
```


Variation Examples

```
VARIATION ipDefaultTTL
   SYNTAX        INTEGER { maxttl(255) }
   DESCRIPTION   "Hard-coded to a specific value"

VARIATION ipRouteType
   SYNTAX        INTEGER { direct(3), indirect(4) }
   WRITE-SYNTAX  INTEGER { invalid(2), direct(3), indirect(4) }
   DESCRIPTION   "Values limited"

VARIATION tcpConnState
   ACCESS        read-only
   DESCRIPTION   "Unable to to set this"

VARIATION ipNetToMediaEntry
   CREATION-REQUIRES { ipNetToMediaPhysAddress }
   DESCRIPTION   "Implementation limited to 25 entries"
```


More Examples

Marshall Rose with more examples...

Reactions to Proposal

- Timely - can use today
- Important - solves some major problems
- Only a few problems/controversial items in proposal

Recommendations

- Write agent descriptions on a small number of agents
  - determine if proposal is a good fit
- Consider how proposal would be used by network management station platform and applications.
- Widen scope of people to review document
- Send suggestions and questions to authors

Controversial Issues

- Problems/Issues:
  - missing support for Traps
  - refinements for "syntax" not completely specified
  - can objects be augmented
  - row creation description has problems
  - missing definition for row deletion
  - groups are not well defined
  - can not handle duplicate object names
  - large of number of MODULE-CONFORMANCE instances
  - overloading of value for sysObjectID

Issues/Problems Detailed

- Traps not supported
  - Traps are part of the definition of an Agent.
  - Need to specify which Traps are generated by an Agent.
  - Need mechanism to define trap variations, that is, which variables are not returned and which additional variables are returned with a trap.

Issues/Problems Detailed, continued...

- Refinements for "syntax" not completely specified
  - As currently defined, there is a only given a list of "not allowed" syntax refinements.
  - What are the allowed refinements other than those given in the examples?
  - Is the size of an OCTET STRING allowed to be shortened?
  - Is the range of an INTEGER allowed to be specified?
  - Why are Gauges, NetworkAddresses, and Opaque syntax types not restricted from being refined?

Issues/Problems Detailed, continued...

- Can objects be augmented, not just refined?
  - Proposal has rules only for refining objects
  - Are there valid ways to augment syntax?
    * more enumerated values
    * longer OCTET STRINGs
    * greater ranges for INTEGERs
    * semantics extended
  - Can ACCESS of read-only be augmented to read-write?
  - Can DEFVAL be changed if one is given in MIB?

Issues/Problems Detailed, continued...

- Row creation has "bugs" in specification
  - ASN.1 macro has missing "empty" clause
  - Description of clause and examples of macro don't match.
  - Easy to fix

Issues/Problems Detailed, continued...

- Can row deletion clause be added?
  - Older MIBs may have "strange" methods for row deletion.
  - Newer MIB objects typically use "status" column. If so, then what value to use.

Agent Description

Issues/Problems Detailed, continued...

- Groups are not well defined
  - No formal definition of a "group" in SMI.
  - Obvious for simple MIB organizations.
  - What are the definitions for "complex" or "twisted" MIBs?
  - Examples:

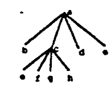

Group within a group

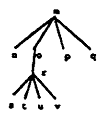

Table within a group

Agent Description

Issues/Problems Detailed, continued...

- As currently written, different objects that have the same textual name can not be specified.
  - a not common case, but still legal in proprietary MIBs.

Agent Description

Issues/Problems Detailed, continued...

- Large of number of MODULE-CONFORMANCE instances:
  - becomes very apparent when a managed device can have different configurations - new instance for each.
  - any difference means a new instance even though the change is very small.
  - current mechanism requires compile/LINK time knowledge of agent's objects.
  - special mechanisms needed for runtime definition of agent's objects.
  - Example: four collections become seven instances

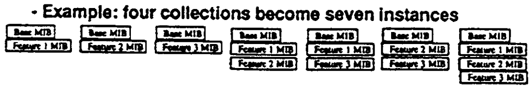

Agent Description

Issues/Problems Detailed, continued...

- The value for object sysObjectID is overloaded
  - sysObjectID is already used.
  - adds to confusion of definition of sysObjectID.
  - sysObjectID definition:

"The vendor's authoritative identification of the ~~network management subsystem~~ [i.e., Agent] contained in the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining 'what kind of box ' is being managed. For example, if vendor 'Flintstones, Inc.' was assigned the subtree 1.3.6.1.4.1.4242, it could assign the identifier 1.3.6.1.4.1.4242.1.1 to its 'Fred Router'."

... is the agent being identified or the hardware?

Agent Description

Proposed Changes

- Solution for Trap variations:
  - Add "TRAP" clause that specifies included traps.
  - Add "TRAP-VARIATION" clause with sub-clauses "OMITS" and "ADDS" to indicate a list of omitted or added variables to a trap.
  - Syntax in BNF would be:

```
trapVariation ::=
    "TRAP-VARIATION" trapName
        ["ADDS"     "{" objectName ["," objectName]... "}"]
        ["OMITS"    "{" objectName ["," objectName]... "}"]
        "DESCRIPTION" variationDescr
```

Agent Description

Proposed Changes, continued...

- Solution for complete definition of refinements to "syntax" and for augmenting objects:
  - Disallow refinements to Gauges, NetworkAddresses, and Opaque syntax types.
  - Add list of allowed refinements including:
    - restriction of values for enumerated integers
    - restriction of size of varying length OCTET STRINGs
    - restriction of range of INTEGERs
  - Add list of augmentations
    - access changes, where appropriate
    - no new enum values
    - size/range change
    - DEFVAL change

Agent Description

Proposed Changes, continued...

- Solution for row creation/deletion:
  - Fix syntax and examples in current proposal
  - Add new deletion clause that specifies pair(s) of (object,value)

Agent Description

Proposed Changes, continued...

- Solution for group definition:
  - define groups
  - modify and extend examples to show groups

Agent Description

Proposed Changes, continued...

- Solution for duplicate object names
  - define to not use IMPORTs
  - use "module.object" name for variables in traps

Agent Description

Proposed Changes, continued...

- Solution for LARGE number of MODULE-SUPPORT instances:
  - Allow multiple "MIB Descriptions" per agent.
  - Don't use object sysObjectID to specify the "MIB description".
  - Create a new table of "MIB Descriptions" in the SNMP group. Columns would be: index of row in table, and conformance id which is an OID of conformance macro instance.
  - Remove the "PRODUCT-RELEASE" clause or change its meaning since it no longer describes an agent.

CURRENT MEETING REPORT

Reported by Larry J. Blunk/Merit

Minutes of the Terminal Server Accounting and Authentication BOF (TERMACCT)

Editors Note (md): These Minutes were inadvertantly omitted from the Santa Fe Proceedings. We wish to extend our apologies to the Chair and to the members of the IETF for the oversight.

The meeting began with a presentation of the Authentication, Authorization, and Accounting services currently provided in Merit terminal servers, and how these features are lacking in commercial terminal server equipment.

Authentication was discussed and there seemed to be a consensus that Kerberos would be the way to go. There was some question about whether terminal servers with limited resources would be able to implement Kerberos (such as, how much ROM would it take?).

Authorization was mentioned as being a difficult issue. Kerberos V5 has hooks for authorization, but currently provides no definitions. OSF DCE apparently provides some authorization capabilities using Kerberos V5, but it is not clear how suitable it would be for terminal servers.

Accounting and billing issues were discussed among which was the need to define accounting and billing variables. There also may need to be interaction between the authorization and accounting systems (to deny authorization for someone who has exceeded a usage quota, for example). It was mentioned that the cost, in resources and real dollars, of accounting needs to be weighed against the actual value of the service.

There was much interest in the notion of a "connection manager" which could provide a common or customizable user interface. Such a manager would be run on a host machine and would likely interact with the authentication, authorization, and accounting services.

The consensus of the BOF participants seemed to be that Merit should come up with a requirements document for further discussion. It could then be determined whether a new working group should be formed.

Attendees

Steve Alexander	`stevea@i88.isc.com`
James Barnes	`barnes@xylogics.com`
Charles Bazaar	`bazaar@emulex.com`
Larry Blunk	`ljb@merit.edu`
Philip Budne	`phil@shiva.com`
Stephen Crocker	`crocker@tis.com`

Erik Fair	fair@apple.com
Barbara Fraser	byf@cert.sei.cmu.edu
Jack Hahn	hahn@sura.net
Neil Haller	nmh@bellcore.com
Ellen Hoffman	esh@merit.edu
Borka Jerman-Blazic	jerman-blazic@ijs.ac.mail.yu
Charles Kaufman	kaufman@dsmail.enet.dec.com
Mark Knopper	mak@merit.edu
Kenneth Laube	laube@bbn.com
Mark Lewis	mlewis@telebit.com
John Linn	linn@zendia.enet.dec.com
Brian Lloyd	brian@ray.lloyd.com
Glenn McGregor	ghm@merit.edu
Dave Monachello	dave@pluto.dss.com
J. Bradford Parker	brad@cayman.com
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Allan Rubens	acr@merit.edu
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
William Simpson	Bill_Simpson@um.cc.umich.edu
Frank Slaughter	fgs@shiva.com
Eric Smith	eric@telebit.com
Richard Smith	smiddy@pluto.dss.com
Michael St. Johns	stjohns@umd5.umd.edu
Theodore Tso	tytso@mit.edu
Wengyik Yeong	yeongw@psi.com

CURRENT MEETING REPORT

Reported by Larry J. Blunk/Merit

Minutes of the Terminal Server Accounting and Authentication BOF (TERMACCT)

Discussion began with the distinguishing features of a Network Access Server (NAS). The concept of a NAS is considered to be an abstraction. For example, a Unix host with async ports could very well be considered a NAS. The difference between a NAS and a router is the notion of session based services which can be authenticated and authorized.

It was questioned whether the Authentication, Authorization, and Accounting (AAA) servers would be running as separate servers or perhaps in the NAS itself. Again AAA servers were viewed as a logical abstraction. The AAA servers could indeed be separate or in fact all run on the same machine. Mention was made of the possibility of providing for interdomain AAA services. Some thought that this should be of primary concern in the design process. The DNS was used as example of a hierarchical domain of servers.

Propagation of authentication information was discussed. It would be desirable to not have to re-authenticate the user for each service requested.

There were questions asked concerning how Kerberos could be used as the authentication mechanism. While it would work fine for dumb terminals and PPP's PAP protocol, PPP's CHAP protocol presents difficulties.

There was discussion of authorization and how configuration parameters are retrieved. Authorization needs to be kept distinct from configuration. Authorization information could be retrieved using a query and response mechanism or all at once. This is an implementation issue.

The purpose of a NAS working group was discussed. Should it define the necessary standards, or use a liaison structure (similar to the Security Working Group)? While authentication and accounting are currently being addressed, there are no groups currently working on authorization. This is a big issue. A NAS working group could specify NAS specific authorization, but it would be desirable to make it extensible rather than limit it to NAS use only. Some discussion was given to providing a mechanism for a common user interface. It was generally agreed that this would be outside the scope of the group.

There was some speculation that the requirements for dumb terminal access and framed serial line services differed substantially enough to warrant independent sub-groups. However, there were many who thought that there was enough common overlap to require a single group. The name NAAAG was suggested as possible acronym for the group.

The consensus of the BOF was that a NAS working group is needed and that the requirements document needs to be further refined. It was also mentioned that those areas outside the scope of the working group should be defined. There is also need for communication and coordination with existing working groups.

Attendees

Serge Aumont	aumont@cicb.fr
James Barnes	barnes@xylogics.com
Charles Bazaar	bazaar@emulex.com
Larry Blunk	ljb@merit.edu
Christopher Bucci	bucci@dss.com
David Carrel	carrel@munin.com
Cerafin Castillo	cec@emulex.com
Henry Clark	henryc@oar.net
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Karl Fox	karl@morningstar.com
Maria Gallagher	maria@nsipo.nasa.gov
Jeff Hayward	j.hayward@utexas.edu
Ittai Hershman	ittai@nis.ans.net
Frank Kastenholz	kasten@europa.clearpoint.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Rajeev Kochhar	rxk@3com.com
Linda Leibengood	ldl@ans.net
Mark Lewis	mlewis@telebit.com
Hock-Koon Lim	lim@po.cwru.edu
Brian Lloyd	brian@lloyd.com
Bill Manning	bmanning@rice.edu
Scott Marcus	smarcus@bbn.com
Glenn McGregor	ghm@merit.edu
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Stephanie Price Marasciullo	price@cmc.com
Paul Raison	raison@xylogics.com
Jim Reinstedler	jimr@ub.com
Robert Reschly	reschly@brl.mil
Allan Rubens	acr@merit.edu
Miguel Sasson	sasson@xylogics.com
Brian Shiflett	rcollett@icmi.icp.net
William Simpson	bsimpson@vela.acs.oakland.edu
Bob Stewart	rlstewart@eng.xyplex.com
Carol Ward	cward@westnet.net
Steve Willens	steve@livingston.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Steven Wong	wong@took.enet.dec.com

3.3.1 Chassis MIB (chassis)

Charter

Chair(s):
Bob Stewart, `rlstewart@eng.xyplex.com`
Jeffrey Case, `case@cs.utk.edu`

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

Description of Working Group:

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

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

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

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

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

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

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

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

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

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

Minutes of the Chassis MIB Working Group (CHASSIS)

Agenda

- Welcome
- Introductions
- Review Charter
- Chassis Information Model
- Define Scope of Work
- Review Contributed MIBs
  - Keith McCloghrie, Hughes LAN Systems
  - Donna McMaster, Synoptics
  - Manu Kaycee, Cabletron
- Plan For Producing Baseline Documents
- Next Meeting Date and Location

The Group discussed the Charter, concentrating on the three work areas: mapping logical functions to physical devices, power supplies, and aggregation. This discussion was limited to the meaning of the Charter with technical discussion deferred to later in the meeting.

The major points regarding the Charter for logical to physical mapping (the Chassis MIB, proper) were:

- This is a "meta-MIB", pointing to other MIBs.
- Multiple instances of the same device may have "virtual agents."
- Any system in any slot may implement the Chassis MIB.
- One agent may point all slots to the same agent.

The major points regarding the Charter for a power supply MIB were:

- "Power supply" may include environment, such as fans and temperature.
- "Power supply" most likely does not include items such as interrupt vectors and memory jumpers.
- Environment perhaps should be a separate MIB.
- Discussion should stay focussed on a "network" chassis, not general VME, Multibus, PC bus or such.

Little was said at this point regarding aggregation.

Regarding the general constraints, the major points were:

- This is NOT the place for a VME MIB.
- Large companies, such as IBM, are not considered as standards bodies.
- For the sake of robustness, reliance on third parties is to be avoided.

The Charter was accepted as written.

The Group discussed the scope of work for a Power Supply MIB. The major points were:

- Many are interested having such a MIB, few are interested in working on it.
- A document is not useful if it does not result in widespread implementations.
- A poll of what current implementations provide obtained: state, backup, voltage, current, etcetera ad nauseum.
- A Power Supply MIB might point to an Environment MIB.
- A Power Supply MIB is applicable outside a chassis, but a power supply in a chassis is more important than in a single system.
- What is available across implementations resulted in a consensus for on/off status and an average of 4/5 variables from about 25 companies.
- Who is actually using this information resulted in responses from Hughes, Digital, Synoptics, Chipcom, Fibronics, NCR, and several others.
- Everyone who has such variables is to send relevant MIB segments to Bob Stewart for compilation, including temperature, fans, and such.

The Group discussed the scope of work for aggregate information. The major points were:

- Widespread confusion over what this topic means concluded that this is to be an assessment of "how is it (a Group of entities) working".
- The answer to "Why with the Chassis MIB?" was because a chassis creates a well-identified Group.
- The Group agreed the definition is still too vague, what constitutes a Group?
- Is there anything like this now? There is some CMIP work which will be one of our proposals, along with two others.
- Examples of aggregation are total errors, total packets, and such.

- Jeff wants to do this.
- Some specific examples are traffic in and out of a regional network, or the sum of ifInOctets for a Group.
- Might this solve the problem of SNMP management for non- SNMP devices not necessarily in a chassis? No.
- This is an interesting next step in network management, but may be beyond the reasonable scope of this Working Group.
- Synoptics has a MIB for the health of a device, set by rules.
- The RMON MIB totals things, but is LAN-bound.
- This looks like artificial intelligence.
- In favor of this work, but shouldn't be called "chassis."
- Does the rule of not defining objects deducible from others preclude this? No, that was a rule for initial definition of MIB I.
- Conclusion was that this is useful, important work, many want to work on it and want the product of that work.

The Working Group concluded that the Chassis MIB is of primary importance, a Power Supply MIB is worth working on if there is enough common ground, and an Aggregation MIB may combine in some ways with the Chassis MIB or may need to be separated so as not to adversely effect delivery of a Chassis MIB, which is the primary deliverable. The Working Group then turned to presentations of possible Chassis MIBs.

Keith McCloghrie presented a proposed Chassis MIB that he and Donna McMaster prepared. The major points presented were:

- Purpose is to manage a box with multiple modules. The box comprises physical modules (slots), logical devices (repeaters, bridges, etc.), backplane "wires" (Ethernet, Token Ring, FDDI, etc.), and power supply.
- Physical devices are indexed by slot number. They have an object identifier for board type (including empty and unknown), and a time of last insertion or removal.
- Logical devices are integer indexed. They have a function (a sum of values such as repeater, bridge, or terminal server), the device's sysObjectId, and, for SNMP access, an SNMP party object identifier or a community string and IP address. Issues here included relationship to SMUX and UPD ports, non-IP addressing, and multiple communities for get and set.
- Backplane wires are integer indexed and have an object identifier to indicate type.
- A configuration table has an entry for each relation with a slot number, a logical

device index, and a backplane wire index, meaning that (part of) the logical device is in the slot and connected to the indicated backplane wire. Several such entries may make up a single logical device.

- Concepts in the document but not on the original slides include a status object and a null index to indicate lack of relevance, such as no backplane wire for a power supply.
- Additional issues include definition of "chassis", generalization of "slot" to include "physical device," more information such as ifIndex or ifOperStatus, inclusion of external "wires," what is a proper device (such as a host), a directory of devices beyond a chassis, and the number of tables (done to be concise).

Manu Kaycee presented a Chassis MIB being implemented as a private extension by Cabletron. The major points presented were:

- Requirements are to support hub-based products, many-to- many associations, logical and physical representations, physical partitioning of components and tables, MIB discovery, multiple component instances, virtual chassis.
- MIB is very similar to Keith and Donna's.
- Lacks map from logical device to backplane wires.
- Adds chassis type for agent-supporting device.
- Backplane includes VME or such.
- Has a slot count.
- Component table includes adminStatus (needs operStatus), string to pass with initialize command, name, software version, access policy.
- Slot table indexed by slot and component to give map. It includes slot class for restricted slots, a unique module ID, and is empty if "chassis" is slotless.
- Includes a (controversial) MIB group table, indexed by slot, component, and Group that can distribute the MIB-II ifTable across slots and could be VERY big.
- This is currently being implemented, Cabletron will share experience if that is helpful.

Jeff called for additional proposals. Two were offered to be submitted by mid April. The first draft MIB is to appear as soon as possible after final call for proposals on mailing list.

The Working Group decided not to have an interim meeting, especially not at INTEROP. Discussion will be via email.

Attendees

Steve Bostock	steveb@novell.com
Niels Brunsgaard	nob@dowtyns.dk
Jeffrey Case	case@cs.utk.edu
Carson Cheung	carson@bnr.ca
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Gary Ellis	garye@hpspd.spd.hp.com
David Engel	david@ods.com
Michael Erlinger	mike@lexcel.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Shawn Gallagher	gallagher@quiver.enet.dec.com
Michael Grieves	mgrieves@chipcom.com
Pete Grillo	pl0143@mail.psi.net
Walter Guilarte	70026.1715@compuserve.com
Tom Hemp	themp@ralvm11.vnet.ibm.com
Ronald Jacoby	rj@sgi.com
George Kajos	kajos@coral.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Yoav Kluger	ykluger@fibhaifa.com
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@synoptics.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
Venkat Prasad	vsp@3com.com
Jim Reinstedler	jimr@ub.com
Sam Roberts	sroberts@farallon.com
Kary Robertson	kr@concord.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Timon Sloane	peernet!timon@uunet.uu.net
Evan Solley	solley@farallon.com
Bob Stewart	rlstewart@eng.xyplex.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com

Geoffrey Thompson	thompson@synoptics.com
Dean Throop	throop@dg-rtp.dg.com
Chuck Townsend	townsend@ctron.com
Maurice Turcotte	dnmrt@interlan.com
Dono van-Mierop	dono_vanmierop@3mail.3com.com
David Waitzman	djw@bbn.com
Timothy Walden	tmwalden@saturn.acc.com
Jay Westmark	westmark@dsi.gw.scri.fsu.edu
Bert Wijnen	wijnen@vnet.ibm.com
Brian Wyld	brianw@spider.co.uk

3.3.2 DS1/DS3 MIB (trunkmib)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Mar 1992 Submit the DS1 document for the Network Management Directorate Review.

Apr 1992 Submit the DS1 MIB to the IESG for Draft Standard Status.

Mar 1992 Submit the DS3 MIB to the Network Management Directorate for review.

Apr 1992 Submit the DS3 MIB to the IESG for approval as a Draft Standard.

Feb 1992 Post a draft version of the new DS1 MIB to the Internet-Drafts Directory.

Feb 1992 Post a revised version of the DS3 MIB to the Internet-Drafts Directory.

CURRENT MEETING REPORT

Reported by Tracy Cox/Bellcore

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

The Group discussed implementation experience of the DS1 and DS3 MIB. The following items resulted from the discussion.

Feedback

- Add Far End Information – as optional tables
- Add more alarm information – ds1AlarmState
- Consistency with standards – updated Terminology section
- Can't "CSU" MIB be used to manage other DS1 interfaces?

Implementation experience was received on RFC1232 and RFC1233. Vendors wanted the definitions of the counters to be consistent with T1M1 standards. The Working Group agreed that the definitions should be updated. However, if the documents should conflict, vendors should follow the definitions in the Internet DS1 and DS3 MIBs. Text was added to the Internet Drafts to reflect this consensus.

Next, the need for far end information was discussed. Vendors requested that the far end information received from the DS1 and DS3 signal be collected in the MIBs. The Working Group agreed that this information should be added as an optional group. The Working Group agreed to structure the MIBs into two groups, the:

- DS* Near End Group which is mandatory and
- DS* Far End Group which is optional.

The Near End Group contains Configuration, Interval, Current, and Total tables. The Far End Group contains Configuration, Interval, Current, and Total tables.

The Working Group also reviewed the request from a vendor that more configuration information be added to the Configuration Table. The Working Group agreed that this information is important; however, it should not be contained on the SNMP agent on the device. The Network Management Station should have this information in its database. Therefore, the configuration information will not be added to the MIBs. This is only true for the Near End Group. Since, the configuration information from the Far End is received from the incoming signal, the Far End Configuration Table does contain this information. Therefore, the Far End Group does contain configuration information. Only the circuitID object is contained in the Near End Configuration Table.

Based on vendor requests and consistency with T1M1 standard, some objects were deprecated, and new objects were added. This is true for both DS1 and DS3 MIBs.

- ds*Loopback has been deprecated.
- A new object has been added called ds*NewLoopback, which better describes the loopback capabilities of a DS* interface on a device.
- ds*YellowAlarm has been deprecated.
- ds*RedAlarm has been deprecated.
- A new object has been added called ds*LineStatus. This object better describes the status (e.g., alarm state and loopback state) of a DS* interface.
- Only the ds3IntervalCSSs, ds3CurrentCSSs, and ds3TotalCSSs have been deprecated, because these counts are not collected on DS3 interfaces. They are retained in the DS1 MIB.
- Additional objects and status are necessary to fully support E1; NewBridge will supply details, to be edited into the DS1 MIB.

The Internet Draft will reflect these changes.

Also, vendors requested that the DS1 and DS3 MIBs be used to manage devices other than CSUs. Therefore, the MIBs are updated to reflect this request. The MIB manages DS1/DS3 interfaces.

The following objects have been changed to reflect this request:

- ds*CSUIndex has been renamed ds*LineIndex. This object is the identifier of a DS* Interface on a device. If there is at least one ifEntry directly associated with the DS* interface (e.g., if the DS* interface is used to communicate with the Network Layer), it should have the same value as ifIndex. Otherwise, its value should exceed ifNumber.
- ds*Index has been renamed ds*IfIndex. This value for this object is equal to the value of ifIndex from the Interfaces table of MIB II (RFC 1213). The utility of this object is under discussion.

The fractional table was deprecated from the DS1 MIB, because no one implemented it or wanted it.

Since the changes to RFC1232 and RFC1233 were on the borderline of being "too much", the Group agreed to cycle at the Proposed Standard status. This implies that the Working Group will have a longer life cycle than intended, probably on the order of a year.

New Internet Drafts reflecting these changes will be sent to the trunk-mib mailing list and posted in the Internet Drafts directories; when consensus is achieved on the mailing list, they will be forwarded to the IESG.

Attendees

Fred Baker	fbaker@emerald.acc.com
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Tom Hemp	themp@ralvm11.vnet.ibm.com
George Kajos	kajos@coral.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Maurice Turcotte	dnmrt@interlan.com
James Watt	jamesw@newbridge.com

3.3.3 Bridge MIB (bridge)

Charter

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

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

Description of Working Group:

The Bridge MIB Working Group is a subgroup of the SNMP Working Group, and is responsible for providing a set of SNMP/CMOT managed objects which IEEE 802.1 Bridge Vendors can and will implement to allow a workstation to manage a single bridged domain. This set of objects should be largely compliant with (and even draw from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

Goals and Milestones:

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

Request For Comments:

RFC 1286 "Definitions of Managed Objects for Bridges"

3.3.4 Character MIB (charmib)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

Goals and Milestones:

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

Request For Comments:

RFC 1316 "Definitions of Managed Objects for Character Stream Devices"

RFC 1317	"Definitions of Managed Objects for RS-232-like Hardware Devices"
RFC 1318	"Definitions of Managed Objects for Parallel-printer-like Hardware Devices"

3.3.5 DECnet Phase IV MIB (decnetiv)

<u>Charter</u>

Chair(s):

Jonathan Saperia, `saperia@tcpjon.ogo.dec.com`

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

Done	Review and approve the Charter and description of the Working Group, making any necessary changes. At that meeting, the scope of the work will be defined and individual working assignments will be made.
Done	Mailing list discussion of Charter and collection of concerns.
Done	Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. If possible, prototype implementation to begin after revisions have been made.
Done	Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.
Done	Review final draft and if OK, give to IESG for publication as RFC.
Jul 1991	Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

Request For Comments:

RFC 1289	"DECnet Phase IV MIB Extensions"

3.3.6 Ethernet MIB (ethermib)

Charter

Chair(s):

Frank Kastenholz, kasten@europa.clearpoint.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

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

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

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

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

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

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

Goals and Milestones:

TBD	Draft Variable Status Rationale document.
TBD	Develop Implementation Report.

Internet Drafts:

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

"Definitions of Managed Objects for the Ethernet-like Interface Types", 03/24/1992, Frank Kastenholz <draft-ietf-ethermib-objects-00.txt>

CURRENT MEETING REPORT

Reported by Frank Kastenholz/Clearpoint

Minutes of the Ethernet MIB Working Group (ETHERMIB)

The Ethernet MIB Working Group met in San Diego on Tuesday, March 17th.

The Working Group reviewed the report on MIB variable implementation that had earlier been posted to the mailing list (a copy is attached for the record). As a result of this review, the Working Group has decided to make the following changes to the MIB:

1. The dot3TestTdrValue object will be deprecated from the standard mib. There are effectively no implementations of this object, and some chips were reported to return an incorrect value for the TDR count.
2. The dot3StatsInRangeLengthErrors object and the dot3StatsOutOfRangeLengthFields object will be deprecated from the MIB. These objects were not widely implemented and their utility in diagnosing network problems was strongly questioned.
3. In addition to the the dot3InitializeMac object, and the dot3MacSubLayerStatus object, the dot3MulticastReceiveStatus object, and the dot3TxEnabled object will be deprecated from the MIB. These objects were not widely implemented and their utility in diagnosing network problems was strongly questioned.
4. The dot3StatsExcessiveDeferrals object will be deprecated from the MIB. Only one system implemented this object. Furthermore, its exact definition was called into question.
5. The dot3StatsSQETestErrors object received few implementations. However, the Working Group strongly supported its retention in the MIB on the basis that certain forms of transceiver and cable errors that are not uncommon can only be detected with this counter.
6. The collision histogram table (dot3CollTable) will be kept as an optional group, even though the objects are not widely implemented nor is there hardware support on all reported chips.

The implementation data presented at the meeting is:

MIB Variable	Implementation 1	2	3	4	5	6	7	8	9	10	11	Yesses
dot3InitializeMac	C	C	Y	Y	Y	Y	Y	C7	C7	N	Y	6
dot3MacSubLayerStatus	C	C	Y	Y	Y	Y	Y	C7	C7	N	C	5
dot3MulticastReceiveStatus	C	C	Y	C3	Y	C	C	C7	C7	N	C	2

dot3TxEnabled	C	C	Y	Y	Y	Y	Y	C7	C7	N	C	5
dot3TestTdrValue	C	1	C	C4	C	C	C	C4	C4	N	C	1
dot3StatsAlignmentErrors	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	11
dot3StatsFCSErrors	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	11
dot3StatsSingleCollisionFrames	Y	Y	Y	N	Y	Y	Y	Y	Y	Y	Y	10
dot3StatsMultipleCollisionFrames	Y	Y	Y	N	Y	Y	Y	Y	Y	Y	Y	10
dot3StatsSQETestErrors	Y	C	C	C	Y	C	C	C	C	Y	C	3
dot3StatsDeferredTransmissions	Y	C	Y	N	Y	Y	Y	Y	Y	Y	Y	9
dot3StatsLateCollisions	C	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	10
dot3StatsExcessiveCollisions	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	11
dot3StatsInternalMacTransmitErrors	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	11
dot3StatsCarrierSenseErrors	Y	C	Y	Y	Y	Y	Y	Y	Y	Y	Y	10
dot3StatsExcessiveDeferrals	C	C	Y	C	C	C	C	C	C	N	C	1
dot3StatsFrameTooLongs	Y	Y2	Y	Y	C	Y	Y	Y	Y	Y	Y	10
dot3StatsInRangeLengthErrors	C	C	C	N5	C	Y	Y	C	C	N	C	2
dot3StatsOutOfRangeLengthFields	C	C	C	C6	C	C	C	C	C	N	C	0
dot3StatsInternalMacReceiveErrors	Y	Y	Y	Y	Y	C	C	Y	Y	Y	C	8
dot3CollCount	Y	Y	C	N	N	N	N	C	C	N	Y	3
dot3CollFrequencies	Y	Y	C	N	N	N	N	C	C	N	Y	3
Yesses	13	11	16	11	15	14	14	11	11	12	13	

```
Y     Fully implemented, reports a truthful count, or
      indication of state. All values may be written to the
      variable with the expected action occurring.

N     Not implemented at all. Would return a noSuchName error
      if accessed.

C     Implemented but returns a constant value for gets and
      returns a badValue error for any set attempt to set the
      variable to a value other than this constant (writable
      variables only).
```

Notes:

- Does not implement TDR test, but reports TDR from last collision!
- Not supported by the chip, detected solely in software.
- But set to disabled(2) -> badValue
- Underlying TDR function not implemented on this chip.
- Only counts frames too short though.
- Due to Ethernet encapsulation

- Implementation does not support set operations but reports the correct value for these.

Implementation	Vendor	Chip
1	1	Intel 82586
2	1	Fujitsu 86950
3	2	Sonic
4	3	AMD Lance
5	4	National NIC 8390C
6	4	Intel 82596
7	4	AMD Lance
8	5	AMD Lance
9	5	AMD ILACC
10	6	AMD Lance
11	7	Intel 82586

Attendees

James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
Niels Brunsgaard	nob@dowtyns.dk
Christopher Bucci	bucci@dss.com
Philip Budne	phil@shiva.com
Lida Carrier	lida@apple.com
Jeffrey Case	case@cs.utk.edu
Carson Cheung	carson@bnr.ca
Dave Cullerot	cullerot@ctron.com
Gary Ellis	garye@hpspd.spd.hp.com
David Engel	david@ods.com
Bill Fardy	fardy@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Michael Grieves	mgrieves@chipcom.com
Pete Grillo	pl0143@mail.psi.net
Walter Guilarte	70026.1715@compuserve.com
William Jackson	jackson@manta.nosc.mil
Frank Kastenholz	kasten@europa.clearpoint.com
Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
David Lindemulder	dcl@mtung.att.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com

Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Dan Romascanu	dan@lannet.com
Rick Royston	rick@lsumvs.sncc.lsu.edu
Mark Schaefer	schaefer@davidsys.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
Steve Thompson	steve@ods.com
Geoffrey Thompson	thompson@synoptics.com
Dean Throop	throop@dg-rtp.dg.com
Chuck Townsend	townsend@ctron.com
Dono van-Mierop	dono_vanmierop@3mail.3com.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Timothy Walden	tmwalden@saturn.acc.com
Preston Wilson	preston@i88.isc.com
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Henry Yip	natadm!henry@uunet.uu.net

3.3.7 IEEE 802.3 Hub MIB (hubmib)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

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

Internet Drafts:

"Definitions of Managed Objects for IEEE 802.3 Repeater Devices", 07/23/1991, Donna McMaster, Keith McCloghrie <draft-ietf-hubmib-mib-02.txt>

CURRENT MEETING REPORT

Reported by Donna McMaster/SynOptics

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

Agenda

- IEEE 802.3 Repeater Management Report
- Outstanding Issues from Chapter 8 of latest Draft
- Issues from Mailing List
- Any Other Issues on Latest Internet Draft
- Discussion of MAU MIB
- MAU MIB Strategy
- Plans for Progression of Document(s)

A new (March 6) version of the Repeater MIB draft was distributed. This incorporated the text, updated in light of IEEE 802.3 Repeater Management Task Force's February meeting, and was previously distributed to the Working Group's mailing-list.

IEEE Report

Donna presented the following report on the progress of the IEEE 802.3 Repeater Management Task Force (formerly known as "Hub Mgmt TF"):

- The confirmation ballot closed January 31, 1992.
- There were eighty-two comments, primarily requesting clarification.
- At the interim meeting February 24-26, 1992 the section on "Port Functions to Support Management," was rewritten and provided initial resolution for all comments.
- At the IEEE 802 plenary earlier in March, there were some minor tweaks. Results are being sent for a 2nd confirmation ballot and the prognosis is very good.
- The March 6, 1992 draft of the SNMP Repeater MIB is based on output from the interim meeting.
- For the next draft, the editors plan to do "tweaks" from last week's plenary along with other changes that come out of this meeting.
- MAU MIB is now the hot topic.

Geoff Thompson reported on the minor "tweaks" from the IEEE meeting in Irvine the previous week. These edits will be incorporated into the next draft of the Repeater MIB.

Outstanding Issues From Chapter 8

The meaning of the enumerated value notPresent for the MIB object rptrGroupOperState was discussed. It was questioned whether the "has been physically removed" wording used in the document implied that the removal must have occurred since the last reboot. After lengthy discussion and after several votes, a consensus eventually emerged to add more definitive/ instructive text while leaving the enumerations as they were. In order to make progress, two attendees volunteered to draft additional text while the next items were discussed.

The next item was whether rptrPortAutoPartitionState should be combined with rptrPortOperState into a single MIB object. After some discussion, the MIB objects were left as being separate.

Next, discussion took place on the seriousness of autoPartition and the overhead in polling every port's autoPartition state on a regular basis. The Group does not want to issue a trap when this happens. Instead they agreed to add a new repeater-level object "total partitioned ports" with syntax of Gauge. This object will represent the total number of ports in the repeater that are currently enabled and present but autoPartitioned.

On the issue of Total Counters, it was agreed that while a total counter was redundant in the sense that it was a sum of other counters already represented as MIB objects, it was most beneficial in reducing the amount of network traffic, particularly on repeaters with many (e.g., over a hundred) ports.

Two such counters were suggested: one was Total Errors per port, as suggested by the editors in the draft. It was agreed that the errors included in this total would be:

- rptrMonitorPortFCSErrors, rptrMonitorPortAlignmentErrors
- rptrMonitorPortFrameTooLongs, rptrMonitorPortShortEvents
- rptrMonitorPortLateEvents, rptrMonitorPortDataRateMismatches
- and rptrMonitorPortVeryLongEvents

The other total counter was the number of frames across all ports. The difficulty was observed of how this counter would behave when one or more of the ports were removed. A decrease in the counter's value was not consistent with the syntax of Counter. Various suggestions were made, including:

1. Count the total number since the last group (re-)configuration, adding a timestamp to record when that occurred;
2. Add a "virtual port" which would conceptually be a promiscuous monitor on all ports.
3. Have a total counter per group rather than per repeater.

The consensus emerged to have three total counters associated with each group:

- groupTotalFrames
- groupTotalOctets
- groupTotalErrors

On the issue of counting FramesTooLong and VeryLongEvents, the consensus was to align with IEEE, and count them all.

The new text from the two volunteers for rptrGroupOperState was reviewed. The consensus was that either would be acceptable though a slight majority preferred the following text:

> "notPresent(x) indicates that the group is temporarily or permanently physically and/or logically not a part of the repeater. It is an implementation-specific matter as to whether the agent effectively removes notPresent entries from the table".

The Group also agreed to change rptrGroupUpTime to be rptrGroupOperStateLastChange (or some abbreviation of this) with the customary semantics.

Discussion of MAU MIB

A first draft of the MAU MIB was distributed. (This document was also mailed to the hubmib mailing list on Friday, March 13.) Donna presented the following overview of MAU management status and issues:

- 802.3 Medium Attachment Unit (MAU) attaches repeater port or Ethernet- like interface to the local network medium.
- MAU types include 10BASE5 (thick coax), 10BASE2 (thin coax), 10BASE-T (twisted pair), FOIRL and 10BASE-F (fiber optic).
- MIB information includes MAU type, link status, jabbering.
- Discussions in IEEE 802.3 Hub Management Group over past year, postponed MAU work to finish Repeater Management
- The draft proposal was brought to the interim meeting and was well-received. More work done was done at the plenary.
- The draft of the SNMP MAU MIB was mailed to the mailing list last Friday, and was based on output from IEEE 802.3 plenary.
- The Group will discuss how the IETF Hub MIB Working Group wants to handle the MAU MIB.

MAU MIB Objects

1. MAU Type
2. Administrative State (operational, standby, shutdown)
   - Option to implement as read/write, reset MAU
3. Media Available
   - Link status (link integrity/low light) for link media (10BASE-T or 10BASE-F), loopback normal for coax media
   - Lost media counter indicates stability of medium
4. Jabber state, jabbers counter
   - Jabbering (continuous transmission) indicates serious problem in host, not as interesting for repeater ports
5. Jabber trap

MAU MIB Questions

- MAU can attach to repeater port or DTE (Ethernet interface), therefore related to both Repeater MIB and Ethernet-like Itfs MIB
- Most objects are common to both port MAUs and interface MAUs
- Multiple MAUs can be attached to a single port or interface
- How to instantiate? For rptr ports, "group.port.mau" is desirable, for interfaces, "interface.mau". Stay tuned...

MAU MIB Options (none perfect!)

1. Add MAU tables to Repeater and Ethernet-like Interfaces MIBs:
   - MAU table in Repeater MIB indexed "group.port.mau"
   - MAU table in Ethernet-like Ifs MIB indexed "interface.mau"
   - > Destabilizes both drafts, bad timing
2. Create new MAU MIB document with MAU table, indexed 1..n.
   - Add two tables that give mappings from port -> MAU, interface -> MAU.
   - > Awkward instantiation when using MIB browser
3. Create two new MAU MIBs in separate documents (or combine)
   - Repeater MAU MIB with table indexed "group.port.mau"
   - Etherlike Ifs MAU MIB with table indexed "interface.mau"
   - > Duplicates much information

Some discussion. People agreed that the application of the MAU information to Repeaters comes within the Charter of this Working Group. However, it was suggested that we didn't want to slow down the progress of the current Repeater MIB draft, and so the meeting agreed to treat this as a separate MIB document to be produced by the Working Group.

With little time remaining in the meeting, the Group also agreed to deal with MAUs separately for repeaters and for interfaces, but there was no time for any other discussion of the MAU MIB at this meeting. Attendees were encouraged to raise any/all issues on the mailing-list.

Issues From the Mailing-List

The issue of the interaction between rptrPortOperStatus and rptrPortAdminStatus had been raised on the mailing-list since the meeting in Santa Fe. All agreed that they should have the same interaction as MIB-II's ifOperStatus and ifAdminStatus, but there was confusion of ifOperStatus's semantics. Explanation that ifOperStatus was defined to become "down" as soon as possible after ifAdminStatus was set to "down" resolved the confusion.

Any Other Issues

No other issues were raised.

Progression of Documents

The editors were chartered to update the Repeater MIB draft in light of the agreements at this meeting, and to distribute it to the mailing list within two weeks. Thereafter, the Working Group would have two weeks to review it. If no concerns were raised on the mailing-list within the following two weeks (or if all raised concerns were satisfactorily resolved), then the Repeater MIB would be done, and should be forwarded to the IESG with a recommendation for being progressed to Proposed Standard status. Meanwhile, the MAU MIB would be discussed on the mailing-list.

Attendees

James Barnes	`barnes@xylogics.com`
Steve Bostock	`steveb@novell.com`
Jack Brown	`jbrown@huachuca-emh8.army.mil`
Niels Brunsgaard	`nob@dowtyns.dk`
Lida Carrier	`lida@apple.com`
Jeffrey Case	`case@cs.utk.edu`
Carson Cheung	`carson@bnr.ca`
Dave Cullerot	`cullerot@ctron.com`
Billy Durham	`durham@fishery.honeywell.com`
David Engel	`david@ods.com`

Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Shawn Gallagher	gallagher@quiver.enet.dec.com
Michael Grieves	mgrieves@chipcom.com
Walter Guilarte	70026.1715@compuserve.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Yoav Kluger	ykluger@fibhaifa.com
David Lindemulder	dcl@mtung.att.com
Richard McBride	
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Edison Paw	esp@3com.com
Jim Reinstedler	jimr@ub.com
Sam Roberts	sroberts@farallon.com
Dan Romascanu	dan@lannet.com
Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
L. Michael Sabo	lms@denver.ssds.com
Mark Schaefer	schaefer@davidsys.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
Geoffrey Thompson	thompson@synoptics.com
Timothy Walden	tmwalden@saturn.acc.com
Steven Wong	wong@took.enet.dec.com
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Brian Wyld	brianw@spider.co.uk
Henry Yip	natadm!henry@uunet.uu.net

3.3.8 Internet Accounting (acct)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

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

Request For Comments:

RFC 1272 "Internet Accounting: Background"

CURRENT MEETING REPORT

Reported by Gregory Ruth/BBN

Minutes of the Internet Accounting Working Group (ACCT)

Internet Accounting Architecture

The Tuesday session began with a presentation by Jim Alfieri of Bellcore. Jim spoke on Bellcore's plans for SMDS (Switched Multi-megabit Data Service) usage accounting. He mentioned that packet counting was done upon egress from the switch (in our last session, the issue of where to count was raised and the Working Group agreed that the architecture should support both counting on entrance and counting on exit).

Jim explained that the SMDS accounting service does not use the Meter -> collector -> application model, where the collector aggregates frequent meter reports into a usage file. In the SMDS case, the meter (inside the SMDS node) itself does the aggregation of usage data to a file. This file, in AMA TPS (Automatic Message Accounting TeleProcessing System) format, is periodically transferred directly to the billing system application via FTP. A drawback of this architecture is that FTP service provides no security. Further details can be found in Bellcore Technical Advisory TA-TSV-001062, "Generic Requirements for SMDS Customer Network Management Service".

The rest of the session focussed on the most recent version of the Internet Accounting Architecture Document. Working Group participants offered comments, criticisms and suggestions that will be incorporated into the next draft. In particular the following modifications will be made.

- Explain that it is the collector's responsibility to manage time and determine/interpret the meaning of meter ticks.
- Discuss space reclamation strategies for dealing with memory exhaustion in the meter.
- Discuss the problem of trap storms that may occur when the meter is trying to avoid losing data due to memory exhaustion.
- Discuss the possibility of trapping to other/multiple managers.
- Discuss the (in)applicability of the RMON MIB.

We believe that when the Working Group input has been incorporated in the Document and it has been reviewed by the mailing list it will be ready to become an Internet Draft.

Internet Accounting MIB

On Wednesday the Working Group discussed the first real draft of the Internet Accounting MIB. Many detailed (but important) changes were suggested and will be incorporated into the next version of the Architecture Document. Among them:

- Put a general explanatory section up front.
- Add/improve the descriptions of special values for parameters. As a special case, zero values should be used in certain cases (e.g., high water mark or flood mark) to disable a feature.
- We need to add text explaining flows.
- The MIB's "address type" should be expanded to include (a) layer number and (b) type (e.g., if layer = 2, allow all interface designations from the standard MIB).
- Use standard SNMP timeticks (1/100 second).
- Explain what max lifetime is good for.
- Recommend UDP checksums.

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

The Working Group discussed what should be done when the architecture and MIB documents are completed. It was generally felt that implementation of a prototype accounting system was next. This could be done either through a revision/extension of the Working Group's Charter, or by forming a new, successor Working Group.

Attendees

James Alfieri	jdal@troy.cc.bellcore.com
Larry Blunk	ljb@merit.edu
Robert Bosen	bbosen@world.std.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Christopher Bucci	bucci@dss.com
Douglas Carson	carson@utcs.utoronto.ca
John Chang	jrc@uswest.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Henry Clark	henryc@oar.net
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Kathleen Dodd	kathy@gateway.mitre.org
Alisa Hata	hata@cac.washington.edu
Yigal Jacoby	jacoby@Tekelec.com
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELE MAIL/C=US/@sprint.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Hock-Koon Lim	lim@po.cwru.edu
Patrice Lyons	3432266@mcimail.com
Scott Marcus	smarcus@bbn.com
Cyndi Mills	cmills@bbn.com

Clifford Neuman	bcn@isi.edu
Chandy Nilakantan	csn@3com.com
Lars Poulsen	lars@cmc.com
Kary Robertson	kr@concord.com
Gregory Ruth	gruth@bbn.com
Eve Schooler	schooler@isi.edu
Chris Thomas	chris_thomas@ccm.hf.intel.com
David Waitzman	djw@bbn.com
Timothy Walden	tmwalden@saturn.acc.com
Abel Weinrib	abel@bellcore.com
Frank Welch	frank_welch@ccm.hf.intel.com
Chris Wheeler	cwheeler@cac.washington.edu
Allen Wiersma	wiersma@nsipo.nasa.gov
Bert Wijnen	wijnen@vnet.ibm.com

3.3.9 Remote LAN Monitoring (rmonmib)

Charter

Chair(s):

Mike Erlinger, mike@lexcel.com

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discussion and final approval of Charter; discussion and agreement on models and terminology. Make writing assignments.
Done	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1991	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

Internet Drafts:

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

Request For Comments:

RFC 1271 "Remote Network Monitoring Management Information Base"

INTERIM MEETING REPORT

Reported by Mike Erlinger/Lexel

Minutes of the Remote LAN Monitoring Working Group (RMONMIB)

February 5-7, 1992

Activities

The Group spent the three days discussing Token Ring management via a remote monitoring device. Two preliminary MIBs (one from Novell and one from Protools) were used to determine pertinent variables and activities for the Token Ring RMONMIB. Steve Waldbusser was designated as MIB author and took detailed notes. The Group decided that Steve should put his notes into MIB form and release a MIB document to the entire Working Group prior to the San Diego IETF. Once released the document would be put into the Internet-Drafts directory.

The San Diego IETF will be spent reviewing the document resulting from the February meeting.

Attendees

Steve Bostock	steveb@novell.com
Gary Ellis	garye@hpspd.spd.hp.com
Michael Erlinger	mike@lexcel.com
Darrell Furlong	drf@concord.com
Rob Graham	graham@cs.pdx.edu
Steve Grau	steveg@novell.com
To-Choi Lau	tlau@novell.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Brian Wyld	brianw@spider.co.uk
Henry Yip	natadm!henry@uunet.uu.net

CURRENT MEETING REPORT

Reported by Michael Erlinger/Lexcel

Minutes of the Remote LAN Monitoring Working Group (RMONMIB)

The RMON Working Group met in two formal sessions during the San Diego IETF meeting. Issues from both RFC 1271 and the developing Token Ring MIB were discussed. What follows is an encapsulation of the discussions.

Interoperability Interoperation between two or more independent implementations is a requirement for an RFC to become a full Internet standard (RFC 1310). Accordingly, Mike Erlinger and Steve Waldbusser discussed their attempt to organize two interoperability testing sessions for Ethernet RMON probes and managers.

1. Week of May 25th (after InterOp) at CMU in Pittsburgh.
2. Week of Jul 20th (after IETF) at Frontier Software in Boston.

   It was noted that these interoperability activities were being organized outside of the auspices of the IETF and that Mike and Steve were reporting on their progress as part of the IETF Working Group meeting.

The first draft of a Token Ring RMON MIB was created about six-weeks prior to the IETF at a Working Group session in Fullerton, CA., by combining the Novell and ProTools MIBs. This draft was reviewed with the results described below. The goal is for the editor to have an updated draft on the mailing list by the middle of April, and an Internet Draft by the end of April.

- ringOrder Table

  The MIB editor had left out the ring-order table because he believed the information could be easily obtained from the ring station table. The Working Group felt that the table was so useful that it was more than worth the small amount of extra complexity it added. It will be put back in.

- ringStationControl Table

  Text needs to be added to indicate that the ringStationControl Table and the ringStation Table are associated (ifInterface).

- The MAC Address of the active monitor will be added to the ringStationControl Table.

- NAUN

  Needs to be added to the ringOrder Table and the ringStation Table.

- Data Packet

Text needs to be added indicating that a Data Packet is NOT a MAC Packet.

- tokenStats

  remove the words "including those in bad MAC Packets"

- The AllRoutesBroadcastPckts variable was deleted from the basic stats because it already exists in the source routing group.

- host Table

  Define a Station Address as "Station MAC Address minus Source Routing Bits"

- matrix Table

  Make sure to indicate that Station Address (MAC Address) does NOT include Source Routing Bits

- MAC vs LLC

  Make sure that MAC vs. LLC vs. Other is well defined as far as packet types. The Token Ring standard will be checked for the correct wording.

- In/out line errors

  Add burst and in/out line errors (similar to beacons) back into error list in the ringStation Table, since this can improve the ability to correctly identify a problem domain.

There was a lot of discussion about whether the Token Ring MIB should be usable by a non-promiscuous probe. After much debate it was decided that Token Ring was significantly different from Ethernet and that all of the useful fault management and configuration information could be acquired by just examining MAC frames (without the overhead of promiscuous mode). Accordingly, the Working Group decided to repartition the Token Ring MIB into (essentially) four groups:

1. Promiscuous stats (frames, octets, size distribution, etc.)
2. MAC layer stuff (ring station table, ring order table, augmented with MAC layer counts which were previously in the stats group - these would be added to the ring station control table).
3. Ring configuration information that required active gathering methods.
4. Source routing stats

For each entry in the History control table both a promiscuous history and a MAC layer history (with the same parameters) will be generated if appropriate. The Config table will be split into a control table (containing the "push buttons" update-stats and remove-station and a time of last update), and a data table which contains the actual data. All columns must be present in the data table, but the row only comes into existence (or is updated) when the button is actively pushed in the control table.

Attendees

Andy Adams	ala@merit.edu
Ken Boggs	boggs@ralvm12.vnet.ibm.com
Steve Bostock	steveb@novell.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Niels Brunsgaard	nob@dowtyns.dk
Douglas Carson	carson@utcs.utoronto.ca
Jeffrey Case	case@cs.utk.edu
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
Carson Cheung	carson@bnr.ca
Isaac Choi	ichoi@synoptics.com
Dave Cullerot	cullerot@ctron.com
Dave Dalton	dave_dalton@ccm.hf.intel.com
Billy Durham	durham@fishery.honeywell.com
Gary Ellis	garye@hpspd.spd.hp.com
David Engel	david@ods.com
Michael Erlinger	mike@lexcel.com
Jeff Erwin	Jeff_Erwin@protools.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
Karen Frisa	karen.frisa@andrew.cmu.edu
Shawn Gallagher	gallagher@quiver.enet.dec.com
Rob Graham	rob@protools.com
Michael Grieves	mgrieves@chipcom.com
Pete Grillo	pl0143@mail.psi.net
Walter Guilarte	70026.1715@compuserve.com
Daniel Hansen	
Tom Hemp	themp@ralvm11.vnet.ibm.com
Ronald Jacoby	rj@sgi.com
Yigal Jacoby	jacoby@Tekelec.com
Manu Kaycee	kaycee@ctron.com
Mark Kepke	mak@cnd.hp.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Yoav Kluger	ykluger@fibhaifa.com
Tim Lee-Thorp	ngc!tim@uunet.uu.net
David Lindemulder	dcl@mtung.att.com
Carl Madison	carl@startek.com
Scott Marcus	smarcus@bbn.com
Richard McBride	
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Rohit Mital	rm@protools.com

Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Chan-Ok O'Brien	chan@bilbo.jdssc.dca.mil
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
Kary Robertson	kr@concord.com
Jonathan Rodin	rodin@ftp.com
Dan Romascanu	dan@lannet.com
Mark Schaefer	schaefer@davidsys.com
Anil Singhal	rnadkarn@cs.ulowell.edu
Mark Sleeper	mws@sparta.com
Timon Sloane	peernet!timon@uunet.uu.net
Brad Steinka	brad@python.eng.microcom.com
Mark Therieau	markt@python.eng.microcom.com
Chris Thomas	chris_thomas@ccm.hf.intel.com
Steve Thompson	steve@ods.com
Chuck Townsend	townsend@ctron.com
Chiaohuey Tung	tung@ashland.intel.com
Dono van-Mierop	dono_vanmierop@3mail.3com.com
Huyen Vu	vi@polaris.dca.mil
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Timothy Walden	tmwalden@saturn.acc.com
Frank Welch	frank_welch@ccm.hf.intel.com
Jay Westmark	westmark@dsi.gw.scri.fsu.edu
Bert Wijnen	wijnen@vnet.ibm.com
Steven Wong	wong@took.enet.dec.com
Paul Woodruff	biccdn!paulw@eros.uknet.ac.uk
Brian Wyld	brianw@spider.co.uk
Henry Yip	natadm!henry@uunet.uu.net

3.3.10 Token Ring Remote Monitoring (trmon)

Charter

Chair(s):

Michael Erlinger, mike@lexcel.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Feb 1992 Discussion and agreement on models and terminology. Comparison of RMON architecture and Token Ring requirements. Assign author and editor responsibilities.

Mar 1992 Working Group Meeting at IETF. Present and confirm results of February meeting. Develop MIB draft. Publish initial version as Internet Draft.

Jul 1992 Working Group Meeting at IETF to discuss and revise draft of Token Ring Extensions. Publish revised version as Internet Draft.

Nov 1992 Working Group meeting to discuss and reach closure on Token Ring MIB Extensions MIB. Publish agreed version MIB as Internet Draft. Make Working Group recommendation on Token Ring Extensions MIB.

3.3.11 X.25 Management Information Base (x25mib)

Charter

Chair(s):

Dean Throop, `throop@dg-rtp.dg.com`

Mailing Lists:

General Discussion: `x25mib@dg-rtp.dg.com`
To Subscribe: `x25mib-request@dg-rtp.dg.com`
Archive: `dg-rtp.dg.com:x25mib/Current.Mail`

Description of Working Group:

This Working Group will produce a set of three documents that describe the Management Information Base for X.25. The first document will specify the objects for the X.25 Link Layer. The second document will specify the objects for the X.25 Packet Layer. The third document will specify the objects for managing IP over X.25. The Working Group need not consider the Physical Layer because the "Definition of Managed Objects for RS-232-like Hardware Devices" already defines sufficient objects for the Physical Layer of a traditional X.25 stack. Any changes needed at the Physical Layer will be addressed as part of that activity.

The X.25 object definitions will be based on ISO documents 7776 and 8208 however nothing should preclude their use on other similar or interoperable protocols (i.e., implementations based on CCITT specifications).

The objects in the Link and Packet Layer documents, along with the RS-232-like document, should work together to define the objects necessary to manage a traditional X.25 stack. These objects will be independent of any client using the X.25 service. Both of these documents assume the interface table as defined in MIB-II contains entries for the Link and Packet Layer interfaces. Thus these documents will define tables of media specific objects which will have a one to one mapping with interfaces of ifType ddn-x25, rfc877-x25, or lapb. The objects for the IP to X.25 convergence functions will be defined analogously with the ipNetToMedia objects in MIB II.

The Working Group will endeavor to make each layer independent from other layers. The Link Layer will be independent of any Packet Layer protocol above it and should be capable of managing an ISO 7776 (or similar) Link Layer provider serving any client. Likewise the X.25 Packet Layer objects should be independent of the Link Layer below it and should be capable of managing an ISO 8208 (or similar) Packet Layer serving any client.

The Working Group will also produce a third document specifying the objects for managing IP traffic over X.25. These objects will reside in their own table but will be associated with the X.25 interfaces used by IP. These objects will not address policy decisions or other implementation specific operations associated with X.25 connection management decisions except as explicitly described in existing standards. These objects will manage the packet flow between IP

and the X.25 Packet Layer specifically including observation of packet routing and diagnosis of error conditions. Progress on the Link and Packet Layer documents will not depend on progress of the IP over X.25 document. The IP over X.25 document will proceed on a time available basis after work on the Link and Packet Layer documents and as such the Link and Packet Layers may be completed before the IP over X.25 work.

All documents produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions (such as Concise MIB format). To the extent feasible, the object definitions will be consistent with other network management definitions. In particular ISO/IEC CD 10733 will be considered when defining the objects for the X.25 Packet Layer.

Goals and Milestones:

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

Internet Drafts:

"SNMP MIB extension for LAPB", 10/07/1991, Dean Throop, Fred Baker <draft-ietf-x25mib-lapbmib-02.txt>

"SNMP MIB extension for MultiProtocol Interconnect over X.25", 10/07/1991, Dean Throop <draft-ietf-x25mib-ipox25mib-02.txt>

"SNMP MIB extension for the X.25 Packet Layer", 10/07/1991, Dean Throop <draft-ietf-x25mib-x25packet-02.txt>

CURRENT MEETING REPORT

Reported by Dean Throop/Data General

Minutes of the X.25 Management Information Base Working Group (X25MIB)

The x25mib Working Group met to review the LAPB, X25 Packet, and IP over X.25 MIB drafts. The group suggested and approved many changes to all the documents.

Editors Note (md): A detailed listing of these changes is available via ftp under x25mib-minutes-92mar.txt. Refer to Section 1.2 of the Proceedings for retrieval instructions

While there were many changes approved for the LAPB and X.25 MIBs, the group felt the documents were becoming stable. As such the group plans to review the revised documents when they are released and attempt to propose advancement on the mailing list.

The IP over X.25 MIB was changed to the Multiprotocol Interconnect over X.25 MIB to match the RFC coming from the IPLPDN Working Group. There were many changes to this MIB and the Working Group will meet again to review the revised document.

Attendees

Fred Baker	fbaker@emerald.acc.com
Bill Bowman	bill@hpprsd.mayfield.hp.com
Christopher Bucci	bucci@dss.com
Jeffrey Case	case@cs.utk.edu
Carson Cheung	carson@bnr.ca
James Davin	jrd@ptt.lcs.mit.edu
Billy Durham	durham@fishery.honeywell.com
Walter Guilarte	70026.1715@compuserve.com
Jon Infante	jon@rv.icl.com
Frank Kastenholz	kasten@europa.clearpoint.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Andrew Malis	malis@bbn.com
Chandy Nilakantan	csn@3com.com
David Perkins	dperkins@synoptics.com
Christopher Ranch	cranch@novell.com
Timon Sloane	peernet!timon@uunet.uu.net
Brad Steinka	brad@python.eng.microcom.com
Bob Stewart	rlstewart@eng.xyplex.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@interlan.com

3.4 OSI Integration Area

Director(s):

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

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

The OSI Integration Area of the IETF is attempting to bring OSI to the Internet. OSI applications, especially The Directory, offer features that are necessary to sustain the astonishing growth of the Internet; Message Handling Services and Office Documentation offer features that are both new and complementary to existing Internet applications. Working groups in the IETF, in collaboration with RARE working groups, continue to experiment with and deploy these applications in production networks, over OSI as well as TCP/IP.

During the early deployment of OSI applications, hybridization of OSI and TCP/IP was necessary and has proven to be useful in obtaining experience and acceptance for OSI. Experimentation with OSIs transport services is important as well; the continued efforts of the NOOP Group to expand CLNP connectivity across regional and international networks may well prove essential to the growth of the Internet. Currently, this effort requires hybridization of a different sort: SNMP management over OSI is as desirable here as X.500 over TCP/IP. For this reason, and the more far-reaching need to integrate additional protocols and architectures into the Internet, the first multi-protocol Working Group was formed under the joint directorship of the SNMP and OSI Integration Areas. Based on the success of the Working Group, it may not be the last of its kind.

Introduction

The OSI Integration Area currently consists of the following working groups, (working group Chairs/affiliation in parentheses):

- Network OSI Operations (Sue Hares/Merit)
- X.400 Operations (Alf Hansen/SintefDELAB, Rob Hagens/Univ.Wisconsin)
- OSI Directory Services (Steve Hardcastle-Kille/UCL)
- Office Document Architecture (Peter Kirstein/UCL)
- MHS-DS (Kevin Jordan/CDC, Harald Alvestrand/Sintef DELAB)
- SNMP over a Multiprotocol Internet (Theodore Brunner/Bellcore)
- OSI General (Ross Callon/DEC)

The DISI Working Group, in the User Services Area, also has strong ties into the OSI Integration Area.

The following BOFs related to the OSI Integration Area were held in San Diego:

- MIME to MHS Mapping BOF (Marshall Rose/DBC)
- Wais and Directory integration (Steve Hardcastle-Kille/UCL)

Network OSI Operations Working Group (NOOP)

The NOOP Group continues work towards compiling a compendium of OSI support – applications and in particular, support of CLNP – in regional networks. The existing questionaire was reviewed and revised. Work continues on the compilation of a list of available OSI diagnostic tools (Tools RFC), and requirements for OSI support in the SNMP. The Working Group heard presentations from John Curran of NEARNet (New England Academic & Research Network) and Steve Deering NEARNets on OSI NSAP address assignment plans. Work continues as well on the Security RFC, in particular, on an OSI Packet Filtering document that discusses the issues associated with filtering OSI by application type in the context of using packet filtering to restrict OSI connections (i.e., to establish firewalls).

The Working Group is investigating sharing test suites and coordinating test pilots with RARE, especially for inter-domain routing and applications.

X.400 Operations Working Group (X400OPS)

The composition of the Working Group that met in San Diego illustrates the increasingly international flavor of the IETF. Among the 29 participants were 7 Europeans and 3 Koreans.

The routing coordination document has been reviewed again and was now judged suitable for submission as an experimental (or if possible as a prototype) RFC.

Three other documents are still under discussion and need another round of review:

- "Operational Requirements for X.400 Management Domains"
- "Mapping between X.400 (1984/1988) and Mail-11 (DECnet mail)"
- Use of DNS to store RFC -987 mapping data

Experiments with the protocols described in these last two drafts are under way.

The Working Group will produce several documents in the very near term, including:

- Minimum Level of Service
- Table update procedures
- National Character set usage in X.400

In addition, two documents in the standards track are under review by the IESG/IAB and on the verge of becoming proposed standards:

- "X.400 1988 to X.400 1984 downgrading"
- "Mapping between X.400(1988) / ISO 10021 and RFC 822"

During the Working Group session, there was an extensive discussion on MHS communities. Noteworthy was that there was participation from a public service provider.

OSI Directory Services Working Group (OSIDS)

The Working Group met in San Diego with the following results:

Several documents were reviewed and are to be submitted to the IESG for consideration a draft RFCs:

- "Naming Guidelines for Directory Pilots"
- User Friendly Naming
- String representation of Distinguished Names

Experiments continue (and reports were given) in the following areas:

- Quality of service parameters
- JPEG photo attribute

New experiments to be started:

- European character sets in X.500
- DIT counting

Other noteworthy decisions and discussions:

- Naming Schema document maintenance is now moved to a small committee
- The naming schema document will be restructured
- Discussion on deployment of DNS through X.500
- O=internet will be put directly under the root
- Discussion on registration vs Listing
- Discussion on Skinny OSI stack, Lightweight protocols and Simple OSI Stack.

Office Document Architecture Working Group (ODA)

The ODA Working Group met at San Diego with the following results:

It was reported that 5 Implementations are available and under test. So far these implementations run only on top of X.400. Service is to be launched during next few months.

MHS - DS Working Group (MHSDS)

This Working Group met for the first time in San Diego.

There were two input documents:

- PP use of directory Services (implementation specific)
- The ISO proposal on this issue.

After an extensive debate, it was proposed that seven documents be drafted based upon the PP input document:

1. Table and tree handling by DS
2. O/R name representation as Distinguished Name
3. Routing info in DS
4. X.400/822 mapping tables in DS
5. tMailing list expansion using DS
6. RFC-822 routing based on DS
7. Simple application profile

The next Working Group meeting will be in Innsbruck, Austria in May 1992.

SNMP over a Multiprotocol Internet Working Group (MPSNMP)

The Working Group was chartered with defining the mapping and SNMP encapsulation for three transport domains – OSI, Appletalk, and XNS/IPX. Since it was envisioned that at a future time, additional transport domains might be identified, the Working Group agreed to produce a how to RFC, identifying a checklist of issues to consider in specifying an encapsulation of SNMP.

The Working Group considered three existing documents:

- The informational RFC 1298, entitled "SNMP over IPX",
- The internet draft (draft-ietf-appleip-snmp-00.txt) entitled "SNMP over AppleTalk"
- The experimental RFC 1283, entitled "SNMP over OSI"

In all cases, the committee determined that a connectionless TS similar to UDP was desirable (architecturally appropriate). Security, maximum packet size, and addressing considerations for each transport domain were discussed. Authors of each of the documents were present, and agreed to make appropriate changes, and further agreed to post the documents to the mailing list for a three week review. If a consensus is reached following the posting period, the Working Group agreed that the three SNMP over foo documents be submitted to the IESG for consideration as draft RFCs.

OSI General Working Group (OSIGEN)

This Working Group did not meet in San Diego.

The Area Directors will propose that this Group will disband.

MIME to MHS Mapping BOF (MIMEMHS)

There were two input documents for this BOF on mapping between MIME and X.400. A third one was announced but not tabled.

There was a consensus that the mapping could be easily constructed for some of the body-parts and not be constructed for others. It was therefore proposed to create a short lived Working Group to create two documents:

- Basic mapping
- Specific bodypart mapping

This deals with an initial set of registered conversions. The registration will have to be maintained.

WAIS and Directory Integration BOF (WAIS)

This meeting followed discussion at the "Living Documents" BOF the previous evening, and was more focussed in its discussion.

The WAIS, World-Wide Web, Prospero systems for network information retrieval (NIR) were presented (the Gopher protocol was presented in plenary the following day). The X.500 Directory was presented in the light of NIR needs, as were two proposals to use the directory to refer to documents. A discussion followed as to how to allow these systems to interoperate, and on requirements for name spaces. A working group was proposed to define the format for a generalized printable format for a name or address in any of these systems.

CURRENT MEETING REPORT

Minutes of the MIME to MHS Mapping BOF (MIMEMHS)

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

Attendees

Harald Alvestrand	harald.alvestrand@delab.sintef.no
Tim Berners-Lee	timbl@info.cern.ch
Ronald Broersma	ron@nosc.mil
Erik Fair	fair@apple.com
Jill Foster	jill.foster@newcastle.ac.uk
James Galvin	galvin@tis.com
Jisoo Geiter	geiter@gateway.mitre.org
Terry Gray	gray@cac.washington.edu
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Erik Huizer	huizer@surfnet.nl
Kevin Jordan	kej@udev.cdc.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Larry Masinter	masinter@parc.xerox.com
Daniel Molinelli	moline@gumby.dsd.trw.com
Keith Moore	moore@cs.utk.edu
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Marshall Rose	mrose@dbc.mtview.ca.us
Michael St. Johns	stjohns@umd5.umd.edu
Einar Stefferud	stef@nma.com
Steve Thompson	sjt@gateway.ssw.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us

CURRENT MEETING REPORT

Reported by Tim Berners-Lee/CERN

Minutes of the WAIS and Directory Integration BOF (WAIS)

This meeting followed discussion at the "Living Documents" BOF[2] the previous evening, and was more focussed in its discussion.

The WAIS, World-Wide Web (WWW), Prospero systems for Network Information Retrieval (NIR) were presented (the Gopher protocol was presented in plenary the following day). The X500 Directory was presented in the light of NIR needs, as were two proposals to use the Directory to refer to documents. A discussion followed as to how to allow these systems to inter-operate, and on requirements for name spaces. A Working Group was proposed to define the format for a generalized printable format for a name or address in any of these systems.

WAIS

John Curran of BBN presented the WAIS protocol, in the absence of anyone from Thinking Machines Corporation who was originally responsible for it. The WAIS model is of a number of servers, each of which serves a number of databases, each of which contains a number of documents. Client software allows many databases to be searched at the same time. The server keeps an inverted full text index for each database, so the search is very fast. Non-text files may also be served: recent extensions allow indexing of text files in new formats. The files indexed need not be copied, but the index is of the same order of size as the files.

Many databases exist, but there is no scalable way of finding them (TMC currently keeps a master index). Use of X500 was discussed.

The WAIS protocol is an extended subset of Z3950. The differences were discussed: WAIS allows relevance feedback ("Give me a document like this one") , and specifies how a query should be formulated. WAIS and Z39.50 have the same presentation layer.

Documents in the Directory

Wengiyk Yeong presented his paper OSI-DS-22, "Representing public archives in the directory"[4]. His project puts information about documents, including the network address for retrieval, into the Directory. He currently has RFCs and FYI documents in, but would like to move on to other internet archives. He concluded that he needed a more sophisticated approach. It was difficult to characterize arbitrary archives, with too little information about them. (See IAFA Working Group[5]).

The World-Wide Web

Tim Berners-Lee presented the World Wide Web (WWW) and discussed requirements for interworking between the systems. The WWW project was initially funded to provide an information infrastructure to the world-wide community of high energy physicists. The data model is of documents which are hypertext and/or searchable indexes. The philosophy behind it is that a user should be able to point and click on phrase or a word within a document and the associated document would be retrieved from wherever in the world and presented to the user in an appropriate format - without the user having to be aware of where the document is located or what the access method is. These details are hidden in the hypertext links. There were server programs for many information servers, gateways to WAIS, Archie and Gopher and client programs for various user machines.

The WWW clients use several protocols for accessing documents (FTP, NNTP, WAIS, Gopher, and WWW's own "HTTP") although this is hidden from the user. The HTTP protocol is a simple stateless search/retrieve protocol running over TCP. As originally conceived but not yet implemented, it included authentication and data format negotiation. Tim discussed the differences between World Wide Web (WWW), WAIS, Archie, Gopher and Prospero systems.

The need for a Universal Document Identifier (UDI) for describing the address or, given a directory, name, for a document whatever is access protocol was discussed, as outlined in OSI-DS-XX. Each application uses a "handle" for a file which can be prefixed by the particular protocol name to generate a universal address.

Most systems (WAIS excepted) are extensible, entertaining document addresses which refer to other systems. WAIS indexes currently can only refer to documents in the same database, let alone with other retrieval methods. There is a need for WAIS to be more flexible. John Curran said he would bring this to the attention of the WAIS community.

Addresses would not in the long-term be suitable for references to documents, so it was hoped that some sort of directory service, operating within the UDI framework, would be incorporated.

More information: telnet info.cern.ch. Client and server code is available by anonymous FTP from info.cern.ch.

Mailing lists: www-talk@info.cern.ch, www-interest@info.cern.ch

Discussion document: OSI-DS-29[6]

Representing the Real World in the Directory

Paper: OSI-DS-25[7]Steve Kille discussed this paper "Representing the Real World in an X.500 Directory".

A Listing Service may be used to group like information items together, for example, to provide a Yellow Pages Service. Such a service could for example provide for members of a special interest group, or could group documents on a particular subject. Services such as Archie could be considered to be Listing Services. One imagines an information Universe in which Information Brokers provide different subject based (say) views via their listing service. One would then need to locate the various listing services (using a mechanism such as a directory?)

UK British Library Project

Paul Barker described a project, sponsored by the British Library, to represent grey literature (unpublished research papers) in the Directory. The project is thought to be unlikely to succeed - but one of the aims is to demonstrate whether or not it is possible. They will take the (UK) MARC records and model these within X.500. They might also consider trying to provide a listing service so that the documents might be retrieved more readily by subject area.

Prospero

Cliff Neuman described Prospero. It follows a file system model, rather than the hypertext model. It is built on UDP for speed. It has the notion of a Directory which contains links to other objects (other directories or files). It returns the link to the information object and then automatically retrieves the file by another mechanism by the appropriate access method (Archie, WAIS, NNTP, WWW - soon!, NFS, FTP, etc.) It has been used very successfully to access the Archie database.

Cliff stated that he expected to be able to use X.500 to translate between the document ID and how to get the document.

With Prospero the user has his own view of the global information base (or has a view built for him). Cliff thought there should be multiple name spaces - but the difficulty would be that these would need representing near the top of the directory tree. With multiple user chosen views - this would be difficult to manage. Also two users might refer to an object by different handles which would be relative to their individual name spaces - difficult when passing references (say in a mail message) from one person to the other.

The concept of "Closure": Each object has a related name space. All references within the object are resolved using the context of the name space. Name spaces themselves have global network addresses, but the user doesn't see that. More information: info-prospero@isi.edu

System 33

Larry Masinter talked about a project at Xerox PARC with the following concepts:

HANDLE	32 byte number (is a content ID). In fact this contains hints for finding the document.

FILE Location (6 part)	Protocol; Host; Path; piece; format; timeout
Description	(normal "Catalogue" information: Name, Author, etc.)

There is format negotiation when a document is retrieved. It is not simple in reality to categorize data formats as there is such a plethora of different varieties.

Gateways provide access between systems not sharing transport protocols.

Also considered Access Control. ACL is part of description. The Server exploits multiple protocols for Search and retrieve.

There is a problem with dealing with different types of document (applications for jobs, product specifications, memos, contracts, faxes, etc.) It is difficult to normalize the attributes of a general document.

Summing up

Tim Berners-Lee summed up by saying that all applications described used resolvable document address, and so for interworking, we need a universal representation for such a network object address. With the coming of directories, names should increasingly be used in place of network addresses. The UDI was intended to be able to hold either a name or address for any access protocol. (This is not the same as "USDN" a document serial number which is not resolvable, but only one of which exists for each document).

In discussion, Steve Kille suggested there should be a working group on details of UDIs and a separate one for USDN. A comment was that the WWW data model encompasses those of the other systems. John Curran insisted on a better term than "UDI", suggesting "Document Access Token".

Peter Deutch's need for a USDN is to be able to determine the equivalence of two USDNs. Chris Weider agreed to co-author a document on the issues. Jill Foster suggested a pilot project to put UDI's in the directory for a set of documents and to have the Gopher, Prospero, and Archie people try to utilize these.

[These Minutes have been largely built from Jill Foster's report[8] and Karen Sollins' notes[9] for which I am most grateful, though errors in the above are probably mine.

References:

[1]http://info.cern.ch/hypertext/Conferences/IETF92/IETF-9203.html
[2]http://info.cern.ch/hypertext/Conferences/IETF92/LivingDocuments.html

[3]http://info.cern.ch/hypertext/WWW/Administration/Mailing/ietf-wwx-bof
[4]file://cs.ucl.ac.uk/osi-ds/osi-ds-22-00.txt
[5]http://info.cern.ch/hypertext/Conferences/IETF92/IAFA-BOF.html
[6]file://cs.ucl.ac.uk/osi-ds/osi-ds-29-00.txt
[7]file://cs.ucl.ac.uk/osi-ds/osi-ds-25-00.txt
[8]http://info.cern.ch/hypertext/Conferences/IETF92/WWX_BOF.html
[9]http://info.cern.ch/hypertext/Conferences/IETF92/WWX_BOF_Sollins.html

These Minutes are available in hypertext form using WWW as:

http://info.cern.ch./hypertext/Conferences/IETF92/WWX_BOF_ mins.html

as well as through the normal channels.

Attendees

Harald Alvestrand	harald.alvestrand@delab.sintef.no
Serge Aumont	aumont@cicb.fr
Paul Barker	p.barker@cs.ucl.ac.uk
Tim Berners-Lee	timbl@info.cern.ch
J. Nevil Brownlee	nevil@aukuni.ac.uz
Susan Calcari	calcaris@cerf.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Peter Danzig	danzig%warthog.usc.edu@usc.edu
Osmund de Souza	osmund.desouza@att.com
Erik Fair	fair@apple.com
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Jill Foster	jill.foster@newcastle.ac.uk
Cliff Frost	cliff@cmsa.berkeley.edu
Anders Gillner	awg@sunet.se
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Alisa Hata	hata@cac.washington.edu
Tim Howes	Tim.Howes@umich.edu.
Lew Jenkins	lew@enlil.premenos.sf.ca.us
Peter Kirstein	kirstein@cs.ucl.ac.uk
Sylvain Langlois	sylvain@cli53an.edf.fr
Hock-Koon Lim	lim@po.cwru.edu
Kenneth Lindahl	lindahl@violet.berkeley.edu
April Marine	april@nisc.sri.com
David Martin	dem@fnal.fnal.gov
Larry Masinter	masinter@parc.xerox.com
Daniel Matzke	matzked@cerf.net
Cyndi Mills	cmills@bbn.com

Daniel Molinelli	moline@gumby.dsd.trw.com
Keith Moore	moore@cs.utk.edu
Robert Morgan	morgan@jessica.stanford.edu
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Rakesh Patel	rpatel@rutgers.edu
Geir Pedersen	geir.pedersen@usit.uio.no
Mel Pleasant	pleasant@rutgers.edu
Lars Poulsen	lars@cmc.com
Michael Schwartz	schwartz@cs.colorado.edu
Mark Smith	mcs@umich.edu
Patricia Smith	psmith@merit.edu
Karen Sollins	sollins@lcs.mit.edu
Chris Weider	weider@ans.net
Russ Wright	wright@lbl.gov
Wengyik Yeong	yeongw@psi.com

3.4.1 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 is expected.

Goals and Milestones:

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

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.

CURRENT MEETING REPORT

Reported by Kevin Jordan/CDC

Minutes of the MHS-DS Working Group (MHSDS)

The first meeting of the MHS-DS Working Group took place on March 17th, 1992 in San Diego. The duration of the meeting was two hours, from 10:00 until 12:00. Despite the shortness of time, very good progress was made.

The meeting was Co-Chaired by Harald Tveit Alvestrand of Delab Sintef, Norway and Kevin Jordan of Control Data Corporation, Arden Hills, Minnesota, USA. After brief introductions and a review of the Working Group Charter, Steve Hardcastle-Kille provided an overview of his paper entitled, "PP Use of Directory".

This paper defines a comprehensive approach for using X.500 Directory Services to support X.400 routing, RFC1148 address mapping, distribution list expansion, and various other purposes suited to electronic mail. The paper establishes the basis for further work by the MHS-DS Working Group.

The Working Group decided that the following seven draft RFC's will be derived from "PP Use of Directory":

1. An RFC which describes how to represent tables in the directory and which also defines the concept of X.400 routing trees.
2. An RFC which defines the mechanism for mapping X.400 O/R addresses onto X.500 distinguished names. This RFC will define the basic mapping rules, and it will also describe how knowledge information and cross references can be used to avoid unnecessary chaining and referrals through DSA's managed by ADMD service providers.
3. An RFC which defines the mechanism for using X.500 Directory Services to support X.400 routing. This RFC will draw heavily from the concepts defined in the first two RFC's.
4. An RFC which defines the mechanism for using X.500 Directory Services to support mapping between RFC822 addresses and X.400 addresses, in compliance with RFC1148bis. This RFC will also draw heavily from the concepts defined in the first two RFC's.
5. An RFC which defines mechanisms for using X.500 Directory Services to support practical implementations of distribution list expansion and management.
6. An RFC which defines mechanisms for using X.500 Directory Services to support RFC822 mail routing and distribution.
7. An RFC which defines a simple profile of the other RFC's, especially RFC's 1 through 4. This RFC could be used as a guide to producing minimal implementations of the first six RFC's above.

We agreed that these RFC's should be released initially as Experimental Drafts. As working implementations become available and practical experience establishes proof of concept, the RFC's will be evolved into Internet Standards. The Working Group openly expressed optimism that working implementations would be available by the end of this year. In fact, a PP-based prototype is already underway.

The first four RFC's have top priority. They will be created and distributed before the last three. Our goal is to bless the final drafts of the first four documents at the next general IETF meeting in Boston this summer.

It was generally agreed that well performing implementations of these RFC's will need to implement mechanisms for caching information obtained from the directory. This will be necessary in order to minimize the number of directory operations requested during normal operations, thereby optimizing response time in critical functions such as route discovery. A recommendation was made that at least one of the RFC's should provide guidelines for the implementation of caching. In particular, guidelines should be provided for recommended time-to-live values of cache entries.

Harald reminded/informed the Working Group of related work occurring within ISO. Specifically, R.H. Willmott has written a paper concerned with using X.500 Directory Services in support of X.400. This paper is being circulated within the ISO standardization community. Harald brought a copy of the paper to San Diego and provided copies to the MHS-DS membership. We agreed to review this paper and evaluate its relevance to our work plan.

Before the meeting was closed, Steve Hardcastle-Kille summarized some new features which will be added to the RFC's in response to comments he has received recently. These features include:

1. Improvement in the mechanism for defining and managing MTA passwords in the directory. Specifically, a mechanism will be defined for using the normal X.500 compare operation to compare MTA passwords. It is likely that this will enhance the security of MTA passwords stored in the directory.
2. Added support for explicitly defining a UA which always causes nondelivery reports to be generated. It is not clear how useful this feature truly is, but it does provide a mechanism for catching undeliverable X.400 addresses as early as possible.
3. Added a hook for discovering MTA's which are common to a community of recipients. This feature allows mutually remote MTA's to optimize their usage of network resources by detecting that a large group of recipients can be reached by relaying mail through a common MTA which will perform a distribution function. For example, an MTA in the US can use this feature to send a single copy of a message to a WEP in France, and the French WEP will distribute the message to a group of MTA's within France. This optimizes utilization of the network link between the US and France.

Action items:

1. Kevin Jordan will update the MHS-DS Charter such that it specifies the RFC's we intend to produce and the time-frame in which we intend to produce them.
2. Steve Hardcastle-Kille will generate the first four draft RFC's from his "PP Use of Directory" paper and make them available to the Working Group for review and comment. He will accomplish this prior to the upcoming JENC-3 conference in May.
3. Erik Huizer will provide Steve with an OID to be used for identifying the new draft RFC's.

Future meetings:

The next meeting of the MHS-DS Working Group will take place at the upcoming JENC-3 (3rd annual Joint European Networking Conference) in Innsbruck, Austria. The meeting will probably take place on Friday, May 15.

The third meeting of the MHS-DS Working Group will take place at the next general IETF meeting in Boston, Massachusetts, in July.

Attendees

Harald Alvestrand	harald.alvestrand@delab.sintef.no
Serge Aumont	aumont@cicb.fr
William Biagi	bbiagi@cos.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Alan Clegg	abc@concert.net
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppen@verw.switch.ch
Mark Fox	m_fox@took.enet.dec.com
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@es.net
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Max Hillson	hillson@koala.enet.dec.com
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Barbara Jennings	bjjenni@sandia.gov
Kevin Jordan	kej@udev.cdc.com
Eva Kuiper	eva@hpindda.cup.hp.com
Sylvain Langlois	sylvain@cli53an.edf.fr
Triet Lu	trietl@sparta.com
David Miller	dtm@mitre.org
Daniel Molinelli	moline@gumby.dsd.trw.com

Michael Newell	mnewell@nhqvax.hg.nasa.gov
William Nichols	nichols@took.enet.dec.com
David Piscitello	dave@sabre.bellcore.com
Mark Sleeper	mws@sparta.com
Einar Stefferud	stef@nma.com
Huyen Vu	vi@polaris.dca.mil
Chris Weider	weider@ans.net
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov
Wengyik Yeong	yeongw@psi.com

3.4.2 Network OSI Operations (noop)

Charter

Chair(s):
Susan Hares, skh@merit.edu

Mailing Lists:
General Discussion: noop@merit.edu
To Subscribe: noop-request@merit.edu
Archive: merit.edu:pub/noop-archive

Description of Working Group:

The working group is chartered to work on issues related to the deployment of CLNP in the Internet. The first area of this group's work has been the learning necessary to start deploying OSI in internet networks. This phase has includes planning for OSI deployment by creating routing plans for regional networks and education on using OSI routing protocols.

This first area of the group's work will be on-going as we continue to deploy OSI in the Internet. This step has lead to people deploying OSI for Pilot projects and demonstrations of OSI.

The second step of deploying OSI will be the transition of OSI from a pilot service to a production service. During this phase we will work on specifying the network debugging tools and test beds. We will need to track the level of OSI support in the Internet. We will need to provide documentation for new users of OSI on the Internet.

Goals and Milestones:

Jan 1992	Post as an Internet Draft of a tutorial for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.
Jul 1992	Post as an Internet Draft a collection of regional Routing and Addressing plans.
Ongoing	Provide a forum to discuss OSI routing plans by email or in group discussions.
Jul 1992	Post as an Internet Draft a list of OSI Network Utilities available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.
Jul 1992	Post as an Internet Draft a description of OSI network layer debugging methods.
Jul 1992	Post as an Internet Draft a list of OSI Network Layer NOC tools available in the public domain and from vendors. This list will be passed over to the NOC tools group effort for joint publication.

Apr 1992 Post as an Internet Draft a requirements document specifying what OSI network tooks are needed on every host and router.

Jul 1992 Submit to the IESG for Proposed Standard a requirements document specifying what network tools are needed on every OSI host and router.

Aug 1992 Submit to the IESG as an Informational RFC a description of OSI network layer debugging methods.

CURRENT MEETING REPORT

Reported by Sue Hares/Merit, David Bolen/ANS and Dave Miller/MITRE

Minutes of the Network OSI Operations Working Group (NOOP)

The Network OSI Operations Group met three times during the week of the San Diego IETF. The first meeting took place on Monday morning. Steve Deering presented his ideas behind his paper "City Codes is an Alternative to Topological NSAP Allocation (RFC 1237)." In addition, Ross Callon presented the basics of RFC 1237. Both presentations prompted a great deal of discussion. Sally Tarquinio took very detailed notes of the discussion. Due to the length of both the notes and the discussion, the notes will not be available in the Proceedings. Notes may be retrieved via anonymous ftp from merit.edu in /pub/iso/noop/notes/notes.03.16.92.am).

In addition to the City Code discussion, the following topics were discussed:

- Mobile Hosts
- Comparison Geographical Area Code (GARP) with NAT and CNAT
- GARP vs RFC 1237
- Whether asymmetric pathways are acceptable.

The second NOOP session occurred on Wednesday morning at 9:00 a.m. and covered several different items. Notes were taken by David Bolen.

Usage Questionnaire

Sue Hares made available copies of the OSI usage questionnaire and requested that anyone involved with OSI work try to complete one. This was a copy of the same questionnaire that was previously distributed electronically.

A question was raised as to why DECnet was considered different than CLNP (p. 9) - the answer was that it wasn't really, but the goal was to see if DECnet usage was pushing CLNP usage (here, DECnet really means DECnet Phase V traffic).

NEARnet OSI Routing/Addressing Plan

- Introduction

  John Curran from NEARnet (New England Academic & Research Network) made a presentation of NEARnet's OSI Plan. NEARnet is comprised of 120 members in six states, and coordinates with other New England service providers to provide service in that area. Cisco routers are used throughout NEARnets network.

- Address Assignment Plan

  When researching an address assignment plan, NEARnet found that area codes were a nice match for population density, and therefore for assignment beneath NEARnet's AAI.

NEARnets final address format breakdown assumed the following limits:

- Total of 16 COs per area code.
- Total of 256 RDs per CO. This could be a real problem in a fairly short-term (~ two years). It is hard to gauge demand though, and NEARnet isn't the only network assigning in the New England area.

CO codes are assigned to aggregate at other boundaries as well:

```
.00-.3F = Massachusetts    (.10 = 508 area code, .20 = 617, etc..)
.40-.4F \
.50-.5F  \  Assign to different states - allows room for
.60-.6F  /  expansion according to area code.
.70-.7F /
.80-    still some slack available for future expansion
```

The next step is then to assign RDs logically under this scheme. The multiple aggregation points within this scheme helps to limit the routing table size.

- Routing (ala RFC1237)

  The following points were raised with respect to routing issues under this sort of an address assignment scheme:

  - Routes will often have to be manually injected.
  - IGRP is not currently collapsing routes - hopefully newer protocols will begin to do this.
  - NEARnet doesn't like the fact that they have to accept all routes as it allows NEARnet's routing tables to grow without bounds.
  - NEARnet sees NSAP changes as a lot of work currently - thus, if a customer has already been assigned an RD, NEARnet lets the customer keep it.
  - NEARnet will accept any other assignments, but isn't sure what will happen with other networks - will they be as accepting as NEARnet?

- Questions

  John brought up the general question as to what the general opinion of this plan was. Could the approach be viewed as dangerous? The following points were raised:

  - Something has to be done, and at least this policy allows future aggregation - it's very hard to take back what we give out today.
  - There could be a problem if all service providers don't become as accepting as NEARnet. If routes begin to be refused, it might cause everyone to bail out to their own AAI which would create a flat address space once again.

- If we are allowing entropy at all (as this plan does), then there needs to be some sort of entropy reduction in the system. A possible recourse might be the need to start charging for the announcement of a special route to other networks.

A general point was made that at the moment, the whole area of addresses and assignment represents more of a controlled economy than a true market economy. Moving from one to the other is always tough.

NEARnet's addressing and routing plan may be found (via anonymous ftp) in: nic.near.net:/docs/osi-routing-plan.txt or merit.edu:/pub/iso/noop/papers/nearnet.osi-routing-plan.txt

John also gave credit to CICNet for their previously released OSI plan, and said that NEARnet's plan borrowed a lot from CICNet's.

OSI Pilot Projects

The discussion of OSI pilot projects centered around some documentation that Sue supplied describing some of the work that RARE is doing in this area. RARE has a suite of tests that they are requesting users at their sites to perform, sending them results back to a central site to be summarized. Sue was interested in whether or not their was enough interest in trying the same sort of pilot plan here in the U.S., as well as trying to get together another OSI demo for Interop East.

The general consensus of the Group was that, yes, it would be useful to try the same sort of pilot project here, and that the RARE approach seemed a reasonable way to proceed. It would also be nice to see about some coordination with RARE, although mostly for inter-domain and application level than at the ES-IS and IS-IS level. The application-level portion of the RARE plan was a little weak, and may need to be augmented for our tests.

A possible problem was brought up in that a number of sites have beta implementations of OSI code, and may not be able to publish the results of tests. Sue suggested that at least saying "I've tested" is useful, even if the exact results of the test cannot be released.

NIST was brought up as an organization that was already handling some OSI testing in the same vein as what was being discussed. NIST has provided open labs and a test environment for multiple vendors to come together in order to test interoperability. There has been no automated or documented test procedures followed, however - just vendor engineers running particular tests. Richard Collela will be sending some further information about this testing to the NOOP mailing list.

The fundamental difference between the testing that has been performed at NIST, and the type of pilot projects being run by RARE is that the latter case involves actual end users, while the former is run by the vendors and their engineers.

John Curran suggested that the request for tests be sent out to the NOOP list. While many midlevels may not run the tests themselves, they may have clients that can. Sue Hares agreed to send the test plan, and a request for volunteers to the NOOP mailing list.

Document Review: (reported by Dave Miller)

- TOOLS RFC

  Their was a brief discussion on the "Tools" Document and the need for tools to provide OSI ping, OSI traceroute, and OSI routing table dump utilities. The discussion then focused on what routing table information should be available via SNMP. It was noted that the document should specify a minimal set of objects as MUST requirements for this specification.

  For CLNS MIB objects, no one took exception to the list in the draft document. For IS-IS, the document currently states the object in very general terms. Dino Farinacci was asked to write-up a minimal list of IS-IS objects and send them to the NOOP mailing list. For IDRP, no one took exception to the list in the draft document.

  There was no review of objects for CMIP management at this time.

- SURVEY FORM

  Sue Hares gave a status overview of the OSI in the Internet Survey. The plan to maintain the survey is to perform a monthly revision to incorporate any new information received. Sue also encouraged others to respond to the survey since only about ten responses have been received to date.

  There were a few suggestions to modify the survey (it was noted that the changes should be highlighted when new surveys are sent out):
  - DECnet traffic should refer to DECnet Phase V traffic.
  - Text should refer to RFC 1006 as opposed to a TCP/IP stack.
  - There should be a context section added to identify who's responding to the survey and what their role in OSI is.

  Cathy Wittbrodt suggested that the surveys be stored individually on-line so particular responses could be retrieved as desired. Sue agreed to do this.

  Dave Farber suggested sending the survey out to the IETF mailing list to reach a broader community. Sue agreed to do this.

- SECURITY RFC

  Walt Lazear gave an overview of the OSI Packet Filtering document. The document discusses the issues associated with filtering OSI by application type in the context of using packet filtering to restrict OSI connections (establishing firewalls).

  Walt noted that he has not received comments on the current document and solicited feedback.

  There was some discussion of what were the security requirements that were trying to be met. MITRE was asked if they could put together a short paper on OSI security policy to set the context for this work and to stimulate further discussion. Bill Barns volunteered himself and Walt to pursue this.

- NIST IS-IS TEST LAB

  Rich Colella gave a quick overview of the IS-IS test lab established at NIST. Two Open Lab sessions have been conducted to date to perform vendor interoperability testing.

  A report of the router testing is being prepared. Rich agreed to get a copy of the completed report sent to the NOOP list.

The third NOOP session took place during the 7:00 p.m. session on Wednesday, March 18th, and was devoted to discussion of the IDRP for IP document. A new working group will be formed to discuss the IDRP issues. Detailed notes are available via ftp, cd ietf, get noop-minutes-92mar.txt.

Attendees

Andy Adams	ala@merit.edu
Vikas Aggarwal	aggarwal@jvnc.net
Robert Aiken	raiken@nsf.gov
Steve Alexander	stevea@i88.isc.com
Nagaraj Arunkumar	nak@3com.com
Zorica Avramovic	zorica@sparta.com
William Barns	barns@gateway.mitre.org
Tony Bates	tony@ean-relay.ac.uk
William Biagi	bbiagi@cos.com
David Bolen	db31@nis.ans.net
Scott Brim	swb@cornell.edu
J. Nevil Brownlee	nevil@aukuni.ac.uz
Eric Brunner	brunner@practic.com
Niels Brunsgaard	nob@dowtyns.dk
Ross Callon	callon@bigfut.enet.dec.com
John Chang	jrc@uswest.com
A. Lyman Chapin	lyman@bbn.com
Henry Clark	henryc@oar.net
Alan Clegg	abc@concert.net
Richard Colella	colella@osi.ncsl.nist.gov
Rob Coltun	rcoltun@ni.umd.edu
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Dave Cullerot	cullerot@ctron.com
John Curran	jcurran@bbn.com
Steve Deering	deering@xerox.com
Kathleen Dodd	kathy@gateway.mitre.org
Tom Easterday	tom@cic.net
Dino Farinacci	dino@cisco.com

Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Vince Fuller	vaf@stanford.edu
Eugene Geer	ewg@nvuxr.bellcore.com
Robert Hagens	hagens@cs.wisc.edu
Susan Hares	skh@merit.edu
Jeff Hayward	j.hayward@utexas.edu
Juha Heinanen	juha.heinanen@datanet.tele.fi
Randy Hoebelheinrich	rho@rowan.lanl.gov
Kathleen Huber	khuber@bbn.com
Bob Jeckell	rrj@3com.com
Barbara Jennings	bjjenni@sandia.gov
Dan Jordt	danj@nwnet.net
Dave Katz	dkatz@merit.edu
H. A. Kippenhan	kippenhan@fndcd.fnal.gov
Yoav Kluger	ykluger@fibhaifa.com
Deidre Kostick	dck2@sabre.bellcore.com
Eva Kuiper	eva@hpindda.cup.hp.com
John Larson	jlarson@parc.xerox.com
Walter Lazear	lazear@gateway.mitre.org
Eliot Lear	lear@sgi.com
Kurt Lidl	lidl@sura.net
Susan Lin	suelin@ibm.com
Tony Li	tli@cisco.com
Triet Lu	trietl@sparta.com
Gary Malkin	gmalkin@xylogics.com
Bill Manning	bmanning@rice.edu
Glenn McGregor	ghm@merit.edu
Milo Medin	medin@nsipo.nasa.gov
David Miller	dtm@mitre.org
Henry Miller	henrym@sacusr.mp.usbr.gov
Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.dca.mil
Donald Morris	morris@ucar.edu
Brad Passwaters	bjp@sura.net
David Piscitello	dave@sabre.bellcore.com
Jon Postel	postel@isi.edu
Rex Pugh	pugh@hprnd.rose.hp.com
Yakov Rekhter	yakov@watson.ibm.com
Ron Roberts	roberts@jessica.stanford.edu
Jonathan Rodin	rodin@ftp.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel.rodrigues@att.com
Sharad Sanghi	sharad@ans.net

Erik Sherk	sherk@sura.net
William Simpson	Bill_Simpson@um.cc.umich.edu
Keith Sklower	sklower@cs.berkeley.edu
Mark Sleeper	mws@sparta.com
Frank Solensky	solensky@clearpoint.com
Walter Stickle	wls@ftp.com
William Streilein	bstreile@wellfleet.com
Sally Tarquinio	sally@gateway.mitre.org
Richard Thomsen	rgt@lanl.gov
Claudio Topolcic	topolcic@nri.reston.va.us
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Huyen Vu	vi@polaris.dca.mil
Carol Ward	cward@westnet.net
Bert Wijnen	wijnen@vnet.ibm.com
Steve Willens	steve@livingston.com
Scott Williamson	scottw@nic.ddn.mil
Walter Wimer	walter.wimer@andrew.cmu.edu
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Peter Yee	yee@ames.arc.nasa.gov

3.4.3 OSI Directory Services (osids)

Charter

Chair(s):

Steve Hardcastle-Kille, `s.kille@cs.ucl.ac.uk`

Mailing Lists:

General Discussion: `ietf-osi-ds@cs.ucl.ac.uk`
To Subscribe: `ietf-osi-ds-request@cs.ucl.ac.uk`
Archive:

Description of Working Group:

The OSI-DS Group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this Group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

Done	Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.
Done	Study the relationship of the OSI Directory to the Domain Name Service.
Ongoing	Maintain a Schema for the OSI Directory on the Internet.
Ongoing	Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

Internet Drafts:

"Building an Internet Directory using X.500", 11/19/1990, S. Kille <draft-ietf-osix500-directories-01.txt, or .ps>

"Using the OSI Directory to Achieve User Friendly Naming", 11/26/1990, S. Kille <draft-ietf-osids-friendlynaming-03.txt, or .ps>

"Handling QOS (Quality of service) in the Directory", 03/20/1991, S.E. Kille <draft-ietf-osids-qos-01.txt, or .ps>

"Naming Guidelines for Directory Pilots", 03/21/1991, P. Barker, S.E. Hardcastle-Kille <draft-ietf-osids-dirpilots-04.txt, .ps>

"DSA Naming", 03/21/1991, S.E. Hardcastle-Kille <draft-ietf-osids-dsanaming-02.txt, or .ps>

"Schema for Information Resource Description in X.500", 06/14/1991, Chris Weider <draft-ietf-osids-resdescripx500-00.txt>

"Schema for NIC Profile Information in X.500", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-osids-nicprofilex500-00.txt>

"Interim Directory Tree Structure for Network Infrastructure Information", 06/14/1991, Chris Weider, Mark Knopper, Ruth Lang <draft-ietf-osids-treestructure-00.txt>

"Directory Requirements for COSINE and Internet Pilots (OSI-DS 18)", 07/09/1991, S.E. Hardcastle-Kille <draft-ietf-osids-requirements-00.txt, .ps>

"Generic Security Service Application Program Interface Overview and C bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-00.txt>

"An Access Control Approach for Searching and Listing", 09/23/1991, S.E. Hardcastle-Kille, T. Howes <draft-ietf-osids-accesscntrl-00.txt, .ps>

"Representing Public Archives in the Directory", 12/04/1991, Wengyik Yeong <draft-ietf-osids-archdirectory-00.txt>

"A String Representation of Distinguished Names", 01/30/1992, S. E. Hardcastle-Kille <draft-ietf-osids-distnames-00.txt, .ps>

Request For Comments:

RFC 1275 "Replication and Distributed Operations Extensions to Provide an Internet Directory using X.500"

RFC 1276 "Replication and Distributed Operations Extensions to Provide an Internet Directory"

RFC 1277 "Encoding Network Addresses to Support Operation Over Non-OSI Lower Layers"

RFC 1278 "A String Encoding of Presentation Address"

RFC 1279 "X.500 and Domains"

CURRENT MEETING REPORT

Reported by Steve Hardcastle-Kille/UCL and Justin Walker/Apple

Minutes of the OSI Directory Services Working Group (OSIDS)

There were no comments on the Minutes of the San Jose meeting; they were accepted as written.

Steve Hardcastle-Kille was to have prompted George Brett to circulate documents. It was not known if this had been done, so the action was dropped.

Richard Collela was to send a current list of the OIW documents to the OSIDS mailing list. The question was asked whether this was done, and no one knew for sure. Subsequent to the meeting, Rich did distribute the OIW document list. It is appended here.

Other items of business were to be discussed as specific points on the Agenda.

Liaison Reports

1. RARE WG3: (Erik Huizer)

   A number of documents were discussed. The "character set" issue was also discussed. On a sad note, the January meeting for WG3 was cancelled, due to restructuring within RARE. In the future, it will be more like IETF (from May onwards). There will be a followon to WG3, but the form has not yet emerged.

2. ISO/CCITT - (No liaison was present.)

   Availability of the Directory root over CONS has been requested by JANET. This will cause reachability problems for CLNS use. The issues haven't been fully addressed yet.

3. OIW: (Russ Wright)

   Agreements on replication have gone stable (1992); 1988 documents on replication are stable. Trying to distinguish between '88, '92 items. The X.400 and X.500 SIGs met. The X.400 folks complained about lack of attribute types for routing. EWOS sent a statement about adding transport requirement (NSAPs don't specify transports). Major work on international standard profiles (dealing with DAP) is underway; this should be out by December.

4. NADF: (Einar Stefferud)

   The pilot proposed for February 1992 is "underway" , with all NADF members participating. Due to agreements between NADF members, a "utopian" view of the pilot will be presented to the world outside the NADF in that no details will be discussed as to which pilot member is doing what. There are interworking issues between this pilot and the White Pages Pilot (WPP), due to different naming schemes and the listing vs. registration models. Discussions have been held at NADF to determine that two pilots could *not* be connected. According to Stef, there is no common naming of schema. The major problem is operational (naming of DSAs, etc.). PSI can not act as broker (there are knowledge and data sharing problems). Desire is there, so it seems that meetings are needed to discuss this. The NADF pilot work needs to stabilize before these can reasonably proceed. The NADF wants to push knowledge sharing (open DIT; global system).

WPP was being run as a registration tree, so that the it had to be the national registration authority for c=US by virtue of holding the c=US MASTER. While none of the principals ever claimed to be the US registration authority per se, we just ended up doing that as a consequence of the registration model. It was pointed out that these assumptions were necessary for early deployment.

NADF is waiting for the 1992 changes to the directory (X.500) to be published to determine what membership will do about compliance.

The NADF has issues of competitiveness, tariffs, etc., guiding its pilot development. These are real world assumptions. The WPP assumptions were simplifying. NADF documents are available, modulo media issues.

5. DISI: (Chris Weider)

   Three new RFCs are out: 1292, 1308, 1309 (a "real executive summary"). They now have a clean slate, so if new documents are needed, speak up.

6. AARNet: (Steve Hardcastle-Kille)

   Report to the IETF OSI-DS WG from the AARNet Directory Project

   (a) Australian Networkshop in last December.

       We conducted a demonstration of the Directory at the recent Networkshop which attracted considerable interest, and as resulted in 3 more AARNet members joining the pilot.

       The demonstration was spoiled somewhat by the failure of our frame grabber and where we had hoped to use colour images, JPEG encoded, we had to make do with greyscale imagines (still using JPEG). The DIT used for the Networkshop is still available, as "c=AU@o=Australian Networkshop", having been migrated from the loan machine we had at the Networkshop to one of our project machines.

   (b) Future of the AARNet Directory Project.

       Officially the project has concluded, except for the submission to AARNet of our report, but we expect that the Project will continue, hopefully with additional funds from AARNet.

We will continue to champion the Directory as an information resource and encourage AARNet members to run their own directories. We also intend to use of our machines to provide a service where AARNet members can experiment with the Directory without having to run their own, as well as providing a registration point for any organisation connected to AARNet so that basic information about their organisation can be made available through the Directory.

(c) Binary distribution of DUAs and DSAs.

The AARNet Directory Project have made available a number of binary kits (SPARC, RISC/Ultrix, Sun3 and Pyramid) of the Quipu distribution for anonymous ftp on ftp.adelaide.edu.au in the pub/white_pages/KITS directory. The main purpose of this is to allow other sites to easily access the the pilot, either by making access to the Directory available at their site or allow them to easily configure a DSA of their own. The kit has been tailored for sites wishing to join the pilot in Australia but the binaries could be used anywhere.

(d) Current state of the Directory in Australia.

There are currently 25 DSAs in Australia, and they master 45,975 entries. After checking the sites that have fetched a copy of one of our binary kits I would hope that there will be 3 more sites in Australia starting to run their own DSA shortly.

Status Reports of Operational Pilots

- FOX: (Tom Tignor)

  FOX is waiting on NSF funding; final reports have been submitted, and nothing is happening now. Individual efforts:

  - SRI

    x5whois - whois information in a DSA. Conversion problems overcome, but DSA loading is taking a long time (they have added more memory, reduced the number of attributes held). There are 150000 entries now. Interoperability testing (between QUIPU and CUSTO) is underway.

  - PSI

    Three commands are being developed at PSI. x5ftp is under development. x5rfc is done, development-wise, but is awaiting on x5ftp for release of both to assure no changes to x5rfc due to a problem discovered in x5ftp. usconfig is done and released. It should be in future ISODE releases (it's basically the core of the wpp-addon stuff right now).

  - MERIT

    Working on making information resources (e.g., k-12, NIC) available on X.500; schema documents on these are available. They are looking at storing data as pointers to original information. The University of Michigan is developing a Macintosh DUA (maX.500, by Mark Smith). This isn't related to FOX in terms of funding or anything. The connection with FOX is that Merit, a FOX member, likes it and has been helping to promote it.

- ISI

  Currently, they are in a cheerleading mode, and acting as a central switchboard for these efforts. They are just moving to QUIPU 7.0. They are looking at a lightweight version of x5whois.

  A question was asked regarding the transition to X.500 in Europe: have there been real directories mapped into x500? The consensus is no, that most directory efforts have focused on creating new X.500 databases. We should then look at any problems arising from moving the "whois" base to X.500.

- White Pages

  According to Wengyk, the transition to the NADF naming scheme is going unusually slowly due to opposition/apathy on the part of pilot project members. There are 91 organizations in the pilot now. Operationally, heavy use of the wp.psi.net machine reported; this appears to be causing the 'dad' server to fail sporadically. Also, as a result of heavy use of wp.psi.net, an auxiliary DSA "c=US@cn=Horned Frog" was created to be the service DSA on wp.psi.net. So the "Fruit Bat" DSA is now back to being a c=US (and other things) slave only. During Wengyk's report, the NADF/WPP differences were discussed again.

- PARADISE: (Paul Barker)

  There have been problems with (large) getedbs. PARADISE is moving to ISODE 8.0, and this is causing some service upset. Use of central DUA services on a central ULCC system is rising. It was requested that we all please take some of the lush documents from PARADISE. These describe the services supplied, as well as the user interface alternatives provided. Revisions are being planned for the DUA (e.g., loosening up the hierarchy). Multilingual versions of I/F are becoming available. Among others, a management interface for simple maintenance; for small or disinterested users (e.g., for those with a simple o=, or for lower level updates). A probe (written in C++) is being produced, with better post processing of results. One partner (the Dutch PTT) has sent query to other PTTS on attitudes on X500 (most said "X.What"?). Steve Hardcastle-Kille and Paul Barker are producing 3 metric documents - for DUA, DSA, and Pilots. These will be in the form of questionnaires, and they are looking for details on each.

The operational reports being given, we plunged into the individual items from the Agenda.

Security

The NADF started looking at it last year. A Directory Bill Of Rights has been published as RFC 1295. Each word of the Bill of Rights has been lovingly crafted to say exactly what it says, nothing more and nothing less. Also, security for competitive products has been under study. A revision of this is expected after NADF meeting, when it will be revised and published as an RFC (the week of 4/21).

Need for a Directory Operations Group - Does the IETF need a WG for Operations, dealing practical issues of running a directory service on the Internet. This group could work on a benchmark document, operating specifications, interoperability issues. During the discussion, a question was raised regarding the difference between the new group and DISI; the latter was described as an educational provider. Suggested differences: the OSI-DS provides implementations of the directory; DISI is for users; and the new group is for operators. It was pointed out that this obeys the Narrow Focus admonition of IETF WGs. A straw poll indicated low interest in both having and not having a separate WG for operations (a majority abstained from the voting; only a handful cast votes), so the issue was put aside for now.

Strategy Document

Some issues need to be resolved, privately, before getting closure on this document. Concern has been raised that Steve H-K is generating documents faster than the rest of us can read them. The protagonists are looking for insight on what should go into and what should not go into the document. The problems are: the document describes the registration model; attention needs to be paid the work of the NADF and listing model. The document also doesn't address deployment issues, e.g., where the resources come from. A section on security is wanting, but should be filled in from Steve. A version is promised by the end of April. Anyone with views should speak with Erik Huizer.

New Object Models

Three papers on new object models have been published. The object models are described therein as schemas. Comments are solicited. One comment - this doesn't match the X.500 model of having "objects" that have "real" significance. What is "service" (called "resource" in the papers)? A subgroup meeting was suggested for further resolution of the subclass/object definition. Another comment: how does one search, based on schema? One must distinguish between DIT structure and object models. The former is to be considered in the WAIS BOF. There followed a discussion of how to represent network infrastructure information in the directory. A previous paper thought now to be wrong (by the author) It was suggested that IP representation should be widened to include host parts, AD, other information. Concern was expressed that the representation of network addresses not lose information (e.g., net masks).

OSI-DS-12 Discussion (and the "list vs. registration" debate).

Earlier, after discussion on the osi-ds mailing list, the document was modified to add a note on an alternative (championed by Christian Huitema). In discussions at RARE, the WG3 folks have suggested removing the alternative (i.e., going back to the form prior to Christian's suggestions were added). There followed a lively discussion on the two alternative positions, although no one was present to support the alternative. The position taken in the paper was well and eloquently defended. Note that the document hasn't been through the IESG/IAB process yet. Note also that the disputed section is really an advisory one, dealing with countries without current registration authorities. A straw poll was taken on the question of removing the "alternative": lots in favor; two abstentions; none against.

Also, it was observed that Sec 3: the UFN statement makes it (this particular UFN syntax) special. After discussion, it was accepted that this section should be deleted.

The subject of "Who Owns The Root?" arose, relating to an ongoing concern with resolving the differences between the listing and the registration models. A discussion ensued regarding the effects of putting in things to the root, willy-nilly. o=Internet, small numbers of "l="s, , and a small number of DSAs were examples used to highlight some of the issues. No conclusions were reached by the meeting.

Registration vs. Listing Discussion

In the Listing corner were Einar Stefferud and Marshall Rose. The NADF is leveraging off US civil authority (in particular, that resting with the states, counties, and "localities"). There is a problem of looking for a company (or a person) without knowing its state of incorporation (that is, Delaware, not Confused) (or, in the case of a person, the organization chart of his (s/he/it's) company). From this view, the DIT should be organized based on search needs. Therefore, we need to do this at national level. A basic issue is the mapping from civil authority to DIT (need not be 1-1). This is the Listing view.

It is claimed that registration authentication already exists, except for registration under c=US. ANSI does allow registration here (at the c=US level) at $2500 a pop; the details have appeared on the net a number of times. Control of the directory (i.e., assuring that we don't pollute the directory at too high a level) comes with listing charges.

The listing model, following an anecdote from Einar Stefferud, emphasizes the need to lose your keys under the light. The point is that you are more likely to find your keys where there is light (even if you didn't lose them there). Similarly, one needs to list oneself in the Directory where one would be expected. Where one is actually registered is less of an issue, and depends on vagaries of the domain administering your neck of the woods (or DIT).

The membership was advised that no lunch break would be forthcoming until this discussion is done. As a result, our focus narrowed.

The registration side view was detailed by Steve Hardcastle-Kille. The Directory should leverage off existing civil authority. It is important to separate directory and registration (at least at high levels; at lower levels, convenience of the DNS approach may override). Multiple providers are needed, as is naming coherency (tied in later). NADF requires all providers to assure naming coherency. There are 3 kinds of registration: ANSI, civil, and derived. These are the listings. The point was made that listings are actually a form of registration, in that a listing takes up "name space" and that listing agents must work to assure that collisions don't occur. A counter argument was made that collisions will naturally clean themselves up as a result of the competitive nature of the Directory provision market. The problem seems to be the issue of recursive listing authorities.

The debate continued with no clear winner, although the weight of evidence seemed to favor the listing folks. The point was made that the NADF model had no implications for components of the DIT outside c=US (other than those inherent in its adoption beyond those boundaries). The Listing view starts with the observation that names are intellectual property, sanctioned by civil authority within some (e.g., c=) boundary. Listings (following NADF) are algorithmically derived from names, hence (at least within the domain covered by NADF), no chance for collision. There was disagreement on the issue of listing being an implicit registration.

In the end, the sense of the meeting (by show of hands) was to push 12 to an RFC. The Listing vs. Registration debate will continue, with efforts being made to align the various pilots for interoperability. Steve Hardcastle-Kille will reread NADF-175, to help determine what can be done, while Wengyk will continue to work on the interoperability issues between the NADV and WPP pilots.

User Friendly Naming (UFN)

Per a suggestion from the IAB, the UFN document will be split. The specific string representation (the use of ";" vs. ",") has gotten a lot of discussion. The use of UFN itself has received little comment. Discussion on the string rep: use ';' or ',' or "not" both.

On the vote to forward the UFN document, the 'ayes carried (so it will be forwarded). Steve Hardcastle-Kille will post the resulting documents.

Quality of Service (QOS)

There has been no progress on the Quality Of Service issue. The QUIPU implementation now agrees with "the documentation" (the RFC, not the QUIPU manual). There are two pieces: the user interface and the deployment. Deployment underway. To date, there has been no user interface defined to allow a user to invoke this capability. A dissenting view on the utility of QOS is that it is up to the guy who provides the service to describe QOS, and there is little or no uniformity to allow this. For example, for the provider using the ISODE-provided DSA, he may describe it as experimental if he is a commercial provider, or as non-experimental if he is a university researcher. The experiments will continue.

JPEG

Support for this should be in the next version of QUIPU. A schema for JPEG photos is not yet ready. Currently, this is specified as an octet string. There is a conflict with G3Fax, which will be resolved by separating attributes (per last meeting).

Character Sets

The paper is partly from discussions in RARE WG3 and RARE/COSINE groups. Current DUAs don't support national characters and the T61 data type very well. Europe (at least) has a requirement for national characters. The providers need to add this support in a coordinated way. The directory should have national versions of names (I18N). The solution proposed by the author is:

- Store national characters using T61 string syntax
- DSA string search algorithms must account for I18N'd names
- Mapping table
- DUA presentation to user dictated by the user (to use or not use I18N)

Issues include:

- What are precise requirements?
- What are the implications for UFN?
- The necessity to agree on conversion at a national level

Note that UFN is assumed to be defined on abstract character set, so I18N not an issue(?).

Remarks:

- Is this only an "operational" issue, or are there other issues?
- How are I18N strings stored, searched? X.500 discusses this briefly, but that discussion does not seem acceptable.
- No experimentation is underway, but should be started (e.g., between France and Norway).

Counting the DIT

Current work is DSA-specific and is very implementation specific. A suggested new approach is to add new attributes (integers all) that count appropriate things at each level. Counts can be done manually or automatically. The question arose: do we count the # of registered or listed entries? The sense of the meeting was to progress with the experiment Tim Howes volunteered to write some quipu syntax handlers for the counting attributes.

He did not volunteer to change quipu to make use of these features, but would be willing to at least look into that.

Steve Hardcastle-Kille volunteered to write some quipu syntax handlers for the counting attributes. He did not volunteer to change quipu to make use of this stuff, but he'd be willing to at least look into that.

RFC1274

The original intent was for Steve and Colin to maintain this document. Problems have arisen with the time needed to maintain it (keep it up to date) and how to maintain it. A suggestion is that we try a structured approach a la SNMP. We need to document each "object class" as with a MIB: what are the mandatory, optional, and experimental entries. Another problem is expressed concern over the openness of the process to extend attribute and class lists. We could either establish a small committee or a new WG to oversee the development of the Directory. The consensus was for a small committee. The IAB was previously asked about their feeling. The thought was put forward that this could be more like Host Requirements, than like SNMP. A show of hands called for an attempt to tack down what the committee would do. Five brave souls stepped forward. Paul Barker volunteered to restructure the main document. The new structure will include procedures for extending the current definition, a list of other documents and general purpose attributes; and a mechanism for generating other documents as needed.

Schema Publishing

An alternative to the preceding approach is "don't write RFCs". Instead, just write a new schema into the DIT. Tim Howes and Mark Smith volunteered to write this up for public consumption. Code to do this is also needed. There followed a discussion of machine generated schema descriptions, e.g., by automatically culling appropriately prepared documents from the new RFC1274 structure. Stay tuned.

preferredName Attribute Discussion

Others deferred to the committee. The attribute type preferreddisplayname is a subtype of CN (for 1988 Directories, this would be a duplicate of the CN). A DUA could use this as the display value for CN. The attribute would not be mandatory.

Administrative Limits

In a note sent out in January, the idea of size and time limits on searches was proposed. Also, it would be nice to have a value to limit the number of DSAs to which to refer during a search. This is thought to be related to issues of QOS. A document discussing these values was proposed for the next meeting. Note that this puts information about Directory use in the Directory. Doing this may require the use of security above that currently available. Should these be represented in MIBs? Steve Hardcastle-Kille discussed the use of SNMP as a tool for the management of directories.

Adding DNS information to directory - Software has been created to load DNS information into the DIT. "dnsconfig" will create the initial EDB hierarchy for the DNS part of the tree. "dnsupdate" loads DNS information into the directory. "fred" has been modified to resolve user@domain. "dnsconfig" and "dnsupdate" are under test. The modified "fred" has not been released yet. The work is being done at PSI, by Wengyik Yeong.

Some comments were offered on RFC1279, based on this work: using case insensitive string to represent the values of all types of DNS records is too simplistic. However, defining separate attribute syntaxes for every DNS record is both impractical and wasteful. It doesn't scale, and the effort is wasted for those less frequently used record types. As a compromise, one can special case those DNS records with their own syntaxes. The others can continue to use case insensitive string values.

It was suggested that DNS records that use case insensitive string values need to have the sequence in which the TTL, Class, and Type fields occur, standardized. One could fix the sequence (e.g., in the order of Class, TTL and Type) with all three mandatory in every record; or fix the sequence as above, but let the class be optional and default to IN. Steve noted that the flexibility present in the current draft of RFC1279 is due to similar flexibility in the DNS RFCs themselves (which allow different orderings for TTL, Class and Type). Steve and Wengyk took an action item to find out why such flexibility exists in the DNS.

Some concern was expressed that "leaves" in the DNS can be interior nodes in the DIT. This could be a problem, since QUIPU is very slow when loading non-leaf entries.

During the wrapup of this discussion, some open questions were posed. Can we make o=Internet the final resting place for the DNS tree in the DIT? It was felt that the group had consensus on this issue, and that the answer is "yes".

Can we load up all the top level domains (from DNS) without explicit consent from domain owners? With reference to the discussion of zone transfers below, the consensus was that the answer to this question is "yes".

Further questions:

- Where do we put the o=Internet tree? We assumed this had already resolved, i.e., place it under the root. It was noted that we have no authority to do that, hence perhaps we should place it under a c= node. It is possible that, e.g., CNRI could pay ANSI to register it. One camp says "just do it"; another says ="put it where it is safe", so we won't have to change further down the road. A straw poll regarding where to place the root was taken. There were lots for the "under c=somewhere" position, only a few for the "under root [few]", and a number of abstentions.

  The debate on placement continued for a while, with lots of back and forth regarding the effects of each choice of placement of "o=Internet". We ended on the comment that, if we chose to place it under the root, this would be one of the few times that Stef would later say he told us so.

- Do DSA operators have right to load top level DNS zone files into the Directory? One argument is that, if you permit zone transfers, then the door is open. The counter argument is that DNS never agreed to this (X.500) usage, that this is a new usage, and thus assumptions should not be made about acceptability. It was uniformly agreed that the following applied: ZONE TRANSFERS SHOULD BE NOTED AS LEADING TO POTENTIAL HARM.
- Further discussion of the interaction between X.500 attributes and DNS records. It was suggested that attribute syntax for common DNS records be changed (to fit more neatly into X.500), while less common DNS records be standardized using string records.
- According to Steve Hardcastle-Kille, it doesn't matter where OID's come from. A top level, OID, 0=Internet, is desirable. Vint Cerf suggests that the Internet Society may be able to register "Internet" for OID and RDN.

Common Authentication Technology (CAT)

This discussion was concerned with integrating security in a variety of technologies. The presenter (John Linn) was from the CAT WG, and wanted to raise the consciousness of the OSI-DS WG, since many of the issues that confronted them involve naming in the X.500 sense. The CAT depends on global naming, in that they are using X.500 Distinguished Names (DNs) in X.509 certificates. They are encountering early adopter penalties.

A major issue is that the CAT folks don't want DNs that are used for authentication to diverge from those used for other purposes within the Directory. CAT needs to deal with hosts, users, processes (for authentication purposes) within the environment and protocols used by DNS, accommodating mismatches. Hosts are currently handled, but not users or processes. This is, fundamentally, a naming issue.

Implementations must support Directory access routines for security purposes (i.e., parsing is not needed; the only requirement is for matching).

Our respective areas could benefit from naming coexistence (API definitions, available support libraries. The main question the presenter had was: is this part of OSI-DS charter or current plans?

In the discussion that followed, to comment on the presentation and answer John's question, the issue of what problem was being solved arose. It is important to not replace DNS host names, but to provide unique names for authentication usage. John will post a brief description of his work, with pointers to his documents.

The sense of the meeting was that this was best pursued in the context of the discussion list rather than during the meeting, because a clear understanding of the issues is wanting.

Lightweight Protocols

Dixie and DAS are recent alternatives to the full OSI stack implementation of a Directory user agent. They are incompatible. Wengyik Yeong, Tim Howes, and Steve Hardcastle-Kille have designed other alternatives.

LDBP

This is the first in a series of protocols. It is targeted to browsing (no name service). It is session oriented and supports few operations. Its error structures have been flattened (they contain only a code and a string). There is no Basic Encoding Rules (BER) (and no authentication!!), and instead, uses its own binary rules (the so-called "string encoding"). Passwords are sent in the clear. This could be used in a DUA, bridge, or it could be embedded in a DSA for direct support.

SOS

Based on the observation that many applications make little use of the upper layers of OSI, SOS allows direct mapping to a transport layer(either CONS or CLNS). The (optional) streaming approach is incompatible with X.400/X.500 security issues (because access to the full PDU is required) – this is fundamental.

Steve Hardcastle-Kille hasn't looked at related OSI work on application environments and on reorganizing application layers, putting the session goo in the transport layer, removing the presentation layer.

Note that SOS really transcends X.500 - it is a much more general OSI issue. This differs from the "Skinny Stack", a multi-layer collapsing of a local stack, due to Van Jacobson. Note that the presence of the skinny stack is not seen at the remote end, unlike the case for SOS. X is the first to use the skinny stack (cf. implementors agreements). Concern was expressed that, because of the above observation, this is beyond the scope of this WG.

Colin's message (on DS 26) is: is LDBP needed in the face of SOS? Also, Colin had question on "DAP lite" - why not work on this, providing interoperability with existing DSAs, rather than new protocols (e.g., add skinny stacks).

How do these affect "homogeneity" for having a global directory? What are motivations, tradeoffs, payoffs.

Final Agenda Item: Next meeting - it was decided that the next OSI-DS meting will be at the 24th IETF meeting (July, 1992, in Boston). We will postpone DSA Naming discussions until then.

Actions

Huizer	Produce revised strategy document by the end of April.
Weider	Revise OSI-DS 14, 16, 17, 19 in light of meeting and offline discussions.
Hardcastle-Kille	Revise OSI-DS 12, and submit to IESG as proposed standard.
Hardcastle-Kille	Submit OSI-DS 24 to IESG as experimental.
Hardcastle-Kille	Revise OSI-DS 23, and submit to IESG as proposed standard.
Hardcastle-Kille	Review NADF 175 in light of Rose's comments.
Yeong	Study issues of NADF / WPP interoperability.
All	Continue QOS experiment.
Schema Group	Sort out JPEG Schema.
Geir Pederson et. al.	Start Character set experiment.
Howes	Provide DIT Counting Syntax Handlers.
All (?)	Start DIT Counting experiment.
Barker	Establish Schema Group.
Howes/Smith	Write OSI-DS note on Schema Publishing.
Schema Group	Record decision on preferred name.
Pays	Write OSI-DS note on storing limited information in the DIT.
Hardcastle-Kille	Revise RFC 1279 in consultation with Yeong.
All	Continue DNS in X.500 experiment.
Yeong/Howes/Hardcastle-Kille	Revise LDBP note.
Hardcastle-Kille	Examine ISO proposed alternatives to SOS.

Appendix: OIW Documents

The output of the NIST Workshop for Implementors of OSI (OIW) is a pair of aligned documents, one representing Stable Implementation Agreements (SIA), the other containing Working Implementation Agreements (WIA) that have not yet gone into the stable document. Material is in either one or the other of these documents, but not both, and the documents have the same index structure.

The SIA is reproduced in its entirety at the beginning of each calendar year, with an incremented version number. Replacement page sets are distributed subsequently three times during each year (after each Workshop), reflecting errata to the stable material, as well as new functionality declared stable. In this way an up-to-date document is maintained.

Retrieving OIW Documents

The documents are available from osi.ncsl.nist.gov via anonymous ftp (129.6.48.100) or anonymous ftam: user = anon, realstore unix,

osi.ncsl.nist.gov filestore NULL #1/#1/#1/NS+47000580005a0000000001e137080020079efc00

All documents are in the directory: ./pub/oiw/agreements

Attendees

Claudio Allocchio	claudio.allocchio@elettra-ts.infn.it
Harald Alvestrand	harald.alvestrand@delab.sintef.no
John Ballard	jballard@microsoft.com
Paul Barker	p.barker@cs.ucl.ac.uk
William Biagi	bbiagi@cos.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
Alan Clegg	abc@concert.net
Richard Colella	colella@osi.ncsl.nist.gov
James Conklin	conklin@bitnic.educom.edu
Urs Eppenberger	eppen@verw.switch.ch
Stefan Fassbender	stf@easi.net
Mark Fox	m_fox@took.enet.dec.com
James Galvin	galvin@tis.com
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@es.net
Sang-Chul Han	schan@garam.kreonet.re.kr
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Alton Hoover	hoover@ans.net
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl

Ole Jacobsen	ole@csli.stanford.edu
Barbara Jennings	bjjenni@sandia.gov
Darren Kinley	kinley@crim.ca
Mark Knopper	mak@merit.edu
Eva Kuiper	eva@hpindda.cup.hp.com
Sylvain Langlois	sylvain@cli53an.edf.fr
Kenneth Lindahl	lindahl@violet.berkeley.edu
Triet Lu	trietl@sparta.com
Scott Marcus	smarcus@bbn.com
Daniel Matzke	matzked@cerf.net
David Miller	dtm@mitre.org
Daniel Molinelli	moline@gumby.dsd.trw.com
Robert Morgan	morgan@jessica.stanford.edu
William Nichols	nichols@took.enet.dec.com
Tracy LaQuey Parker	tracy@utexas.edu
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Rakesh Patel	rpatel@rutgers.edu
Geir Pedersen	geir.pedersen@usit.uio.no
David Piscitello	dave@sabre.bellcore.com
Jon Postel	postel@isi.edu
Marshall Rose	mrose@dbc.mtview.ca.us
Ursula Sinkewicz	
Mark Sleeper	mws@sparta.com
Mark Smith	mcs@umich.edu
Einar Stefferud	stef@nma.com
Tom Tignor	tpt%@isi.edu
Justin Walker	justin@apple.com
Chris Weider	weider@ans.net
Brien Wheeler	blw@mitre.org
Cathy Wittbrodt	cjw@nersc.gov
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov
Wengyik Yeong	yeongw@psi.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.4.4 OSI General (osigen)

Charter

Chair(s):

Robert Hagens, `hagens@cs.wisc.edu`
Ross Callon, `callon@bigfut.enet.dec.com`

Mailing Lists:

General Discussion: `ietf-osi@cs.wisc.edu`
To Subscribe: `ietf-osi-request@cs.wisc.edu`
Archive: `janeb.cs.wisc.edu:/pub/archives/ietf-osi`

Description of Working Group:

Help facilitate the incorporation of the OSI protocol suite into the Internet, to operate in parallel with the TCP/IP protocol suite. Facilitate the co-existence and interoperability of the TCP/IP and OSI protocol suites.

Goals and Milestones:

Done — Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet. Coordinate addressing format with GOSIP version 2 and possibly other groups.

TBD — Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley.

TBD — Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.

TBD — Determine what should be used short-term for (i) intra-domain routing; and (ii) inter-domain routing.

TBD — For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Determine if there are any outstanding issues here.

TBD — Review short-term issues involved in adding OSI gateways to the Internet. Preferably, this should allow OSI and/or dual gateways to be present by the time that Berkeley release 4.4 comes out.

Request For Comments:

RFC 1139 "Echo function for ISO 8473"

3.4.5 Office Document Architecture (oda)

Charter

Chair(s):
Peter Kirstein, kirstein@cs.ucl.ac.uk

Mailing Lists:
General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk
Archive:

Description of Working Group:

The ODA Working Group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g., Spreadsheets. The Working Group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liaison with the SMTP and X.400 Working Groups.

This Working Group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Done	Inaugural meeting.
Done	Produce a paper stating what ODA standards or profiles still need completing.
Jul 1991	Produce paper on how both SMTP and X.400 message systems should be supported.
Done	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Dec 1991	Provide first feedback on the ODA Pilot.
Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.

CURRENT MEETING REPORT

Reported by Peter Kirstein/UCL

Minutes of the Office Document Architecture Working Group (ODA)

The Chair repeated that for the ODA Working Group, some papers are available in electronic form on the UCL-CS info-server: "info-server@cs.ucl.ac.uk". The documents are accessed by standard message systems, giving a message body of the form:

request:ietf-osi-oda

topic:xxxx

where xxxx is the name of the document required.

A number of documents are available in the info-server - all in text form. Many of the documents are also stored in ODA/ODIF format. The list of documents currently in the collection is listed in a document called INDEX.

The latest document defining the current status is [1] below. It is available on the info-server. It gives details about the current implementations available for the Pilot. These are listed below:

PRODUCT	Supplier	Status	Product	Availability	Source
SLATE/ODA V1.2	BBN/UCL	OK	Slate-yes	Now	UCL
Word-for-Windows/ODA v3	Bull	OK	Yes	4/1	Bull
DECWRITE/ODA	DEC	OK	Yes	Now	DEC
Global View	Xerox	Testing	Yes	Now	Xerox

UCL started shipping SLATE/ODA v1.2 in February 1992; this is based on SLATE v1.2. Because of the late shipment, few have had a chance to test it. BBN is planning to bring out V2 of SLATE during the second quarter of 1992; we expect SLATE/ODA v2 to be ready for shipping by the end of that quarter. The Bull product is being released for shipping on 4/1; UCL has tested it for conformity. The DEC product is available now, and has been tested for conformity; no other site had tried a recent version of that software. The Xerox product was delivered to UCL only just before the meeting. It is still being tested for interoperability by UCL.

At earlier meetings, it had been hoped that a WORDPERFECT/ODA system from University Politechnic of Catalonia would become available. Unfortunately for two reasons this now seems less likely. UCL tested it during the summer of 1991, and it was found non-conformity in a number of ways; no new version has been received since. Second, there seemed to be some contractual problems in getting that version used in the Pilot, because it included some software from ICL - which UPC had not had permission to release for the

Pilot. Earlier UPC had stated that they intended to have software entirely of their own available for the summer of 1992 - but nothing has been heard further from them on that subject. The Chairman could only assume that the WORDPERFECT software could not become available for a Pilot in the near future.

During the discussion, it became clear that there was a requirement to use a large number of National characters; for instance support for the specific Nordic ones were requested by one participant, and of Greek ones by another. It was considered desirable to use character set switching according to ISO 8859/4 in the long-term; mention was made also of ISO 10041. Neither the Bull nor BBN implementations support such characters; no information was available on the DEC or Xerox ones.

A number of additions to the Q112 profile of the Standard are under discussion; one of these is tables. If this became standardized, it is probable that several of the implementations would include it shortly.

A number of small pilots were mentioned - a Nordic one, one based on Mitre, one organized by UCL in Europe, and a small one at NASA-Ames. There was interest in broader pilots. The Chair promised to set up a list of ODA and normal mailboxes. This list will be available from ietf-osi-oda, and the interested parties will be requested to provide their details to the Chair.

So far, most use of the systems were using X.400 for transport. Interest was expressed in working with the Multi-media Mail format (MIME); this is in accordance with the Charter, since this is the output of the SMTPEXT Group. The Chairman mentioned that it had already been assured that ODA was a recognized Content Type in MIME, and he promised to investigate when there could be integration of MIME with some of the ODA products. There is already a MIME implementation available to UCL, and he hopes to add it to the UCL PP mail system during the next quarter.

It was expected that there would be substantial experience with the Pilot prior to the next meeting which would take place during the July IETF meeting in Boston.

Reference

1. D. Sadok et al: The ODA Document Convertors, UCL Internal Report No 2, Version 3, March 1992.

People So Far Receiving SLATE/ODA from UCL

FIRST NAME	SURNAME	ORGANISATION	CITY	COUNTRY
Velu	Sinhu	OKIDATA	Framingham, Mass	USA
Jim	Knowles	NASA-AMES	Mountain View, CA	USA
Mark	Stansfield	Paisley Coll	Paisley	Scotland

Mark	Savela	Telecom Res	Otakaari	Finland
Erik	Lillevold	NTRA	Kjeller	Norway
Encarna	Pastor	U. Pol. Madrid	Madrid	Spain
Knut	Holm	SINTEF	Trondheim	Norway
Trevor	Hales	CSIRO	Carlton, Vic	Australia
Naresh	Kumar	Touch Com	Campbell, CA	USA
Andrew	Macpherson	BNR-Europe	Harlow, Essex	UK
Paul	Kennedy	DEC	Galway	Ireland

Attendees

Harald Alvestrand	harald.alvestrand@delab.sintef.no
David Crocker	dcrocker@mordor.stanford.edu
Kathleen Dodd	kathy@gateway.mitre.org
Lew Jenkins	lew@enlil.premenos.sf.ca.us
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@binky.arc.nasa.gov
Daniel Molinelli	moline@gumby.dsd.trw.com
Brien Wheeler	blw@mitre.org
Linda Winkler	lwinkler@anl.gov
Peter Yee	yee@ames.arc.nasa.gov

3.4.6 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 comunicate 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 amongst multiple protocols?

Goals and Milestones:

Done	Post an Internet Draft describing operation of SNMP over OSI.
Done	Post an Internet Draft describing operation of SNMP over IPX.
Done	Post an Internet Draft describing operation of SNMP over Appletalk.

Nov 1992 Submit a document describing the operation of SNMP over OSI as a Proposed Standard.

Nov 1992 Submit a document describing the operation of SNMP over IPX as a Proposed Standard.

Nov 1992 Submit a document describing the operation of SNMP over Appletalk as a Proposed Standard.

Internet Drafts:

"SNMP over OSI", 04/10/1992, Marshall Rose <draft-ietf-mpsnmp-overosi-01.txt>

CURRENT MEETING REPORT

Reported by Ted Brunner/Bellcore

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

The meeting started by considering the Working Group Charter. It had not been sent out over the ietf mailing list, nor made available for anonymous ftp prior to the first meeting of the Working Group. It was read aloud (enough paper copies were available for about half the participants) and accepted as written.

The Charter names three transport domains (OSI, Appletalk, and XNS/IPX) over which SNMP can be carried, and tasks the Working Group (among other things) with developing suitable encapsulation techniques. Questions were raised as to the appropriateness of considering other transport domains. In particular running SNMP over SNA was brought up as an interesting candidate for consideration. The Working Group decided it did not have the necessary expertise to pursue such an undertaking and it was dropped. Running SNMP directly over Ethernet was suggested and also dropped. There is already an RFC that deals with such a case (RFC1089). Running SNMP over TCP was suggested. Here the sense of the Working Group was that this was outside the scope of the Charter. The Charter speaks of environments where the recommended method for exchanging SNMP messages (UDP/IP) is not available. It does not speak of changing the recommended method of communicating SNMP messages. The rational for choosing UDP/IP rather that TCP/IP is expressed in RFC1270.

The informational RFC (RFC1298) entitled "SNMP over IPX," was considered first. SNMP encapsulation in this case is relatively straightforward, and all of the relevant points addressed by RFC1298 were discussed.

- It is recommended that the minimum maximum packet size supported by the SNMP/IPX agent be raised to 546 - the limit under IPX.
- Two sockets are assigned: for get-next/sets and for traps.
- The agent address field in the trap pdu is left as 0.0.0.0 and the agent identified by the source address in the IPX header.
- The IPX addresses will be represented by an OCTET STRING.
- In OBJECT DISCRIPTOR is defined for the ipx transport domain.

The Working Group identified a small omission. An OBJECT IDENTIFIER for an initial party ID was added to the RFC.

With this change, and assuming customary time for consideration of Internet Drafts and no further controversy, the Working Group expressed its intention to propose this RFC for promotion to "Proposed Standard."

A show of hands was made of companies who had implemented or had some inclination of implementing this standard. Present were Synoptics, Spider, MicroComm and Shiva.

The Working Group next considered the Internet Draft (draft-ietf-appleip-snmp-00.txt) entitled "SNMP over AppleTalk." SNMP encapsulation in this case was a little less straightforward. All the relevant issues raised by the draft were discussed.

- Appletalk (DDP) datagrams contain 0 to 586 octets of data. The Working Group recommended that the SNMP agent increase its minimum maximum packet size to 586.
- DDP Socket numbers and protocol types are assigned to SNMP requests, responses and traps.
- In Appletalk, network elements advertize themselves using the Name Binding Protocol (NBP) which dynamically binds names to addresses. A NBP type is assigned to SNMP agents (to receive requests), and SNMP managers (to receive traps).
- The agent address field in the trap PDU is left as 0.0.0.0. In Appletalk the name advertized by the NBP is unique and constant, but the address is not. So the agent inserts its name (object and zone) in the VarBind list of the trap PDU.
- Names are represented as OCTET STRINGs.
- There is discussion of some implications for robust service with this use of names and the NBP to identify managers and agents. Caching is suggested.

The Working Group expressed possible implementation concerns at the maximum name length of 96 Bytes.

The Working Group observed that this reliance on NBP is susceptible to denial of service attacks, but this is not a *further* security hole to SNMP.

The Working Group recommended that "SNMP Security Widgets" be added to this draft: Transport Domain OIDs and Default Party.

With these changes, and assuming customary time for consideration of Internet Drafts and no further controversy, the Working Group expressed its intention to propose this Internet Draft for promotion to "Proposed Standard."

Again a show of hands was made of companies who had implemented or had some inclination of implementing this standard. Present were Apple, Novell and Shiva. Mentioned, but not present, were Cayman, Neon, Interconn, etc.

Next the Working Group discussed the experimental RFC (RFC1283) "SNMP over OSI."

- Three transport mappings are included: Connectionless Transport (CLTS), Connection Oriented Transport (COTS) over TP4 and over TP0 with X.25.

- In OSI, locally meaningful "selectors" are used where IP uses "well known ports." T-selectors for request/response and for trap are specified.
- An address representation convention known as string encoding is adopted from RFC1278.

The Working Group agreed (with the author) to drop this convention.

- The trap pdu specifies a network address, following the syntax defined in the CLNP mib (RFC1238).

The Working Group observed that this convention is not the same as that of the previous two documents considered. It was recommended (by the author) that this document be changed to be consistent: the network address be left as 0.0.0.0.

The Working Group also observed that the "SNMP Security Widgets" were not defined here. It was recommended that they be included. The Working Group also observed that two transport selectors are needed, one for CLTS carried over connectionless network service, and one for COTS carried over connection oriented network service.

Further discussion evolved around the COTS, with the concern being that the description was too vague. It was also observed that COTS is a painful way to support SNMP - and a full description would likewise be painful. A poll was taken of implementation experience: one implementation under BSD.

It was suggested by the Working Group that CLTS is the architecturally appropriate way to support SNMP (see RFC1270), and that COTS should be dropped from the document.

Other contributing factors being implementation costs of COTS compared to CLTS, and interoperability issues with COTS compared with CLTS. The suggestion was carried.

Further discussion concerned whether this Working Group had the authority to make this recommendation. An informal sounding was taken of individuals present with some knowledge of other OSI efforts. A general agreement seemed to be that though CLTS was the better approach, there was potential conflict with some efforts which concentrated solely on COTS (predominantly over X.25.) The Area Director for OSI Integration suggested that this Working Group did have the authority to set such a standard, and should seize the moment to deliver the best solution. He also offered to check with the NOOP Group, which is developing OSI technology for use in the Internet, to get their reaction (expected to be positive.)

With these changes, and assuming customary time for consideration of Internet Drafts and no further controversy, the Working Group expressed its intention to propose this RFC for promotion to "Proposed Standard."

Finally the Working Group suggested that a short how-to RFC be generated to describe the checklist of issues to consider in specifying an encapsulation of SNMP. These issues are:

- Connectionless Mode Mapping
- Choosing Addresses and Sockets
- Packet Size
- Trap Pdu Network Address (0.0.0.0)
- Transport Address Representation (OCTET STRING)
- "Security Widgets" - Transport Domain OID, Default Party
- Identify Reliance on Other "Servers"
- Check Implementation Experience

Two volunteers came forward to help write this RFC. Another was identified as potentially being interested given his experience with related RFCs.

Attendees

Miriam Amos Nihart	miriam@ltning.enet.dec.com
Steve Bostock	steveb@novell.com
Theodore Brunner	tob@thumper.bellcore.com
Niels Brunsgaard	nob@dowtyns.dk
Christopher Bucci	bucci@dss.com
Philip Budne	phil@shiva.com
Lida Carrier	lida@apple.com
Peter Caswell	pfc@pacer.soft.com
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
Richard Cherry	rcherry@wc.novell.com
Carson Cheung	carson@bnr.ca
Cyrus Chow	cchow@ames.arc.nasa.gov
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Michael Grieves	mgrieves@chipcom.com
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Tim Howes	Tim.Howes@umich.edu.
William Jackson	jackson@manta.nosc.mil
Bob Jeckell	rrj@3com.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
Deidre Kostick	dck2@sabre.bellcore.com
Carl Madison	carl@startek.com
Scott Marcus	smarcus@bbn.com
Evan McGinnis	bem@3com.com
Greg Minshall	minshall@wc.novell.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
Edison Paw	esp@3com.com
David Perkins	dperkins@synoptics.com
David Piscitello	dave@sabre.bellcore.com
Mike Ritter	mwritter@applelink.apple.com

Marshall Rose	mrose@dbc.mtview.ca.us
Rick Royston	rick@lsumvs.sncc.lsu.edu
Tom Simons	nes!simons@phx.mcd.mot.com
Ursula Sinkewicz	sinkewic@netrix.nac.dec.com
Mark Therieau	markt@python.eng.microcom.com
Chris Thomas	chris_thomas@ccm.hf.intel.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jay Westmark	westmark@dsi.gw.scri.fsu.edu
Linda Winkler	lwinkler@anl.gov
Steven Wong	wong@took.enet.dec.com
Brian Wyld	brianw@spider.co.uk
Wengyik Yeong	yeongw@psi.com
Henry Yip	natadm!henry@uunet.uu.net

3.4.7 X.400 Operations (x400ops)

Charter

Chair(s):
Alf Hansen, Alf.Hansen@delab.sintef.no

Mailing Lists:
General Discussion: ietf-osi-x400ops@pilot.cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@pilot.cs.wisc.edu
Archive:

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this Group is to insure interoperability between Internet X.400 management domains and the existing Internet mail service. The specific task of this Group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

Done Initial meeting, produce internal outline.

Done Working draft, circulate to interested people.

Jul 1991 Internet Draft available.

Dec 1991 Document ready for publication.

CURRENT MEETING REPORT

Reported by Alf Hansen/SINTEF DELAB

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

Review Minutes and Liaison Report.

Liaison report: (given on March 18, 1992)

MHS-MD subcommittee of Study Group D of the US Department of State (the administrator of C=US) liaison report from Einar Stefferud (Stef).

The US backbone will exist as a virtual ADMD, all ADMDs must be able to send mail to all other ADMDs. The ADMDs, however, do not have to be directly connected to the other ADMDs. PRMD names do NOT have to be unique in the US. You can register with a service provider using ADMD=USBB (rather than the service provider's ADMD), provided that you are registered in the national registry.

Action List from Last Meeting.

- EMA (Electronic Mail Association) member present: John Sherburne, SPRINT, gave report: Liaison from EMA (non official) - Full connectivity with commercial world is important. Biggest problem is naming of domains (in particular ADMD='blank').
- Alf - Tell Working Group that mapping coordination procedures should be published as an RFC. Working Group Chair, Urs Eppenberger, at this meeting was notified, and volunteered the COSINE MHS Project Team to submit the coordination procedures as an Experimental RFC.
- All other action items done.

Review of "Routing coordination for X.400 MHS services within a multi protocol /multi network environment" by Urs Eppenberger.

Urs gave an overview of his document. The following terms were defined: WEP, MHS Community, MHS subtree.

It was noted that all WEPs must know about all other WEPs within a given community.

The Routing Coordination document was created, in part, to facilitate the connection of domains that do not share the same lower layer stacks.

Selection of the WEP is determined by the priority and delay parameters in the DOMAIN document.

- **General Discussion Starts:**

  Tony Genovese had a number of operational questions that were not answered in the document. Urs said that was not part of the scope of the document.

  Urs volunteered the COSINE MHS Project Team to be the global MHS coordinator.

  The COSINE MHS-Managers meetings will have to be funded by the regional networks after this next meeting. Tony expressed his concern that we may loose global coordination if this group goes away.

  Urs stated that there should be no more than 100 WEPs per community.

  Conclusion: Urs document looks good. This Working Group recommends that this document be published as a Draft RFC with the assumption that it will be moved to "Experimental RFC" status.

Alf's View of "Our Community".

This community consists of at least three sub-communities: The COSINE Community, North American community, Pacific Rim community.

Stef floated an idea - MIX (mail exchange point) should be created to allow mail systems to connect at a (possible virtual) central MTA. Stef also noted that "OUR" is not a good name for a community. Will probably create misunderstandings.

Conclusion: We need a single global community. If we need more sub-communities, we will deal with that when the need surfaces.

The Working Group agreed that there should be mandatory support of X.25, RFC1006 and CLNS. THIS DOES NOT IMPLY THAT EACH WEP IS REQUIRED TO SUPPORT ALL THREE STACKS. Bilateral agreements must be made where support of one of the mandatory stacks is not present. It was noted that some people may not want to go to the trouble of making these agreements. This (hopefully small) group will have to form a sub-community of our global community.

Review of "Operational Requirements for X.400 Management Domains"

The Working Group discussed the use of ADMD=<blank>. The Working Group decided to add editors note that the semantics of ADMD=<blank> are not yet understood.

Rob Hagens discussed the following comments from Working Group:

NOTE: Some editorial changes not included in Minutes: See new version of document.

Section 1.2: Profiles

Which profiles should we support? There was a request to add UK Gosip. It was decided that the section should be phrased more generically. The basic idea is that there are many different profiles. Each country will have to support their own profile.

Section 3.1.7: Domain Defined Attributes

Request to soften requirement of support of DDAs. Something to the effect of "old MTAs don't have to support DDAs; new ones MUST support DDAs" Working Group agreed that DDA support should remain mandatory.

Request to make automatic return of contents mandatory. The Working Group decided to add a recommendation to support automatic return of contents.

Global substitute "The Internet X.400 Community" for "International X.400 Service."

Add a section to the document that defines "Our X.400 Community".

Section 2.1: Management Domains

Question: "Should an MD be part of a community". The answer is yes.

It was suggested that a new section that specifies the minimum requirement of WEPs be added. Alf suggested that a separate document be created to address this issue.

Section 2.2: WEP

Should this section be re-written using the "community" concept? NO.

Last sentence: Replace "shall operate" with "shall route"

The Working Group agreed to add a statement that says that one level of OUs SHOULD be used.

Section 3.1.6: Given Name, Initials, Surname

Add a sentence that we recommend using:

1. Given name + surname OR
2. Initials + Surname

Add statement that you SHALL NOT use dots between initials.

Section 3.5: Minimum Statistics/Accounting

It was decided that we get a list of the data elements that are required.

Working Group Business:

- Review of Charter: Alf will update and send to list.
- List of documents:
  - Routing document
  - Operational requirements document
  - 7 documents from MHS-DS Working Group
  - 2 documents form MIME/MHS Working Group
  - Mapping table update procedure from Working Group
  - MAIL11 Gateway
  - DNS
  - 1148bis
  - 88/84 Downgrading
  - X.400 and International character sets.

  General discussion of problem with documents originated outside of IETF. There have been problems convincing people to publish these documents as RFCs.

- Claudio Allocchio presented his paper on "Mapping between X.400 and Mail-11"

  Stef commented that there should be a change in the way BCC is handled. Don't just treat it as a regular CC. Use the method that MH uses (enclose original message in envelope and deliver to BCC recipient).

  Claudio discussed how to handle hidden areas and different domains: DD.Dnet will contain the community name.

- Claudio presented his experiments with using DNS to store X.400 to RFC822 mapping information using DNS to store X.400 routing information.

  A number of Working Group members volunteered to put routing and mapping information into the experimental DNS subtree (under .it).

- Rob Hagens - Status of multi-stack connectivity.

  Tony Genovese noted that ESNET was very close to having a production CLNP service.

- Harald Tveit Alvestrand - International Character sets.

  Harald will turn the draft document into a draft RFC.

- Milestones

  COSINE - reports available on the COSINE fileserver (anonymous FTP from nic.switch.ch)

  ESNET- A white paper on x.400/x.500 available.

  UNINETT - Have tested X.400 to Word Perfect and Banyan gateways. Also have contracted with a company to create a user interface for PP that runs under X windows. There will also be a PC version. This will be available to all educational sites.

XNREN - Has made a fax gateway available.

The next meeting will take place during the week of the July IETF in Boston (July 13-17, 1992).

Action Items

- John Sherburne (SPRINT) will work with Tony Genovese to figure out how the US can provide an MTA that has X.25 connectivity.
- Urs will ask the COSINE MHS Project Team to submit the address mapping table procedures as a draft RFC.
- Stef - Start a discussion on X.400 OPS and Working Group lists about ADMD name in the US. See section 3.1.2.
- Alf will send the updated Charter to the list.
- Claudio will produce a draft document that will propose a method for using DNS to store X.400 to RFC 822 mapping and routing.
- Claudio will follow up the MAIL 11 mapping document.
- Harald will follow up the International Character set document.

Attendees

Claudio Allocchio	claudio.allocchio@elettra-ts.infn.it
Harald Alvestrand	harald.alvestrand@delab.sintef.no
Serge Aumont	aumont@cicb.fr
William Biagi	bbiagi@cos.com
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppen@verw.switch.ch
Jisoo Geiter	geiter@gateway.mitre.org
Tony Genovese	genovese@es.net
Robert Hagens	hagens@cs.wisc.edu
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Max Hillson	hillson@koala.enet.dec.com
Erik Huizer	huizer@surfnet.nl
Lew Jenkins	lew@enlil.premenos.sf.ca.us
Barbara Jennings	bjjenni@sandia.gov

Kevin Jordan	kej@udev.cdc.com
Eva Kuiper	eva@hpindda.cup.hp.com
Triet Lu	trietl@sparta.com
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
John Sherburne	john.sherburne@sprintintl.sprint.com
Mark Sleeper	mws@sparta.com
Einar Stefferud	stef@nma.com
Brien Wheeler	blw@mitre.org
Linda Winkler	lwinkler@anl.gov
Russ Wright	wright@lbl.gov
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.5 Operational Requirements Area

Director(s):

- Susan Estrada: estradas@cerf.net
- Phill Gross: pgross@nis.ans.net
- Bernhard Stockman: boss@sunet.se

Area Summary reported by Bernhard Stockman/NORDUnet

Operational Statistics Working Group (OPSTAT)

The OPSTAT Working Group met two times during this IETF. At the first session the document "An Internet Model for Operational Statistics" was reviewed. Only minor changes were approved and the document is now ready to be submitted as an Internet Draft.

At the second session a review was made on which NOC's are able to adopt the OPSTAT model. NOC's in the US, in Europe and in the Pacific expressed interest in participating in a test of the model.

Finally the Group discussed future activities for the OPSTAT Working Group. A client/server based retrieval system may be useful to offload routing equipment from extensive SNMP-querying and to enforce access control to statistical data.

As some of the thinking in the OPSTAT model is based on common practises there is a need for a theoretical model verifying the assumptions made. Research in this direction was presented at the BOF on Wide Area Network Measurement at this IETF. The outcome of this research may show very fruitful for the OPSTAT work.

BGP Deployment and Application BOF (BGPDEPL)

BGP deployment status

The current status of BGP deployment was reviewed. There are today around 21 regional networks using BGP towards NSFnet in production mode. Europe is actively doing a BGP pilot and some sites are already running BGP. Some of the MILnet sites are using BGP.

Drawbacks in today cisco implementation of BGP:

- Only one BGP process.
- The box choked after 9 routing process.
- BGP/EGP conflicts.
- Not a good idea to run BGP and EGP to the same AS.

Routing Policy Description

The Group recognized the need of sharing routing policy between networks in the Internet in order to avoid contradictory routing policies and therefor artificial routing disconnections. The Group discussed mechanisms to describe routing policy and share them. As a first cut, a routing policy description form will be developed. Merit will collect such forms and install them as a nis.nsf.net database.

The next step is to develop a mechanism for using the the policy database as input to a routing processor. A very interesting approach is the possibility of using configuration compilers. The idea is that input are parsable forms like the above described routing policy description which are processed into loadable configuration files.

Routing policy description processing:

```
                             Routing policies
                             from other relevant domains
                               !  !
!---------------------------!  !  !
!                           !  !  !
!                           V  V  V
!                      !----------------!
!                      ! Common Routing !<------- Query
!                      ! Policy storage !
!                      !-------!--------!
!                              !            !-------------!
!                              !    !--------! Negotiation !
!                              V    V       !------!------!
!  !---------------!   !----------------!          ! YES
!  !  Domain A's   !   ! Routing Policy !    /-----!----\
!  !    Routing    !-->!   Processor    !----< Conflict ? >
!  ! Requirements  !   !----------------!    \-----!----/
!  !---------------!                               ! NO
!                      !----------------!          !
!                      !   Domain A's   !          !
!----------------------! Routing Policy !<---------!
                       !-------!--------!
                               !
                               V
                       !----------------!
                       !  Configuration !   Configuration
                       !    Compiler    !--> File
                       !----------------!
```

By using gated it would be possible to modify sources to accommodate for increased hash tables. If a fast unix host is used gated may be used in production traffic networks. It is also possible to trace almost everything in the routing process. Finally, gated could be used in stand by mode just logging the incoming routing updates and by this be used as a routing traffic analyzer. For example when changes to the routing configuration are being installed it is possible to verify that everything works as intended before starting up gated in active mode. A version with, besides the usual roting protocols, support for IS-IS, OSPF and BGP-II/III was recently announced as an alpha-version and is ftp:able from gated.cornell.edu.

Peter Ford of Los Alamos National Laboratory (LANL) presented the latest thinking on the NSFnet future. A resolicitation will be made for the NSFnet backbone. It is foreseen that at least two different carriers will provide the basic NSFnet infrastructure.

```
!--------------------------------!
!          Carrier #1            !
!--!----!-----!-----!----!----!--!
   !    !     !     !    !    !
   O    O     O     O    O    O
   !    !     !     !    !    !
!--!----!-----!-----!----!----!--!
!          Carrier #2            !
!--------------------------------!

O = Network Access Point (NAP)
```

Regional networks will then be interconnected at one or more NAP's. This may also be the right place for networks like EBONE to hook on to the US infrastructures, especially if there will be no AUP in the NAP's but only in the backbones circuits.

Benchmarking Methodology Working Group (BMWG)

The document on "Testing methodology" was reviewed. It will need one more iteration on the body of the material before it is ready to be submitted as an Internet Draft.

There are some appendices that need more work and a video conference will be set up within one month to accomplish this.

There is a companion document on the test frame formats that shall be ready by the Boston IETF. Next on the Agenda are two things:

1. A document giving advice on how to interpret the test results.
2. Definition of host based protocol testing.

Wide Area Network Measurement BOF (WAIS)

Mike Schwartz of the University of Colorado presented part of his research in network measurements with regards to metrics and tools as compared to the resources being measured. His work includes measurements of telnet and ftp connection durations, modeling and generation of random topologies, measurement of availability and bandwidth, etc. The intention is to create a model of traffic sources, i.e., a work load model of the Internet and by this be able to predict network growth and requirements.

A PhD thesis is under preparation by Kim Claffy at the San Diego Supercomputer Center. 4 million packet headers, collected during 10pm December 24 and 2am December 25 at fix-west, were analyzed with regards to application distributions, flows, protocol distributions and performance. Resource consumption and latency behavior where investigated. Performance degradation under resource starvation was studied. A remarkable discovery was that to make an analysis with regards to application distribution, it was not necessary to collect every packet. A collection of every 10:th packet or a collection of continuous 10 packets at every 1000:th packet gave almost identical patterns. Diagrams were presented on interagency traffic and network to network traffic.

CURRENT MEETING REPORT

Reported by Paul Tsuchiya/Bellcore

Minutes of the Routing Table Lookup Algorithm BOF (RTGTBL)

The purpose of the meeting was simply to discuss the various known approaches to software-driven routing table lookup algorithms. There is no intention to meet again, and no documentation is expected. Presentations were given by the following:

Paul Tsuchiya	Cecilia Algorithm
Joel Halpern	Patricia Algorithm
Keith Sklower	Radix Algorithm (in Unix BSD)
Rob Coltun	Binary Search Algorithm
Fred Baker	16-wide Radix Algorithm
Dino Farinacci	Hash Algorithm

Very briefly, Cecilia, Patricia, Sklower Radix, and 16-wide radix are all radix-type algorithms, meaning that they follow the search tree by branching according to the value of certain bits of the "key" (the address being looked up). The three former algorithms check one bit at a time, while the 16-wide checks 4 bits at a time. The three former algorithms can check bits in any order, while the 16-wide checks bits strictly MSB to LSB.

The Binary Search Algorithm does a greater than/less than compare to determine how to traverse the search tree. The Hash Algorithm is not a tree-based search algorithm as the first five are, and won't be further discussed in these Minutes.

Of the six algorithms, only Cecilia and Patricia handle non-contiguous masks efficiently (meaning, a tree of small size and depth). Sklower's Radix handles non-contiguous masks, but at the expense of a larger depth. Cecilia has an efficient delete operation, whereas Patricia does not. Patricia uses roughly one half the memory of Cecilia. Note that in most router applications, a delete operation is not that important because routes don't often appear and then disappear forever. Occasionally reforming the tree from scratch for the purpose of garbage collection suffices. All of the algorithms can handle contiguous masks, but the hash algorithm performance decreases as the number of different sized masks increases.

Of the first five algorithms, only the Binary Search is balanced. However, note that this is only with respect to the number of elements in the tree, not with respect to the frequency with which each element is searched. Usually the 16-wide Radix will have the smallest depth, because it covers 4 bits with a single compare rather than just 1. However, since it goes strictly left to right, if the left bits do not differ in any of the routing table entries, the 16-wide Radix can be deeper than the first 4 algorithms.

Attendees

Fred Baker	fbaker@emerald.acc.com
David Bridgham	dab@epilogue.com
Scott Brim	swb@cornell.edu
Dean Cheng	dean@sun2.retix.com
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Osmund de Souza	osmund.desouza@att.com
Dino Farinacci	dino@cisco.com
Vince Fuller	vaf@stanford.edu
Joel Halpern	jmh@network.com
Jeffrey Honig	jch@risci.cit.cornell.edu
Kathleen Huber	khuber@bbn.com
Bob Jeckell	rrj@3com.com
Bill Manning	bmanning@rice.edu
Milo Medin	medin@nsipo.nasa.gov
Andy Nicholson	droid@cray.com
Stephanie Price Marasciullo	price@cmc.com
Paul Raison	raison@xylogics.com
Guenter Roeck	roeck@conware.de
Keith Sklower	sklower@cs.berkeley.edu
Frank Solensky	solensky@clearpoint.com
William Streilein	bstreile@wellfleet.com
Paul Traina	pst@cisco.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Kannan Varadhan	kannan@oar.net
Carol Ward	cward@westnet.net
Linda Winkler	lwinkler@anl.gov
Jessica Yu	jyy@merit.edu

CURRENT MEETING REPORT

Reported by Kim Claffy/SDSC

Minutes of the Traffic Collection, Measurement and Characterization BOF (TRAFCHAR)

Introduction

The purpose of this BOF is to initiate discussion and information exchange within the community concerning research in wide-area network traffic measurements. Five brief presentations of related research were made, followed by discussion of each.

One theme of the BOF was to discuss exactly what kind of network instrumentation, measurement facilities, and types of measurements should be recommended to the Internet community. Many of us would like to encourage the managers of stub networks and routers to collect and make available information similar in spirit to the statistics that NSFnet makes available through Merit/NSFnet Information Services (NIS.NSF.NET). We hope this effort eventually evolves into an RFC, and eventually leads to a widespread cooperative effort. We freely admit that the road to success will be an iterative process, fraught with plenty of challenging technical details.

The amount of space consumed by this data completely depends on the type of measurement. For example, collecting TCP SYN/FIN/RST packets could lead to hundreds of megabytes a day, depending on the collection site. Other methods, like sampling or recording the quantity of bytes sent to particular destination networks might require less than a hundred kilobytes a month. In the first case, the volume of trace data can be on the order of one to two percent of the traffic itself, with the resulting data possibly having to be sent by tape rather than electronic means to the location where the network analysis will happen.

The Internet Activities Board (IAB) recently announced guidelines for measurement activities. RFC 1262 lists bounds that should be commonly acceptable. However RFC 1262 directly addresses invasive measurement activities, and is only marginally applicable to passive data collection. We believe we will have to face many new issues hitherto unaddressed. What we propose must honor the concerns and restrictions that individual networks may impose, yet be thorough enough to capture the data that we need to accomplish the research goals. It should also allow for flexibility. An example of a difficult issue to resolve is privacy when using network addresses, in particular as workstations with their own IP addresses frequently map to individual users. Our efforts should address privacy measures, that still allow professional research to be conducted.

Most likely, each of us has a different idea as to the data we need to have measured to achieve our various objectives. Below, we summarize these motivations and give a preliminary list of the measurements and trace data that we believe should be collected or capturable.

Everyone is encouraged to add to both the motivation list and chart of traces and measurements, and mail them back to wanchar@usc.edu for inclusion in this document.

Motivations

- Artificial Workload Models (Danzig and Jamin)

  Good artificial workload models are needed to drive simulations of new resource management algorithms, flow control algorithms, and routing algorithms. The artificial workload models that we are developing consist of an application specific model (ftp, telnet, nntp, etc.) and an application arrival rate model that is stub network dependent. So far we have been able to identify applications from their port numbers. As new transport protocols emerge, we may need other mechanisms. Creating the application specific model requires full traces of TCP/IP packet headers. Creating the stub network specific model requires traces of TCP SYN/FIN/RST packets only. Most of our data has been collected with statspy or tcpdump from a machine on the same Ethernet segment as the stub network's gateway to the backbone. We would like to collect SYN/FIN/RST traces from hundreds of stub networks. Given current network bandwidth and usage, these traces can range to 200MB/day.

- Network Planning (Braun and Claffy)

  SDSC and UCSD are undertaking a network analysis effort with multiple goals of immediate applicability and interest to the Internet environment, with respect to both performance and ubiquity.

  Areas of current investigation include: measurements and analysis of resource consumption and latencies, network performance degradation under resource starvation, and end-to-end performance testing. We have determined, for selected data sets, characteristics of network usage by application, bandwidth requirements, and geographic distribution. We are also exploring the role that granularity plays in traffic analysis, both in statistical sampling of traffic on an operational basis, and in the level of detail one presents data to optimize the information/noise ratio.

  We are currently analyzing data from a variety of sources, including national networks as well as federal network interconnection points of multiple agencies. Statistical examination and manipulation of data reveals significant traffic correlations, trends, and dependencies.

  We are also undertaking collaborative efforts with Toshiya Asaba and the WIDE statistics working group in Japan. In particular, Asaba is largely responsible for the analysis scripts which facilitated statistical examination and data presentation. We first intended the scripts for use in a study of international traffic between Japan and other nations. We were able to adapt the script for use in subsequent studies. Building a public library of usable scripts for different analysis tasks requires agreement on data formats in multiple phases of collection and analysis. We would like to see a collaborative effort within the community toward accomplishing such a task.

Further information and slides are available by sending requests to the SDSC Applied Network Research Group, via hwb@sdsc.edu or kc@sdsc.edu

- Stateful Router Studies (Estrin and Mitzel)

  [Related information, though not participated at the BOF.]

  The current Internet is based on a stateless (datagram) architecture. However, many recent proposals rely on the maintenance of state information within network routers, leading to our interest in the implications of a "stateful" network layer. We wish to collect internetwork traffic traces at the border routers of stub and transit networks, and use this data to evaluate, or predict, the effects of design alternatives for stateful architectures.

  An important design decision is the level at which conversations are defined. This determines the granularity of control over the network traffic, and affects the scalability of the system. We are interested in several granularities of conversations, ranging from a single TCP application association, up to aggregation of all traffic between two communicating networks. We will use the data to estimate the number of active conversations at a router, and derive the storage requirements for the associated conversation state table. We will analyze the feasibility of fine grain control at the network periphery and deeper within the network.

  In conventional IP, the only lookup function normally required for packet forwarding is a routing table lookup. This has been recognized as a bottleneck in the forwarding process [Feldmeier, Jain]. It has been shown that the introduction of an LRU cache can substantially improve the efficiency of the packet forwarding process. Route caching is used in many existing routers. However, unlike the stateful schemes investigated here, which require lookup based on source–destination pairs, current route caches are based only on destination host or network. It is not intuitively obvious whether the solutions developed for routing table caches can be applied here. We will use our network traffic traces to perform trace driven simulations of an LRU cache, for different conversation granularities, and thereby assess traffic locality and the benefits of caching.

- Network Monitoring (Schwartz and Pu)

  Schwartz proposed that a group of a dozen of us or so agree to collaborate to collect traces and measurements. He also described his recent study of FTP traffic which showed that tools to locate copies of large, replicated files may reduce wide area network traffic due to FTP. The unique aspect of Schwartz's traces was that it actually peered at application level data in a way that preserved privacy.

- Host Reliability and Availability (Long)

  Long summarized his study of internet host reliability and availability. This was the only active form of tracing discussed during the BOF.

Measurements and Traces

Here is a first pass at the type of data we would like to see collected, and what studies would use this data. These categories need to be detailed, and new categories probably need to be filled in. The table identifies four types of data to collect. These include captured packets and packet headers (excluding data), headers of selected packets, summary data, and routing and congestion data. The first three types of data are pretty well defined, while the last is much less so. Although we can collect such data from anywhere in the Internet, we classify it into three classes: entrances to stub networks, regional and backbone routers, and international gateways.

TYPE OF DATA

MEASUREMENT POINT	Captured Packets & Packet Headers	TCPDUMP Conversation SYN/FIN/RST Traces	NSF.NIS.NET LIKE DATA Data	Router Timing and Queue length (MIB)
STUB NETWORKS	Workload models Workload Planning	Workload models	Workload Planning	Congestion studies
REGIONAL AND BACKBONE NETWORKS	Stateful Routers Workload Planning		Workload Planning	Congestion studies
INTER-NATIONAL GATEWAYS	Workload Planning		Workload Planning	Congestion studies

Table 1.

Trace Formats and Tools

We need to define the storage format for trace and statistical data. For some formats, like tcpdump or statspy, the format is already pre-defined. Almost certainly we should adopt NSFnet's current format for the type of data they collect. We also need to define "sanitizer" programs that implement the security concerns of particular networks.

There is an operations area in IETF which has been defining some standard transport and storage formats for various kinds of operational data.

Dealing with gigabytes of data results in a serious resource impact. An effort has to be undertaken to identify schemes to make such large quantities of data useful, possibly via multiple levels of data reduction.

Attendees

Andy Adams	ala@merit.edu
Vikas Aggarwal	aggarwal@jvnc.net
Cynthia Bagwell	cbagwell@gateway.mitre.org
Hans-Werner Braun	hwb@sdsc.edu
J. Nevil Brownlee	nevil@aukuni.ac.uz
Eric Brunner	brunner@practic.com
Randy Butler	rbutler@ncsa.uiuc.edu
Douglas Carson	carson@utcs.utoronto.ca
John Chang	jrc@uswest.com
Bilal Chinoy	bac@sdsc.edu
Kimberly Claffy	kc@sdsc.edu
Henry Clark	henryc@oar.net
Robert Cole	robert.g.cole@qsun.att.com
Peter Danzig	danzig%warthog.usc.edu@usc.edu
Peter Deutsch	peterd@cc.mcgill.ca
Tom Easterday	tom@cic.net
Erik Fair	fair@apple.com
Kevin Fall	kfall@ucsd.edu
Joseph Godsil	jgodsil@ncsa.uiuc.edu
Kenneth Goodwin	goodwin@psc.edu
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Olafur Gudmundsson	ogud@cs.umd.edu
Ronald Jacoby	rj@sgi.com
Mark Kepke	mak@cnd.hp.com
Eliot Lear	lear@sgi.com
Susan Lin	suelin@ibm.com
Kenneth Lindahl	lindahl@violet.berkeley.edu
Daniel Long	long@nic.near.net
Pushpendra Mohta	pushp@cerf.net
Donald Morris	morris@ucar.edu

Chris Myers	chris@wugate.wustl.edu
K. Ramakrishnan	rama@erlang.enet.dec.com
Robert Reschly	reschly@brl.mil
Kary Robertson	kr@concord.com
Manoel Rodrigues	manoel.rodrigues@att.com
Gregory Ruth	gruth@bbn.com
Anthony Rutkowski	amr@sprint.com
Andreas Schachtner	afs@germany.eu.net
Michael Schwartz	schwartz@cs.colorado.edu
Erik Sherk	sherk@sura.net
Mark Sleeper	mws@sparta.com
Bernhard Stockman	boss@sunet.se
Claudio Topolcic	topolcic@nri.reston.va.us
Huyen Vu	vi@polaris.dca.mil
David Waitzman	djw@bbn.com
Timothy Walden	tmwalden@saturn.acc.com
Carol Ward	cward@westnet.net
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Frances Yeh	fyeh@rafael.jpl.nasa.gov

3.5.1 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 enviroment 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.

Mar 1993 Post an Internet Draft defining a mechanism to share policy information between Administrative Domains.

Dec 1992 Post as an Internet Draft, a report of BGP deployment status.

CURRENT MEETING REPORT

Reported by Jessica Yu/Merit

Minutes of the Border Gateway Protocol Deployment and Application Working Group (BGPDEPL)

Agenda

- BGP Deployment Status and Issues
- Routing Policy Sharing Mechanism
- Next Phase NSFnet Architecture Discussion

BGP Deployment Status and Issues

There are more sites/regionals using BGP in production or testing since the last IETF meeting when the Group started. There are twenty-two regionals/sites (it was eight during the last IETF) connecting to NSF/ANSnet which are using BGP in production mode. Most of them are using cisco BGP2 (BGP version 2) with their collocated NSS/ENSS. Two or three sites are using gated with BGP1 or BGP2.

MILnet started beta test BGP3 on T-20 router. There are six or seven MILnet sites using cisco to BGP3 with the T-20 router.

The BGP pilot project for the Pan-European network (Ebone) is making progress. There are BGP peer sessions within the European network. There is also a lot of BGP testing taking place within European networks.

- New developments in BGP implementation
  - Gated supports BGP2 and BGP3 now. It is in alpha version. CA*net is using this version to BGP2 with the NSS.
  - cisco has 9.0 beta which supports BGP2 and BGP3.
  - Tony Li, at cisco, did an interoperate test between the cisco and other types of routers with BGP implementation. The result is listed in Appendix C.
  - Matt Mathis of PSCnet gave a presentation of "BGP usage at PSC". His slides are included at the end of these Minutes.

Routing Policy Sharing Mechanism

The Routing Policy Sharing Mechanism was discussed at the meeting. The purpose of developing such a mechanism is for each network sharing its routing policy, to keep a global routing view so each network can design and implement their policy routing to avoid being inconsistent with the global routing.

We started by developing a routing policy description form. It will be evolved to an information base to install this kind of information and make it available for general access. It will also define the processor to process the information and generate inconsistency should it occur. At the meeting, an initial draft routing policy form was discussed and agreed upon. The routing form is included in Appendix B. Merit will create a directory on host nis.nsf.net to install this form for anonymous ftp. Merit will also create a mailing list for people sending the filled forms and will install them in the directory for anonymous ftp as well.

The next step is to develop a mechanism for using the the policy database as input to a routing processor. A very interesting approach is the possibility of using configuration compilers. The idea is that the routing policy descriptions are in a parsable form, which can be processed into loadable configuration files.

Routing policy description processing:

```
                             Routing policies
                             from other relevant domains
                               !  !
!---------------------------!  !  !
!                           !  !  !
!                           V  V  V
!                      !----------------!
!                      ! Common Routing !<------- Query
!                      ! Policy storage !
!                      !-------!--------!
!                              !             !-------------!
!                              !    !--------! Negotiation !
!                              V    V        !------!------!
!  !---------------!   !----------------!           ! YES
!  !  Domain A's   !   ! Routing Policy !     /-----!----\
!  !   Routing     !-->!   Processor    !----< Conflict ? >
!  ! Requirements  !   !----------------!     \-----!----/
!  !---------------!                                ! NO
!                      !----------------!           !
!                      !   Domain A's   !           !
!----------------------! Routing Policy !<----------!
                       !-------!--------!
                               !
                               V
                       !----------------!
                       !  Configuration !   Configuration
                       !    Compiler    !--> File
                       !----------------!
```

Next Phase NSFNET Architecture Discussion

Peter Ford led a discussion of the NSFnet recompetition architecture. His report is included in Appendix A below.

Appendix A: NSFnet recompetition, architectural considerations.

The National Science Foundation (NSF) reported on architectural considerations for the NSFnet recompetition which will occur during mid-1992.

Leading the discussion were:

- Robert Aiken (NSF)
- Hans-Werner Braun (SDSC)
- Peter Ford (LANL)
- Stephen Wolff (NSF)

NSF expressed thanks to the many people who provided input into the process of developing an architectural model for the future NSFnet including Merit, the FEPG, the operators of the U.S. regional networks, the FARNET membership, and the Intercontinental Engineering and Planning Group members.

The architecture discussion focused on an interconnection architecture for networks which either provisioned research and education networks or for those who needed to connect to the research and education networks. NSF stated that they believed that it is critical for the U.S. R & E networks to get provisioned out of the growing telecommunications base so that they could take advantage of the economies of scale available in the telecommunications industry. This should allow for maximizing the flexibility in choosing within a parameter space defined by bandwidth, cost, reliability, etc.

The model was based on the notion of "network access points" (NAPs) which would be a chunk of shared media where networks (regionals, national, international, multination corporate, etc.) could interconnect. NSF will provide support for developing a management and routing coordination function at the NAPs. This would be done through a collaborative agreement under the name of "Routing Arbiter". The Routing Arbiter would provide routing support at the NAPs through the use of a route server box which could peer with the attached networks and send an "homogenized" picture of the routing topology. This picture would be dependent on the picture a network would want to see as previously negotiated with the Routing Arbiter. The NAPs would be the focal point for evolving the internet architecture as new technologies become available and need to be incrementally introduced. The NAPs would be "NSF Acceptable Use Policy free" so commercial traffic from a U.S. regional could go onto the NAP in the process of going to another network which carries commercial traffic.

There would be greater than four and less than seventeen NAPs and they would be geographically dispersed. The NSF recompetition will have at least two providers for national scale R & E traffic and they will connect to all the NAPs. The providers are free to connect to things other than NAPs, including regionals.

There was an extensive question and answer session, where the core topics included:

- Asymmetric routes and the problems these posed (Vince Fuller Stanford/Barrnet, Matt Mathis/PSC, Milo Medin/NASA).
- The issue of who could connect to a NAP (Dan Long, John Curran BBN/NEARnet).
- How the routing arbiter would be managed, with emphasis placed on the observation that the current Merit/ANS/NSF configuration did not have the regional networks in a position where they were the customers of the NSFnet backbone (Scott Bradner, Harvard/NEARnet). Steve Wolff commented that he would like to hear input on this topic.
- Acceptable Use Policy issues.
- Would regionals have to connect to a NAP? – No.
- Overall management of the 2+ providers and the NAPs?

NSF also stated that the need for maintaining the routing database that Merit currently manages would fall under the auspices of the Routing Arbiter.

The session was well attended and the discussions were vigorous. The NSF would like to thank Jessica Yu for allowing them to barge into the BGP Deployment Group meeting on short notice, and would like to encourage any interested parties to contact the NSF with ideas, questions, and concerns (steve@nsf.gov, raiken@nsf.gov, peter@lanl.gov).

Appendix B: Routing Policy Description Form (4/6/92)

```
A. General Description of your Autonomous System (AS) *

a. List the name of your AS and the AS number(s)

b. List the Routing administrator of the AS

            Name:
            Organization:
            E-mail:
            Phone number:

c. List the networks that belong to your AS and mark them with RE or
   non-RE type if applicable
```

```
d. List the name of your direct neighbor AS(s) and its AS number(s)

e. List each IP address of your border router(s) which interface with your
   neighbor border router(s) and the Exterior Routing Protocol(s) used
   respectively.

f. Is your AS a
                Stub AS?
                Multi-homed Stub AS?
                Transit AS?
                Pure Transit AS?

g. List the IGP used within your AS (optional for a non-transit AS)

h. Describe the maximum and minimum bandwidth of the transit portion of your
   AS (optional for a non-transit AS).

i. Describe the delay characteristic of physical links of transit portion
   of you AS, e.g., satellite, terrestrial (optional for a non-transit AS).
```

B. Policy Descriptions

a. For all the AS's.

- Outbound advertisement filtering:
  1. list the set of nets belong to your AS that you do not advertise to your neighbor(s).
  2. list the set of nets belong to your AS that you do not wish to be advertised to certain ASs. List the AS numbers.
- Inbound acceptance filtering: list the set of ASs whose nets you do or (do not) accept from your neighbor(s) If AS number is not a satisfactory granularity, list the set of nets.
- Describe your routing policy based on your Acceptable Use Policy (AUP):

  Does your AS accept:
  1. All types of traffic?
  2. Only RE type traffic?
  3. Other? (please specify)
- Does your border router default to your neighbor border router? If yes, described the mechanism.

b. For Multi-homed Stub ASes only:

- List the preference/denial of the neighbor ASs which can route your traffic to the same destination (Multi-homed AS only).

c. For Transit AS only:

- List the paris of source/destination ASs or nets that your domain does not provide transitivity.
- List the preference/denial of the ASs which can route your traffic to a certain destination.

Note:

* The notion of Autonomous System (AS) here means the following:

a. An AS is a network blob which has a coherent Interior routing plan and under single administration.

b. The AS number will be in the BGP AS path

Appendix C: BGP Interoperability Matrix

```
Versions tested:
a) cisco 8.2,8.3 - v2
b) cisco 9.0 (beta) - v2,v3
c) BBN T/20 2.0 - v2,v3
d) NSS/eNSS (???) - v1,v2
e) gated 2.x - v1
f) gated (alpha) - v2,v3
```

	a	b	c	d	e	f
a	same					
b	ok	same				
c	ok	ok	same			
d	ok	ok	?	same		
e	vers	vers	?	ok	same	
f	ok	ok	?	ok	vers	same

```
Status:
ok - interoperability tested, found to work
same - same version, interoperability assumed
vers - lack common version
? - unknown

Tests:
Establish connection.
Negotiate version (if applicable).
Exchange routes.
```

Thanks to Bernhard Stockman for providing his note to this summary.

Attendees

Andy Adams	ala@merit.edu
Robert Aiken	raiken@nsf.gov
Nagaraj Arunkumar	nak@3com.com
William Barns	barns@gateway.mitre.org
Jordan Becker	becker@nis.ans.net
David Bolen	db31@nis.ans.net
Scott Bradner	sob@harvard.edu
Ronald Broersma	ron@nosc.mil
Susan Calcari	calcaris@cerf.net
John Chang	jrc@uswest.com
Henry Clark	henryc@oar.net
Richard Cogger	rhx@cornell.cit.bitnet
James Conklin	conklin@bitnic.educom.edu
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Osmund de Souza	osmund.desouza@att.com
Billy Durham	durham@fishery.honeywell.com
Tom Easterday	tom@cic.net
Susan Estrada	estradas@cerf.net
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Jisoo Geiter	geiter@gateway.mitre.org
Elise Gerich	epg@merit.edu
David Getchell	dag@llnl.gov
Terry Gray	gray@cac.washington.edu

Keunhee Han	khhan@gaya.kreonet.re.kr
Susan Hares	skh@merit.edu
Kenneth Hayward	crm57d@bnr.ca
Dan Jordt	danj@nwnet.net
Daniel Karrenberg	daniel.karrenberg@cwi.nl
Dave Katz	dkatz@merit.edu
Thomas Kile	tgkile@brl.mil
Mark Knopper	mak@merit.edu
Rajeev Kochhar	rxk@3com.com
Eliot Lear	lear@sgi.com
Linda Leibengood	ldl@ans.net
Tony Li	tli@cisco.com
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
Bill Manning	bmanning@rice.edu
Matt Mathis	mathis@psc.edu
Milo Medin	medin@nsipo.nasa.gov
Ed Menze	menze@cs.arizona.edu
Donald Merritt	don@brl.mil
Dennis Morris	morrisd@imo-uvax.dca.mil
Lars Poulsen	lars@cmc.com
Yakov Rekhter	yakov@watson.ibm.com
Robert Reschly	reschly@brl.mil
Tom Sandoski	tom@concert.net
Andreas Schachtner	afs@germany.eu.net
Tim Seaver	tas@concert.net
Brian Shiflett	rcollett@icmi.icp.net
Pyda Srisuresh	suresh@3com.com
Bernhard Stockman	boss@sunet.se
Sally Tarquinio	sally@gateway.mitre.org
Paul Traina	pst@cisco.com
Kannan Varadhan	kannan@oar.net
Huyen Vu	vi@polaris.dca.mil
Carol Ward	cward@westnet.net
Chris Wheeler	cwheeler@cac.washington.edu
Stephen Wolff	steve@nsf.gov
Jessica Yu	jyy@merit.edu

BGP Usage at PSC

Matt Mathis
mathis@psc.edu

Outline

- Context
- Prototype Configuration
- Routing Policy
- Configuration compiler
- Gated "Nanny"
- Routing Design

Context: Other Projects at PSC ca. Jan 1991

- Ultranet MAN
- Cray YMP / HIPPI / TMC CM-2 800 Mb/s distributed application
- Build up local infrastructure from Ethernet/T1 to FDDI/fiber extender (25km) and T3
- Demo @ National Net '91 with LBL Real-Time Systems Group (Van Jacobson) and Visualization Group
  - TCP large window "mod" to Unicos
  - NSC DX8130 router deployed (Cray LSC to FDDI)
  - FDDI Single Mode Fiber Extenders
  - DEC DS5000 FDDI + gated
  - NSFnet ENSS FDDI
  - BGP
  - Test ALL at Once!

Fig 1: LBL Demo Overview

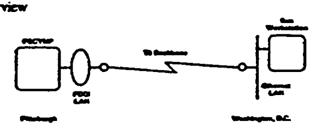

Prototype Configuration

Fig 2: LBL Demo Plan A

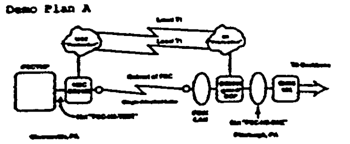

Fig 3: LBL Demo Plan B

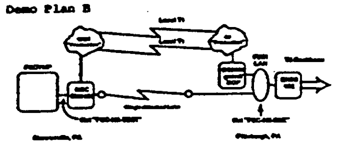

Routing Policy

- Whaler Primary for PSC, Secondary for PREPnet. Peers with ENSS13
- Oday Primary for PREPnet, Secondary for PSC. Peers with ENSS132
  Full load splitting and redundancy between Whaler and Oday.
- PSC-GW1 "last resort" - vestigial from T1 backbone.
  Generates high metric default if EGP up with T1 backbone.
- Whaler and Oday import all T1 routes (3500+ nets) via BGP.
- Import all T3 routes except from T1 Crossovers (1350 nets) via BGP
  Always prefer T3 over T1.
- Generate Default if BGP up with T3 (or T1).
  Install static route to peer (gated code change).
- Complete per net, per interface filters for All interior nets (including PREPnet).

Fig 4: PSC Internal Topology

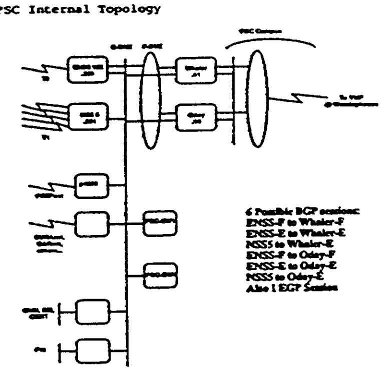

Configuration (gated) Compiler

Why?

- Assures consistent Configuration.
- Trivial to make Global policy changes.
- Support standby mode for in-place testing.
- Whaler, Oday, Comet, Bunter, Snipe, etc. are all clones of each any can be used to spare another.

Fig 5:

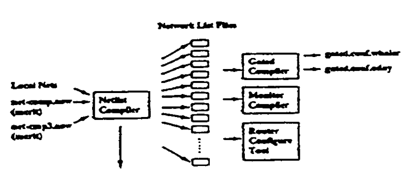

Eventless Operation

- Recompile gated.conf for standby.
- Restart gated - traffic moves elsewhere.
- Wait for idle.
- Stop gated and perform system maintenance.
- Restart standby gated, inspect routing.
- Restart production gated, inspect traffic.

Can be performed any time except during peak load.

Gated Nanny

- Attempt to open RIP listener port.
- If successful, gated is dead - kill and restart.

Routing Design

BGP Downside
- New - Still bugs as of March, 1991.
- Can't hide kills from peers.

BGP Upside
- Large Routing Tables.
- Fast Convergence.
- "Silent" reroute between peers.

3.5.2 Benchmarking Methodology (bmwg)

Charter

Chair(s):

Scott Bradner, sob@harvard.edu

Mailing Lists:

General Discussion: bmwg@harvisr.harvard.edu
To Subscribe: bmwg-request@harvisr.harvard.edu
Archive:

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the Group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

Done	Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.
Done	The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.
TBD	Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

Request For Comments:

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices"

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

Minutes of the Benchmarking Working Group (BMWG)

The Benchmarking Methodology Working Group met on Wednesday, March 18th.

The testing portion of the current draft of the testing methodology document was reviewed and a number of small changes were suggested. A revised version will be submitted as an Internet Draft before the July meeting.

Attendees

Nagaraj Arunkumar	nak@3com.com
Scott Bradner	sob@harvard.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Henry Clark	henryc@oar.net
Kathleen Dodd	kathy@gateway.mitre.org
Constance Fleenor Lloyd	cfleenor@lloyd.com
Kenneth Goodwin	goodwin@psc.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Alisa Hata	hata@cac.washington.edu
George Kajos	kajos@coral.com
Hock-Koon Lim	lim@po.cwru.edu
Carl Madison	carl@startek.com
Rina Nathaniel	rina!rnd!rndi@uunet.uu.net
K. Ramakrishnan	rama@erlang.enet.dec.com
Robert Reilly	rreilly@athena.mit.edu
Ron Roberts	roberts@jessica.stanford.edu
Miguel Sasson	sasson@xylogics.com
Tom Simons	nes!simons@phx.mcd.mot.com
Michael Smith	mes5674@hertz.njit.edu
Walter Stickle	wls@ftp.com
Bernhard Stockman	boss@sunet.se
Glenn Trewitt	trewitt@pa.dec.com
Frances Yeh	fyeh@rafael.jpl.nasa.gov
Bahman Zargham	zargham_bahman@tandem.com

3.5.3 Network Joint Management (njm)

Charter

Chair(s):
Gene Hastings, hastings@psc.edu

Mailing Lists:
General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu
Archive:

Description of Working Group:

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the Internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the Agenda of this Group will include the IAB and its task forces, and groups within FARNET. In particular FARNET has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the Group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing.

Goals and Milestones:

None specified

CURRENT MEETING REPORT

Minutes of the Network Joint Management Working Group (NJM)

Report not submitted.

Attendees

Andy Adams	ala@merit.edu
Vikas Aggarwal	aggarwal@jvnc.net
Scott Bradner	sob@harvard.edu
David Carrel	carrel@munin.com
Douglas Carson	carson@utcs.utoronto.ca
Henry Clark	henryc@oar.net
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Robert Gutierrez	gutierre@nsipo.nasa.gov
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Daniel Long	long@nic.near.net
Pushpendra Mohta	pushp@cerf.net
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Ron Roberts	roberts@jessica.stanford.edu
Tom Sandoski	tom@concert.net
Mark Sleeper	mws@sparta.com
Sandra Vest	vests@imo-uvax.dca.mil
Carol Ward	cward@westnet.net
Chris Wheeler	cwheeler@cac.washington.edu
Linda Winkler	lwinkler@anl.gov
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.5.4 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, boss@sunet.se
Phillip Gross, pgross@nis.ans.net

Mailing Lists:

General Discussion: oswg-l@wugate.wustl.edu
To Subscribe: oswg-l-request@wugate.wustl.edu
Archive:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The Working Group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Done	Agreement on a model.
Done	Survey for most useful and popular metrics.
Done	Survey for most useful and popular presentation formats.
Dec 1990	Identify similar efforts being performed by other groups.
Done	Define a common minimal set of metrics.
Mar 1991	Propose a MIB for metrics not already there.
Done	Define a common storage format to facilitate data sharing.
Done	Define common presentation formats to make data comparable.
Mar 1991	Develop outline, and make writing assignments for paper (Opstat1) documenting March 91 milestones.
May 1991	Complete paper Opstat1.

May 1991 Possible mid-term meeting to review Opstat1.

May 1991 Submit Opstat1 as Internet Draft.

Jul 1991 Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?

Jul 1991 Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.

Jul 1991 Propose old tools to be retrofitted.

Jul 1991 Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools.

Sep 1991 Complete paper Opstat2.

Sep 1991 Possible mid-term meeting to review Opstat2.

Sep 1991 Submit Opstat2 as Internet Draft.

Dec 1991 Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Bernhard Stockman/NORDUnet

Minutes of the Operational Statistics Working Group (OPSTAT)

Tuesday's Session

1. Review of the OPSTAT Document.

   The document on "A Internet Model for Operational Statistics" was reviewed. Decisions were made regarding the following:

   - In listings of MIB variables suggested for gathering the name shall be the last part of the fully qualified MIB name as long as this uniquely defines the variable.
   - A suggestion of adding InErrors and OutErrors was not approved as these variables are counted differently in different MIB implementations.
   - Valid types for bandwidths and protocol types should be explicitly enumerated.
   - The timezone part of the timestamp in the datasection is redundant and is moved to the device section.
   - It shall be more clearly stated that document gives recommendations and polling and saving periods are not mandatory.

   The Group decided that after these changes have been included the document shall be submitted as an Internet Draft. Additional suggestions, which would make major changes necessary, can be presented, discussed and decided on during the six month Internet Draft reviewing period.

2. SQL Database

   The Group discussed possible use of SQL database technique. As SQL is optimized for other purposes than retrieval of serialized data the meeting concluded that SQL was not appropriate to use in retrieval of statistical data.

3. Implementations of the OPSTAT Model

   The Group made a review of NOC's prepared to implement the operational statistical model when the Internet Draft is submitted. Below are NOC's who were represented at this meeting and who expressed interest:

   - Merit (Dale Johnsson)
   - New Zealand (Nevil Brownbe)
   - RIPE NCC (Daniel Karrenberg)
   - EUnet (Daniel Karrenberg on behalf of Joy Marino)

4. Need of Theoretical Framework

   The Group discussed the need for a theoretical statistical framework as current thinking to some extent is based in practical experiences. Kim Cluffy, SDSC, is writing a PhD thesis on analysis of wide area network. The output may have a lot in common with the OPSTAT work. Bellcore has developed models for statistical analysis of data from packet switch networks which could also show beneficial for the OPSTAT work. Contact shall be made with Bellcore to investigate if their models might also fit in the Internet architecture.

Wednesday's Session

1. Review of Future Activities.

   At an early stage the OPSTAT Group discussed possibilities of using a client/server based system for retrieval of statistical data. The Group agreed that such a model would be useful to offload network equipment from SNMP processing and to enforce access control of statistical data.

   Some similar system may already exist. For example, a system developed by DEC was mentioned.

2. Review of Existing Tools.

   NASA has tools that are currently configured for 1 minute polling. The 1 minute polls are stored internally in the program and 15 minutes average and peaks are being stored onto secondary storage. The NASA statistical tools may be made publically available.

   Milo Medin, NASA, expressed the need for differentiating between statistical tools and monitoring tools. SNMP access is not the same as access to statistical data. A client/server based system may show useful in accessing logged data when there is no "public" SNMP access. Another method would be to give access to the tables and diagrams produced from statistical data.

   OARnet has tools that currently do weekly loggings of the data. The tools are more oriented towards logging of error conditions.

   RIPE NCC has tools developed from the ISODE SNMP code using gawk with SNMP capabilities. These tools are already adopted to the OPSTAT thinking and changes to reflect the lastest storage formats may easily be included. Daniel Karrenberg gave a short presentation of these tools.

   RICE University is using enhanced Merit tools. The tools are written for AIX and SUN Sparc and could be made publically available.

Attendees

Vikas Aggarwal	aggarwal@jvnc.net
James Alfieri	jdal@troy.cc.bellcore.com
Miriam Amos Nihart	miriam@ltning.enet.dec.com
Bill Bliss	billbl@microsoft.com
J. Nevil Brownlee	nevil@aukuni.ac.uz
Henry Clark	henryc@oar.net
Tom Easterday	tom@cic.net
Stefan Fassbender	stf@easi.net
Kenneth Goodwin	goodwin@psc.edu
Dale Johnson	dsj@merit.edu
Linda Leibengood	ldl@ans.net
Hock-Koon Lim	lim@po.cwru.edu
Bill Manning	bmanning@rice.edu
Pushpendra Mohta	pushp@cerf.net
Chris Myers	chris@wugate.wustl.edu
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com
Brian Shiflett	rcollett@icmi.icp.net
Ursula Sinkewicz	sinkewic@netrix.nac.dec.com
Frank Solensky	solensky@clearpoint.com
Bernhard Stockman	boss@sunet.se
David Waitzman	djw@bbn.com
Allen Wiersma	wiersma@nsipo.nasa.gov
Frances Yeh	fyeh@rafael.jpl.nasa.gov

3.5.5 User Connectivity (ucp)

<u>Charter</u>

Chair(s):

Dan Long, long@nic.near.net

Mailing Lists:

General Discussion: ucp@nic.near.net
To Subscribe: ucp-request@nic.near.net
Archive:

Description of Working Group:

The User Connectivity Working Group will study the problem of how to solve network users' end-to-end connectivity problems.

Goals and Milestones:

Done	Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.
TBD	Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate.

Internet Drafts:

"FYI on an Internet Trouble Ticket Tracking System for addressing Internet User Connectivity Problems", 02/11/1991, M. Mathis, D. Long <draft-ietf-ucp-connectivity-01.txt>

Request For Comments:

RFC 1297 "NOC Internal Integrated Trouble Ticket System Functional Specification Wishlist ("NOC TT REQUIREMENTS")"

CURRENT MEETING REPORT

Minutes of the User Connectivity Working Group (UCP)

Report not submitted.

Attendees

Andy Adams	ala@merit.edu
Vikas Aggarwal	aggarwal@jvnc.net
J. Allard	jallard@microsoft.com
Ken Boggs	boggs@ralvm12.vnet.ibm.com
Henry Clark	henryc@oar.net
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Robert Gutierrez	gutierre@nsipo.nasa.gov
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Ittai Hershman	ittai@nis.ans.net
Dale Johnson	dsj@merit.edu
Daniel Karrenberg	daniel.karrenberg@cwi.nl
Thomas Kile	tgkile@brl.mil
John Klensin	klensin@infoods.mit.edu
Walter Lazear	lazear@gateway.mitre.org
John Leddy	jleddy@ibm.com
Linda Leibengood	ldl@ans.net
Hock-Koon Lim	lim@po.cwru.edu
Daniel Long	long@nic.near.net
Matt Mathis	mathis@psc.edu
Pushpendra Mohta	pushp@cerf.net
Donald Morris	morris@ucar.edu
Emmanuel Pasetes	ekp@enlil.premenos.sf.ca.us
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Tom Sandoski	tom@concert.net
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Kannan Varadhan	kannan@oar.net
Carol Ward	cward@westnet.net
Allen Wiersma	wiersma@nsipo.nasa.gov
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov

Ki-Sung Yoo	`ksyu@garam.kreonet.re.kr`
Jessica Yu	`jyy@merit.edu`
Paul Zawada	`zawanda@ncsa.uiuc.edu`

3.6 Routing Area

Director(s):

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun

Routing continued to be an active area at the San Diego IETF meeting. Six Routing working groups met and there were three routing related BOF's. Their reports are as follows:

Border Gateway Protocol Working Group (BGP)

The Working Group reviewed the BGP-OSPF and BGP-ISIS interactions documents. The other main activity was discussing extensions to BGP3 to accommodate the CIDR work. This will result in BGP Version 4.

IP Over Large Public Data Networks Working Group (IPLPDN)

The IPLPDN Group completed work on the IP over X.25 Specification. Andy Malis presented the proposal to the IETF plenary. The Working Group plans to submit this for proposed standard.

The Group discussed the text of a draft document for IP over Circuit ISDN. No final consensus was reached but work will continue on this topic.

A draft document was presented on Frame Relay Parameter Negotiation. There was agreement to make this similar to the PPP parameter negotiation. Work on this will continue.

A proposal for Shortcut Routing was made by Paul Tsuchiya. Several issues were uncovered and Paul agreed to revise his proposal for the next meeting.

IP Routing for Wireless/Mobile Hosts BOF (MOBLHOST)

The Group met and reviewed the proposed charter. A permanent Working Group will be formed. Presentations were made by John Ioannidis on the Columbia University Mobile Host Protocol and by Karl Auerbach on his thoughts on alternative approaches. Additions to the Dynamic Host Configuration Protocol (DHCP) to support mobile hosts was also discussed.

A detailed review of the Columbia Specification was completed. This is intended to be published as an Internet Draft.

Multicast Extensions to OSPF Working Group (MOSPF)

The Group discussed several suggestions for simplifications of the protocol based on the Proteon implementation experience. Several clarifications were made to the multicast specification including implementations suggestions and examples. The Working Group plans to submit MOSPF for Proposed Standard.

Open Shortest Path First IGP Working Group (OSPF)

The Working Group discussed changes to the OSPF V2 documents to handle specific virtual link configuration problems and MIB variable changes. They are close to completing an OSPF Trap MIB and a description of the Not So Stubby Area (NSSA) option. Also discussed were approaches to carry BGP attributes and how to run OSPF over Frame Relay.

RIP Version II Working Group (RIPV2)

The newly formed RIP Working Group finalized new field formats for RIP V2 using the MBZ fields in the RIP header. This version of the protocol will include extendible authentication support. RIP V2 should be backward compatible with RIP V1.

The Working Group plans to publish an Internet Draft of the RIP V2 specification in the near future. A first draft of a RIP V2 MIB was distributed.

CURRENT MEETING REPORT

Reported by Peter Ford/LANL

Minutes of the Internet Routing and Addressing BOF (ROAD)

The ROAD Group held a two session BOF with the discussions centered on the problems of IP address space depletion and routing in the Internet with a large number of autonomous systems. The sessions were well attended with 90-100 attendees. The first session was organized around talks by Ross Callon and Bob Hinden. Ross discussed the issues surrounding the use of CLNP in the Internet. Bob discussed a plan involving the encapsulation of IP within IP. The second session was an open mike session led by Peter Ford and Philip Almquist.

Ross Callon discussed the general issues of a large Internet, pointing out that the current IP address does not meet the needs of a really large global Internet. He enumerated requirements for such an address which included the requirement of being able to handle up to 10**9 networks in the Internet. Ross noted that there already was a standard which shares many of the familiar traits of IP but with a larger, multi-level hierarchical address: CLNP. Ross discussed likely migration paths from IP to CLNP and in turn introduced the notion of a Network Address Translator (NAT). This migration did not subscribe to a wholesale transition to OSI; the focus was on replacing IP with CLNP and using UDP/TCP along with familiar Internet protocols (SMTP,FTP,TELNET,etc.,).

Bob Hinden presented a scheme for routing and addressing in the Internet based on the ideas of encapsulation and interdomain routing based on routing by Autonomous Domains (ADs). The encapsulating packet would be an IP packet with addresses which denote ADs. During initial deployment the encapsulating and decapsulating would be done at the borders of ADs. The hosts would still use the current IP packet format with addresses drawn from the current 32 bit IP address space. The machinery for determining the encapsulation header addresses depends on interactions between the DNS and interdomain routing.

During the second session there was vigorous discussion and Q&A of the two proposals, along with other proposed changes to Internet routing and addressing. It was clear that both proposals needed more work and the authors ask for possible "partners in crime" for BOFS at the Boston IETF.

There was an open discussion of what should be done about addressing the growth of the Internet. This led to a far ranging discussion of the problem space:

- When/how will the IP address space be depleted.
- What are the technology trends in routers (memory size, speed, caching)?
- Should host changes be required, over what time period, etc.
- What about resource allocation, flows, ATM?
- Should the Internet header be "fixed" or should it be self describing and extensible.
- Route servers vs. global information distribution.

Much concern was expressed that the I**3 (IETF,IESG,IAB) will have a difficult time addressing these issues effectively and in a timely matter. It was decided these discussions should continue on the 'big-internet@munnari.oz.au' mailing list.

Copies of early writeups by Ross Callon on CNAT and Bob Hinden on IP-IP encaps can be found in the gated.cornell.edu:pub/road directory.

Attendees

Robert Aiken	raiken@nsf.gov
Guy Almes	almes@ans.net
Philip Almquist	almquist@jessica.stanford.edu
Michael Anello	mike@xlnt.com
Nagaraj Arunkumar	nak@3com.com
Fred Baker	fbaker@emerald.acc.com
David Bolen	db31@nis.ans.net
David Borman	dab@cray.com
Scott Brim	swb@cornell.edu
Gregory Bruell	gob@shiva.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Douglas Carson	carson@utcs.utoronto.ca
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
John Chang	jrc@uswest.com
Dean Cheng	dean@sun2.retix.com
Isaac Choi	ichoi@synoptics.com
Richard Colella	colella@osi.ncsl.nist.gov
Rob Coltun	rcoltun@ni.umd.edu
John Curran	jcurran@bbn.com
Osmund de Souza	osmund.desouza@att.com
Kathleen Dodd	kathy@gateway.mitre.org
Daniel Duchamp	djd@cs.columbia.edu
Susan Estrada	estradas@cerf.net
Dino Farinacci	dino@cisco.com
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Peter Ford	peter@lanl.gov
Karl Fox	karl@morningstar.com
Vince Fuller	vaf@stanford.edu
Richard Gallup	rpg@sdsc.edu
Eugene Geer	ewg@nvuxr.bellcore.com
David Getchell	dag@llnl.gov
Walter Guilarte	70026.1715@compuserve.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Susan Hares	skh@merit.edu

Robert Hinden	hinden@eng.sun.com
Kathleen Huber	khuber@bbn.com
Bob Jeckell	rrj@3com.com
Dan Jordt	danj@nwnet.net
George Kajos	kajos@coral.com
Prakash Kashyap	prakash.kashyap@sun.com
Dave Katz	dkatz@merit.edu
Manu Kaycee	kaycee@ctron.com
Tom Kessler	kessler@eng.sun.com
Zbigniew Kielczewski	zbig@eicom.qc.ca
John Leddy	jleddy@ibm.com
John Lekashman	lekash@nas.nasa.gov
Susan Lin	suelin@ibm.com
Tony Li	tli@cisco.com
Daniel Long	long@nic.near.net
Bryan Lyles	lyles@parc.xerox.com
John Lynn	lynn@jalps2.cit.cornell.edu
Kent Malave	kent@chang.austin.ibm.com
Gary Malkin	gmalkin@xylogics.com
Bill Manning	bmanning@rice.edu
Ed Menze	menze@cs.arizona.edu
David Miller	dtm@mitre.org
Dennis Morris	morrisd@imo-uvax.dca.mil
Donald Morris	morris@ucar.edu
John Moy	jmoy@proteon.com
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Michael O'Dell	mo@bellcore.com
Ayal Opher	aopher@synoptics.com
David Oran	oran@sneezy.lkg.dec.com
David Oran	oran@sneezy.lkg.dec.com
Lars Poulsen	lars@cmc.com
Stephanie Price Marasciullo	price@cmc.com
Rex Pugh	pugh@hprnd.rose.hp.com
Scott Quinn	smq@aries.bellcore.com
Paul Raison	raison@xylogics.com
Yakov Rekhter	yakov@watson.ibm.com
Jonathan Rodin	rodin@ftp.com
Allyn Romanow	allyn@eng.sun.com
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Miguel Sasson	sasson@xylogics.com
Tim Seaver	tas@concert.net
John Seligson	johns@ultra.com
Erik Sherk	sherk@sura.net
Tom Simons	nes!simons@phx.mcd.mot.com

William Simpson	Bill_Simpson@um.cc.umich.edu
Frank Solensky	solensky@clearpoint.com
Martha Steenstrup	msteenst@bbn.com
Claudio Topolcic	topolcic@nri.reston.va.us
Kannan Varadhan	kannan@oar.net
Huyen Vu	vi@polaris.dca.mil
Lee Wade	wade@nsipo.arc.nasa.gov
Justin Walker	justin@apple.com
Yuan Wang	natadm!ycw@uunet.uu.net
Carol Ward	cward@westnet.net
Jonathan Wenocur	jhw@shiva.com
Chris Wheeler	cwheeler@cac.washington.edu
Douglas Williams	dougw@ralumg.unet.ibm.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr
Jessica Yu	jyy@merit.edu

CURRENT MEETING REPORT

Reported by Kevin J. Rowett/Tandem

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

Agenda

- Status of Group to change from BOF to Working Group (not completed, pending Charter submission by Chair to Area Director).
- Reports from related groups.
- Review Columbia specification for mobile internetworking goal to publish as an Internet Draft.
- DHCP implications and requirements.
- Karl Auerbach wants to speak about some vague notions.

No report from the IEEE 802.11 Working Group. To get on the mail list for 802.11 documents, contact the 802.11 Chair, Vic Hayes (Vic.Hayes@Utrecht.ncr.com).

Charter Presented for Benefit of Recent New Members

- Develop/adopt architectures and protocols to support mobile hosts in the Internet.
- Convey Internet mobility concerns/ideas to other relevant working groups and standards bodies.

Scope of Work

- Issues above media access layer: addressing/naming/routing/bridging.
- Issues beyond Dynamic Host Configuration: roamers and temporarily relocated hosts.
- Mobile hosts, mobile networks, mobile internetworks.
- Cellular and general topologies, with and without wired infrastructure.
- Mobility across multiple link layers (wired and wireless).
- Multi-protocol, as well as IP-only.
- Impact on higher layers, e.g., transport.
- Accommodation of sleeping (battery-saving) and off-line hosts.

Outside of Scope

- Solutions that do not interoperate with existing Internet systems.
- Issues of delay or jitter-sensitive traffic, or other quality-of-service concerns (should be orthogonal to mobility).

- Congestion avoidance/control (ditto).
- Header or packet compression (ditto).
- Privacy (ditto); however, authentication is within scope.

John Ioannidis presented his mobile networking architecture (see accompanying viewgraphs). John, Dan Duchamp, and Gerald Maguire have produced a draft specification for the architecture and protocols, for possible publication as an RFC. (A subset of the attendees of this BOF met later in the week for a thorough review of the document. After an editing pass, the document will be made available as an Internet Draft, which the Group may then choose to recommend for submission as an RFC.)

Karl Auerbach spoke for about fifteen minutes about how he saw the realm of mobile computing being used and some of his experiences at Sun.

- Traffic is mobile to a fixed host.
- Optimized base station to base station to allow easier "mobiling" of a users environment.
- It's really people that are mobile, not the machines. People want to take their environment with them.

Steve Deering discussed possible DHCP implications of the Columbia mobile routing scheme:

- Concept of portable IP addresses.
- Which subnet to assign mobile host to.
- Vendor extensions to be used for IP address attributes.

These concerns will be presented to the DHC Working Group for their consideration.

Attendees

Karl Auerbach	karl@empirical.com
William Barns	barns@gateway.mitre.org
John Chang	jrc@uswest.com
Dean Cheng	dean@sun2.retix.com
Andrew Cherenson	arc@sgi.com
Steve Deering	deering@xerox.com
Ralph Droms	droms@bucknell.edu
Guy Dubuc	guyd@newbridge.com
Daniel Duchamp	djd@cs.columbia.edu
Julio Escobar	jescobar@bbn.com
Karen Frisa	karen.frisa@andrew.cmu.edu

Cliff Frost	cliff@cmsa.berkeley.edu
Robert Gilligan	gilligan@eng.sun.com
Terry Gray	gray@cac.washington.edu
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Robert Gutierrez	gutierre@nsipo.nasa.gov
Robert Hinden	hinden@eng.sun.com
Nat Howard	nrh@bellcore.com
Thomas Hutton	hutton@opus.sdsc.edu
John Ioannidis	ji@cs.columbia.edu
Ronald Jacoby	rj@sgi.com
Phil Karn	karn@qualcomm.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Rajeev Kochhar	rxk@3com.com
Fidelia Kuang	kuang@apple.com
David Kumpf	kumpf@hprnd.rose.hp.com
John Larson	jlarson@parc.xerox.com
Dory Leifer	del@terminator.um.cc.umich
Joshua Littlefield	josh@cayman.com
Leo McLaughlin	ljm@wco.ftp.com
Robert Morgan	morgan@jessica.stanford.edu
Michael O'Dell	mo@bellcore.com
David Oran	oran@sneezy.lkg.dec.com
Radia Perlman	perlman@radia.enet.dec.com
Stephanie Price Marasciullo	price@cmc.com
Rex Pugh	pugh@hprnd.rose.hp.com
James Rees	jim.rees@umich.edu
Kevin Rowett	kevinr@tandem.com
Anthony Rutkowski	amr@sprint.com
Andreas Schachtner	afs@germany.eu.net
Tom Simons	nes!simons@phx.mcd.mot.com
Brad Steinka	brad@python.eng.microcom.com
Dave Thompson	daveth@ingate.microsoft.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Chin Yuan	cxyuan@pacbell.com

MI: Motivation

- Powerful, affordable portable computers *notebooks, laptops, etc.*.
- Wireless networking.
- Ubiquitous computing/networking *people want the same enviroment wherever they are*.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 2 of 37

Mobile Internetworking:

Provide Mobile Hosts with continuous network access throughout the Internet, in the presence of:

- Change of interface.
- Change of physical location.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 3 of 37

MI: Goals

- Migrating hosts keep IP address.
- Transport &higher protocols untouched.
- Stay compatible with current Internet.
- Non-participating entities *e.g., routers* not affected.
- Low overhead/scalability.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 4 of 37

Addressing:

Duality of IP addresses:

- Connection identifiers.
- Route identifiers.

⇒ Mobility intimately involves handling this duality.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 5 of 37

Handling Mobility

- Address not a unique ID:
  - Changes as host migrates
  - Directory service or forwarding pointers map new address
  - Higher-level protocols have to adapt.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 6 of 37

Handling Mobility cont'ed

- Address has no structure (flat a/s)
  - Per-host routing.
  - Distribution of routes problematical vis-a-vis scaling.
  - e.g., multicast routing.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 7 of 37

Handling Mobility cont'ed

- Keep structured address
  - Handle mobility within lowest hierarchical level(s).
  - Scaling benefits of hierarchical addressing maintained.
  - Works best when mobility exhibits spacial locality.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 8 of 37

Routing

Routing in a mobile environment:

- Tracking the mobile's *location*as it moves.
  - Handoffs.
- Updating the routing information.
  - Distributed routes(consistency).
  - Centralized routes(ask KGB). depending on the approach, these tasks can be made easier or harder

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 9 of 37

Route discovery

- Fix-The-Routes (FTR):
  - No special treatment of mobiles.
  - Update per-host routes in gateways on-demand or otherwise.
  - All routers are involved.
- Heal-The-Partition (HTP):
  - Mobiles on a virtual network
  - Special routers 'healing' the partition
  - Only the special routers are involved.

CU/CS

'Hiding' mobility:

- Transport layer and up:
  - Should be left unaffected.
- Network layer:
  - Provides uniform access to interfaces.
  - Handles routing and internetworking.
- Data link layer:
  - Exchange of routing info awkward.
  - Heterogeneity in interfaces.

CU/CS

The Setup:

- Mobile Hosts (MHs)
  - Addresses on reserved subnet(s).
  - Maintain one address (home address).
- Mobile Support Stations (MSSs)
  - Each defines a *cell.*
  - Double as wireless-to-wired gateways.
  - MSSs 'conspire' to 'heal' the partitioned MH network.
  - minimum 1MSS per campus.

CU/CS

Two new protocols:

- IPIP: IP encapsulation protocol:
  - IPPROTO_IPIP = 94.
  - Used to tunnel traffic between MSSs.
    thus healing the partition
- MICP: Mobile Internetworking Control Protocol:
  - IPPROTO_MICP = 95
  - analogous to ICMP for MSS messages
    convenient anagram of ICMP.

CU/CS

Intra- vs. Inter- campus:

- Structure of IP address:
  (<network> <host-within-network>)

- Naturally leads to two cases:
  - *Intra*-campus mobility.
  - *Inter*-campus mobility.

CU/CS

Intracampus operation:

- MSS beacon identifies cell.

- MH-MSS 3-way handshake.

- MSSs exchange MH location info.

- Lazy updates of MH location.

CU/CS

Sample campus

describe (MSSs form cells, all cells in one subnet, mention partitioned subnet and healing. mention static hosts and routes have the nearest MSS as their route to mobile subnet)

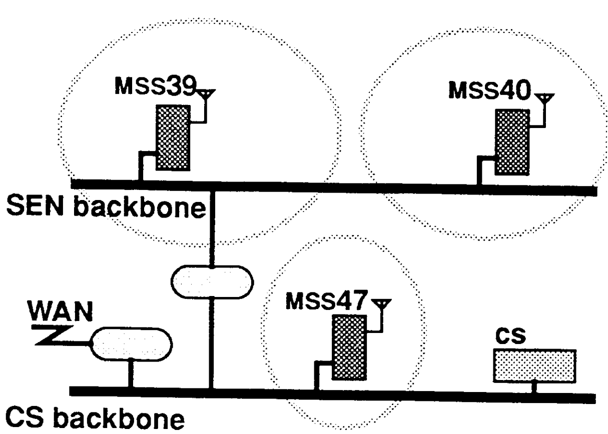

CU/CS

MH moves in MSS39's cell:
- MSS is broadcasting beacon
- 3-way handshake
- MSS39 is the MH's default gw

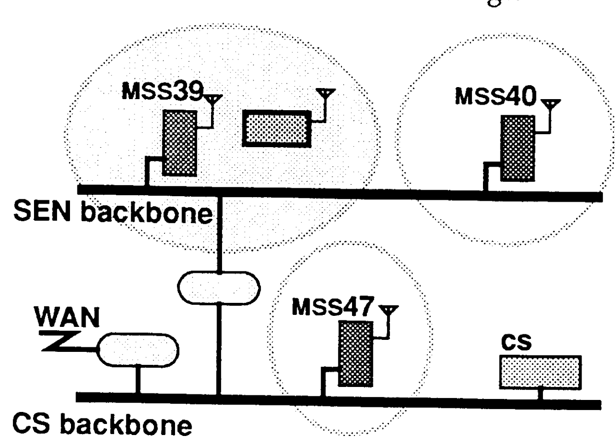

CU/CS

MH-MH communication within a cell

- ARP for destination address. *if no direct MH-MH comm (i.e. through hub), case is same as before*

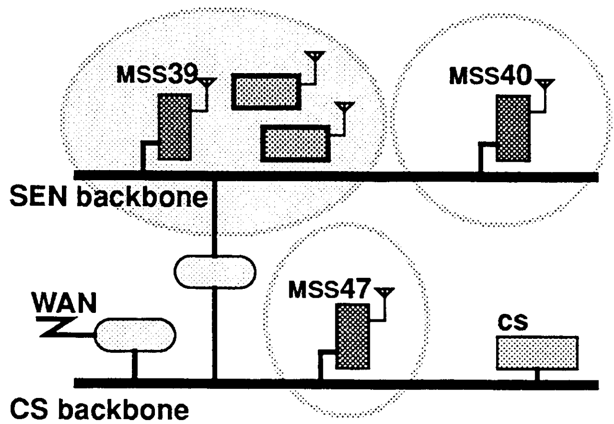

CU/CS

jj. 18MAR92 — 23rd IETF, Mobile Internetworking BOF — Slide 18 of 37

MH to regular host communication:

- outgoing pkts routed thru MSS39.
- incoming pkts picked up by MSS47 tunneled to MSS39 and delivered.

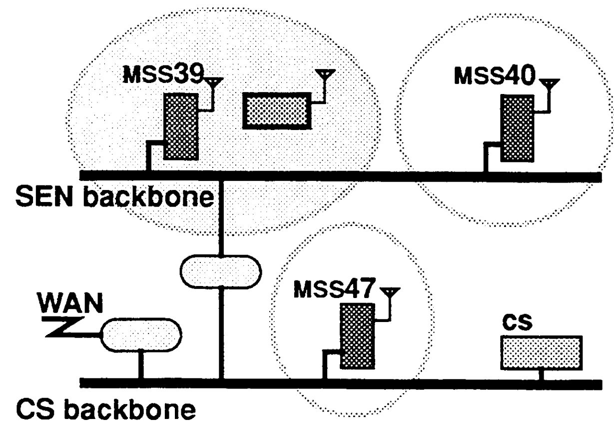

CU/CS

jj. 18MAR92 — 23rd IETF, Mobile Internetworking BOF — Slide 19 of 37

Remote MH-MH communication:

- MSS39 proxy-ARPs for MHb
- MHa now sends packets to MSS39
- MSS39 tunnels them to MSS47

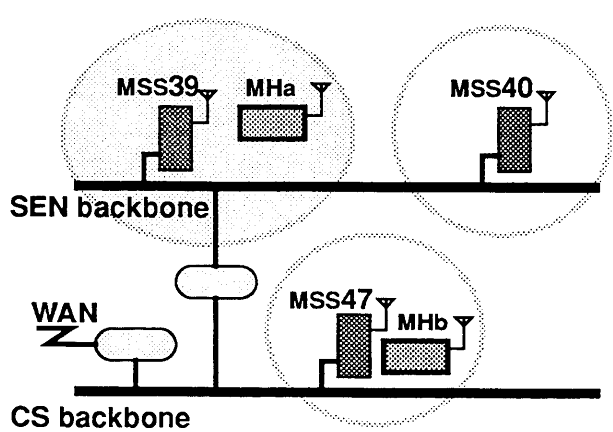

CU/CS

jj. 18MAR92 — 23rd IETF, Mobile Internetworking BOF — Slide 20 of 37

MH migration.

- lose old or hear new beacon.
- 3-way handshake.
- MSS40 notifies MSS39.

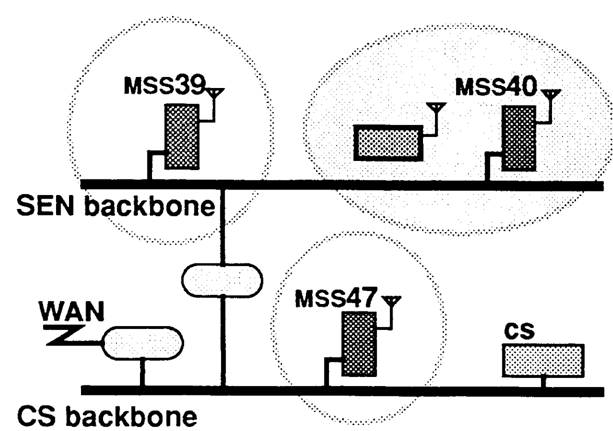

CU/CS

jj. 18MAR92 — 23rd IETF, Mobile Internetworking BOF — Slide 21 of 37

- cs sends packet to MH thru MSS47
- MSS47 tunnels it to MSS39
- MSS39 tunnels it to MSS40 but also sends redirect to MSS47

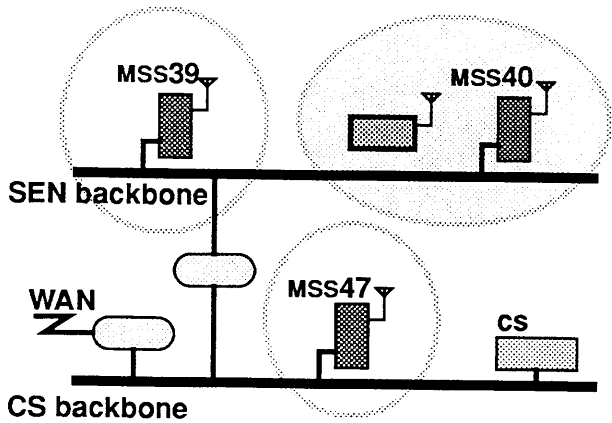

CU/CS

Intercampus operation:

- MH ('popup') acquires a nonce address.
  - e.g., using DHCP
  - real i/f gets this address...
  - and registers with home MSS.
- MH home address assigned to virtual i/f.
- MH acts as its own MSS.

CU/CS

Outgoing traffic from popup

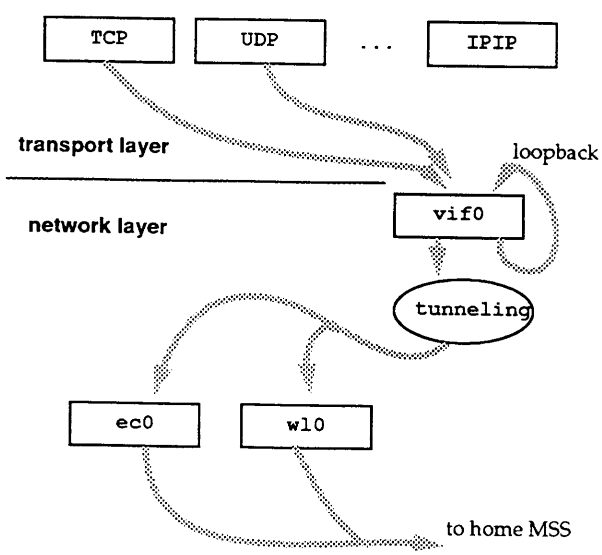

CU/CS

Incoming traffic to popup

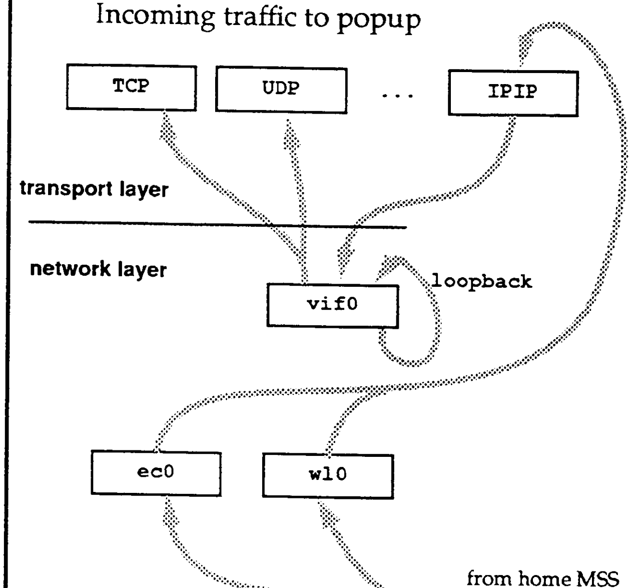

CU/CS

Intercampus, cont'd

- Outgoing packets have a source address in their home network.
- Incoming packets are first routed to home network, then tunneled.
- Not possible to distinguish 'roamer' from a regular mobile (think of security).
- Optimizations (e.g., talk to fellow popups) possible.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 26 of 37

Security considerations

- MSSs only accept traffic (IPIP or MICP) from other known MSSs.
- Pop-ups (local or remote) authenticate themselves *(really, it's just like a remote new MSS authenticating itself)*.
- Anything else handled at a higher level.
  - Host security remains a host issue.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 27 of 37

Scale & Performance

- Redirects keep hop count close to 1.
- Timeouts keep per-MH info low.
- Lazy MH location keeps control traffic low
- Negligible network overhead to tunnel.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 28 of 37

ARPing:

if (*target is not an MH*)
    use the regular ARP code.

else if (*source is a remote or unknown MH* or *both source and target are local*)
    drop ARP request.

else if (*target is unknown*)
    attempt to locate MSS handling MH.

else
    proxy-ARP for target.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 29 of 37

Routing:

if (*target is not an mh*, or *is a local MH*)
deliver/route using regular IP code

else if (*target is remote and known*)
encapsulate in IPIP and deliver to remote MSS.

else
drop packet
attempt to locate MSS handling MH.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 30 of 37

Software Complexity:

MSS:

- ~1800 lines of kernel code.
  - 12 bytes/MH static data.
- ~1600 lines of user-level code.

MH (local and popup):

- ~2400 lines of user-level code.

Code running on i386 machines under Mach 2.6 since 4/91.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 31 of 37

Software Performance: MH

Local MH:

- Control process (`mhmicp`): ~0.2% of CPU
- Cell switch ~2s (limiting factor: beacon)
- 1s beacons, 30s pings, 30 mobiles: need 1kbps of bandwidth.

Popup:

- Control process (`pumicp`): ~0.1% of CPU.
- Mobile pays the overhead to tunnel.

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 32 of 37

Software Performance: MSS

- Control process (mssmicp): <2% of CPU
- MH location: ~10ms
- Perceived lag from MH: <0.5s
- MICP redirects: 2 packets + 1 per hop per MSS that knows about it.
- Tunneling overhead:
  - add/remove a header
  - no extra hops
  - 24 extra bytes per packet

CU/CS

ji. 18MAR92 23rd IETF, Mobile Internetworking BOF Slide 33 of 37

Additional work

- Improve routing interaction between MSSs and routers. *(only msss involved)*
- Incorporate more robust authentication mechanisms.
- Automate popup configuration with DHCP or BOOTP.

CU/CS

Future Work

- Addressing with more hierarchical levels.
- Handle multiple service providers.
- Define service interfaces that know about mobility:
  - replicated file systems.
  - other 'directory' services.
  - allow prolonged separation from the infrastructure.
  - structure client/server splits that minimize traffic and exploit coherence.
- Security and privacy issues.

CU/CS

Central Points

- Locality that follows address hierarchy boundaries.
  - IP address structure provides a natural cleavage line.
  - *Heal-The-Partition* is the right approach.
  - Separating inter- from intra-campus case makes the task manageable.
- Switching at network layer (rather than in the fabric) results in cleaner design.
- Minimalist design is easy to implement and performs well.

CU/CS

Value of the work

- Produce a much-needed, working system.
- Shows the interaction of hierarchical addressing/routing and locality of mobility.
- Defines performance metrics for evaluating such mobile systems.
- Paves the way for an Internet-based personal computing and communications network.

CU/CS

3.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):
Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:
General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu
Archive:

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done — Complete development of Version 2 of the Border Gateway Protocol (BGP).

Ongoing — Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.

Done — Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.

Done — Develop a MIB for BGP.

Done — Work with the Security Area to enhance the provision for security in BGP.

Done — Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

Internet Drafts:

"Border Gateway Protocol NEXT-HOP-SNPA Attribute", 04/15/1991, Paul Tsuchiya <draft-ietf-bgp-nexthop-01.txt>

"Default Route Advertisement In The Border Gateway Protocol", 08/09/1991, Dimitry Haskin <draft-ietf-bgp-defaultroute-00.txt>

"Multicast Communications Using BGP", 08/26/1991, Scott Brim <draft-ietf-bgp-multicast-01.txt>

"BGP OSPF Interaction", 10/25/1991, Kannan Varadhan <draft-ietf-bgp-ospfinteract-02.txt>

"A Unified Approach to Inter-Domain Routing", 12/06/1991, D. Estrin, Y. Rekhter, S. Hotz <draft-ietf-bgp-unirouting-00.txt>

Request For Comments:

RFC 1105 "Border Gateway Protocol BGP"

RFC 1163 "A Border Gateway Protocol (BGP)"

RFC 1164 "Application of the Border Gateway Protocol in the Internet"

RFC 1265 "BGP Protocol Analysis"

RFC 1266 "Experience with the BGP Protocol"

RFC 1267 "A Border Gateway Protocol 3 (BGP-3)"

RFC 1268 "Application of the Border Gateway Protocol in the Internet"

RFC 1269 "Definitions of Managed Objects for the Border Gateway Protocol (Version 3)"

CURRENT MEETING REPORT

Reported by Peter Ford/LANL

Minutes of the joint session of the Border Gateway Protocol Working Group (BGP) and the CIDR Supernetting BOF (CIDR)

The BGP Working Group met jointly with people interested in Classless InterDomain Routing (CIDR) to discuss the development of an addressing plan which can be used for IP. CIDR would allow collapsing adjacent network addresses into a single prefix, and that prefix would be passed within the routing system as the route to all the "collapsed" networks. CIDR is proposed to mitigate the scaling problems in the Internet's routing system which are due to "flat routing" and the fact that the Internet will shortly (1 to 3 years) run out of Class B addresses. When the Internet runs out of Class B addresses, the current available option is to allocate Class C network addresses which will require networks which have more than 255 end systems to advertise multiple network addresses to the global Internet routing system. The purpose of this session was to discuss various schemes for assigning and collapsing addresses, including collapsing along a multi-level hierarchy, what the hierarchies would look like (size and placement), what the mapping between network providers and collapsed prefixes would look like. There was a significant turnout of interested people and the discussion was quite spirited.

Yakov Rekhter led-off the discussion with a brief overview of CIDR and an explanation of the goals of the session. He then presented a proposal for Address Assignment Authorities (AAAs). (see attached slides).

Yakov went on to propose a possible allocation of AAAs, which was to assume a top down allocation of 1000 AAAs which would require coding top level AAA coding of 10 bits. Within the Class C address space this would imply that each AAA would have a maximum of 1000 Class C network addresses.

This was a good starting point for discussing network topology issues, and "who would be candidates for being AAAs?".

Several people disagreed with Yakov's proposal for picking a fixed size breakout for top level AAAs. Several people proposed an allocation of top level AAAs which was scaled by the size of the community one was trying to serve, perhaps using the population size or the size of the telephone networks as scaling factors. It was noted that using a Kampei style address assignment scheme might be a good thing to do here.

There was concern expressed for deploying CIDR too soon, before a sufficient technology base was deployed for aggregating multiple Class C network addresses. Several people noted that this may severely impact intra-domain routing protocols since an aggregated prefix would have to be exploded to its constituent Class C networks if the routing protocols did not handle aggregation correctly (RIP and EGP).

There was significant discussion of how to carve up Class A and Class B network addresses effectively. There was general concurrence that for the time being Class A's should not be allocated. This would be until there is a technology base which can be used with carved up A's. It was noted that this would be feasible once most routers "knew how to do variable length subnets".

There was discussion on how Class C# (Solensky and Kastenholz) could coexist with CIDR.

Jon Postel gave a short description of what the IANA does and how it decides who gets what network addresses.

The discussion clearly overran the time allotment and future discussion of this issue was proposed to continue on Email using the big-internet@munnari.oz.au mailing list. Yakov Rekhter agreed to discuss with the IESG about forming a Working Group to work on an IP addressing plan.

Attendees

Andy Adams	ala@merit.edu
Nagaraj Arunkumar	nak@3com.com
John Ballard	jballard@microsoft.com
Tom Benkart	teb@saturn.acc.com
David Bolen	db31@nis.ans.net
Scott Brim	swb@cornell.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Ross Callon	callon@bigfut.enet.dec.com
Isaac Choi	ichoi@synoptics.com
Rob Coltun	rcoltun@ni.umd.edu
Robert Enger	enger@ans.net
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Richard Gallup	rpg@sdsc.edu
Eugene Geer	ewg@nvuxr.bellcore.com
Leonard Gomelsky	gomelsky@ee.ucla.edu
Paul Griffiths	griff@chang.austin.ibm.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Susan Hares	skh@merit.edu
Eugene Hastings	hastings@psc.edu
Ittai Hershman	ittai@nis.ans.net
Randy Hoebelheinrich	rho@rowan.lanl.gov
Jeffrey Honig	jch@risci.cit.cornell.edu
Kathleen Huber	khuber@bbn.com
Dan Jordt	danj@nwnet.net

Dave Katz	dkatz@merit.edu
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELE MAIL/C=US/@sprint.com
Thomas Kile	tgkile@brl.mil
Sylvain Langlois	sylvain@cli53an.edf.fr
John Leddy	jleddy@ibm.com
Tony Li	tli@cisco.com
Susan Lin	suelin@ibm.com
Kenneth Lindahl	lindahl@violet.berkeley.edu
Daniel Long	long@nic.near.net
John Lynn	lynn@jalps2.cit.cornell.edu
Kent Malave	kent@chang.austin.ibm.com
Michael Marcinkevicz	mdm@csu.net
April Marine	april@nisc.sri.com
Matt Mathis	mathis@psc.edu
Ed Menze	menze@cs.arizona.edu
Donald Merritt	don@brl.mil
Dennis Morris	morrisd@imo-uvax.dca.mil
Jon Postel	postel@isi.edu
Lars Poulsen	lars@cmc.com
Paul Raison	raison@xylogics.com
Yakov Rekhter	yakov@watson.ibm.com
Jonathan Rodin	rodin@ftp.com
Benny Rodrig	4373580@mcimail.com
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Sharad Sanghi	sharad@ans.net
Miguel Sasson	sasson@xylogics.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Erik Sherk	sherk@sura.net
Frank Solensky	solensky@clearpoint.com
Martha Steenstrup	msteenst@bbn.com
William Streilein	bstreile@wellfleet.com
Glenn Trewitt	trewitt@pa.dec.com
Kannan Varadhan	kannan@oar.net
Reed Wade	wade@cs.utk.edu
Chris Wheeler	cwheeler@cac.washington.edu
Bert Wijnen	wijnen@vnet.ibm.com
Cathy Wittbrodt	cjw@nersc.gov
Jessica Yu	jyy@merit.edu

Goal: "Recommended Guidelines for IP Address Assignment."
To Achieve:

Consistency
Efficiency
Ease of Management and Coordination

"Address Assignment Authority" (AAA)

Distributed way of managing address space
Promote routing information efficiency
Recursive => delegation of AAA

Need to Extend CIDR

Pure Class "C" supernetting provides _limited_ extension with respect to the IP Address Space.

Large portions of A & B are still unused!

"AAA concept needs to be applied to the _whole_ IP address space.

How to Carve Address Space?

Top-Down to ensure feasible routing (wrt scaling)

How many levels

Branching Factor at each level
(deep trees v. bushy trees)

Need to determine number of top level ASSs.

CIDR & "NSAP Address Guideline"

Attempt to solve the same problem

May benefit from coordination between NSAP address assignment and IP address assignment

Single AAA --> IP & NSAP

IP and NSAP topology is likely to be congruent
Address administration boundaries are likely to be congruent:

- a service provider provides both IP and CLNP services
- the same geographical area provides both IP and CLNP services

CURRENT MEETING REPORT

Reported by David Bolen/ANS

Minutes of the joint session of the Border Gateway Protocol (BGP) and the ISIS for IP Internets (ISIS) Working Groups

Discussion of BGP-4 Extensions

The largest portion of the BGP meeting was concerned with the discussion of the BGP4 Document authored by Vince Fuller and Tony Li. The additions were made primarily to add the necessary support into BGP for Classless Inter-Domain Routing (CIDR).

IP Prefixes vs. Masks

The need for carrying some form of network masks as part of BGP was discussed in the light of the need for CIDR. The necessary information could be carried either as a combination of network and netmask, or it could be encoded as a prefix with a length value.

A key difference between the two proposals is that carrying an entire mask allows the use of non-contiguous masks, while a prefix requires that the mask be contiguous.

The general consensus of the Group was that non-contiguous masks presented several problems, especially in routing table lookups (where multiple entries can match), and in the automatic aggregation of such masks (which we aren't sure how to do yet, and it's critical for CIDR). Therefore prefixes, being more deterministic, were a better choice for BGP-4.

It was agreed that masks could prove useful later when some of the trickier issues have been dealt with. In that case, they could always be added to a later version of the protocol.

>> Use prefixes rather than masks.

Aggregation Rules from BGP-4 Document

The Group discussed the various rules presented in the BGP-4 Document to handle aggregation:

- Rule 1 - always do longest match.
- Rule 2 - Inject "poison" routes to avoid loops. In a multi-homed case, if the aggregator is the primary provider, the aggregator must also announce the longer prefix for the client (to override the same announcement via that client's other provider). If the aggregator is not primary, this additional announcement is not necessary.
- Rule 3 - Punch hole to sever old route when switching providers. This requires an announcement withdrawal mechanism in BGP.

In particular, Rule 3 was discussed in that it required the addition of a withdrawal mechanism in BGP to withdraw a previous announcement (along the lines of the facility provided within IDRP).

The largest concern, if this wasn't provided, was that packets could flow partially down a bad path before they were either bounced or black-holed. Also, traceroute would no longer function properly in such a case.

The general consensus was that these problems were not critical enough to warrant the added complexity of the withdrawal mechanism, especially when interoperability with older implementations of BGP that didn't have such a mechanism was taken into account.

>> Rules 1 & 2 ok - removing Rule 3.

Support for AS Sets

In order to be able to handle multi-level aggregation, the ability to specify an AS_PATH that included AS sets rather than simply a sequence is very important. If AS-1 and AS-2 both flow packets through AS-3, AS-3 would like to be able to aggregate routes if AS-1 and AS-2 fall under the same prefix. Since they represent different AS_PATHs, that is currently not possible.

IDRP was brought up as an example of using sets. In IDRP, the RD_PATH (like BGP's AS_PATH) is an overall sequence of elements, where each element is either an RD sequence, or an RD set. An RD set implies that the packets flow through one or more of the RDs in the set, but not necessarily all of them, and no order is specified.

IDRP also provides for the concept of routing confederations, which is a method for aggregating several routing domains into a single routing domain confederation (RDC) which is generally treated just like an RD when it appears in the RD_PATH.

Using confederations was considered more powerful than just sets, but also more complicated, and not required for the CIDR support we want to include in BGP-4.

A possibility was just to turn BGP's AS_PATH into a set, which would imply no particular order of AS traversal. However, this would prevent any route filtering based on AS order. Matt Mathis suggested having an AS_PATH that started with a sequence, and was followed by a set.

This discussion also brought up the fact that the AS_PATHs would probably be growing as this structure would encourage the size of an AS to be small, which led to thinking about the assignment of AS numbers hierarchically in order to allow them to be aggregated as well. Finally, the discussion turned to whether AS values should be increased to 32-bit rather than 16-bit. Some people strongly felt it should be increased in size, especially now as we were making these other changes.

>> General consensus for requiring sets/sequences, but not confederations. Stick with 16-

bit AS, although moving to 32-bit at the same time as the AS set change was strongly discussed.

Parallel BGP Sessions

During the discussion, the issue of being able to run parallel BGP sessions turned out to primarily address the problem of wanting to filter on an AS wherever it shows up in the AS_PATH, rather than just the first AS (as is common in current applications). This is an implementation problem rather than one with the protocol.

The point was made that it can be dangerous to run two BGP connections in the same host with the only exchange of information between them being a routing table handled through a common IGP.

Therefore, it was decided that it wasn't worth changing the protocol to handle what was essentially an implementation issue.

>> Dropped

NEXT_HOP Handling

The NEXT_HOP discussion centered around the issue of allowing one BGP peer to pass along the address of some other host on a common wire as the NEXT_HOP value rather than using its own address. The knowledge that the other host was available would commonly have been determined via an IGP (such as OSPF), or by the fact that the BGP peer was maintaining sessions with both hosts.

Yakov brought up IDRP as an example of this functionality but under control of the source, which can choose to add additional options to a NEXT_HOP announcement controlling whether the recipient can propagate the NEXT_HOP value. This addresses the multiple BGP session, but not the IGP case.

The concern with allowing this optimization was that the new host was being included in the NEXT_HOP announcement without it being aware of its involvement. Under certain circumstances this can be very useful, but it can also easily violate policy or routing rules. It deserved some further investigation, but for now (and at least for BGP-4), it was decided to keep the current definition of NEXT_HOP.

>> Keep definition/use same as is currently specified.

BGP/OSPF Interaction Document

A short discussion of the current BGP/OSPF Interaction Document by Kannan Varadhan took place. The Document is nearing completion, and only had a few changes since the previous IETF.

Tag Bits

If the upper ("trusted") bit of the tag is set, the tag was system generated or configured, and the following 3 bits are used to encode the completeness of the route and how it should be handled:

		pl 00	01	10	11
c	0	<EGP> <l>	<EGP><l,nh>	never export	reserved
	1	<IGP><l>	<IGP><l,nh>	out of band	reserved

Matt Mathis suggested that the term "trusted" may not be the most appropriate (it could be taken to imply that the network administrator isn't trusted). Tony Li suggested the substitution of the term "automatic" instead.

>> Change "trusted" keyword to something like "automatic".

NEXT_HOP Handling with Subnets

This issue dealt with an optimization for handling the case where you learn a set of routes through OSPF that represented an entire subnet for a network, and what to assign NEXT_HOP to in that case. The optimization in question was whether or not you still had to always place yourself in NEXT_HOP rather than the node through which the subnets were routed.

It was agreed that this represented an implementation optimization and should not be dictated by this Document, but be left to the choice of the developer.

>> Leave optimization to developers.

A new revision of the Document will be published with a fairly short deadline for comment, so that the Document can then proceed through the standards process.

IBGP and OSPF Discussion

Jeff Honig held a discussion about possible problems with the use of IBGP, and how the same information might be propagated through an IGP (specifically OSPF) rather than with IBGP connections. The main concerns with IBGP were the need for scaling of N^2 connections, and the bandwidth utilized for IBGP traffic.

Enhancements to OSPF that were being discussed to be used to carry the external information included:

- OSPF Variable Tags (to allow the encoding of a complete AS_PATH)

- New OSPF Path Attribute LSA to propagate path/attribute pairs (to send unique path/attribute information around in separate packets that were referenced by the route announcements)

During the discussion, the general feeling was that the difference between IBGP resource usage (N-1 TCP connection blocks on each border router (BR), and bandwidth) and the resource required to distribute the same information via an IGP was not significantly different. Also, trying to store external information in internal routers could cause physical resource problems (such as memory) for those routers.

One agreed issue with IBGP was router discovery. Using the IGP to aid the BRs in discovering each other was considered an important feature. At a minimum, all that would be required is a single flag bit to be carried by the IGP to indicate that a host was a BR. Ideally several bits to encode additional information would be useful.

>> General discussion felt that IBGP resource utilization was not significantly more than that introduced by having the IGP carry EGP information.

>> Agreement was reached that border router discovery was important and that some bits (at least 1) should be requested from IGPs such as OSPF to support this.

BGP <-> IS-IS Interaction Document

Sharad Sanghi held a discussion on the BGP/IS-IS Interaction Document that he and Atul Bansal had authored.

The general issues were similar to the BGP/OSPF Interaction Document. The basic question discussed was in regards to the advantages and disadvantages of doing route injection vs. piggybacking vs. just using IBGP.

As with the previous IBGP vs. IGP (OSPF) discussion, the Group felt that it was not clear that the savings in resources by eliminating IBGP and using IS-IS to carry external routing information was worth the work to transfer the routes.

Router discover was still very useful, so adding bits to IS-IS (which it already has room to support) is definitely desired. Using IBGP with these bits for discovery should be fine for stub domains. For transit systems, it could prove useful to consider the injection/piggybacking cases, which should be studied further.

Attendees

Nagaraj Arunkumar	nak@3com.com
Syed Ashraf	irfan@pluto.dss.com
William Barns	barns@gateway.mitre.org
David Bolen	db31@nis.ans.net
Scott Brim	swb@cornell.edu
Ross Callon	callon@bigfut.enet.dec.com

Rob Coltun	rcoltun@ni.umd.edu
Osmund de Souza	osmund.desouza@att.com
Julio Escobar	jescobar@bbn.com
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Eugene Geer	ewg@nvuxr.bellcore.com
Terry Gray	gray@cac.washington.edu
Susan Hares	skh@merit.edu
Robert Hinden	hinden@eng.sun.com
Jeffrey Honig	jch@risci.cit.cornell.edu
Kathleen Huber	khuber@bbn.com
Bob Jeckell	rrj@3com.com
Dan Jordt	danj@nwnet.net
Manu Kaycee	kaycee@ctron.com
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELE MAIL/C=US/@sprint.com
Yoav Kluger	ykluger@fibhaifa.com
Bill Manning	bmanning@rice.edu
Matt Mathis	mathis@psc.edu
Dennis Morris	morrisd@imo-uvax.dca.mil
David Oran	oran@sneezy.lkg.dec.com
Rex Pugh	pugh@hprnd.rose.hp.com
Paul Raison	raison@xylogics.com
Yakov Rekhter	yakov@watson.ibm.com
Benny Rodrig	4373580@mcimail.com
L. Michael Sabo	lms@denver.ssds.com
Sharad Sanghi	sharad@ans.net
Miguel Sasson	sasson@xylogics.com
Tom Simons	nes!simons@phx.mcd.mot.com
Frank Solensky	solensky@clearpoint.com
William Streilein	bstreile@wellfleet.com
Sally Tarquinio	sally@gateway.mitre.org
Paul Traina	pst@cisco.com
Glenn Trewitt	trewitt@pa.dec.com
Chris Wheeler	cwheeler@cac.washington.edu
Jessica Yu	jyy@merit.edu

3.6.2 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: iplpdn@nri.reston.va.us
To Subscribe: iplpdn-request@nri.reston.va.us
Archive: /ietf.mail.archives/iplpdn.mail.archive

Description of Working Group:

The IP over Large Public Data Networks Working Group (IPLPDN) will specify the operation of the TCP/IP protocol suite over public data networks (PDNs) such as SMDS, ISDN, X.25 PDNs, and Frame Relay. The Working Group will develop and define algorithms for the resolution of IP addresses and for the routing of IP datagrams over large, potentially global, public data networks.

The IP over SMDS Working Group has defined the operation of the Internet protocols when SMDS is used to support relatively small virtual private networks, or Logical IP Subnets (LISs). Issues arising from public and global connectivity were delegated to the IPLPDN Working Group.

The IPLPDN Working Group will also continue the work of the Private Data Network Routing Working Group (pdnrout) on X.25 PDNs. This work will be extended to include call management and the use of the ISDN B channels for the transport of IP datagrams.

Address resolution and routing over Frame Relay will also be discussed.

Goals and Milestones:

Done	Establish priorities and dates of completion for documents.
TBD	Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.
TBD	Routing of IP datagrams across very large internets implemented SMDS and on other PDNs.
TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.

Internet Drafts:

"Discovery and Routing over the SMDS Service", 06/17/1991, Paul Tsuchiya <draft-tsuchiya-routingsmds-01.txt>

"Management Information Base for Frame Relay DTEs", 06/17/1991, Caralyn Brown, Fred Baker, Charles Carvalho <draft-ietf-iplpdn-frmib-05.txt>

Request For Comments:

RFC 1293 "Inverse Address Resolution Protocol"

RFC 1294 "Multiprotocol Interconnect over Frame Relay"

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

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

The IPLPDN Working Group made the following progress:

- Andy Malis lead a careful review of the "IP over X.25" draft. All remaining issues were resolved, and Andy volunteered to incorporate the modifications in the weeks after the IETF meeting. The revised draft will be made available, and, if there are no objections, the Chair will submit the final draft to the IESG and IAB for release as an RFC on the standards track.
- The status of the "IP over Circuit ISDN" draft was discussed. No text had been generated during the interim, but Brian Lloyd volunteered to write a first draft.
- Work on parameter negotiation for Frame Relay Permanent Virtual Circuits (PVCs) was reviewed. To avoid unnecessary duplication of the work done for Point-to-Point Protocol (PPP), the Group agreed to coordinate this work with that of the PPP Extensions Working Group.
- Paul Tsuchiya presented an expansion of the "shortcut" approach to address resolution. Issues were uncovered, particularly an issue concerning security, which Paul volunteered to address in a future revision of the draft.

Before adjourning, the Group agreed that more time would be required at the next IETF meeting.

Attendees

J. Allard	jallard@microsoft.com
Syed Ashraf	irfan@pluto.dss.com
Fred Baker	fbaker@emerald.acc.com
John Ballard	jballard@microsoft.com
William Barns	barns@gateway.mitre.org
Caralyn Brown	cbrown@wellfleet.com
Rick Bubenik	rick@arl.wustl.edu
John Burnett	jlb@adaptive.com
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
Dilip Chatwani	dilip@synoptics.com
Dean Cheng	dean@sun2.retix.com
George Clapp	clapp@ameris.center.il.ameritech.com
Robert Cole	robert.g.cole@qsun.att.com
Guy Dubuc	guyd@newbridge.com

Karl Fox	karl@morningstar.com
Cliff Frost	cliff@cmsa.berkeley.edu
Eugene Geer	ewg@nvuxr.bellcore.com
Daniel Grossman	dan@merlin.dev.cdx.mot.com
Joel Halpern	jmh@network.com
Kenneth Hayward	crm57d@bnr.ca
Frank Heath	heath@cmc.com
Juha Heinanen	juha.heinanen@datanet.tele.fi
Darryl Henry	henryd@imo-uvax.dca.mil
Zbigniew Kielczewski	zbig@eicom.qc.ca
Darren Kinley	kinley@crim.ca
Roman Kochan	roman@doelz.timeplex.com
Rajeev Kochhar	rxk@3com.com
Lawerence Lang	llang@cisco.com
Dory Leifer	del@terminator.um.cc.umich
Mark Lewis	mlewis@telebit.com
Susan Lin	suelin@ibm.com
Triet Lu	trietl@sparta.com
Andrew Malis	malis@bbn.com
David Martin	dem@fnal.fnal.gov
Chandy Nilakantan	csn@3com.com
Tushar Patel	tpatel@doelz.temeplex.com
Lars Poulsen	lars@cmc.com
Venkat Prasad	vsp@3com.com
Rex Pugh	pugh@hprnd.rose.hp.com
K. Ramakrishnan	rama@erlang.enet.dec.com
Christopher Ranch	cranch@novell.com
Jonathan Rodin	rodin@ftp.com
Benny Rodrig	4373580@mcimail.com
Manoel Rodrigues	manoel.rodrigues@att.com
L. Michael Sabo	lms@denver.ssds.com
Miguel Sasson	sasson@xylogics.com
Andreas Schachtner	afs@germany.eu.net
Erik Sherk	sherk@sura.net
Keith Sklower	sklower@cs.berkeley.edu
Pyda Srisuresh	suresh@3com.com
Brad Steinka	brad@python.eng.microcom.com
Walter Stickle	wls@ftp.com
Sally Tarquinio	sally@gateway.mitre.org
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Mark Therieau	markt@python.eng.microcom.com
Dean Throop	throop@dg-rtp.dg.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Carol Ward	cward@westnet.net
James Watt	jamesw@newbridge.com

Jay Westmark	westmark@dsi.gw.scri.fsu.edu
Cathy Wittbrodt	cjw@nersc.gov
Brian Wyld	brianw@spider.co.uk
Chin Yuan	cxyuan@pacbell.com
Paul Zawada	zawanda@ncsa.uiuc.edu

3.6.3 ISIS for IP Internets (isis)

Charter

Chair(s):

Ross Callon, callon@bigfut.enet.dec.com

Mailing Lists:

General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu
Archive:

Description of Working Group:

The IETF IS-IS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

Done — Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.

TBD — Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.

TBD — Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

Internet Drafts:

"Integrated IS-IS Management Information Base", 11/05/1991, Chris Gunner
<draft-ietf-isis-mib-00.txt>

Request For Comments:

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

CURRENT MEETING REPORT

Minutes of the ISIS for IP Internets Working Group (ISIS)

The ISIS Working Group met jointly with the Border Gateway Protocol Working Group (BGP). Please refer to the BGP Minutes for a summary of the session.

Attendees

Nagaraj Arunkumar	nak@3com.com
Syed Ashraf	irfan@pluto.dss.com
William Barns	barns@gateway.mitre.org
David Bolen	db31@nis.ans.net
Scott Brim	swb@cornell.edu
Ross Callon	callon@bigfut.enet.dec.com
Rob Coltun	rcoltun@ni.umd.edu
Osmund de Souza	osmund.desouza@att.com
Julio Escobar	jescobar@bbn.com
Dino Farinacci	dino@cisco.com
Stefan Fassbender	stf@easi.net
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Eugene Geer	ewg@nvuxr.bellcore.com
Terry Gray	gray@cac.washington.edu
Susan Hares	skh@merit.edu
Robert Hinden	hinden@eng.sun.com
Jeffrey Honig	jch@risci.cit.cornell.edu
Kathleen Huber	khuber@bbn.com
Bob Jeckell	rrj@3com.com
Dan Jordt	danj@nwnet.net
Manu Kaycee	kaycee@ctron.com
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELEMAIL/C=US/@sprint.com
Yoav Kluger	ykluger@fibhaifa.com
Bill Manning	bmanning@rice.edu
Matt Mathis	mathis@psc.edu
Dennis Morris	morrisd@imo-uvax.dca.mil
David Oran	oran@sneezy.lkg.dec.com
Rex Pugh	pugh@hprnd.rose.hp.com
Paul Raison	raison@xylogics.com
Yakov Rekhter	yakov@watson.ibm.com
Benny Rodrig	4373580@mcimail.com
L. Michael Sabo	lms@denver.ssds.com
Sharad Sanghi	sharad@ans.net
Miguel Sasson	sasson@xylogics.com
Tom Simons	nes!simons@phx.mcd.mot.com

Frank Solensky	solensky@clearpoint.com
William Streilein	bstreile@wellfleet.com
Sally Tarquinio	sally@gateway.mitre.org
Paul Traina	pst@cisco.com
Glenn Trewitt	trewitt@pa.dec.com
Chris Wheeler	cwheeler@cac.washington.edu
Jessica Yu	jyy@merit.edu

3.6.4 Inter-Domain Policy Routing (idpr)

Charter

Chair(s):
Martha Steenstrup, `msteenst@bbn.com`

Mailing Lists:
General Discussion: `idpr-wg@bbn.com`
To Subscribe: `idpr-wg-request@bbn.com`
Archive:

Description of Working Group:

The Inter Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Ongoing	Gain experience with the prototype in "real networks".
TBD	Develop gated version.
TBD	Add a small set of additional features and submit protocol into IETF standards process.

Internet Drafts:

"An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-idpr-architecture-03.txt>

"Inter-Domain Policy Routing Protocol Specification and Usage: Version 1", 03/05/1991, M. Steenstrup <draft-ietf-idpr-specv1-00.txt, or .ps>

"Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)", 07/22/1991, R.A. Woodburn <draft-ietf-idpr-mib-00.txt, .ps>

"Inter-Domain Policy Routing Configuration and Usage", 07/25/1991, H. Brown, M. Steenstrup <draft-ietf-idpr-configuration-00.txt>

Request For Comments:

RFC 1126 "Goals and functional requirements for inter-autonomous system routing"

3.6.5 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):

Steve Deering, deering@xerox.com

Mailing Lists:

General Discussion: mospf@devvax.tn.cornell.edu
To Subscribe: mospf-request@devvax.tn.cornell.edu
Archive:

Description of Working Group:

This Working Group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done — Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.

Done — Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.

Done — We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using the existing BSD OSPF code, written by Rob Coltun of the University of Maryland, as an example.

Done — Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations. The specification should now be ready to submit as an RFC.

Internet Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-01.ps>

CURRENT MEETING REPORT

Reported by Kevin Rowett/Tandem

Minutes of the Multicast Extensions to OSPF Working Group (MOSPF)

Agenda

- IETF Audio Multicast Experiment
- John Moy - Implementation Experiences of MOSPF
- Inter-Domain Multicast Routing
- Internet Draft to Proposed RFC?

Steve Deering described the efforts to provide an audiocast of various IETF sessions to people unable to attend. The equipment used was a Sun workstation, taking advantage of the Sparc built-in 79C30 codec. IP Multicasting was being used to relay audio via IP tunneling around routers which didn't support IP Multicast Routing. Destinations included Hawaii, Australia, Sweden, and the United Kingdom. Steve noted that the package provides for talk-back from the distant participants, but locally they have been unable to interconnect with the audio PA system. (Subsequently demonstrated at the Thursday Plenary).

John Moy's experiences with implementing MOSPF.

As a result of implementing, several changes were suggested:

1. There should be only one wildcard bit in router-LSA.

   ```
   0   0   0   0   wildcard  VL  ASBR  ABR
   ```

   Only one wildcard option was ever really needed.

2. Remove IGMP enable/disable switch:

   ```
   DL unicast -> IGMP off
   MC fwd off -> IGMP off
   Else IGMP on.
   ```

   IGMP switch doesn't provide any new information.

3. No more one directional links in OSPF (LSinfinity loses meaning in router-LSA). This is actually a subject of OSPF Working Group, but John wants both groups in synch. Eliminates disparity between Unicast and Multicast Routing Tables. Scott Brim questioned if OSPF can force IGMP off, especially if BGP is also present.

Multicast Forwarding Models

John proposed a mixture of the BSD and MOSPF models (see illustrations in accompanying viewgraphs). The proposed scheme could result in duplicate packets to multi-homed hosts.

For packets originated in a local machine, a check is required to see if local network looped back the packet. i.e., a LAN (or LAN interface) that does forward originated packets. This forces MOSPF forwarding decision to check source address does not equal the local address. TTL value is one less after MOSPF hop (seems reasonable).

John's Plans for Document Re-organization:

1. More on initialization of DIJKSTRA SPF algorithm.
2. Multicast portion of DIJKSTRA algorithm not optional; everyone must do identical DIJKSTRA for consistent tie-breaking.
3. Add an appendix with tie-breaking examples, just to make sure everyone gets this part right.

MOSPF Implementation Notes:

John noted that the hardest part of his MOSPF implementation was keeping the Unicast and Multicast Dijkstra Algorithm in synch. Multicast Routing Table starts with the Unicast table. If the Unicast table is not right, then multicast won't ever be. Must tie multicast forwarding cache maintenance to Unicast Routing Tables changes.

John asked if IGMP queries should be done over point-to-point links? Steve suggest that yes, they should be done on point-to-point links because, for example, the link might be to a host or to a non-MOSPF router.

John also suggested that it might be possible to do multipath multicast routing by using more creative tie-breaking during tree construction. The idea is intriguing, but needs more thought. Could possibly use a hash of the source address, or the (source, destination) pair, to select among multiple route choices.

John estimated that his implementation took him the equivalent of thirty full-time working days to complete; it added 20 percent to the size of the OSPF code.

It was agreed that, after John makes his final pass through the MOSPF draft, we would submit for publication as a Proposed Standard RFC.

The Agenda topic on Inter-Domain Multicasting was not addressed, due to time limitations and lack of enthusiasm. (We have gone over that territory at every previous meeting.)

Attendees

Nagaraj Arunkumar	nak@3com.com
John Chang	jrc@uswest.com
Dilip Chatwani	dilip@synoptics.com
Dean Cheng	dean@sun2.retix.com
Steve Deering	deering@xerox.com
Dino Farinacci	dino@cisco.com
Sang-Chul Han	schan@garam.kreonet.re.kr
Keunhee Han	khhan@gaya.kreonet.re.kr
Jeffrey Honig	jch@risci.cit.cornell.edu
John Leddy	jleddy@ibm.com
John Lynn	lynn@jalps2.cit.cornell.edu
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Rex Pugh	pugh@hprnd.rose.hp.com
Benny Rodrig	4373580@mcimail.com
Kevin Rowett	kevinr@tandem.com
Tom Simons	nes!simons@phx.mcd.mot.com
William Streilein	bstreile@wellfleet.com
Dave Thompson	daveth@ingate.microsoft.com
Yuan Wang	natadm!ycw@uunet.uu.net
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

Proposed MOSPF Changes

1. Only one wild-card bit in router-LSA:

O	O	O	O	Wild Card	VL	ASBR	ABR

2. Remove IGMP enable/disable switch:

   DL unicast ====> IGMP off
   MCfwd off ====> IGMP off
   Else IGMP on

3. No more one directional links in OSPF (LSInfinity loses meaning in router-LSA)

Mix BSD & MOSPF Membership Models

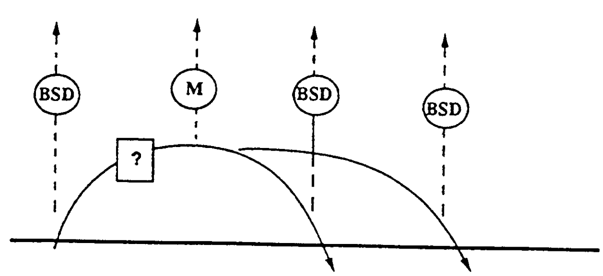

Note: All specified as part of MOSPF forwarding loop.

Internally-originated packets

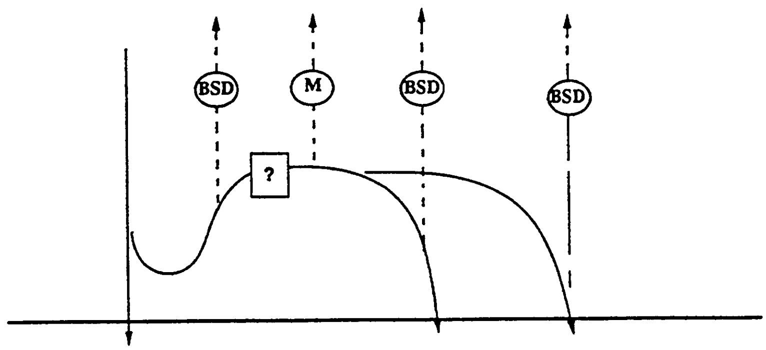

Note: packet marked as either received from net or direct from application (looped back)

Note: forwarding loop checks source against own addresses

Document Reorg

- More on Dijkstra initialization
- Multicast Dijkstra not optional
- Appendix with tie-breaking examples

Hardest to implement

- Keeping unicast and multicast Dijkstra in sync
- Must tie forwarding cache maintenance to unicast routing table changes

Questions/Ideas

- Should you send IGMP queries over P-P lines?
- Multipath through more creative tie-breaking during tree construction

3.6.6 Open Shortest Path First IGP (ospf)

Charter

Chair(s):

Mike Petry, petry@ni.umd.edu
John Moy, jmoy@proteon.com

Mailing Lists:

General Discussion: ospfigp@trantor.umd.edu
To Subscribe: ospfigp-request@trantor.umd.edu
Archive:

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

Done Design the routing protocol, and write its specification.

Done Develop multiple implementations, and test against each other.

Done Obtain performance data for the protocol.

Done Make changes to the specification (if necessary) and publish the protocol as a Draft Standard RFC.

TBD Gather operational experience with the OSPF protocol and submit the document as a Standard.

Internet Drafts:

"OSPF Version 2 Traps", 07/23/1991, Rob Coltun <draft-ietf-ospf-trapmib-00.txt>

Request For Comments:

RFC 1131 "OSPF specification"

RFC 1245 "OSPF Protocol Analysis"

RFC 1246 "Experience with the OSPF Protocol"

RFC 1247 "OSPF Version 2"

RFC 1248 "OSPF Version 2 Management Information Base"

RFC 1252 "OSPF Version 2 Management Information Base"

RFC 1253 "OSPF Version 2 Management Information Base"

CURRENT MEETING REPORT

Reported by John Moy/Proteon

Minutes of the Open Shortest Path First IGP Working Group (OSPF)

The meeting began with an announcement that IP multicast has been assigned an 802.5 functional address. This affects OSPF over 802.5, since many OSPF control packets (Hellos, etc.) are sent as IP multicasts. The functional address assignment will be documented in a short RFC, probably written by Steve Deering. It was suggested that to aid in the transition from the 802.5 broadcast address to the new functional address, a configuration knob be provided in OSPF implementations to select between broadcast/functional address. In any case, by the principle of "be liberal in what you receive", you should not reject a packet whose link level destination is the broadcast if you are instead expecting the functional address (this would enable a staged transition to the functional address, where you first enable reception of the functional address, and then in some future release, start sending it).

There are now a number of OSPF documents that will soon be ready for publication as RFCs: re-issues of the base specification and the OSPF MIB, the OSPF Trap MIB, and the OSPF NSSA Document. It is likely that these documents will want to reference each other, which may cause some logistical problems since you can't reference Internet Drafts (maybe they'll have to be issued as a set). In any case, it was decided to delay the publication of these documents until there were at least two interoperable implementations.

During the main part of the meeting, the following issues were discussed:

- We reviewed proposed changes to the base OSPF specification (RFC 1247). These changes are: a fix to a bug found in certain virtual link configurations, updating the TOS representation to include the new monetary cost bit, making summarization of routes into stub areas, optional, and an optimization to summarizing routes into transit areas. At the previous IETF a different fix to virtual link problem was discussed and rejected due to its complexity. The present fix, suggested independently be several people, is much simpler. Part of the fix involves removing the ability to assign a cost of LSInfinity to router interfaces. The fans of Strong TOS (e.g., Milo Medin and John Lekashman) were against this. John Moy was assigned the action item of further explaining why LSInfinity was a problem, and then negotiating with Milo and John. Fred Baker also mentioned that he had come across a situation where he wanted to condense inter-area routing information (not just intra-area, as specified in the current specification), but that the provision making summarization into stub areas optional would serve just as well.

  The proposed changes to RFC 1247 will be published shortly as an Internet Draft, followed by a revised version of the OSPF specification. All changes are backward compatible; there will be no need to increase the OSPF version number.

- Rob Coltun presented a collection of backward-compatible additions to the OSPF MIB. These additions included three variables to deal with OSPF Database Overflow: LSDBHiWater (read-only; the largest number of LSAs that have ever been in the router's database), LSDBOverflowWarning (read/write; when the number of LSAs hits this number a trap is generated) and LSDBLimit (read/write; when the number of LSAs hits this number the router takes further action to limit the size of the database as specified in a document to be written by John Moy). Also, a new table for type 5 external-LSAs is to be included, since in the current MIB it is not clear in which area ospfLsdbTable these LSAs should be reported. Fred Baker explained that, in order to be backward compatible, it would still be legal to report the type 5 externals LSAs in the old ospfLsdbTable.

  It was also noted that there should be two new ospfLsdbType values: 6 (for the group-membership-LSAs) and 7 (for the LSAs used by NSSA areas). In addition, since the interface cost LSInfinity is being removed, the comment in the ospfIfMetricMetric entry ("The value FFFF is distinguished to mean 'no router via this TOS' ") should be removed.

  Jeff Honig brought up a list of MIB variables that were named inconsistently. According to Fred Baker, we do not have to maintain the ASCII text representation of MIB variables to qualify for backward-compatibility, even though this may be an inconvenience to certain network management stations. Rob, Fred and Jeff are to go though the MIB to see which variables warrant renaming.

- Rob Coltun summarized the state of the OSPF Trap MIB (see Slide 2) which is very near to being finalized. There was some discussion on the best strategy for inhibiting traps when a router first starts, with the arguments for and against inhibiting traps on a per-interface basis being rehashed once again.

- Jeff Honig brought up the issue on how the OSPF MIB could handle multiple instances of OSPF running in the same box. While there is a straightforward technical solution to this problem (basically adding another index to all the MIB's tables), this is not backward-compatible and was viewed by several people as making the MIB overly complicated. Fred Baker suggested that this was a larger issue than just for OSPF, and suggested that we pass the problem (namely, how to monitor several instances of a protocol) on to the Network Management Working Group.

- Rob Coltun and Vince Fuller have completed a document describing the OSPF Not-so-stubby-area (NSSA) option, adding a motivational section to the outline that was presented at the previous meeting (Santa Fe), and completing the technical details. Rob presented an overview of the NSSA support, together with some of the more non-obvious details (see slide 3). Basically, NSSAs are a new type of area, similar to OSPF stub areas in that they do not handle type 5 external-LSAs (and so routers internal to these areas require less resources). However, NSSAs are capable of importing external information of their own, which will be converted to normal type 5 LSAs at the NSSA boundary. This enables, for example, RIP clouds to be hung off of NSSA areas.

Discussion centered upon whether we should be multiplexing several functions onto a single OSPF options bit (now that they are getting scarce), and the correct way to model the translation of external information that takes place at the NSSA boundary.

- Rob Coltun and Jeff Honig presented a proposal for another new OSPF option, which they called the PRI (Peripheral Router Interconnect) option (see Slides 4-7). This would provide a way to configure a set of distinguished OSPF routers, which would automatically discover each other and then be able to exchange additional information formatted as new LSA types. Jeff Honig explained an application of this whereby the PRI routers could exchange AS path information, obviating the need for IBGP. Rob and Jeff intend to write this up in more detail.

- Osmund deSouza led a discussion on how to run OSPF over Frame Relay (slides 8-12). One concern was that, since in real Frame Relay networks you are unlikely to have full mesh connectivity for PVCs, the NBMA model does not apply. In these cases, the Frame Relay would have to be treated as a collection of point-to-point links. A number of people thought that it might be possible to model Frame Relay as a collections of some number of NBMAs and serial lines, to achieve maximum efficiency (slide 11). To aid in this, Fred Baker thought that the Frame Relay MIB already had the provision to allocate particular sets of PVCs to particular IP networks.

  People agreed that, in order to guarantee interoperability, a document is needed to discuss the options for running OSPF over Frame Relay. This document could also discuss ways of detecting configuration errors (e.g., when some routers are configured for NBMA support and others are configured to see the Frame Relay as serial lines).

  Osmund also discussed a possibility whereby routers connected to a Frame Relay network could be grouped so that the groups were fully interconnected (slide 12). It was thought that the NBMA and Designated Router functions could be generalized to optimize running OSPF over such a configuration, although exactly how to implement this was unclear.

Attendees

Fred Baker	fbaker@emerald.acc.com
Dean Cheng	dean@sun2.retix.com
Rob Coltun	rcoltun@ni.umd.edu
Vince Fuller	vaf@stanford.edu
Leonard Gomelsky	gomelsky@ee.ucla.edu
Sang-Chul Han	schan@garam.kreonet.re.kr
Keunhee Han	khhan@gaya.kreonet.re.kr
Jeffrey Honig	jch@risci.cit.cornell.edu
Thomas Hutton	hutton@opus.sdsc.edu
George Kajos	kajos@coral.com
Yoav Kluger	ykluger@fibhaifa.com

John Lekashman	lekash@nas.nasa.gov
John Moy	jmoy@proteon.com
Michael Reilly	reilly@nsl.dec.com
Benny Rodrig	4373580@mcimail.com
Tom Simons	nes!simons@phx.mcd.mot.com
William Streilein	bstreile@wellfleet.com
Paul Traina	pst@cisco.com
Huyen Vu	vi@polaris.dca.mil
Yuan Wang	natadm!ycw@uunet.uu.net
Douglas Williams	dougw@ralumg.unet.ibm.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

TRAP UPDATE

- → LSDBHiWaterExceeded + LSDBOverflow TRAPS
  - Never supressed
- → Initial trap wait interval == largest dead interval (may want upper limit)
- → Only DR sends traps when appropriate

MIB ADDITIONS

- → LSDBHiWater + LSDB Limit + LSDB Overflow Warning READ/WRITE
- → TABLE FOR TYPE-5 LSAs

A FEW NSSA DETAILS

- → ABR must be ASBR
  - BECAUSE OF TYPE-7 to TYPE-5 TRANSLATION
- → P-bit option is only used by TYPE-7 LSA (too much of a hack?)
- → Must re-originate TYPE-5 LSA if TYPE-7 LSA forwarding addr changes.
  - ASBR must re-originate TYPE-7 if forwarding address intf changes to down
- → Router with highest Router ID must always translate TYPE-7 LSA, even if it uses a TYPE-5 LSA as its route.

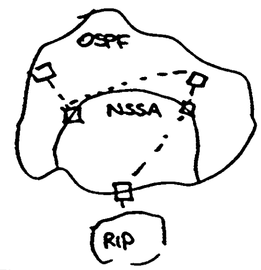

OSPF PRI OPTION

- → Peripheral Router Interconnect
- → Like NBMA network
  - Hello pkts sent as unicast to PRI neighbors
  - Adjacency formed between NBRS
  - Only certain types of LSAs are exchanged + flooded on PRI intf.
    - TYPE 8 LSAs

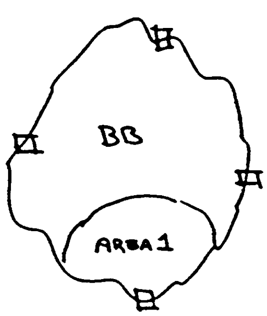

- DR is elected
- NET + Router LSA Equivalent not needed
  - Can be used to do inter-AS routing?
- PRI interface must be part BB if more than 1 area
- → Can work the same for "multicast network"
- → NBR Discovery
  - P-bit in Router LSA (must have E-bit set)
  - Option bit in TYPE-4 LSA (sum ASB)
  - Destination (NBR's) IP address is found in RTR LSA
  - Found in unused mask field of TYPE-4 LSA.
- → Can use new LSA to inform PRI NBRs of what type of LSAs it will accept.
- → PRI Packet Recognition
  - Auth type? (not really associated w/area)

→ APPLICATIONS
- ROUTE SERVER
  - PRI ROUTER IN CENTER OF AS
- EXCHANGE ADDRESS TRANSLATION INFO (DOMAIN → IP ADDRESS)
- TUNNEL (FOR MULTICAST)
- REPLACE IBGP

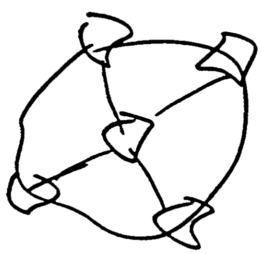

TYPE 8

Key | rtid | ~~path ID~~ | Attr ID

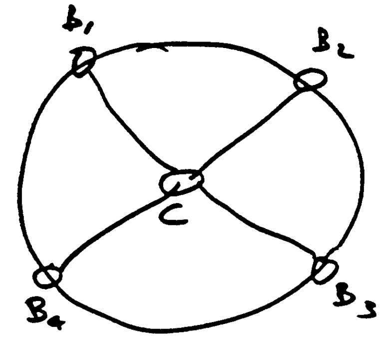

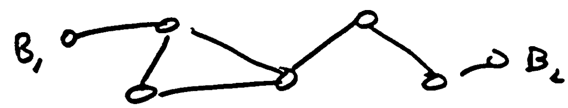

OSPF over FR

* By OSPF definition, FR networks are non-broadcast multiaccess
* OSPF v2 guidelines for NBMA nets
  - full interconnection required
    spec is not explicit about procedures
  - implementations may vary
* Will FR services provide the necessary capabilities?

FR = NBMA?

* FR PVC's are associated with "real" network resources (e.g., CIR), and may be expensive / scarce.
  - full mesh connectivity is not always realistic
  - OSPF over NBMA may not be the right model for FR nets.

What Is Realistic?

* POSSIBLE MODEL: PVC = Point-to-Point link
  - A collection of P2P OSPF links defined in a single physical FR interface

* QUESTIONS: what changes are needed
  - interface data structures, states
  - MIB variables

* The spec must be explicit and unambiguous

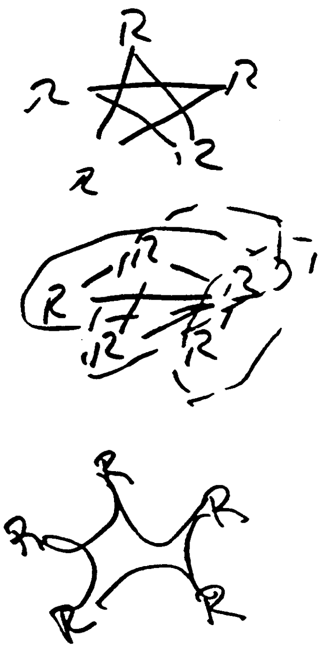

Another NBMA Issue

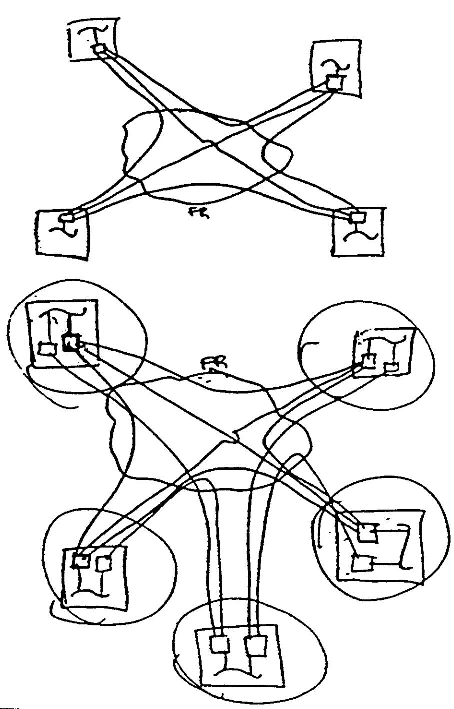

3.6.7 RIP Version II (ripv2)

Charter

Chair(s):
Gary Malkin, gmalkin@ftp.com

Mailing Lists:
General Discussion: ietf-rip@ftp.com
To Subscribe: ietf-rip-request@ftp.com
Archive: gmalkin/rip/archive@vax.ftp.com

Description of Working Group:

The RIPV2 Working Group is chartered to expand the RIP protocol, as defined in RFC 1058. The expansion will include the addition of subnet masks to the routing entries. The expansion may also include authentication, AS numbers, next hop address, MTU, or linkspeed. Since all routing protocols are required to have a MIB, one will be defined. The primary issue is the maintainance of backwards compatibility, which must be preserved.

The purpose of improving RIP is to make a simple, widely available protocol more useful. It is not intended that RIP-II be used in places where OSPF would be far better suited.

Goals and Milestones:

Mar 1991 Review of RIP-II Internet Draft to ensure the additions are useful and backwards compatible. Also ensure that the additions cannot cause routing problems.

Jul 1991 Final review of RIP-II Internet Draft and submission into the standards track. First review of RIP-II MIB.

TBD Review of implementations. Final review of MIB.

TBD Given successful implementation experience, advancement of RIP-II to Draft Standard. Submission of MIB into the standards track.

TBD Final meeting to achieve closure on any pending issues.

Internet Drafts:

"RIP Version 2 Addition of Subnet Masks", 08/14/1991, Gary Malkin <draft-ietf-malkin-rip-01.txt>

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the RIP Version II Working Group (RIPV2)

The RIPV2 Group settled on the final packet format and new field definitions, determined that no backwards compatibility issues exist and made a first pass at listing the objects needed in the MIB.

The Charter was accepted unchanged.

There were two changes to the packet format due to some confusion about the Routing Domain (RD) field. The RD field, as defined, is a per packet, user configurable parameter. It has, therefore, been moved into the Must Be Zero field in the RIP packet header. The RD field which was in the RIP entry has been renamed to the Route Tag. It is used to indicate that the route was learned from an external source. The exact use of this field is still under discussion; however, it has been determined that the contents of the RT field must be preserved when that route is propagated.

The subsumption of routes, made necessary by the addition of the subnet masks is still under work. The issues accompanying supernetting were also discussed, but no final solutions were reached. The Group did determine that there should be no user controls for this, since this could lead to black holes if routers were dis-similarly configured. It was decided that supernetting could not be used in the presence of RIP-I routers.

It should be explicitly mentioned that next hop is an advisory value. Next hop may also only be used for the directly connected network over which it was received.

It was decided that addressless links would not be considered.

The Group will need a new route type for MIB-II.

It should be mentioned, if RFC 1058 does not, that split horizon does not apply to routes learned via routing protocols other than RIP.

Attendees

Steve Alexander	stevea@i88.isc.com
Fred Baker	fbaker@emerald.acc.com
David Bridgham	dab@epilogue.com
Fwu-Tsair Chang	chang@n2sun1.ccl.itri.org.tw
Dean Cheng	dean@sun2.retix.com
Andrew Cherenson	arc@sgi.com
Vince Fuller	vaf@stanford.edu
Leonard Gomelsky	gomelsky@ee.ucla.edu

Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Jeffrey Honig	jch@risci.cit.cornell.edu
Darren Kinley	kinley@crim.ca
Rajeev Kochhar	rxk@3com.com
Mark Lewis	mlewis@telebit.com
Gary Malkin	gmalkin@xylogics.com
Glenn McGregor	ghm@merit.edu
Mely Pacheco	mely@tdd.sj.nec.com
Lars Poulsen	lars@cmc.com
Rex Pugh	pugh@hprnd.rose.hp.com
Paul Raison	raison@xylogics.com
Jonathan Rodin	rodin@ftp.com
Benny Rodrig	4373580@mcimail.com
John Seligson	johns@ultra.com
Tom Simons	nes!simons@phx.mcd.mot.com
Frank Solensky	solensky@clearpoint.com
Brad Steinka	brad@python.eng.microcom.com
William Streilein	bstreile@wellfleet.com
Kathleen Wilde	wilde@decvax.dec.com
Preston Wilson	preston@i88.isc.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker

The Security Area within the IETF is responsible for development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet.

Much of the work of the Security Area is performed in coordination with working groups in other areas. The Security Area Advisory Group (SAAG) is a group of security experts which provides both consulting help to other areas and direct management of working groups within the Security Area.

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 working groups within the Trusted Systems Interoperability Group.
3. Security relevant developments within working groups in areas other than security. These are marked according to the relevant area, viz., Applications, Internet Services, 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 SAAG met during the first and last working group period of the San Diego IETF. The first meeting was used to coordinate the activities for the week and the second meeting was used to report on the activities that have occurred.

During the week, of the twenty-two open work items on Monday, two work items were closed and two new work items were opened. The key activities for the week to report are working groups and work items in the security area: SNMP Security, Common Authentication Technology, Privacy Enhanced Mail, RFC 931 Revision, and Architectural Discussions.

SNMP Security Working Group (SNMPSEC)

There were three documents, published in January 1992, which are currently under consideration by the IAB.

Common Authentication Technology Working Group (CAT)

The basic idea is you have a set of applications that want access to one or more authentication mechanisms, for example Kerberos or the Distributed Authentication Security Service (DASS). There is a common program interface, a General Security Services Application Program Interface (GSS-API), that has been defined such that these applications can be written to be neutral with respect to which mechanism is actually employed. The binding with a mechanism takes place at some later time, currently compile time. This raises the question of how two applications each bound to a different mechanism would interoperate. In particular, if one peer supported Kerberos and the other peer DASS, would they be able to authenticate each other?

This question was the principal focus of the meetings during this week. Although the GSS-API was not designed with hybrid/common mechanism in mind, it was discovered that it would support such an objective through a number of different technical solutions. Most of the meeting time this week was spent identifying the requirements of a solution. It is believed that the objective is both technically feasible and achievable.

Privacy Enhanced Mail Working Group (PEM)

The specification of the key management infrastructure has been the principal source of controversy during the last few meetings. A revised document was prepared and distributed prior to this meeting, and was well received during this meeting. Along with the three other documents associated with PEM (Message Encryption and Authentication Procedures; Algorithms, Modes and Identifiers; Key Certification and Related Services), it will be submitted to the IESG by June in hopes of achieving publication as a Proposed Standard by the Boston IETF.

The publication of the documents stabilizes the specifications and sets the stage for the deployment of the Internet reference implementation of PEM. A set of action items predicating the deployment of PEM were identified and assigned. These items include establishing the necessary database mechanisms and software at the Internet Certification Authority (ICA) for resolving distinguished name conflicts (this is necessary in the absence of Directory Services), drafting an agreement to be used between the ICA and the Policy Certification Authorities (PCA), and facilitating the creation of PCAs (only one PCA proposal has been submitted to the ICA for review; others are expected soon). All of these items are non-technical and do not effect the publication of the specifications nor Beta testing the deployment of PEM, which is expected to begin soon.

RFC 931 Revision

RFC 931 is a specification of a protocol for a receiving peer of TCP connection request to ask a server on the originating host of the originating peer for an identifier associated with the originator of the request. The identifier would typically be the login name of the user initiating the request. This protocol was called an authentication server. As far as security is concerned, the value returned by the server is only meaningful in the context of that host, and is informational only since there are no assurances that a valid value is being returned.

There is an effort to revise the document to tighten up the syntax of the protocol and put it on the standards track. In addition, a public domain implementation exists that is currently being used by a modest number of sites.

Previously this effort was being led by Dan Bernstein, the author of the revised document and the implementation. In order to give the protocol the discussion it needs the effort has been restructured and a working group created with Mike St. Johns as the Chair, the author of the original RFC 931. In addition the revised protocol has been renamed to be called the Identity Server to better reflect its functionality.

Architectural Discussions

The SAAG in its two meetings spent a significant about of time discussing a security architecture for the Internet. Since the Privacy and Security Research Group (PSRG) is currently addressing the long-term objective(s) in this area, the majority of the discussion focussed on what the SAAG role could be in this area.

A number of action items were identified as a result of these discussions. First, Barbara Fraser from the CERT has agreed to draft a document identifying some near-term security goals that the IETF, in particular the SAAG, could be concerned about. This will help to focus SAAG discussions and guide interactions with working groups in other areas. We expect to have the document in time for discussions at the Boston IETF SAAG meeting.

Second, two Birds of a Feather sessions will be scheduled at the Boston IETF. One will be for Lower Layer Security and it will probably focus on IP layer authentication and encryption, although some discussion about the OSI TLSP and NLSP, and the SDNS SP3 and SP4 is expected.

The other BOF will be to discuss access control. Given the existence of authentication, in particular the strong authentication work of the CAT Working Group, the next question is what to do with the knowledge that you know who your peer is.

Finally, the routing area has received very little attention from security to date. With all of the activity in routing it has become essential that the Security Area become much more directly involved. Radia Perlman will be the liaison to the SAAG for the routing area. We will be discussing a routing area security plan during our Boston meeting.

3.7.1 Commercial Internet Protocol Security Option (cipso)

Charter

Chair(s):
Ron Sharp, rls@neptune.att.com

Mailing Lists:
General Discussion: cipso@wdl1.wdl.loral.com
To Subscribe: cipso-request@wdl1.wdl.loral.com
Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Commercial Internet Protocol Security Option (CIPSO) Working Group is chartered to define an IP security option that can be used to pass security information within and between security domains. This new security option will be modular in design to provide developers with a single software environment which can support multiple security domains.

The CIPSO protocol will support a large number of security domains. New security domains will be registered with the Internet Assigned Numbers Authority (IANA) and will be available with minimal difficulty to all parties.

There is currently in progress another IP security option referred to as IPSO (RFC 1108). IPSO is designed to support the security labels used by the U.S. Dept of Defense. CIPSO will be designed to provide labeling for the commercial, U.S. civilian and non-U.S. communities.

The Trusted Systems Interoperability Group (TSIG) has developed a document which defines a structure for the proposed CIPSO option. The Working Group will use this document as a foundation for developing an IETF CIPSO specification.

Goals and Milestones:

Done	Review and approve the Charter for the IETF CIPSO Working Group. Review revised TSIG CIPSO Specification.
Done	Review outstanding comments/issues from mailing list. Continue work on specification and prepare it for submission as an Internet Draft by the end of May.
Jul 1991	Review outstanding comments/issues from mailing list. The specification will be submitted to the IESG for consideration as a Proposed Standard.
Mar 1992	Submit specification to the IESG for consideration as a Draft Standard. There must be at least two interoperable implementations by this time.

Ongoing Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard.

Internet Drafts:

"Commercial IP Security Option", 12/03/1991, Trusted Sys Interop. Group (TSIG) <draft-ietf-cipso-ipsecurity-00.txt>

3.7.2 Common Authentication Technology (cat)

Charter

Chair(s):
John Linn, linn@zendia.enet.dec.com

Mailing Lists:
General Discussion: cat-ietf@mit.edu
To Subscribe: cat-ietf-request@mit.edu
Archive: /cat-ietf/archive@bitsy.mit.edu

Description of Working Group:

The goal of the Common Authentication Technology Working Group is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security-oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accommodating evolution of underlying technologies.

In support of these goals, the Working Group will pursue several interrelated tasks. We will work towards agreement on a common service interface allowing callers to invoke security services, and towards agreement on a common authentication token format, incorporating means to identify the mechanism type in conjunction with which authentication data elements should be interpreted. The CAT Working Group will also work towards agreements on suitable underlying mechanisms to implement security functions; two candidate architectures (Kerberos V5, based on secret-key technology and contributed by MIT, and X.509-based public-key Distributed Authentication Services being prepared for contribution by DEC) are under current consideration. The CAT Working Group will consult with other IETF working groups responsible for candidate caller protocols, pursuing and supporting design refinements as appropriate.

Goals and Milestones:

Done	Preliminary BOF session at IETF meeting, discussions with Telnet and Network Printing Working Groups.
Done	Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet Draft process.

Done First IETF meeting as full Working Group: review Charter distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.

Oct 1991 Update mechanism-independent Internet Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.

Nov 1991 Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.

Dec 1991 Submit service interface specification to RFC standards track.

Ongoing Progress Internet Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms.

Internet Drafts:

"Generic Security Service Application Program Interface", 06/12/1991, John Linn <draft-ietf-cat-genericsec-00.txt, .ps>

"The Kerberos Network Authentication Service", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-00.txt, .ps>

"Distributed Authentication Security Service", 11/04/1991, Charles Kaufman <draft-ietf-cat-dass-00.txt, .ps>

CURRENT MEETING REPORT

Reported by John Linn/DEC

Minutes of the Common Authentication Technology Working Group (CAT)

The March CAT meetings included discussion of standards advancement plans, and of interface extension requests made by ICL in support of ECMA authorization architecture. Most of the discussion was spent, however, on the evolving topic of a unified Internet authentication mechanism hybridizing Kerberos secret-key and DASS public-key technologies.

Standards and Rollout Plan

John Linn led a standards plan discussion, the result of which was a decision to recommend the GSS-API interface specifications for advancement to proposed standards. We anticipate that Kerberos and DASS specifications, as well as a specification for the planned unified mechanism, will follow in succession onto the standards track.

Two previously-cited technical topics regarding GSS-API were raised in this discussion: (1) the prospect of additional interfaces oriented to stream-oriented integration (as with UNIX(tm) sockets), tabled as being separately definable later in an upwardly-compatible fashion, and (2) the prospect of adding callouts so that user input (e.g., for passwords or hand-held authenticator information) could be collected at context establishment time. The latter was tabled because of lacking implementation experience, possible OS-specificity of approaches, and consideration that such data might more securely be acquired through end system trusted path facilities than via application mediation.

ICL Comments and ECMA Security Architecture

P. Rajaram stood in for Piers McMahon of ICL (who was unable to attend the meeting) in leading a discussion based on Piers' message as forwarded to the mailing list. Piers' message proposed interface extensions (GSS_Modify_Cred and GSS_Get_Attributes primitives) to support authorization features of the ECMA security architecture (as described in ECMA reports TR/46, TR/138), and Raj presented an overview of that architecture, the slides from which are included at the end of these Minutes. Interest was expressed in the prospect of having the ECMA reports available in FTP-accessible on-line form.

In Group discussion, it was recognized (consistent with discussion at the SAAG) that specific authorization support features, and related extensions in support thereof, would comprise a likely area for future IETF security work. Such work would consider not only ECMA inputs but also contributions from the Kerberos community as well as other sources, selecting an approach or defining a core intersection of multiple approaches. Any and all relevant inputs would be solicited. As with other Internet standards, prototyping results would be necessary for advancement. Lacking a concrete Internet community decision to adopt the ECMA architecture, no decision to incorporate the ECMA extension requests at this time was taken. Raj suggested that it might be useful to convene a BOF at a subsequent IETF meeting to further familiarize interested IETF participants with the ECMA architecture.

Specific points raised in ECMA-related discussion: To acquire a Privilege Attribute Certificate (PAC), a subject contacts a server. The PAC contains a sequence of attribute triples type, authority, value which govern the ways in which the PAC can be used, and an audit ID which allows audit accountability for actions independent of the privileges on which access controls are based, among other elements. Confusion was expressed about the circumstances under which a PAC must be confidentiality-protected in transfer, and about whether concurrent and separate authentication was necessary in order to demonstrate oneself as an authorized user of a particular PAC. Some of the answers were thought to depend on the particular attributes bound into the PAC, per definitions in ECMA TR/138.

Unified Authentication Mechanism

John Linn gave an overview of goals for the effort, Charlie Kaufman and Cliff Neuman presented alternative technical options, and Jeff Schiller led a discussion to collect requirement and priorities inputs to be considered in selecting among available alternatives.

Overview of Effort

John's overview slides contained the following points:

- DASS-Kerberos Unification: How Did We Get Here?
  - Cross-mechanism portability addressed in CAT.
  - Suggestion at Santa Fe SAAG: support universal interoperability for strong Internet authentication.
  - Kerberos and DASS designers and architects met in a series of interim meetings.
- Where are we going?
  - Internet-Draft documentation of hybrid mechanism to fit under CAT/GSS-API framework.
  - Ability to migrate applications from already-defined mechanisms to hybrid when available.
  - Common token format which can accommodate both public-key (PK) and secret-key (SK) authentication processes.
- Premises
  - Domains and endpoints can be built native to Kerberos-like SK and DASS-like PK technologies; all endpoints can interoperate.
  - Support for user, host, and process principals, represented by cryptographic keys.

  - Global naming (plan: X.500 Distinguished Names as basis within mechanism), trust path tied to naming hierarchy.

- Goals
  - PEM X.509 certificate infrastructure usable as a basis for scaling.
  - Domains equipped with public-key technology can operate without establishing on-line authentication servers.
  - Domains can be constructed without public-key technology.
  - Self-sufficient startup: can form a domain in isolation and later incorporate it into the broader hierarchy.
  - Can transport user-provided data (undefined by us) restricting the use of authentication tokens.
  - Avoid need for endpoints to contact foreign-mode support servers (KDCs, certificate stores, ...).

- Strong Authentication
  - Successful authentication requires either:
    * Current possession of principal's key.
    * Principal's authorization to act for principal with other (short-term) key + demonstration of that key.
  - Intercepted tokens can't be used by attackers to build new tokens for masquerade, or be successfully replayed outside narrow window.

- Four Directions

  We believe all can be made to work, and seek to resolve priorities and tradeoffs.

  - SK endpoints add complexity to interwork with PK.
  - PK endpoints add complexity to interwork with SK.
  - "Client makes right".
  - "Server makes right".

- Issues and Tradeoffs
  - Interoperability with existing/emerging technology bases.
  - What entities can accommodate complexity and performance demands?
  - What entities can and can't be changed feasibly?
  - What entities must perform what crypto-functions?

Alternative Technical Approaches

Charlie noted that support for interoperability between Kerberos-native and DASS-native authentication peers wasn't (unlike cross-mechanism portability) a chartered goal of GSS-API, and that it was a positively surprising result to discover, upon investigation, that such support within a unified mechanism below the interface in fact appeared to be possible, via any of several approaches. We confirmed the fact that the ability to support global scaling was intended.

Charlie's presented approach has the following characteristics:

- SK endpoints need not perform RSA operations or communicate with certificate stores.
- PK endpoints need not communicate with KDCs and the security of authentication between PK endpoints cannot be compromised by faulty KDCs.

It imposes the following impacts on particular system components:

- No impact on SK client.
- PK clients and servers must be able to end treewalks at a GKDC and use that GKDC's key in token generation and processing.
- SK server must interact with KDC to process incoming tickets arriving from PK domains.
- GKDC must be able to create and open PK tickets.

The fact of crossing from a public-key to a secret-key domain (or vice versa) needs to be determinable in a trusted fashion; naming prefix rules play an important part in this determination.

Cliff Neuman began the second CAT session by presenting an approach which matched Charlie's for the case of an SK client accessing a PK server, using tickets signed with the private key of a GKDC and integrable into the unified ticket format. It was observed that adoption of the unified format (in contrast, e.g., to use of Kerberos V5 tokens for SK cases) would require some level of change to all presently-extant peer systems.

Cliff presented an alternative approach for the case of a PK client accessing an SK server. A goal of this alternative was to avoid the need for an SK server to contact the GKDC, since such communication requires that the SK server be stateful in a manner divergent from the current Kerberos operational model. Cliff's proposal included a "Gateway Certificate Distribution Center" or GCDC, to which PK clients would DASS-authenticate and would receive, in response, a Kerberos ticket for the target SK server along with an associated encrypted session key. The GCDC, not the target server, would mediate interactions with intermediary SK authentication servers. In order to support both SK->PK and PK->SK accesses under this model, both GKDCs and GCDCs would be required; while these functions are logically distinct, they could likely be collocated.

Cliff summarized the impact which his proposal would impose on particular system components as follows:

- No impact on SK client.
- PK client must use different protocol in interacting with the last CDC in an outbound chain.
- No impact on SK server.
- PK server impact equivalent to Charlie's proposal.
- GKDC and GCDC must be able to create and open PK and SK tickets on behalf of clients.

Requirements/Priorities Evaluation

Jeff Schiller led a discussion at the end of the meeting with a goal of soliciting Group inputs on requirements and priorities for the unified mechanism. We created a comparison matrix, and the exercise's results served to validate many of the assumptions adopted by the designers. In particular, the Group showed popular acceptance for the idea of a dual-mode approach which employs PK and SK techniques for different cases. It was noted that allocation of computationally-intensive functions among components would be an additional useful metric, though not one which was included within this analysis.

Criteria included:

- Free availability of software, in terms of licensing, anonymous FTP-ability (ranked #3 criterion).
- Availability of source code implementations (ranked #2 criterion).
- Ability of approach to scale to world (by a broad margin, ranked as #1 criterion).
- Avoidance of on-line trusted KDCs (ranked #4 criterion).
- Simplicity/elegance of approach (construed by some attendees as equivalent to verifiability of protocol).
- Client simplicity (ranked #5 criterion, more important than server simplicity).
- Server simplicity.
- Compatibility with existing Kerberos (relatively low priority).
- Compatibility with SPX (lower priority than Kerberos compatibility).

Note: Jeff Schiller took an additional set of meeting notes which include pictures and which are available (in PostScript form) by anonymous FTP from bitsy.mit.edu with pathname: /cat-ietf/cat-wg-mar92-jis-picurenotes.ps

Attendees

Richard Basch	basch@mit.edu
Mark Baushke	mdb@nsd.3com.com
Larry Blunk	ljb@merit.edu
David Borman	dab@cray.com
Robert Bosen	bbosen@world.std.com
David Carrel	carrel@munin.com
Barbara Fraser	byf@cert.sei.cmu.edu
Bob Friesenhahn	pdrusa!bob@uunet.uu.net
James Galvin	galvin@tis.com
Terry Gray	gray@cac.washington.edu
Michael Grieves	mgrieves@chipcom.com
Neil Haller	nmh@bellcore.com
Alton Hoover	hoover@nis.ans.net
Thomas Kile	tgkile@brl.mil
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Kent Malave	kent@chang.austin.ibm.com
Scott Marcus	smarcus@bbn.com
Richard McBride	
Glenn McGregor	ghm@merit.edu
Clifford Neuman	bcn@isi.edu
Brad Passwaters	bjp@sura.net
P. Rajaram	rajaram@eng.sun.com
Michael Reilly	reilly@nsl.dec.com
Jim Reinstedler	jimr@ub.com
Gregory Ruth	gruth@bbn.com
Jeffrey Schiller	jis@mit.edu
Tom Simons	nes!simons@phx.mcd.mot.com
William Streilein	bstreile@wellfleet.com
Theodore Tso	tytso@mit.edu
Huyen Vu	vi@polaris.dca.mil
Preston Wilson	preston@i88.isc.com
Steven Wong	wong@took.enet.dec.com

rajaram@sun speaking-for piers-mcmahon@icl

Motivation:

- Extend GSS-API to enhance
  - o Authorization (useful to Kerberos-5)
  - o Delegation (more than just ON/OFF)
- Strategy
  - o Don't modify existing API
  - o Add a few new interfaces

SUMMARY of ECMA SECURITY ARCHITECTURE

- European Computer Manufacturers Assoc.
- This framework developed by a working group TC-32 / TG-9
- Supports many security models
  - o ACL based
  - o Capability based
  - o Label based (MAC)
  - & o Extensible combinations of above
- 10 security facilities
  - o promote modularity
  - o support interdomain security
  - o allow interoperability

THE 10 SECURITY FACILITIES

- o Authentication Service
- o Privilege Service
- o Subject Sponsor
- o Cryptographic Service
- o Secure Association
- o Authorization Service
- o Interdomain Service
- o Audit
- o Security Recovery
- o Security State

SUBJECTS and OBJECTS

Subjects access Objects

Subjects and Objects have Security Attributes

Subject Privilege Attributes
- o Identity, Group
- o Role
- o Clearance

Object Control Attributes
- o ACLs
- o Information Labels
- o Classifications

Ultimately, access is granted only if the Subject's privilege attributes "dominate" the Object's control attributes.

PRIVILEGE ATTRIBUTE CERTIFICATE

Contains:

- o a Sequence of Atributes
  { Type, Authority, Value }
- o Validity times
- o Contained PACs
- o Audit Identity
- o Signing Authority name (Domain authority)
- o Signature

The last two are required.

SIMPLE MODEL

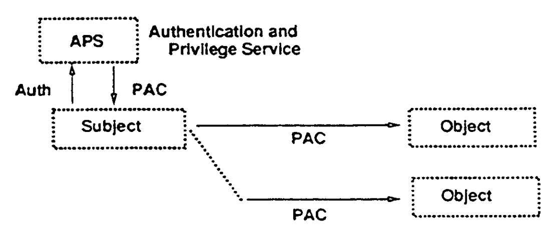

1) Subject authenticates itself to APS
2) Subject receives PAC
3) Subject offers PAC to Object and receives requested service.

SIMPLE MODEL - (cont'd)

A PAC contains both Identification AND Authorization Info.

- A PAC need not contain a Subject Identifier.
- An anonymous PAC may contain only a security clearance, and an audit ID.
- This may be enough to authorize access.
- Kerberos 5 and Kerberos/DCE can benefit from proposed API

PROPOSED API (greatly simplified)

GSS-Modify-Credentials

- [In] Cred handle
- [In] {Attributes & Values}...
- [In/Out] Credentials

GSS-Get-Attributes

- [In] Cred handle
- [In] Requested Attributes
- [Out] Returned {Attributes & Values}

DISCUSSION:

- GSS-API: for authentication only?
- Allow for ECMA, in addition to Kerberos & DASS?
- CAT ⟶ CAAT
- ECMA BOF?

3.7.3 Privacy-Enhanced Electronic Mail (pem)

Charter

Chair(s):
Stephen Kent, `kent@bbn.com`

Mailing Lists:
General Discussion: `pem-dev@tis.com`
To Subscribe: `pem-dev-request@tis.com`
Archive: `pem-dev-request@tis.com`

Description of Working Group:

PEM is the outgrowth of work by the Privacy and Security Research Group (PSRG) of the IRTF. At the heart of PEM is a set of procedures for transforming RFC 822 messages in such a fashion as to provide integrity, data origin authenticity, and optionally, confidentiality. PEM may be employed with either symmetric or asymmetric cryptographic key distribution mechanisms. Because the asymmetric (public-key) mechanisms are better suited to the large scale, heterogeneously administered environment characteristic of the Internet, to date only those mechanisms have been standardized. The standard form adopted by PEM is largely a profile of the CCITT X.509 (Directory Authentication Framework) recommendation.

PEM is defined by a series of documents. The first in the series defines the message processing procedures. The second defines the public-key certification system adopted for use with PEM. The third provides definitions and identifiers for various algorithms used by PEM. The fourth defines message formats and conventions for user registration, Certificate Revocation List (CRL) distribution, etc. (The first three of these were previously issued as RFCs 1113, 1114 and 1115. All documents have been revised and are being issed first as Internet Drafts.)

Goals and Milestones:

Done	Submit first, third, and fourth documents as Internet Drafts.
Done	Submit second document as Internet Draft.
Done	First IETF Working Group meeting to review Internet Drafts.
Sep 1991	Submit revised Internet Drafts based on comments received during Working Group meeting, from pem-dev mailing list, etc.
Nov 1991	Submit Internet Drafts to IESG for consideration as Proposed Standards.

Ongoing Revise Proposed Standards and submit to IESG for consideration as Draft Standard, and repeat for consideration as Internet Standard.

Internet Drafts:

"Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", 03/26/1991, John Linn <draft-ietf-pem-msgproc-01.txt>

"The MD5 Message-Digest Algorithm", 07/08/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md5-01.txt>

"The MD2 Message-Digest Algorithm", 07/10/1991, B. Kaliski <draft-rsadsi-kaliski-md2-00.txt>

"The MD4 Message-Digest Algorithm", 07/10/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md4-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV: Notary, Co-Issuer, CRL-Storing and CRL-Retrieving Services", 07/10/1991, B. Kaliski <draft-ietf-pem-notary-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", 07/17/1991, Steve Kent <draft-ietf-pem-keymgmt-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", 08/22/1991, David Balenson <draft-ietf-pem-algorithms-00.txt>

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

Minutes of the Privacy-Enhanced Mail Working Group (PEM)

The PEM Working Group met twice during the San Diego IETF meeting. The written comments provided by John Linn (and distributed to the PEM-DEV list) were reviewed and adopted. Several other changes are being made to the specification based on additional written comments from Jeff Schiller, plus PEM-DEV messages regarding encoding of the signature as described at the end of Appendix A. One significant change deals with the CRL database which was to be maintained by the ICA. Now, each PCA will be responsible for publishing a mail address to which CRL queries can be sent, and another for CRL updates from CAs. It is the responsibility of the PCAs, working with the ICA, to provide access to the complete CRL database via these interfaces. Details of the management of this database will not be part of the RFC, as this is exclusively a PCA-ICA interface issue. The description of the DN conflict detection database has been revised and the specification of the protocol for accessing this database will not appear as Appendix B, but rather will become a separate document.

Generally the version of the document (1114E) distributed prior to the meeting was well received and the consensus is that it is almost ready for publication as an Internet Draft. A revised version of this document (1114F) will be distributed immediately after the IETF meeting, via the PEM-DEV mailing list. New versions of 1113 and 1115 should be available very soon, and an updated version of FORMS (removing all the text not generally relevant to UAs) should be forthcoming as well. The current plan calls for publication of all four, updated documents as Internet Drafts as they become available, and then submission of all four as standards track RFCs.

Attendees

Robert Austein	sra@asylum.sf.ca.us
Mark Baushke	mdb@nsd.3com.com
Robert Bosen	bbosen@world.std.com
David Carrel	carrel@munin.com
Vinton Cerf	vcerf@nri.reston.va.
Stephen Crocker	crocker@tis.com
Steve Dusse	spock@rsa.com
Barbara Fraser	byf@cert.sei.cmu.edu
Maria Gallagher	maria@nsipo.nasa.gov
James Galvin	galvin@tis.com
Jeff Gumpf	jag@po.cwru.edu
Neil Haller	nmh@bellcore.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr

Randy Hoebelheinrich	rho@rowan.lanl.gov
Alton Hoover	hoover@nis.ans.net
John Ioannidis	ji@cs.columbia.edu
William Jackson	jackson@manta.nosc.mil
Stephen Kent	kent@bbn.com
Thomas Kile	tgkile@brl.mil
John Linn	linn@erlang.enet.dec.com
Steven Lunt	lunt@bellcore.com
Ellen McDermott	emcd@osf.org
Kevin Rowett	kevinr@tandem.com
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
Michael St. Johns	stjohns@umd5.umd.edu
R.B. Swenson	rswenson@igc.org
Theodore Tso	tytso@mit.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Wengyik Yeong	yeongw@psi.com
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.7.4 SNMP Security (snmpsec)

Charter

Chair(s):

James Galvin, galvin@tis.com
Keith McCloghrie, kzm@hls.com

Mailing Lists:

General Discussion: snmp-sec-dev@tis.com
To Subscribe: snmp-sec-dev-request@tis.com
Archive: snmp-sec-dev-request@tis.com

Description of Working Group:

The SNMP Security Working Group is chartered to determine the set of security services needed by the SNMP. The specification of those services, the supporting mechanisms, and the adjunct infrastructure will become an enhancement to the SNMP and eventually an Internet standard.

The specification must not alter the fundamental SNMP network management philosophy and must not entail changes to existing SNMP standards or framework.

Goals and Milestones:

Done	Publish Internet Draft specifications.
Jul 1991	Submit specification to IESG for consideration as a Proposed Standard.
Dec 1991	Submit specification to IESG for consideration as a Draft Standard.
Ongoing	Submit specification to IESG for consideration as a Standard.

Internet Drafts:

"SNMP Administrative Model", 04/09/1991, James Davin, James Galvin, Keith McCloghrie <draft-ietf-snmpsec-admin-02.txt, .ps>

"Definitions of Managed Objects for Administration of SNMP Parties", 04/09/1991, Keith McCloghrie, James R. Davin, James M. Galvin <draft-ietf-snmpsec-mib-02.txt>

"SNMP Security Protocols", 04/09/1991, James M. Galvin, Keith McCloghrie, James R. Davin <draft-ietf-snmpsec-protocols-02.txt, .ps>

3.8 Transport and Services Area

Director(s):

- David Borman: dab@cray.com

Three working groups met at the San Diego IETF meeting. They were the Audio/Video Transport Working Group, the Domain Name System Working Group, and the Services Location Protocol Working Group. In addition, a presentation was made on the new TCP options that have now been approved by the IAB for publication as a Proposed Standard.

Audio/Video Transport Working Group (AVT)

The Audio/Video Transport Working Group meet three times. One of the more exciting activities that happened was the bi-directional "audiocast" of the IETF plenaries and the AVT sessions from Australia to the United Kingdom. Since the purpose of the Working Group is to specify one or more protocols for doing audio/video experiments, they dove right into defining data packet header formats for real-time audio/video. For quick deployment, a UDP base was chosen, rather than running right over IP. Some of the fields considered in an additional header are timestamps, sequence numbers, decryption checksum, version number, encoding type, energy level and cumulative delay. A set of criteria was established to help decide which fields should be included, and discussion will continue via email. Addressing was another topic that was discussed, covering IP multicast address and UDP port numbers, and which parts can be dynamically allocated, and which parts need to be administratively assigned. Linkage between data and control was the next topic; simple applications can be done without a control protocol, but more complex applications will probably need some form of a control protocol. The Working Group may choose to specify a simple, interim control protocol once the transport protocol is defined. The last topic was software encoding. Although the Working Group is concerned with the transport layer, it is felt that in the interest of interoperability it would be useful to agree on some software compression techniques until hardware becomes generally available.

Domain Name System Working Group (DNS)

The DNS Working Group discussed two main topics. The first topic was brought up by Mike St. Johns, on the issue of getting a policy statement on adding sub- domains to the EDU and COM domains, so that names can be added in a consistent manner without overloading the top level domain. The rest of the meeting dealt with a proposed DNS MIB by Jon Saperia, which was distributed prior to the meeting, and a counter proposal by Rob Austin that was distributed at the meeting. Most of the discussion revolved around the differences between these two documents, and in the end both authors were asked to work together to produce a single proposal.

Service Location Protocol Working Group (SVRLOC)

The Services Location Working Group discussed an architecture proposal, which was presented by Scott Kaplan of FTP Software. Some of the issues outstanding are data representation and an RPC mechanism, both of which it was felt could be borrowed from someplace else. Multilingual support is another issue that the Working Group wants to address, and John Veizades will be persuing this topic for discussion at the next IETF meeting.

For more information on each of these sessions, please refer to the individual working group Minutes.

3.8.1 Audio/Video Transport (avt)

Charter

Chair(s):

Stephen Casner, casner@isi.edu

Mailing Lists:

General Discussion: rem-conf@es.net
To Subscribe: rem-conf-request@es.net
Archive: rem-conf/rem-conf-archive:nic.es.net

Description of Working Group:

The Audio/Video Transport Working Group was formed to specify protocols for real-time transmission of audio and video over UDP and IP multicast. The result may be independent protocols specific to each medium, or a common, lightweight, real-time transport protocol may be extracted.

UDP transmission of audio/video is only sufficient for small-scale experiments over fast portions of the Internet, but the transport protocols produced by this Working Group should be useful on a larger scale in the future when network-level resource management mechanisms are deployed to provide low-delay service and to guard against unfair consumption of bandwidth by audio/video traffic.

Similarly, initial experiments can work without any connection establishment procedure so long as a priori agreements on port numbers and coding types have been made. To go beyond that, we will need to address simple control protocols as well. Since IP multicast traffic may be received by anyone, the control protocols must handle authentication and key exchange so that the audio/video data can be encrypted. More sophisticated connection management is the subject of current research, and should be the topic of a follow-on working group.

Goals and Milestones:

Nov 1991 Define the scope of the Working Group, and who might contribute. Our first step will be to solicit contributions of potential protocols from projects that have already developed packet audio and video. From these contributions we will distill the appropriate protocol features.

Jan 1992 Conduct a teleconference Working Group meeting using a combination of packet audio and telephone. The topic will be a discussion of issues to be resolved in the process of synthesizing a new protocol. Make writing assignments for first-draft documents.

Mar 1992 Review first draft documents, determine necessary revisions. Follow-up discussion will occur on mailing list. Plan implementations.

May 1992 Teleconference meeting using implementations of draft protocols. Discuss draft revisions based on implementations, submit as Internet Drafts.

Jul 1992 Review updated draft, and assess whether these protocols should enter the standards track or be published only as experimental protocols. Make final revisions to drafts and give to IESG for publication as RFCs of appropriate type.

CURRENT MEETING REPORT

Reported by Steve Casner/ISI

Minutes of the Audio/Video Transport BOF (AVT)

Introduction: Goals, Scope of this Working Group

The AVT Working Group met for three sessions on Tuesday in San Diego. Audio from the presentations and discussions at these sessions was "audiocast" via UDP and IP multicast to participants at a number of locations ranging from Australia to the United Kingdom, and the remote participants were able to ask questions over the return path.

The purpose of this Working Group is to specify one or more experimental protocols to foster interoperation among multiple packet audio/video implementations in experiments such as this audiocast. The focus of the Working Group is short-term (see the Charter). Our first goal is to have the protocols defined and experimental implementations running in time for use in a second audiocast at the July, 1992 IETF meeting. Therefore, in this meeting we dove right in to a discussion of what the protocol should look like.

Data Packet Header Formats for Real-time Audio and Video

We need a "transport" protocol for real-time, continuous media. That means we don't want the retransmission and flow control of TCP, but we do want sequencing and checksumming. We could define a new protocol to fit directly over IP, but in keeping with the short-term scope of this Working Group, we choose to fit a new protocol over IP+UDP so it can be deployed quickly. Alternatively, another protocol that provides the necessary functions, such as ST-II, can be used. Those functions are port addressing, length, and (optional) checksumming.

The missing function is sequencing. Steve Casner described the data packet format of the Network Voice Protocol (NVP-II) which was serving this function for the audiocast of this meeting. The header is efficient (only 4 octets), but that makes some of the fields too small to support current requirements. To begin discussion of a replacement, the following strawman protocol, with only two fields, was proposed:

- 32-bit Timestamp (16 bits of seconds + 16-bit fraction)
- Sequence Number (could be less than 32 bits)

There was substantial discussion of the nature of the timestamp. It must have sufficient range to cover any network delay (segment lifetime) that might be expected, and it must have sufficient resolution to allow the desired degree of superposition and coordination among media streams. The bit allocation shown has a range of 18 hours and a resolution of 16 microseconds. The timestamp could be synchronous with the media sampling clock, in which case it would tick at the nominal sampling rate and drift with respect to real time, or it could be synchronous with real time. In the latter case, the timestamp could represent absolute real time if it were defined to be the middle 32 bits of a Network Time Protocol (NTP) timestamp, or it could be merely relative to real time.

For purposes of synchronization among multiple media sources, real-time timestamps should be used, though they need not be absolute. Julio Escobar from BBN gave a presentation on the Synchronization Protocol. It is based on globally synchronized clocks (e.g., using NTP) and defines a set of control protocol exchanges to establish an equalization delay for synchronized playback. It's only requirement on the data packet format is that a real-time-synchronous, relative timestamp be carried.

Although the timestamp field can be used to sequence the packets, it cannot be used to detect lost packets for media, such as voice, that suppress transmission when there is no activity. The sequence number serves that function. It could be smaller than 32 bits because the timestamp disambiguates wrap-around within the maximum segment lifetime. The number of bits should be large enough that the loss of exactly one sequence space of packets is a rare-enough event that failure to detect it is acceptable.

Steve Deering proposed some additional fields/functions to be included in the data packet header:

- Checksum (to validate decryption)
- Version Number
- Encoding Type

The UDP checksum cannot be used to validate decryption because it must be applied after encryption, so a separate checksum would be required. An alternative that would not require an additional field but would require more complex processing is to use the successful decryption of several properly sequenced packets as the validation of the key. On the other hand, including a checksum at this level, covering either just the header or header plus data, would also be useful with the ST-II protocol that does not checksum higher-layer protocols.

A version number would allow implementations to distinguish among multiple versions of the protocol.

The encoding type field might be used for several purposes. It could identify the particular compression algorithm used so that the receiver could select the correct decompression. However, if that selection would be constant over the life of the session, it could be communicated in an out-of-band control protocol.

If multiple media are sent on one port number, then an additional level of demultiplexing would be needed and the encoding field could serve that purpose. For layered (embedded) coding schemes, a field is needed to identify the separate layers, but this field might be here or might be consigned to the application-layer protocol. For the network to process the separate layers at different priorities, it is expected that some priority field would be needed in the network layer.

Finally, two fields from other packet audio protocols were considered:

- Energy Level (from Xerox PARC Phoenixphone)
- Cumulative Delay (from CCITT G.764)

For audio packets, the energy level is an indication of the sound volume in the packet. This may be useful to the receiver when mixing audio streams, for example. It could be recalculated by the receiver rather than being carried in the packet.

The CCITT recommendation G.764 Packetized Voice Protocol includes a field that records the cumulative variable queueing delays experienced by a packet in traversing the network. This may be useful for deadline-scheduling of packet forwarding, but it was decided that those experimenting with such algorithms would need to add the field in some lower layer.

Field Inclusion Criteria

We did not attempt to decide "in real time" what fields/functions should be included or excluded. Further discussion is expected via email. Instead, we established some criteria for inclusion of these and other fields in a real-time transport protocol:

- What percentage of applications would require the field? If only a small percentage, the field should be left to the application layer.
- What application functions we are trying to support with these fields? We may be able to combine functions by choosing the fields right.
- How should we tradeoff network bandwidth vs. processing and complexity of control algorithms? (The discussions of the checksum and energy fields are examples.)
- Would the field be constant in all packets at a given demultiplexing level? If so, that information could be implicit and carried in an out-of-band control protocol. Or is there a need for the data to be self-describing?
- Does the field/function "belong" at this level? Considerations include overlap with other layers, aesthetics, common practice and understanding.

Addressing

In the third session we discussed how addressing (multiplexing) should be divided among the layers. Steve Deering explained:

- The IP multicast address should identify a particular session or set of recipients. Two different sets of recipients should have two different addresses.

- The destination port address must be the same for all recipients if the packets are to be multicast, so the destination port must be administratively, not dynamically, assigned. Since the space of well-known port numbers is small, we can't assign separate ports for each kind of data in a multimedia session. It may be appropriate to have a control port and a data port, or perhaps to distinguish major data types, such as audio and video. Source port numbers are dynamically assigned and can distinguish multiple participants at one IP address.
- If there are multiple flows (e.g., audio and video) to one multicast address, it may be necessary to include another level of demultiplexing in the audio/video transport layer. This relates to the "encoding" field mentioned earlier.

Further discussion is needed to decide how much multiplexing should occur at each layer. There are considerations both of address space and of implementation (whether it is better to read multiple media on one socket or separate sockets, for example).

Linkages Between Data and Control

Flexible management of multimedia connections or sessions is the subject of current research and beyond the short-term scope of this Working Group. For simple application modes, such as an audiocast on an advertized "channel" (e.g., IP multicast address), operation is possible with no control protocol at all.

For spontaneous communication, there is pool of 2^16 IP multicast addresses from which an address may be chosen, but then that address must be communicated to the participants. This Group may define a simple interim protocol for this purpose as a second step (after the transport protocol). Some inputs to this process would be the "session protocol" used by the vat program, the Connection Control Protocol from ISI, and the DVC control protocol (see next section).

Software Encoding

Listed as a bonus topic on the agenda was a discussion of algorithms and protocols for software encoding of real-time media. This is not a main topic because such protocols should be at a layer above the transport. However, in keeping with the Working Group goal to foster interoperation and experimentation with packet audio and video, it may be valuable to agree on some (perhaps low performance) software compression techniques for use until hardware is generally available.

For this purpose, Paul Milazzo from BBN gave an update on the protocol used in the Desktop Video Conference program. DVC uses the low-cost VideoPix frame-grabber card for SPARCstations plus software compression to generate video at about 5 frames per second. The DVC protocol communicates sequences of video subimage blocks over UDP and uses TCP for the control connection. A recent enhancement is the ability to decode multiple streams (up to 6 so far).

Further Discussion

Thanks to Karen Sollins and Eve Schooler for taking the notes from which these Minutes were prepared. A longer report of the meeting with more detail will be posted to the mailing list rem-conf@es.net to stimulate discussion of the issues raised above. It is proposed that we also hold some packet audio teleconference meetings as needed to augment the e-mail discussion.

Attendees

Cynthia Bagwell	cbagwell@gateway.mitre.org
Ronald Broersma	ron@nosc.mil
Stephen Casner	casner@isi.edu
Andrew Cherenson	arc@sgi.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Richard Cogger	rhx@cornell.cit.bitnet
Steve Deering	deering@xerox.com
John Denune	jdenune@saturn.sdsu.edu
Jack Drescher	drescher@concert.net
Ralph Droms	droms@bucknell.edu
Julio Escobar	jescobar@bbn.com
Kevin Fall	kfall@ucsd.edu
Darryl Henry	henryd@imo-uvax.dca.mil
Russ Hobby	rdhobby@ucdavis.edu
Alton Hoover	hoover@nis.ans.net
Hung Huang	huang@ics.uci.edu
Thomas Hutton	hutton@opus.sdsc.edu
Ronald Jacoby	rj@sgi.com
Neil Katin	katin@eng.sun.com
Jim Knowles	jknowles@binky.arc.nasa.gov
Dory Leifer	del@terminator.um.cc.umich
Mark Lewis	mlewis@telebit.com
Arthur Lin	yalin@srv.pacbell.com
Bill Manning	bmanning@rice.edu
Ed Menze	menze@cs.arizona.edu
Paul Milazzo	milazzo@bbn.com
Mark Needleman	mhn@stubbs.ucop.edu
Ari Ollikainen	ari@es.net
Geir Pedersen	geir.pedersen@usit.uio.no
James Rees	jim.rees@umich.edu
Mike Ritter	mwritter@applelink.apple.com
Ron Roberts	roberts@jessica.stanford.edu
Allan Rubens	acr@merit.edu
Eve Schooler	schooler@isi.edu
Sam Sjogren	sjogren@tgv.com

Frank Solensky	solensky@clearpoint.com
Karen Sollins	sollins@lcs.mit.edu
Pyda Srisuresh	suresh@3com.com
Scott Stein	scotts@apple.com
Chiaohuey Tung	tung@ashland.intel.com
Vernor Vinge	vinge@sdsu.edu
Jil Westcott	Westcott@BBN.COM
Jonathan Whaley	whaley@concert.net
Peter Yee	yee@ames.arc.nasa.gov

3.8.2 Distributed File Systems (dfs)

Charter

Chair(s):
Peter Honeyman, honey@citi.umich.edu

Mailing Lists:
General Discussion: dfs-wg@citi.umich.edu
To Subscribe: dfs-wg-request@citi.umich.edu
Archive:

Description of Working Group:

Trans- and inter-continental distributed file systems are upon us. The consequences to the Internet of distributed file system protocol design and implementation decisions are sufficiently dire that we need to investigate whether the protocols being deployed are really suitable for use on the Internet. There's some evidence that the opposite is true, e.g., some distributed file systems protocols don't checksum their data, don't use reasonable MTUs, don't offer credible authentication or authorization services, don't attempt to avoid congestion, etc. Accordingly, a Working Group on DFS has been formed by the IETF. The Working Group will attempt to define guidelines for ways that distributed file systems should make use of the network, and to consider whether any existing distributed file systems are appropriate candidates for Internet standardization. The Working Group will also take a look at the various file system protocols to see whether they make data more vulnerable. This is a problem that is especially severe for Internet users, and a place where the IETF may wish to exert some influence, both on vendor offerings and user expectations.

Goals and Milestones:

May 1990 Generate an RFC with guidelines that define appropriate behavior of distributed file systems in an internet environment.

3.8.3 Domain Name System (dns)

Charter

Chair(s):
Michael Reilly, reilly@nsl.dec.com

Mailing Lists:
General Discussion: dns-wg@nsl.dec.com
To Subscribe: dns-wg-request@nsl.dec.com
Archive:

Description of Working Group:

The DNS Working Group is concerned with the operation of name servers on the Internet. We do not operate name servers but serve as a focal point for the people who do operate them. We are also concerned with the Domain Name System itself. Changes to the existing RFC's, for example, are discussed by the Working Group. If changes to the RFC's or additional DNS related RFC's are deemed necessary the Working Group will propose them and will prepare the associated documents.

Because we intend to serve as the focal point for people operating name servers, one of our projects will be to assist anyone bringing up a name server by publishing a collection of useful hints, tips and operational experience learned by the people already running name servers.

The DNS Working Group will also take an active role in the dissemination of solutions to problems and bugs encountered while running various name server implementations. We will also provide guidance to anyone writing a new name server implementation, whenever possible.

Goals and Milestones:

TBD	Adding DNS variables to the MIB.
TBD	Hints, tips, and operations guide for DNS software.
TBD	Implementation catalog for DNS software.
TBD	Discussion of adding load balancing capability to the DNS.
TBD	Discussion of adding a Responsible Person Record.
TBD	Discussion of adding network naming capability to the DNS.

CURRENT MEETING REPORT

Reported by Michael Reilly/DEC

Minutes of the Domain Name System Working Group (DNS)

The announced purpose of the DNS Working Group meeting held at the San Diego IETF was to discuss the DNS MIB proposed by Jon Saperia. A draft of the proposed MIB was circulated on the namedroppers mailing list approximately two weeks before the San Diego meeting.

The meeting opened with a proposal from Mike St. Johns that the Working Group write a policy statement describing the addition of subdomains to one of the existing top level domains. The policy would help answer such questions as "Should elementary and/or high schools be included in the EDU domain?" or, "Should all businesses, regardless of size, be subdomains of the COM domain?". He proposed that the Group work with Jon Postel to insure that domains which didn't meet the criteria for inclusion within a domain were registered in a more general domain. Mike indicated that the MIL domain already had such a policy statement. He was asked to send a copy to the namedroppers mailing list as a starting point for further discussions. On behalf of the DNS Working Group, the Working Group Chair will ask the IESG for help in determining if this is something which should be pursued.

At the same time as discussion of the proposed MIB was begun, Rob Austein passed around copies of an alternative proposal. Since this new proposal had not previously been circulated, the Working Group was not able to compare it with the original proposal.

Much of the meeting was taken up by a discussion of the differences between the two documents. At the conclusion of the discussion the Working Group asked if a single proposal could be produced incorporating features of both of the original proposals. Rob and Jon agreed to work on a common proposal which they both could support.

Since the Working Group meeting the principle authors of the two MIB proposals have been Working together to resolve the differences between the proposals. It is expected that a revised proposal will be ready for discussion at the next Working Group meeting. Any remaining differences will be worked out at that time.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Robert Austein	sra@asylum.sf.ca.us
David Borman	dab@cray.com
David Bridgham	dab@epilogue.com
Eric Brunner	brunner@practic.com
Andrew Cherenson	arc@sgi.com

John Curran	jcurran@bbn.com
Kathleen Dodd	kathy@gateway.mitre.org
Billy Durham	durham@fishery.honeywell.com
Scott Kaplan	scott@wco.ftp.com
John Larson	jlarson@parc.xerox.com
Walter Lazear	lazear@gateway.mitre.org
Kenneth Lindahl	lindahl@violet.berkeley.edu
Henry Miller	henrym@sacusr.mp.usbr.gov
Michael Newell	mnewell@nhqvax.hg.nasa.gov
William Nowicki	nowicki@legato.com
Michael Reilly	reilly@nsl.dec.com
Robert Reschly	reschly@brl.mil
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Tim Seaver	tas@concert.net
Michael St. Johns	stjohns@umd5.umd.edu
Craig Summerhill	craig@cni.org
Justin Walker	justin@apple.com
Walter Wimer	walter.wimer@andrew.cmu.edu

3.8.4 Service Location Protocol (svrloc)

Charter

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

Mailing Lists:
General Discussion: srv-location@apple.com
To Subscribe: srv-location-request@apple.com
Archive: pub/srv-location/svr-loc-archive

Description of Working Group:

The Service Location Working Group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based services location protocols. Protocols that would be designed by this Group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the services location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this Group to understand the operation of services location and then determine the correct action in their view whether it be to use current protocols to suggest a services location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

Done — Open discussion and determine if a Working Group should be formed.

Done — Continue discussion trying to refine the problem statement and possible resolutions.

Jul 1991 — Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Minutes of the Service Location Protocol Working Group (SVRLOC)

Scott Kaplan of FTP Software made a presentation of an architecture proposal he has been working on. Some of the principals that have been made in the design of this protocol have been that the end nodes (these are the nodes searching for services) would have no user interface configuration. This means that the information that is reported in the user interface of the end node is information that is contained in the network, in the services access points or in surrogates for these access points. Other protocol design principals are that a data representation language can be borrowed or stolen, that some RPC mechanism can also be borrowed. Some issues that were brought up in the discussion of the protocol is that this protocol is not SNMP even though there are similarities between the two protocols (both protocols can transmit a list of attributes that are supported by the service provider). This protocol should deal with both the dentist office case (plug and play IP in a small unconnected network) as well as in a well connected global internet. The protocol design document will be refined and posted to the mailing list.

Issues of multilingual support will be worked on by John Veizades and will be presented at the next Working Group meeting in Boston. It is assumed that there is some character encoding standard that can be used for this issue.

There will be a BOF session on this topic Thursday night, May 21st at the INTEROP in Washington, DC.

Scott and John met with the IRTF Working Group on resource location. There are many similar issues in the work that this Group is trying to solve.

Attendees

Gregory Bruell	gob@shiva.com
Andrew Cherenson	arc@sgi.com
Cyrus Chow	cchow@ames.arc.nasa.gov
Ralph Droms	droms@bucknell.edu
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Jeff Gumpf	jag@po.cwru.edu
Nat Howard	nrh@bellcore.com
Scott Kaplan	scott@wco.ftp.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Holly Knight	holly@apple.com
Joshua Littlefield	josh@cayman.com
Leo McLaughlin	ljm@wco.ftp.com

Greg Minshall	minshall@wc.novell.com
J. Bradford Parker	brad@cayman.com
Stephanie Price Marasciullo	price@cmc.com
Christopher Ranch	cranch@novell.com
Mike Ritter	mwritter@applelink.apple.com
Andreas Schachtner	afs@germany.eu.net
Ursula Sinkewicz	sinkewic@netrix.nac.dec.com
Mark Smith	mcs@umich.edu
Walter Stickle	wls@ftp.com
Dave Thompson	daveth@ingate.microsoft.com
John Veizades	veizades@apple.com
Lee Wade	wade@nsipo.arc.nasa.gov
Scott Williamson	scottw@nic.ddn.mil
Walter Wimer	walter.wimer@andrew.cmu.edu

3.8.5 TCP Large Windows (tcplw)

Charter

Chair(s):
David Borman, dab@cray.com

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

Description of Working Group:

The TCP Large Windows Working Group is chartered to produce a specification for the use of TCP on high delay, high bandwidth paths. To this end, this Working Group recommended RFC 1072 "TCP extensions for long-delay paths" and RFC 1185 "TCP Extension for High-Speed Paths" be published jointly as a Proposed Standard. Deficiencies in the technical details of the documents were identified by the End-to-End Research Group of the IRTF. Rather than progress the standard with known deficiencies, the IESG tasked the End-to-End Research Group to fix and merge these two documents into a single protocol specification document. This review was done on the eze-interest@isi.edu mailing list.

The TCP Large Windows Working Group is being resurrected for a one time meeting, to review and if appropriate, approve this new document.

Goals and Milestones:

Nov 1991 Review the TCP Extended Window Size proposal from the IRSG End to End Research Group and if acceptable, recommend it for standards status.

Internet Drafts:

"TCP Extensions for High Performance", 11/12/1991, V. Jacobson, R. Braden, D. Borman <draft-ietf-tcplw-tcpext-01.txt>

3.8.6 Trusted Network File Systems (tnfs)

Charter

Chair(s):
Fred Glover, fglover@decvax.dec.com

Mailing Lists:
General Discussion: tnfs@wdl1.wdl.loral.com
To Subscribe: tnfs-request@wdl1.wdl.loral.com
Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Trusted Network File System (TNFS) Working Group is chartered to define protocol extensions to the Network File System (NFS) Version 2 protocol which support network file access in a multilevel secure (MLS) Internet environment. MLS functionality includes mandatory access control (MAC), discretionary access control (DAC), authentication, audit- ing, documentation, and other items as identified in the Trusted Computer System Evaluation Criteria (TCSEC) and Com- partmented Mode Workstation (CMW) documents.

The primary objective of this Working Group is to specify extensions to the NFS V2 protocol which support network file access between MLS systems. It is intended that these extensions should introduce only a minimal impact on the existing NFS V2 environment, and that unmodified NFS V2 clients and servers will continue to be fully supported.

Transferring information between MLS systems requires exchanging additional security information along with the file data. The general approach to be used in extending the NFS V2 protocol is to transport additional user context in the form of an extended NFS UNIX style credential between a Trusted NFS (TNFS) client and server, and to map that con- text into the appropriate server security policies which address file access. In addition, file security attributes are to be returned with each TNFS procedure call. Other- wise, the NFS V2 protocol remains essentially unchanged.

The Trusted System Interoperability Group (TSIG) has already developed a specification which defines a set of MLS exten- sions for NFS V2, and has also planned for the future integration of Kerberos as the authentication mechanism. The TNFS Working Group should be able to use the TSIG Trusted NFS document as a foundation, and to complete the IETF TNFS specification within the next 3-6 months.

Goals and Milestones:

Done	Review and approve the TNFS Working Group Charter, review revised TSIG TNFS Specification, and publish a proposed standard following the July meeting.

Jul 1991 Review revised TSIG TNFS Specification.

Nov 1991 Publish a Proposed Standard following the July meeting.

Oct 1991 Review outstanding comments/issues from mailing list.

Oct 1991 Make any final revisions to TNFS document based on comments, issues, and interoperability testing.

Mar 1992 Request IESG to make the revised document a Draft Standard.

Mar 1991 Verify the interoperability of TNFS implementations at the 1992 NFS Connectathon.

Internet Drafts:

"A Specification of Trusted NFS (TNFS) Protocol Extensions", 07/23/1991, Fred Glover <draft-ietf-tnfs-spec-00.txt, .ps>

INTERIM MEETING REPORT

Reported by Fred Glover/DEC

Minutes of the Trusted Network File Systems Working Group (TNFS)

General Summary - January 1992

The TNFS Working Group met in January as a joint IETF/TSIG Working Group. By the end of this meeting, we achieved closure on the TNFS document; all outstanding issues were resolved. The updated document will be archived, and a request will be made to advance the document from Internet Draft to Proposed Standard.

Meeting Summary

During the January meeting, we:

- Inspected (page by page review) the modifications to the TNFS document.
- Reviewed the TKM specification.
- Discussed the DNSIX token mapping mechanism.
- Discussed plans for associated TNFS documentation.
- Reviewed interoperability opportunities, future plans.
- Discussed TSIG document numbering; assigned TSIG document numbers.

TNFS Document Review

The IETF TNFS document has been available for comment in the IETF Draft directory and TNFS archive since July, 1991. During the January meeting, the Working Group completed work on the resolution of all of the outstanding draft comments, and voted to advance the draft to that of Proposed Standard. Conforming implementations are being encouraged in order to support future interoperability testing.

Final updates to the TNFS document include:

- The distinguished value will be changed from ZERO to "all bits on".
- The document will be updated to clarify the use of process and file privileges.
- A single privilege token will be included in the credential and file attribute structures; this token may be used to represent either a single or multiple privilege sets.
- Client side auditing will be enabled by default; a note regarding auditing of non-MLS clients will be included.
- File name labeling and multi-level directories will be included in the TNFS specification, along with new protocol operations to support them.

- The client caching section will be updated to reflect additional considerations in the use of cached information after a modification to a process' security attributes.

The updated document will be included in the IETF and TNFS archives.

Token Manager Review

Closure was also reached for TKM document, in its support for TNFS, with the following updates:

- Include new protocol operation for inverse mapping (attribute to token).
- Update the document to use the AUTH_UNIX credential; this is required to eliminate initialization deadlock.

The updated TKM document will be placed into the IETF Draft directory and the TSIG TNFS archive.

DNSIX Token Mapping

Charlie Watt presented an overview of the DNSIX token mapping mechanism. The Working Group provided a few editorial comments back to him. The major issue identified was whether this Token Mapping model would be made public, and thus available to the IETF community. Charlie believed that this would happen in the future, and a representative from the government also confirmed that this was planned. At the present time, however, the document is not publicly available. So any possibilities for potential IETF use are delayed until the document can be distributed.

Associated TNFS Documentation

The Working Group recommended that the TNFS Implementation and TNFS Administration guides be updated based on the October '91 reviews of these documents, and then placed in the TNFS archive and the IETF Draft directory as informational RFCs.

Interoperability Testing

The Working Group reviewed the progress of implementations, and discussed the possibility of interoperability testing at the April IETF/TSIG meeting. A proposed test plan was reviewed, which would be used for this purpose.

TSIG Document Numbering

During the plenary session, a document numbering scheme was selected. Using this scheme, the Working Group assigned the following document numbers:

- TNFS Specification:TSIG-TNFS-001.02.01

- TNFS Test Plan:TSIG-TNFS-002.01.01
- TNFS Test Attributes:TSIG-TNFS-003.01.01
- TNFS Implementation Guide:TSIG-TNFS-004.01.01
- TNFS Administration Guide:TSIG-TNFS-005.01.01
- TNFS TKM Specification:TSIG-TNFS-006.01.01
- TNFS tnfs.hTSIG-TNFS-007.01.01

Next Meeting

The TNFS Working Group will plan to meet as both a TSIG and an IETF Working Group at the April meeting in Mountain View, California. At that meeting, we plan to:

- Review the "final" version of the TNFS documents (updated documents placed into the TNFS archive and IETF Drafts directory: Fred, Fran, Carl, Ali).
- Review the interoperability test plan (all).
- Update/develop NFS test suite extension for TNFS (Fran).
- Identify conforming implementations to support our request to transition our TNFS document (all).
- Investigate NFS lock manager and status monitor for B1/CMW extensions (Charlie).
- Commence identification of auditable TNFS events (Mark).
- Place "tnfs.h", test plan, test attributes into TNFS archive (Fred).

The next meeting is planned for April 28-30, 1992 at Silicon Graphics in Mountain View, California.

Attendees

Fran Fadden	fran@decvax.dec.com
Jonathon Fraser	
Fred Glover	fglover@decvax.dec.com
Ali Gohshan	
Narayan Makaram	
Mark Saake	saake@llnl.gov
Carl Smith	cs@eng.sun.com
Charles Watt	

3.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Eight working groups in the User Services Area of the IETF met in San Diego. Two of the eight are newly formed working groups: Internet School Networking (ISN) and Internet Anonymous FTP Archives (IAFA). Below is a summary of ISI's User Services IETF activities.

Directory Information Services Infrastructure Working Group (DISI)

The DISI Working Group, Chaired by Christopher Weider was established to provide a forum for defining user requirements in X.500. It is an offshoot of the OSI Directory Services Group and is a combined effort of the User Services Area and the OSI Integration Area of the IETF. Three papers were published as FYI RFCs:

- Lang, R., and R. Wright, "A Catalog of Available X.500 Implementations", FYI 11, RFC 1292, January 1992.
- Weider, C., and J. Reynolds, "Executive Introduction to Directory Services Using the X.500 Protocol", FYI 13, RFC 1308, March 1992.
- Weider, C., Reynolds, J., and S. Heker, "Technical Overview of Directory Services Using the X.500 Protocol", FYI 14, RFC 1309, March 1992.

In DISI's Charter, they had some additional documentation they wanted to produce. There was a discussion of what other documents to write in addition to those that are listed in the charter. There were five different topics presented by Working Group attendees. An advanced usages document, how to get registered, where do I belong, how to keep your X.500 up to date, your directory up to date, a pilot project catalogue and a DSA setup guide.

The general consensus was an "advanced usages" draft, a "how to get registered" document and a "where do I belong" draft would be the next endeavors that they would like to do. Second in priority would be the "pilot project catalogue" document. Writing assignments and volunteers were tasked to go off and create drafts before the next IETF in Boston.

Internet Anonymous FTP Archives Working Group (IAFA)

This Group is Chaired by Peter Deutsch and Alan Emtage. This is a new Working Group which met for the second time in San Diego.

The IAFA 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 IAFA Working Group brought in a rough draft of a guide to FTP site administration which had been worked on via electronic mail in the last few months. This is what this Working Group primarily focussed on during their session. Discussion included data formatting issues and FTP extensions, caching, mirroring, redundancy and resources.

Internet School Networking Working Group (ISN)

The ISN Working Group is Chaired by John Clement, Art St. George, and Connie Stout. This is also a new Working Group which met for the second time in San Diego.

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.

ISN's session gathers educators and Internet folks together. ISN also had a rough working draft. The draft presents multi-generic connectivity models for schools. Their focus is primarily on IP connectivity. ISN's document and intent is not to recommend to schools any one model, but to present various suggestions and various models for school systems to look at. Then the school systems can take these models, choose one for their needs, their students' needs, and also how much they have in their budgets to contribute to equipment and software.

Internet User Glossary Working Group (USERGLOS)

USERGlOS is Chaired by Gary Malkin and Tracy LaQuey Parker. The USERGLOS Working Group is chartered to create an Internet specific glossary of networking terms and acronyms for the Internet community.

A draft document was ready for review at this session. USERGLOS had a two marathon sessions, one in the morning, and since they were going so well, they took a break and came back from 4:00 to 6:00 and continued to work on the glossary. This Group got an enormous amount accomplished in one day and those participants are to be commended for their stamina! They specifically worked on finding Internet specific terms that are needed in this glossary and weeded out a lot of words that were not pertinent.

Network Information Services Infrastructure Working Group (NISI)

This Group is exploring the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. One paper was published as an FYI RFC:

- Sitzler, D., Smith, P., and A. Marine, "Building a Network Information Services Infrastructure, FYI 12, RFC 1302, February 1992.

NISI's session focused on discussing what more there is to do. The NISI chairs feel that there's a lot more that they could contribute and they did not necessarily did not want to go dormant and come back out at a later time. The two topics that emerged were nethelp and a list of services. The two groups were tasked with volunteers, about five or six in each group, that are going to discuss what they need to develop in these two areas before the next IETF.

NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

NOCTOOL2, is Chaired by Robert Enger and Darren Kinley. The "Son of NOCTools" Working Group are updating and revising their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

The NOCTOOL2 Working Group has been a little slow in getting the entries in. Entries are continuing to slowly arrive, but Bob and Darren are going to push hard on final document compilation and completion for an Internet-Draft. Bob and Darren have pledged to get a document out in the next two months.

User Documents Revisions Working Group (USERDOC2)

USERDOC2, chaired by Ellen Hoffman and Lenore Jackson. The User-Doc 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.

USERDOC's original bibliography was published in August of 1990. It is sorely out of date. This is the new revised Working Group. It has a revised charter. They met during the User Services Working Group session as they had time constraint problems. Items discussed included what items should included in the outdated document, and what items should be deleted in revising the bibliography.

User Services Working Group (USWG)

The USWG Group is Chaired by Joyce K. Reynolds. 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.

User-Doc Working Group took up about half the session in their discussions. We also had a two reports, one on the RIPE meetings in Amsterdam in January, and the FARNET meeting in February. RIPE was originally chartered as a technical community. They have realized that user services is now something very important they should focus on, especially with

East Central Europe and what is going on in Russia right now in user service needs. They have now just formed a User Information Services Working Group in the RIPE community. The FARNET meeting focussed on, "Hardening of the Interim NREN". FARNET meeting participants that were at the USWG session were asked to briefly describe the meeting in their own words. We also had round table discussions on how to connect to the Internet. ISOC has tasked the USWG to look upon ISOC Secretariat training programs. We also discussed DNS cookbooks and other types of topics the User Services Area could do in the future.

CURRENT MEETING REPORT

Reported by Peter Deutsch/McGill and Alan Emtage/McGill

Minutes of the Living Documents BOF (LIVDOC)

The Preliminary Agenda called for discussion on a wide range of topics related to the creation and implementation of Living Documents but in practice the majority of the discussion revolved around data representation issues for network-based information discovery and delivery systems.

Much of the discussion centered upon the characteristics needed to implement a practical scheme for Universal Document Identifiers, contrasting these with a proposal for Unique Document Serial Numbers. UDIs have been proposed to allow multiple information systems to communicate location and access information. Initial proposals that had been circulated by Tim Berners-Lee, Brewster Kahle and others were discussed and these were compared to the information needed and currently provided by such systems as Prospero, W3, WAIS and others. No firm conclusions were reached, but it was agreed that a mailing list (nir@cc.mcgill.ca) would be created to pursue this issue with a goal of discussions about a possible standardization of UDIs for Internet use and, by extension, the purpose for which this BOF was initially called: Living Documents. Initially, all attendees of this BOF are to be placed on the list, and existence of the list is to be announced to the Internet community.

Discussion concerning Unique Document Serial Numbers centered around the perceived need to identify and compare the _contents_ (in contrast to the location) of documents in an internet environment. Ideally, we would have a means for:

- Identifying the contents of a document and comparing it with other documents without copying and comparing them directly.
- Identifying derivative works and ancestral links between documents.
- Identifying documents that contain the same information despite representational changes that do not add or delete information contents.

It was generally accepted that the first of these could probably be met with relatively straightforward signature schemes, but that the last two would be difficult or impossible using strictly syntactic means.

The discussion continued across a range of topics, examining the other issues to be addressed in implementing Living Documents and network-based information systems. The following list was drawn up outlining some of the issues to be addressed in subsequent work:

- Univeral Document Identifiers
  - Design, documentation and deployment. Issues involved include the need to encode individual access methods and specific location information within a

specified access method. An initial proposal for such a scheme had been circulated by Tim Berners-Lee prior to the meeting.

- Unique Document Serial Numbers
  - Design, documentation and deployment. Issues involved include identifying specific documents, version control and derivation information.

- Cataloguing Information
  - Librarians already make use of far more cataloguing information than any of the experimental systems currently in use on the Internet. Work with those directly involved in library science working with extending MARC records, ISBN and ISSN numbers is called for.

- Discovery Mechanisms
  - There remains a large open problem in rapidly and efficiently discovering the existence and location of information in a large distributed computing environment. The proposed UDIs and UDSNs may enable such systems to be built but additional work is still needed. There are problems both in locating individual service providers and specific pieces of information.

- Authentication and Access Control
  - Security issues were not discussed in depth, but it was agreed that such issues would become more important as large-scale systems are developed and deployed.

- Editorial Control
  - Again a topic touched upon only briefly, it was suggested by one participant that true Living Document systems would have to include some method of imposing editorial control.

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Mark Baushke	mdb@nsd.3com.com
Tim Berners-Lee	timbl@info.cern.ch
Susan Calcari	calcaris@cerf.net
Andrew Cherenson	arc@sgi.com
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Clement	clement@educom.edu
David Crocker	dcrocker@mordor.stanford.edu
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@ans.net
Elizabeth Feinler	jake@nsipo.arc.nasa.gov

Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov
Anders Gillner	awg@sunet.se
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	esh@merit.edu
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Darren Kinley	kinley@crim.ca
Walter Lazear	lazear@gateway.mitre.org
Kurt Lidl	lidl@sura.net
Clifford Lynch	calur@uccmvsa.ucop.edu
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Larry Masinter	masinter@parc.xerox.com
Daniel Matzke	matzked@cerf.net
Mark McCahill	mpm@boombox.micro.umn.edu
Mark Needleman	mhn@stubbs.ucop.edu
Clifford Neuman	bcn@isi.edu
Torben Nielsen	torben@hawaii.edu
Geir Pedersen	geir.pedersen@usit.uio.no
Cecilia Preston	cpreston@info.berkeley.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Anthony Rutkowski	amr@sprint.com
Mark Smith	mcs@umich.edu
Karen Sollins	sollins@lcs.mit.edu
Craig Summerhill	craig@cni.org
Kevin Walsh	kw@ucsd.edu
Moira West	mjw@cert.sei.cmu.edu
Scott Williamson	scottw@nic.ddn.mil
Wengyik Yeong	yeongw@psi.com

3.9.1 Directory Information Services Infrastructure (disi)

Charter

Chair(s):

Chris Weider, clw@merit.edu

Mailing Lists:

General Discussion: disi@merit.edu
To Subscribe: disi-request@merit.edu
Archive: pub/disi-archive@merit.edu

Description of Working Group:

The Directory Information Services (pilot) Infrastructure Working Group (DISI) is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

The DISI Working Group is an offshoot of the OSI Directory Services Group, and, accordingly, is a combined effort of the OSI Integration Area and User Services Area of the IETF. The current OSIDS Working Group was chartered to smooth out technical differences in information storage schema and difficulties in the interoperability and coherence of various X.500 implementations. The DISI Group is concerned solely with expanding the Directory Services infrastructure. As DISI will be providing infrastructure with an eye towards truly operational status, DISI will need to form liaisons with COSINE, Paradyse, and perhaps the RARE WG3.

As a final document, the DISI Working Group shall write a Charter for a new working group concerned with user services, integration, maintenance, and operations of Directory Services, the Internet Directory User Services Group.

Goals and Milestones:

Done First IETF Meeting: review and approve the Charter making any changes necessary. Examine needs and resources for the documentation to be produced, using as a first draft a document produced by Chris Weider, Merit, which will be brought to the IETF. Assign writing assignments. Further work will be done electronically.

Jul 1991 Second IETF Meeting: review and approve documentation; review and approve Charter for the IDUS Group.

Aug 1991 Electronically review final draft of documentation, and, if acceptable, submit to IESG for publication.

Dec 1991 Third IETF Meeting: Declare success and reform DISI Group as IDUS group.

Internet Drafts:

"Interim Schema for Network Infrastructure Information in X.500 New name: Encoding Network Addresses to support operation ov", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-disi-netinfrax500-00.txt>

"An Executive Introduction to Directory Services Using the X.500 Protocol", 12/18/1991, Chris Weider, Joyce Reynolds, Sergio Heker <draft-ietf-disi-execdir-01.txt>

Request For Comments:

RFC 1292 "A Catalog of Available X.500 Implementations"

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Directory Information Services Infrastructure Working Group (DISI)

Agenda:

1. Review Old Minutes
2. Review RFC Progress
3. Review and Discuss New Documents

The meeting began with a review and discussion of the old Minutes. No modifications were deemed necessary so the Group moved onto a discussion of RFC Progress. All the documents currently in the Working Group have come out as RFCs. Ruth Lang and Russ Wright's "A catalog of Available X.500 Implementations" is out as FYI 11, RFC 1292. The Executive Overview was split into two documents, and is out as:

```
FYI 13, RFC 1308: Executive Introduction to Directory Services
Using the X.500 Protocol, by Chris Weider and Joyce Reynolds
```

and

```
FYI 14, RFC 1309: Technical Overview of Directory Services
Using the X.500 Protocol, by Chris Weider, Joyce Reynolds,
and Sergio Heker.
```

Five new documents were discussed:

1. A Pilot Project Catalog.
2. A DSA Setup Guide.
3. An Advanced Usages Document.
4. A document on how to get yourself or your company registered in the X.500 Directory.
5. A document on the naming philosophy of X.500, which would answer questions such as: Where do I really belong in the Directory, and why would I want a name?

Two of these were assigned, with the rest to wait for assignment until Boston: Document three will be written by Chris Weider, Russ Wright, and Jake Feinler. Document four will be written by April Marine, Tim Howes, Mark Smith, and Pat Smith.

Attendees

Paul Barker	p.barker@cs.ucl.ac.uk
William Biagi	bbiagi@cos.com
Susan Calcari	calcaris@cerf.net
Robert Cooney	cooney@wnyose.nctsw.navy.mil
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppen@verw.switch.ch
Susan Estrada	estradas@cerf.net
Carol Farnham	cfarnham@unl.edu
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Sallie Fellows	sallie%ed@psc.plymouth.edu
Jisoo Geiter	geiter@gateway.mitre.org
Jeff Gumpf	jag@po.cwru.edu
Steve Hardcastle-Kille	s.kille@cs.ucl.ac.uk
Tim Howes	Tim.Howes@umich.edu.
Barbara Jennings	bjjenni@sandia.gov
Darren Kinley	kinley@crim.ca
Sylvain Langlois	sylvain@cli53an.edf.fr
Arthur Lin	yalin@srv.pacbell.com
Triet Lu	trietl@sparta.com
April Marine	april@nisc.sri.com
Daniel Matzke	matzked@cerf.net
Cyndi Mills	cmills@bbn.com
Rakesh Patel	rpatel@rutgers.edu
Dan Pederson	pederson@venera.isi.edu
Jon Postel	postel@isi.edu
Joyce K. Reynolds	jkrey@isi.edu
Ursula Sinkewicz	sinkewic@netrix.nac.dec.com
Patricia Smith	psmith@merit.edu
Mark Smith	mcs@umich.edu
Chris Weider	weider@ans.net
Scott Williamson	scottw@nic.ddn.mil
Russ Wright	wright@lbl.gov
Wengyik Yeong	yeongw@psi.com

3.9.2 Internet Anonymous FTP Archives (iafa)

Charter

Chair(s):

Peter Deutsch, peterd@cc.mcgill.ca
Alan Emtage, bajan@cc.mcgill.ca

Mailing Lists:

General Discussion: iafa@cc.mcgill.ca
To Subscribe: iafa-request@cc.mcgill.ca
Archive: pub/iafa-archive@archive.cc.mcgill.ca

Description of Working Group:

The Internet Anonymous FTP Archives Working Group is chartered to define a set of recommended standard procedures for the access and administration of anonymous ftp archive sites on the Internet. Such a set of procedures will provide a framework for:

(a) allowing the inexperienced Internet user the ability to more easily navigate the hundreds of publically accessible archive sites; and,

(b) allowing users and network-based tools to retrieve specific site information such as access policies, contact information, possible areas of information specialization, archived package descriptions, etc., in a standardized manner.

Particular emphasis will be placed on the possible impact of these procedures on the FTP site administrators.

Attention will be paid to the impact of newer archive indexing and access tools on the operation of such archive sites. A set of suggestions will be offered to allow archive site administrators to better integrate their offerings with such tools as they are developed.

The security of the anonymous FTP site configuration will also be considered to be an integral part of this document. It is expected that remote management of the archives will be adequately handled by existing network management procedures.

Goals and Milestones:

Nov 1991 First IETF Meeting: review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended procedures and impact on site administrators. Assign writing assignments for the first draft of the documents.

Mar 1991 Review first draft and determine necessary revisions. Follow up discussion will occur on mailing list.

Jun 1991 Make document an Internet Draft. Continue revisions based on comments at IETF and on the mailing list.

Nov 1992 Fourth IETF meeting. Review final drafts and if OK, give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Alan Emtage/McGill

Minutes of the Internet Anonymous FTP Archives Working Group (IAFA)

The Minutes of the previous meeting in Santa Fe were reviewed and accepted and the Agenda for the current meeting was approved. The Group then considered the question of whether a User Document for Anonymous FTP was still needed, given the publication of several new user guides in recent months. The general consensus was that it was since there is no current Internet RFC (FYI) covering the topic. April Marine, Ellen Hoffman and John Curran have agreed to continue working on this document. Finally, a detailed discussion of the Site Administrator's Document draft was held and its general ideas and focus approved. The following points were made:

- The additional information for the archive suggested in the documentation should not reside in fixed locations. While recognizing that this places additional burdens on the information tool designers, it was thought that having fixed locations for this information would reduce its usability. Files with special names should be able to be deposited at any place in the archive and be retrieved for their information.
- Some mechanism for a Time To Live (TTL) should be associated with the additional information being asked for in the archives to help control the problem of outdated information.
- Investigations should be made into Universal Document Serial Numbers and Universal Document Identifiers and contact maintained with those working on these in the IRTF and other IETF areas.
- A section should be included describing and summarizing existing projects which would be of use to site administrators, however care must be taken not to endorse any specific project as "required".
- Contact should be made and maintained with those doing work in this area: OSI (X.500) and CNI projects. The OSI document "Representing Public Archives in the Directory" was mentioned as one source of information for IAFA. It was recognized that the Library community is also very much involved in similar work.
- Facilities should be included for the case where multiple logical archives are being maintained on one physical archive.
- Possibility for a "Quality of Service" index to be included in the additional information. Problems with this were also noted since we do not want to be seen to be "endorsing" specific archives, but an approach that rates archives as "hobby", "official regional or midlevel archive", "university campus-wide archive" etc. may be of use.
- A cross reference mechanism should be included for users to locate other related information.

- Security concerns about the current draft were raised by CERT representatives. It was agreed that the authors should work with CERT in resolving these concerns.

It was agreed that the authors would work to incorporate these changes and submit a follow-up draft to the IAFA mailing list as soon as possible. No new business was presented.

Attendees

Farhad Anklesaria	fxa@boombox.micro.umn.edu
Mark Baushke	mdb@nsd.3com.com
Tim Berners-Lee	timbl@info.cern.ch
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Carol Farnham	cfarnham@unl.edu
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Sallie Fellows	sallie%ed@psc.plymouth.edu
Jill Foster	jill.foster@newcastle.ac.uk
Anders Gillner	awg@sunet.se
Walter Guilarte	70026.1715@compuserve.com
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	esh@merit.edu
Darren Kinley	kinley@crim.ca
Kurt Lidl	lidl@sura.net
Patrice Lyons	3432266@mcimail.com
Gary Malkin	gmalkin@xylogics.com
Larry Masinter	masinter@parc.xerox.com
Daniel Matzke	matzked@cerf.net
Mark McCahill	mpm@boombox.micro.umn.edu
Chris Myers	chris@wugate.wustl.edu
Clifford Neuman	bcn@isi.edu
Brad Passwaters	bjp@sura.net
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
L. Michael Sabo	lms@denver.ssds.com
Michael Schwartz	schwartz@cs.colorado.edu
Patricia Smith	psmith@merit.edu
Timothy Walden	tmwalden@saturn.acc.com
Moira West	mjw@cert.sei.cmu.edu
Scott Williamson	scottw@nic.ddn.mil

3.9.3 Internet School Networking (isn)

Charter

Chair(s):

John Clement, clement@educom.edu
Arthur St. George, stgeorge@bootes.unm.edu
Connie Stout, cstout@tenet.edu

Mailing Lists:

General Discussion: cosndisc@bitnic.educom.edu
To Subscribe: listserv@bitnic.educom.edu (Sub: cosndisc)
Archive: listserv@bitnic.educom.edu

Description of Working Group:

The Internet School Networking Working Group is chartered to facilitate the connection of the United States' K-12 (Kindergarten-12th Grade) schools, public and private, to the Internet, and school networking in general.

It is critically important that national networking for K-12 education proceed along established lines of protocol, using existing network structures. The Working Group's first priority will be to establish guidelines for specialized user interfaces. K-12 networking will also require other support services, such as directories, online and hotline help, specialized training programs and collaborative projects with instructional and curriculum groups, disciplinary groups and postsecondary institutions.

While the initial focus is school networking in the U.S., the Working Group will coordinate its efforts with similar activities in other countries and regions of the world.

Goals and Milestones:

Nov 1991 Meet for the first time at IETF and establish approval of Charter. Examine the status of projects in process when Working Group was created. Begin work on list of deliverables.

Jan 1992 Release X.500 "K-12 People Directory" version in collaboration with Merit. Develop plans and milestones for K-12 Resources Directory.

Mar 1992 First draft of information packet document for computing directors to assist them in connecting K-12 schools. First draft of user interface guideline statement.

May 1992 Release X.500 K-12 Resource Directory version in collaboration with Merit. Present final draft guideline statement.

CURRENT MEETING REPORT

Reported by John Clement/EDUCOM

Minutes of the Internet School Networking Working Group (ISN)

Specific discussion of the major Agenda topics is summarized below in sections 1) and 2). A number of general issues were addressed during the meeting, and are summarized briefly first.

Joyce Reynolds mentioned that there is considerable interest in international circles, especially in Europe, in what this Working Group is doing, and that products from the Working Group will be scrutinized attentively overseas.

Among the more general issues (exceeding technical boundaries) raised at least briefly during the session were: the need to provide persuasive arguments for educators to use with school administrators, to support both initial access and expansion of networking capabilities; the continuing need for postsecondary institutions to remain involved in helping schools with connectivity, whether as sources of guest accounts, for technical expertise and support where other mechanisms are not yet in place, or as sources of collaboration on connectivity or content problems.

The issue of having the ISN Working Group contribute to providing technical consulting capabilities to K-12 groups and institutions was discussed briefly. Although no definite resolution was offered, after the meeting it was suggested that a roster of persons willing to respond to technical questions might be prepared and offered as part of the FAQ resource (cf. item 2 below), with the possibility that these individuals might provide further technical consultation. One condition of listing in the roster would be that good questions to roster members would be added by them to the FAQ list.

Denis Newman offered that there are many more (both local and wide-area) networks in place in schools for administrative uses than there are instructional- use networks. He stipulated that, if administrative networks can be linked into the Internet – as FIRN is in Florida – overall connectivity can be substantially advanced. Denis suggested that the Working Group examine the issues involved in extending administrative LANs and WANs for instructional use; in particular, the Group might consider addressing security questions, and look at the reality of risks entailed in carrying administrative and instructional traffic over the same networks, especially ways to minimize the risk of unauthorized access.

Review of connectivity alternatives and growth paths. Mailing list for this discussion: connect@unmvm.unm.edu

A draft document prepared by Pat Burns and Ed Zachmann of Colorado State University (available on connect archives in Postscript format, or contact Art St.George for fax) was distributed to attendees, and served as a basis for extensive discussion during the meeting.

John Clement presented a networking growth path from the viewpoint of educator practices. Ensuing discussion pointed out that the Burns and Zachmann paper lacked an initial connection model, what might be called Level 0 (the present paper begins with Level 1).

Burns and Zachmann's Level 1 connectivity models (cf. especially Figures 3a and 3b) led to a consensus that both figures and their explanatory text needed relatively minor but significant revisions.

Extensive discussion of the need for continued connection capability for the existing base of computers in schools (often machines of reduced capacity) led to an expressed consensus that such capability should be maintained in the near future (under Level 0 connection options), but should be eliminated over time as more advanced connection levels predominate. The minimum hardware capability for full Internet access to be supported would be Apple Macintosh machines or MS- DOS machines running Windows.

Joe Blackmon summarized his experiences in putting together full (56kbps) Internet connections for finalist schools in the SuperQuest competition at lowest possible costs, and offered to share a document he is preparing for SuperQuest on the subject.

Discussion of Levels 2-4 (Figures 4-6 of Burns and Zachmann paper) connection models was relatively brief. There seemed to be general agreement that, on first reaction (most attendees were seeing the paper for the first time at the session), the models were acceptable with relatively minor corrections, although considerable extension and clarification were requested (see below). One point made was that the models as presented in the paper were overspecified with regard to their technical components, for the level of generality needed for a paper on connectivity alternatives. At the very least, mention should be made of other technical alternatives.

[This was in no way presented as a criticism of the Burns and Zachmann paper, since they offered a starting point and it was appropriate to list a complete technical implementation so as to estimate costs. But perhaps it could be presented as one example.]

Among the points raised that implied expanding the descriptions of the models:

- Guidance on the boundaries of application of each model: when each becomes inappropriate in given school contexts. For example, model 1 might apply to 3-8 machines in a school, and would not work for more than eight machines.
- Specification of what kinds of Internet capabilities are available for each model: i.e., model 1 offers telnet (remote login) via an interface menu item, but no direct telnet capability, and would allow access to WAIS via Simple WAIS but not the installation of WAIS client software, etc.
- More complete specification of what is required at the Internet node/connection point for each model;
- Cost estimates for moving from one model to the next, and stipulation of what equipment would be no longer needed and could be used elsewhere;

- Specification of personnel, time and training requirements for installation, support (including time for administration and security protection), and maintenance.

Jeff Hayward agreed to lead an effort to respecify the models on the basis of the Burns and Zachmann paper and its discussion, and to amplify their written description, and to prepare a draft before the end of

It was agreed that a revised draft document would be made available and that it would be revised once more from feedback on the list. The document is available on <connect@unmvm.unm.edu>. A twice-iterated draft might be available in advance of the Twenty- Fourth IETF, in Boston in July.

Development of a FAQ ("frequently-asked-questions") archive on school connectivity issues. A new mailing list will be formed for this issue.

Discussion initially focused on identifying sources of information and materials from which a set of FAQs and candidate answers might be developed. KIDSNET and EDTECH lists were mentioned, and there are many others. Gene Hastings, Peter Deutsch and Rob Reilly offered to help round up relevant materials. Art St. George and John Clement offered to assemble these materials into a preliminary archive. Volunteers were solicited to edit the archive and review/revise answers to the questions. The following were dragooned into volunteering: Tony Rutkowski, Joe Blackmon, Tracy LaQuey-Parker, and Michael Marcinkevicz. John Clement will also ask Al Rogers of FrEdMail to volunteer or suggest one of the FrEdMail sysops to volunteer.

By the next IETF meeting, it was agreed that a preliminary FAQ archive would have received an editing pass and would be available for examination and revision.

The following preliminary set of categories is offered for consideration for the FAQ list:

- Why should K-12 educators and schools connect to the Internet?
- Modes of access
- Costs
- Support issues, including maintenance, access rules, security.
- Sources of information
- Troubleshooting
- Glossary of essential terms
- Technical specifications
- Technical consultants roster: people willing to respond to questions, or possibly provide further consultation.

Attendees

Joe Blackmon	blackmon@ncsa.uiuc.edu

Susan Calcari	calcaris@cerf.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Clement	clement@educom.edu
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Carol Farnham	cfarnham@unl.edu
Sallie Fellows	sallie%ed@psc.plymouth.edu
Constance Fleenor Lloyd	cfleenor@lloyd.com
Maria Gallagher	maria@nsipo.nasa.gov
Joan Gargano	jcgargano@ucdavis.edu
Robert Gutierrez	gutierre@nsipo.nasa.gov
Eugene Hastings	hastings@psc.edu
Alisa Hata	hata@cac.washington.edu
Jeff Hayward	j.hayward@utexas.edu
Russ Hobby	rdhobby@ucdavis.edu
James Keller	/PN=JAMES.H.KELLER/O=US.SPRINT/ADMD=TELE MAIL/C=US/@sprint.com
Mark Lewis	mlewis@telebit.com
Kurt Lidl	lidl@sura.net
Hock-Koon Lim	lim@po.cwru.edu
Michael Marcinkevicz	mdm@csu.net
April Marine	april@nisc.sri.com
Daniel Matzke	matzked@cerf.net
Mark Needleman	mhn@stubbs.ucop.edu
Tracy LaQuey Parker	tracy@utexas.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Anthony Rutkowski	amr@sprint.com
Patricia Smith	psmith@merit.edu
Arthur St. George	stgeorge@bootes.unm.edu
Sandra Vest	vests@imo-uvax.dca.mil
Moira West	mjw@cert.sei.cmu.edu
Michael Wrona	mwrona@eis.calstate.edu

3.9.4 Internet User Glossary (userglos)

Charter

Chair(s):
Tracy LaQuey Parker, tracy@utexas.edu
Gary Malkin, gmalkin@ftp.com

Mailing Lists:
General Discussion: usergloss@ftp.com
To Subscribe: usergloss-request@ftp.com
Archive:

Description of Working Group:

The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the Charter as necessary. Discussion of Userglos Working Group Chair nominations submitted by USWGers.
TBD	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
TBD	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.

CURRENT MEETING REPORT

Reported by Gary Malkin/FTP Software

Minutes of the Internet User Glossary Working Group (USERGLOS)

Editor's Note (md): These Minutes were inadvertantly omitted from the Santa Fe Proceedings. We wish to extend our apologies to the Chair and to the members of the IETF for the oversight.

Agenda

- Review of the previous meeting
- Setting milestones for Glossary completion
- Discussion of glossary format
- Determination of Editorial Board
- Worker Bee assignments

Since Usergloss did not meet in Atlanta, this was the second meeting of the Working Group. The objectives of the Working Group remain unchanged since it's inception, the creation of a glossary of networking terms and acronyms for the Internet community.

The primary sources for the glossary are existing glossaries. They are: Carl Malamud's online, all inclusive glossary; the glossary from FYI 4 (RFC 1206); Ole Jacobson's Interop glossary; and NCAR's online glossary. Additional sources may be NASA's online glossary and the Internet Tour glossary.

The introduction will include a section which describes how to get more information. It will mention the FYI series and the Official Standards RFC.

The entries will be organized in alphabetical order (case insensitive). Acronyms will be listed independently of the descriptive text. For example:

```
IANA - see Internet Assigned Numbers Authority
...
Internet Assigned Numbers Authority (IANA)
<descriptive text>
```

The document will be available in ASCII format. In addition, it may be made available in PostScript, Hypertext or Websters formats. Additional format availability is dependent upon User Services Working Group resources.

Some index categories have been tentatively identified. They are: applications, protocols, Internet organizations, networks (e.g., mid-level networks), and address types (e.g., internet address). It is possible that a glossary entry may appear in multiple categories.

The editorial board has the following members: Tracy LaQuey Parker, Gary Malkin, Carl Malamud, Karen Roubicek, Peter Liebscher, April Marine, and Laura Bollettino.

Attendees

Martyne Hallgren	martyne@nr-tech.cit.cornell.edu
Ellen Hoffman	esh@merit.edu
Edward Krol	e-krol@uiuc.edu
Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Michael Patton	map@lcs.mit.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@faxon.com

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics and Tracy LaQuey Parker/UTexas

Minutes of the Internet User Glossary Working Group (USERGLOS)

Having compiled a very large glossary from many sources, we removed those entries which we decided were either too general or too esoteric for the intended audience of the glossary.

The Group will do a final review of the glossary at the Boston IETF. The object will be to ensure that the entries are catagorized properly in the index. The glossary will then be submitted to the RFC editor within a one month period following the Boston IETF.

Agenda

1. Review Working Group Description, Goals and Milestones
2. Progress to Date
   (a) Collected a list of glossaries to use as a base
   (b) Editorial committee has edited and merged these glossaries
3. Review the document we now have
4. Weeding-out process (which definitions should be excluded?)
5. Refining the indexes/categories
6. Final format of the document
   (a) Review format proposed at last meeting
   (b) What to do about the indexes
7. Assign New Tasks
   (a) Editing draft
   (b) Writing/researching better definitions
   (c) Searching for definitions we've left out.

We started with some discussion about copyright issues. We believe that definitions are not copyrighted (except by the creater of a word). We will be giving credit to the sources of the definitions.

We spent considerable time determining who our audience was so that we could decide which words were too generic for an Internet specific glossary and which words were too esoteric for a general audience. The Group decided that the glossary was not for extreme beginners and that an effort to create a comprehensive dictionary should be undertaken in the future.

The remainder of the first session, and all of the second session, was devoted to pruning the glossary. We removed 351 entries from the original 1247 and added 10. We still need to resolve the duplicates (by merging them together).

Agenda items 5-7 will be done over email before the next IETF meeting in Boston.

Attendees

Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Clement	clement@educom.edu
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Carol Farnham	cfarnham@unl.edu
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov
Jeff Hayward	j.hayward@utexas.edu
Ellen Hoffman	esh@merit.edu
Ole Jacobsen	ole@csli.stanford.edu
Kurt Lidl	lidl@sura.net
Gary Malkin	gmalkin@xylogics.com
April Marine	april@nisc.sri.com
Daniel Matzke	matzked@cerf.net
Tracy LaQuey Parker	tracy@utexas.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Patricia Smith	psmith@merit.edu
Scott Williamson	scottw@nic.ddn.mil

3.9.5 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@ans.net
Darren Kinley, kinley@crim.ca

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu
Archive:

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.
- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.
- Identify additional new or improved tools that have become apparent since the last compilation of the reference document.
- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Done	Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.
Aug 1991	Draft of catalog will be prepared, draft to be reviewed and modified. Initiate IETF Internet Draft review process by submission of a "Son of NOCTools" catalog draft to IESG Secretary.
Dec 1991	Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI RFC for publication.

CURRENT MEETING REPORT

Reported by Darren Kinley/CRIM

Minutes of the NOC-Tool Catalogue Revisions Working Group (NOCTOOL2)

Agenda

The NOCtool2 Working Group met briefly at the 23rd IETF in San Diego. This meeting served primarily to report on the current status of the document revision process. After said report, new and interesting tools were discussed for inclusion in the revised NOCtools Catalog.

Current Status

Entries continue to arrive at a slow rate. Other commitments have delayed Enger and Kinley in generating the revised document. An Internet Draft will be made available by the end of May 1992.

New Additions to the Catalog

A list of new tools for inclusion in the catalog was compiled.

Action Items

Robert Enger, Darren Kinley: Draft a new document comprising the received entries, follow up on list of new tools, solicit new entries.

Darren Kinley, Joyce Reynolds: Compile list where NOCtools catalog availability announcement will get the widest possible distribution.

Robert Enger, Darren Kinley, Chris Myers: Define Usenet newsgroups to be created and write charters for these groups as required.

Chris Myers: Make a home for anonymous FTP and automated electronic submissions and distribution tool.

Attendees

Larry Blunk	ljb@merit.edu
Douglas Carson	carson@utcs.utoronto.ca
Robert Enger	enger@ans.net
Tom Hemp	themp@ralvm11.vnet.ibm.com
Darren Kinley	kinley@crim.ca
Linda Leibengood	ldl@ans.net

Hock-Koon Lim	lim@po.cwru.edu
Henry Miller	henrym@sacusr.mp.usbr.gov
Chris Myers	chris@wugate.wustl.edu
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Joyce K. Reynolds	jkrey@isi.edu
L. Michael Sabo	lms@denver.ssds.com

3.9.6 Network Information Services Infrastructure (nisi)

Charter

Chair(s):
April Marine, april@nisc.sri.com
Pat Smith, psmith@merit.edu

Mailing Lists:
General Discussion: nisi@merit.edu
To Subscribe: nisi-request@merit.edu
Archive:

Description of Working Group:

The NISI Working Group will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. The work will begin with existing NIC functions and services and should build upon work already being done within the Internet community. It should address areas such as common information formats, methods of access, user interface, and issues relating to security and privacy of Internet databases.

Goals and Milestones:

Done	Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of 'cooperative agreements' which will allow NICs to work together more effectively to serve users.
Done	Complete draft for phase 2 suggesting cooperative agreements for NICs.
Done	Revised draft document ready for Working Group review. Document defines NIC functions and suggests some standardizations for NIC services, as well as offers new mechanisms for exchanging information between NICs.
Done	Document submitted as Internet Draft for comment from a wider internet audience.
Done	Working Group discussed current Internet draft and suggested minor revisions. Decision made to continue Working Group activity beyond this document.
Nov 1991	First document released as informational RFC. Outline and discuss new NISI tasks at IETF meeting.
Jul 1992	Write a document explaining the security issues of privacy and accuracy in Internet databases. Publish as an informational RFC.

Internet Drafts:

"Building a Network Information Services Infrastructure", 07/15/1991, D. Sitzler, P. Smith, A. Marine <draft-ietf-nisi-infrastructure-00.txt>

CURRENT MEETING REPORT

Reported by April Marine/SRI

Minutes of the Network Information Services Infrastructure Working Group (NISI)

Agenda

- Announcements/Status Reports
- Information Discovery
- Next Project

Announcements/Status Reports

The "Building a Network Information Services Infrastructure" document was issued as an FYI 12 in February.

The document on "Privacy and Integrity Issues in NIC Databases" is now in draft form and will be reviewed by the volunteer authors. (Basically, all the authors had seen the outline, but not the draft at the time of the meeting.) Once they are satisfied, it will be sent to the whole NISI list for review and comments.

Information Discovery

This was left from the last meeting. There was great interest in the question "Once we have information, how do we let users know it's there?" Particularly since the last meeting, a lot of other groups have also become interested in this question. We decided that this was an implementation question beyond the scope of the NISI Group. We would stay informed of the work of other groups and provide input to them, from a user point of view. (After NISI met, April attended the WAIS and Directory Integration (WAIS) BOF, which turned out to include many other "info servers.") That Group decided to form an official Working Group to work on matters related to allowing existing information services to interact. This is something worth following.

Next NISI Project

- Brainstorming

  The Group tossed around a few ideas for future projects, which are listed here. As this was a brainstorming session, detailed descriptions of each idea are not available.

  - Service Discovery (finding out what services are available on the Internet).
  - Procedures for NICs to deposit information into the various information services.

- A nethelp utility that would give users fast, uniform information on what the Internet is and how to use it.
- A list of useful services to do.
- A list of services each site offers.
- A mailing list archive.
- A clarification of the core of information about the Internet that all Internet NICs should provide.
- Tools to provide information.

Of these, the most interest was in the lists of services and the Nethelp (Martyne Hallgren must get credit for that coinage) service.

- List of Services Document

  This document will follow the IAFA recommendations and publish a list of the services available from that administrative subdomain (not necessarily just from that one site, but from the various service providers at that administrative entity). The list of services is something like a description of the various services and information on how to access them.

  This document may, in fact, overlap the Nethelp document described next. It may have been another approach to the same information. Any overlap should come to light as work progresses on the documents. For now, this document should be thought of as the description of the services that a nethelp tool could access (i.e., the tool could access these descriptions). The next document describes what the Group wants the tool to look like (generically) and do. Peter Deutsch, Jake Feinler, John Clement, Cyndi Mills, and April volunteered to work on this document. A draft of this document is due by the next IETF meeting in July.

- Nethelp Document

  This document is aimed at specifying the type of information a nethelp service would provide. A nethelp service is envisioned to be something a new user can very easily access and which will orient him toward what the Internet is and where he can get more information about it. Below is some of the information such a service might provide:

  - Network Use Etiquette and Guidelines
  - New User Guide Information
  - Usage Policy
  - Information on the Net
  - Net Access Information
  - Resources and Services
  - Answers to Commonly Asked Questions (which we can glean from NICs)

  - Who to Call for Help

  Other thoughts on this document included whether it should have a standard interface or not (the vote was no, but an easy interface is necessary); it should include help for individual sites regarding what they can add/customize; it should address the fears a new user has about breaking stuff if he tries new things; it could be scalable and include hooks to add information that might benefit more experienced users as well, or include pointers to more detailed, technical information on different subjects. Specifying concrete ideas in this area might add to the momentum of getting funding to implement such a utility or help those with information servers to extend them to this use.

  Another idea regarding the interface was that it could start with standard questions, such as who-is?, what-is?, where-am-I?, how?, where-is?, etc.

  Although the scope of the document was not fully defined in the meeting, it will concentrate on what information the NISI Group would recommend that such a program contain, possibly even including explicit text, what services such a program should provide, and, generally speaking, what the program could look like. It will not necessarily cover implementation strategies, although those can be considered as well.

  Martyne Hallgren, Ellen Hoffman, Scott Williamson, Marsha Perrot, Tim Berners-Lee, Pat, and April will be working on this document. Suggestions online are welcome!

  A draft of this document is due by the next IETF meeting in July.

- Other

  There was another suggestion that came up during the nethelp discussion that is worth noting, although it is probably out of the scope of NISI. It was the suggestion that a controlled usage environment be created where users can test out their knowledge of how to use different Internet resources, such as FTP. It would be a closed environment in that it would not affect Internet use or operations, but would be realistic in that it would simulate the Internet environment. It might have slightly more user friendly error messages and more help/guidance in using the net. Such an environment could be used to allay fears new users have about using the net which cause them to experiment less than they might otherwise if they were confident they couldn't break anything. It would be especially useful for pre-college aged users.

Action Items

- April and Pat to write new goals and milestones for the Working Group.
- Volunteers start working on the documents.
- Working Group to follow new goals and milestones.

Attendees

Andy Adams	ala@merit.edu
Tim Berners-Lee	timbl@info.cern.ch
Susan Calcari	calcaris@cerf.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Clement	clement@educom.edu
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Carol Farnham	cfarnham@unl.edu
Sallie Fellows	sallie%ed@psc.plymouth.edu
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.nasa.gov
Anders Gillner	awg@sunet.se
Jeff Gumpf	jag@po.cwru.edu
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Deborah Hamilton	debbie@qsun!att.com
Keunhee Han	khhan@gaya.kreonet.re.kr
Sang-Chul Han	schan@garam.kreonet.re.kr
Alisa Hata	hata@cac.washington.edu
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
Kurt Lidl	lidl@sura.net
Hock-Koon Lim	lim@po.cwru.edu
April Marine	april@nisc.sri.com
Daniel Matzke	matzked@cerf.net
Cyndi Mills	cmills@bbn.com
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Patricia Smith	psmith@merit.edu
Sandra Vest	vests@imo-uvax.dca.mil
Moira West	mjw@cert.sei.cmu.edu
Scott Williamson	scottw@nic.ddn.mil
Ki-Sung Yoo	ksyu@garam.kreonet.re.kr

3.9.7 User Services (uswg)

Charter

Chair(s):
Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:
General Discussion: us-wg@nnsc.nsf.net
To Subscribe: us-wg-request@nnsc.nsf.net
Archive:

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF working groups created to perform certain projects, or outside organizations such as SIGUCCS.

- Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should:
  - Clearly address user assistance needs;
  - Produce an end-result (e.g., a document, a program plan, etc.);
  - Have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
  - Not duplicate existing or previous efforts.
- Create working groups or other focus groups to carry out projects deemed worthy of pursuing.
- Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

Ongoing This is an oversight group with continuing responsibilities.

Request For Comments:

RFC 1150 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. notes"

RFC 1177 "FYI on Questions and Answers - Answers to Commonly Asked "New Internet User" Questions"

RFC 1206 "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions"

RFC 1207 "Answers to Commonly asked "Experienced Internet User" Questions"

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

Minutes of the User Services Working Group (USWG)

The User-Doc2 Working Group took up about half of the USWG session in their discussions. They met during the USWG session as we had time constraint problems. User-Doc2 is Chaired by Ellen Hoffman and Lenore Jackson. User-Doc's original bibliography was published in August of 1990. It is sorely out of date. This is the new revised Working Group. It has a revised Charter. Items discussed included what items should be included in the outdated document, and what items should be deleted in revising the bibliography.

We also had two reports, one on the RIPE meetings in Amsterdam in January, and the FARNET meeting in February. RIPE was originally chartered as a technical community. They have realized that user services is now something very important they should focus on, especially with East Central Europe and what is going on in Russia right now in user service needs. They have just formed a User Information Services Working Group in the RIPE community. The FARNET meeting focussed on, "Hardening of the Interim NREN". FARNET meeting participants that were at the USWG session were asked to briefly describe the meeting in their own words.

We also had round table discussions on how to connect to the Internet. There are already efforts going on in producing documentation on "How to Connect to the Internet", including ACM SIGUCCS and the SRI NISC.

ISOC has tasked the USWG to look upon ISOC Secretariat training programs. General consensus is that ISOC should poll the regionals, and draw on already existing resources. We also discussed DNS cookbooks and other types of topics the User Services Area could do in the future.

Peter Deutsch volunteered to produce an "Internet Quick and Dirty" document on descriptions of each network service with pointers on where to obtain additional information. Pat Smith, Cyndi Mills, Dan Matzke, Ellen Hoffman, Anders Gilner, April Marine, Jill Foster and Joyce Reynolds volunteered to assist Peter in writing this document. This project will be developed within the USWG.

Included in the USWG announcements were User Services Area FYI RFC publications since the Santa Fe IETF:

- Lang, R., and R. Wright, "A Catalog of Available X.500 Implementations", FYI 11, RFC 1292, January 1992. (A product of the DISI Working Group.)
- Sitzler, D., Smith, P., and A. Marine, "Building a Network Information Services Infrastructure", FYI 12, RFC 1302, February 1992. (A product of the NISI Working Group.)

- Weider, C., and J. Reynolds, "Executive Introduction to Directory Services Using the X.500 Protocol", FYI 13, RFC 1308, March 1992. (A product of the DISI Working Group.)
- Weider, C., Reynolds, J., and S. Heker, "Technical Overview of Directory Services Using the X.500 Protocol", FYI 14, RFC 1309, March 1992. (A product of the DISI Working Group.)

Attendees

Tim Berners-Lee	timbl@info.cern.ch
Susan Calcari	calcaris@cerf.net
Jodi-Ann Chu	jodi@uhunix.uhcc.hawaii.edu
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgill.ca
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@ans.net
Michael Erlinger	mike@lexcel.com
Carol Farnham	cfarnham@unl.edu
Elizabeth Feinler	jake@nsipo.arc.nasa.gov
Sallie Fellows	sallie%ed@psc.plymouth.edu
Jill Foster	jill.foster@newcastle.ac.uk
Anders Gillner	awg@sunet.se
Martyne Hallgren	martyne@mitchell.cit.cornell.edu
Deborah Hamilton	debbie@qsun!att.com
Alf Hansen	Alf.Hansen@delab.sintef.no
Alisa Hata	hata@cac.washington.edu
Ellen Hoffman	esh@merit.edu
Lenore Jackson	jackson@nsinic.gsfc.nasa.gov
Darren Kinley	kinley@crim.ca
Kurt Lidl	lidl@sura.net
April Marine	april@nisc.sri.com
Daniel Matzke	matzked@cerf.net
Cyndi Mills	cmills@bbn.com
Michael Newell	mnewell@nhqvax.hg.nasa.gov
Jack Pope	pope@uscdsv.acusd.edu
Jon Postel	postel@isi.edu
Robert Reilly	rreilly@athena.mit.edu
Joyce K. Reynolds	jkrey@isi.edu
Andreas Schachtner	afs@germany.eu.net
Patricia Smith	psmith@merit.edu
Claudio Topolcic	topolcic@nri.reston.va.us
Moira West	mjw@cert.sei.cmu.edu
Chris Wheeler	cwheeler@cac.washington.edu
Scott Williamson	scottw@nic.ddn.mil

Chapter 4

Network Status Briefings

4.1 CERFnet

Presented by Pushpendra Mohta/CERFnet

CIX ROUTING

PUSHPENDRA MOHTA
pushp@cerf.net

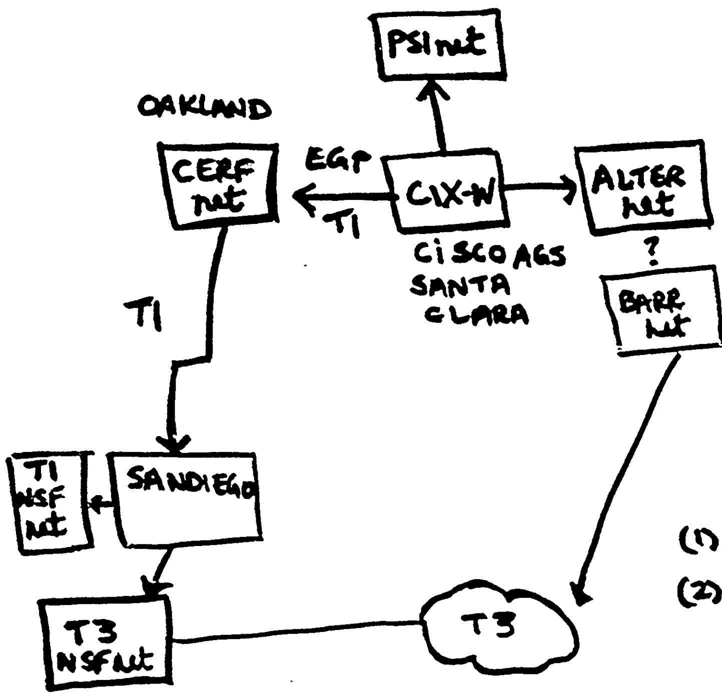

CERFnet
ROUTING PREFS

(1) MEMBER NETS
(2) CSUnet
ESnet
Los Nettos
KREOnet
(3) CIX
(4) T3 NSFnet
(5) T1 NSFnet

CIX.org
info@CIX.org

4.2 EBONE Report

Presented by Bernhard Stockman/NORDUnet

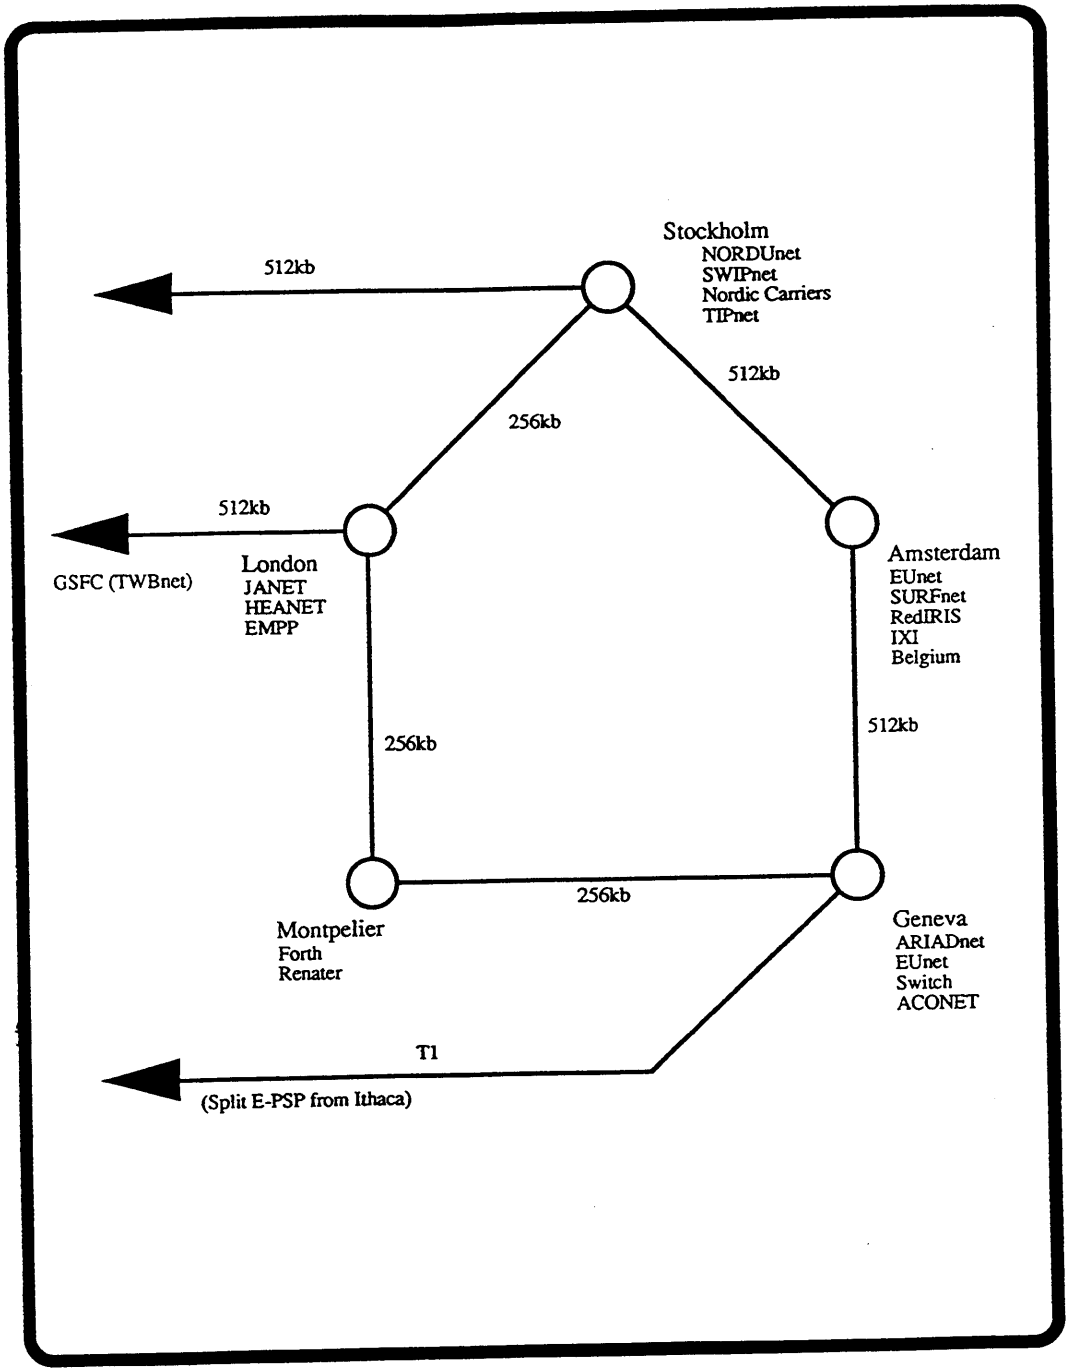
Stockholm
NORDUnet
SWIPnet
Nordic Carriers
TIPnet
512kb
256kb
512kb
512kb
GSFC (TWBnet)
London
JANET
HEANET
EMPP
Amsterdam
EUnet
SURFnet
RedIRIS
IXI
Belgium
256kb
512kb
256kb
Montpelier
Forth
Renater
Geneva
ARIADnet
EUnet
Switch
ACONET
T1
(Split E-PSP from Ithaca)

4.3 ESnet Report

Presented by Tony Hain/LLNL

IETF & ESCC - San Diego

March 1992

Anthony L. Hain

Associate Network Manager
ESnet / NERSC

(510)422-4200 42158::EAGLE::hain hain@es.net

STATUS March 1992

Domestic Connectivity

Foreign Connectivity

Statistics

Other Activities

STATUS March 1992

Domestic Connectivity

Circuit Installations

FIE	56k	IP	>DOE	Sep.'91
SAN-OPS	T1	IP	>LBL	Aug.'91
CHI-OPS	Eth	IP	>ANL	Jan.'92

Circuit Moves

DOE	T1	MP	>FNAL	Jan.'92

Circuit Terminations

CCC	T1	IP	>LLNL	Jan.'92

Network Connections

JVNCnet	T1	IP	>PPPL	Dec.'91

STATUS March 1992

Domestic Connectivity

Planned Circuit Installs

Nevis	T1	MP	>BNL	Mar.'92
Harvard	T1	MP	>MIT	Mar.'92
Bates	T1	MP	>MIT	In Progress
ARM	T1	IP	>ORN	In Progress
LBL	T3	MP	>LLNL	In Progress
EML	56k	MP	>BNL	Pending
Westinghouse	T1	IP	>UTA	Pending
Westinghouse	T1	IP	>CBF	Pending

Planned Circuit Moves

PNL	T1	>LLNL	Mar.'92
SSC	T1	>FNAL	Apr.'92
SSC	T1	>ORNL	Apr.'92
SSC	T1	>UTA	Apr.'92
SSC	T1	>FSU	Apr.'92

STATUS MARCH 1992

INTERNATIONAL CONNECTIVITY

Planned Circuit Installs

KEK	192k	MP`	>Ames	Mar.'92
JAERI	64k	IP	>Ames	Mar.'92
China	64k	MP	>SLAC	Apr.'92
Italy	256k	MP	>?	In Progress

Planned Circuit Upgrades

Germany	256k	IP	>PPPL	In Progress
Brazil	64k	MP	>FNAL	Mar.'92

STATUS MARCH 1992

CURRENT STATISTICS

35 Routers Managed

2471 IP Networks Advertised
- 99 Direct Connect
- 1201 Regional Connect
- 1097 Other Backbones
- 75 International

1.6G packets accepted (Feb.'92)
- ~ 90% IP
- ~ 10% DECnet
- ~ .05% X.25

325G bytes accepted (Feb. '92)

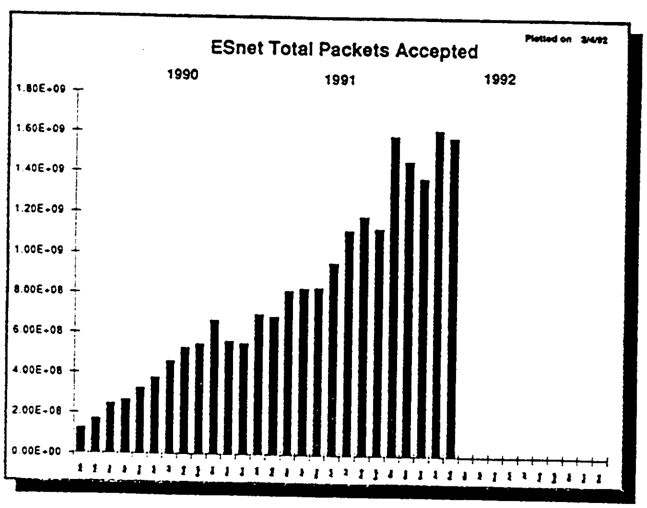

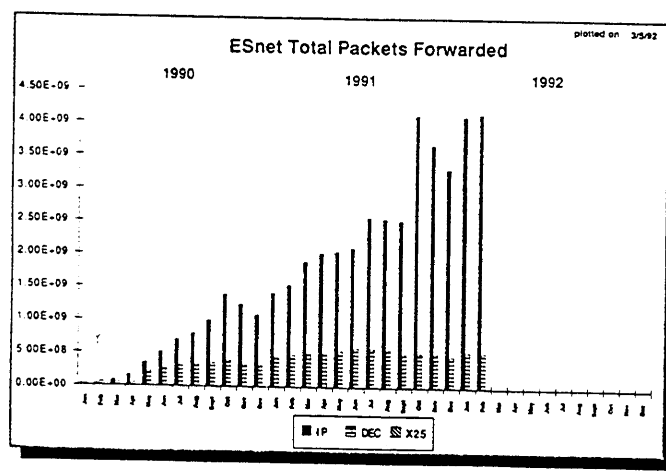

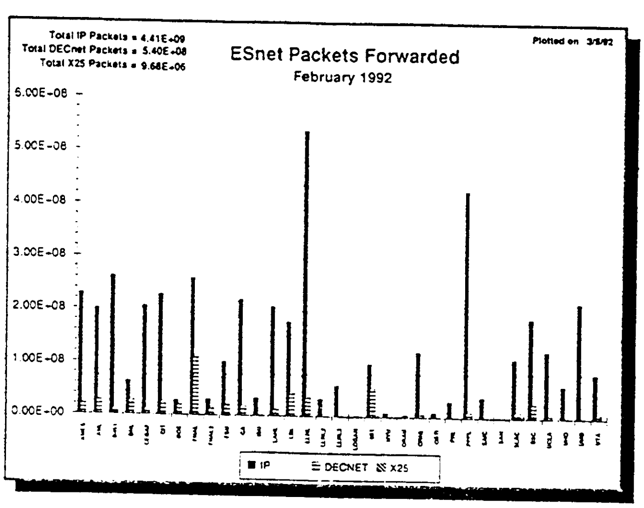

STATUS MARCH 1992

OTHER ACTIVITIES

Frame Relay

3 Site Test-bed in place

- Williams Telecom service
- Test traffic only now
- Friendly Users next

CLNS

Supported in All backbone routers

Site peers; NERSC, LBL, LANL, FNAL, ANL, DOE

Other peers; NSFnet, NSI, CICnet, Suranet

X.400 - MHS/E-mail

2 HP 750 serviers installed

Peers; XNREN, NASA Ames, CDC, Hughes

STATUS MARCH 1992

OTHER ACTIVITIES

X.500

'White Pages' Service operational

ESnet DSA one of several Authoritative DSA's for Country=US

ESnet will master 'WP' info for any backbone site as desired.

Video-confernece support

Request for move to "production level" support

Multiplexor procurement under consideration

DCE Core Services

Definition of DCE and Core to be determined

Vendor market survey in progress

STATUS MARCH 1992

OTHER ACTIVITIES

Publications

Reprints of Network World article ordered

Personnel Changes

Johnathan Brown
Joe Ramus
Greg Schoenberger
Keith Fitzgerald (1/2)
John Renyolds (1/2)
Todd Waybrew (contract 1/2)

Advanced Services RFP (T3 / Cell Relay)

RFP released Feb. '92
Proposals due April 15,'92

4.4 NSI Report

Presented by Jeff Burgan/NASA

NASA SCIENCE INTERNET

STATUS REPORT

MARCH 1992

IETF

SAN DIEGO

WHAT'S NEW

- T1's between ARC, JPL, GSFC MSFC
- Interconnect to T3 NSFnet at SURA
- ENSS 144 installed at Ames Research Center
- Paccom link to Hawaii increased to T1.
- Paccom link to Australia increased to 512 K.

4.5 PACnet

Presented by Torben Nielsen/U Hawaii

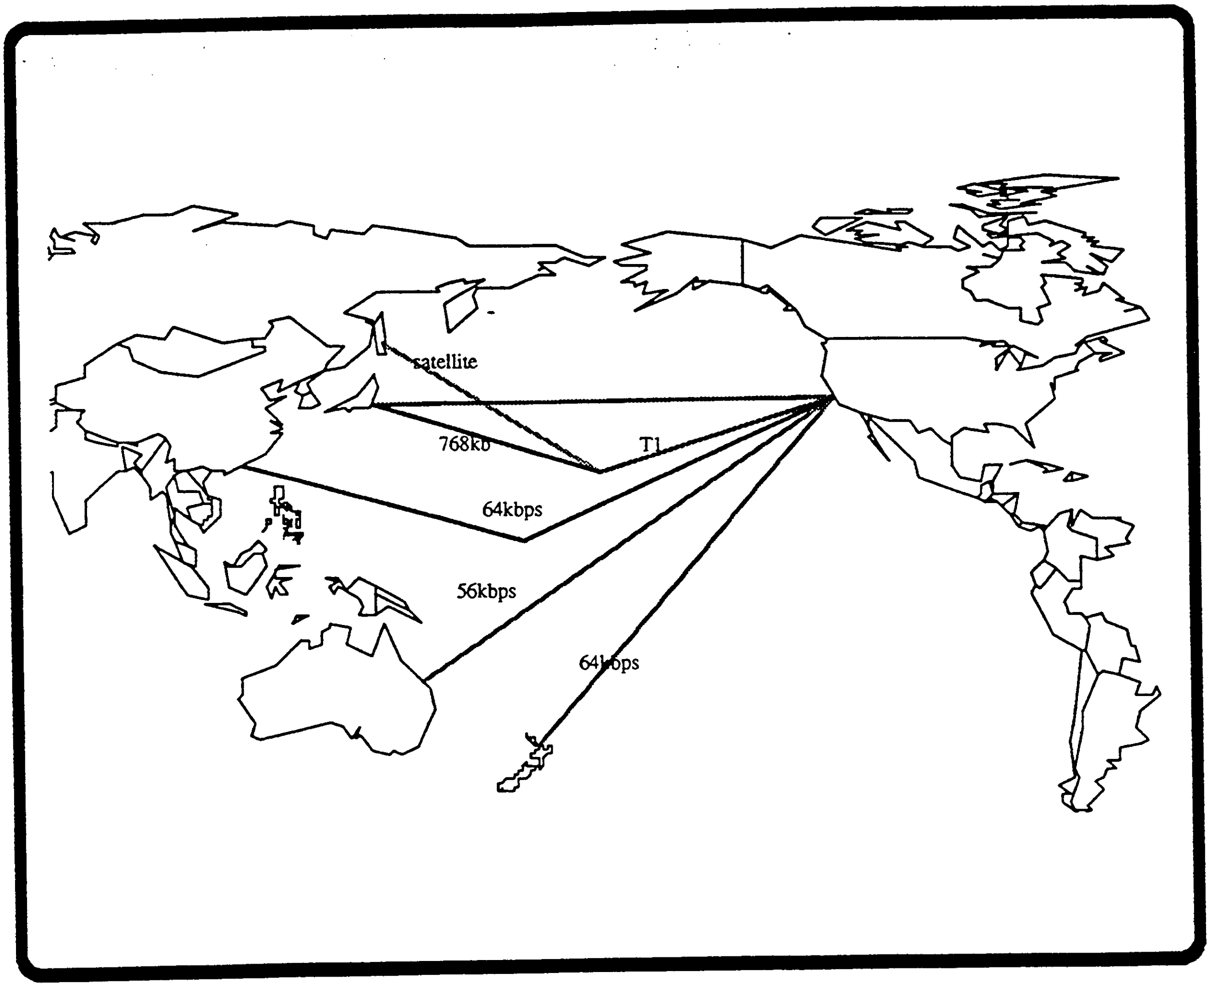
satellite
768kb
T1
64kbps
56kbps
64kbps

4.6 RIPE Report

Presented by Daniel Karrenberg/RIPE

Managing European IP-Networks

Introduction

Daniel Karrenberg

- Works at CWI, Amsterdam, Netherlands
- Responsibilities: internal & external networking
- Includes: running EUnet European node "`mcsun.eu.net`" & EUnet IP backbone
- Vice chairman of RIPE (Reseaux IP Europeens)
- Experience: Unix, service, UUCP, IP, X.25
- **<Daniel.Karrenberg@cwi.nl>**

The Environment

IP networking in EUrope is a huge success:

- more than 100.000 directly reachable hosts
- 767 connected networks
- about 400 organisations
- about 25 network operators
- no managed backbone network
- naturally grown connectivity

It works !

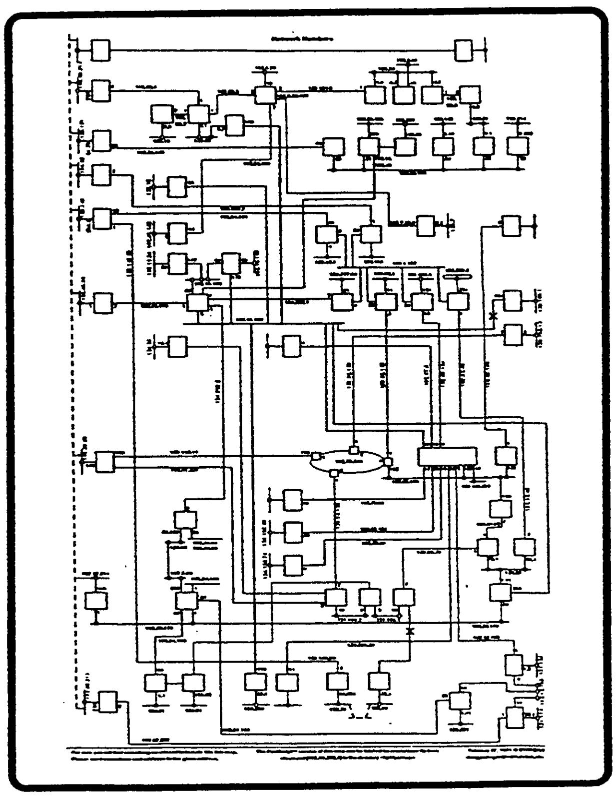

RIPE

RIPE: Reseaux IP Europeens

- founded in 1989
- open for *all* IP network operators in Europe
- coordinates
- makes recommendations
- all activities/documents are open and public
- does ***not*** operate networks

R I P E

RIPE coordinates the management of IP networks in Europe by

- producing (non-binding) recommendations
- maintaining a management information database
- creating and maintaining a "human network"
- defining and managing the activities of the RIPE Network Coordination Centre (NCC)

R I P E

RIPE Recommendations

- Router Management
- Operational Contacts
- DNS name server Operation
- Routing Coordination (being finalised)

R I P E

RIPE Network Management Database

- Networks
- Contact Persons
- DNS Domains
- Name servers (in preparation)
- Routers (in preparation)

R I P E

Network Management Database:

------> WHOIS.RIPE.NET

```
whois karrenberg
person: Daniel Karrenberg
address: Kruislaan 413
address: NL-1098 SJ Amsterdam
address: Netherlands
phone: +31 20 5924112
fax-no: +31 20 5924199
e-mail: dfk@cwi.nl
nic-hdl: DK58
changed: piet@cwi.nl 901115
```

TEMPLATES: FTP.RIPE.NET RIPE

R I P E

```
whois 192.36.148
inetnum: 192.36.148.0
netname: NORDUNET-BACKBONE
descr: NORDUNET,SE
country: SE
admin-c: Bjorn Eriksen
tech-c: Anders Hillbo
connect: NORDU RIPE NSF
gateway: kth
changed: ber@sunet.se 910521
```

R I P E

Network of People

- RIPE acts as a forum for information exchange
- RIPE meetings are focal points
- non time critical management contacts during meetings
- building of common spirit and mutual trust
- coordinate volunteer activity for maximum benefit of all
- maintain contacts to other continents

RIPE

```
whois imr.no
domain: imr.no
descr: Havforskningsinstituttet
descr: Postboks 1870,     Nordnes,
       5024 Bergen
admin-c: Knut Hestenes
tech-c: Roald Lygre
zone-c: Kenneth Hostland
nserver: lilje.uib.no
nserver: eik.ii.uib.no
nserver: nac.no
dom-net: 129.177.44
```

RIPE

Because of the tremendous growth some RIPE activities can no longer effectively be done on a volunteer basis!

Therefore RIPE decided to call for the establishment of the RIPE Network Coordination centre.

The NCC is not a NOC.

The RIPE NCC

The RIPE NCC performs those coordination functions which cannot reliably be provided by voluntary efforts of the RIPE members:

- collect network management information
- collect and develop specific tools
- support RIPE activities administratively
- report regularly

The RIPE NCC

- necessity agreed within RIPE early 1990
- activities defined and agreed September 1990
- structure and setup agreed May 1991
- decision to formally set up NCC within RARE
- some initial funding secured at the same time
- first year funding secured at RARE CoA Rome

The RIPE NCC

Collect Network Management Information

- IP connectivity in Europe: "The Maps"
- network management database
  - networks
  - contact persons
  - DNS domains
  - routers
  - links
- directory of operational contacts
- traffic statistics: "The European View"

The RIPE NCC

Make Network Management Information Available

- maintain a "whois" server
- network Management Database FTP-able
- query Tools
- exchange with other organisations
  - US NIC
  - NSF
  - other continents
- common statistics format / exchange

The RIPE NCC

Collect, create and make available tools

- to maintain the management database
- to query the management database
- statistics gathering tools
- statistics presentation tools
- DNS quality control tools
- mapping tools

The RIPE NCC

Administrative Support

- keep the RIPE document store
- technical support for RIPE projects
- secretarial support for RIPE projects
- secretarial support for RIPE meetings

Reports

- quarterly and yearly about NCC activities
- quarterly and yearly about IP in Europe

The RIPE NCC

People (formally hired by RARE)

- Manager (1 person)
- Technical Staff (1 person)
- Administrative staff (1 person +)

Location

- near an operating NOC AMSTERDAM, NETHERLANDS
- good connectivity

Reporting

- to RIPE (chairman) on all activities
- to RARE on administrative matters

Chapter 5

IETF Protocol Presentations

5.1 Dynamic Host Configuration Protocol

Presented by Ralph Droms/Bucknell

The Dynamic Host Configuration Protocol (DHCP) provides configuration parameters to Internet hosts. DHCP consists of three components: a protocol for delivering host–specific configuration parameters from a DHCP server to a host; a mechanism for allocation of network addresses to hosts; and a protocol through which a collection of DHCP servers can cooperatively allocate network addresses from a shared pool of network addresses.

DHCP is designed to supply hosts with configuration parameters as defined in the Host Requirements RFCs. After obtaining parameters from DHCP, a host should be able to exchange packets with any other host in the Internet. The parameters supplied by DHCP are listed in an appendix of the DHCP Internet Draft..

DHCP allows but does not require the configuration of host parameters not directly related to the IP protocol. A site may choose to use vendor extensions to return information about higher level protocols to the host. DHCP also does not address registration of newly configured hosts with DNS. Also, DHCP is not intended for use in configuring routers.

DHCP is built on a client–server model, where designated DHCP server hosts allocate network addresses and deliver configuration parameters to dynamically configured hosts. DHCP uses only designated DHCP servers, rather than allowing any Internet host to act as a DHCP server. The diversity of hardware and protocol implementations in the Internet would preclude reliable operation if random hosts were allowed to respond to DHCP requests. The format of DHCP messages is based on the format of BOOTP messages, to capture the BOOTP relaying agent behavior described as part of the BOOTP specification and to allow interoperability of existing BOOTP clients with DHCP servers.

DHCP – Dynamic Host Configuration Protocol

Ralph Droms
Computer Science Department
Bucknell University
Lewisburg, PA 17837

droms@bucknell.edu

Introduction

Configuration of Internet hosts is not well–automated. Compared to other protocols, notably AppleTalk, computers using the TCP/IP protocol suite require significant manual configuration.

The *Dynamic Host Configuration Protocol* (DHCP) automates the configuration procedure for Internet hosts. DHCP can provide "plug–and–play" operation of Internet hosts, including delivery of local network configuration paramaters and automatic allocation of Internet addresses.

Topics

- Overview of DHCP
- Design decisions
- Protocol details
- Status
- Future work

Problem statement

- Internet hosts require configuration parameters
- Hosts without non–volatile storage must reacquire configuration parameters with each initialization
- "Out–of–the–box" hosts must be assigned a new network address
- DHCP must supply network configuration parameters through the network before the host can fully use TCP/UDP/IP

Overview of DHCP

- Provides three services:
  - Persistent repository of configuration information for clients (hosts)
  - Dynamic allocation of configuration resources such as network layer addresses
  - Distribution of configuration information among protocol servers
- Supplies all network configuration parameters specified in RFC 1122
- Can also supply some application protocol parameters such as DNS, NTP and NIS servers
- Provides dynamic allocation of network addresses
- Client–server architecture with designated servers

DHCP and BOOTP

- DHCP message format is based on BOOTP message format, with many new vendor extensions
- DHCP hosts and servers can use BOOTP relay agents
- BOOTP and DHCP network address allocation mechanisms:
  - pre–configured address assignment
  - automatic address assignment
  - dynamic address assignment

Design constraints

- Guarantee that no network address is assigned to multiple hosts
- Retain host configuration (wherever possible) across host and server reboots
- Avoid manual configuration by system administrators
- Allow servers to manage multiple subnets
- Use existing BOOTP *relay agent* service
- Coexist with non–participating hosts and BOOTP hosts
- ASSUMPTIONS - DHCP IS ONLY GUARANTEED FOR INITIALIZATION UP TO AND INCLUDING THE TRANSPORT LAYER
- HIGHER-LAYER PARAMETERS MAY BE CARRIED IF SPACE IS AVAILABLE IN THE DHCP PACKET

Protocol specification

- Message formats
  - Based on BOOTP
  - Expanded message size
  - New configuration parameters
- Host behavior
- Server behavior

Message formats

- All messages use the same format:

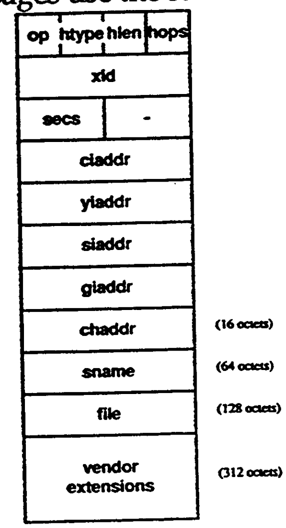

Message details

- "Vendor extensions" include all parameters specified in RFC 1122–1123, RFC 1191 and RFC ~~[illegible]~~ 1256
- Operations are specified in "DHCP message type" vendor extension:
  - DHCPDISCOVER
  - DHCPOFFER
  - DHCPREQUEST
  - DHCPDECLINE
  - DHCPACK
  - DHCPRELEASE

Host–server interaction (acquiring a network address)

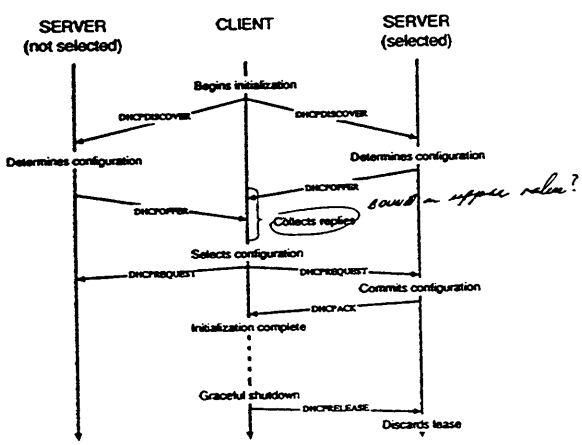

Host–server interaction (reusing an assigned network address)

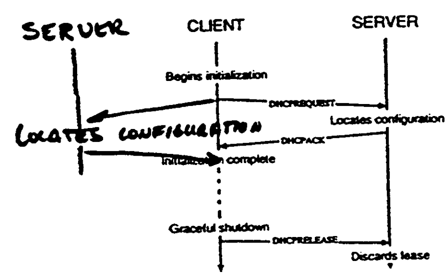

Dynamic allocation of network addresses

- Dynamic allocation of network address based on "lease"
  - Host agrees to terminate use of address at expiration of lease
  - Server can reallocate address after expiration
  - Infinite duration lease is equivalent to permanent allocation
- System administrator configures server with list of addresses to be dynamically allocated to DHCP hosts

Host behavior

- Broadcasts DHCPDISCOVER to locate DHCP server or servers
- Selects one set of network parameters
- Broadcasts DHCPREQUEST to request parameters from selected server (and inform other servers of selection)
- Records parameters and begins network communication
- Can release address before lease termination

Server behavior

- Responds to DHCPDISCOVER messages
  - DHCPOFFER contains network– and host–specific parameters
  - Allocates new network address (if needed)
- Responds to DHCPREQUEST messages
  - Records address allocation to non–volatile storage
  - Returns parameters in DHCPACK
  - Discards any hints if not selected

Client state–transition diagram

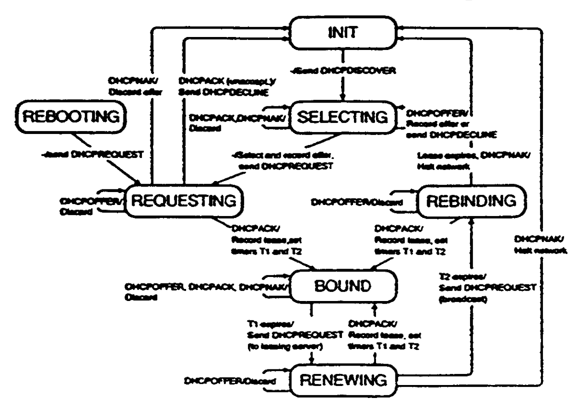

Inter-server protocol

- Servers cooperate to provide master–slave distributed database
- Each network address is controlled by a single server
- Other servers may hold replicated copies
- Database master periodically polls servers to reconcile database and reallocate network addresses

Management interface

- Administrators select site–, network– and host–specific configurations such as:
  - Use of dynamic allocation
  - Network address
  - Subnet mask
  - Broadcast address type
  - Use keep–alives
  - Use router discovery
- Servers notify administrators of error conditions:
  - Network address conflicts
  - Exhaustion of dynamic addresses

What's next for DHCP?

- Protocol status:
  - Current specification in Internet Draft
  - Several implementations complete or in progress: BUCKNELL, BOEING
- Server–to–server protocol being defined

5.2 TCP Large Windows

Presented by Dave Borman/Cray Research and Bob Braden/ISI

TCP EXTENSIONS FOR HIGH-PERFORMANCE

Van Jacobson, LBL

Bob Braden, USC-ISI

Dave Borman, Cray Research

IETF Presentation

March 18, 1992

TCP EXTENSIONS

- Belief: Modification of TCP header not acceptable to community.

  => TCP Options — "legal" extension facility.

  => Backwards compatible (even with sub-standard TCPs).

- Belief: The importance of these extensions justifies their standardization ASAP.

  This should not foreclose efforts to incorporate these ideas into a revised or wholly new reliable transport protocol.

TCP EXTENSIONS

- RFC-1072: LARGE BANDWIDTH*DELAY PRODUCTS
  - A. Large Windows
  - B. RTTM: RTT Measurement using timestamps
  - C. SACK: Selective Acknowledgments [Deferred]
- RFC-1185: HIGH SPEEDS (> O(100 Mbps))
  - D. PAWS: Protect Against Wrapped Sequence numbers

History

- End-to-End Research Group —
  - June 1987: Msg from Van proposing RTTM, Large Windows, and SACK
  - Feb 1988: Draft of RFC-1072
  - Oct 1988: Publish RFC-1072
  - Nov 1989, Mar 1990: Msgs from Van proposing PAWS
  - March 1990: Draft of RFC-1185
  - Oct 1990: Publish RFC-1185
  - Oct 1991: Combined draft.
- IETF —
  - 1990: TCP Large Windows WG formed.
  - Nov 1991 (Sante Fe): TCP Large Windows WG revised draft.
  - March 1992: IESG Recommendation to IAB: Proposed Standard.

A. LARGE WINDOWS

TCP Window Scale Option

3	L=3	S

$0 \leq S \leq 14$

Effective Window = $2^S \times$ Transmitted Window

- Both sides must send Window Scale Option in SYN segments.
- Run out of 64K maximum TCP window at throughputs of:

  O(1 Mbps) for satellites,

  O(10 Mbps) for terrestial links.

A. LARGE WINDOWS

- Q: How big should window scale option be?

  Ans: Big enough so entire user receive [BSD: socket] buffer can be advertised.

- Q: When send window scale option?

  Ans: When receiver's buffer > 64K;

  but only option send in <SYN,ACK> segment if received one in <SYN> segment.

B. RTTM: MEASURING ROUND-TRIP TIME

TCP Timestamps Option: TSopt

1	1	4	4
8	L=10	Timestamp Value TSval	TS Echo Reply TSecr

- Symmetry: to match TCP's symmetry.
- Both sides must send TSopts in SYNs, to allow TSopts in data and ACK segments.
- When useful ? Always! More reliable RTO measurement.

RTTM Example

TCP A — TCP B

1. —> < SYN, TSval=1 > —>
2. <— < SYN, ACK(SYN), TSval=17, TSecr=1 > <—
3. —> < ACK(SYN), TSval=2, TSecr=17 > —>

4. —> < data1, TSval=5, *TSecr=17* > —> (ignore TSecr)
5. <— < ACK(data1), TSval=29, TSecr=5 > <—

6. —> < data2, TSval=8, *TSecr=29*> —> (ignore TSecr)
7. <— < ACK(data1), TSval=33, TSecr=8 > <—

RTTM SPECIAL CASES

TCP A — TCP B

1. —> < dataA, TSval=1 > —> (ACK delayed)
2. —> < dataB, TSval=2 > —>
3. (A) <— < ACK(dataB), TSecr=1 > <—
4. —> < dataC, TSval=3> —> X (lost)
5. —> < dataD, TSval=4 > —> (Queued)
6. (B) <— < ACK(dataB), TSecr=1 > <—
7. (Retransmit) —> < dataC, TSval=8 > —>
8. (C) <— < ACK(dataD), TSecr=8 > <—

D. PAWS: TCP CORRECTNESS AT HIGH SPEEDS

Discard old duplicate segments, despite sequence number wrap-around.

Bandwidth	Time to wrap 32-bit sequence #
1.5 Mbps	> 3 hrs
10 Mbps	~ 30 mins
45 Mbps	~ 6.5 mins
100 Mbps	~ 3 mins (Getting risky...)

Old data segments: can cause undetected errors.

Old ACK segments: can cause connection lock-up/failure.

TCP accepting old duplicate data segments (circles). Sun OS 4.1.1 TCP, modified to wrap sequence space in 16K bytes instead of 2**32.

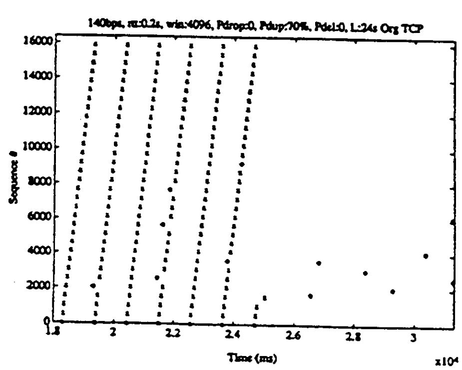

Same, but with PAWS mechanism implemented => Old duplicates rejected.

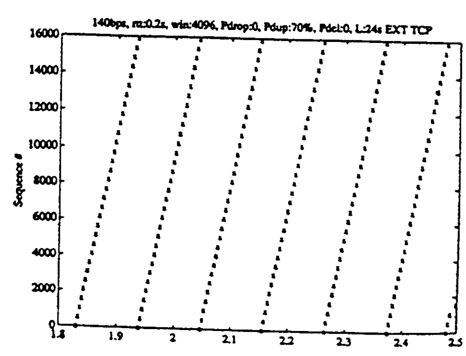

PAWS — Protect Against Wrapped Sequence #s

Jacobson: Use the TCP timestamp option from RFC-1072 as a logical extension of the sequence number field.

Discard segments whose timestamps are earlier than the timestamp of the last segment received at the left window edge.

- Timestamp initialized by TSval in SYN segment.
- Send timestamps on both ACK and data segments.
- Still use 3-way handshake to validate SYN segments.

PAWS & RTTM USE SAME TIMESTAMP CLOCK

- Timestamp clock:
  - Must be proportional to real time (for RTTM)
  - Must be MONOTONIC within same connection.
  - Frequency: **1/sec - 1/ms.**
  - Not monotonic across crashes or new connections.

    *[assuming TIME-WAIT and 'Quiet Time' of 2*MSL].*

TIMESTAMP CLOCK

- Making **monotonic** timestamp clock:
  - Use a hardware clock, or
  - Count clock interrupts, or
  - Add a variable offset to a system (software) clock that is subject to random resets.

  Example: BSD: count SLOWTIMO ticks => 2 per second.

TIMESTAMP CLOCK

- OLD-AGE TIMESTAMP VALUES
  - If connection is IDLE for 25 days, the timestamp value saved in a TCPCB wraps around, and then all new timestamps will be rejected.
  - => Invalidate PAWS timestamp if connection is idle for more than 25 days.

C. SELECTIVE ACKNOWLEDGMENTS

- Questions were raised => defer for more simulation and testing.

  [Work in progress at LBL and ISI].Interim results:

- Interim results:
  - When multiple packets sometimes lost per round trip time, SACK can give significant throughput improvement.
  - Augments Fast Retransmit and Fast Recovery algorithms of 4.3Reno.
  - Can avoid delivery bursts
  - Not very complicated to implement.

SACK Encoding

- Problem: Size of SACK option.

  [With Timestamps option, 28 bytes available for SACK.]

  Want compact and efficient encoding.

- Want compact and efficient encoding.
  - WG rejected interaction with window scaling => SACK support large windows.
  - "Reasonable" encodings: describe max 3 or 4 holes in sequence space.
  - If cannot describe all losses in one SACK, take additional RTT's to recover [but still don't need RTO usually].

TCP Extensions — Implementations

Implementations include:

- Cray Research [Dave Borman will describe]
- 4.3BSD Reno (2–3)
- Sun OS 4.1
- SGI Irix

Mailing List: tcplw@cray.com

TCP/IP Performance at Cray Research

David A. Borman

Cray Research, Inc
655F Lone Oak Drive
Eagan, MN 55121

dab@cray.com

CRAY Networking — IETF — San Diego, March 18, 1992

History

SNQ1: Single Processor Cray 2
Software Loopback Driver
32K MTU
Memory to Memory Transfer

9th IETF meeting, March 1-3, 1988, San Diego, CA
Kernel Buffer: 128K
User to Kernel copy size: 4 K
118 Mbits/sec

10th IETF meeting, June 15-17, 1988, Annapolis, MD
Kernel Buffer: 512K
User to Kernel copy size: 32 K
247 Mbits/Sec

Computer Communications Review, April, 1989
Kernel Buffer: 512K
User to Kernel copy size: 32 K
310 Mbits/Sec

YMP: 550 Mbits/sec

CRAY Networking — IETF — San Diego, March 18, 1992

History

SNQ1: Single Processor Cray 2
Software Loopback Driver
64K MTU
Memory to Memory Transfer
Machine is Dedicated

19th IETF meeting, Dec 3-7, 1990, Boulder CO
Kernel Buffer: 378K
User to Kernel copy size: 64 K
461 Mbits/Sec
YMP: 795 Mbits/Sec

Changes since Boulder Meeting

TCP Option Prediction Code
Two Forms: Header Prediction & Regular TCP

Combined Copy/Checksum Routine
Used in tcp_output()
Unicos 7.0

In-line Header Checksum Routine
Short, word-aligned checksums
Unicos 8.0

CRAY Networking — IETF — San Diego, March 18, 1992

Today

SNQ1: Single Processor Cray 2
Software Loopback Driver
64K MTU
Memory to Memory Transfer
Machine is Dedicated
Kernel Buffer: 378K
User to Kernel copy size: 64 K
Window Scale Option

SNQ1: 522 Mbits/Sec

YMP: 907 Mbits/Sec

CRAY Networking — IETF — San Diego, March 18, 1992

HIPPI

Rated at 800 Mbits/Sec
YMP to YMP Model E IOS
User to Kernel copy size: 64 K
Machines are Dedicated
Memory to Memory Transfer
Unicos 8.0
Through Switch Requires Disconnects After Each Packet

Through Switch	Kernel Buffer	MTU	Window Shift	Application Read/Write	Transfer Rate
Yes	378K	33K	4	100@1512K	416 Mbits/sec
Yes	576K	49K	4	100@1512K	525 Mbits/sec
Yes	.378K	64K	4	100@1512K	605 Mbits/sec
No	378K	64K	None	100@1512K	199 Mbits/sec
No	378K	64K	1	100@1512K	381 Mbits/sec
No	378K	64K	2	100@1512K	537 Mbits/sec
No	378K	64K	3	100@1512K	775 Mbits/sec
No	378K	64K	4	100@1512K	781 Mbits/sec
No	288K	49K	4	100@1152K	571 Mbits/sec

781.345 Mbits/sec is 98.96% of theoretical maximum

(781.345*((63*1024+96)/(63*1024)))/(800*(256/259)) = .9896

Multi Threaded Kernel

2 Processor YMP
Locking on per-stream basis
Semaphore Debugging Enabled
Read/Write: 100@1512K per stream
Wall Clock time includes process startup, connection establishment, and statistic reporting

Single TCP Stream, Software Loopback

Single Threaded: 907 Mbits/sec
Multi Threaded: 1.08 Gbits/sec

Dual TCP Streams

Single Threaded: 456/441 Mbits/sec
Wall Clock: 6.21 secs, 798 Mbits/sec
Multi Threaded: 687/679 Mbits/sec
Wall Clock: 4.25 secs, 1.16 Gbits/sec

Three TCP Streams

Single Threaded: 315/301/302 Mbits/sec
Wall Clock: 9.35 secs, 794 Mbits/sec
Multi Threaded: 494/484/431 Mbits/sec
Wall Clock: 6.40 secs, 1.16 Gbits/sec

FUTURES

Replace ECHO/ECHO REPLY options
TIMESTAMPS code waiting for RFC

Multi Thread the Kernel
2 Processor YMP, Software Loopback
With Debugging On: 1.08 Gbits/sec

YMP-C90
Single Threaded: expect at least 50% over YMP
Multi Threaded: Who Knows?

64 bit Hippi Channel
Rated at 1.6 Gbits/sec

Copy/Checksum to/from User Space

Streamline Device Drivers


```
[sn1703c] ftp localhost
Connected to localhost.cray.com.
220 sn1703c.cray.com FTP server (Version 5.2 Fri Sep 7 14:09:58 CDT 1990) ready.
Remote system type is UNIX.
Using binary mode to transfer files.
Name (localhost:dab):
331 Password required for dab.
Password:
230 User dab logged in.
ftp> cd /tmp
250 CWD command successful.
ftp> get dab.tmp dab.tmp2
200 PORT command successful.
150 Opening BINARY mode data connection for dab.tmp (38371072 bytes).
226 Transfer complete.
38371072 bytes received in 1.4 seconds (2.7e+04 Kbytes/s)
ftp> get dab.tmp dab.tmp2
200 PORT command successful.
150 Opening BINARY mode data connection for dab.tmp (38371072 bytes).
226 Transfer complete.
38371072 bytes received in 0.5 seconds (7.5e+04 Kbytes/s)
ftp> get dab.tmp2 dab.tmp3
200 PORT command successful.
150 Opening BINARY mode data connection for dab.tmp2 (38371072 bytes).
226 Transfer complete.
38371072 bytes received in 0.59 seconds (6.3e+04 Kbytes/s)
ftp> get dab.tmp dab.tmp2
200 PORT command successful.
150 Opening BINARY mode data connection for dab.tmp (38371072 bytes).
226 Transfer complete.
38371072 bytes received in 0.52 seconds (7.3e+04 Kbytes/s)
ftp> get dab.tmp2 dab.tmp3
200 PORT command successful.
150 Opening BINARY mode data connection for dab.tmp2 (38371072 bytes).
226 Transfer complete.
38371072 bytes received in 1.4 seconds (2.8e+04 Kbytes/s)
ftp> quit
221 Goodbye.
[sn1703c] ls -l dab.tmp
-rw-r--r--  1 dab      38371072 Mar 14 13:36 dab.tmp
[sn1703c] time cp dab.tmp dab.tmp2

real    0m1.16s
user    0m0.00s
sys     0m0.38s
[sn1703c] time cp dab.tmp dab.tmp2

real    0m1.21s
user    0m0.00s
sys     0m0.42s
[sn1703c] time cp dab.tmp dab.tmp2

real    0m1.18s
user    0m0.00s
sys     0m0.43s
[sn1703c] bc
7.5*10*8
600.0
2.8*10*8
224.0
38371072*8/1000000/1.16
263
```

5.3 The Internet Gopher Protocol

Presented by Mark McCahill/UMinn

THE INTERNET GOPHER PROTOCOL
a distributed document search and retrieval protocol

Bob Alberti, Farhad Anklesaria, Paul Lindner,
Mark McCahill, Daniel Torrey

Microcomputer and Workstation Networks Center
University of Minnesota
Minneapolis, MN 55455

What is Internet Gopher?

The Internet Gopher protocol is designed primarily to act as a distributed document search and retrieval system. The protocol and software follows a client-server model. Documents reside on many autonomous servers on the Internet. Users run client software on their desktop systems, connecting to a server and sending the server a selector (a line of text, which may be empty) via a TCP connection at a well-known port. The server responds with a block of text terminated with a period on a line by itself and closes the TCP connection. No state is retained by the servers. The simple nature of the protocol stems from the need to quickly and efficiently implement servers and clients for the slow, smaller desktop computers (1MB Macintoshes and DOS machines.) For all the simplicity, Internet Gopher yields a surprising amount of functionality.

The protocol is designed to permit users on a heterogeneous mix of desktop systems to browse, search, and retrieve documents residing on multiple distributed server machines. While documents (and services) reside on many servers, Gopher client software presents users with a hierarchy of items and directories much like a file system. Submitting a query to a search server yields "virtual directory listings" that contain files matching the search criteria. The Gopher interface is designed to resemble a file system since a file system is a good model for organizing documents and services; the user sees what amounts to one big networked information system containing document items, directory items, and full-text searching capabilities across subsets of the information base.

Servers return either directories or documents. Each item in a directory is identified by a type (the kind of object the item is), user-visible name (used to browse and select from menu listings), an opaque selector string (typically containing a pathname used by the destination host to locate the desired object), a host name (which host to contact to obtain this item), and an IP port number (the port at which the server process listens for connections.) The user only sees the user-visible name. The client software can locate any item by the trio of selector, hostname, and port.

Services Available Via Gopher.

Currently, besides browsing and searching files and directories on a network of Gopher servers, users can obtain information from Archie servers, WAIS servers, and FTP servers without leaving the familiar Gopher user interface. Users can also search for names and phone numbers on campus phone-book servers, and access Telnet-session based information servers such as the myriad of library card catalogs on the Internet. The campus phone-book servers and telnet-session based services appear as special items in directory listings. The protocol is extensible by adding new types to the existing collection of Gopher objects. The protocol is also extensible by writing gateways to other services (as has been done for Archie, FTP, etc.)

Obtaining Gopher Software.

Client software is available for Macintosh, DOS-based machines, generic curses-based UNIX, X, NeXT (browser style interface), IBM VM/CMS, and VAX/VMS. Server software is available for Macintosh, UNIX, IBM VM/CMS, and VAX/VMS; a server should soon be completed for DOS machines. Full-text search server software is available for generic UNIX (based on the public domain WAIS search engine) and for NeXT (using their native indexing libraries). Full-text search servers should soon be completed for Macintosh. All software is available via anonymous ftp (or via Gopher) from boombox.micro.umn.edu.

The Internet Gopher development team can be reached via e-mail at the first address below. To be added to the Gopher-News mailing list, send e-mail to the second address.

- mailing list: gopher@boombox.micro.umn.edu.
- request list: gopher-news-request@boombox.micro.umn.edu.

Other gopher concerns are often discussed in the USENET newsgroup: alt.gopher.

Internet Gopher

Bob Alberti, Farhad Anklesaria, Paul Lindner, Mark McCahill, Daniel Torrey

e-mail: gopher@boombox.micro.umn.edu

University of Minnesota

What is Gopher?

- A distributed document serarch and retrieval protocol
- Supports browsing through hierarchical collections of menus
- Supports full-text searches

What is Gopher?

- Gateway servers to WAIS, FTP, Archie (ie. not leaving the Gopher interface)
- Supports Telnet sessions to non-Gopher Telnet-based information services
- Servers for Macintosh, UNIX, VMS, VM/CMS
- Clients for Macintosh, UNIX, PC, VMS, VM/CMS

Gopher Philosophy

- A simple and extensible protocol
- Information producers can easily put up their own server
- System provides seamless navigation from one server to another
- Different sites can organize information in Gopher to suit their needs

How does it work?

- A client/distributed server architecture
- Server returns a list of objects to the client
- A server can return "links" to objects on other servers to the client
- Search engines can be separate from data

Sample Gopher session

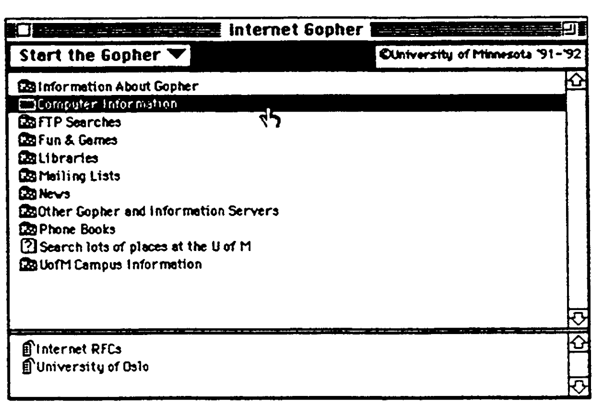

Browsing in Gopher

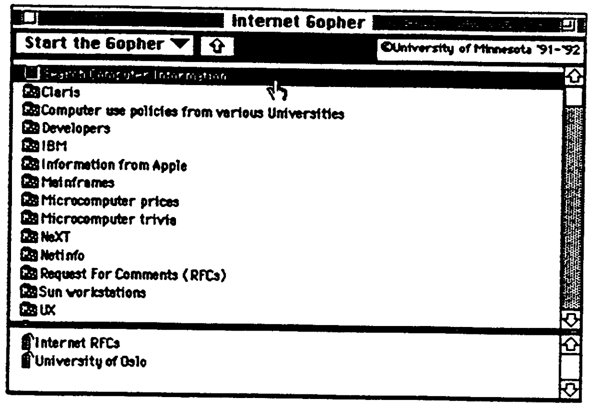

Searching for documents

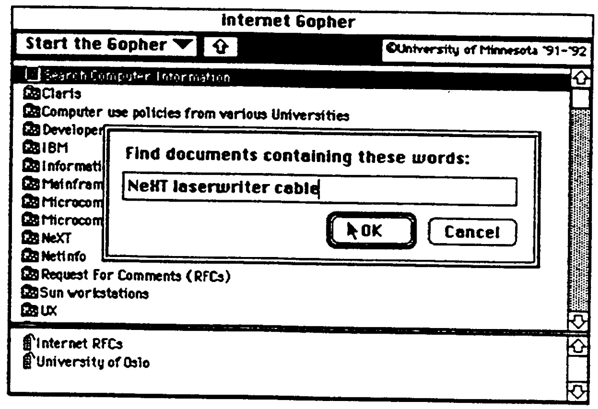

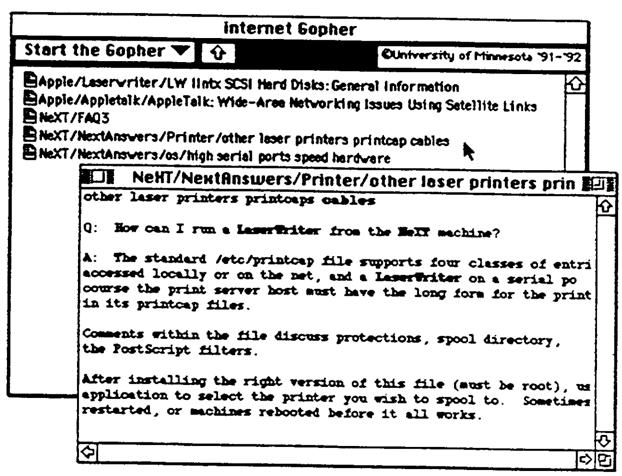

Other Gopher servers...

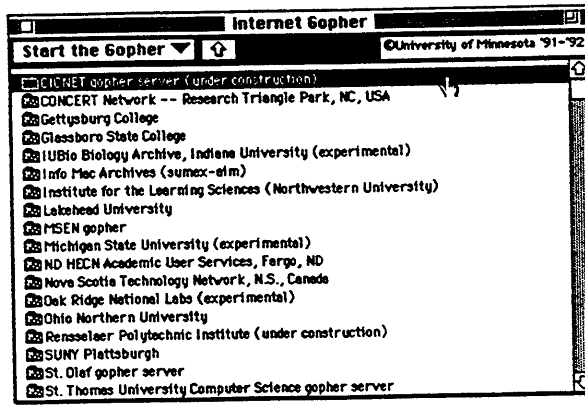

Gateways to other services

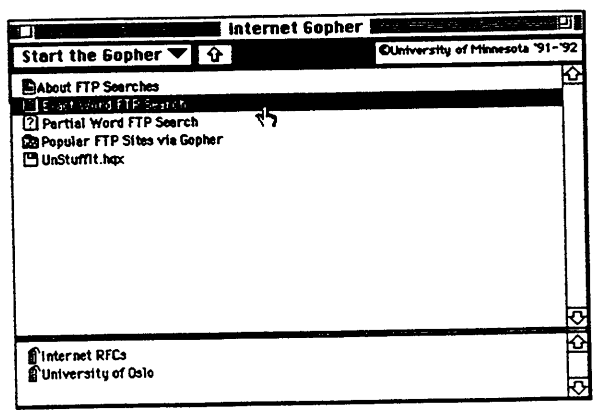

Telnet based information

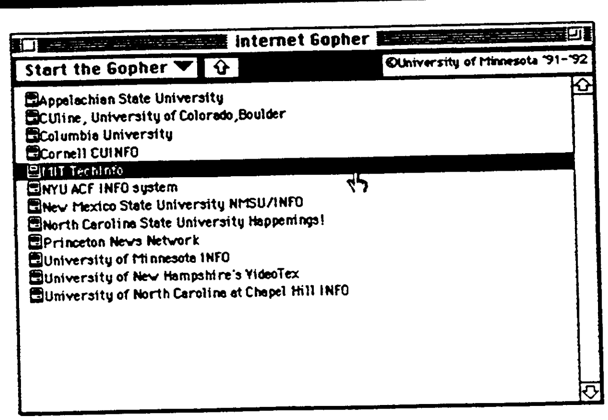

Library card catalogs

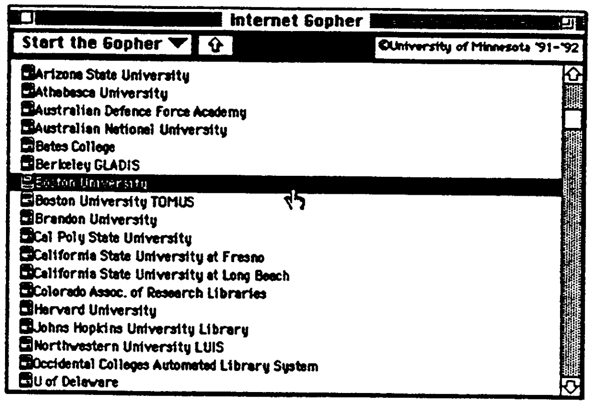

How is an object described?

- object type
- object name
- selector string
- server machine's domain name
- server port number

What kinds of objects are there?

- text documents
- directories
- search engines
- telnet session descriptors
- CSO (ph) phonebook server descriptors
- some experimental types. Gopher is extensible.

5.4 IP Over X.25

Presented by Andy Malis/BBN

Multiprotocol Interconnect over X.25

Goals and Requirements

- Interoperable upgrade to RFC 877
- Conform with CCITT and ISO standards
- Allow interoperable multiprotocol operation between routers and bridges over X.25
- Allow both multiplexing multiple protocols over a single circuit and specifying individual protocols for each circuit
- Extend IP MTU to match current router implementations (1500+ octets)
- Define method for 802.1d bridging over X.25

San Diego IETF

Multiprotocol Interconnect over X.25

Solution

- Use ISO/IEC TR 9577 to define encapsulation format - conforms with both existing RFC 877 and ISO X.25 implementations
- Defines Network Layer Protocol Identifier in first octet of X.25 Call Request Call User Data Field:

  CC - IP
  81 - CLNP, ES-IS, IS-IS, IDRP
  80 - SNAP (followed by 5-octet SNAP header)
  00 - Null Encapsulation: Each X.25 data packet carries a NLPID (Allows multiplexing multiple protocols on one VC)

- Ethertypes are available via SNAP (OUI 00-00-00)
- Bridging is available via SNAP, using 802.1's OUI, as specified in RFC1294
- NLPID of CC is the only required encapsulation; all others are optional

San Diego IETF

Multiprotocol Interconnect over X.25

Encapsulations

IP Encapsulation Call Request:

GFI, LCN, type	Addresses	Facilities	CC

Data Packet:

GFI, LCN, seq. #s	IP Datagram

CLNP Encapsulation Call Request:

GFI, LCN, type	Addresses	Facilities	81

Data Packet:

GFI, LCN, seq. #s	CLNP Datagram

San Diego IETF

Multiprotocol Interconnect over X.25

SNAP Specific Protocol Encapsulation Call Request:

GFI, LCN, type	Addresses	Facilities	80	SNAP (5 octets)

Data Packet:

GFI, LCN, seq. #s	Protocol Data Unit (no SNAP header)

Multiplexed Encapsulation Call Request:

GFI, LCN, type	Addresses	Facilities	00

Data Packet:

GFI, LCN, seq. #s	NLPID (1 octet)	Protocol Data Unit

San Diego IETF

Multiprotocol Interconnect over X.25

Multiplexed IP Data Packet:

GFI, LCN, seq. #s	CC	IP Datagram

Multiplexed CLNP Data Packet:

GFI, LCN, seq. #s	81	CLNP Datagram

Multiplexed SNAP Data Packet:

GFI, LCN, seq. #s	80	SNAP (5 octets)	Protocol Data Unit

San Diego IETF

Chapter 6

Technical Presentations

6.1 ROAD Presentation

Presented by Phill Gross/ANS and Peter Ford/LANL

Routing and Addressing Report
The ROAD group

Presented by
Peter S. Ford, LANL
peter@lanl.gov

Phill Gross, ANS
pgross@ans.net

16 March, 1992
San Diego IETF

ROAD group: Routing and Addressing

- Chartered by IAB
  - Follow up San Diego IAB retreat on Architecture of Internet, June 1991
  - Vision: Large Global Public Data Network
  - Goal: Scale Internet to 10**9 Networks
  - Must not have a Flag Day
  - Hip Pocket Solution
- Members primarily from Santa FE BGP working group

ROAD group methodology

Met 5 times: Santa Fe, Reston, Boston, Tucson.
Output of ROAD group intended to be input for:

- IAB, IETF
- Intercontinental Engineering Planning Group and CCIRN
- Federal Networking Council/FEPG
- Autonomous Networks Research Group (ANRG
- Network Operators, Vendors, developers of software

ROAD Group members

Dave Bolen (ANS)
Scott Brim (Cornell)
Ross Callon (DEC)
Noel Chiappa (IESG)
Steve Deering (Xerox Parc)
Deborah Estrin (USC)
Dino Faranacci (Cisco)
Peter Ford (Los Alamos)
Vince Fuller (BARRnet)
Elise Gerich (Merit)
Phill Gross (ANS)
Sue Hares (Merit)
Bob Hinden (Sun Microsystems)
Van Jacobson (LBL)
Tony Li (Cisco)
Dave Oran (DEC)
Yakov Rekhter (IBM)
Martha Steenstrup (BBN)
Paul Tsuchiya (Bellcore)
Kannan Varadnan (OARnet)
Jessica Yu (Merit)

Acknowledgements:

- Vint Cerf, CNRI, and especially Terry Weigler
- Lyman Chapin, BBN
- Joel Snyder, U of Arizona
- NSF, DARPA, DOE
- all of our employers ...

ROAD Group as Consultants
Customer is the Internet.
Believe that the Internet will evolve into one of the major global public networks.
Problems are those of success:

- Getting Large
- Resource Depletion
- How to manage

Typical of almost all successful enterprises.

Problems ROAD addressed

- IP address space will run out
- Class B runout, Real Soon Now
- Routing Tables are getting big
- Some management issues:
  - address assignment
  - routing management

Note: current immediate solution when class B network addresses runout is to hand out class C network addresses, aggravating the problem of routing table size.

Major Problems not addressed by ROAD

- Potential of major Routing Architectural changes
- ATM is coming
- LPDN are coming
- Policy is here and won't go away
- "centrally administered routing"
- TOS/QOS
- Mobility
- Is AD model the "right way" to evolve.

Overview of rest of talk

- Assumptions in routing architecture and possible futures
- Strategies for IP address space runout
- A Strategy (tactic) for managing current IP address space
- ROAD activities at San Diego IETF

Assumptions in Routing Architecture and Possible Futures

Routing Architecture – Evolutionary

- Current major transition is to break out of hierarchical rut (EGP-2 → BGP)
- Will facilitate richer interconnection topology
- Still hop by hop
- BGP adds Path Vector
- Don't expect major changes in the immediate future.

Routing Architecture (cont.)

- Operate in arbitrary ID topology
- Support scaling of forwarding and routing information
- More hierarchical addressing scheme

Changes to expect in the future, some sooner than others

- Policy-based routing with selection at the source and support for transit restrictions.
- Confederations (ADs of ADs, nesting)
- Need to study impact of using QOS/TOS on IGPs, ID routing protocols.
- Source Demand Routing

Long term future

- Basic infrastructure of public network supporting "no brainer routes"
- More Special routes, installed by sources
- Major overhaul? NIMROD
- Use of Multicast
- Mobility
- Dynamic Topology (Switched Services)
- ATM, LPDN, etc.

Strategies for IP address space runout

Large Global Internet is now a reality

- We believe it will continue to grow and evolve, but it must overcome fundamental limits in the IP address.
- Immediate scaling problems can be mitigated by "Planning the Internet"
- Last major breakthrough technology in scaling of Internet was DNS which had nice property of hierarchical delegation of assignment authority

The Themes of the Plan:

- Better Addresses
- Better Use of Addresses
- Address Administration can be improved at the same time

Fundamental Problem: IP address space limits

- Need more bits in the address so we can scale to 10**9 networks!
- Need more hierarchy
  - Helps reduce memory load on routing system
  - Helps in delegation of address assignments
- Transition plan must not kill off old IP only hosts.

Possible Solutions to IP address exhaustion

- Use bigger addresses in packet header
  - CLNP
  - IP addresses with more bits
  - Encapsulate
- Make current IP address non global
  - Will happen as IP address space really runs out
  - use encapsulation to record stuff such as AD or Transit network
  - Rewrite/Map addresses at AD borders

Possible Solutions (cont.)

- Remember transition requirement
- Remember we need to route the system so routing technology and infrastructure i must be available.
- Be sure we don't get bit by current IP address space running out on us

Bigger Addresses:

- IP-IP Encapsulation
- Migrate to CLNP
- What about others?
  - The above minimize amount of brand new technology required.
  - The CLNP infrastructure is being deployed today, and router and host vendors will be (or already are) supporting it.

A Migration to CLNP

Takes advantage of success of Internet suite:

- Ubiquitous Network Service
- Datagram Network
- Globally significant address
- BUT, 32 bits is too small
- BUT, too little hierarchy in addresses

Propose using CLNP (ISO-IP), already available. Use rest of IP suite of protocols(TCP, SMTP, UDP, etc.) running over CLNP. Uses current architecture of Internet. In this proposal, migration replaces only the network layer.

Issues in migration to CLNP

- Changing host software, never done.
- Focus on inside out strategy, work at border routers first.
- Current host software (IP) working with updated hosts (CLNP).
- DNS
- Encapsulation vrs. Translation at borders of ADs.
- NSAP addressing schemes? Deering City Codes?

IP Encapsulation

- Encapsulate IP in IP.
- Encapsulating IP used to represent AD info.
- Interdomain routing by AD
- Encaps/Decaps done at 1st and last border routers.
- No Host changes.
- Assumes we do not run out of IP addresses for a long time.
- Maximizes use of current architecture.

IP encapsulation Issues.

- Encapsulation header would use a reserved portion of current IP address space.
- Host part of this space would be AD.
- Need DNS to return mapping from IP address to AD.
- Will shrink routing tables by separating Inter/Intra AD routing.

A Strategy for managing current IP address space

What about the immediate future?

Class B exhaustion looks imminent

- Assignment of class C addresses to preserve class B's adds to routing table explosion
- New address classes (i.e. "Class E") problematic because of old host software.

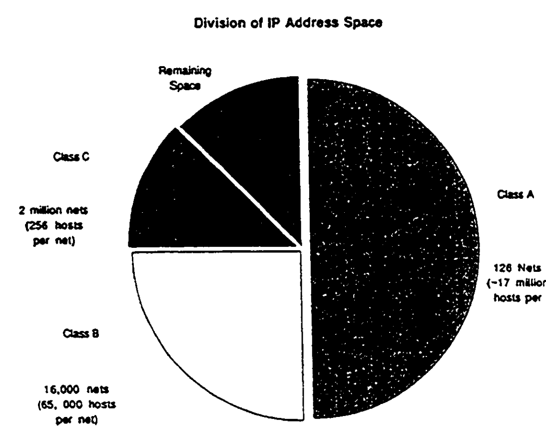

Solution is Classless Interdomain Routing (CIDR)

- Plan the usage and assignment of remaining IP addresses.
- Addresses assigned along routing topology considerations.
- Interdomain routing protocols will aggregate network numbers along topological assignments
- Note: Will not solve IP address space exhaustion problem!

CIDR address and routing plan

- Sites will get contiguous class C network addresses with the maximal number of initial bits being the same (ie, powers of 2)
- ⇒ can use a single prefix for routing (network address w/mask).
- Sites adjacent to each other in the routing hierarchy will be assigned out of the same power-of-two block
- ⇒ multiple sites routed with a single routing table entry.

What is needed to support CIDR Plan

- IETF must develop and standardize the technology base used, there a couple of possibilities:
  - BGP-4
  - IP routed by IDRP (ISO-BGP)
- Operators meet at IETF to establish usage profiles, plans, establish operational requirements.
- Internet Infrastructure (Network Operators) already evolving to support CLNP and use of IDRP.

CIDR at IETF

- IP address plan → ORAD → IANA, IAB
- CLNP address and routing plans already being discussed by operators in
  - NOOP and ORAD
  - Alternate NSAP assignment plan based on geographic lines proposed by Steve Deering
- Using and enhancing BGP in light of CIDR → BGP WG
- Transition to larger IP? → ROAD BOF → WGs?

Future of the Internet is bright

- Crisis == Danger + Opportunity (in Chinese)
- IP address crisis, surmounting it will lead to better routing and easier administration of Internet.

This week at the IETF

- Monday 9:30-12:00 – Geographic Routing and Addressing Assignment (in NOOP)
- Monday 1:30-3:30pm – CIDR Addressing and Routing Plan (in ORAD)
- Tuesday 1:30-6:00pm – Routing and Addressing BOF
- Wednesday 1:30-3:30pm – CIDR Addressing and Routing (joint with BGP WG)
- Thursday 4:00-6:00pm – Discussion in Open Plenary

6.2 Copyright Claims and Standards

Presented by Patrice Lyons/Attorney

I have been working with Vint Cerf on various legal matters such as copyright as they affect the procedures and operation of the IAB and IETF. In preparing for this meeting, I reviewed the many e-mail messages I have received that touch on various aspects of intellectual property rights in documentation produced in the Internet community. These might be working documents or notes submitted for distribution to persons interested in the development of Internet standards. It is clear that some effort should be made to protect the works of individuals and groups in this process.

Perhaps a reference to a recent case involving the right to claim authorship, Weissmann v. Freeman, 868 F.2d 1313 (2d Cir. 1989), may serve to illustrate a possible abuse that may arise where there is no respect for copyright claims in documentation that becomes part of the Internet standards process. The case involved two accomplished medical doctors in the field of nuclear medicine. The parties had jointly authored many articles while Dr. Weissmann was a research assistant to Dr. Freeman. Sometime after this association had ceased, Dr. Weissman wrote a paper reviewing the state of the art of hepatobiliary imaging techniques. While the paper was derived from the former joint work, the court found that there was sufficient original expression for the work to be subject to copyright as a derivative work.

Dr. Freeman used the paper in his course and a foreign lecture, but, in doing so, he deleted Dr. Weissman's name and replaced it with his own. Noting that the total deletion of the original author's name and the substitution of the copier's weighed against a defense of fair use, the court stressed that copyright protection was not limited to economic rewards, but covered gain derived from recognition in a chosen profession, credit that often influences professional advancement and academic tenure.

In the Internet standards process, where I understand from Vint there is a robust and cooperative spirit among the volunteers, I am not inferring that there would be widespread abuse of any copyrights in works generated in this process; however, it is advisable to protect against the isolated case where individual rights, including credit for substantive contributions, may be neglected. Establishing intellectual property procedures that are clear and widely accepted will assist in continuing the productive efforts that have been carried out by the Internet community over the last twenty years.

I have also been impressed with the notion that most members of the Internet community would like to minimize the need to introduce any more complexity in the process than already exists. However, it would be desirable to prevent the standards process from being abused by a single individual who, for personal or other reasons, may seek to delay or even derail the dissemination of a particular standard or a report on work in progress. To reconcile these concerns, it may be advisable to make certain adjustments in the administration of the standards process. I'd like to spend the next few minutes with you addressing these concerns from a copyright perspective. While there are aspects of the standards process that raise patent, trade secret or trademark issues, my comments today will be limited only to copyright considerations.

I have had a chance to read "The Internet Standards Process" circulated as an Internet Draft, dated December 4, 1991. I noted that a distinction is drawn between specifications that are on a standards track, namely those likely to become Internet Standards, often evolving through an experimental maturation process, and those not yet, and perhaps never likely to become standards, namely the non-standards track. A distinction is also drawn between documents that are published as "Internet Drafts" and entered on the IETF's Internet-Draft directory, and documents published as a Request for Comments or RFCs. Various descriptive labels are attached to these documents. From a copyright perspective, I believe there may be good reason to treat standards track documents differently from those not on such a track, whether labeled "Historic," "Experimental," or "Informational."

1. Internet Drafts and other non-standards track documents: Documents prepared and submitted under normal IAB/IETF procedures, which contain no copyright notice, may not be assumed to be in the public domain. Lack of a copyright notice may not just be a slip up on the part of the authors. The placement of a notice on published copies of a work is not a condition of copyright in the United States as well as other countries that are party to the Berne Convention. While it is still advisable to place a copyright notice on all published copies to protect against a defense of "innocent infringement" under U.S. law, and for purposes of the Universal Copyright Convention, absence of a notice will not inject a work into the public domain.

This leads to a question: what's copyrightable in a standards-related document? One may argue, in a given case, that there is not sufficient original authorship to support a copyright claim. Just to give an example. There was a case decided in 1989, Secure Services Tech. v. Time and Space Processing, 722 F.Supp. 1354 (E.D.Va. 1989), involving rights in a manufacturer's variations on a "handshake protocol" known as T-30 that was developed by CCITT. CCITT called the T-30 protocol "Procedures for Document Facsimile Transmission in the General Switched Telephone Network." The T-30 protocol allowed users a limited opportunity to vary the content and timing of various signals used in the protocol.

A particular user claimed that his choice of variations in timing and content constituted an original, derivative work, and claimed copyright in that work. The court found that the variations did not contain sufficient choice and selection to qualify for copyright protection, and, more generally, that the opportunities for original authorship within the constraints of the T-30 protocol were too limited to merit copyright status. There was no reference in the decision to any copyright claims by the CCITT in the T-30 protocol itself, although copyright by CCITT in this protocol did not appear to be in question.

Unlike the T-30 situation, for purposes of this discussion, I assume that most documents submitted to the IAB or IETF, or generated in the course of the Internet standards process more generally, will contain at least some expression subject to copyright protection. This is particularly true where the material may represent computer program code.

Memos or other documents may be submitted, either in paper or electronic form, for dissemination as an "Internet Draft," "RFC" or other publication. As a practical matter, it may not be reasonable or time effective for an IETF Editor or other person receiving non-standards track documents to determine on a case-by- case basis whether, and how much of a particular document is protected by copyright; or, if a work is protected by copyright, who is the owner of rights (for example, whether a contributor is in fact an employee for hire or an independent contractor). The need to go behind the information on the face of a document to determine its copyright status may make the process more complicated than it needs to be.

In order to simplify the publication of these documents, it is advisable to specify in the IETF policies that, by submitting an input, the author or other copyright owner therein shall be deemed to have granted the relevant coordinating body, an implied, non-exclusive right and license to reproduce, distribute, and transmit the work to the public, and to make, distribute and transmit derivative works based on or incorporating such work, and to authorize others to do so. This implied grant of a non- exclusive license under copyright would be followed up soon thereafter with a confirmation of the license, perhaps by e-mail. The confirmation should also contain an expression of willingness on the part of the submitter to furnish written assurances to the Editor on the provision of licenses to users on a reasonable, nondiscriminatory basis in the event their input becomes part of the standards track, and I will come back to this point.

It may be helpful to say a few words about the requirements of the copyright law, at least in the United States, with respect to rights and permissions, recognizing that there will surely be local requirements in each country that may vary somewhat from the practice in the U.S. Generally, under the U.S. copyright law, for a transfer of copyright ownership to be valid, the instrument of conveyance, or a note or memorandum of transfer, must be in writing and signed by the owner of the rights conveyed, or such owner's duly authorized agent. However, this writing and signature requirement does not extend to the grant of a nonexclusive license. It is possible to execute a nonexclusive license using e-mail technology, without the need for a handwritten, or "wet," signature.

It may be argued that a nonexclusive license is implied from the conduct of the parties; and, in the case of the Internet community, there is experience in standards setting going back over twenty years to draw on for a course of conduct. There is some legal precedent for this position. For example, in a case decided in 1990, Effects Associates, Inc. v. Larry Cohen, Larco Productions and New World Entertainment, 1990 Copyright Law Decisions ¶26,605 (9th Cir. 1990), the Court of Appeals found that a filmmaker was entitled to use certain special effects footage without a written or oral license from the copyright owner, because the owner's conduct gave the filmmaker and the production company implied, nonexclusive licenses to incorporate the special effects footage into the film entitled "The Stuff," and to distribute the film.

You may ask: why, if it is possible to recognize the grant of an implied, nonexclusive license, is there a need for an e-mail confirmation? The need may arise when an Editor is asked to authorize other persons to reproduce, transmit or make certain other uses of the documents.

There may be situations that require some evidence in the form of an e-mail confirmation. For example, a company may wish to reproduce certain RFCs on a CD-ROM. The company comes to the RFC Editor and asks permission to reproduce an RFC for profit. The first reaction: it's an historical document, certainly. However, the company's legal counsel wants to see some evidence that the RFC Editor is authorized to grant such permission. If there is no authorization on file, it may be difficult to convince the company that a nonexclusive license covers the situation. There may also be situations where, if they go back to the originator of the document, the originator may place restrictions on further dissemination. While a nonexclusive license is not an ideal approach, it's certainly adequate, particularly where a document is circulated among the Internet community for information purposes.

2. Standards Track Specifications: As outlined in the 1991 "Internet Standards Process" draft, a specification that may lead to an Internet Standard may originate from at least three different sources: (a) an IAB-sponsored effort (typically an IETF Working Group), (b) independent activity by individuals, or (c) an external organisation. Internet Draft, at 9 (12/91). In the second and third cases, the work may be tightly integrated with the work of an existing IETF Working Group, or it may be offered for standardisation without prior IETF involvement. I understand that, in most cases, a specification resulting from an effort that took place outside of an IETF Working Group context will be submitted to an appropriate Working Group for evaluation and refinement; and, if necessary, an appropriate Working Group will be created.

In the case of standards track publications, it is advisable to obtain written assurances from all contributors to cover possible copyright claims, patents and patent applications, and other intellectual property rights. This should be done at the time an RFC is entered on the standards track at the Proposed Standard level, or otherwise incorporated in a specification that is proposed or adopted as an Internet standard. This would also include a written and signed confirmation of the nonexclusive license discussed in connection with non-standards track publications, or the grant of such a license in the event it has not already been obtained from the rightsholder. Again, I will limit my remarks to any copyright claims.

Without clear understandings among the parties prior to the adoption of an Internet standard, it is always possible for an individual claiming copyright in say a piece of computer code embodied in a specification to block or delay the dissemination of the standard. I know these kinds of problems have not occurred very often, if at all, in the past, but they could arise in the future. It is important to defend against such potential abuses in order to continue the grand tradition of openness and accessibility fostered over the years in connection with Internet standards. I noted that there is already a proposed requirement in the Internet Standards memo circulated in December 1991 that: in order for a vendor-proprietary specification to be incorporated within the Internet standards process, the proprietor must agree in writing that licenses will be available on a non-discriminatory basis and at a reasonable cost.

Let me give you a specific example to help clarify the need for such written assurances on the availability of licenses in connection with possible copyrights, patents and patent applications or other intellectual property rights, or assurances that there are no such claims.

I understand that there is something called a Management Information Base or MIB that consists of descriptions in a formal language of managed objects. The descriptions may be compiled and incorporated in software. Documentation in standards specifications are sometimes used in that way. Absent a written assurance, the author of a MIB could prevent such uses and inhibit the practice of the standard.

3. Working Groups: A word about the copyright status of Working Group members. It is evident that Working Groups in various forms and at different times play a major role in advancing a specification along the Internet standards track. Therefore, it is important to clarify the copyright status of the work performed by such Groups. Assuming that a nonexclusive license is obtained from the copyright owner in any specification to reproduce, distribute, transmit, display and prepare derivative works, and to authorize others to do so, members of a Working Group could be permitted to dissect a particular document and reformulate it into a new derivative work. The question would then arise: who has a claim to copyright in this derivative work, particularly where the work takes the form of a proposed Internet standard specification?

In order to encourage the fair and equitable distribution of a specification prepared by a Working Group, and to prevent a single member of the Group from holding up the standards process, it is important to require all members of the Group who actually contribute to the drafting of a document, to provide written assurances with respect to intellectual property rights, including the grant of a nonexclusive license to the coordinating body along the lines of the licenses required for other standards track publications. There would also have to be some provision for clearing rights in works of third parties that may be incorporated in any work produced by the Group.

Why would a nonexclusive license be advisable; and what form would it take? Under U.S. law, material produced by a member of a Working Group may be considered a contribution to a collective work, and, if there is a written agreement that the work is to be considered a "work-for-hire," the hiring organization may be deemed to be the author and owner of copyright. However, it is not certain whether such an agreement would be enforceable in other countries. Many countries, such as France, would not generally recognize that the commissioning party is the author, especially where the agreement is entered into in France. To avoid such uncertainty, it is advisable to obtain a nonexclusive license from members of each Working Group who actually participate in the drafting of a standards document. This would not include persons merely attending Working Group meetings. To avoid possible legal challenges in certain jurisdictions, the agreement must be in writing and signed by the parties. The Editor of the Group could be charged with obtaining such releases.

Since a signed writing is advisable in this case, I suggest that some thought be given to using the emerging digital signature technology, which I know you are developing, for authenticating transactions. While I understand that this may take a while to implement, it is not too early to start planning for it. The Internet community would appear to be ideally suited for experimentation with on-line clearance of rights using digital signatures. Of course, it would also be necessary to seek a change in the regulations of the Copyright

Office to recognize and accept such signatures in lieu of actual handwritten signatures for recordation purposes.

While members of a Working Group would be required to provide written assurances on intellectual property rights, the members of the Group would be assured that they would be listed as contributors to the particular specification developed by the Group. The principal Editor could also be recognized as such.

4. Notices concerning intellectual property: When a standard (either proposed, adopted or maintained) is distributed, transmitted over computer networks or otherwise made available to users, it should bear a notice at or near the title of the standards document indicating who is claiming copyright, and authorizing reproduction, distribution or transmission solely for informational purposes, provided that appropriate credit is given to any rightsholders. Depending on the nature of the specification, there may also be a statement relating to the written assurances with respect to any copyright claims, patents and patent applications, or other intellectual property rights, to the effect that any party will be able to obtain, under reasonable, nondiscriminatory terms, the right and license to practice the standard covered by such rights. The exact wording may vary from case-to-case, but the general intent would remain constant.

As noted in the description of the Internet Standards Process: "The process works because the IETF Working Groups display a spirit of cooperation as well as a high degree of technical maturity; most IETF members agree that the greatest benefit for all members of the Internet community results from the cooperative development of technically superior protocols and services." Internet Draft, at 3 (12/91). Allowing private intellectual property rights to undermine this effort would have unfortunate consequences for the entire Internet community. We should all strive to make this work in the smoothest possible way.

6.3 First IETF Internet Audiocast

Presented by Steve Casner/ISI and Steve Deering/Xerox PARC

First IETF Internet Audiocast

Steves Casner & Deering

- multicast, interactive audio from plenaries and some WGs, across the Internet to remote participants spanning 15.5 time zones
- demonstration / transfer of technology developed / tested in the DARTnet research testbed
  - — DARTnet multicast backbone initially extended to include researchers at Mitre and UMass
  - — More tunnels added for AVT working group meeting
  - — Allison Mankin asked, "Why not transmit the IETF meetings?"

vat — Visual Audio Tool,
by Van Jacobson & Steve McCanne, LBL

one of four interoperable audio-conferencing programs used in DARTnet (others are from ISI, MIT, UMass)

IP	UDP	VT	PCM Voice Samples
20	8	4	180

audio data: 64 Kbit/sec PCM,
from Sun Sparcstation audio device

packets: 75 Kbit/sec = 44.4 pkt/sec
(1 pkt every 22.5 msec)

sent only when input > silence threshold

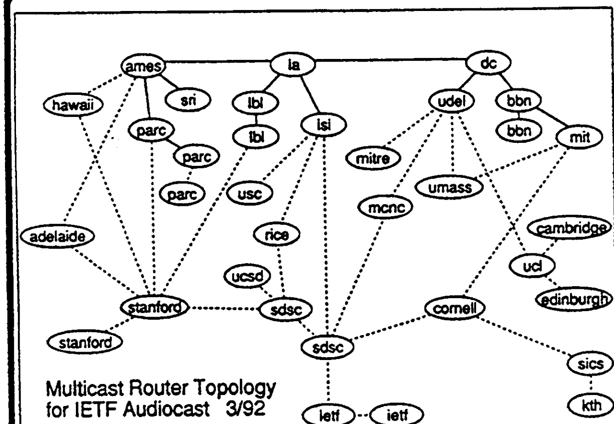

Multicast Router Topology
for IETF Audiocast 3/92

34 routers, linking 40 subnets

You too can play!

- vat from ftp.ee.lbl.gov : vat.tar.Z
- IP multicast extensions to BSD/SunOS/Ultrix from gregorio.stanford.edu : vmtp-ip/ipmulticast.*
- microphone from Radio Shack

Still to come:

- lower-bandwidth audio
- video, images, shared "whiteboards"
- resource management in routers
- multiple, interoperable multicast routing protocols

Thanks To:

Allison Mankin

Paul Love, Tom Hutton, and rest of SDSC crew

Phill Gross, Steve Coya, Megan Davies

Van Jacobson

Walt Prue, Milo Medin, Jeff Burgan

DARPA & the DARTnet research community

The remote vat participants / guinea pigs

VAT Participants

hutton@opus.sdsc.edu
rrosales@bbn.com
yhc@happy.concert.net
mmconf@guy.isi.edu
kink@fride.rice.edu
klemets@sam.sics.se
van@ee.lbl.gov
mankin@192.80.55.61
iwakeman@kant.cs.ucl.ac.uk
jon@kant.cs.ucl.ac.uk
hgshulz@sparc1.ecs.umass.edu
peter@sarapis.sics.se
lixia@dartvader.parc.xerox.com
beers@zooey.cit.cornell.edu
steve@garuda.sics.se
uffe@moray.tos.kth.se

perkins@beauregard.udel.edu
swinehar@ptarmigan.parc.xerox.com
clynn@clynn.bbn.com
lwei@bigsur.usc.edu
wbe@crystal.bbn.com
mccanne@horse.ee.lbl.gov
hackett@odin.arch.adelaide.edu.au
wkd@dorsai.net.hawaii.edu
saavedra@pismo.usc.edu
davidc@sirius.net.hawaii.edu
boonkong@bigsur.ucl.edu
floyd@owl.ee.lbl.gov
adrian@kate.internode.com.au
steig@dragonlance.stanford.edu
frederic@ladora.parc.xerox.com
ietf@oingo.ietf.cerf.net

6.4 Traffic Characterization in Wide Area Networks

Presented by Peter Danzig/USC

Data Analysis

in

Wide Area Network

Environments

Host Centric Environment

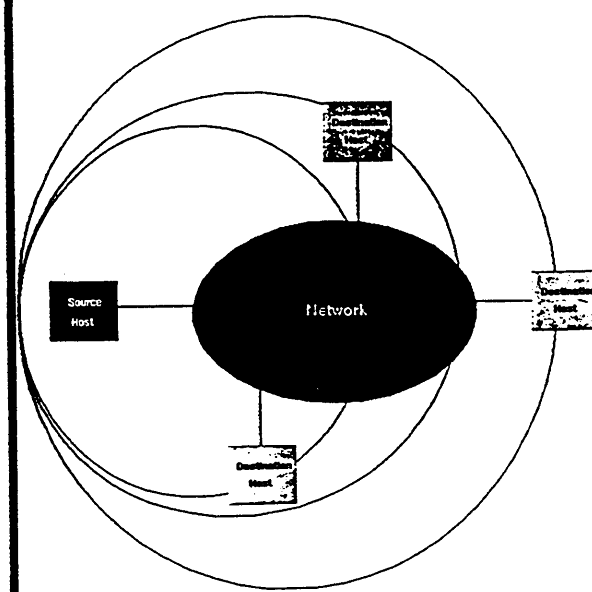

Network Centric Environment

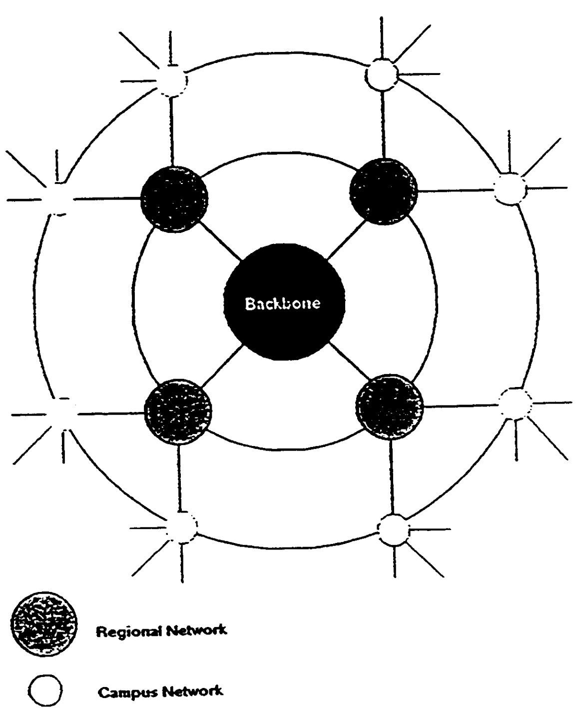

Dimensions of Analysis

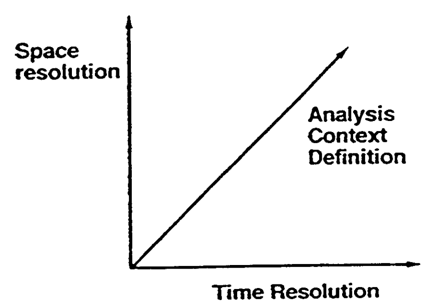

Space: Application – User – Host – Network – Domain – External Interface – Service Provider – (Core) network – Internet

Time: Year – Month – Day – Hour – Second – Subsecond – Packet

Analysis Context Definition: Applications – Flows – Traffic distribution – Performance

Initial Areas of Investigation:

- resource consumption and latencies
- degradation under resource starvation
- performance metrics
  - end–user
  - network
- levels of granularity
  - analysis
  - sampling
- application distribution and flows
- geographic traffic distribution
- routing methodology considerations
- determination of bandwidth requirements

Instrumentation

Data Collection

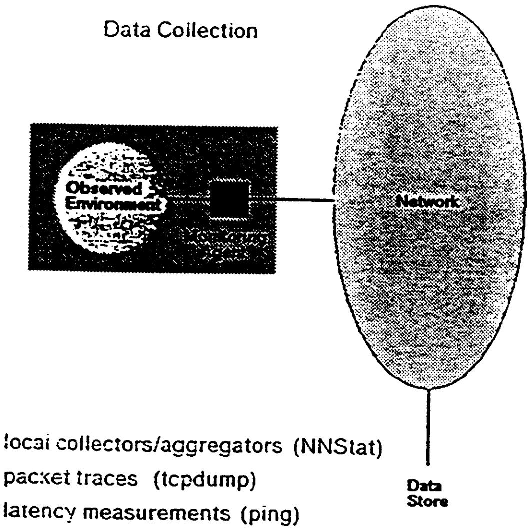

- local collectors/aggregators (NNStat)
- packet traces (tcpdump)
- latency measurements (ping)
- local counters (SNMP queryable)
- specific kernel modifications

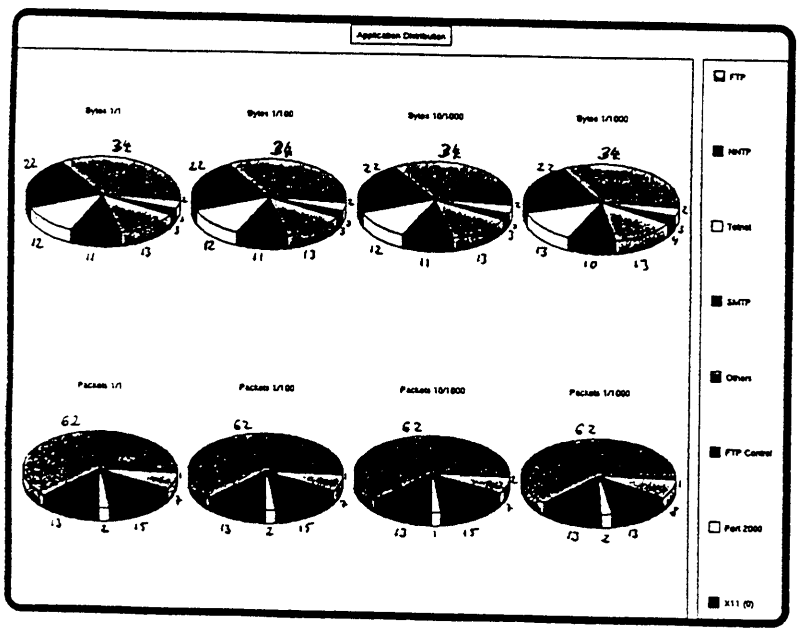

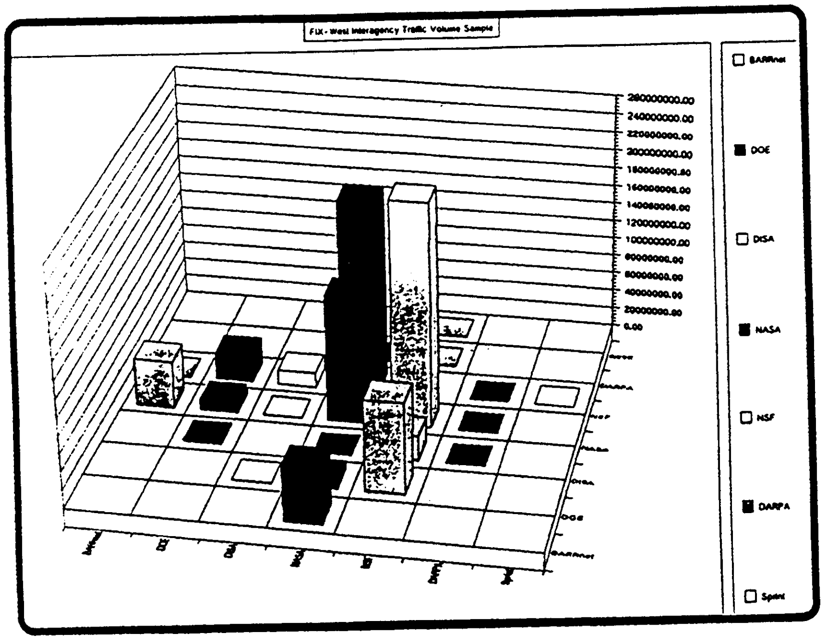
FIX-West Interagency Traffic Volume Sample
BARRnet
DOE
DISA
NASA
NSF
DARPA
Sprint

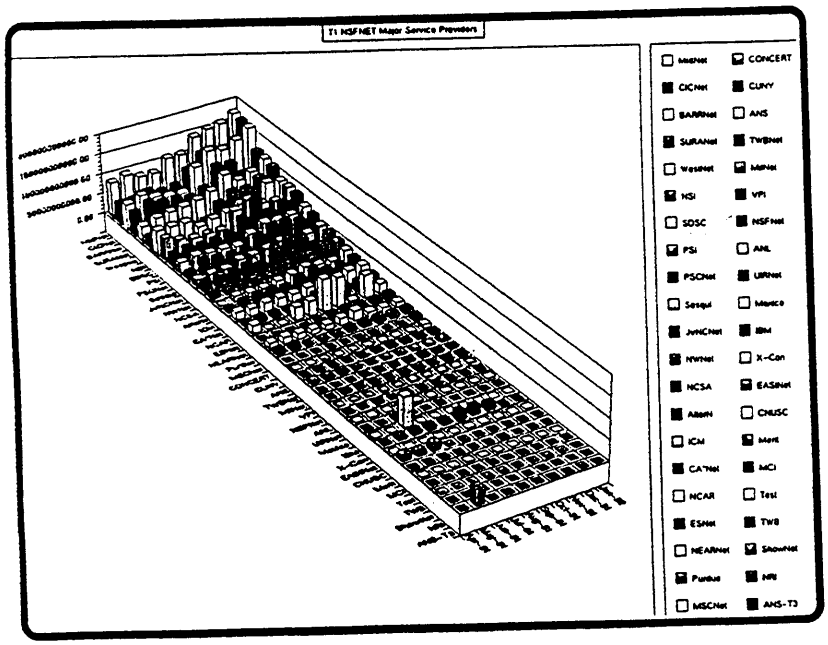
T1 NSFNET Major Service Providers
MidNet
CONCERT
CICNet
CUNY
BARRNet
ANS
SURANet
TWBNet
WestNet
MRNet
NSI
VPI
SDSC
NSFNet
PSI
ANL
PSCNet
Sesqui
Mexico
JvNCNet
IBM
NWNet
X-Con
NCSA
EASINet
CNUSC
ICM
Merit
CA*Net
MCI
NCAR
Test
ESNet
TWB
NEARNet
ShowNet
Purdue
MSCNet
ANS-T3

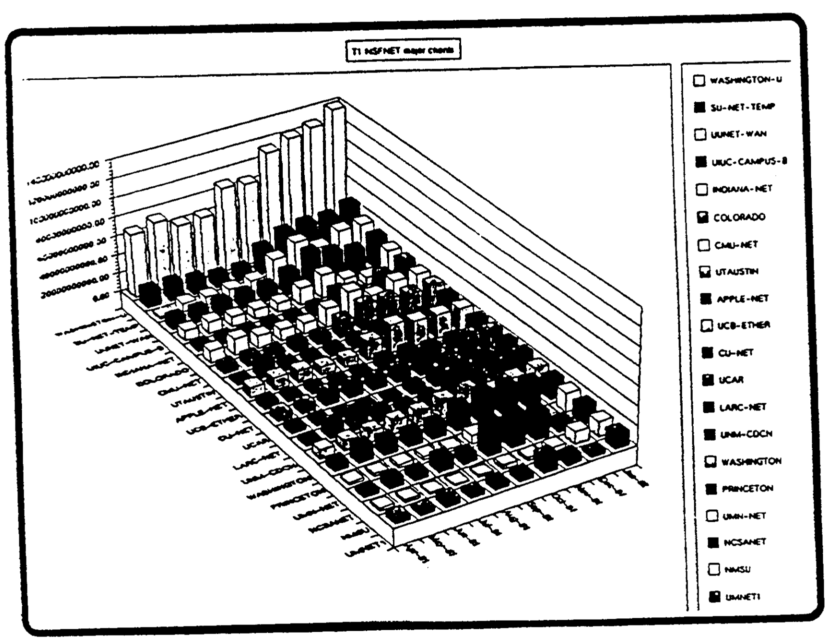
T1 NSFNET major clients
WASHINGTON-U
SU-NET-TEMP
UUNET-WAN
UIUC-CAMPUS-B
INDIANA-NET
COLORADO
CMU-NET
UTAUSTIN
APPLE-NET
UCB-ETHER
CU-NET
UCAR
LARC-NET
UNM-CDCN
WASHINGTON
PRINCETON
UMN-NET
NCSANET
NMSU
UMNET1

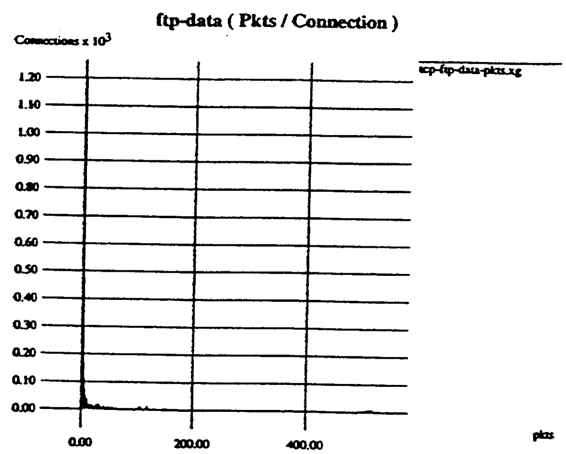
ftp-data (Pkts / Connection)
Connections x 10^3
tcp-ftp-data-pkts.xg
pkts

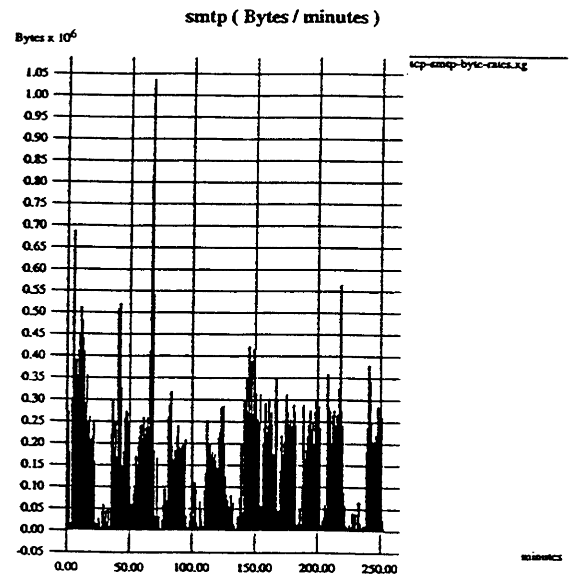
smtp (Bytes / minutes)
Bytes x 10^6
tcp-smtp-byte-rates.xg
minutes

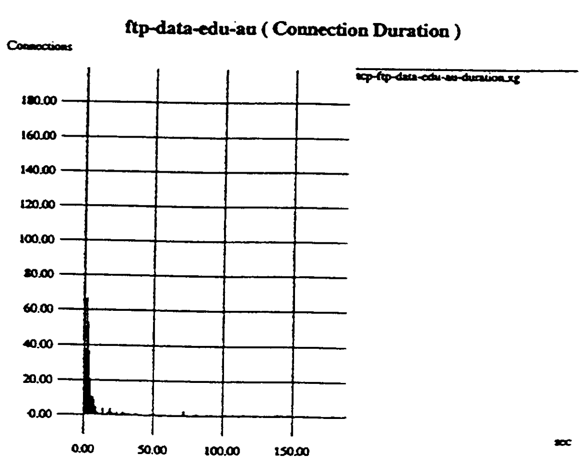
ftp-data-edu-au (Connection Duration)
Connections
tcp-ftp-data-edu-au-duration.xg
sec

6.5 Multi-Media Mail Extensions

Presented by Greg Vaudreuil/CNRI

RFC 822 Message Extensions
Greg Vaudreuil/ CNRI

Overview

Working Group Goals

- Multi-Lingual/ Multi-Character Set
- Standard Attachments
- Multi-Media
- Rich Text Format
- Within Current RFC822/ 821 Framework

Scope

- Message Body
- Headers

Overview

The Problems and Constraints

- Heterogeneous e-mail "I"nternet using RFC822 based mail
- Multiple transport environments
  - SMTP
  - Bitnet
  - UUCP
  - NNTP
  - Commercial carriers

Overview

The Problems and Constraints (cont)

- Multiple mail relays/multiple mail readers
  - Name your favorite
  - Header stripping, reordering, and reformatting
  - Line wrapping, line truncating

Overview

The Problems and Constraints (cont)

- Multiple character sets/ lauguages
  - US ASCII
  - ISO 646 national varient ASCII
  - Many EBCDIC's
  - ISO 2022 / ECMA registry
  - IBM-PC codepages
  - ISO DIS2 10646
  - Unicode
  - Mnemonic

Corporation for National Research Initiatives

Overview

Mail Today

```
From:    Keld J|rn Simonsen <keld@dkuug.dk>
To:      Patrik F{ltstr|m <paf@nada.kth.se>
Cc:  Greg Vaudreuil <gvaudre@nri.reston.va.us>
--------

Try this neat routine....

begin 664 neat_thing
L4V\L('E0=2!W97)E(&)O<F5D(&5
    N;W5G:"!T;R!D96-09&4@=&AI<R!E:#\
end
```

Corporation for National Research Initiatives

The Simple Solution

Content-Type

The content-type header labels the information contained in the body part. It is based on RFC934 content-type, with the addition of subtype values.

Content-Transfer-Encoding

The content-transfer-encoding header identifies the transformation required to make the body part transportable over the choosen transport path.

Corporation for National Research Initiatives

Encoding for Transport

`content-encoding: quoted-printable`

Default Types

- 7bit
- 8bit
- binary

7 Bit Encoding Types

- quoted-printable
- base64

* Encoding is orthogonal to content-type

Corporation for National Research Initiatives

Encoding for Transport

Quoted-Printable

- For text-like content types.
- Designed to preserve (some) readability.
- Escape the 8 bit bytes.
- 3 to 1 expansion on high order bytes.

```
content-type: text/plain;
   charset=8859-1
content-transfer-encoding:
   Quoted-Printable

This is a bodypart encoded in quoted
printable. Can you read it Keld
J=F8rn Simonsen?
```

Corporation for National Research Initiatives

Encoding for Transport

Base64

- Designed for arbitrary binary.
- 3/4 encoding scheme
- Based on RFC 1113 - PEM.

```
Content-type: image/xwd
Content-Transfer-Encoding: base64
```


```
AAAAawAAAAcAAAACAAAAAQAAAXcAAAB6AAAA
   AAAAAAEAAAAIAAAAAQAAAAgAAAABAAAALwAA
AAAAAAAAAAAAAAAAAAAAAAAABAAAAAgAAAAIA
   AAF3AAAAegAAAAAAAAAAAAAAAHJhc3RlcgAA
...
```


Corporation for National Research Initiatives

Basic Content Types

Seven Basic Content Types

- text
- multipart
- message
- application
- image
- audio
- video
- x- private values

Corporation for National Research Initiatives

Basic Content Types

text

No special software is required to get the full meaning of the text, aside from support for the indicated character set.

Subtypes

- text/plain
- text/richtext

Corporation for National Research Initiatives

Basic Content Types

text/richtext

Hello, ietf-822'ers!

I'm sending this message out because I believe that I have finished creating a modified version of *Andrew* that is *MIME*-smart.

```
content-type: text/richtext
content-transfer-encoding:
   Quoted-printable
```

```
<bold><bigger>Hello, ietf-822'ers!
</bigger></bold>  I'm sending this message
out because I believe that I have finished
creating a modified version of
<italic>Andrew</italic> that is
<bold><italic>MIME</italic></bold>-smart.
```

Corporation for National Research Initiatives

Character Sets

There is not just one!

Current Usage

- National varient ASCII (ISO 646)
- ISO 8859-1 (8 Bit)
  - Raw 8 Bit SMTP
  - Text-Hex
- ISO 2022 codepage switching

Approach Taken

- No global sets widely available.
- Mimimum number of "common" formats needed.
- Only well documented character sets.

Corporation for National Research Initiatives

Character Sets

Initial Character Sets

- 8859-(<10)
  (Replacement for ISO 646s)
- US-ASCII (current practice)

Coming Soon

- ISO 2022
  (profiled for Asian use)
- Mnemonic

Corporation for National Research Initiatives

Basic Content Types

multipart

A body which consisting of multiple body parts, each of independent data types. Multipart is fully recursive.

Subtypes

- multipart/mixed
- multipart/digest
- multipart/alternative
- multipart/parallel

Corporation for National Research Initiatives

MultiPart

Multipart Format

- Based on RFC 934.
- Changed Boundary definition to eliminate quoting.

```
content-type: Multipart/Mixed;
   boundary="Interpart_Boundary_1"

--Interpart_Boundary_1
content-type: Text/plain; charset=us-ascii
content-transfer-encoding: 7bit

This is part 1.
--Interpart_Boundary_1
content-type: Text/plain; charset=us-ascii
content-transfer-encoding: 7bit

This is part 2.
--Interpart_Boundary_1--
```

Corporation for National Research Initiatives

Basic Content Types

message

A body of content-type message is itself a fully formatted RFC 822 conformant message which may contain its own different content-type header field.

Subtypes

- message/822
- message/partial (Allows split messages)
- message/external-body (Body pointer)

```
content-type: Message/822
content-transfer-encoding: 7bit

to: gregv@nri.reston.va.us
Subject: Hi!

This new MIME format makes it easy
to forward messages!
```

Corporation for National Research Initiatives

Basic Content Types

application

Raw data, typically uninterpreted binary data or information to be processed by a mail-based application.

Subtypes

- application/oda
- application/postscript
- application/octet-stream

Corporation for National Research Initiatives

Basic Content Types

image

Data which requires a graphics display device (such as a graphical terminal, a printer, or a FAX machine) to view the information.

Subtypes

- image/gif
- image/jpeg

Corporation for National Research Initiatives

Basic Content Types

audio

Data which requires a sound output device.

Subtype

- audio/basic
  - 8000 hz
  - 8 bit
  - 1 channel
  - U-law companding

Corporation for National Research Initiatives

Basic Content Types

video

Video requires the capability to display moving images, typically including specialized hardware and software.

Subtype

- video/mpeg

Corporation for National Research Initiatives

Message Headers

Goals

- No modification to transport software.
- No changes required to old-style mail readers to work as they did.
- New headers should be as usable as possible on old-style mail readers.
- As robust as possible through the current email "Internet".

Corporation for National Research Initiatives

Message Headers

Approaches

- Encoded variable
- Real-header (real-from, real-subject)
- Header charset/header encoding
- Encoded word

Corporation for National Research Initiatives

Message Headers

Example Header Lines

```
To: Keld Jørn Simonsen <keld@dkuug.dk>,
    Keith Moore <moore@cs.utk.edu>
From: André Pirard  <PIRARD@vm1.ulg.ac.be>
-----------
```

Corporation for National Research Initiatives

Message Headers

Encoded Variable

```
To: $1 <keld@dkuug.dk>,
    Keith Moore <moore@cs.utk.edu>
From: $Andr?_Pirard <PIRARD@vm1.ulg.ac.be>
Encoded-variable:
    $1:8859-1:q-p:Keld J=F8rn Simonsen
Encoded-variable:
    $Andr?_Pirard:8859-1:q-p:Andr=E9 Pirard
-----------
```

- Multiple character sets in each header line.
- Interline dependency

Corporation for National Research Initiatives

Message Headers

Real-headers

```
Real-Encoding: quoted-printable
To: <keld@dkuug.dk>, <moore@cs.utk.edu>
From: <PIRARD@vm1.ulg.ac.be>
Real-from: Andr=E9 Pirard \\
     <PIRARD@vm1.ulg.ac.be>
Real-to: Keld J=F8rn Simonsen \\
    <keld@dkuug.dk>,
     Keith Moore <moore@cs.utk.edu>
-----------
```

- Preserves semantics of actual headers once decoded.
- Single character set per header
- Inter-header dependency

Corporation for National Research Initiatives

Message Headers

Header-Charset-Encoding

```
Header-charset: 8859-1
Header-encoding: quoted-printable
To: Keld J=F8rn Simonsen <keld@dkuug.dk>,
    Keith Moore <moore@cs.utk.edu>
From: Andr=E9 Pirard  <PIRARD@vm1.ulg.ac.be>
-----------
```

- Pleasing appearance to old mail readers.
- Inter-header dependency.
- Only one character set allowed per message header.

Corporation for National Research Initiatives

Message Headers

Encoded Word

```
To: =?ISO-8859-1?Q?Keld_J=F8rn_Simonsen?=
     <keld@dkuug.dk>,
     Keith Moore <moore@cs.utk.edu>
From:=?ISO-8859-1?Q?Andr=E9_?= Pirard
     <PIRARD@vm1.ulg.ac.be>
-----------
```

- Word, character set and encoding is an atomic entity to current RFC 822 parsers.
- Robust against line wrapping and most header processing software.
- Each word can be it's own character set.
- No inter-header dependency.
- Somewhat readable on old Mail Readers.

Corporation for National Research Initiatives

6.6 NSFnet Update

Presented by Mark Knopper/Merit and Jordan Becker/ANS

ANSNET/NSFNET Backbone: Activities and Status

Mark Knopper, Merit
Jordan Becker, ANS
March 20, 1992

- T1 -> T3 traffic cutover
- Traffic statistics
- T3 backbone status
- Phase III technology upgrade
- T1 backbone status, problems

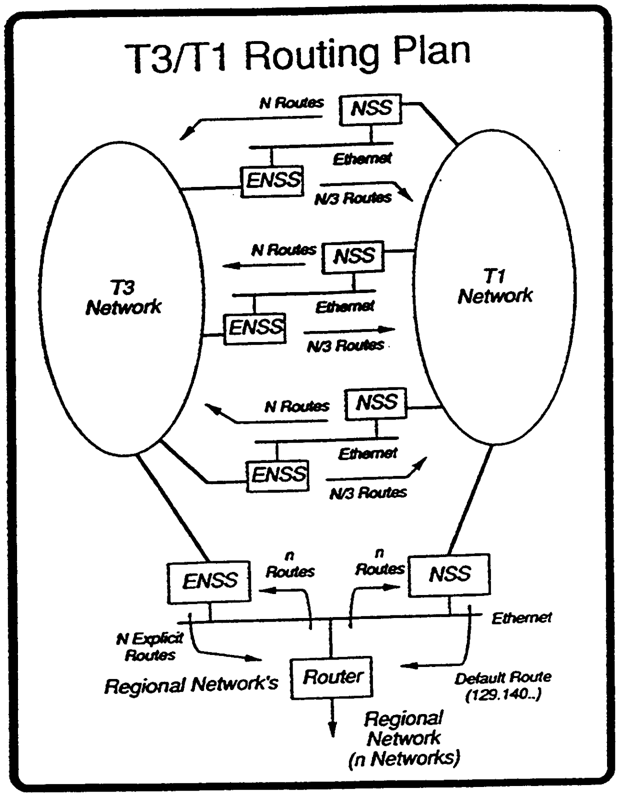

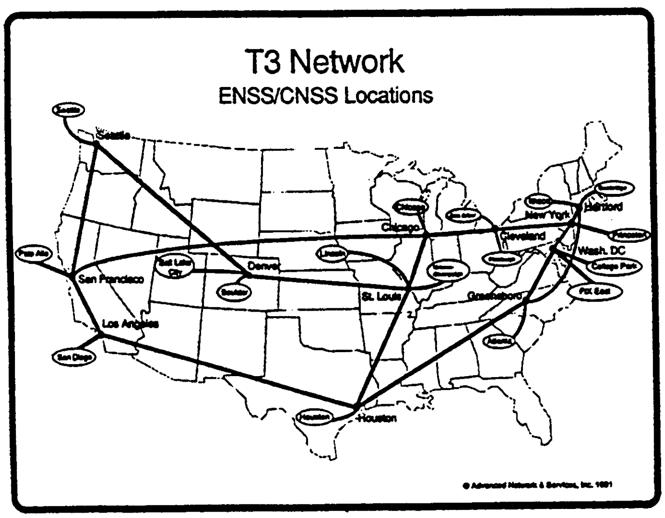

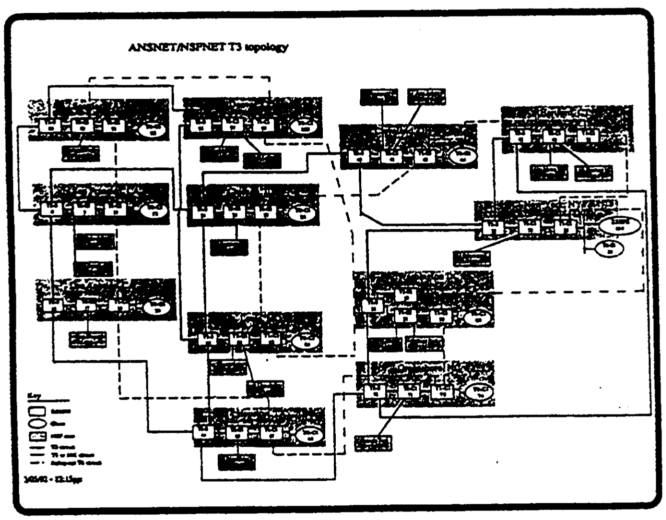

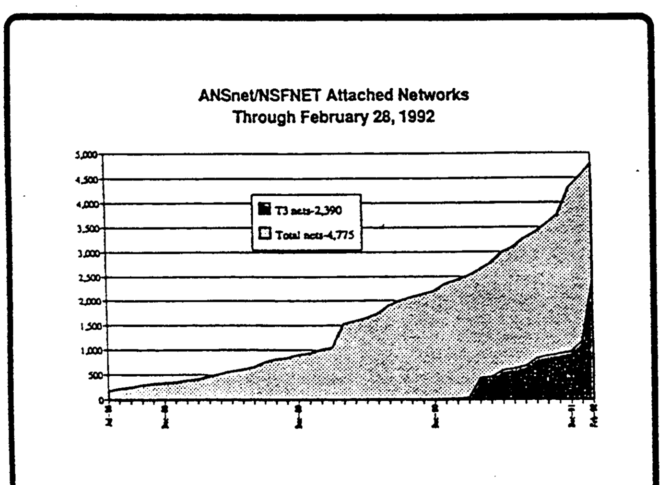

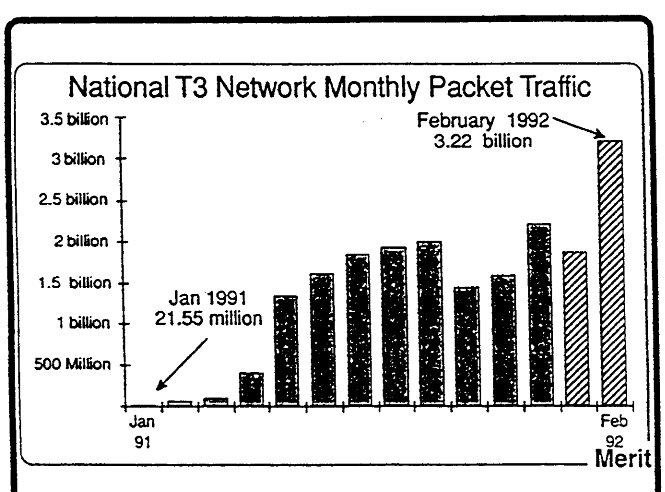

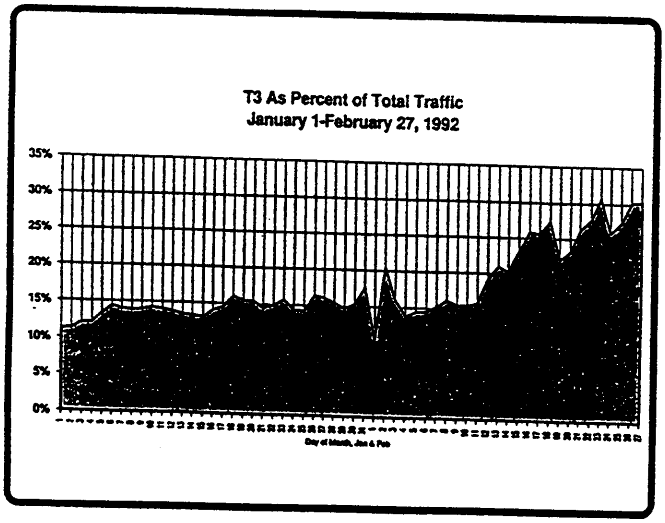

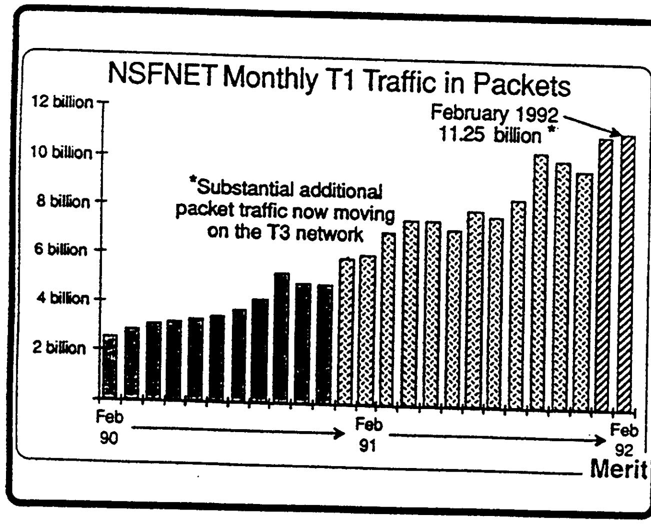

NSF Midlevel T3 Traffic Migration Status

Network	Location	T3 Status
Merit	Ann Arbor	T3
NEARNet	Boston	T3
BARRNet	Palo Alto	T3
CERFNet	San Diego	T3
SURANet	College Park	T3
SURANet	Atlanta	T3
NYSERNet	Ithaca	T3
SESQUINet	Houston	T3
NASA/FIX-East	College Park	T3
JVNCNet	Princeton	March T.B.D
UIUC	Champaign	3/18
MIDNet	Lincoln	3/25
WESTNet-E	NCAR	4/1
CICNet	Ann Arbor	4/7
WESTNet-W	Salt Lake City	Not Yet Scheduled
NorthWestNet	Seattle	Not Yet Scheduled
Argonne	Chicago	Not Yet Scheduled
NASA-Ames	San Jose	Not Yet Scheduled

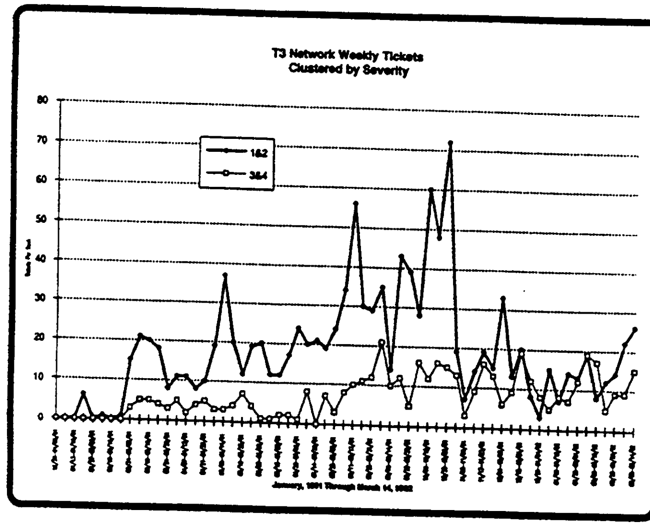

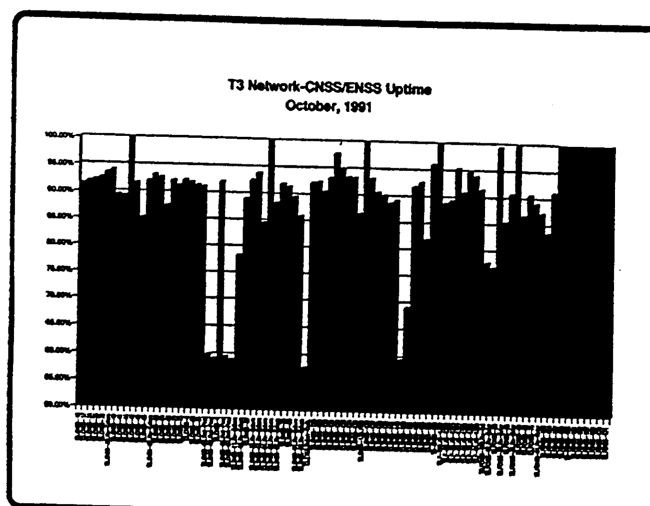

T3 Network-CNSS/ENSS Uptime
November, 1991

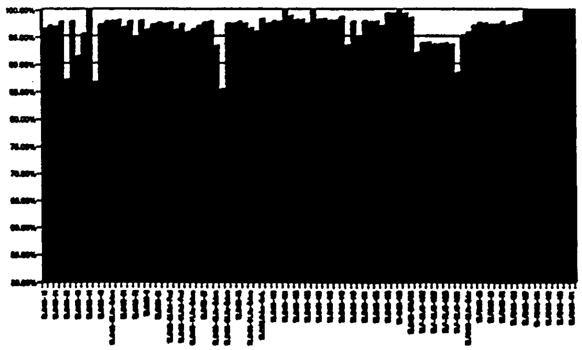

T3 Network-CNSS/ENSS Uptime
January, 1992

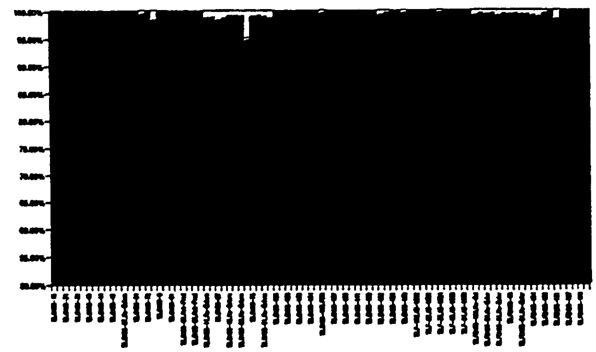

T3 Network-CNSS/ENSS Uptime
February, 1992

T3 Network Status

- o T3 Routing Software Enhanced
  - Added Digital Signature Based on MD4 to Routing Software
  - Improved BGP Route Aggregation to Reduce Size of Updates
  - Optimal T3 Routing for Same AS at Multiple Interconnects
  - Load Splitting At Interconnects Through External Metrics
- o Installed 3rd T1/T3 Interconnect Gateway at Houston

Outstanding Problems

- o T1/Ethernet Interface Freeze Bug Fixes
- o Intermittent T3 Interface Freeze (Black Link) Problem Remains
  - Safety-Net T1 Links Help Avoid Black Link Problems
- o FDDI Card Freeze Problem

Network Status Summary

- T3 Network Stabilized
  - Network is Very Reliable
  - Significant Traffic Load Cutover from T1 -> T3 Network
  - Remaining Intermittent Black Links Manageable So Far
  - Traded Performance for Stability
- T1 Network Stability Improving with More Engineering Focus
  - Bug Fixes, Performance Enhancements
- RS/960 Adapter Upgrade To Begin on 4/10
  - Improved Reliability & Performance
- Other Software & Infrastructure Enhancements to Follow

Other Proposed Network Changes in 1992

Proposed Software Enhancements

- BGP3 Support On T3 Network
- Connectionless OSI Network Protocol Support (CLNP, ES-IS)
- Migrate to Common IGP Within T3 System
  - Intra-Domain IS 10589 Routing (IS-IS)
- Inter-Domain Routing (IDRP)
- Dual-mode Routing
  - Deploy Dual IS-IS and Dual IDRP (OSI and IP routing)
- Routing Software Being Implemented in GATED & Cisco

Network Infrastructure Enhancements

- RS/960 T3 Interface Upgrade Begins
- IBM FDDI Adapter Upgrade (with AIX 3.2 Software Upgrade)
- Cisco T3 & FDDI Interface Support
- T3 ENSS Backup Support via Redundant ENSS Connections
- Dismantling of T1 Backbone
  - Requires Full Cutover to T3, OSI Support, ENSS Backup

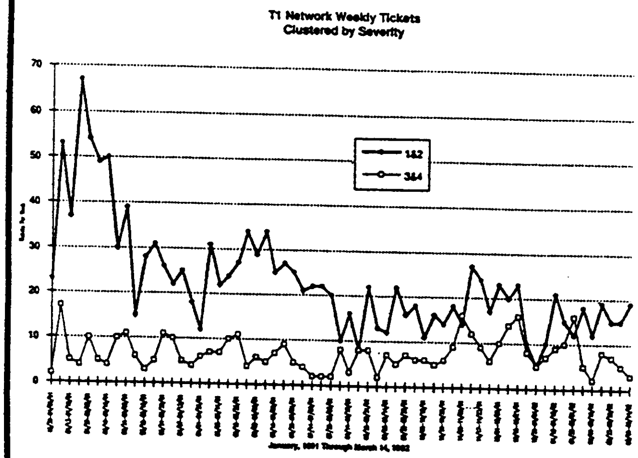

T1 Network Stability Status

- Congestion Related Problems
- LSP Packet Size Exceeded
- CPU Starvation on RCP RT/PC Nodes
- ICMP Network Unreachable Messages
- Router Response to Circuit Flapping
- Remaining T1 Network Stability Problems

Phase-III T3 Network Plan Status

RS/960 "Smart" Adapter

- More Reliable T3 Adapter for RS/6000 Router
- Major Performance Improvement Over Existing Adapters
- 5 Adapters Per RS/6000 Router Possible

RS/960 Test Plan Goals & Strategy

- Now Deployed on T3 Testnet
- Routing Real User Generated Traffic onto Test Network
- Stressing Testnet to Point of Collapse
- Generate Test Outages and Measure Network Response
- Test Routing Software for Function, Performance, Scalability
- 28 Tests with 2-3 Test Cases Being Run

DS3 Transmission Facility Monitoring

- New MCI DS3 Monitoring Units Being Installed
- C-bit Parity Standard Supported on Routers

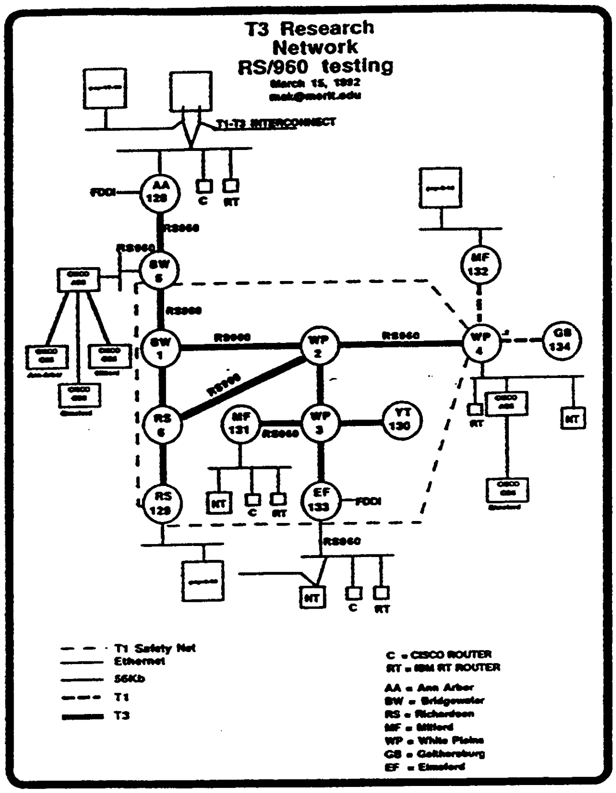
T3 Research
Network
RS/960 testing
T1-T3 INTERCONNECT
AA
128
FDDI
C
RT
RS960
BW
5
BW
1
WP
2
WP
4
MF
132
GB
134
RS
6
MF
131
WP
3
YT
130
RS
129
EF
133
FDDI
RT
C
RT
T1 Safety Net
Ethernet
56Kb
T1
T3
C = CISCO ROUTER
RT = IBM RT ROUTER
AA = Ann Arbor
BW = Bridgewater
RS = Richardson
MF = Milford
WP = White Plains
GB = Gaithersburg
EF = Elmsford

6.7 NEARnet Presentation

Presented by Scott Bradner/Harvard

NEARnet & NSFnet

(& MERIT)

(& ANS)

NEARnet

- 121 members
- cream of the NE educational & research institutions
- high speed core (10MB)
- 24 hour usage
- self supporting
- liberal AUP

NEARnet Service Goals

- highly reliable network
- production level service
- user perception critical

NEARnet Service Requirements

- must have same goals in service providers
- immediate attention to reported problems
- understanding of importance of perception
- production level service

NEARnet & NSFNET

Oct 2 1989
- NEARnet responded to NSF RFP

May 1990
- awarded T3 ENSS

- 1st all-T3 connection because of traffic load etc.

- we were happy

NEARnet & NSFNET, contd.

April 1991
- ENSS installed

a bit later in 1991
- traffic started flowing over ENSS

early summer 1991
- we weren't so happy

late summer/early fall 1991
- "happy" did not quite describe the way we felt

IETF meeting, Nov 1991
- "interesting" discussion during MERIT report

Some historical data

MIT Network Monitor

- "pings" nodes on Internet

- tries every *T* sec
- if net-unreachable
  node marked *unreachable*
- if no response in *S* sec
  sends *N* pings at *R* sec intervals
  if no response - node marked *timeout*
- try again in *T* sec

node	*T*	*S*	*N*	*R*
t3peer.merit.edu	20	10	10	10
ihtfp-ether.mit.edu	8	3	3	4
nss.mit.edu	15	3	4	3
pnss.mit.edu	15	3	4	3
tn-ext-gw.tn.cornell.edu	20	10	10	10
su-a.barrnet.net	20	10	10	10

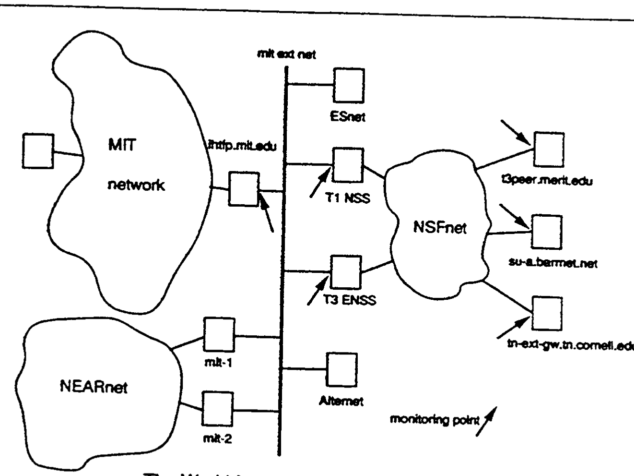

The World from the MIT Point of View

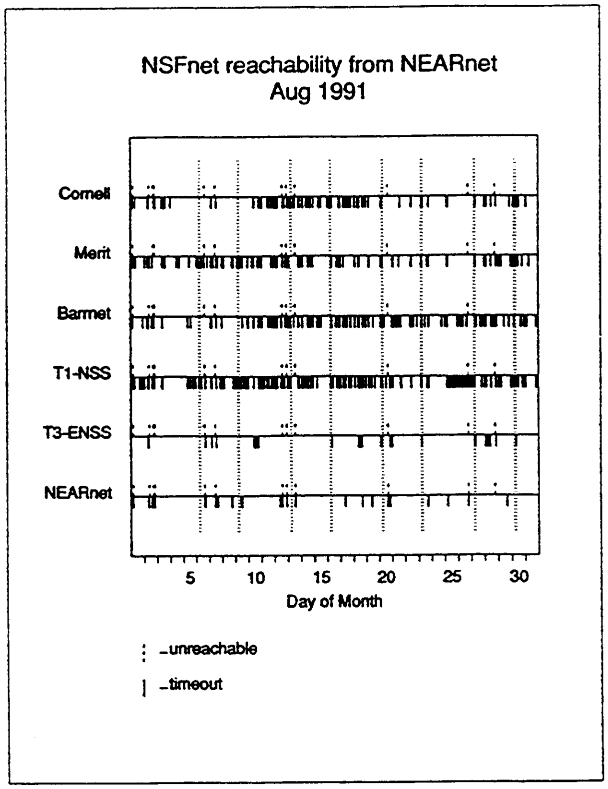
NSFnet reachability from NEARnet
Aug 1991
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

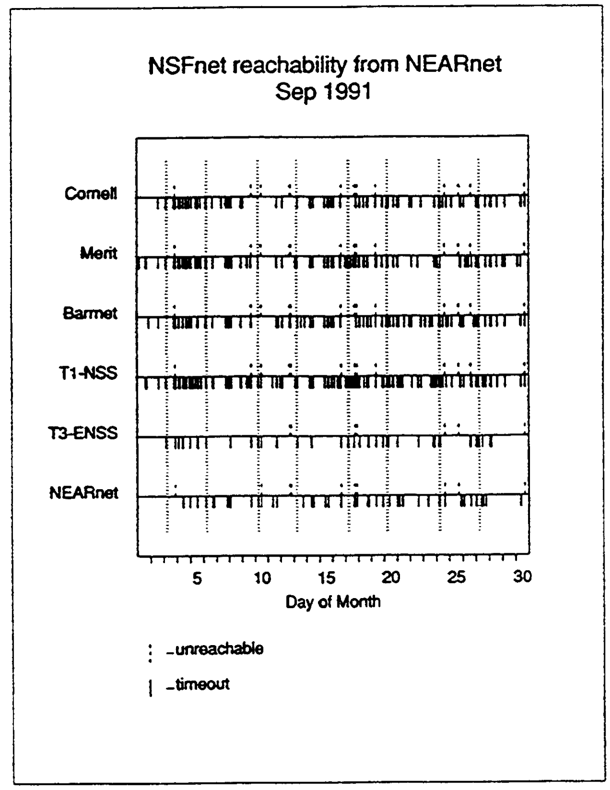
NSFnet reachability from NEARnet
Sep 1991
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

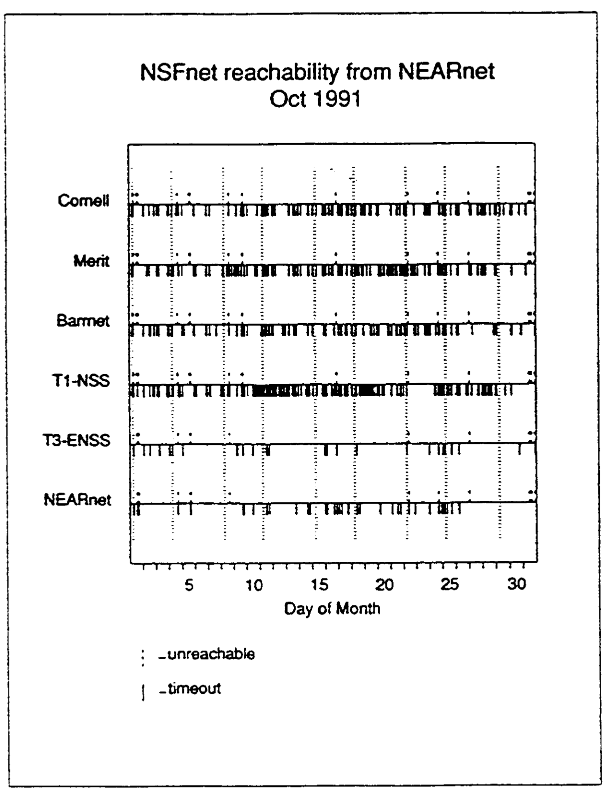
NSFnet reachability from NEARnet
Oct 1991
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

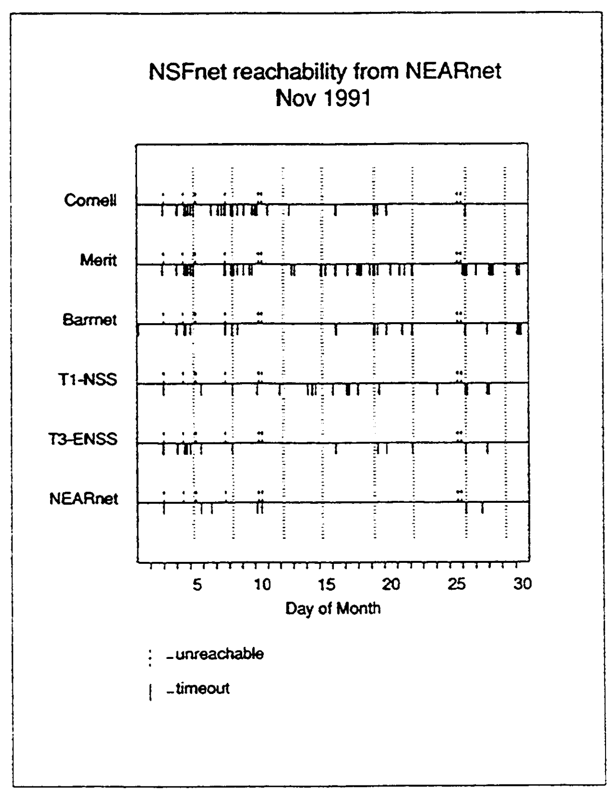
NSFnet reachability from NEARnet
Nov 1991
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

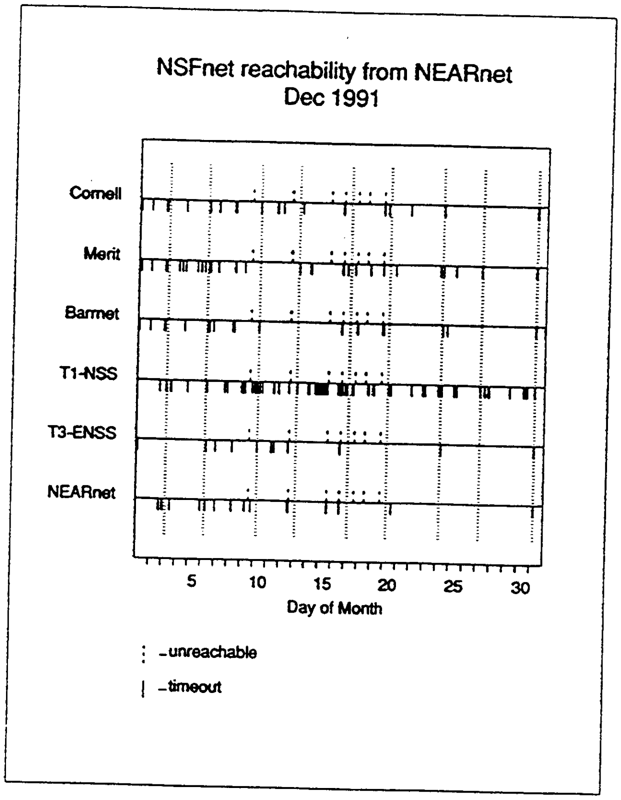
NSFnet reachability from NEARnet
Dec 1991
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

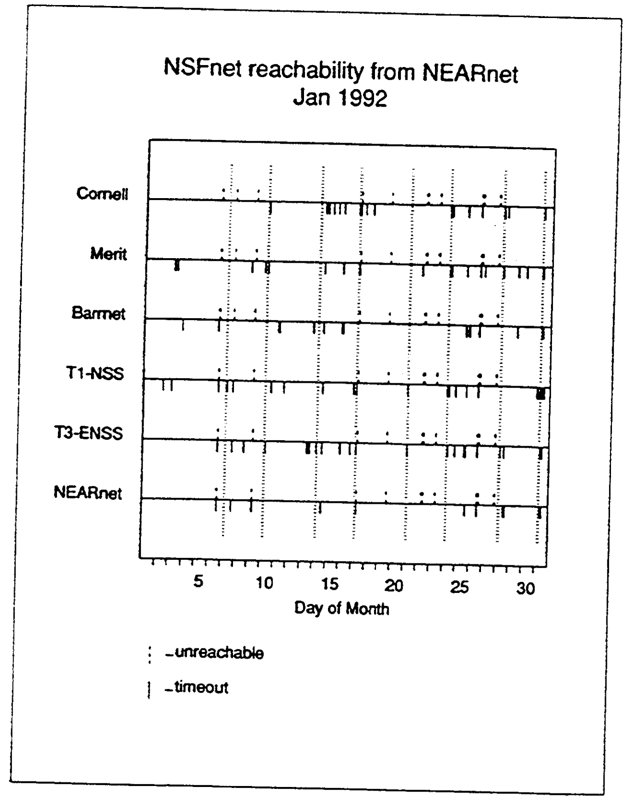
NSFnet reachability from NEARnet
Jan 1992
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
30
Day of Month
-unreachable
-timeout

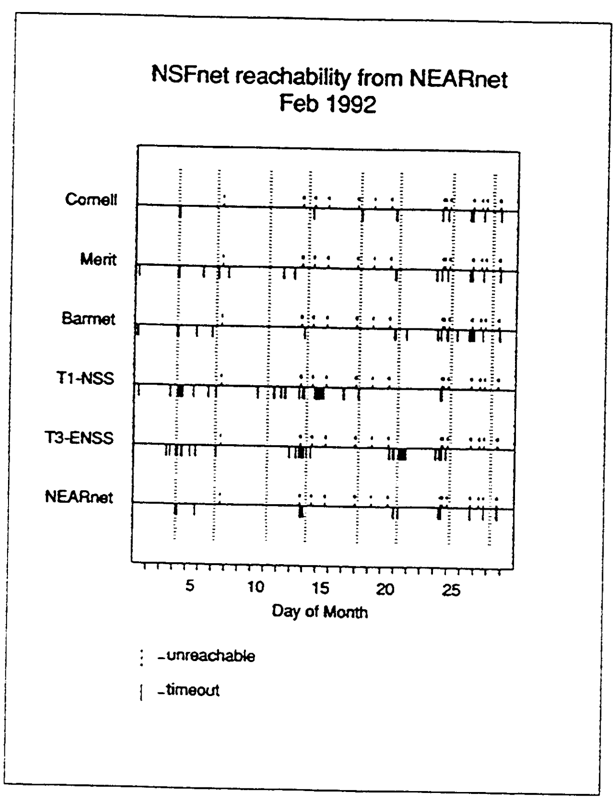
NSFnet reachability from NEARnet
Feb 1992
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
Day of Month
-unreachable
-timeout

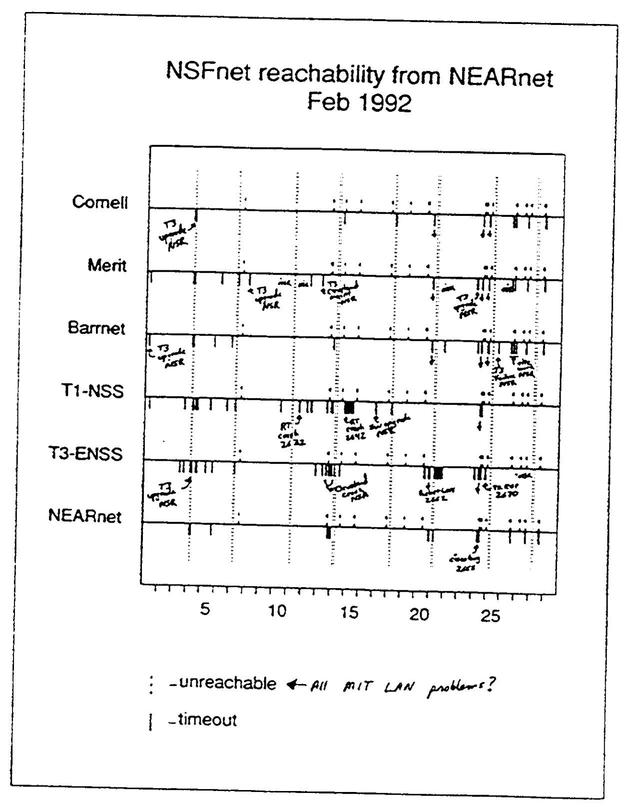
NSFnet reachability from NEARnet
Feb 1992
Cornell
Merit
Barrnet
T1-NSS
T3-ENSS
NEARnet
5
10
15
20
25
-unreachable ← All MIT LAN problems?
-timeout

NSFnet reachability from NEARnet
Aug 1991 - Feb 1992

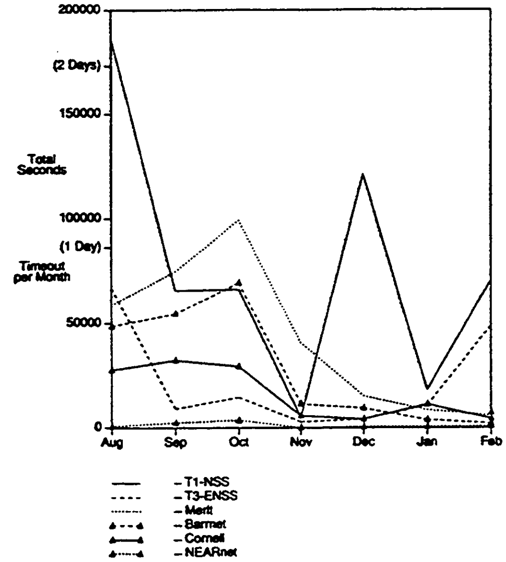

NSFNET Network Service

- NEARnet expected a network *service*
  not just a network
- includes:
  real-time network monitoring
  reliable, cooperative NOC
  good coordination on elective outages
  good coordination on technology changes
  feedback on problems
  i.e., not just paths through a cloud
- includes service relationship

Information exchange examples

```
From: Merit Network Operations Center <noc@merit.edu>
Date: Sat, 8 Feb 92 20:23:51 EST
Message-Id: <9202090123.AA15078@home.merit.edu>
To: nsr@merit.edu
Subject: 02/08/92 ENSS132 and NSS5, Pittsburgh, unreachable 19:59 - ? EST
Status: R

   02/08/92 ENSS132 and NSS5, Pittsburgh, unreachable 19:59 - ? EST.

                                        Riaz Taherzadeh-Yazdian
                                        Merit Network Operations


From: Merit Network Operations Center <noc@merit.edu>
Date: Sat, 8 Feb 92 20:40:33 EST
Message-Id: <9202090140.AA15123@home.merit.edu>
To: nsr@merit.edu
Subject: 02/08/92 ENSS132 and NSS5, Pittsburgh, unreachable 19:59 - 20:29
Status: R

   02/08/92 ENSS132 and NSS5, Pittsburgh, unreachable 19:59 - 20:29 EST.

                                        Riaz Taherzadeh-Yazdian
                                        Merit Network Operations
```

The Cost to NEARnet

- intangible
  unhappy users
  customers not signed
  lost upgrades
- tangible
  labor on NSFnet related problems May - Aug
    50% of operations labor
    25% of analysis labor
  $$

Operations	$175K	* .5	=	$ 87.5
Analysis	$ 80K	* .25	=	$ 20.0
Total				$107.5

NEARnet - MERIT relationship

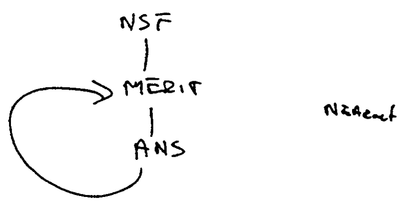

>> NEARnet reported a problem to Merit on 11/26/91...
>> Merit opened a ticket (#17125).
>> Nine days later an MCI ticket was opened.
>> Meanwhile...seven more NEARnet ticket notes were entered.
>> MCI misses two scheduled maintenance windows.
>> Three days after opening the MCI ticket, testing begins and continues until 12 Dec when MCI replaces a ESFMU card to allow statistics to be gathered.
>> On 13 Dec NEARnet switches over to the T3...
>> We are now 17 days into a major problem with no end in sight.
>> On 16 Dec MCI repairs a cable and replaces a smart jack and a repeater...
>> On 17 Dec Merit and MCI reset the ACSUs and discovered that clocking had been set inconsistently.
>> On 17 Dec NEARnet Operations reports that pings look good.
>> 21 days and 28 NEARnet ticket notes later and we may be out of the woods...

finger ticket-2420@nic.near.net

02/04/92 14:22
skennedy

morning to SAT morning. We have not been apprised of the results as of yet, but will update the ticket when we are.

We are investigating continuing daily user compaints about connectivity across this link.

Note: 42
02/19/92 18:49
skennedy

Never was apprised of the results from the MCI testing. Merit IE personnel will be speaking with NEARnet personnel tomorrow.

Note: 43
02/28/92 12:41
skennedy

Merit deployed some new software on NSS-17 at Ann Arbor which is supposed to help with the DCD problem. In this new software the rcp_routed is less sensitive, or even ignores DCD fluctuation for routing purposes. Also the PSPs will not send Net Unreachables in the new software. No word yet on it's performance at NSS-17, and whether it will be deployed to Princeton NSS-8 or Pittsburgh NSS-5, NEARnet's two T1 NSFnet connections.

Note: 44
03/10/92 15:59
perfetti

Contacted MERIT Operations for a status on the DCD waffling problem. Awaiting a response as to whether the RT's at BBN and MIT got the latest change for the DCD problem.

Note: 45
03/11/92 13:15
perfetti

Received response from MERIT staff and the the DCD waffling problem is no longer an issue since a software configuration change was made to ignore DCD transitions. When the DCD transitions were acknowledged, short line outages of 2-3 minutes were experienced thus causing connectivity problems for NEARnet users.
This is an interim change until there is a new system to detect short duration outages.

Problem Solution:

A software configuration change was made by MERIT staff to ignore DCD transitions which, when acknowledged would cause a brief line outage of 2-3 minutes. Also, a new version of ricd has been deployed. This corrects the T1 interface freeze problem which also caused connectivity problems.

For more information, please contact the NEARnet Network Operations Center at nearnet-ops@nic.near.net or (617) 873-8730. Thanks!

NSF, MERIT & NEARnet meeting - Feb 7

action items

1 FARNET should address hardening of regionals
2 this talk
3 NEARnet to put its ticket system into public domain
4 Merit to circulate the IETF Internet Engineering report to all regionals & share bypass recovery plans
5a should there be a 3rd net? research net, production net + *test net*?
5b Merit needs to circulate plans for introduction of new technologies to mid-levels & get their management support
6 Merit to write & circulate white paper on escalation procedures
7 Merit should call for help from ANS, IBM & MCI on establishing quality approaches for network management re: above

Current

- connectivity *much* better
- still problems users see
  routing transitions = zapped sessions
  NEARnet Planning committee report
  'users not relying on Internet
  for interactive use'
- have the underlying technology
  problems been fixed?

  don't know - no data from Merit messages
- need work on service relationship

Opinion

- new technology can be a b...
- not a bad job of implementing T3
- not a good job of relating to mid-levels
- not a good migration strategy

Future

- technology problems *look* mostly solved
- real data in messages
- developing relationship with Merit
  e.g. participation in IETF, FARNET etc

imponderables
- whither ANS, PSI, Alternet, CIX, ...
- congress is getting into the act
  (back to a dark cloud hiding paths)

6.8 NIC Services

Presented by Scott Williamson/Network Solutions

NIC SERVICES

Government Systems, Inc. and Network Solutions, Inc.

Presented By

Scott Williamson

scottw@nic.ddn.mil

CHANGES AT THE NIC

- Replaced FTP and Telnet Servers
- Sorted Out Mail List Problems
- Improved Query

  @Hostname Search For Mailbox Users

TODAYS FINAL WHOIS SOLUTION

PHASE II

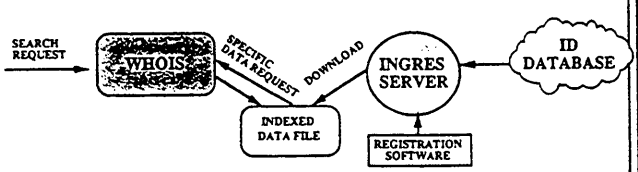

REGISTRATION REMINDER

- Domain Registration

  Unique Name

  Servers Registered And Active

  New POC Needs To Be Registered In Whois

- Internet Number Registration

  Need To Fully Justify Class B Request

  Class A Addresses Need Extensive Documentation

CONNECTED-VS-NON-CONNECTED

- Final Decision Required From DISA
- Trying To Process Requests The Same Day
- Ready To Merge The Databases

Chapter 7

Distributed Papers

7.1 Flow and Congestion Control in Networks

Flow and Congestion Control in Networks Using Deterministic Models

Presented by Ashok K. Agrawala/UMD

Traditional designs of flow and control techniques is based on stochastic models of the network environment and steady state analysis of the models. Recently, we formulated an approach to the analysis of store and forward networks using deterministic techniques which has been used to design a flow control scheme. Instead of a window or a rate, this scheme is based on using *send time control* in which the time to send each packet is determined and used. The time to send is determined based on the model used.

Based on this model we have developed a new transport level protocol, DTP, which gives the same functionality as TCP. The flow control scheme uses a few estimators for the crucial parameters of the network which are obtained from the sending time and the acknowledgement time of the packets. implemented to run on the internet.

Measurements from experiments on the Internet show that DTP gives a significant (4x-8x) improvement in throughput as compared to the current version of TCP. These results will be presented along with the results on the network dynamics obtained through the measurements.

Send Time Flow Control Experiments

Ashok K. Agrawala
Dheeraj Sanghi
Olafur Gudmundsson

Department of Computer Science
University of Maryland
College Park, MD 20742
(301) 405 2665
agrawala@cs.umd.edu

Outline

- Motivation
- Model of a Connection
- Send Time Control
- DTP Design
- Experimental Results
  - September/Nov 1991 | T1 Backbone
  - March 1992 | T3 Backbone
- Concluding Remarks
- Future Work

Motivation

- Flow control is significant for bulk transfers
  - FTP
  - Multi Media
- Not as important for sporadic transfers
  - Telnet

Motivation

- Best time to send a packet?
- Time to send controlled by flow control scheme
  - Window Based
  - Rate Based
  - *Time Based*
- Designs often based on studies of Steady State
  - Reality is that steady state is rarely present
- Performance
  - Throughtput
  - Transit Time

TRANSIENT BEHAVIOR

Structure of a Store and Forward Path

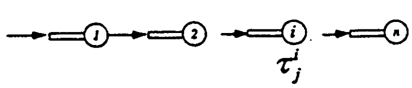

a_j^i Arrival Time of Packet j at Node i

d_j^i Departure Time of Packet j at Node i

τ_j^i Service Time of Packet j at Node i

π^i Delay from Node $(i-1)$ to Node i

Sender is treated as server 0

Performance Measure

- Require measures which reflect dynamic behavior

T_j Transit Time $= d_j^n - d_j^0$

R_j Interpacket Time $= d_j^n - d_{j-1}^n$

P_j Packet Performance Index

$= T_j R_j$

- The inverse of packet performance index is indicative of POWER

POWER = E { $1/Pj$ }

RT Plot

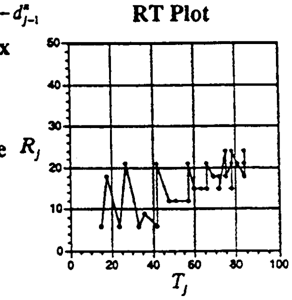

Send Time Control Scheme

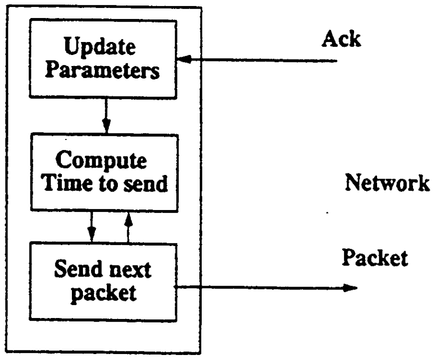

Send Time Control Scheme

- Determine the time to send each packet
- Receive the ack and time stamp it
- Based on the received ack time and the send time of previous packets, calculate the time to send the next packet
- Estimate and update parameters as necessary

Simple Send Time Control Policy

- Simple service time model
- Simple cross traffic model

$$\delta_x = d_x^n - Max(d_x^0 + \tau, d_{x-1}^n + \tau^b)$$

$$\hat{e}_x = \alpha \hat{e}_{x-1} + (1-\alpha)\delta_x$$

$$d_{x+k}^0 = Max(d_x^n + k\tau^b - \tau + (k-1)\hat{e}_x, \; d_{x+k-1}^0 + \tau^b + \hat{e}_x)$$

Send Time Control Implications

- Determine the time to send each packet
- Have a mechanism for sending the packet at this time

- Performance depends on the technique used for determining the time to send
- Overhead depends on the implementation of the mechanism
- Estimators have to take into account the assumptions about the network dynamics
- Minimal added complexity

DTP Design

Same Functionality as TCP

- Send Time Control
- Selective and Cumulative Acknowledgements
- Immediate Acknowledgements
- Use Packet Sequence Numbers
- Instrumentation is integrated with the design

DTP Implementation

- On 4.3 BSD Unix and SUN OS 4.1 in IP environment
- Use 10 ms timer

- Data is packetized on arrival at source
  - Retransmission timer for each packet
  - Table for per-packet information, e.g. send time
  - Store information about packets and not stream
- Memory management done by protocol and not by socket layer

Error Handling

- Packet Loss
  - Selective repeat
  - No drastic action taken to change flow control parameters
- Ack Loss
  - Due to selective and cumulative acks - has little effect
- Out of sequence packet
  - Reordered by the receiver
  - Limitation is the receiver buffer size

Experimental Setup

- Source and destination are the same host
- Use IP Option to send packets via another host
- Packets routed via many different gateways
- Used 32K window size for TCP

Design of Experiments

- 1000 packets of 512/1024 bytes - received correctly
- Varied
  - Time of day
  - Protocol Used
  - Gateway
- Studied
  - Losses
  - Throughput
  - RT Plots

Results -September 1991

September 10, 1991, 11:45 AM

Source: Huggin.cs.umd.edu
Packets sent via: UCB

	TCP	DTP
Transfer Time	81.26s	14.88s
Losses	21	22
Avg T	367.8ms	405.1ms
Avg R	81.26ms	14.88ms

Time Plots

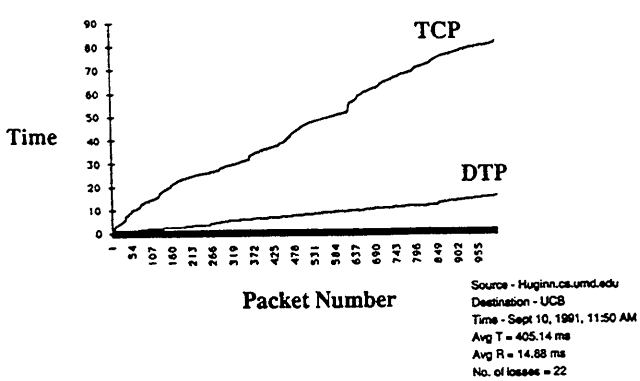

RT Plot for DTP Connection to UCB

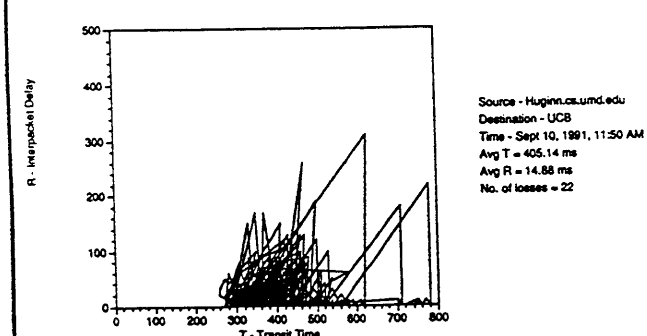

RT Plot for TCP Connection to UCB

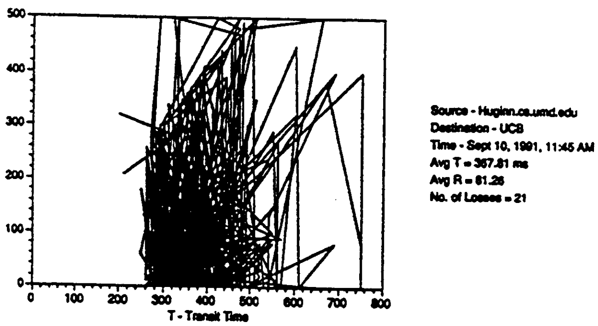

All R values > 500 were replaced by 500 for the plot
Rmax = 1390 ms

October Experiments

T1 Backbone

- 1000 packets of 1024 Bytes
- TCP buffer size 32K
- Run from SparcStation(s)
- Routed via 5 gateways:
  1. Local Connection
  2. AVWHub
  3. NSS.Suranet
  4. MIT
  5. UCBerkeley
- Time reported is the time to receive the ack for the 1000th correctly received packet

OCT/NOV 1991 Experiments

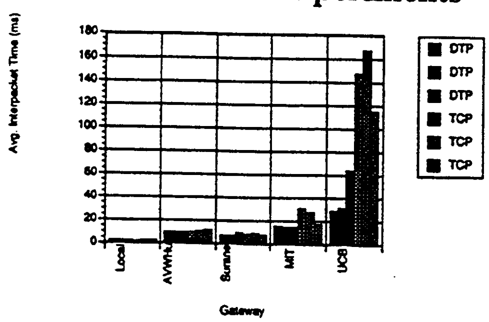

March 1992 Experiments

T3 Backbone

- 1000 packets of 1024 Bytes
- TCP buffer size 32K
- Run from SparcStation(s)
- Routed via gateways:
  1. AVWHub
  2. NSS.Suranet
  3. UCBerkeley
- Time reported is the time to receive the ack for the 1000th correctly received packet

March 1992 Experiments

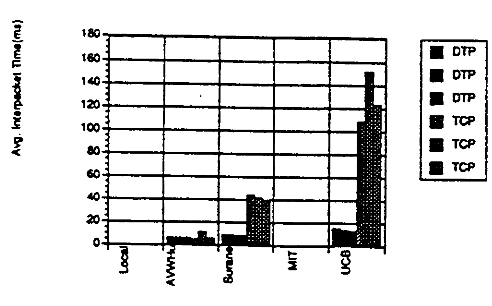

Other Experiments

- Overhead - 200 microsec/packet
  with no optimization of code and with instrumentation
- Two hosts - No change in behavior
- Selective Ack - Does not account for all the improvement

Some Observations

- SuraNet Gateway still giving significant losses
- Losses observed on connections not through SuraNet (Not as high but not negligible either)
- Some variation in T3 results is due to the use of IP Options

Concluding Remarks

- Based on Experimental Observations
  - DTP shows significant improvement over current Internet protocols in its ability to adapt
  - Protocol behaves better in unstable/lossy environments
- Math model relatively straightforward
- Time based scheme is well suited for many environments
- Software available for experimentation

Future Work

- Experiment with
  - Variety of source destinations
  - Networks, technologies, speeds, topologies
- Optimize Implementation
  - Timers
  - Instrumentation
- Study
  - Cross traffic models
  - Loss models
  - Estimators

Appendix A

Attendees

Andy Adams
Merit Network, Inc.
(W)313-763-4897
(WF)313-747-3745
ala@merit.edu

Vikas Aggarwal
JvNCnet
(W)609-258-2403
(WF)609-258-2424
aggarwal@jvnc.net

Robert Aiken
National Science Foundation
(W)202-357-9717
(WF)202-357-0320
raiken@nsf.gov

Edward Alcoff
The Wollongong Group, Inc.
(W)415-962-7100
(WF)415-969-5547
oldera@twg.com

Steve Alexander
INTERACTIVE Systems
Corporation
(W)708-505-9100
(WF)708-505-9133
stevea@i88.isc.com

James Alfieri
Bell Communications Research
(W)908-758-5460
(WF)908-758-4331
jdal@troy.cc.bellcore.com

J. Allard
Microsoft Corporation
(W)206-936-9031
(WF)206-936-7329
jallard@microsoft.com

Claudio Allocchio
National Institute for
Nuclear Physics
(W)+39 40 3758523
(WF)+39 40 226338
claudio.allocchio@elettra-ts.infn.it

Guy Almes
Advanced Network & Services, Inc.
(W)917-789-5333
(WF)914-789-5310
almes@ans.net

Philip Almquist
Consultant
(W)415-752-2427
almquist@jessica.stanford.edu

Harald Alvestrand
UNINETT
(W)+47 7 59 70 94
(WF)+47 7 53 25 86
harald.alvestrand@delab.sintef.no

Miriam Amos Nihart
Digital Equipment Corporation
(W)206-865-8549
(WF)206-865-8890
miriam@ltning.enet.dec.com

Michael Anello
XLNT Designs, Inc.
(W)619-487-9320
(WF)619-487-9768
mike@xlnt.com

Farhad Anklesaria
University of Minnesota
(W)612-624-9360
(WF)612-625-6817
fxa@boombox.micro.umn.edu

Susan Arnold
CERFnet
(W)619-455-3905
(WF)619-455-3990
susie@cerf.net

Nagaraj Arunkumar
3Com Corporation
(W)408-764-5104
(WF)408-764-5091
nak@3com.com

Syed Ashraf
Datability, Inc.
(W)201-438-2400
(WF)201-438-7767
irfan@pluto.dss.com

Karl Auerbach
Empirical Tools and Technologies
(W)408-427-5280
(WF)408-427-5281
karl@empirical.com

Serge Aumont
CICB
(W)+33 99 84 71 47
(WF)+33 99 84 71 11
aumont@cicb.fr

Robert Austein
Epilogue Technology Corporation
(W)617-942-0915
(WF)617-942-7846
sra@asylum.sf.ca.us

Zorica Avramovic
SAIC
(W)703-448-1683
zorica@sparta.com

Cynthia Bagwell
The MITRE Corporation
(W)703-883-5438
(WF)703-883-7142
cbagwell@gateway.mitre.org

Fred Baker
Advanced Computer Communications
(W)805-963-9431
(WF)805-962-8499
fbaker@emerald.acc.com

Alan Bales
Sprint Government Systems Division
(W)816-854-2303
(WF)816-854-2301
rcollett@icmi.icp.net

John Ballard
Microsoft Corporation
(W)206-936-4688
(WF)206-883-8101
jballard@microsoft.com

Paul Barker
University College London
(W)+44 71 380 7366
(WF)+44 71 387 1397
p.barker@cs.ucl.ac.uk

James Barnes
Xylogics, Inc.
(W)617-272-8140
(WF)617-272-3159
barnes@xylogics.com

William Barns
The MITRE Corporation
(W)703-883-6832
(WF)703-883-7142
barns@gateway.mitre.org

Richard Basch
Massachusetts Institute of Technology
(W)617-253-0100
(WF)617-258-8736
basch@mit.edu

Mark Baushke
3Com Corporation
(W)408-764-5110
(WF)408-764-5003
mdb@nsd.3com.com

Charles Bazaar
Emulex Corporation
(W)714-668-5416
(WF)714-668-6818
bazaar@emulex.com

Jordan Becker
Advanced Network & Services, Inc.
(W)914-789-5308
(WF)914-789-5310
becker@nis.ans.net

Tom Benkart
ACC Systems
(W)410-290-8100
(WF)410-290-8106
teb@saturn.acc.com

Tim Berners-Lee
CERN
(W)+41 22 767 3755
(WF)+41 22 767 7155
timbl@info.cern.ch

William Biagi
Corporation for Open Systems International
(W)703-883-2806
(WF)703-848-4572
bbiagi@cos.com

Richard Binder
Corporation for National Research Initiatives
(W)703-620-8990
(WF)703-620-0913
rbinder@nri.reston.va.us

Joe Blackmon
National Center for Supercomputing Applications
(W)217-244-4019
(WF)217-244-1987
blackmon@ncsa.uiuc.edu

Bill Bliss
Microsoft Corporation
(W)206-822-8080
(WF)206-936-7329
billbl@microsoft.com

Larry Blunk
Merit Network, Inc.
(W)313-763-6056
(WF)313-747-3745
ljb@merit.edu

Ken Boggs
IBM Corporation
(W)919-543-5588
(WF)919-543-0723
boggs@ralvm12.vnet.ibm.com

David Bolen
Advanced Network & Services, Inc.
(W)914-789-5327
(WF)914-789-5310
db31@nis.ans.net

David Borman
Cray Research, Inc.
(W)612-683-5571
(WF)612-683-5599
dab@cray.com

Robert Bosen
Enigma Logic
(W)415-827-5707
(WF)415-827-2593
bbosen@world.std.com

Steve Bostock
Novell, Inc.
(W)408-473-8203
(WF)408-435-1706
steveb@novell.com

Bill Bowman
Hewlett-Packard
(W)415-691-5678
(WF)415-691-5681
bill@hpprsd.mayfield.hp.com

Robert Braden
Information Sciences Institute
(W)310-454-8852
(WF)310-823-6714
braden@isi.edu

Scott Bradner
Harvard University
(W)617-495-3864
(WF)617-495-0914
sob@harvard.edu

Hans-Werner Braun
San Diego Supercomputer Center
(W)619-534-5035
(WF)619-534-5113
hwb@sdsc.edu

David Bridgham
Epilogue Technology Corporation
(W)617-942-0915
(WF)617-942-7846
dab@epilogue.com

Scott Brim
Cornell University
(W)607-255-9392
(WF)607-255-5771
swb@cornell.edu

Ronald Broersma
Naval Ocean Systems Center
(W)619-553-2293
(WF)619-553-6811
ron@nosc.mil

Caralyn Brown
Wellfleet Communications, Inc.
(W)617-275-2400
(WF)617-275-5001
cbrown@wellfleet.com

Jack Brown
United States Army
Computer Engineer Center
(W)602-533-2834
(WF)602-533-5676
jbrown@huachuca-emh8.army.mil

Jordan Brown
Quarterdeck Office Systems
(W)310-314-4260
jbrown@qdeck.com

J. Nevil Brownlee
University of Auckland
(W)+64 9 373 7599
(WF)+64 9 373 7425
nevil@aukuni.ac.uz

Gregory Bruell
Shiva Corporation
(W)617-252-6300
(WF)617-252-6852
gob@shiva.com

Eric Brunner
Tule Network Systems
(W)415-965-7710
brunner@practic.com

Theodore Brunner
Bell Communications Research
(W)201-829-4678
(WF)201-267-9298
tob@thumper.bellcore.com

Niels Brunsgaard
Dowty Network Systems
(W)+45 44 53 01 00
(WF)+45 44 53 14 15
nob@dowtyns.dk

Rick Bubenik
Washington University
(W)314-935-7530
(WF)314-935-7302
rick@arl.wustl.edu

Christopher Bucci
Datability, Inc.
(W)201-438-2400
bucci@dss.com

Philip Budne
Shiva Corporation
(W)617-252-6300
(WF)617-252-6852
phil@shiva.com

Jeffrey Burgan
NASA Ames Research Center
(W)415-604-5705
(WF)415-604-0063
jeff@nsipo.nasa.gov

John Burnett
Adaptive
(W)415-780-7768
(WF)415-364-0123
jlb@adaptive.com

Randy Butler
National Center for
Supercomputing Applications
(W)217-244-8285
(WF)217-244-1987
rbutler@ncsa.uiuc.edu

Susan Calcari
CERFnet
(W)619-455-3904
(WF)619-455-3900
calcaris@cerf.net

Ross Callon
Digital Equipment Corporation
(W)508-486-5009
(WF)508-486-5279
callon@bigfut.enet.dec.com

David Carrel
Munin Systems, Inc.
(W)412-481-8133
(WF)412-661-6837
carrel@munin.com

Lida Carrier
Apple Computer, Inc.
(W)408-974-8325
(WF)408-974-0892
lida@apple.com

Douglas Carson
University of Toronto
(W)416-978-3855
(WF)416-978-6620
carson@utcs.utoronto.ca

Jeffrey Case
University of Tennessee
(W)615-573-5931
(WF)615-573-9197
case@cs.utk.edu

Stephen Casner
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
casner@isi.edu

Cerafin Castillo
Emulex Corporation
(W)408-452-4777
(WF)408-452-4751
cec@emulex.com

Peter Caswell
Pacer Software, Inc.
(W)508-898-3300
(WF)508-366-1356
pfc@pacer.soft.com

Vinton Cerf
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
vcerf@nri.reston.va.us

Fwu-Tsair Chang
Industrial Technology
Research Institute
(W)408-946-3015
(WF)408-946-3019
chang@n2sun1.ccl.itri.org.tw

John Chang
US West Advanced Technologies
(W)303-541-6255
(WF)303-541-6250
jrc@uswest.com

A. Lyman Chapin
Bolt Beranek and Newman
(W)617-873-3133
(WF)617-873-4086
lyman@bbn.com

Dilip Chatwani
SynOptics Communications, Inc.
(W)408-764-1624
dilip@synoptics.com

Chih-Hui Chen
Industrial Technology
Research Institute
(W)408-946-3015
(WF)408-946-3019

Dean Cheng
Retix
(W)310-828-3400
(WF)310-828-2255
dean@sun2.retix.com

Andrew Cherenson
Silicon Graphics, Inc.
(W)415-962-3486
(WF)415-969-2314
arc@sgi.com

Richard Cherry
Novell, Inc.
(W)510-975-4520
(WF)510-947-1238
rcherry@wc.novell.com

Carson Cheung
Bell Northern Research
(W)613-763-3562
(WF)613-763-8581
carson@bnr.ca

J. Noel Chiappa
(H)804-898-8183
(HF)804-898-7663
jnc@lcs.mit.edu

Bilal Chinoy
San Diego Supercomputer Center
(W)619-534-5035
(WF)619-534-5113
bac@sdsc.edu

Isaac Choi
SynOptics Communications, Inc.
(W)408-764-1546
(WF)408-988-5525
ichoi@synoptics.com

Cyrus Chow
NASA Ames Research Center
(W)415-604-6843
(WF)415-604-4003
cchow@ames.arc.nasa.gov

Jodi-Ann Chu
University of Hawaii
(W)808-956-2400
(WF)808-956-2412
jodi@uhunix.uhcc.hawaii.edu

Kimberly Claffy
San Diego Supercomputer Center
(W)619-534-5035
(WF)619-534-5113
kc@sdsc.edu

George Clapp
Ameritech Services, Inc.
(W)708-248-5462
(WF)708-248-3775
clapp@ameris.center.il.ameritech.com

Cynthia Clark
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
cclark@nri.reston.va.us

Henry Clark
OARnet
(W)614-292-6483
(WF)614-292-7168
henryc@oar.net

Alan Clegg
MCNC
(W)919-248-1120
(WF)919-248-1405
abc@concert.net

John Clement
EDUCOM
(W)202-872-4200
clement@educom.edu

Richard Cogger
Cornell University
(W)607-255-5510
(WF)607-255-5771
rhx@cornell.cit.bitnet

Richard Colella
National Institute of
Standards and Technology
(W)301-975-3627
(WF)301-590-0932
colella@osi.ncsl.nist.gov

Robert Cole
AT&T Bell Laboratories
(W)908-949-1950
(WF)908-949-1726
robert.g.cole@qsun.att.com

Robert Collet
US Sprint
(W)703-904-2230
(WF)703-904-2527
rcollet@icm1.icp.net

Rob Coltun
Consultant
(W)301-240-2168
rcoltun@ni.umd.edu

Garth Conboy
Pacer Software, Inc.
(W)508-898-3300
(WF)508-366-1356
gc@pacer.soft.com

James Conklin
EDUCOM
(W)202-872-4200
(WF)202-872-4318
conklin@bitnic.educom.edu

Robert Cooney
Naval Computer and
Telecommunications Station
(W)202-433-2448
(WF)202-433-3188
cooney@wnyose.nctsw.navy.mil

Curtis Cox
Naval Computer and
Telecommunications Station
(W)202-433-4026
(WF)202-433-5786
ccox@wnyose.nctsw.navy.mil

Tracy Cox
Bell Communications Research
(W)908-758-2107
(WF)908-758-4177
tacox@sabre.bellcore.com

Stephen Coya
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
scoya@nri.reston.va.us

David Crocker
The Branch Office
(W)408-246-8253
(WF)408-249-6205
dcrocker@mordor.stanford.edu

Stephen Crocker
Trusted Information Systems, Inc.
(W)301-854-6889
(WF)301-854-5363
crocker@tis.com

Dave Cullerot
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-335-3467
cullerot@ctron.com

John Curran
Bolt Beranek and Newman
(W)617-873-4398
jcurran@bbn.com

Dave Dalton
Intel Corporation
(W)602-869-4226
(WF)602-869-4001
dave\verb+_+dalton@ccm.hf.intel.com

Peter Danzig
University of Southern California
(W)213-740-4780
(WF)213-740-7285
danzig@pollux.usc.edu

Megan Davies
Corporation for National Research Initiatives
(W)703-620-8990
(WF)703-620-0913
mdavies@nri.reston.va.us

James Davin
Massachusetts Institute of Technology
(W)617-253-6020
(WF)617-258-8682
jrd@ptt.lcs.mit.edu

Osmund de Souza
AT&T Bell Laboratories
(W)908-949-1393
(WF)908-949-4673
osmund.desouza@att.com

Buzz Dean
Enterprise Systems Division
(W)408-974-2509
(WF)408-974-5610

Steve Deering
Xerox Corporation
(W)415-812-4839
(WF)415-812-4471
deering@xerox.com

John Denune

(H)619-698-9340
(WF)619-594-6191
jdenune@saturn.sdsu.edu

Peter Deutsch
McGill University
(W)514-398-3709
(WF)514-398-6876
peterd@cc.mcgill.ca

Kurt Dobbins
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-335-8121
dobbins@ctron.com

Kathleen Dodd
The MITRE Corporation
(W)703-883-6664
(WF)703-883-7142
kathy@gateway.mitre.org

Bob Doolittle
Thinking Machines Corporation
(W)617-234-2734
(WF)617-876-1823
rad@think.com

Jack Drescher
MCNC
(W)919-248-1412
(WF)919-248-1405
drescher@concert.net

Ralph Droms
Bucknell University
(W)717-524-1145
(WF)717-524-3760
droms@bucknell.edu

Guy Dubuc
Newbridge Networks Corporation
(W)613-591-3600
(WF)613-591-3680
guyd@newbridge.com

Daniel Duchamp
Columbia University
(W)212-854-8115
(WF)212-666-0140
djd@cs.columbia.edu

Billy Durham
Honeywell Corporation
(W)602-345-6468
(WF)602-436-2252
durham@fishery.honeywell.com

Steve Dusse
RSA Data Security, Inc.
(W)415-595-8782
(WF)415-595-1873
spock@rsa.com

Tom Easterday
Ohio State University
(W)614-292-4843
(WF)614-292-7081
tom@cic.net

Raymond Edgerton
Bell Atlantic
(W)215-963-6740
(WF)215-963-6776
bellmf!edgerton@uwm.edu

Bill Edwards
US Sprint
(W)913-967-5334
(WF)913-967-2526

Gary Ellis
Hewlett-Packard
(W)415-857-8199
(WF)415-852-3762
garye@hpspd.spd.hp.com

Alan Emtage
McGill University
(W)514-398-8117
(WF)514-398-6876
bajan@cc.mcgill.ca

David Engel
Optical Data Systems, Inc.
(W)214-234-6400
(WF)214-234-4059
david@ods.com

Robert Enger
Advanced Network & Services, Inc.
(W)914-789-5355
(WF)914-789-5310
enger@ans.net

Urs Eppenberger
Switch
(W)+41 1 2623141
(WF)+41 1 2618138
eppen@verw.switch.ch

Johnny Eriksson
NORDUnet
(W)+46 8 790 6517
(WF)+46 8 24 11 79
bygg@sunet.se

Michael Erlinger
Lexcel
(W)714-870-2117
(WF)714-680-0279
mike@lexcel.com

Julio Escobar
Bolt Beranek and Newman
(W)617-873-4579
(WF)617-873-3776
jescobar@bbn.com

Susan Estrada
CERFnet
(W)619-455-3902
(WF)619-455-3990
estradas@cerf.net

Erik Fair
Apple Computer, Inc.
(W)408-974-1779
(WF)408-974-3103
fair@apple.com

Kevin Fall
San Diego Supercomputer Center
(W)619-534-1149
(WF)619-534-7029
kfall@ucsd.edu

Bill Fardy
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-335-8121
fardy@ctron.com

Dino Farinacci
cisco Systems
(W)415-688-4696
(WF)415-688-8282
dino@cisco.com

Carol Farnham
MIDnet
(W)402-472-8971
(WF)402-472-8486
cfarnham@unl.edu

Stefan Fassbender
EASInet
(W)+49 228 819 9654
(WF)+49 228 819 9635
stf@easi.net

Elizabeth Feinler
Sterling Software
Federal Systems
(W)415-604-1238
jake@nsipo.arc.nasa.gov

Sallie Fellows
New Hampshire State
Department of Education
(W)603-271-2778
(WF)603-271-1953
sallie%ed@psc.plymouth.edu

Constance Fleenor Lloyd
Lloyd & Associates
(W)916-676-1147
cfleenor@lloyd.com

Peter Ford
Los Alamos National Laboratory
(W)505-665-0058
(WF)505-665-4939
peter@lanl.gov

Jill Foster
University of
Newcastle-Upon-Tyne
(W)+44 91 222 8250
(WF)+44 91 222 8232
jill.foster@newcastle.ac.uk

Karl Fox
Morning Star Technologies, Inc.
(W)614-451-1883
(WF)614-459-5054
karl@morningstar.com

Mark Fox
Digital Equipment Corporation
(W)508-486-5577
(WF)508-486-7417
m\verb+_+fox@took.enet.dec.com

Barbara Fraser
CERT
(W)412-268-7090
(WF)412-268-6989
byf@cert.sei.cmu.edu

Jeff Fried
Proteon, Inc.
(W)508-898-2800
(WF)508-898-3176
jmf@relay.proteon.com

Bob Friesenhahn
PUREDATA Research/USA
(W)512-690-1024
(WF)512-690-0597
pdrusa!bob@uunet.uu.net

Karen Frisa
Carnegie Mellon University
(W)412-268-8576
(WF)412-268-4987
karen.frisa@andrew.cmu.edu

Cliff Frost
University of California,
Berkeley
(W)510-643-5385
(WF)510-643-5385
cliff@cmsa.berkeley.edu

Vince Fuller
BARRNet
(W)415-723-6860
(WF)415-723-0010
vaf@stanford.edu

Maria Gallagher
NASA Science Internet
(W)415-604-3601
(WF)415-604-0063
maria@nsipo.nasa.gov

Shawn Gallagher
Digital Equipment Corporation
(W)508-952-3519
gallagher@quiver.enet.dec.com

Richard Gallup
San Diego Supercomputer Center
(W)619-534-5064
(WF)619-534-5152
rpg@sdsc.edu

James Galvin
Trusted Information Systems, Inc.
(W)301-854-6889
(WF)301-854-5363
galvin@tis.com

Joan Gargano
University of California,
Davis
(W)916-752-2591
(WF)916-752-9154
jcgargano@ucdavis.edu

Eugene Geer
Bell Communications Research
(W)908-758-2149
(WF)908-758-4184
ewg@nvuxr.bellcore.com

Jisoo Geiter
The MITRE Corporation
(W)703-883-6571
(WF)703-883-7142
geiter@gateway.mitre.org

Tony Genovese
ESnet
(W)415-423-2471
(WF)415-422-0435
genovese@es.net

Elise Gerich
University of Michigan
(W)313-936-2120
(WF)313-747-3745
epg@merit.edu

David Getchell
Lawrence Livermore
National Laboratory
(W)510-423-0338
(WF)510-423-8534
dag@llnl.gov

Robert Gilligan
Sun Microsystems, Inc.
(W)415-336-1012
(WF)415-968-4298
gilligan@eng.sun.com

Anders Gillner
Swedish University
Computer Network
(W)+46 8 7906518
(WF)+46 8 241179
awg@sunet.se

Joseph Godsil
National Center for
Supercomputing Applications
(W)217-244-7715
(WF)217-244-1987
jgodsil@ncsa.uiuc.edu

Leonard Gomelsky
Timeplex, Inc.
(W)310-443-5760
gomelsky@ee.ucla.edu

Kenneth Goodwin
Pittsburgh Supercomputing
Center
(W)412-268-7034
(WF)412-268-5832
goodwin@psc.edu

Rob Graham
ProTools, Inc.
(W)503-645-5400
(WF)503-645-3577
rob@protools.com

Terry Gray
University of Washington
(W)206-543-3282
(WF)206-543-3909
gray@cac.washington.edu

Michael Grieves
Chipcom Corporation
(W)508-460-8900
(WF)508-460-8952
mgrieves@chipcom.com

Paul Griffiths
IBM Corporation
(W)512-838-3544
(WF)512-838-3509
griff@chang.austin.ibm.com

Pete Grillo
Network Innovations
(W)503-526-9766
(WF)503-526-9766
pl0143@mail.psi.net

Daniel Grossman
Motorola Codex
(W)617-827-7333
(WF)617-821-4218
dan@merlin.dev.cdx.mot.com

Phillip Gross
Advanced Network & Services, Inc.
(W)914-789-5341
(WF)914-789-5310
pgross@nis.ans.net

Olafur Gudmundsson
University of Maryland
(W)301-405-2767
(WF)301-405-2707
ogud@cs.umd.edu

Walter Guilarte
Wandel & Goltermann
Technologies, Inc.
(W)919-941-5730
(WF)919-941-5751
70026.1715@compuserve.com

Jeff Gumpf
Case Western Reserve University
(W)216-368-2982
(WF)216-368-4693
jag@po.cwru.edu

Robert Gutierrez
NASA Science Internet
(W)415-604-3655
(WF)415-604-0063
gutierre@nsipo.nasa.gov

Robert Hagens
University of Wisconsin-Madison
(W)608-262-5776
(WF)608-262-9777
hagens@cs.wisc.edu

Tony Hain
Lawrence Livermore
National Laboratory
(W)510-422-4200
(WF)510-422-0435
hain@es.net

Neil Haller
Bell Communications Research
(W)201-829-4478
(WF)201-829-2504
nmh@thumper.bellcore.com

Martyne Hallgren
Cornell University
(W)607-255-5510
(WF)607-255-5771
martyne@mitchell.cit.cornell.edu

Joel Halpern
Network Systems Corporation
(W)612-424-1606
(WF)612-424-2853
jmh@network.com

Deborah Hamilton
AT&T Bell Laboratories
(W)908-949-9459
(WF)908-949-1726
debbie@qsun!att.com

Keunhee Han
KIST/SERI
(W)+82 42 829 1359
(WF)+82 42 861 1999
khhan@gaya.kreonet.re.kr

Sang-Chul Han
KIST/SERI
(W)+82 42 829 1357
(WF)+82 42 861 1999
schan@garam.kreonet.re.kr

Alf Hansen
SINTEF DELAB
(W)+47 7 592982
(WF)+47 7 532586
Alf.Hansen@delab.sintef.no

Daniel Hansen
Network General Corporation
(W)201-616-0986

Steve Hardcastle-Kille
University College London
(W)+44 71 380 7294
(WF)+44 71 387 1397
s.kille@cs.ucl.ac.uk

Susan Hares
Merit Network, Inc.
(W)313-936-2095
(WF)313-747-3745
skh@merit.edu

Eugene Hastings
Pittsburgh Supercomputing
Center
(W)412-268-4960
(WF)412-268-5832
hastings@psc.edu

Alisa Hata
University of Washington
(W)206-543-3282
(WF)206-543-3909
hata@cac.washington.edu

Jeff Hayward
The University of Texas
(W)512-471-2444
(WF)512-471-2449
j.hayward@utexas.edu

Kenneth Hayward
Bell Northern Research
(W)613-763-4042
(WF)613-763-2626
crm57d@bnr.ca

Frank Heath
Rockwell CMC
(W)714-833-6864
(WF)714-833-6863
heath@cmc.com

Juha Heinanen
Telecom Finland
(W)+358 49 500958
(WF)+358 31 2432211
juha.heinanen@datanet.tele.fi

Tom Hemp
IBM Corporation
(W)919-543-7857
(WF)919-543-0723
themp@ralvm11.vnet.ibm.com

Darryl Henry
Defense Information
Systems Agency
(W)703-487-3035
(WF)703-487-3190
henryd@imo-uvax.dca.mil

Ittai Hershman
Advanced Network & Services, Inc.
(W)914-789-5337
(WF)914-789-5310
ittai@nis.ans.net

Max Hillson
Digital Equipment Corporation
(W)603-881-2956
hillson@koala.enet.dec.com

Robert Hinden
Sun Microsystems, Inc.
(W)415-336-2082
(WF)415-336-6015
hinden@eng.sun.com

Russ Hobby
University of California,
Davis
(W)916-752-0236
(WF)916-752-9154
rdhobby@ucdavis.edu

Randy Hoebelheinrich
Los Alamos National Laboratory
(W)505-667-5421
(WF)505-665-4361
rho@rowan.lanl.gov

Ellen Hoffman
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
esh@merit.edu

Duke Hong
University of California,
Irvine
(W)714-725-3097
(WF)714-856-4056
duke@ics.uci.edu

Jeffrey Honig
Cornell University
(W)607-255-5510
(WF)607-255-5771
jch@risci.cit.cornell.edu

Alton Hoover
Advanced Network & Services, Inc.
(W)914-789-5325
(WF)914-789-5310
hoover@ans.net

Nat Howard
Bell Communications Research
(W)201-829-5196
(WF)201-829-2645
nrh@bellcore.com

Tim Howes
University of Michigan
(W)313-747-4454
(WF)313-764-5140
Tim.Howes@umich.edu.

Hung Huang
University of California,
Irvine
(W)714-856-4105
(WF)714-856-4056
huang@ics.uci.edu

Kathleen Huber
Bolt Beranek and Newman
(W)617-873-2520
(WF)617-873-4086
khuber@bbn.com

Erik Huizer
SURFnet
(W)+31 30 310290
(WF)+31 30 340903
huizer@surfnet.nl

Thomas Hutton
San Diego Supercomputer Center
(W)619-534-5136
hutton@opus.sdsc.edu

Jon Infante
ICL
(W)714-458-7282
(WF)714-458-6257
jon@rv.icl.com

John Ioannidis
Columbia University
(W)212-854-8120
(WF)212-666-0140
ji@cs.columbia.edu

Lenore Jackson
NASA Science Internet
(W)301-286-4574
(WF)301-286-5251
jackson@nsinic.gsfc.nasa.gov

William Jackson
NRaD
(W)619-553-2291
(WF)619-533-6811
jackson@manta.nosc.mil

Ole Jacobsen
Interop, Inc.
(W)800-INTEROP
(WF)415-949-1779
ole@csli.stanford.edu

Ronald Jacoby
Silicon Graphics, Inc.
(W)415-335-1371
(WF)415-969-2314
rj@sgi.com

Yigal Jacoby
Tekelec
(W)818-880-7900
(WF)818-880-6993
jacoby@Tekelec.com

Bob Jeckell
3Com Corporation
(W)408-764-6049
rrj@3com.com

Lew Jenkins
Premenos Corporation
(W)510-602-2000
(WF)510-602-2133
lew@enlil.premenos.sf.ca.us

Barbara Jennings
Sandia National Laboratory
(W)505-815-8554
(WF)505-844-2067
bjjenni@sandia.gov

Dale Johnson
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
dsj@merit.edu

James Jones
Digital Equipment Corporation
(W)508-635-8612
jones@regent.enet.dec.com

Kevin Jordan
Control Data Corporation
(W)612-482-6835
(WF)612-482-4455
kej@udev.cdc.com

Dan Jordt
NorthWestNet
(W)206-562-3000
(WF)206-562-4822
danj@nwnet.net

Ed Juskevicius
Northern Telecom
(W)613-763-7763
(WF)613-763-9225
cnv29@bnr.ca

George Kajos
Coral Network Corporation
(W)508-460-6010
(WF)508-481-6258
kajos@coral.com

Brian Kantor
University of California,
San Diego
(W)619-534-6865
(WF)619-534-1746
brian@ucsd.edu

Scott Kaplan
FTP Software, Inc.
(W)415-543-9001
(WF)415-543-9002
scott@wco.ftp.com

Phil Karn
Qualcomm, Inc.
(W)619-597-5501
(WF)619-455-0571
karn@qualcomm.com

Daniel Karrenberg
RIPE
(W)+31 20 592 0303
daniel.karrenberg@cwi.nl

Prakash Kashyap
Sun Microsystems, Inc.
(W)415-336-4779
prakash.kashyap@sun.com

Frank Kastenholz
Clearpoint Research
Corporation
(W)508-435-2000
(WF)508-435-7504
kasten@europa.clearpoint.com

Neil Katin
SunSoft
(W)415-336-3718
(WF)415-961-0992
katin@eng.sun.com

Dave Katz
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
dkatz@merit.edu

Charles Kaufman
Digital Equipment Corporation
(W)508-486-7329
(WF)508-486-5279
kaufman@dsmail.enet.dec.com

Manu Kaycee
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-424-8121
kaycee@ctron.com

James Keller
Sprint Data Group
(W)703-689-7016
(WF)703-689-6998

Stephen Kent
Bolt Beranek and Newman
(W)617-873-3988
(WF)617-873-4086
kent@bbn.com

Mark Kepke
Hewlett-Packard
(W)303-229-4201
(WF)303-229-2692
mak@cnd.hp.com

Tom Kessler
Sun Microsystems, Inc.
(W)415-336-3145
(WF)415-969-9131
kessler@eng.sun.com

Zbigniew Kielczewski
Eicon Technology, Inc.
(W)514-631-2592
(WF)514-631-3092
zbig@eicom.qc.ca

Thomas Kile
Ballistic Research Laboratory
(W)410-278-8676
(WF)410-278-5075
tgkile@brl.mil

Darren Kinley
RISQ
(W)514-340-5768
(WF)514-340-5777
kinley@crim.ca

H. A. Kippenhan
Fermi National Accelerator Laboratory
(W)708-840-8068
(WF)708-840-2783
kippenhan@fndcd.fnal.gov

Peter Kirstein
University College London
(W)+44 71 380 7286
(WF)+44 71 387 1397
kirstein@cs.ucl.ac.uk

John Klensin
Massachusetts Institute of Technology
(W)617-253-1355
(WF)617-491-6266
klensin@infoods.mit.edu

Yoav Kluger
Fibronics Ltd.
(W)+972 4 313313
(WF)+972 4 550260
ykluger@fibhaifa.com

Holly Knight
Apple Computer, Inc.
(W)408-974-5110
(WF)408-974-0892
holly@apple.com

Mark Knopper
Merit Network, Inc.
(W)313-763-6061
(WF)313-747-3745
mak@merit.edu

Jim Knowles
Sterling Software Federal Systems
(W)415-604-3056
(WF)415-604-6999
jknowles@binky.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
stev@ftp.com

Roman Kochan
ascom Timeplex
(W)714-770-1221
(WF)714-583-7509
roman@doelz.timeplex.com

Rajeev Kochhar
3Com Corporation
(W)408-764-5176
(WF)408-764-5003
rajeev\verb+_+kochhar@3com.com

Deidre Kostick
Bell Communications Research
(W)908-758-2642
(WF)908-758-4177
dck2@sabre.bellcore.com

Fidelia Kuang
Apple Computer, Inc.
(W)408-974-4425
kuang@apple.com

Eva Kuiper
Hewlett-Packard
(W)408-447-3163
(WF)408-447-3660
eva@hpindda.cup.hp.com

David Kumpf
Hewlett-Packard
(W)916-785-5162
kumpf@hprnd.rose.hp.com

Lawerence Lang
cisco Systems
(W)415-688-4638
(WF)415-326-1989
llang@cisco.com

Sylvain Langlois
Electricite de France
(W)+33 47 64 44 02
(WF)+33 47 65 30 01
sylvain@cli53an.edf.fr

Sylvain Langlois
Electricite de France
(W)+33 47 64 44 02
(WF)+33 47 65 30 01
sylvain@cli53an.edf.fr

Tracy LaQuey Parker
University of Texas at Austin
(W)512-471-2844
(WF)512-471-2449
tracy@utexas.edu

John Larson
Xerox Corporation
(W)408-662-9755
(WF)408-662-9756
jlarson@parc.xerox.com

Anthony Lauck
Digital Equipment Corporation
(W)508-486-7644
(WF)508-486-5279
lauck@tl.enet.dec.com

Joseph Lawrence
Bell Communications Research
(W)908-758-4146
jcl@sabre.bellcore.com

Walter Lazear
The MITRE Corporation
(W)703-883-6515
(WF)703-883-7142
lazear@gateway.mitre.org

Eliot Lear
Silicon Graphics, Inc.
(W)415-335-1056
lear@sgi.com

John Leddy
IBM Corporation
(W)203-877-7012
jleddy@ibm.com

Tim Lee-Thorp
Network General Corporation
(W)415-688-2848
(WF)415-327-2303
ngc!tim@uunet.uu.net

Debra Legare
Corporation for National
Research Initiatives
(W)703-620-8990
(WF)703-620-0913
dlegare@nri.reston.va.us

Linda Leibengood
Advanced Network & Services, Inc.
(W)914-785-5347
(WF)914-789-5310
ldl@ans.net

Dory Leifer
University of Michigan
(W)313-763-4896
(WF)313-747-3745
del@terminator.um.cc.umich

John Lekashman
NASA Ames Research Center
(W)415-604-4359
(WF)415-604-4377
lekash@nas.nasa.gov

John Leong
Carnegie Mellon University
(W)412-268-6722
(WF)412-268-4987
john.leong@andrew.cmu.edu

Mark Lewis
Telebit Corporation
(W)408-745-3232
(WF)408-745-3810
mlewis@telebit.com

Kurt Lidl
SURAnet
(W)301-982-4600
(WF)301-982-4605
lidl@sura.net

Hock-Koon Lim
Case Western Reserve University
(W)216-368-2982
(WF)216-368-4693
lim@po.cwru.edu

Arthur Lin
Pacific Bell
(W)510-823-1480
(WF)510-867-1405
yalin@srv.pacbell.com

Kenneth Lindahl
Data Communication and
Network Services
(W)510-642-0866
(WF)510-643-5385
lindahl@violet.berkeley.edu

David Lindemulder
NCR
(W)908-576-2276
(WF)908-576-2276
dcl@mtung.att.com

John Linn
Digital Equipment Corporation
(W)508-486-7384
(WF)508-486-6014
linn@erlang.enet.dec.com

Susan Lin
IBM Corporation
(W)203-783-4561
(WF)203-783-5046
suelin@ibm.com

Tony Li
cisco Systems
(W)415-688-8186
(WF)415-688-8282
tli@cisco.com

Joshua Littlefield
Cayman Systems, Inc.
(W)617-494-1999
(WF)617-494-5167
josh@cayman.com

Brian Lloyd
Lloyd & Associates
(W)916-676-1147
brian@lloyd.com

Daniel Long
Bolt Beranek and Newman
(W)617-873-2766
(WF)617-873-5620
long@nic.near.net

E. Paul Love
San Diego Supercomputer Center
(W)619-534-5043
(WF)619-534-5152
loveep@sdsc.edu

Steven Lunt
Bell Communications Research
(W)908-699-4244
(WF)908-336-2943
lunt@bellcore.com

Triet Lu
SAIC
(W)703-448-1683
trietl@sparta.com

Bryan Lyles
Xerox Corporation
(W)415-812-4466
(WF)415-812-4471
lyles@parc.xerox.com

Clifford Lynch
University of California
(W)510-987-0522
(WF)510-839-3573
calur@uccmvsa.ucop.edu

John Lynn
Cornell University
(W)607-255-5510
(WF)607-255-5771
lynn@jalps2.cit.cornell.edu

Patrice Lyons
Attorney at Law
(W)202-939-9666
(WF)202-232-4757

Carl Madison
Star-Tek, Inc.
(W)508-393-9393
(WF)508-393-6934
carl@startek.com

Kent Malave
IBM Corporation
(W)512-838-3486
(WF)512-838-3509
kent@chang.austin.ibm.com

Andrew Malis
Bolt Beranek and Newman
(W)617-873-3419
(WF)617-873-4086
malis@bbn.com

Gary Malkin
Xylogics, Inc.
(W)617-272-8140
gmalkin@xylogics.com

Bill Manning
Rice University
(W)713-285-5415
(WF)713-527-6099
bmanning@rice.edu

Michael Marcinkevicz
CSUnet
(W)714-773-3009
(WF)714-773-3990
mdm@csu.net

Scott Marcus
Bolt Beranek and Newman
(W)617-873-3075
(WF)617-873-4086
smarcus@bbn.com

April Marine
SRI International
(W)415-859-5318
(WF)415-859-6028
april@nisc.sri.com

David Martin
Fermi National Accelerator Laboratory
(W)708-840-8275
(WF)708-840-2783
dem@fnal.fnal.gov

Larry Masinter
Xerox Corporation
(W)415-812-4365
(WF)415-812-4333
masinter@parc.xerox.com

John Mason
Apple Computer, Inc.
(W)408-974-6194
(WF)408-862-1159
MASON1@applelink.apple.com

Matt Mathis
Pittsburgh Supercomputing Center
(W)412-268-7034
(WF)412-268-5832
mathis@psc.edu

Daniel Matzke
CERFnet
(W)619-455-3909
(WF)619-455-3990
matzked@cerf.net

Richard McBride
Axon Networks, Inc.
(W)+44 31 664 0409

Mark McCahill
University of Minnesota
(W)612-625-1300
(WF)612-625-6817
mpm@boombox.micro.umn.edu

Keith McCloghrie
Hughes LAN Systems
(W)415-966-7934
(WF)415-960-3738
kzm@hls.com

Ellen McDermott
Digital Equipment Corporation
(W)617-621-7266
(WF)617-225-2782
emcd@osf.org

Evan McGinnis
3Com Corporation
(W)408-764-6064
(WF)408-764-6135
bem@3com.com

Glenn McGregor
Merit Network, Inc.
(W)313-936-3000
(WF)313-747-3745
ghm@merit.edu

Leo McLaughlin
FTP Software, Inc.
(W)415-543-9001
(WF)415-543-9002
ljm@wco.ftp.com

Donna McMaster
SynOptics Communications, Inc.
(W)408-764-1206
(WF)408-764-1299
mcmaster@synoptics.com

Milo Medin
NASA
(W)415-604-6440
(WF)415-604-0063
medin@nsipo.nasa.gov

Ed Menze
University of Arizona
(W)602-621-4889
(WF)602-621-4246
menze@cs.arizona.edu

Marjo Mercado
Hewlett-Packard
(W)408-447-2826
(WF)408-447-3660
marjo@hpda.cup.hp.com

Greg Merrell
Digital Equipment Corporation
(W)415-691-4658
merrell@caldec.enet.dec.com

Donald Merritt
Ballistic Research Laboratory
(W)410-278-6808
(WF)301-278-5075
don@brl.mil

Paul Milazzo
Bolt Beranek and Newman
(W)617-873-4325
(WF)617-873-5620
milazzo@bbn.com

David Miller
The MITRE Corporation
(W)617-271-3993
(WF)617-271-8771
dtm@mitre.org

Henry Miller
COMSIS Corporation
(W)916-973-1685
(WF)916-488-8012
henrym@sacusr.mp.usbr.gov

Scott Miller
Qualcomm, Inc.
(W)619-597-5030
(WF)619-452-9096
smiller@qualcomm.com

Cyndi Mills
Bolt Beranek and Newman
(W)617-873-4143
(WF)617-873-5620
cmills@bbn.com

Greg Minshall
Novell, Inc.
(W)510-975-4507
(WF)510-975-1238
minshall@wc.novell.com

Rohit Mital
ProTools, Inc.
(W)503-645-5400
(WF)503-645-3577
rm@protools.com

Pushpendra Mohta
CERFnet
(W)619-455-3908
(WF)619-455-3990
pushp@cerf.net

Daniel Molinelli
TRW Space and Defense
(W)310-812-2161
(WF)213-812-7111
moline@gumby.dsd.trw.com

Lynn Monsanto
Sun Microsystems, Inc.
(W)415-336-2946
(WF)415-336-3156
monsanto@sun.com

Keith Moore
University of Tennessee
(W)615-974-3735
(WF)615-974-4404
moore@cs.utk.edu

John Moreland
San Diego Supercomputer Center
(W)619-534-5170
(WF)619-534-5152
mareland@sdsc.edu

Robert Morgan
Stanford University
(W)415-723-9711
(WF)415-723-0010
morgan@jessica.stanford.edu

Dennis Morris
Defense Information
Systems Agency
(W)703-487-3329
(WF)703-487-3190
morrisd@imo-uvax.dca.mil

Donald Morris
National Center for
Atmospheric Research
(W)303-497-1282
(WF)303-497-1137
morris@ucar.edu

John Moy
Proteon, Inc.
(W)508-898-2800
(WF)508-366-8901
jmoy@proteon.com

Chris Myers
Washington University
(W)314-935-7390
(WF)314-935-7302
chris@wugate.wustl.edu

Rina Nathaniel
RAD Network Devices
(W)+972 3 5447361
(WF)+972 3 5447368
rina!rnd!rndi@uunet.uu.net

Mark Needleman
University of California
(W)510-987-0530
(WF)510-839-3573
mhn@stubbs.ucop.edu

Clifford Neuman
Information Sciences Institute
(W)310-822-1511
bcn@isi.edu

Michael Newell
PRC, Inc.
(W)202-453-8752
(WF)202-755-2637
mnewell@nhqvax.hg.nasa.gov

Andy Nicholson
Cray Research, Inc.
(W)612-683-5473
(WF)612-683-5599
droid@cray.com

William Nichols
Digital Equipment Corporation
(W)508-486-5266
(WF)508-486-7417
nichols@took.enet.dec.com

Torben Nielsen
University of Hawaii
(W)808-956-3499
(WF)808-956-3548
torben@hawaii.edu

Chandy Nilakantan
3Com Corporation
(W)408-764-5213
(WF)408-764-5003
csn@3com.com

William Nowicki
Legato Systems, Inc.
(W)415-329-7880
(WF)415-329-8898
nowicki@legato.com

Chan-Ok O'Brien
DISA
(W)703-487-8675
(WF)703-487-8677
chan@bilbo.jdssc.dca.mil

Michael O'Dell
Bell Communications Research
(W)201-829-5128
(WF)201-455-1931
mo@bellcore.com

Ari Ollikainen
Lawrence Livermore
National Laboratory
(W)510-423-5962
(WF)510-422-0435
ari@es.net

Ayal Opher
SynOptics Communications, Inc.
(W)408-764-1559
aopher@synoptics.com

Alan Oppenheimer
Apple Computer, Inc.
(W)408-974-4744
(WF)408-862-1159
oppenheime1@applelink.apple.com

David Oran
Digital Equipment Corporation
(W)508-486-7377
(WF)508-486-5279
oran@sneezy.lkg.dec.com

Mely Pacheco
NEC America
(W)408-433-2946
mely@tdd.sj.nec.com

J. Bradford Parker
Cayman Systems, Inc.
(W)617-494-1999
(WF)617-494-5167
brad@cayman.com

Emmanuel Pasetes
Premenos Corporation
(W)510-602-2000
(WF)510-602-2133
ekp@enlil.premenos.sf.ca.us

Brad Passwaters
SURAnet
(W)301-982-4600
(WF)301-982-4605
bjp@sura.net

Rick Patella
San Diego Supercomputer Center
(W)619-534-5095
(WF)619-534-5152
patellar@sdsc.edu

Rakesh Patel
Rutgers University
(W)908-932-3436
(WF)908-932-2968
rpatel@rutgers.edu

Tushar Patel
Timeplex, Inc.
(W)714-770-1221
(WF)714-583-7509
tpatel@doelz.temeplex.com

Edison Paw
3Com Corporation
(W)415 940-7663
esp@3com.com

Geir Pedersen
University of Oslo
(W)+47 2 453470
(WF)+47 2 453730
geir.pedersen@usit.uio.no

David Perkins
SynOptics Communications, Inc.
(W)408-764-1516
(WF)408-988-5525
dperkins@synoptics.com

Radia Perlman
Digital Equipment Corporation
(W)508-486-7648
(WF)508-486-5279
perlman@radia.enet.dec.com

Marsha Perrott
PREPnet
(W)412-268-7870
(WF)412-268-7875
mlp+@andrew.cmu.edu

Roy Perry
US West Advanced Technologies
(W)303-541-6245
(WF)303-541-6276
rperry@uswat.uswest.com

David Piscitello
Bell Communications Research
(W)908-758-2286
(WF)908-758-4177
dave@sabre.bellcore.com

Mel Pleasant
Rutgers University
(W)908-932-2023
(WF)908-932-2023
pleasant@rutgers.edu

Valerie Polichar
University of California,
San Diego
(W)619-534-1859
(WF)619-534-1746
valerie@ucsd.edu

Jack Pope
University of San Diego
(W)619-260-4810
(WF)619-260-4619
pope@uscdsv.acusd.edu

Jon Postel
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
postel@isi.edu

Lars Poulsen
Rockwell CMC
(W)805-968-4262
(WF)805-968-8256
lars@cmc.com

Venkat Prasad
3Com Corporation
(W)408-764-5226
(WF)408-764-5003
vsp@3com.com

David Preston
Epilogue Technology Corporation
(W)805-650-7107
(WF)805-650-7108
djp@epilogue.com

Stephanie Price Marasciullo
Rockwell CMC
(W)805-968-4262
(WF)805-968-8256
price@cmc.com

Rex Pugh
Hewlett-Packard
(W)916-785-4471
(WF)916-786-9185
pugh@hprnd.rose.hp.com

Scott Quinn
Bell Communications Research
(W)201-829-4147
(WF)201-829-5976
smq@aries.bellcore.com

Paul Raison
Xylogics, Inc.
(W)617-272-8140
(WF)617-273-5392
raison@xylogics.com

P. Rajaram
SunSoft, Inc.
(W)415-336-5194
(WF)415-336-3156
rajaram@eng.sun.com

K. Ramakrishnan
Digital Equipment Corporation
(W)508-486-7267
(WF)508-486-5279
rama@erlang.enet.dec.com

Christopher Ranch
Novell, Inc.
(W)408-473-8667
(WF)408-946-3447
cranch@novell.com

Joe Random
Novell, Inc.

James Rees
University of Michigan
(W)313-763-4174
(WF)313-763-4434
jim.rees@umich.edu

Michael Reilly
Digital Equipment Corporation
(W)415-688-1323
(WF)415-324-2797
reilly@nsl.dec.com

Robert Reilly
K12NET
(W)413-443-0027
rreilly@athena.mit.edu

Jim Reinstedler
Ungermann-Bass, Inc.
(W)408-562-5636
(WF)408-970-7389
jimr@ub.com

Yakov Rekhter
IBM Corporation
(W)914-945-3896
yakov@watson.ibm.com

Robert Reschly
U.S. Army Ballistics
Research Laboratory
(W)410-278-6808
(WF)410-278-5075
reschly@brl.mil

Joyce K. Reynolds
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
jkrey@isi.edu

Mike Ritter
Apple Computer, Inc.
(W)408-862-8088
(WF)408-257-7982
mwritter@applelink.apple.com

Michael Roberts
EDUCOM
(W)202-872-4200
(WF)202-872-4318
roberts@educom.edu

Kary Robertson
Concord Communications, Inc.
(W)508-460-4646
(WF)508-481-9772
kr@concord.com

Ron Roberts
Stanford University
(W)415-723-3352
roberts@jessica.stanford.edu

Sam Roberts
Farallon Computing, Inc.
(W)510-596-9215
(WF)510-596-9024
sroberts@farallon.com

Jonathan Rodin
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
rodin@ftp.com

Benny Rodrig
RAD Network Devices
(W)+972 3 6458555
(WF)+972 3 5447368
4373580@mcimail.com

Manoel Rodrigues
AT&T Bell Laboratories
(W)908-949-4655
(WF)908-949-1726
manoel.rodrigues@att.com

Guenter Roeck
Conware Computer Consulting
(W)+49 721 9495 0
(WF)+49 721 9495 130
roeck@conware.de

Allyn Romanow
Sun Microsystems, Inc.
(W)415-336-1326
(WF)415-336-3156
allyn@eng.sun.com

Dan Romascanu
Lannet Data Communications
(W)+972 3 6458414
(WF)+972 3 5447146
dan@lannet.com

Marshall Rose
Dover Beach Consulting, Inc.
(W)415-968-1052
(WF)415-968-2510
mrose@dbc.mtview.ca.us

Kevin Rowett
Tandem Computers, Inc.
(W)408-285-4325
(WF)408-285-4324
kevinr@tandem.com

Rick Royston
U.S. Robotics
(W)708-933-5430
(WF)703-982-1348
rick@lsumvs.sncc.lsu.edu

Allan Rubens
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
acr@merit.edu

Gregory Ruth
Bolt Beranek and Newman
(W)617-873-3150
(WF)617-873-8202
gruth@bbn.com

Anthony Rutkowski
US Sprint
(W)703-689-5080
(WF)703-689-6715
amr@sprint.com

L. Michael Sabo
SSDS, Inc.
(W)303-790-0660
(WF)303-790-1663
lms@denver.ssds.com

Timothy Salo
Minnesota Supercomputer Center
(W)612-626-0347
(WF)612-624-6550
tjs@msc.edu

Tom Sandoski
MCNC
(W)919-248-1104
(WF)919-248-1405
tom@concert.net

Sharad Sanghi
Advanced Network & Services, Inc.
(W)914-789-5332
(WF)914-789-5310
sharad@ans.net

Jon Saperia
Digital Equipment Corporation
(W)508-496-8333
(WF)508-496-9929
saperia@tcpjon.ogo.dec.com

Miguel Sasson
Xylogics, Inc.
(W)617-272-0924
(WF)617-272-3159
sasson@xylogics.com

Andreas Schachtner
University of Dortmund
(W)+49 231 755 5135
(WF)+49 231 755 2386
afs@germany.eu.net

Mark Schaefer
David Systems, Inc.
(W)408-720-6804
(WF)408-720-1337
schaefer@davidsys.com

Jeffrey Schiller
Massachusetts Institute of Technology
(W)617-253-0161
(WF)617-258-8736
jis@mit.edu

Eve Schooler
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
schooler@isi.edu

Michael Schwartz
University of Colorado, Boulder
(W)303-492-3902
(WF)303-492-2844
schwartz@cs.colorado.edu

Tim Seaver
MCNC
(W)919-248-1973
(WF)919-248-1455
tas@concert.net

John Seligson
Ultra Network Technologies
(W)408-922-0100
(WF)408-433-9287
johns@ultra.com

John Sherburne
Sprint International
(W)703-689-6297
(WF)703-689-5644
john.sherburne@sprintintl.sprint.com

Erik Sherk
SURAnet
(W)301-982-4600
(WF)301-982-4605
sherk@sura.net

Brian Shiflett
Sprint Government
Systems Division
(W)202-416-6213
(WF)202-833-8583
rcollett@icmi.icp.net

Robert Shirey
The MITRE Corporation
(W)703-883-7210
(WF)703-883-1397
shirey@mitre.org

Tom Simons
Motorola GEG
(W)602-441-7510
(WF)602-441-5810
nes!simons@phx.mcd.mot.com

William Simpson
Computer Systems
Consulting Services
(W)313-543-3269
Bill\verb+_+Simpson@um.cc.umich.edu

Anil Singhal
Frontier Software Development, Inc.
(W)508-851-8872
(WF)508-851-6956
rnadkarn@cs.ulowell.edu

Surinder Singh
ICL
(W)714-458-7282
(WF)714-458-6257
singh@irv.icl.com

Ursula Sinkewicz
Digital Equipment Corporation
(W)508-486-5604
(WF)508-486-7417

Sam Sjogren
TGV
(W)408-427-4366
(WF)408-427-4365
sjogren@tgv.com

Keith Sklower
University of California,
Berkeley
(W)510-642-9587
sklower@cs.berkeley.edu

Mark Sleeper
SPARTA, Inc.
(W)703-448-0210
(WF)703-734-3323
mws@sparta.com

Timon Sloane
PEER Networks
(W)408-727-4111
(WF)408-474-4410
peernet!timon@uunet.uu.net

Eric Smith
Telebit Corporation
(W)408-745-3129
(WF)408-745-3810
eric@telebit.com

Mark Smith
University of Michigan
(W)313-764-2277
(WF)313-764-5140
mcs@umich.edu

Michael Smith
Home Box Office, Inc.
(W)212-512-5206
(WF)212-512-7049
mes5674@hertz.njit.edu

Patricia Smith
Merit Network, Inc.
(W)313-936-3335
(WF)313-747-3745
psmith@merit.edu

Frank Solensky
Clearpoint Research Corporation
(W)508-435-7563
(WF)508-435-7504
solensky@clearpoint.com

Evan Solley
Farallon Computing, Inc.
(W)503-292-1374
(WF)503-226-6543
solley@farallon.com

Karen Sollins
Massachusetts Institute of Technology
(W)617-253-6006
(WF)617-258-8682
sollins@lcs.mit.edu

Pyda Srisuresh
3Com Corporation
(W)408-764-6968
suresh@3com.com

Michael St. Johns
Department of Defense
(W)301-688-6741
stjohns@umd5.umd.edu

Martha Steenstrup
Bolt Beranek and Newman
(W)617-873-3192
(WF)617-491-1496
msteenst@bbn.com

Einar Stefferud
Network Management Associates, Inc.
(W)714-842-3711
(WF)714-848-2091
stef@nma.com

Brad Steinka
Microcom, Inc.
(W)617-551-1209
(WF)617-551-1015
brad@python.eng.microcom.com

Scott Stein
Apple Computer, Inc.
(W)408-974-3124
scotts@apple.com

Bob Stewart
Xyplex, Inc.
(W)508-264-9900
(WF)508-264-9930
rlstewart@eng.xyplex.com

Walter Stickle
FTP Software, Inc.
(W)617-246-0900
(WF)617-245-7943
wls@ftp.com

Bernhard Stockman
NORDUnet
(W)+46 8 790 6519
(WF)+46 8 241 179
boss@sunet.se

William Streilein
Wellfleet Communications, Inc.
(W)617-275-7750
bstreile@wellfleet.com

Tatsuya Suda
University of California,
Irvine
(W)714-856-5474
(WF)714-856-4056
suda@ics.uci.edu

Craig Summerhill
Coalition for
Networked Information
(W)202-232-2466
(WF)202-462-7849
craig@cni.org

Michael Swan
Neon Software, Inc.
(W)510-283-9771
neon.sw@applelink.apple.com

R.B. Swenson
Enigma Logic
(W)408-425-8523
(WF)408-425-8533
rswenson@igc.org

Wayne Tackabury
Pacer Software, Inc.
(W)508-898-3300
(WF)508-366-1356
wft@pacersoft.com

Sally Tarquinio
The MITRE Corporation
(W)703-883-7987
(WF)703-883-7142
sally@gateway.mitre.org

Kaj Tesink
Bell Communications Research
(W)908-758-5254
(WF)908-758-4196
kaj@nvuxr.cc.bellcore.com

Mark Therieau
Microcom, Inc.
(W)617-551-1195
(WF)617-551-1015
markt@python.eng.microcom.com

Bill Thomas
Novell, Inc.
(W)408-473-8514
(WF)408-946-3447
jthomas@novell.com

Chris Thomas
Intel Corporation
chris\verb+_+thomas@ccm.hf.intel.com

Dave Thompson
Microsoft Corporation
(W)206-868-3886
(WF)206-936-7329
daveth@ingate.microsoft.com

Geoffrey Thompson
SynOptics Communications, Inc.
(W)408-764-1339
(WF)408-988-5525
thompson@synoptics.com

Steve Thompson
Optical Data Systems, Inc.
(W)214-234-6400
(WF)214-234-4059
steve@ods.com

Richard Thomsen
Los Alamos National Laboratory
(W)505-667-4210
(WF)505-665-4361
rgt@lanl.gov

Dean Throop
Data General Corporation
(W)919-549-8421
(WF)919-248-5942
throop@dg-rtp.dg.com

Tom Tignor
Information Sciences Institute
(W)310-822-1511
(WF)310-823-6714
tpt%@isi.edu

Claudio Topolcic
Corporation for National Research Initiatives
(W)703-620-8990
(WF)703-620-0913
topolcic@nri.reston.va.us

Chuck Townsend
Cabletron Systems, Inc.
(W)603-332-9400
(WF)603-332-1019
townsend@ctron.com

Paul Traina
cisco Systems
(W)415-688-8189
(WF)415-325-4328
pst@cisco.com

Glenn Trewitt
Digital Equipment Corporation
(W)415-688-1324
(WF)415-324-2797
trewitt@pa.dec.com

Mike Truskowski
cisco Systems
(W)314-275-9174
(WF)314-542-3176
truskowski@cisco.com

Theodore Tso
Massachusetts Institute of Technology
(W)617-253-8091
(WF)617-258-8736
tytso@mit.edu

Paul Tsuchiya
Bell Communications Research
(W)201-829-4484
(WF)201-829-2504
tsuchiya@thumper.bellcore.com

Chiaohuey Tung
Intel Corporation
(W)503-696-2214
(WF)503-640-2484
tung@ashland.intel.com

Maurice Turcotte
Racal Datacom
(W)404-804-5190
(WF)404-804-5190
dnmrt@interlan.com

Dono van-Mierop
3Com Corporation
(W)408-764-5930
(WF)408-764-5001
dono\verb+_+vanmierop@3mail.3com.com

Kannan Varadhan
Ohio Academic Resources Network
(W)614-292-4137
(WF)614-292-7168
kannan@oar.net

Gregory Vaudreuil
Corporation for National Research Initiatives
(W)703-620-8990
(WF)703-620-0913
gvaudre@nri.reston.va.us

John Veizades
Apple Computer, Inc.
(W)408-974-2672
(WF)408-974-9083
veizades@apple.com

Sandra Vest
Defense Information
Systems Agency
(W)703-487-3029
(WF)703-487-3190
vests@imo-uvax.dca.mil

Vernor Vinge
San Diego State University
(W)619-594-6191
vinge@sdsu.edu

Huyen Vu
Defense Information
Systems Agency
(W)703-487-3038
(WF)703-487-3190
vi@polaris.dca.mil

Lee Wade
Sterling Software
Federal Systems
(W)415-604-4789
(WF)415-604-0063
wade@nsipo.arc.nasa.gov

David Waitzman
Bolt Beranek and Newman
(W)617-873-4323
(WF)617-873-3776
djw@bbn.com

Steven Waldbusser
Carnegie Mellon University
(W)412-268-6628
(WF)412-268-4987
waldbusser@andrew.cmu.edu

Timothy Walden
ACC Systems
(W)301-290-8100
(WF)301-290-8106
tmwalden@saturn.acc.com

Justin Walker
Apple Computer, Inc.
(W)408-974-4427
(WF)408-974-0892
justin@apple.com

Kevin Walsh
University of California,
San Diego
(W)619-534-6571
(WF)619-534-6206
kw@ucsd.edu

Yuan Wang
Network Application Technology
(W)408-370-4366
(WF)408-370-4222
natadm!ycw@uunet.uu.net

Carol Ward
Westnet
(W)303-492-5860
(WF)303-492-4198
cward@westnet.net

James Watt
Newbridge Networks Corporation
(W)613-591-3600
(WF)613-591-3680
jamesw@newbridge.com

Chris Weider
Advanced Network & Services, Inc.
(W)313-663-2482
(WF)313-663-2927
weider@ans.net

Abel Weinrib
Bell Communications Research
(W)201-829-4454
(WF)201-984-5837
abel@bellcore.com

Frank Welch
Intel Corporation
(W)602-869-4059
(WF)602-869-3509
frank\verb+_+welch@ccm.hf.intel.com

Jonathan Wenocur
Shiva Corporation
(W)617-252-6300
(WF)617-252-6852
jhw@shiva.com

Jil Westcott
Bolt Beranek and Newman
(W)617-873-3148
(WF)617-873-3776
Westcott@BBN.COM

Jehu Westmark
Centel Corporation
(W)904-599-1843
(WF)904-599-1033
westmark@dsi.scri.fsu.edu

Moira West
Carnegie Mellon University
(W)412-268-7090
(WF)412-268-6989
mjw@cert.sei.cmu.edu

Jonathan Whaley
MCNC
(W)919-248-1165
(WF)919-248-1100
whaley@concert.net

Brien Wheeler
The MITRE Corporation
(W)617-271-2567
(WF)617-271-8771
blw@mitre.org

Chris Wheeler
University of Washington
(W)206-543-3282
(WF)206-543-3909
cwheeler@cac.washington.edu

Allen Wiersma
NASA Science Internet
(W)415-604-1031
(WF)415-604-0063
wiersma@nsipo.nasa.gov

Bert Wijnen
IBM International Operations
(W)+31 2975 53316
(WF)31 2975 62468
wijnen@vnet.ibm.com

Kathleen Wilde
Digital Equipment Corporation
(W)603-881-0387
(WF)508-486-5279
wilde@decvax.dec.com

Steve Willens
Livingston Enterprises, Inc.
(W)510-426-0770
(WF)510-426-8951
steve@livingston.com

Douglas Williams
IBM Corporation
(W)919-254-4801
dougw@ralumg.unet.ibm.com

Scott Williamson
Network Solutions, Inc.
(W)703-802-8405
(WF)703-802-8376
scottw@nic.ddn.mil

Preston Wilson
INTERACTIVE Systems Corporation
(W)708-505-9100
(WF)708-505-9133
preston@i88.isc.com

Walter Wimer
Carnegie Mellon University
(W)412-268-6252
(WF)412-268-4987
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
(W)708-972-7236
(WF)708-972-5983
lwinkler@anl.gov

Cathy Wittbrodt
Lawrence Livermore
National Laboratory
(W)415-466-4016
(WF)415-422-0435
cjw@nersc.gov

Stephen Wolff
National Science Foundation
(W)202-357-9717
(WF)202-357-0320
steve@nsf.gov

Steven Wong
Digital Equipment Corporation
(W)508-486-5344
(WF)508-486-7417
wong@took.enet.dec.com

Paul Woodruff
3Com Corporation
(W)+44 442 231000
(WF)+44 442 254701
biccdn!paulw@eros.uknet.ac.uk

Russ Wright
Lawrence Berkeley Laboratory
(W)510-486-6965
(WF)510-486-6363
wright@lbl.gov

Michael Wrona
Far West Laboratory
(W)415-241-2709
(WF)415-565-3012
mwrona@eis.calstate.edu

Brian Wyld
Spider Systems Limited
(W)+44 031 554 9424
(WF)+44 031 554 0649
brianw@spider.co.uk

Peter Yee
NASA Ames Research Center
(W)415-604-3812
(WF)415-604-6999
yee@ames.arc.nasa.gov

Frances Yeh
Jet Propulsion Lab
(W)818-393-6992
(WF)818-354-5629
fyeh@rafael.jpl.nasa.gov

Wengyik Yeong
Performance Systems
International, Inc.
(W)703-620-6651
(WF)703-620-4586
yeongw@psi.com

Henry Yip
Network Application Technology
(W)408-370-4201
(WF)408-370-4222
natadm!henry@uunet.uu.net

Ki-Sung Yoo
KIST/SERI
(W)+82 42 829 1322
(WF)+82 42 861 1999
ksyu@garam.kreonet.re.kr

Chin Yuan
Pacific Bell
(W)510-867-5016
(WF)510-867-1405
cxyuan@pacbell.com

Jessica Yu
Merit Network, Inc.
(W)313-936-2655
(WF)313-747-3745
jyy@merit.edu

Bahman Zargham
Tandem Computers, Inc.
(W)408-285-4358
(WF)408-285-4240
zargham\verb+_+bahman@tandem.com

Paul Zawada
National Center for
Supercomputing Applications
(W)217-333-1189
(WF)217-244-7396
zawanda@ncsa.uiuc.edu