


Proceedings of the Twenty-Seventh Internet Engineering Task Force

SURFnet and RARE
Amsterdam, The Netherlands
July 12-16, 1993

Corporation
for
National
Research
Initiatives



PROCEEDINGS OF THE
TWENTY-SEVENTH
INTERNET ENGINEERING
TASK FORCE

SURFnet and RARE
Amsterdam, The Netherlands
12-16 July 1993

Compiled and Edited by
Debra J. Legare
John W. Stewart, III
Lois J. Keiper

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

This material is based on work supported by the National Science Foundation under grant number NCR 8820945.



"On the Internet, nobody knows you're a dog."

Drawing by P. Steiner; © 1993
The New Yorker Magazine, Inc.

ACKNOWLEDGEMENTS

The Internet Engineering Task Force held its first meeting outside of the North American Continent from July 12 through 16 at the RAI Congress Centre in Amsterdam, The Netherlands. Nearly 500 folks were present. We anticipated a change in demographics and were not disappointed with the final count reflecting attendance almost evenly split between US and international participants.

Erik Huizer of SURFnet was our Local Host for this meeting. Erik and a group of other SURFnet staff put together a terrific terminal room for the attendees. Equipment was generously provided by the following organizations:

| | |
|-------------------------------|--|
| cisco Systems | Routers and Terminal Server |
| Digital Equipment BV | Terminals |
| Geveke Electronics BV | Routers and Terminal Server |
| RIPE NCC | Manpower for Unix Server and AVT Systems |
| SARA | Cabling and Manpower |
| Sun Microsystems Nederland BV | Workstations |
| SURFNET | Manpower, Equipment and Funding |
| Unisource Bussiness Networks | 2Mb/s access line |

Both SURFnet and RARE provided invaluable assistance with logistics arrangements both prior to the meeting and while on-site. Thanks especially to Marieke Dekker, Anne Cozanet and Josefien Bersee for being so helpful to the Secretariat.

Presentations began first thing on Monday morning. Erik Huizer opened with the "IETF from a European Perspective" and was followed by Chuck McManis with a presentation on "Advances on ONC." Daniel Karrenberg closed Monday's session with a presentation on "Setting up a Routing Registry in Europe." Other presentations throughout the week included "Electronic Cash: Theory Towards Application" by David Chaum, "JVTOS for Workstations" by Bernhard Plattner and "Service Location" by John Veizades. The IETF concluded on Friday with a series of presentations. The first was by Marshall Rose on "An Experiment in Remote Printing." Marshall's presentation was followed by status reports by the IPng groups given by Steve Deering

(SIP), Paul Francis (PIP), Mark Knopper (TUBA) and Vladimir Sukonnik (TP/IX). Thanks to all the presentors for their excellent work which contributed to the success of this meeting.

We also wish to thank the staff at the RAI for helping to ensure that our first European meeting ran smoothly. The facility itself helped to ease the work of the Secretariat in particular those individuals who work behind the Registration Desk. Without exception, we are indebted to Debra Legare, Terry Weigler and Cynthia Clark for all their hard work.

Megan Davies Walnut
IETF Meeting Coordinator

Contents

| | |
|---|-----------|
| Director's Message | 1 |
| Agenda of the Twenty-Seventh IETF | 7 |
| 1 IETF Overview | 17 |
| 1.1 Future IETF Meeting Sites | 21 |
| 1.2 On Line IETF Information | 23 |
| 1.3 Guidelines to Authors of Internet-Drafts | 27 |
| 2 Area and Working Group Reports | 29 |
| 2.1 Applications Area | 31 |
| 2.1.1 Interactive Mail Access Protocol (IMAP) | 59 |
| 2.1.2 Internet Message Extensions (822EXT) | 63 |
| 2.1.3 MIME-MHS Interworking (MIMEMHS) | 65 |
| 2.1.4 Network News Transport Protocol (NNTP) | 67 |
| 2.1.5 OSI Directory Services (OSIDS) | 69 |
| 2.1.6 TELNET (TELNET) | 77 |
| 2.1.7 Telnet TN3270 Enhancements (TN3270E) | 81 |
| 2.1.8 X.400 Operations (X400OPS) | 83 |
| 2.2 Internet Area | 87 |
| 2.2.1 Dynamic Host Configuration (DHC) | 111 |
| 2.2.2 IP Over AppleTalk (APPLEIP) | 113 |
| 2.2.3 IP Over Asynchronous Transfer Mode (ATM) | 115 |
| 2.2.4 IP Over Large Public Data Networks (IPLPDN) | 121 |
| 2.2.5 P. Internet Protocol (PIP) | 125 |
| 2.2.6 Point-to-Point Protocol Extensions (PPPEXT) | 133 |
| 2.2.7 Router Requirements (RREQ) | 149 |
| 2.2.8 Simple Internet Protocol (SIP) | 151 |
| 2.2.9 TCP/UDP Over CLNP-Addressed Networks (TUBA) | 167 |
| 2.2.10 TP/IX (TPIX) | 179 |
| 2.3 Network Management Area | 185 |
| 2.3.1 ATM MIB (ATOMMIB) | 195 |
| 2.3.2 Bridge MIB (BRIDGE) | 201 |

| | | |
|--------|--|-----|
| 2.3.3 | Character MIB (CHARMIB) | 203 |
| 2.3.4 | DECnet Phase IV MIB (DECNETIV) | 205 |
| 2.3.5 | FDDI MIB (FDDIMIB) | 207 |
| 2.3.6 | Frame Relay Service MIB (FRNETMIB) | 209 |
| 2.3.7 | Interfaces MIB (IFMIB) | 213 |
| 2.3.8 | Mail and Directory Management (MADMAN) | 217 |
| 2.3.9 | Modem Management (MODEMMGT) | 221 |
| 2.3.10 | Remote LAN Monitoring (RMONMIB) | 229 |
| 2.3.11 | SNA DLC Services MIB (SNADLC) | 231 |
| 2.3.12 | SNA NAU Services MIB (SNANAU) | 233 |
| 2.3.13 | Token Ring Remote Monitoring (TRMON) | 237 |
| 2.3.14 | Uninterruptible Power Supply (UPSMIB) | 243 |
| 2.4 | Operational Requirements Area | 247 |
| 2.4.1 | BGP Deployment and Application (BGPDEPL) | 253 |
| 2.4.2 | Benchmarking Methodology (BMWG) | 261 |
| 2.4.3 | Network Joint Management (NJM) | 263 |
| 2.4.4 | Network OSI Operations (NOOP) | 265 |
| 2.4.5 | Operational Statistics (OPSTAT) | 267 |
| 2.4.6 | User Connectivity (UCP) | 271 |
| 2.5 | Routing Area | 273 |
| 2.5.1 | Border Gateway Protocol (BGP) | 277 |
| 2.5.2 | IP Routing for Wireless/Mobile Hosts (MOBILEIP) | 285 |
| 2.5.3 | IS-IS for IP Internets (ISIS) | 299 |
| 2.5.4 | Inter-Domain Multicast Routing (IDMR) | 303 |
| 2.5.5 | Inter-Domain Policy Routing (IDPR) | 307 |
| 2.5.6 | Multicast Extensions to OSPF (MOSPF) | 309 |
| 2.5.7 | OSI IDR for IP Over IP (IPIDRP) | 311 |
| 2.5.8 | Open Shortest Path First IGP (OSPF) | 315 |
| 2.5.9 | RIP Version II (RIPV2) | 317 |
| 2.5.10 | Source Demand Routing (SDR) | 323 |
| 2.6 | Security Area | 329 |
| 2.6.1 | Authorization and Access Control (AAC) | 335 |
| 2.6.2 | Commercial Internet Protocol Security Option (CIPSO) | 341 |
| 2.6.3 | Common Authentication Technology (CAT) | 343 |
| 2.6.4 | Internet Protocol Security Protocol (IPSEC) | 351 |
| 2.6.5 | Network Access Server Requirements (NASREQ) | 355 |
| 2.6.6 | Privacy-Enhanced Electronic Mail (PEM) | 359 |
| 2.7 | Service Applications Area | 367 |
| 2.7.1 | Domain Name System (DNS) | 373 |
| 2.7.2 | MHS-DS (MHSDS) | 381 |
| 2.7.3 | Minimal OSI Upper-Layers (THINOSI) | 389 |
| 2.7.4 | Network Database (NETDATA) | 395 |

| | | |
|----------|--|------------|
| 2.7.5 | Network Printing Protocol (NPP) | 397 |
| 2.7.6 | Service Location Protocol (SVRLOC) | 399 |
| 2.7.7 | Trusted Network File Systems (TNFS) | 403 |
| 2.8 | Transport Area | 405 |
| 2.8.1 | Audio/Video Transport (AVT) | 415 |
| 2.8.2 | Multiparty Multimedia Session Control (MMUSIC) | 417 |
| 2.8.3 | TCP Large Windows (TCPLW) | 431 |
| 2.9 | User Services Area | 437 |
| 2.9.1 | Integrated Directory Services (IDS) | 443 |
| 2.9.2 | Integration of Internet Information Resources (IIR) | 449 |
| 2.9.3 | Internet Anonymous FTP Archives (IAFA) | 455 |
| 2.9.4 | Internet School Networking (ISN) | 457 |
| 2.9.5 | Network Information Services Infrastructure (NISI) | 463 |
| 2.9.6 | Network Training Materials (TRAINMAT) | 469 |
| 2.9.7 | Networked Information Retrieval (NIR) | 475 |
| 2.9.8 | Uniform Resource Identifiers (URI) | 481 |
| 2.9.9 | User Documents Revisions (USERDOC2) | 487 |
| 2.9.10 | User Services (USWG) | 491 |
| 2.9.11 | Whois and Network Information Lookup Service (WNILS) | 497 |
| 3 | IAB Open Meeting | 503 |
| 3.1 | IAB Open Meeting | 505 |
| 4 | IETF Protocol Presentations | 509 |
| 4.1 | IPng - SIP | 511 |
| 4.2 | IPng - PIP | 515 |
| 4.3 | IPng - TUBA | 519 |
| 4.4 | IPng - TPIX | 523 |
| 5 | Technical Presentations | 525 |
| 5.1 | The IETF from a European Perspective | 527 |
| 5.2 | Advances on ONC | 533 |
| 5.3 | Setting Up a Routing Registry in Europe | 539 |
| 5.4 | Electronic Cash: Theory To Application | 545 |
| 5.5 | JVTOS for Workstations | 549 |
| 5.6 | Service Location | 553 |
| 5.7 | An Experiment in Remote Printing | 557 |
| 6 | Workshop for Working Group Chairs | 563 |
| 6.1 | Workshop for Working Group Chairs | 565 |
| A | Attendees | 571 |

Director's Message

The 27th meeting of the Internet Engineering Task Force was held in Amsterdam, The Netherlands, July 12-16, 1993. The meeting was co-hosted by SURFnet and RARE, and our thanks and appreciation go out to Erik Huizer and to all the others that helped with the terminal room and the social event.

This was the first time an IETF meeting has been held outside of North America. The meeting was well attended with almost 500 attendees during the week, a little over the original estimates of 450 attendees made one year ago during the Cambridge meeting.

The ratio of non-US attendees was, as expected, significantly higher than at the past few meetings which have ranged from between 8% and 11%. For this meeting, 46% of the attendees were from outside the United States. The top five non-US countries were, in terms of the number of individuals attending:

| | |
|-----------------|--------------|
| The Netherlands | 55 attendees |
| United Kingdom | 30 attendees |
| Germany | 25 attendees |
| Sweden | 15 attendees |
| France | 14 attendees |

The number of first time attendees remained close to the 200 mark. There were 187 first time attendees at the Amsterdam meeting. About 110 people showed up for the Newcomers' Orientation on Sunday afternoon. Interestingly enough, the ratio of first-time attendees remained the same. Approximately 37% of the attendees were at their very first IETF meeting.

This is very exciting and encouraging. There might be a tendency to consider Amsterdam a "special" meeting (and indeed it was special) with respect to numbers and ratios, especially comparing it to previous meetings. The total attendance was less than the normal 600, but not by much. While there were not 200 first timers, there were 187. Instead of 150 attendees at the Newcomers' Orientation, there were 110. But we did have almost 500 attendees, and the percentage of first timers was the same as at previous meetings. All things considered, this was just another IETF meeting. I am looking forward to other IETF meetings held outside North America.

From the multicast perspective, things keep growing and growing, and it has become impossible to guess how many people are listening in during the IETF meetings. There were more than 400 hosts on the receiving end of the IETF multicast effort,

up from approximately 330 host at the Columbus meeting. The number of countries "listening in" rose to 16 for the meeting in Amsterdam: Austria, Australia, Canada, Switzerland, Czechoslovakia, Denmark, Estonia, Finland, France, Japan, Korea, The Netherlands, Norway, Sweden, the United Kingdom, and the United States. This technology is expanding the number of virtual attendees at these meetings, and it's not limited to the technical presentations and plenaries. The broadcast system is on wheels, and it is not uncommon to see the multicast volunteers (from the host group) wheeling a cart from meeting to meeting!

The Future of IETF Meetings

What will IETF meetings be like in 1995 and 1996? Anything like 1993? Certainly the planning has got to change, the logistic details will definitely change, and our meeting model may need to be modified. Will the terminal room requirements change as well as expand? What effect will multicasting have on our attendance levels? Where will we be meeting? What will the host be required to do? Will we even have a host?

It seems there are more and more groups being formed, all of whom meet at various times. It is almost impossible to avoid all conflicts when scheduling IETF meetings, and this is even more challenging as we attempt to avoid holidays.

The number of people attending IETF meetings continues to stay well over 600. Gone forever are the days of meeting on college campuses with easy to provide terminal rooms. In fact, gone are the days when the IETF could "take over" a hotel. We've grown so much that we are now at the small conference size for groups, and often have to share the facilities with other groups meeting at the same time. Previous IETF meeting sites may no longer be possible considering the size and levels of effort.

Remember the terminal room in Atlanta? That was when we had 350 attendees. Now think of the terminal rooms we've been used to since we hit the 600+ attendee mark. Now, think instead of the volunteer work provided by the local host group; imagine what *you* would need to do to set up a terminal room? Be honest... which features or capabilities might you be tempted to not provide?

Can we begin scheduling further and further into the future? Can we find semi-stable periods that we will always meet? How much do IETF meetings depend on the host? How can we reduce the workload for hosts while maintaining the excellent connectivity and variety of workstations we have come to expect?

This is what the Secretariat will be working on over the next few months... how we might be operating in the future, what the requirements will be, and what steps must be taken now to be in a position to accommodate those requirements when they become current.

Future Meetings

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

The first IETF meeting in 1994 will be held in Seattle, Washington the last week of March (March 28 - April 1, 1994). This meeting is being hosted by NorthWestNet.

It appears that the summer IETF meeting will be held in Toronto, Ontario, Canada. As I write this, a contract has not been signed, though negotiations are underway.

Note that information on future IETF meetings can be always be found in the file /ietf/0mtg-sites.txt which is located on the IETF shadow directories.

Stephen J. Coya
Executive Director, IETF

IETF Progress Report

The IESG and IETF have been very active since the Columbus, Ohio IETF meeting last March; 60 Internet-Drafts, 22 Protocol Actions, and over 30 RFCs.

Between the IETF meetings in Columbus, Ohio and Amsterdam, The Netherlands, there were seventeen new working groups created:

1. Character MIB (CHARMIB)
2. DECnet Phase IV MIB (DECNETIV)
3. RIP Version II (RIPV2)
4. Internet Protocol Security Protocol (IPSEC)
5. Authorization and Access Control (AAC)
6. TP/IX (TPIX)
7. Modem Management (MODEMMGT)
8. Frame Relay Service MIB (FRNETMIB)
9. Mail and Directory Management (MADMAN)
10. Interactive Mail Access Protocol (IMAP)
11. ATM MIB (ATOMMIB)
12. Telnet TN3270 Enhancements (TN3270E)
13. SNA DLC Services MIB (SNADLC)
14. Interfaces MIB (IFMIB)
15. SNA NAU Services MIB (SNANAU)
16. Multiparty Multimedia Session Control (MMUSIC)
17. Inter-Domain Multicast Routing (IDMR)

and seven working groups were concluded:

1. Internet Accounting (ACCT)
2. Office Document Architecture (ODA)
3. X.25 Management Information Base (X25MIB)
4. SNMP Security (SNMPSEC)
5. SNMP Version 2 (SNMPV2)
6. NOC-Tool Catalogue Revisions (NOCTOOL2)
7. Distributed File Systems (DFS)

Additionally, there were 34 RFCs published since the Columbus IETF meeting in March, 1993:

| RFC | Status | Title |
|---------|--------|---|
| RFC1441 | PS | Introduction to version 2 of the Internet-standard Network Management Framework |

| | | |
|---------|----|--|
| RFC1442 | PS | Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1443 | PS | Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1444 | PS | Conformance Statements for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1445 | PS | Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1446 | PS | Security Protocols for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1447 | PS | Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1448 | PS | Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1449 | PS | Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1450 | PS | Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2) |
| RFC1451 | PS | Manager to Manager Management Information Base |
| RFC1452 | PS | Coexistence between version 1 and version 2 of the Internet-standard Network Management Framework |
| RFC1453 | I | A Comment on Packet Video Remote Conferencing and the Transport/Network Layers |
| RFC1454 | I | Comparison of Proposals for Next Version of IP |
| RFC1455 | E | Physical Link Security Type of Service |
| RFC1456 | I | Conventions for Encoding the Vietnamese Language VISCII: Vietnamese Standard Code for Information Interchange VIQR: Vietnamese Quoted-Readable Specification |
| RFC1457 | E | Security Label Framework for the Internet |
| RFC1458 | I | Requirements for Multicast Protocols |
| RFC1459 | E | Internet Relay Chat Protocol |
| RFC1460 | DS | Post Office Protocol - Version 3 |
| RFC1461 | PS | SNMP MIB extension for MultiProtocol Interconnect over X.25 |
| RFC1462 | I | FYI on "What is the Internet?" |
| RFC1463 | I | FYI on Introducing the Internet—A Short Bibliography of Introductory Internetworking Readings for the Network Novice |
| RFC1464 | E | Using the Domain Name System To Store Arbitrary String Attributes |
| RFC1465 | E | Routing coordination for X.400 MHS services within a multi protocol / multi network environment Table Format V3 for static routing |
| RFC1466 | I | Guidelines for Management of IP Address Space |

| | | |
|---------|----|---|
| RFC1468 | I | Japanese Character Encoding for Internet Messages |
| RFC1469 | PS | IP Multicast over Token-Ring Local Area Networks |
| RFC1471 | PS | The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol |
| RFC1472 | PS | The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol |
| RFC1473 | PS | The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol |
| RFC1474 | PS | The Definitions of Managed Objects for the Bridge Network Control Protocol of the Point-to-Point Protocol |
| RFC1475 | I | TP/IX: The Next Internet |
| RFC1476 | I | RAP: Internet Route Access Protocol |

Final Agenda of the Twenty-Seventh IETF

(12-16 July 1993)

MONDAY, 12 July 1993

- | | | | | | | | | | | | | | | | | | | | |
|-----------|--|-----|---|-----|--|-----|---|-----|---|-----|--|-----|---|-----|--|-----|--|-----|--|
| 0800-0900 | IETF Registration and Continental Breakfast | | | | | | | | | | | | | | | | | | |
| 0900-0930 | Introductions | | | | | | | | | | | | | | | | | | |
| 0930-1200 | Technical Presentations <ul style="list-style-type: none"> • “The IETF from a European Perspective” (Erik Huizer/SURFnet) • “Advances on ONC” (Chuck McManis/Sunsoft) • “Setting up a Routing Registry in Europe” (Daniel Karrenberg/RIPE) | | | | | | | | | | | | | | | | | | |
| Breaks | Coffee available throughout morning. | | | | | | | | | | | | | | | | | | |
| 1330-1530 | Afternoon Sessions I <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">APP</td> <td>OSI Directory Services WG (osids) (Steve Kille/ISODE)</td> </tr> <tr> <td style="vertical-align: top;">INT</td> <td>IP over ATM WG (atm) (Mark Laubach/Hewlett-Packard)</td> </tr> <tr> <td style="vertical-align: top;">INT</td> <td>Network Address Translators BOF (nat) (Kjeld Borch Egevang/Cray Communications)</td> </tr> <tr> <td style="vertical-align: top;">MGT</td> <td>Interfaces MIB WG (ifmib) (Ted Brunner/Bellcore)</td> </tr> <tr> <td style="vertical-align: top;">RTG</td> <td>Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)*</td> </tr> <tr> <td style="vertical-align: top;">RTG</td> <td>OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)*</td> </tr> <tr> <td style="vertical-align: top;">OPS</td> <td>Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET)</td> </tr> <tr> <td style="vertical-align: top;">SEC</td> <td>Security Area Advisory Group (saag) (Steve Crocker/TIS)</td> </tr> <tr> <td style="vertical-align: top;">USV</td> <td>Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/UNC)</td> </tr> </table> | APP | OSI Directory Services WG (osids) (Steve Kille/ISODE) | INT | IP over ATM WG (atm) (Mark Laubach/Hewlett-Packard) | INT | Network Address Translators BOF (nat) (Kjeld Borch Egevang/Cray Communications) | MGT | Interfaces MIB WG (ifmib) (Ted Brunner/Bellcore) | RTG | Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)* | RTG | OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)* | OPS | Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET) | SEC | Security Area Advisory Group (saag) (Steve Crocker/TIS) | USV | Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/UNC) |
| APP | OSI Directory Services WG (osids) (Steve Kille/ISODE) | | | | | | | | | | | | | | | | | | |
| INT | IP over ATM WG (atm) (Mark Laubach/Hewlett-Packard) | | | | | | | | | | | | | | | | | | |
| INT | Network Address Translators BOF (nat) (Kjeld Borch Egevang/Cray Communications) | | | | | | | | | | | | | | | | | | |
| MGT | Interfaces MIB WG (ifmib) (Ted Brunner/Bellcore) | | | | | | | | | | | | | | | | | | |
| RTG | Border Gateway Protocol WG (bgp) (Yakov Rekhter/IBM)* | | | | | | | | | | | | | | | | | | |
| RTG | OSI IDRP for IP over IP WG (ipidrp) (Sue Hares/Merit)* | | | | | | | | | | | | | | | | | | |
| OPS | Operational Statistics WG (opstat) (Phill Gross/ANS and Bernhard Stockman/SUNET) | | | | | | | | | | | | | | | | | | |
| SEC | Security Area Advisory Group (saag) (Steve Crocker/TIS) | | | | | | | | | | | | | | | | | | |
| USV | Uniform Resource Identifiers WG (uri) (Alan Emtage/Bunyip and Jim Fullton/UNC) | | | | | | | | | | | | | | | | | | |
| 1530-1600 | Break (Refreshments provided) | | | | | | | | | | | | | | | | | | |

1600-1745

Monday, 12 July 1993 - Afternoon Sessions II

- APP OSI Directory Services WG (**osids**)
(Steve Kille/ISODE)
- INT IP over Large Public Data Networks WG (**iplpdn**)
(George Clapp/Ameritech)
- INT Simple Internet Protocol WG (**sip**)
(Steve Deering/Xerox PARC and Bob Hinden/Sun)
- MGT Interfaces MIB WG (**ifmib**) (Ted Brunner/Bellcore)
- MGT Uninterruptible Power Supply WG (**upsmib**)
(Jeff Case/UTenn)
- SEC Common Authentication Technology WG (**cat**)
(John Linn/GZA)
- USV Uniform Resource Identifiers WG (**uri**) (Alan Emtage/Bunyip
and Jim Fullton/UNC)

* BGP and IPIDPR will be meeting in joint session.

TUESDAY, 13 July 1993

- 0830-0900 Continental Breakfast
- 0900-1200 Morning Sessions
- APP TELNET WG (**telnet**)
(Steve Alexander/Lachman Technology)
- INT P. Internet Protocol WG (**pip**) (Paul Francis/Bellcore)
- MGT Network Management Area: Open Meeting (**nmarea**)
(Marshall Rose/DBC)
- SAP Network Database Working Group (**netdata**) (Daisy
Rose/IBM)
- SEC Internet Protocol Security Protocol WG (**ipsec**)
(Al Hoover/ANS and Paul Lambert/Motorola)
- TSV Multiparty Multimedia Session Control WG (**mmusic**)
(Eve Schooler/ISI and Abel Weinrib/Bellcore)
- USV Integrated Directory Services WG (**ids**) (Tim Howes/UMich
and Chris Weider/Merit)
- USV User Services WG (**uswg**) (Joyce K. Reynolds/ISI)
- Breaks Coffee available throughout morning.
- 1330-1530 Afternoon Sessions I
- APP Interactive Mail Access Protocol WG (**imap**)
(Terry Gray/UWash)
- INT IP over ATM WG (**atm**) (Bob Hinden/Sun)
- INT P. Internet Protocol WG (**pip**) (Paul Francis/Bellcore)
- MGT Modem Management WG (**modemmgmt**) (Mark Lewis/Telebit)
- RTG RIP Version II WG (**ripv2**) (Gary Malkin/Xylogics)
- SAP MHS-DS WG (**mhsds**) (Kevin Jordan/CDC and
Harald Alvestrand/SINTEF DELAB)
- SEC Common Authentication Technology WG (**cat**)
(John Linn/GZA)
- USV WHOIS and Network Information Lookup
Service WG (**wnils**) (Joan Gargano/UCDavis)
- 1530-1600 Break (Refreshments provided)

1600-1800

Tuesday 13 July 1993 - Afternoon Sessions II

- INT IP over Large Public Data Networks WG (**iplpdn**)
(George Clapp/Ameritech) *
- INT Point-to-Point Protocol Extensions WG (**pppext**)
(Fred Baker/ACC) *
- INT TCP/UDP over CLNP-addressed Networks WG (**tuba**)
(Peter Ford/LANL and Mark Knopper/Merit)
- MGT ATM MIB WG (**atommib**) (Kaj Tesink/Bellcore)
- RTG Border Gateway Protocol WG (**bgp**) (Yakov Rekhter/IBM)*
- RTG Inter-Domain Multicast Routing WG (**idmr**)
(Tony Ballardie/UCL)
- RTG OSI IDRP for IP over IP WG (**ipidrp**) (Sue Hares/Merit)*
- SAP MHS-DS WG (**mhsds**) (Kevin Jordan/CDC and
Harald Alvestrand/SINTEF DELAB)
- TSV TCP Multiplexing BOF (**tmux**) (James Barnes/Xylogics)
- USV Integration of Internet Information Resources WG (**iiir**)
(Chris Weider/Merit)
- USV User Documents WG (**userdoc2**) (Ellen Hoffman/UMich
and Lenore Jackson/NASA)

1930-2200

Evening Sessions

- APP Networking Multimedia Applications BOF (**multiapp**)
(Chris Adie/Edinburgh University)
- GEN IAB Open Meeting.
- INT ST2 BOF (**st2**) (Luca Delgrossi/IBM and
Steve DeJarnett/IBM)
- INT TP/IX WG (**tpix**) (Vladimir Sukonnik/Process Soft-
ware)
- MGT ATM MIB WG (**atommib**) (Kaj Tesink/Bellcore)
- MGT Mail and Directory Management WG (**madman**)
(Steve Kille/ISODE)

* IPLPDN and PPPEXT will be meeting in joint session.

WEDNESDAY, 14 July 1993

0830-0900 Continental Breakfast

0900-1200 Morning Sessions

- APP X.400 Operations WG (**x400ops**) (Alf Hansen/Sintef and Tony Genovese/LLNL)
- INT IP over Large Public Data Networks WG (**iplpdn**) (George Clapp/Ameritech)*
- INT Point-to-Point Protocol Extensions WG (**pppext**) (Fred Baker/ACC) *
- INT TP/IX WG (**tpix**) (Vladimir Sukonnik/Process Software)
- MGT SNA NAU Services MIB WG (**sna nau**) (Zbigniew Kielczewski/Eigon Technology and Deirdre Kostick/Bellcore)
- MGT Token Ring Remote Monitoring WG (**trmon**) (Mike Erlinger/Harvey Mudd College)
- RTG Source Demand Routing Protocol WG (**sdr**) (Deborah Estrin/USC and Tony Li/cisco)
- SAP NFS and ONC IETF Standards Effort BOF (**onc**) (Chuck McManis/SunSoft and Dave Crocker/Silicon Graphics)
- TSV Multiparty Multimedia Session Control WG (**mmusic**) (Eve Schooler/ISI and Abel Weinrib/Bellcore)
- USV Uniform Resource Identifiers WG (**uri**) (Alan Emtage/Bunyip and Jim Fullton/UNC)

Breaks Coffee available throughout morning.

1330-1530 Afternoon Sessions I

- APP X.400 Operations WG (**x400ops**) (Alf Hansen/Sintef and Tony Genovese/LLNL)
- INT Extensions to OSI for use in the Internet BOF (**osiextnd**) (Dave Katz/cisco)
- INT IP over ATM WG (**atm**) (Bob Hinden/Sun)
- SAP Domain Name System WG (**dns**) (Rob Austein/Epilogue Technology)

- 1330-1530 Wednesday, 14 July 1993 - Afternoon Sessions I (cont'd.)
- SAP Service Location Protocol WG (**svrloc**)
 (John Veizades/FTP Software and
 Scott Kaplan/FTP Software)
- SEC Privacy-Enhanced Electronic Mail WG (**pem**)
 (Steve Kent/BBN)
- USV Networked Information Retrieval WG (**nir**)
 (Jill Foster/UNewcastle-Upon-Tyne and George Brett/MCNC)
- 1530-1600 Break (Refreshments provided)
- 1600-1800 Afternoon Sessions II
- APP Internet Mercantile Protocols BOF (**imp**)
 (Taso Devetzis/Bellcore)
- APP UCS Character Set BOF (**ucs**)
 (Borka Jerman-Blazic/Jozef Stefan Institute)
- MGT Frame Relay Service MIB WG (**frnetmib**)
 (James Watt/Newbridge Networks)
- RTG Inter-Domain Multicast Routing WG (**idmr**)
 (Tony Ballardie/UCL)
- RTG IP Routing for Wireless/Mobile Hosts WG (**mobileip**)
 (Steve Deering/Xerox PARC)
- SAP MHS-DS Tutorial (**mhsds**) (Kevin Jordan/CDC and
 Harald Alvestrand/SINTEF DELAB)
- SEC Authorization and Access Control WG (**aac**)
 (Cliff Neuman/ISI)
- TSV TCP Large Windows WG (**tcplw**)
 (Dave Borman/Cray Research)
- USV Network Information Services Infrastructure WG (**nisi**)
 (April Marine/SRI and Pat Smith/Merit)
- 1930-2200 Evening Sessions
- INT IPng Decision Process BOF (**ipdecide**)
 (Brian Carpenter/CERN and Tim Dixon/RARE)
- INT Point-to-Point Protocol Extensions WG (**pppext**)
 (Fred Baker/ACC)

1930-2200

Wednesday, 14 July 1993 - Evening Sessions (cont'd.)

OPS Generic Internet Service Specification BOF (giss)
 (Tony Bates/RIPE and Daniel Karrenberg/RIPE)

USV Integration of Internet Information Resources WG (iir)
 (Chris Weider/Merit)

* IPLPDN and PPPEXT will be meeting in joint session.

THURSDAY, 15 July 1993

- 0830-0900 Continental Breakfast
- 0900-0930 Technical Presentations
- “Electronic Cash: Theory Towards Application”
(David Chaum/DigiCash)
- 0930-1200 Morning Sessions
- APP X.400 Operations WG (**x400ops**) (Alf Hansen/Sintef
and
Tony Genovese/LLNL)
- INT Point-to-Point Protocol Extensions WG (**pppext**)
(Fred Baker/ACC)
- MGT FDDI MIB (**fddimib**) (Jeff Case/UTenn)
- OPS BGP Deployment and Application WG (**bgpdepl**)
(Jessica Yu/Merit)
- RTG IP Routing for Wireless/Mobile Hosts WG (**mobileip**)
(Steve Deering/Xerox PARC)
- RTG ISIS for IP Internets WG (**isis**) (Chris Gunner/DEC)
- SAP Minimal OSI Upper-Layers WG (**thinosi**) (Peter Furniss/Consultant)
- SEC Network Access Server Requirements WG (**nasreq**)
(Jim Barnes/Xylogics)
- USV Network Training Materials WG (**trainmat**)
(Ellen Hoffman/Merit and
Jill Foster/UNewcastle-Upon-Tyne)
- Breaks Coffee available throughout the morning.
- 1330-1530 Afternoon Sessions I
- INT Simple Internet Protocol WG (**sip**)
(Steve Deering/Xerox PARC and Bob Hinden/Sun)
- INT TCP/UDP over CLNP-addressed Networks WG (**tuba**)
(Peter Ford/LANL and Mark Knopper/Merit) *
- MGT IFIP Electronic Mail Management BOF (**emailmgt**)
(Maria Dimou/CERN)

| | |
|-----------|--|
| 1330-1530 | Thursday, 15 July 1993 - Afternoon Sessions I (cont'd.) |
| | MGT Frame Relay Service MIB WG (frnetmib) (James Watt/Newbridge Networks) |
| | OPS Network OSI Operations WG (noop) (Sue Hares/Merit and Cathy Wittbrodt/BARRnet)* |
| | OPS Operational Area Directorate (orad) Scott Bradner/Harvard) |
| | SAP Structured Text Interchange Format BOF (stif) (Dave Crocker/Silicon Graphics) |
| | SEC Security Area Advisory Group (saag) (Steve Crocker/TIS) |
| | USV Internet School Networking WG (isn) (Jennifer Sellers/NASA) |
| 1530-1600 | Break (Refreshments provided) |
| 1600-1700 | Technical Presentations |
| | <ul style="list-style-type: none"> • "JVTOS for Workstations" (Bernhard Plattner/Institut fuer Technische) • "Service Location" (John Veizades/FTP Software and Scott Kaplan/FTP Software) |
| 1700-1930 | Open Plenary and IESG |

* TUBA and NOOP will be meeting in joint session.

FRIDAY, 16 July 1993

0830-0900 Continental Breakfast

0900-1200 Technical Presentations

- "An Experiment in Remote Printing" (Marshall Rose/DBC)
- "IPng Presentations"
- "Final Remarks" (Phill Gross/ANS and Erik Huizer/SURFnet)

Key to Abbreviations

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

Chapter 1

IETF Overview

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

The IETF mission includes:

1. Identifying and proposing solutions to pressing operational and technical problems in the Internet;
2. Specifying the development (or usage) of protocols and the near-term architecture to solve such technical problems for the Internet;
3. Facilitating technology transfer from the Internet Research Task Force (IRTF) to the wider Internet community; and
4. Providing a forum for the exchange of relevant information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

Technical activity on any specific topic in the IETF is addressed within working groups. All working groups are organized roughly by function into ten technical areas. Each is led by one or more area director who has primary responsibility for that one area of IETF activity. Together with the Chair of the IETF, these technical directors (plus, the Director for Standards Procedures) compose the Internet Engineering Steering Group (IESG).

The current areas and directors, which compose the IESG are:

| | |
|--------------------------|---------------------------|
| IETF and IESG Chair | Phill Gross/ANS |
| Applications | Erik Huizer/SURFnet |
| | John Klensin/UNU |
| Internet | Stev Knowles/FTP Software |
| | Dave Piscitello/Bellcore |
| IP: Next Generation | Scott Bradner/Harvard |
| | Allison Mankin/NRL |
| Network Management | Marshall Rose/DBC |
| Operational Requirements | Scott Bradner/Harvard |
| Routing | Robert Hinden/Sun |
| Security | Steve Crocker/TIS |
| Service Applications | Dave Crocker/SGI |
| Transport | Allison Mankin/NRL |
| User Services | Joyce K. Reynolds/ISI |
| Standards Management | A. Lyman Chapin/BBN |

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

| | |
|------------------------------------|---------------------|
| IETF Executive Director | Steve Coya |
| IESG Secretary | John Stewart |
| IETF Meeting Coordinator | Megan Davies Walnut |
| IETF Meeting Registrar | Debra Legare |
| IETF Internet-Drafts Administrator | Cynthia Clark |
| IETF Administrative Support | Lois Keiper |

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

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

Mailing Lists

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

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

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

1.1 Future IETF Meeting Sites

Fall 1993

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

Spring 1994

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

Summer 1994

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

1.2 On Line IETF Information

The Internet Engineering Task Force maintains up-to-date, on-line information on all of its activities. This information is available via FTP and e-mail. Procedures for retrieving the information are described below.

The IETF Directory

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

Files prefixed with a 0 contain information about upcoming meetings. Files prefixed with a 1 contain general information about the IETF. Working group charters and minutes are in sub-directories under the working group acronym. Retrieve and view the `1wg-summary.txt` file for a list of working groups and their acronyms.

FILE NAME

| | |
|------------------------------------|--|
| <code>0tao.txt</code> | This file contains "A Guide for New Attendees of the Internet Engineering Task Force", RFC 1391. |
| <code>0mtg-agenda.txt</code> | The current agenda for the upcoming IETF meeting, containing scheduled working group meetings, technical presentations and network status reports. |
| <code>0mtg-at-a-glance.txt</code> | The announcement for the upcoming IETF meeting, containing specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and meeting costs. |
| <code>0mtg-rsvp.txt</code> | A standardized RSVP form to notify the Secretariat of your plans to attend the upcoming IETF meeting. |
| <code>0mtg-sites.txt</code> | Current and future meeting dates and sites for IETF meetings. |
| <code>1id-guidelines.txt</code> | Instructions for authors of Internet-Drafts. |
| <code>1ietf-description.txt</code> | A short description of the IETF, the IESG and how to participate. |
| <code>1wg-summary.txt</code> | A listing of all current working groups, the working group Chairs and their e-mail addresses, working group mailing list addresses, and where applicable, documentation produced. This file also contains the standard acronym for the working groups by which the IETF and Internet-Drafts directories are keyed. |

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

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

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

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

The Internet-Drafts Directory

The Internet-Drafts directory has been installed to make available, for review and comment, draft documents that may eventually be submitted to the IESG and/or the RFC Editor to be considered for publication as RFCs. These documents are indexed in the file `lid-abstracts.txt` in the Internet-Drafts directory. Comments are welcome and should be addressed to the responsible person(s) whose name and e-mail address are listed on the first page of the respective draft.

FILE NAME

`lid-abstracts.txt` This file lists the current Internet-Drafts and their pathnames.

`lid-index.txt` This file contains an abbreviated listing of Internet-Drafts. This contains only the document title, the filename and the posting date.

For more information on writing and installing an Internet-Draft, see the file `lid-guidelines` in the `ietf` directory, “Guidelines to Authors of Internet-Drafts.”

The IESG Directory

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

FILE NAME

| | |
|-----------------------|---|
| lprotocol_actions.txt | This file contains a list of protocols currently under consideration by the IESG. |
| lold_standards.txt | This file contains a list of Proposed and Draft Standards eligible for advancement. |

The minutes are contained in files named with the pattern:

iesg.YY-MM-DD

e.g.,

iesg.92-11-10

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

FTP Access

IETF Information is available by anonymous FTP from several sites.

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

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

Europe Address: nic.nordu.net (192.36.148.17)

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

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

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

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

Individual files can then be retrieved using the GET command:

```
get lwg-summary.txt
get 822ext/822ext-charter.txt
```

E-mail Access

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

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

In the body, put a command of the form:

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

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

1.3 Guidelines to Authors of Internet-Drafts

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

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

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

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

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

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

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a “working draft” or “work in progress.”

To learn the current status of any Internet-Draft, please check the `lid-abstracts.txt` listing contained in the Internet-Drafts Shadow Directories on `ds.internic.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munnari.oz.au`.

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

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

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

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

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

Chapter 2

Area and Working Group Reports

2.1 Applications Area

Director(s):

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

Area Summary reported by Erik Huizer/SURFnet

The Applications Area currently contains the following working groups:

- Interactive Mail Access Protocol (IMAP)
- Internet Message Extensions (822EXT)
- MIME-MHS Interworking (MIMEMHS)
- Network News Transport Protocol (NNTP)
- OSI Directory Services (OSIDS)
- TELNET (TELNET)
- Telnet TN3270 Enhancements (TN3270E)
- X.400 Operations (X400OPS)

The 822EXT and MIMEMHS Working Groups have finished their work and did not meet in Amsterdam. The NNTP and TN3270E Working Groups also did not meet in Amsterdam.

Three BOFs under the Applications Area were held in Amsterdam:

- Internet Mercantile Protocols (IMP)
- Networking Multimedia Applications (MULTIAPP)
- UCS Character Set BOF (UCS)

Internet Mercantile Protocols BOF (IMP)

The IMP BOF was convened to assess community interest in Internet-based commerce and to explore some concrete ideas on how that might be realized using existing technology.

The session was comprised of two presentations followed by general discussion. Taso Devetzi presented some principles on which protocols for Internet commerce might be based, followed by a detailed example of how such principles might be realized using existing Internet technology (e.g., PEM, MIME). Mitra presented a brief overview of a system being developed by Pandora Systems to support its commercial offerings over the Internet.

From the discussion emerged three issues relating to the requirements for Internet commerce:

1. A protocol that admits bilateral transactions may afford less incentive to aspiring third-party mediators than does a framework that requires all transactions to be mediated.
2. An acceptable protocol must support "real-time," interactive use.
3. An acceptable protocol must be compatible with existing Internet applications (e.g., Gopher).

Brief discussion led to general agreement on the second and third points. Neither point was regarded as necessarily inconsistent with the proposed leveraging of MIME and PEM technology. Time did not permit full discussion of the first point above. Erik Huizer, an Applications Area Director, concluded the meeting by saying that interest in this topic was clearly sufficient to merit further work but that further definition of the task would be valuable before chartering a working group. To this end, specific topics for e-mail discussion were identified.

Networking Multimedia Applications BOF (MULTIAPP)

Chris Adie introduced himself as the leader of the RARE Multimedia Information Services Task Force and described the scope of the BOF, namely covering networked access to multimedia resources from both the users' and providers' points of view. Chris then gave a short presentation of the probable application categories, the requirements, the existing systems and standards and the aims he saw for the BOF. Aims of the BOF were: to identify issues involved in providing access to multimedia data; to identify ways to make progress in addressing these issues; and to agree how to interact with existing groups working in the area. A list of issues was then solicited from the floor. The discussion arising from this was wide ranging, and it took some time before suitable areas for IETF involvement were agreed upon. The discussion eventually settled on three items: attempting some pilot work in mounting (at least) one multimedia application across the network; working up a draft charter for a proposed IETF working group; looking at ways of linking existing or emerging standards (RTP, HTML+ were cited as candidates) to define a multimedia access protocol (loosely described as 'Multimedia X'). All of these, particularly the last, required some study of prior art (SunSoft, DEC, Apple, Microsoft and Bellcore were cited) and liaison with other groups.

UCS Character Set BOF (UCS)

The UCS BOF discussed the common basis for making possible in Internet protocols the interoperable use of characters beyond the US-ASCII repertoire. The following issues were discussed:

- Alternatives to 10646 BMP for Asian ideographic character sets.
- More internationalized coding alternatives.
- Byte-order for languages that do not run left-to-right, and how to represent them.
- Possible feedback and liaison processes from IETF to ISO/IEC JTC1/SC2 and requests for inclusion of special sets of characters as part of 10646.

The BOF identified several items to be worked out within the IETF working procedures. Among many of them the most urgent seems to be the following:

- A document defining the necessary meta-protocol or process which will deal with the items required for the internationalization of the networked services. The document will provide guidance to other protocols dealing with these items over the Internet.
- A document which will specify the byte-ordering of data streams coded with UCS to be used in the Internet. Recommendation regarding the particular encoding (e.g., UTF-2 or some variation) to be used in Internet protocols.
- A document identifying the languages and the characters required for coding text written in natural language (a sort of “guidelines for services” such as NIR based on usage of plain text written in languages different from English).
- A document defining a tool for coded character sets conversion to be provided within some services such as e-mail (i.e. conversion of character set codes that are outside of the supported character repertoire of the receiver).

Interactive Mail Access Protocol Working Group (IMAP)

Twenty people participated. For several it was their first exposure to IMAP, so a few minutes were spent summarizing what IMAP is, how it compares/relates to other alternatives, and what the working group is chartered to do. The working group charter and notes from the Columbus BOF were reviewed and questions were answered. The status of the protocol specification and known IMAP implementations was reviewed. (An Internet-Draft is being composed that integrates and updates RFC 1176 and the imap2bis extensions.) Existing practice on the use of IMAP for news, archive, and document access—in addition to mail—was covered. Discussion on possible IMAP extensions followed. Finally, the next working group meeting, to be held in Seattle on August 30 and 31, was announced.

OSI Directory Services (OSIDS)

- Document progress was discussed. The User Friendly Naming document is still waiting for RFC Editor action. “DSA Metrics” is finished and will be submitted to the RFC Editor as an Informational document.

- Editorial updates will be made to RFC 1278, "A String Encoding of Presentation Address."
- InterNIC progress on formalising X.500 deployment on the Internet was presented and discussed.
- There is new work on representing IP information in the DIT. This work will hopefully lead to a set of RFCs later this year, including an update to RFC 1274.
- A document on representation of bibliographic information in the DIT is progressing. It is expected to be an Experimental RFC later this year.
- RFC 1384, "Naming Guidelines for Directory Pilots," will be updated, based on a draft document by the RARE NAP Working Group.
- With the InterNIC taking the lead, a new schema management group was formed that will look after maintaining the directory schema for the Internet.

TELNET (TELNET)

The primary topic of discussion was the environment option. Originally, a proposal was made to re-issue RFC 1408 with the VAR and VALUE definitions corrected to match the BSD implementation. Additionally, a document was produced that explained heuristics that could be used to handle implementations that did not agree.

The attendees briefly discussed the proposed revisions to the charter. The general consensus is that the group does not want to remain a clearing-house for TELNET documents. The charter will be revised to only cover the authentication work that is going on.

The group reviewed a proposal for a "Telnet Transfer of Control" option. The group feels that there are some security issues with the document. Various people will forward their comments directly to the author. At this time the group does not wish to formally work on this specification.

Finally, the group briefly discussed Dave Borman's new option that merged authentication and encryption. The initial feedback is favorable, but Dave needs to complete an initial implementation using Kerberos IV.

X.400 Operations (X400OPS)

- The session included liaison reports from: MHSDS (Longbud Project status), the RARE MHS Working Group, EMA/EEMA, and Cosine MHS (coordination service update).

- It was recommended that a new working group be formed to handle the topic of file transfer over e-mail.
- The PRMD Requirements document is being held up because some technical work is needed on the postmaster document which is a required component of this document. It has been determined that with the publishing of the PRMD Requirements RFC that the main objectives of this working group will have been met. With this in mind, it is believed that this working group can close down. The group would like this accomplished before the Houston IETF.
- It was proposed that a working group be formed to work on deployment of X.400 within the Internet.
- It was proposed that a new working group be formed to work on the C=US and A=IMX issues.
- A charter was presented for a working group that would be formed to work on ADMD interconnections to the Internet.
- Status reports were given on GO-MHS support and the RARE X400 88 pilot.

CURRENT MEETING REPORT

Reported by Taso Devetzis/Bellcore

Minutes of the Internet Mercantile Protocols BOF (IMP)

Introduction

The IMP BOF session was convened to assess community interest in Internet-based commerce and to explore some concrete ideas on how that might be realized using existing technology. The session comprised two presentations together with some general discussion. Taso Devetzis presented some principles on which protocols for Internet commerce might be based followed by an illustrative example of how such principles might be realized using existing Internet technology (e.g., PEM, MIME). Mitra presented some informal ideas on what a system for Internet commerce might look like.

Kick-Off Presentation

The session began with the circulation of the attendance roster and other administrative. The following agenda was accepted without much discussion:

- Introductory talk
 - The vision
 - Bits, bytes, and examples
- Questions and discussion
- Future directions

Taso began with an introductory presentation. He identified the goal as enabling commerce over the Internet—focusing on “commercial consummation” rather than on “commercial foreplay.” He also attempted to focus discussion by identifying goals that, however worthy, are not the most immediate problems for enabling Internet commerce. Among these non-goals are:

- The “electronic cash” problem
- Automation of today’s billing and collection processes
- The directory services problem
- The resource identification and discovery problem
- Replication of the entire EDI suite
- Reforming society and altering the human condition

Taso identified the motivations for pursuing this work and cited a number of unilateral efforts as evidence of growing interest, and suggested that this technology should be driven

by the needs of Internet users rather than by any of a number of other vested interests. He identified the other benefits to the community that this effort could provide:

- Convenience to Internet users
- Easier vendor access to broader markets
- Internet commerce will help fund Internet infrastructure
- Promotes Internet growth
- Provides incentive for fully automated procurement and its benefits
- Reduces paperwork and bureaucracy

The discussion then turned to the principles on which an overall approach might be based. The emphasis was on policy-free mechanisms and the use of already-standardized Internet technology and infrastructure.

- Allow for bilateral transactions (“Look ma, no trusted third party!”)
- Universal deployment not required
- Simple mechanism
- Leverage existing Internet technology
 - Support for multimedia via MIME
 - Security enhancements via PEM
 - No new or exotic technology is necessary!
- Provide a core mechanism to enable commerce
- Decouples transport accounting from higher-layer services

Taso emphasized the importance of support for bilateral transactions. Bilateral transactions are the simplest case. They represent a mechanism by which commerce is conducted over the Internet today, and new standards in this area should seek to enhance these existing capabilities rather than to restrict them in the service of a particular commercial agenda. To preclude or deprecate support for bilateral transactions is technologically to compel people to accept mediation services for all of their business—even where such mediation may be neither economically warranted nor socially acceptable.

Support for bilateral transactions is not only important as a social principle, but it affords practical advantages as well. Because it represents the mode in which Internet commerce can be conducted today, it serves as a simple reference paradigm by which seemingly complicated legal or social concerns may be placed in proper perspective. Because bilateral transactions represent the mode in which Internet commerce can be conducted today, they may play a significant role in “bootstrapping” deployment—that is, enabling commerce even before total acceptance and deployment of the relevant infrastructures.

An approach that mechanically decouples commercial transactions from transport service accounting not only simplifies the latter but may also admit cost recovery for transport services in ways that enjoy increased social appeal. The dynamics of the user’s interaction with the postal service is completely decoupled from the interaction between the transacting parties. In mail-order transactions, costs for the postal service are recovered in a variety of ways that can be matched to the parties’ accounting overhead and market strategies.

Example Mechanisms

To illustrate these ideas, a possible solution approach that is consistent with the high-level goals was sketched out.

The illustrative mechanisms exploited MIME and PEM technology to provide for secure documentation of agreements among Internet users to exchange goods and services. This mechanism supports both a bilateral transaction model, and a transaction model in which third-party mediators are desirable. Detailed examples of how the mechanism would work in both of these cases were presented.

Basic protocol dynamics were sketched. A message containing an “offer” to make some exchange is sent from one party to the other. Should the latter party agree, that party responds with a message that “accepts” the tendered offer. This acceptance message documents the agreement in a potentially non-repudiable way.

The goods and services being represented in the protocol can be either negotiable or non-negotiable. In this context, negotiable denotes an object of abstract value rather than a concrete object or service. For example, when you hold a negotiable interest in a company, you may not lay claim to a particular desk or paper clip, but you have an abstract claim upon the assets of the company as a whole. Similarly, a dollar is an abstract claim upon the assets of the US Treasury. A non-negotiable good is a concrete object or service, like a bushel of apples or a haircut.

As an example of the simplest case, two mutually-trusting users can consummate the exchange of a tee-shirt in return for a negotiable value of ten guilders. In this case, one party sends an offer message to the other stating a willingness to exchange a specified tee-shirt (described using MIME and/or EDI conventions) for a negotiable obligation in the amount of ten guilders on the part of the other party (essentially, a personal “IOU” from the latter party). The offer message contains an expiration date after which the offer is no longer valid. If the second party accepts the offer, then that party responds with an acceptance message, and the transaction is concluded.

A more complicated example (in which the parties do not trust each other) was also presented. In this case, the transaction is mediated by one or more third parties who are trusted by both principals. This example illustrates a number of distinct functional roles that can be realized by various commercial enterprises:

- Consumers – exchange negotiables for goods or services
- Merchants – vice-versa
- Co-operatives – provide anonymity
- Banks – certify negotiables (like certifying a check)
- Notaries – certify dates (to validate contract acceptances)

In this more complicated case, one principal may not be willing to accept as payment what is essentially a personal “IOU” from the other. Thus, as part of the offer message, the

former specifies that the negotiable instrument must be certified by some trusted financial organization (e.g., Citibank). The policy by which Citibank might certify the user's payment could be related to his current bank balance, some pre-arranged credit line, or simply the cut of his/her jib. It is not a matter for protocol standardization. Because this "check certification" function (and also the notarization function) are themselves modelled as transactions, not only is there an established way for certifiers and notaries to get paid for their service, but also the complexity of the overall protocol is reduced. (A minor "bootstrapping" issue does arise: presumably, a certifier may require a customer to have at least some "hard" funds on account, if only to be assured of payment for the certification service.)

Once a transaction is completed, a user may ask his bank (certifier) to credit his account with any new income he derived from the transaction. The user may send the transaction acceptance message to his bank, and the bank will inspect the transaction to determine what additional credit (if any) it will confer upon the user as a result of the transaction. Again, the policy by which credit is assigned is specific to the institution and not a matter of protocol standardization. Because transactions are numbered, a bank can employ fairly simple strategies to counter efforts to "cash-in" a single transaction multiple times (see Dukach and Sollins, among others). With appropriate protocol design, this tracking of transactions need only occur locally between a user and his bank—thereby providing a solution that is not only relatively low in cost but eminently scalable to large numbers of participants.

One functional role not illustrated in the examples is the "Cooperative." This function is one of obscuring the identity of a party to a transaction by acting as a proxy for some large number of parties. This straightforward strategy can (when desirable) afford an acceptable level of privacy to any transaction at lower cost and complexity than electronic cash systems.

The presentation also included detailed examples of message formats and semantics not included in these minutes.

Observations and Discussion

One participant raised the question of how multi-party transactions might be modelled by bilateral protocol exchanges. There was a brief discussion of this question in which various examples of multi-party transactions were posed and analyzed. One view that was expressed was that, in real life, sometimes what seem to be multi-party transactions (e.g., buying a house) are really collections of bilateral transactions that just happen to be concluded at the same meeting (the closing). Another view that was expressed is that it should always be possible to decompose any prima facie multi-party transaction into multiple bilateral agreements, each of which is explicitly conditioned on the others.

Karen Sollins commented that mail queueing mechanisms might impose unacceptable performance constraints on interactive browse-and-buy applications. Devetzis explained that, although e-mail message formats were being used, this approach implied no necessary dependency on time-shifted e-mail delivery mechanisms: the same message formats could be used in both interactive and non-interactive modes.

When the certification procedure was discussed, one of those present observed that a denial of service attack was possible unless the certification failure message is authenticated. This form of attack was not deemed to be very troubling, but it is also not much trouble to counter.

Rob Shirey commented that the presentation at times used the term "privacy" where the term "confidentiality" might be more appropriate. Rob also commented that a list of what services were being provided (and which were not) would also be useful. Such a list would need to be matched against the perceived requirements.

Second Presentation

Mitra gave the second presentation. He described some informal ideas on what a system for Internet commerce might look like. He contrasted the strategy he described with the strategy currently in use by a prototype server. He invited session participants to contact this server on host "path.net" at port 8001. The described strategy had five components:

1. Information Provider - the party who actually produces some information for distribution.
2. Information Retailer - the party who makes information available for sale to Internet users. The Internet system operated by a retailer is sometimes called a "gateway."
3. Host Operator - the party that operates a host system that is used by Internet users for Internet commerce.
4. User - the party who buys stuff from an Information Retailer using a system provided by a Host Operator.
5. Authentication Server - the party who authorizes charges made by an Information Retailer against the account of a Host Operator.

Mitra explained that, in his model, Host Operators and Retailers are authenticated by IP address. Via traditional, out-of-band channels, the Authentication Server bills the Host Operator who in turn bills the attached User for purchased goods.

Mitra identified three trust relationships that are present in his system: Host to Authentication Server, Gateway to Authentication Server, and User to Host.

Phill Gross asked about the way in which such a system could scale to a large number of users. Mitra suggested that a hierarchy of such servers could address the scaling problem, and he cited the use of a single server for Visa credit card authorizations.

Some Issues

Mitra also identified three points for discussion by the group that he felt were especially important:

1. **Bilateralism:** because a great many transactions will occur between parties that do not trust each other, a protocol that supports only bilateral transactions between trusting parties is not adequate.
2. An acceptable protocol must support "real-time," interactive use.
3. An acceptable protocol must be compatible with existing Internet applications (e.g., Gopher).

Brief discussion led to general agreement on the second and third points. Neither point was regarded as inconsistent with the proposed leveraging of MIME and PEM technology. Although time did not permit full discussion of the first point above, it is not clear that it represented a point of actual disagreement as much as a particular way of expressing generally shared beliefs.

Conclusion

Erik Huizer, IETF Area Director for Applications, concluded the meeting by saying that interest in this topic was clearly sufficient to merit further work but that further definition of the task would be valuable before chartering a working group. To this end, specific topics for e-mail discussion were identified. If these topics lead to clearly identifiable work items, a follow-on BOF session, for discussing these work items in view of the possible creation of a WG, will be considered by the area director.

Action Items

1. David Ginsburg of Alcatel SEL volunteered to compile and post via the mailing list a survey of existing experiences in conducting commerce over the Internet.
2. Taso Devetzis took the action of adding the names of the BOF attendees to the mailing list (imp-interest@thumper.bellcore.com).

3. Devetzis took the action item of continuing discussion over e-mail in order to identify work items for the group in addition to those areas of study that would not be appropriate for the IETF.
4. Devetzis took the action item of organizing a second BOF session at the next IETF meeting in order to crystalize intervening e-mail discussion into agreed work items and a framework for continued work.
5. All present took the action item of contributing descriptions of current mechanisms for Internet commerce as Internet Drafts.

Attendees

| | |
|----------------------|-----------------------------------|
| George Abe | abe@infonet.com |
| Toshiya Asaba | asaba@iiij.ad.jp |
| Josee Auber | Josee_Auber@hpgnd.grenoble.hp.com |
| J. Nevil Brownlee | nevil@ccui.aukuni.ac.nz |
| Ross Callon | rcallon@wellfleet.com |
| David Crocker | dcrocker@mordor.stanford.edu |
| James Davin | davin@thumper.bellcore.com |
| Taso Devetzis | devetzis@bellcore.com |
| Hans Eriksson | hans@sics.se |
| Peter Furniss | p.furniss@ulcc.ac.uk |
| David Ginsburg | ginsb@us-es.sel.de |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Phillip Gross | pgross@ans.net |
| Marc Horowitz | marc@mit.edu |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Frank Kastenholtz | kasten@ftp.com |
| Sean Kennedy | liam@nic.near.net |
| Michael Khalandovsky | mlk@ftp.com |
| Peter Kirstein | P.Kirstein@cs.ucl.ac.uk |
| John Klensin | Klensin@infofoods.unu.edu |
| Andrew Knutsen | andrewk@sco.com |
| John Larson | jl Larson@parc.xerox.com |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| Cynthia Mills | cmills@bbn.com |
| Mitra | mitra@pandora.sf.ca.us |
| Andrew Partan | asp@uunet.uu.net |
| Michael Patton | map@bbn.com |
| Shawn Routhier | sar@epilogue.com |
| Miguel Sanz | miguel.sanz@rediris.es |
| Eve Schooler | schooler@isi.edu |
| Robert Shirey | shirey@mitre.org |

| | |
|------------------|----------------------------|
| A. Velu Sinha | avsinha@attmail.com |
| Timon Sloane | timon@timon.com |
| Karen Sollins | sollins@lcs.mit.edu |
| Kamlesh Tewani | ktt@arch2.att.com |
| Susan Thomson | set@bellcore.com |
| Claudio Topolcic | topolcic@cnri.reston.va.us |
| Mario Vecchi | mpv@thumper.bellcore.com |
| John Veizades | veizades@wco.ftp.com |
| Sandro Wallach | sandro@elf.com |
| Abel Weinrib | abel@bellcore.com |

CURRENT MEETING REPORT

Reported by Sam Wilson/The University of Edinburgh

Networking Multimedia Applications BOF (MULTIAPP)

Chris Adie introduced himself as the leader of the RARE Multimedia Information Services Task Force and described the scope of the MULTIAPP BOF, namely covering networked access to multimedia (MM) resources from both the user's and provider's points of view. In particular it was not intended to cover MM conferencing, which was being addressed elsewhere. Three developments suggested the need for work in this area:

- Users were using MM authoring tools (e.g. Guide, Toolbook) to produce MM applications;
- Existing client/server tools such as Gopher, WAIS and WWW are being extended to cope with MM; and
- Various MM-related network tools and projects have emerged: MICE (in the conferencing field), RTP (Real Time Protocol), ST2 and multicast (though multicast was currently being used largely in the conferencing area: person-to-person rather than machine-to-person).

Chris then gave a short presentation covering the probable application categories, the requirements, the existing systems and standards and the aims he saw for the BOF. The slides from his presentation follow these minutes. A couple of points of interest to go with the slides:

1. Chris is the editor of RARE Technical Report 5 "A Survey of Distributed Multimedia: Research, Standards and Products" available on paper and in various forms for anonymous FTP on ftp.ed.ac.uk in the directory pub/mmsurvey; and
2. Many projects seem to use SGML in one form or another.

A list of relevant issues that might be suitable areas of work was then solicited from the floor. The discussion arising from this was wide ranging and it took some time before suitable areas for IETF involvement was agreed upon. The list eventually comprised:

- Responsiveness
- Reliability
- Isochronism
- Reverse Control Flow
- Hypermedia Linking
- Presentation

Editor's Note: A list of discussion topics and points made is available via FTP or mail server from the remote directories as /ietf/93jul/multiapp-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

At this point Tim Berners-Lee was invited to give a short description of HTML and HTML+, the description languages underlying WWW, and compared it particularly to HyTime. Tim was less than encouraging about HTML+ as a contender for a generalised MM presentation language or interchange format, but Chris Adie was more hopeful.

At this point the group seemed to be inclining towards the view that there was nothing obvious that the IETF could be doing except perhaps to encourage a pilot project to mount a MM application across a network. Discussion re-focussed on the description of interchange formats and 'Multimedia X.' Carsten Bormann of Technische Universitaet Berlin revitalised the meeting by suggesting the following 'taxonomy' for the problem:

- Remote Access to MM Applications

Needs MM analogue of X Window System (note that RTP provides some of the necessary functionality but higher level integration is needed along with some kind of session control).

- Global Hyperlinking

- URI/URL
- searching
- "location address" (this is a HyTime term)

- MM Document/Application Interchange Format

This is already being tackled by, for example, HyTime, MHEG, and HyperODA.

The major area of work here is in the "MM X" area, though there is some work needed in interchange formats. It was decided that some members of the group (which members seemed to be decided informally later) should:

1. Take one or more existing applications;
2. Try to mount them for access across a network;
3. Decide whether it is currently possible to achieve this; and
4. Come back to the IETF within a year or so to see if there was any standardisation effort that should be followed up or whether the existing tools and techniques were sufficient.

Looking further at the X analogy brought Carsten to the overhead projector again to describe the possible structure of "MM X" (or perhaps just "M" a Multimedia X analogue.

| “monomedium X” | “MultiMedia X” |
|---------------------------------------|--|
| Events Drawing requests Windows | Video and audio clips |
| Window manager(s) | QoS manager, Time manager, Schedule manager, etc. |
| TCP | RTP |

Between the two layers in “MM X” would be specifications for synchronisation and playout. Some people (SunSoft, DEC, others that people vaguely recalled but couldn’t name) are already working in this area.

Discussion followed on the details of what work should be done, how the tasks should be limited (need achievable goals), and how to liaise with other groups. The meeting eventually settled on the following items:

- Attempting some pilot work in mounting (at least) one MM application across the network;
- Working up a draft charter for a proposed IETF working group;
- Looking at ways of linking existing or emerging standards (RTP, HTML+, etc.) to define a MM access protocol somewhat analogous to X; and
- Study of prior art and liaison with other groups (SunSoft, DEC, Apple, Microsoft, Bellcore and various other groups and consortia).

Minute-Taker’s Disclaimer: In the tradition of BOFs, the discussion in this one was varied and wide ranging. Trying to make coherent minutes that reflect the feel of the meeting is a tricky task and has no doubt been influenced by what the minute-taker happened to be concentrating on (or not) at any particular moment. I therefore take full responsibility for any errors of fact, substance or emphasis, whilst denying any liability in that regard!

Attendees

| | |
|--------------------|--------------------------------------|
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Axel Belinfante | Axel.Belinfante@cs.utwente.nl |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Steve Buchko | stevebu@newbridge.com |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Kevin Gamiel | kgamiel@cnidr.org |
| Joan Gargano | jcgargano@ucdavis.edu |
| Terry Gray | gray@cac.washington.edu |
| Maria Heijne | maria.heijne@surfnet.nl |
| Rune Hjelsvold | Rune.Hjelsvold@idt.unit.no |
| Xinli Hou | xinli@cs.utwente.nl |
| Sascha Ignjatovic | sascha@veda.co.at |
| Xander Jansen | xander.jansen@surfnet.nl |
| Ola Johansson | ojn@se.ubn.net |
| John Johnston | john@berlioz.nsc.com |
| Thomas Kaepfner | kaepfner%heidelbg.vnet@ibmpa.ibm.com |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| Steen Linden | steen.linden@uni-c.dk |
| John Martin | John.Martin@newcastle.ac.uk |
| Paul Milazzo | milazzo@bbn.com |
| Ronny Nilsen | Ronny.Nilsen@usit.uio.no |
| Lars-Gunnar Olsson | Lars-Gunnar.Olsson@data.slu.se |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Dan Romascanu | dan@lannet.com |
| A. Velu Sinha | avsinha@attmail.com |
| Simon Spero | simon_spero@unc.edu |
| Guido van Rossum | guido@cwi.nl |
| Ton Verschuren | Ton.Verschuren@SURFnet.nl |
| Chris Weider | clw@merit.edu |
| Sam Wilson | sam.wilson@ed.ac.uk |

Networking Multimedia Applications

Chris Adie
Edinburgh University Computing Service (GB)
C.J.Adie@ed.ac.uk

Application categories

Database
Academic publishing
Computer-aided learning
General information services

Networking Multimedia Applications

Slide 1

Requirements

Cross-platform support
Mac, PC/MS Windows, Unix/X

Range of media types
Text, image, sound, video

Hyperlinks

Sophisticated presentation
of Multimedia Beethoven

Searching capabilities

Networking Multimedia Applications

Slide 2

Responsiveness

Management facilities
Authentication
Charging

Good authoring tools

Networking Multimedia Applications

Slide 3

Existing systems

Gopher
WAIS
World-Wide Web
Hyper-G
Microcosm
AthenaMuse 2

Networking Multimedia Applications

Slide 4

Standards

MIME
HyTime
MHEG
SGML
AVIs, SMSL
IMA, MMCF

Networking Multimedia Applications

Slide 5

Aims of BOF

Agree on where we're headed

Statement of "goal"

Agree what the important issues are

Prioritised list ?

Decide how to make progress

Working Group ?

Volunteers for further work ?

How to interact with existing groups ?

Networking Multimedia Applications

Slide 6

Some Issues

Conversion tools?

Feasibility?

Which authoring packages to go for?

Use of intermediate format?

Use of existing tools and/or standards?

WAIS, Gopher, WWW, HyTime, MHEG etc

Isochronous data stream servers?

Relevance of AVT working group and RTP

Security implications?

Networking Multimedia Applications

Slide 7

CURRENT MEETING REPORT

Reported by David Crocker/Silicon Graphics

Minutes of the Structured Text Interchange Format BOF (STIF)

STIF is a proposal for a syntax to permit text-based specification of labeled and hierachically-structured data. It tries to serve the dual goal of information interchange and easy human specification. Personal Contact Information (PCI) is a related specification which uses STIF for encoding business-card (Rolodex) descriptions. The Amsterdam BOF was held to discuss IETF interest in pursuing such work. Two Internet-Drafts have been posted. The BOF consisted mostly of discussion about the nature of STIF and PCI, with some useful technical comments offered by the audience, as well as the beginnings of a debate about the alternatives.

There is clear interest in this realm of work. The Application Area Directors have expressed interest in pursuing only one out of several alternatives, and discussions about the alternatives have begun.

Attendees

| | |
|-------------------|---------------------------------|
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Steve Alexander | stevea@lachman.com |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| David Crocker | dcrocker@mordor.stanford.edu |
| Judith Grass | grass@cnri.reston.va.us |
| Gerd Holzhauser | Holzhauser1@applelink.apple.com |
| Per Jacobsen | p.h.jacobsen@usit.uio.no |
| Brewster Kahle | Brewster@wais.com |
| Neil Katin | katin@eng.sun.com |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Eliot Lear | lear@sgi.com |
| Steen Linden | steen.linden@uni-c.dk |
| John Lindsay | lindsay@kingston.ac.uk |
| Kent Malave | kent@bach.austin.ibm.com |
| Brian May | Brian.May@mel.dit.csiro.au |
| Michael Patton | map@bbn.com |
| Luc Rooijackers | lwj@cs.kun.nl |
| David Sitman | a79vm.tau.ac.il |
| Simon Spero | simon_spero@unc.edu |
| Gregory Vaudreuil | gvaudre@cnri.reston.va.us |

STIF: Structured Text Interchange Format & PCI: Personal Contact Information

D. Crocker
Silicon Graphics, Inc.
dcrocker@sgi.com

Agenda

- THE NEED
- THE ALTERNATIVES
- STIF PROPOSAL
 - STIF goals
 - STIF details
- THE FUTURE

STIF & PCI

2

The need for structured information

- ATTRIBUTE/VALUE CONSTRUCT VERY COMMON
- CONNECTIVITY (E.G., MIME EMAIL) MAKING EASY TO EXCHANGE VARIOUS STRUCTURED INFO
- NEED STANDARD BASE FOR SYNTAX
- NEED STANDARD BASE FOR ASSORTED, SPECIFIC USES

STIF & PCI

3

Usage requirements

- TEXT-BASED
 - Human readable
 - Easily transported
 - Embeds within other text
- SIMPLE
 - Human understandable
 - Human generate-able
 - Minimum representation functions possible

STIF & PCI

4

The alternatives

- ASN.1
- SGML
- RFC822 HEADERS
- ???

STIF & PCI

5

STIF goals

- EASILY TYPED
- VERY EASILY READ
- VERY BASIC FUNCTIONALITY
- EMBED WITHIN "ANY" TEXT ENVIRONMENT
- FOCUS CREATIVITY CAREFULLY
 - Re-use pieces of solution

STIF & PCI

6

STIF details

- ENCODING SYNTAX
 - General-purpose, basic
 - Syntax for encoding structured information
- ATTRIBUTE/VALUE
- CHARACTER SET

STIF&PCI

Attribute/value

- ATTRIBUTE/VALUE PAIRS
- MULTIPLE VALUES
- AGGREGATION/NESTING OF A-V'S

STIF&PCI

Data encoding

- BASE US-ASCII CHARACTER SET
- CONCEPT OF "ALTERNATE" CHARACTER SET
 - Alternate specified outside STIF
- LEXICAL CONSTRUCTS FROM RFC822/MIME
- TRANSFER ENCODING SPECIFIED OUTSIDE STIF

STIF&PCI

Examples

```
phone: +1 408 246 8253
phone: +1 408 246 1234 / +1 408 249 6205
Contact <work < phone: +1 415 246 1234>
      home < phone: +1 408 246 8253;
           fax: +1 408 249 6205> >
```

STIF&PCI

PCI example - informal

```
From: "Ole J. Jacobsen" <ole@Call.Stanford.EDU>
Or: +1 415 550-9427 (Home) or +1 415 990-9427 (Cellular)
Direct: +1 415 962-2515 (Office) +1 415 998-4427 (Pager)
Fax: +1 415 949-1779 (Interop) +1 415 826-2008 (Home)
X-Comment: Ignore error messages for "ole@radiomail.net"
Ole J Jacobsen, Editor & Publisher Connections--The
Interoperability Report
Interop Company, 480 San Antonio Road, Suite 100, Mountain
View, CA 94040
Phone: (415) 962-2515 FAX: (415) 949-1779
Email: ole@call.stanford.edu
Note: Ignore error messages for "ole@radiomail.net"
```

STIF&PCI

PCI example - STIF

```
Ole J Jacobsen:
name: Ole J.V. Jacobsen
email: ole@call.stanford.edu;
work <title: Editor & Publisher;
      org: Interop Company;
      dept: Connections--The Interoperability Report;
      street: 480 San Antonio Rd \, Suite 100;
      geo: Mountain View / CA / US;code: 94040;
      phone: +1 415 962 2515;fax: +1 415 949 1779>
home <phone: +1 415 550 9427;fax: +1 415 826 2008>
mobile <phone: +1 415 990 9427;
       pager <phone: +1 415 998 4427> >
```

STIF&PCI

Citation example - informal

[BORG92] Borenstein, N. & Freed, M., "MIME (Multipurpose Internet Mail Extensions): Mechanisms for specifying and describing the format of Internet Message Bodies. March, 1992, Network Information Center, RFC 1341.

[CROC93] Crocker, D., "Evolving the System", in Internet System Handbook, Lynch & Rose (eds.); Reading, Mass., Addison-

STIF&PCI

U

Citation example - STIF

Borenstein-Freed-MIME-92:
author: N. Borenstein, M. Freed;
title: MIME \((Multipurpose Internet Mail Extensions)\):
Mechanisms for specifying and describing the
format of Internet Message Bodies;
date: 1992 / March / ; id: RFC 1341;
org: Network Information Center

Crocker-Evolving-93:
author: D. Crocker; title: Evolving the System;
in: Internet System Handbook; editor: D. Lynch, M.
Rose;
geo: Reading / Mass / ;
org: Addison-Wesley Publishing Co.; date: 1993//

STIF&PCI

K

Use within Mime

--Boundary-1
Content-Type: MULTIPART/MIXED; boundary=Boundary-2

--Boundary-2
Content-Type: TEXT/x-xxx; charset=US-ASCII
(initial part of content, with no special character set
requirements)

--Boundary-2
Content-Type: TEXT/x-xxx; charset=ISO-8859-1
Content-Encoding: Quoted-Printable
(remaining part of content, with character set
supporting some european languages)

STIF&PCI

U

The future

- STIF AS GENERIC ENCODING CONVENTION
- PCI FOR "BUSINESS CARD" AND ROLODEX™
- PURSUE STANDARDS EFFORT?
- PURSUE OTHER SEMANTIC DATA?

STIF&PCI

K

CURRENT MEETING REPORT

Reported by Borka Jerman-Blazic/Jozef Stefan Institute

Minutes of the UCS Character Set BOF (UCS)

Introduction

A brief introductory tutorial was given by Borka Jerman-Blazic. She described some of the problems which appear on the network due to the lack of support for the national character sets used for inputting, outputting, processing and displaying the text written in languages used all over the world. She stressed the need for proper maintenance of the character integrity over the network. The requirement for processing and interchanging different character sets correctly is especially relevant for some Internet services dealing with names of persons or organizations.

Presentation of the Problems

Peter Svanberg gave a short overview of the level of support for non-ASCII character sets in different Internet protocols. Some of the protocols were identified as hostile to 8-bit characters. Among them are: DNS, SMTP, FTP, NNTP, WAIS, MIME Text/Enhanced, NFS, AFS, Whois, URN, Gopher, etc. The more recently developed protocols such as MIME part 1 and part 2 as well as some currently on-going projects such as Whois++, as mentioned by Simon Spero, support 16-bit coding and the repertoires provided by such coding. He also mentioned, that several IETF groups developing new protocols/services consider the importance of the proper support of the character sets to be a problem. The level of support for extended character sets in some protocols used on the Internet is included in the Annex below.

The next speaker was Masataka Ohta. He presented his view regarding the idea that the International Universal Coding system be recommended for use over the Internet. He identified five properties which are required to be present in the recommended coding system:

1. Identity for encoding and decoding, which he understands as unique mapping between particular graphic character and its code (bit combination);
2. Causality, understood as independence of a processed coded character from the other incoming characters in the data stream;
3. Finite state recognition, state dependence of the code required for presentation/display of multi-octet coded data;

4. Finite resynchronizability, which means that the state of automation can be determined uniquely by reading a fixed, finite number of octets; and
5. Equality, requirement that a character coded with a different coding system can always be recognized as the same character.

Masataka looked for the required properties in ISO 10646 and found out that full ISO 10646 (UCS4) satisfies none of the required properties. He also pointed out that ISO 10646 level 1 satisfies all of the required properties for the European languages.

He proposed an extension to the existing UCS code system consisting of five additional bits which will enable the deficiency of the UCS coding system to be overcome. The discussion showed that the proposed solution is not in the general stream of the development of the standard character set codes and their applications in the computing systems. One of the possible solutions to the problems identified by Masataka could be the use of the whole model of UCS, i.e., the four envisaged octets which define the cell and row position for a character in the Multilingual Basic Plane of ISO 10646 additional planes and groups. There was a proposal that the required five additional bits be coded as a private plane in the UCS scheme. John Klensin noted that such an approach could clash with the reassignment of such a plane in the standardization process of ISO JTC1/SC2. In the discussion the problem of the handling of bidirectional text was also identified. Masataka said that one of the five additional bits in his scheme is intended to be used for indication of bidirectional text.

Harald Alvestrand pointed out that what is happening now is a sort of transition period between 8-bit coding and 16-bit coding provided with UCS. Another parallel stream for support of different national character sets is "character switching" which is enabled by use of the code extension technique of ISO 2022. It was obvious that this scheme is not of practical use for the Internet except for special cases, i.e., the Japanese e-mail solution.

Conclusions

The attendees then discussed possible work items which will result if the IESG approves the formation of a working group. The chair identified several documents which deal with character set problems such as: RFC 1345, "Character Mnemonics & Character Sets," the Internet-Draft, "X.400 use of extended character sets," and the Internet-Draft, "Characters and character sets for various languages." John Klensin pointed out that special precautions have to be taken in the recommendation of UTF-2 as a data interchange method over the Internet in connection with the possible assignments of additional coding planes by JTC1/SC2. He also recommended the use of a mailing list already working within IETF, ietf-charsets@innosoft.com. The mailing list of the RARE working group on character sets could be added to that mailing list. Other items were discussed and proposed by the BOF attendees. It was decided that the IESG will be asked to consider the possibility of setting up a working group to produce the following:

- A document defining how UCS can be used in a uniform way in Internet protocols, especially taking into consideration the UTF-2 encoding of UCS. The document will provide guidance to other protocols which have to deal with these items over the Internet.
- A document identifying the languages and the characters required for coding text written in a particular natural language (a sort of guideline for services dealing with multilinguality such as NIR service based on the usage of plain text).
- A document defining a tool for coded character set conversion to be provided within some services such as e-mail user agent including fall-back representation of incoming characters that are outside the supported character repertoire of the receiver.
- A proposal for extending the mandatory issues which have to be covered in the RFC standardization process to include character set consideration and support.

Annex

The level of support for extended character sets in some Internet Standard protocols.

| CharSet Support | Protocol | CharSet Support | "Next Generation" Protocol |
|-----------------|----------------------|-----------------|----------------------------|
| 1 | SMTP | 3 | ESMTP |
| 1 | RFC822 | 4 | MIME part 1 + part 2 |
| 1 | DNS | | |
| 2 | FTP | | |
| 3 | Telnet | | |
| 2 | NNTP | | |
| 2 | Finger | | |
| 2 | POP3 | | |
| 2 | IMAP2 | 3 | IMAP2bis |
| 1 | NFS | | |
| 1 | AFS | | |
| 2 | MIME Text/Enhanced | | |
| ? | MIME Text/simplemail | | |
| 3 | STIF | | |
| 2 | Gopher | 3 | Gopher + |
| 1 | WAIS | | |
| ? | Prospero | | |
| 2 | HTML | | |
| 2 | Whois | 3 | Whois ++ |
| 2 | URL | | |
| 2 | URN | | |
| 3 | URM | | |

Legend:

- 1 – hostile against 8-bit characters
- 2 – no support for different character sets
- 3 – some support for different character sets
- 4 – well thought-out support for different character sets
- 5 – uniform treatment of all characters

Attendees

| | |
|-----------------------|------------------------------|
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Piet Bovenga | p.bovenga@uci.kun.nl |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Tim Dixon | dixon@rare.nl |
| Olle Jarnefors | ojarnef@admin.kth.se |
| Borka Jerman-Blazic | jerman-blazic@ijs.si |
| Tomaz Kalin | kalin@rare.nl |
| John Klensin | Klensin@infofoods.unu.edu |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Keith Moore | moore@cs.utk.edu |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Luc Rooijackers | lwj@cs.kun.nl |
| Rickard Schoultz | schoultz@admin.kth.se |
| Milan Sova | sova@feld.cvut.cz |
| Simon Spero | simon_spero@unc.edu |
| Peter Svanberg | psv@nada.kth.se |
| Guido van Rossum | guido@cwi.nl |

2.1.1 Interactive Mail Access Protocol (IMAP)

Charter

Chair(s):

Terry Gray, gray@cac.washington.edu

Mailing Lists:

General Discussion: imap@cac.washington.edu

To Subscribe: imap-request@cac.washington.edu

Archive: [ftp.cac.washington.edu:~/imap/imap_archive](ftp://ftp.cac.washington.edu/~imap/imap_archive)

Description of Working Group:

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

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

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

This differs from POP3 (RFC 1225) in that POP is a store-and-forward transport protocol that allows an MUA to retrieve pending mail from a mail drop (where it is then usually deleted automatically), whereas IMAP is focused on remote mailbox manipulation rather than transport. IMAP differs from various vendor-specific remote access approaches in that IMAP is an open protocol designed to scale well and accommodate diverse types of client operating systems. Security-related tasks include how to incorporate secure authentication mechanisms when establishing a session, and possible interactions with Privacy Enhanced Mail.

It is expected that most of the work of this group will be conducted via email. A goal is to integrate and update RFC1176 and the existing IMAP2bis draft,

then submit the result as an Internet-Draft well before the November IETF meeting, which would then focus on detailed review of the text in preparation for submission as a Proposed Standard before the end of 1993.

Goals and Milestones:

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

Internet-Drafts:

“INTERACTIVE MAIL ACCESS PROTOCOL - VERSION 2bis”, 08/16/1993,
M. Crispin <draft-ietf-imap-imap2bis-00.txt>

CURRENT MEETING REPORT

Reported by Terry Gray/University of Washington

Minutes of the Interactive Mail Access Protocol Working Group (IMAP)

Agenda

- Introductions
- IMAP overview
- Comments on the charter
- Status of implementations
- Status of protocol specification
- Comments on Columbus BOF notes
- Additional IMAP change requests
- Seattle meeting

References: /imap/imap* on ftp.cac.washington.edu

Summary

There were 20 people in attendance. For several, it was their first exposure to IMAP, so a few minutes was spent summarizing what IMAP is, how it compares/relates to other alternatives, and what the working group is chartered to do. The working group charter and notes from the Columbus BOF were reviewed and questions were answered. The status of the protocol specification and known IMAP implementations was reviewed. (An Internet-Draft is being composed that integrates and updates RFC 1176 and the imap2bis extensions.) Existing practice on the use of IMAP for news, archive, and document access (in addition to mail) was covered. Discussion on possible IMAP extensions followed. Finally, the next working group meeting (in Seattle, August 30-31) was announced.

Discussion Points

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

Action Items

Terry Gray needs to maintain (or cause to be maintained) an IMAP enhancement/request list, sorted into the following categories:

- Protocol bug fixes
- Upward compatible extensions
 - high priority
 - lower priority
- Non-upward compatible changes
 - high priority
 - lower priority
- Bad, or not clearly good, ideas

A subset of that list must then be defined as the target for the immediate standardization effort, with other ideas being deferred for future consideration. Given the desire to preserve compatibility with the installed base, and move ahead promptly in getting a base IMAP standard defined, extensions will be necessarily limited to those deemed to have an extremely high priority.

Mark Crispin needs to integrate RFC 1176 text with IMAP2BIS text and submit it as an Internet-Draft no later than August 15th.

IMAP implementors/interested parties are encouraged to come to the next meeting in Seattle, August 30-31.

Attendees

| | |
|-----------------------|------------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| James Allard | jallard@microsoft.com |
| Luc Boulianne | lucb@cs.mcgill.ca |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Al Costanzo | al@akc.com |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Christoph Graf | graf@switch.ch |
| Terry Gray | gray@cac.washington.edu |
| Jari Hamalainen | jah@rctre.nokia.com |
| Xander Jansen | xander.jansen@surfnet.nl |
| Scott Kaplan | scott@wco.ftp.com |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| Keith Moore | moore@cs.utk.edu |
| Mel Pleasant | pleasant@hardees.rutgers.edu |
| Robert Reschly | reschly@brl.mil |
| Kenneth Rossen | kenr@shl.com |
| David Sitman | a79vm.tau.ac.il |
| Peter Svanberg | psv@nada.kth.se |
| Gregory Vaudreuil | gvaudre@cnri.reston.va.us |

2.1.2 Internet Message Extensions (822EXT)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: ietf-822@dimacs.rutgers.edu

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

Archive: ietf.cnri.reston.va.us:~/ietf-mail-archive/822ext/*

Description of Working Group:

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

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

Goals and Milestones:

- | | |
|------|---|
| Done | Review the Charter, and refine the Group's focus. Decide whether this is a worthwhile effort. |
| Done | Discuss, debate, and choose a framework for the solution. Assign writing assignments, and identify issues to be resolved. |
| Done | Review exiting writing, resolve outstanding issues, identify new work, and work toward a complete document. |
| Done | Post a first Internet-Draft. |
| Done | Review and finalize the draft document. |
| Done | Submit the document as a Proposed Standard. |
| Done | Post an Internet-Draft for the use of Japanese Characters for Internet Mail. |
| Done | Post a revised version of the MIME document as an Internet-Draft. |
| Done | Submit the revised MIME document to the IESG for consideration as a Draft Standard. |
| Done | Submit the Japanese Character set specification as an Informational document. |

Internet-Drafts:

“MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies”, 02/08/1993, N. Borenstein, N. Freed <draft-ietf-822ext-mime2-04.txt, .ps>

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

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

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

Request For Comments:

RFC 1341 “MIME (Multipurpose Internet Mail Extensions): Mechanisms for Specifying and Describing the Format of Internet Message Bodies”

RFC 1342 “Representation of Non-ASCII Text in Internet Message Headers”

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

RFC 1468 “Japanese Character Encoding for Internet Messages”

2.1.3 MIME-MHS Interworking (MIMEMHS)

Charter

Chair(s):

Steve Thompson, slt@gateway.ssw.com

Mailing Lists:

General Discussion: mime-mhs@surfnet.nl

To Subscribe: mime-mhs-request@surfnet.nl

Archive:

Description of Working Group:

MIME, (Multipurpose Internet Mail Extensions) is currently a Draft Standard. MIME redefines the format of message bodies to allow multi-part textual and non-textual message bodies to be represented and exchanged without loss of information. With the introduction of MIME as a Draft Standard it is now possible to define mappings between RFC 822 content-types and X.400 body parts. The MIME-MHS Interworking Working Group is chartered to develop these mappings, providing an emphasis on both interworking between Internet and MHS mail environments and also on tunneling through these environments. These mappings will be made in the context of an RFC 1148bis environment.

Goals and Milestones:

- Done Post an Internet-Draft describing MIME-MHS Interworking.
- Done Post an Internet-Draft describing the “core” set of Registered conversions for bodyparts.
- Done Submit a completed document to the IESG describing MIME-MHS Interworking as a Proposed Standard.
- Done Submit the “core” bodyparts document to the IESG as a Proposed Standard.

Request For Comments:

- RFC 1494 “Equivalences between 1988 X.400 and RFC-822 Message Bodies”
- RFC 1495 “Mapping between X.400 and RFC-822 Message Bodies”
- RFC 1496 “Rules for downgrading messages from X.400/88 to X.400/84 when MIME content-types are present in the messages”

2.1.4 Network News Transport Protocol (NNTP)

Charter

Chair(s):

Eliot Lear, lear@sgi.com

Mailing Lists:

General Discussion: ietf-nntp@turbo.bio.net

To Subscribe: ietf-nntp-request@turbo.bio.net

Archive:

Description of Working Group:

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

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

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

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

Goals and Milestones:

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

2.1.5 OSI Directory Services (OSIDS)

Charter

Chair(s):

Steve Kille, S.Kille@isode.com

Mailing Lists:

General Discussion: ietf-osi-ds@cs.ucl.ac.uk

To Subscribe: ietf-osi-ds-request@cs.ucl.ac.uk

Archive:

Description of Working Group:

The OSI-DS group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

- | | |
|---------|---|
| Ongoing | Maintain a Schema for the OSI Directory on the Internet. |
| Ongoing | Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum. |
| Done | Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced. |
| Done | Study the relationship of the OSI Directory to the Domain Name Service. |

Internet-Drafts:

“DSA Metrics”, 09/23/1992, P. Barker, R. Hedberg <draft-ietf-osids-dsa-metrics-01.txt>

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

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

Request For Comments:

- RFC 1275 “Replication Requirements to provide an Internet Directory using X.500”
- RFC 1276 “Replication and Distributed Operations extensions to provide an Internet Directory using X.500”
- RFC 1277 “Encoding Network Addresses to Support Operation Over Non-OSI Lower Layers”
- RFC 1278 “A String Encoding of Presentation Address”
- RFC 1279 “X.500 and Domains”
- RFC 1384 “Naming Guidelines for Directory Pilots”
- RFC 1430 “A Strategic Plan for Deploying an Internet X.500 Directory Service”
- RFC 1431 “DUA Metrics”
- RFC 1484 “Using the OSI Directory to achieve User Friendly Naming (OSI-DS 24 (v1.2))”
- RFC 1485 “A String Representation of Distinguished Names (OSI-DS 23 (v5))”
- RFC 1487 “X.500 Lightweight Directory Access Protocol”
- RFC 1488 “The X.500 String Representation of Standard Attribute Syntaxes”

CURRENT MEETING REPORT

Reported by Thomas Lenggenhager/SWITCH

Minutes of the OSI Directory Services Working Group (OSIDS)

OSIDS met jointly with RARE WG-NAP.

Introduction

The agenda and the minutes of the Columbus meeting (March 1993) were reviewed and approved without comment.

Liaison Reports

Editor's Note: Liaison reports are available via FTP or mail server from the remote directories as /ietf/osids/osids-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Progression of Documents to Standard Status (Erik Huizer)

- LDAP was delayed due to comments on the lists. Tim Howes has fixed the document and submitted it directly to the RFC Editor.
- DN and UFN are on the top of the RFC Editor's list of documents to publish.
- DSA-Metrics was not on the agenda, but was sent to the list already some time ago. The document was revised by Paul Barker and Roland Hedberg based on the experience with the Siemens' DSA at SURFnet. Paul Barker will publish DSA-Metrics as an Informational RFC.
- Steve Kille will make editorial changes to RFC 1278.

Schema Group/RFC 1274 Update

Sri Sataluri will take over the Schema group which will propose and implement a mechanism on how to maintain RFC 1274. Further volunteers are Tim Howes, Ken Rossen and Russ Wright; there were no European volunteers.

Presentation of InterNIC Activity (Sri Sataluri)

The Directory Services part of InterNIC (contracted to AT&T) offers services on ds.internic.net.

- Directory of Directories
- Directory Services (information about users, institutions, organizations and resources)

X.500 for White Pages: a position paper is available on ds.internic.net as /pub/internic-info/x500.position.paper.

- Public DUA (with login name x500)
- LDAP server
- E-mail access via mailserv@ds.internic.net
- WAIS server (people.src)
- Gopher access is planned

Organizations can have up to 50 entries at no charge by using the template available on ds.internic.net as /pub/internic-info/org.x500.form.

OSI-DS-41: Guidelines for Directory Structure/RFC 1384 Update (Thomas Lenggenhager)

This draft document comes out of the RARE WG-NAP Task Force for Data Management. It tries to summarize all information an organization needs when it wants to use X.500 as a White Pages service.

- How to structure an organizational DIT
- Naming of entries
- Attribute syntaxes, use of T.61
- Languages (no support in the standard at all)
- How to use selected attribute types for a White Pages service

Privacy and data protection will be covered in a separate RFC by Erik Huizer (RARE WG-NAP TF-Legal and IDS). It was agreed that this kind of information shall be integrated into an update of RFC 1384. Paul Barker, Steve Kille and Thomas Lenggenhager will update RFC 1384.

OSI-DS-40: X.500-based File Archive Searching (Paul Barker)

- It is something like Archie but based on X.500.
- The update since the last meeting dropped many special syntaxes.
- A publically accessible interface will be available within 2 weeks.
- A software package to be available within 3-4 weeks.

- Steve Kille will add this topic to the workplan to be progressed as an Experimental RFC some time in the future.

Index DSAs (Paul Barker)

He presented his ideas on specialized DSAs which replicate just the part of the data the DSA administrator is interested in. For the rest of the information, the DN of the entry is available. This is useful for Yellow Page services like finding all biologists in a country, all file archives, or all organizations in the world.

This method would allow subtree searching also near the top without excessive 'costs'. It is just another view to the same data. This would require the DUAs to 'know' where these special DSAs are.

An alternative approach is to have a web of aliases, then query a special part of the DIT instead of special DSAs.

Currently 3 special DSAs are implemented:

- People in all UK computer departments
- All British librarians
- All companies worldwide

OSI-DS-38: Representing IP information in the DIT (Thomas Johannsen)

There have been some attempts to introduce IP information into the DIT. This solution has been generalized and will now get implemented for IP. Thomas Johannsen will submit OSI-DS-37 and 38 for publication as Experimental RFCs by September of 1993.

Network Information in the Directory: Deployment Strategy (Thomas Johannsen)

- No migration plans
- If the X.500 solution is useful, it will be used
- Timescale is needed before it can be progressed
- Special DUAs are needed for network and system managers
- The White Pages tree shall be used instead of special trees

Representing the DNS in the Directory/Revising RFC 1279 (Thomas Johannsen)

Changes:

- More administrative and technical aspects
- Distinction of DNS record types
- Linking DNS ↔ Network info and DNS ↔ White Pages

Steve Kille will add an update of RFC 1279 to the workplan, and Thomas Johannsen will publish the OID tables.

Multiple Service Providers and Distributed Entries (Paul-Andre Pays)

- There are only questions, no solutions yet. One real life object may have several sources of information, either with more or less complete information each, or with only a few attributes.
- Multiple service providers: proposals by NADF (naming links) and Bellcore (reference link with context and DN).
- Distributed entries: attributes held in different databases for one real world object (e.g. phone number in a PABX).
- There were some doubts whether the NADF solution will scale—this has to be shown first. Could collective attributes solve the problem?
- The RARE WG-NAP will prepare a document to state the problem, and OSI-DS will follow up on it.

Next Meeting

OSIDS plans to meet at the Houston IETF in November, towards the end of the week.

Attendees

| | |
|-------------------|--------------------------------------|
| Claudio Allocchio | Claudio.Allocchio@elettra.trieste.it |
| Paul Barker | p.barker@cs.ucl.ac.uk |
| Piet Bovenga | p.bovenga@uci.kun.nl |
| C. Allan Cargille | allan.cargille@cs.wisc.edu |
| Robert Cooney | cooney@wnyose.nctsw.navy.mil |

| | |
|--------------------------|--|
| Thomas DeWitt | tdewitt@osi.ncsl.nist.gov |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Urs Eppenberger | eppenberger@switch.ch |
| Tony Genovese | genovese@es.net |
| Christoph Graf | graf@switch.ch |
| Jan Hansen | Jan.Hansen@teknologi.agderforskning.no |
| Jeroen Houttuin | houttuin@rare.nl |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Ole Jacobsen | ole@interop.com |
| Thomas Johannsen | thomas@ebzaw1.et.tu-dresden.de |
| Ola Johansson | ojn@tip.net |
| Philip Jones | p.jones@jnt.ac.uk |
| Kevin Jordan | Kevin.E.Jordan@cdc.com |
| Peter Jurg | jurg@surfnet.nl |
| Anders Karlsson | sak@cdg.chalmers.se |
| Steve Kille | S.Kille@isode.com |
| Paul Klarenberg | klarenberg@netconsult.ch |
| Andrew Knutsen | andrewk@sco.com |
| Bruno Koechlin | Bruno.Koechlin@inria.fr |
| Mark Kosters | markk@internic.net |
| Arnold Krechel | krechel@gmd.de |
| Sylvain Langlois | Sylvain.Langlois@exp.edf.fr |
| John Larson | jl Larson@parc.xerox.com |
| Erik Lawaetz | erik.lawaetz@uni-c.dk |
| Jose Legatheaux Martins | jalm@fct.unl.pt |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| John Martin | John.Martin@newcastle.ac.uk |
| Ignacio Martinez | martinez@rediris.es |
| Brian May | Brian.May@mel.dit.csiro.au |
| Linda Millington | l.millington@noc.ulcc.ac.uk |
| Jun Murai | jun@wide.ad.jp |
| Paul-Andre Pays | pays@faugeres.inria.fr |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Jim Romaguera | romaguera@netconsult.ch |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Kenneth Rossen | kenr@shl.com |
| Srinivas Sataluri | sri@internic.net |
| Wolfgang Schneider | schneiw@darmstadt.gmd.de |
| Heiner Schorn | heiner.schorn@umdac.umu.se |
| Robert Shirey | shirey@mitre.org |
| John Stewart | john@bunter.fdc.iaf.nl |
| Catherine Treca | Catherine.Treca@dione.urec.fr |
| Panos-Gavriil Tsigaridas | Tsigaridas@fokus.gmd.de |

| | |
|------------------|-------------------------------|
| Eftimios Tsigros | tsigros@helios.iihe.rtt.be |
| Paul Vetter | |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Russ Wright | wright@lbl.gov |
| Peter Yee | yee@atlas.arc.nasa.gov |
| Steve Zeber | zeber@stc.nato.int |

2.1.6 TELNET (TELNET)

Charter

Chair(s):

Steve Alexander, stevea@lachman.com

Mailing Lists:

General Discussion: telnet-ietf@cray.com

To Subscribe: telnet-ietf-request@cray.com

Archive:

Description of Working Group:

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

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

(2) Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically: Environment variable passing, Authentication, Encryption, and Compression.

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

Goals and Milestones:

- | | |
|----------|---|
| Done | Write an environment option. |
| Done | Post an Internet-Draft describing the authentication option. |
| Done | Post an Internet-Draft describing the encryption option. |
| Mar 1991 | Rewrite RFC 854. |
| Done | Submit the authentication option to the IESG as an Experimental Protocol. |
| Jul 1993 | Submit the encryption option to the IESG as an Experimental Protocol. |

Internet-Drafts:

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

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

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

“TELNET Transfer Control Option”, 06/22/1993, S. Denton <draft-ietf-telnet-transfer-option-00.txt>

Request For Comments:

RFC 1116 “Telnet Linemode option”

RFC 1184 “Telnet Linemode Option”

RFC 1372 “Telnet Remote Flow Control Option”

RFC 1408 “Telnet Environment Option”

RFC 1409 “Telnet Authentication Option”

RFC 1411 “Telnet Authentication: Kerberos Version 4”

RFC 1412 “Telnet Authentication : SPX”

RFC 1416 “Telnet Authentication Option”

CURRENT MEETING REPORT

Reported by Steve Alexander/Lachman Technology

Minutes of the TELNET Working Group (TELNET)

Proposed Agenda

- Charter
- Authentication Option
- Transfer of Control Proposal

Summary

The agenda was augmented with a presentation from John Klensin, one of the Area Directors for the Applications Area. John discussed the IESG's views on the environment option controversy.

This shifted the primary focus of the meeting to the environment option. In Columbus, the group had proposed to re-issue RFC 1408 with the VAR and VALUE definitions corrected to match the BSD implementation. Additionally, a document was produced that explained heuristics that could be used to handle implementations that did not agree.

Upon further reflection it was decided that since the heuristics were not 100% deterministic that it would be better to re-issue RFC 1408 using a different option value. In this way, two systems implementing the new option could determine that heuristics were not needed. Most working group members felt that in the long run this would be cleaner. There is a further advantage in that an earlier change requested by the IESG (USERVAR) means that systems using the current option values and the BSD definitions don't work with implementations deployed prior to the introduction of USERVAR.

This issue will be re-discussed on the mailing list in order to reach the widest possible audience.

The group briefly discussed the proposed revisions to the charter. The general consensus is that we don't want to remain as a clearing-house for TELNET documents. The charter will be revised to only cover the authentication work that is going on.

The group reviewed a proposal for a "TELNET Transfer of Control" option. The group feels that there are some security issues with the document. Various people will forward their comments directly to the author. At this time the group does not wish to formally work on this specification.

Finally, the group briefly discussed Dave Borman's new option that merged authentication and encryption. The initial feedback is favorable, but the group is waiting to see the results of an initial implementation using Kerberos IV.

Attendees

| | |
|-----------------|---------------------------|
| Steve Alexander | stevea@lachman.com |
| Jim Barnes | barnes@xylogics.com |
| Julian Bates | bates@xylogics.com |
| David Borman | dab@cray.com |
| Peter Cameron | cameron@xylint.co.uk |
| Sun-Kwan Kimn | sunkimn@hp.com |
| John Klensin | Klensin@infofoods.unu.edu |
| Kent Malave | kent@bach.austin.ibm.com |
| Marjo Mercado | marjo@cup.hp.com |
| Greg Minshall | minshall@wc.novell.com |
| Clifford Neuman | bcn@isi.edu |
| Sam Sjogren | sjogren@tgv.com |
| Theodore Ts'o | tytso@mit.edu |
| Gordon Young | young@xylogics.com |

2.1.7 Telnet TN3270 Enhancements (TN3270E)

Charter

Chair(s):

Robert Moskowitz, 3858921@mcimail.com

Mailing Lists:

General Discussion: tn3270e@list.nih.gov

To Subscribe: listserv@list.nih.gov

In Body: sub tn3270e <first_name> <last_name>

Archive: listserv@list.nih.gov

Description of Working Group:

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

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

The working group will produce: an Informational RFC documenting current TN3270 terminal practices, an Experimental RFC describing an interim approach to printing and LU name selection (this will address the work that is already under way and implementations of this partial solution that are already in place), and a standards-track RFC specifying the TELNET protocols that support a fully functional 3270 display and printing environment. This RFC will supersede RFC 1041 and the Experimental RFC describing the interim approach to printing and LU name selection.

Goals and Milestones:

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

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

Internet-Drafts:

“TN3270 Enhancements”, 07/26/1993, B. Kelly <draft-ietf-tn3270e-enhancements-01.txt>

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

“TN3270 Current Practices”, 08/18/1993, J. Penner <draft-ietf-tn3270e-current-pract-00.txt>

2.1.8 X.400 Operations (X400OPS)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: ietf-osi-x400ops@cs.wisc.edu
 To Subscribe: ietf-osi-x400ops-request@cs.wisc.edu
 Archive:

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this group is to insure interoperability between Internet X.400 management domains and the existing Internet mail service. The specific task of this group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

| | |
|----------|--|
| Done | Initial meeting, produce internal outline. |
| Done | Working draft, circulate to interested people. |
| Done | Internet-Draft available. |
| Dec 1991 | Document ready for publication. |

Internet-Drafts:

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

“Postmaster Convention for X.400 Operations”, 11/23/1992, C. A. Cargille <draft-ietf-x400ops-postmaster-02.txt>

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

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

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

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

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

Request For Comments:

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

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

RFC 1502 “X.400 Use of Extended Character Sets”

CURRENT MEETING REPORT

Reported by Jim Romaguera/NetConsult AG

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

The following is an outline of items that were covered during the X400OPS sessions that took place at the 27th IETF:

1. Introduction
 - Agenda Approval
 - Minutes Approval
2. Action Items
3. Liaison Reports
 - MHSDS
 - RARE WG-MSG
 - EMA/EEMA
 - COSINE MHS
4. Document Review
5. DNS (Claudio Allocchio)
6. E-mail File Distribution (Marko Kaotia)
7. X400 OPS Charter
8. Aconet and ADA ADMD (Christian Panigil)
9. A=IMX
10. ADMD Interconnection BOF
11. TF88 (Erik Huizer)
12. Support of GO-MHS/Mail Based Server (Erik Huizer)
13. Gateway to LANs

Editor's Note: More detailed itemizations for each of these topics are available via FTP or mail server from the remote directories as /ietf/x400ops/x400ops-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Attendees

| | |
|-----------------------|--|
| Claudio Allocchio | Claudio.Allocchio@elettra.trieste.it |
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Per Andersson | pa@cdg.chalmers.se |
| George Chang | gkc@ctt.bellcore.com |
| Robert Cooney | cooney@wnyose.nctsw.navy.mil |
| Thomas DeWitt | tdewitt@osi.ncsl.nist.gov |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Urs Eppenberger | eppenberger@switch.ch |
| Osten Franberg | euaokf@eua.ericsson.se |
| Tony Genovese | genovese@es.net |
| Tim Goodwin | tim@pipex.net |
| Christoph Graf | graf@switch.ch |
| Alf Hansen | Alf.Hansen@delab.sintef.no |
| Jan Hansen | Jan.Hansen@teknologi.agderforskning.no |
| Jeroen Houttuin | houttuin@rare.nl |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Ola Johansson | ojn@tip.net |
| Kevin Jordan | Kevin.E.Jordan@cdc.com |
| Peter Jurg | jurg@surfnet.nl |
| Thomas Kaepfner | kaepfner%heidelbg.vnet@ibmpa.ibm.com |
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Paul Klarenberg | klarenberg@netconsult.ch |
| Bruno Koechlin | Bruno.Koechlin@inria.fr |
| Arnold Krechel | krechel@gmd.de |
| Erik Lawaetz | erik.lawaetz@uni-c.dk |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| Ignacio Martinez | martinez@rediris.es |
| Ronny Nilsen | Ronny.Nilsen@usit.uio.no |
| Christian Panigl | christian.panigl@cc.univie.ac.at |
| Paul-Andre Pays | pays@faugeres.inria.fr |
| Jim Romaguera | romaguera@netconsult.ch |
| Marjo Rottschaefer | |
| Hugh Smith | h.smith@cs.nott.ac.uk |
| Suzanne Smith | smith@es.net |
| Eftimios Tsigros | tsigros@helios.iihe.rtt.be |
| Russ Wright | wright@lbl.gov |
| Peter Yee | yee@atlas.arc.nasa.gov |

2.2 Internet Area

Director(s):

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

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

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

- IP and multi-protocol operation over emerging wide area technologies (ATM, SMDS, Frame Relay) and point-to-point technologies (including narrowband ISDN).
- Development of a “next generation” IP; i.e., a replacement protocol and addressing/routing architecture for IPv4.
- Miscellaneous (Network Address Translation, Stream Technology 2).

The following BOFs and working groups in the Internet Area met during the Amsterdam IETF:

- Extensions to OSI for use in the Internet BOF (OSIEXTND)
- Internet Stream Protocol V2 BOF (ST2)
- IPng Decision Process BOF (IPDECIDE)
- Network Address Translators BOF (NAT)
- IP Over Asynchronous Transfer Mode Working Group (ATM)
- IP Over Large Public Data Networks Working Group (IPLPDN)
- P. Internet Protocol Working Group (PIP)
- Point-to-Point Protocol Extensions Working Group (PPPEXT)
- Simple Internet Protocol Working Group (SIP)
- TCP/UDP Over CLNP-Addressed Networks Working Group (TUBA)
- TP/IX Working Group (TPIX)

The IPLPDN and PPPEXT Working Groups met individually as well as jointly.

Internet Stream Protocol V2 BOF (ST2)

The ST2 BOF resulted in two decisions. The first was that a working group should be formed to update the existing ST-II specification (RFC 1190). The main motivation was to correct errors in the specification and to make it easier to implement ST-II in a manner which is likely to be interoperable with other ST-II implementations. The second decision

was that there was no clear direction on future standards track status for ST-II. A mailing list will be started in the next week to discuss a proposed charter.

IPng Decision Process BOF (IPDECIDE)

About 200 people attended, plus about 100 MBONE auditors. Discussion focussed on the decision process for IPng rather than on technical criteria or the proposals. A clear consensus did not emerge, but constant themes in the discussion included these:

1. Vendors and operators look to the IETF to reach a clear decision.
2. It would be bad to offer the market an ambiguous decision.
3. The market will resist any IPng that does not just look like a new release of IP.
4. It is unclear how to prove that any proposal truly scales to a billion nodes.
5. Timescales for IPv4 address depletion and for IPng deployment are not well understood.
6. The IESG needs to figure out how to pursue the decision process and avoid wasted effort on competing proposals.

Network Address Translators BOF (NAT)

Kjeld Borch Egevang's NAT implementation is described in the NAT Internet-Draft. The scheme in that draft is not dynamic in that the addresses used for translation are statically assigned to single hosts for long periods of time. It is possible, however, to re-assign them to other hosts. Another aspect of the scheme described is that the addresses on the backbone side of the translator must be globally unique. It was pointed out that other NAT schemes do not have these characteristics (for instance, one proposed by Van Jacobson).

It was generally felt that it would be useful to the IP community to have more knowledge of the pitfalls of NAT. This is particularly true because anybody can install a NAT box independent of anybody else, and in the absence of any NAT standard.

IP Over Asynchronous Transfer Mode Working Group (ATM)

The first session began with an announcement by Bob Hinden that Mark Laubach will be the new ATM chair. An agenda was presented and agreed upon for the three days. Steve Willis presented and led a review of recent ATM Forum activities. Their User Network Interface (UNI) Specification Version 3.0 document should be ratified in August. Juha Heinanen presented an overview of the European ATM pilot project. Joel Halpern presented the topic of "routing IP over the switched virtual cloud." He volunteered to write a proposal. Consensus is that ATM will host the proposal but actual work will be moved to another working group that will deal with routing over large public networks. A general discussion

was held to collect comments on Ran Atkinson's MTU draft document. Ran was not present at this meeting. In the last half hour was spent on Mark Laubach's Classical IP and ARP Over ATM draft and discussion and consensus building continued over the next two meetings.

The second session opened with a discussion of the time-table of working group activities over the next half year. In order to fast-track this document, technical review and final consensus on the draft will be collected via email.

Unfortunately, discussion of the classical draft and related issues took up most of the time of the working group. On the last day, Juha was given twenty minutes to lead the discussion of his NBMA draft. Clearly this was not enough time as much discussion was generated. Juha is getting together with Joel Halpern to work on the issues raised in the discussion.

IP Over Large Public Data Networks Working Group (IPLPDN)

IPLPDN met individually and with the PPPEXT Working Group and reached the following decisions:

- A request will be made to advance RFC 1356, "Multiprotocol over X.25" from Proposed to Draft Standard.
- The default encapsulation for circuit switched services will be PPP.
- Concensus was reached for the PPP over X.25, ISDN, and Frame Relay documents. They will be updated in the coming weeks.
- Progress was made on the definition of multilink transfer.

The IPLPDN group met for the last time, but the mailing list will remain for unfinished business. Remaining topics will go to ATM, PPPEXT, and newly created working groups.

P. Internet Protocol Working Group (PIP)

A specification overview was presented at the meeting. The specification of forwarding has remained unchanged for the past three months. The DNS architecture to support PIP has been revised. The PIP identifier structure has been revised. IDRIP routing support for PIP has revisions in progress. The host operations specifications has been revised. The PIP Control Message Protocol is new, and is currently incomplete. The PIP transition specification is new. Missing from the specification is a MIB definition. Routing still requires further definition.

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

- RFC 1171 should be Historical. When updated, the current PPP LCP draft should go to Draft Standard.
- The HDLC Framing draft is a direct extraction from the older PPP LCP document, and is ready for elevation to Draft Standard.
- The PPP LCP Extensions draft is recommended for consideration as a Proposed Standard.
- The PPP requirements document will be reorganized and posted as an Informational RFC.
- A separate breakout meeting was held for PPP Compression, and the slides from the two presentations by Dave Rand and Dave Langley are included with the minutes. They contain a lot of information.

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

- The question was seriously discussed whether a default way to send IP over circuit switched services such as ISDN B channel is needed. It was observed that the question is malformed; a default way to send IP over a V.35 or V.11 interface, for example, is not needed. A way to speak to a peer system at the data link layer, which might be a Frame Relay or X.25 switch, or a peer host or router is needed.
- Various discussions were held about configuring multi-link PPP groups. Discussions were also held about adding, removing, and controlling separate lines in a multilink group.
- An applicability statement for PPP over Frame Relay is needed.
- An applicability statement for PPP over X.25 in view of RFCs 877 and 1356 is needed.
- Bill Simpson presented his paper on PPP over ISDN. PPP must have the same default MRU (and any other defaults) on ISDN as in other environments. Keith Sklower will publish his IPLPDN document, "Determination of Encapsulation of Multi-Protocol Datagrams in Circuit Switched Environment," and Bill indicated that he would like to copy some of the technical material from them into this document. It was decided that he would reference Keith's document.

- Keith and Bill will merge their documents. The resulting document should be separate from the PPP over *foo* documents, as it is desired to be placed on the standards track, and the PPP over *foo* documents may not be placed on that track.

Simple Internet Protocol Working Group (SIP)

The group reviewed recent work on system discovery, auto configuration, SIP RIP and SIP DNS. Translation issues raised by Ross Callon were also reviewed. Gary Malkin presented comments received from Garcia Luna Aceves (JJ) on SIP RIP. Bill Simpson led the discussion on the system discovery draft. Bob Gilligan presented a set of “preliminary ideas” that he proposed to the mailing list on auto configuration.

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

Marcel Wiget (Switch) reported on the RARE CLNS Pilot. Applications tested included: X.400, X.500, FTAM, DECnet, VT, TUBA, OSI Ping, traceroute, etc. Current activities include a task force for CLNS routing coordination. A spirited discussion was held on the use of IS-IS for routing the global CLNS network.

TP/IX Working Group Working Group (TPIX)

The TP/IX Working Group conducted its first meeting. There were two sessions. Two new Internet-Drafts were presented in the first session, “Initial AD Assignment Plan” and “Transit Policy Routing in TP/IX.” In the second session, the TCP large window performance options and a new experimental TCP record marking option, both documented in the new TCP options draft, were discussed. All items on the agenda were covered by the working group. The status of TP/IX and RAP protocols was reported, and a lot of questions were answered concerning the transition from IPv4 to TP/IX.

CURRENT MEETING REPORT

Reported by Mark Knopper/Merit Network

Minutes of the Extensions to OSI for use in the Internet BOF (OSIEXTND)

The group's discussion was based on information presentation by Dave Katz.

Editor's Note: The text of Dave Katz's presentation slides is available via FTP or mail server from the remote directories as /ietf/93jul/osiextnd-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The group agreed that a letter should be written to ISO, based on the presentation and discussion, stating that the IETF has sufficient expertise to contribute to these areas. Dave Katz and Dave Piscitello will draft a document.

This work spans about six IETF working groups.

BGP and IPIDRP Working Groups are meeting jointly (IDRP is BGP5). IDRP for SIP is being considered.

Phill Gross expressed concern over the apparent danger of having different IETF and ISO standards for protocols. ISO has more precedence for accepting standards from other groups without change. IETF tends to change or rewrite protocols before acceptance. Peter Furniss said that the Internet-Draft process will be appreciated by ISO.

Mark Knopper has created a discussion list, `osiextnd@merit.edu`. Those interested in being added to the list should send a request to `osiextnd-request@merit.edu`.

Phill Gross suggested to identify work that could be done by existing IETF working groups, as well as that which could be done by this group if it is to become a working group.

CLNP over Large Public Data Networks (LPDN) is an area which needs consideration. Much of the work is done. ES-IS and IS-IS protocols over LPDNs needs further work. Note that as of this IETF, the IPLPDN Working Group has ended their work. Perhaps CLNPLPDN could be handled as a BOF with identified base documents.

The consensus of the BOF attendees was that a working group should be formed from these ideas, and relationships should be pursued with ISO. Dave Piscitello and Dave Katz have drafted a letter and will send it to the IESG, IAB, and ISO (through Jack Houldsworth).

Editor's Note: A copy of the letter to the IESG and IAB is available via FTP or mail server from the remote directories as /ietf/93jul/osiextnd-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Attendees

| | |
|-------------------|--|
| Nick Alfano | alfano@mpr.ca |
| Bernt Allonen | bal@tip.net |
| Rebecca Bostwick | bostwick@es.net |
| Jim Bound | bound@zk3.dec.com |
| Ross Callon | rcallon@wellfleet.com |
| George Chang | gkc@ctt.bellcore.com |
| A. Lyman Chapin | lyman@bbn.com |
| Richard Colella | colella@nist.gov |
| Dave Cullerot | cullerot@ctron.com |
| Toerless Eckert | Toerless.Eckert@informatik.uni-erlangen.de |
| Dino Farinacci | dino@cisco.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Peter Furniss | p.furniss@ulcc.ac.uk |
| Phillip Gross | pgross@ans.net |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Susan Hares | skh@merit.edu |
| Denise Heagerty | denise@dxcoms.cern.ch |
| Jack Houldsworth | J.Houldsworth@ste0906.wins.icl.co.uk |
| Chris Howard | chris_howard@inmarsat.org |
| Geoff Huston | g.huston@aarnet.edu.au |
| David Jacobson | dnjake@vnet.ibm.com |
| Philip Jones | p.jones@jnt.ac.uk |
| Cyndi Jung | cmj@3com.com |
| Anders Karlsson | sak@cdg.chalmers.se |
| Dave Katz | dkatz@cisco.com |
| Sean Kennedy | liam@nic.near.net |
| Mark Knopper | mak@merit.edu |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Ton Koelman | koelman@stc.nato.int |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Robin Littlefield | robin@wellfleet.com |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| David O'Leary | doleary@cisco.com |
| Christian Panigl | christian.panigl@cc.univie.ac.at |
| Alex Reijnierse | a.a.reijnierse@research.ptt.nl |
| Victor Reijs | reijs@surfnet.nl |
| Georg Richter | richter@uni-muenster.de |
| John Scudder | jgs@merit.edu |
| Keith Sklower | sklower@cs.berkeley.edu |
| John Stewart | john@bunter.fdc.iaf.nl |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Marcel Wiget | wiget@switch.ch |

Douglas Williams
Rachel Willmer

dougw@ralvmg.vnet.ibm.com
rachelw@spider.co.uk

CURRENT MEETING REPORT

Reported by Steve DeJarnett/IBM

Minutes of the Internet Stream Protocol V2 BOF (ST2)

The ST2 BOF met during the 1930-2200 session, opposite the Networking Multimedia Applications BOF (MULTIAPP) and the Open IAB Meeting (among others). The overlap with MULTIAPP in particular, and the multiple reschedulings of the BOF, may have limited the attendance, but in spite of those problems there were still 27 attendees.

Presentations and Discussion

The BOF started with presentations on existing ST-II implementations. Frank Hoffmann presented a description of the IBM ST-II implementation in the Heidelberg Transport System (HeiTS). Lou Berger presented information on BBN's work with ST-II in their T/20 router which is used in the backbone of the ARPA Defense Simulation Internet (DSI). Lou also spoke briefly about other work BBN was doing with ST-II. Luca Delgrossi discussed work going on in the BERKOM project in Germany, which has chosen ST-II as their network layer protocol for use in a multimedia teleservices pilot.

After the presentations, discussion about interest in a working group to revise and clarify the ST specification was started. There was a great deal of discussion during the BOF about the problems that are known with ST-II as currently specified, (e.g., the lack of a formal definition of ST-II states and the problems associated with efficiently supporting very large numbers of destinations in a single connection). Additional discussion focused around whether standardization of ST was warranted at this time, given work in other areas (e.g., resource reservation); no clear consensus was reached.

During the presentations on the various implementations, questions were raised which mostly dealt with implementation details. Those details have been left out of this summary except as they relate to the proposed mission of the working group.

Results

In the end, consensus was achieved on having a working group whose purpose was to clarify the current ST protocol in order to correct errors in the specification, and to make it easier to implement ST-II in a manner which is likely to be interoperable with other ST-II implementations. The future of that protocol, whether continuing as an Experimental Protocol or moving forward as a Proposed Standard is still open for discussion. That discussion should be held at a later date when the results of the working group are demonstrable.

Attendees

| | |
|------------------|-------------------------------|
| Lou Berger | lberger@bbn.com |
| Stefan Braun | smb@cs.tu-berlin.de |
| Michael Brescia | |
| James Davin | davin@thumper.bellcore.com |
| Steve DeJarnett | steve@ibmpa.awdpa.ibm.com |
| Luca Delgrossi | luca@ibmpa.awdpa.ibm.com |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Deborah Estrin | estrin@usc.edu |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Marcel Graf | graf%dhdibm1.bitnet@vm.gmd.de |
| Ken Hayward | crm57d@bnr.ca |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| John Larson | jl Larson@parc.xerox.com |
| Susan Lin | suelin@vnet.ibm.com |
| Cynthia Martin | martin@spica.disa.mil |
| Zbigniew Opalka | zopalka@agile.com |
| Jorg Ott | jo@cs.tu-berlin.de |
| Michael Patton | map@bbn.com |
| John Penners | jpenners@advtech.uswest.com |
| J. Mark Pullen | mpullen@cs.gmu.edu |
| Robert Reschly | reschly@brl.mil |
| Kenneth Smith | kensmith@bnr.ca |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Claudio Topolcic | topolcic@cnri.reston.va.us |
| Werner Vogels | werner@inesc.pt |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |

CURRENT MEETING REPORT

Reported by Brian Carpenter/CERN and Tim Dixon/RARE with additional text from Phill Gross/ANS

Minutes of the IPng Decision Process BOF (IPDECIDE)

The IPng Decision Process BOF was intended to help re-focus attention on the very important topic of making a decision between the candidates for IPng. The BOF focused on the issues of who should take the lead in making the recommendation to the community and what criteria should be used to reach the recommendation. The discussion ranged widely, but some key points emerged:

- Vendors and operators look to the IETF to reach a clear decision.
- It would be bad to offer the market an ambiguous decision.
- The market will resist any IPng that does not just look like a new release of IP. Co-existence, and ease and cost of transition, should be key decision criteria.
- It is unclear how to prove that any proposal truly scales to a billion nodes.
- Timescales for IPv4 address depletion and for IPng deployment are not well understood.
- The IESG needs to figure out how to pursue the decision process and avoid wasted effort on competing proposals. Making a reasonable well-founded decision earlier was preferred over taking longer to decide and allowing major deployment of competing proposals.

In the end, the BOF led very productively to a follow-up discussion in the Thursday afternoon open plenary. During the open plenary, a proposal that the IESG should take the lead responsibility for recommending an IPng choice to the IETF community met with strong consensus. This proposal included a series of steps that the IESG should take, with strong community involvement, toward a recommendation.

We now give a more detailed review of the BOF discussion, in the interest of recording the wide range of opinions expressed.

Meeting Goals

The purpose of the BOF was to focus on the decision process for IPng rather than on technical criteria, the proposals themselves, or on the working group process.

Attendance

About 200 people attended, plus about 100 MBONE auditors. Members of the audience represented the IETF's typical wide community of service providers, equipment vendors and engineers.

The Need for a Decision

The view was frequently expressed that a decision was needed. Vendors and operators looked to the IETF to reach a clear decision. The IPng issue had been widely publicized for some time and the expectation clearly was that it was the responsibility of the IETF to decide. Operators simply reacted to the demands of their customers: the IETF must set the technical standards. The IETF was doing a disservice to the community by appearing to be indecisive.

The alternative of "letting the market decide" (whatever that may mean) was criticised on several grounds:

- There are infrastructural issues, like DNS, which go hand-in-hand with the choice of a protocol and which cannot reasonably be expected to deal with 4 protocols.
- There are already enough other choices (both proprietary and otherwise) in the marketplace.
- The decision was too complicated for a rational market-led solution.

The fact that the Internet is doubling in size about every 11 months means that the cost of transition to IPng (in terms of equipment and manpower) is also increasing. The longer it takes to reach a decision, the more costly the process of transition and the more difficult it is to undertake.

There were some minority views expressed, including:

- The decision will inevitably be controlled by the pricing policy of vendors.
- Router vendors are already supporting multiple network-layer protocols; in principle it would not be significantly more difficult to support several IPng solutions at the same time.

Should there be a decision to recommend *one* proposal, or simply to eliminate some of the candidates? Concern was expressed about the feasibility of conducting reasonably-sized trials of more than one selected protocol and of the confusing signals this would send the market: IETF decisions now have an enormous potential economic impact on suppliers of

equipment and services. It was also likely that uncertainty would lead to customers holding back on their purchases of networking equipment until the situation was clearer.

A straw poll showed a clear majority view that there should be a decision for *one* solution.

The Time Scale for a Decision

The best guesstimates for the remaining lifetime of the IPv4 address space put the figure at around five to seven years, assuming CIDR is widely deployed. A margin of potential error in these figures is to be expected—one suggestion was that they could be out by a factor of four in either direction. However, the address space is only five doublings away from exhaustion.

It was strongly recommended that more work be done on investigating the feasible remaining lifetime of IPv4.

It is also difficult to estimate the time taken to implement, test and then deploy any chosen solution: it was not clear who was best placed to do this. The ordering of the decisions might also have a different priority for customers and vendors than for the IETF. For example, it might be necessary to have a decision about DNS changes early in order to deploy the infrastructure necessary to support IPng in advance of the availability of the IPng protocol itself. The IETF work was not proceeding in this order.

The Evaluation Process

Concern was expressed that the evaluation criteria which had so far been discussed were too general to support a defensible choice on the grounds of technical adequacy. The criteria had emerged in parallel with the protocol designs, and had so far not gelled enough to eliminate any candidate. There were also potential legal difficulties if the IETF appeared to be eliminating proposals on arbitrary grounds.

It was stated frequently and forcibly that the transition costs should be a significant factor in the selection criteria. Concerns were expressed by several service providers that the developers had little appreciation of the real-world networking complexities that transition would force people to cope with. If the cost of transition outweighed the pain of other solutions (application gateways or address translators) customers would not deploy IPng.

It was suggested a couple of times that the working groups should be invited to evaluate each others' proposals in order to investigate their weaknesses, or that the proposals should be vetted by disinterested parties. It was suggested that the proposals were too similar for any reasonable choice to be made on the grounds of technical strength. However there was no consensus on these points.

Although one of the goals of IPng had been to use the inevitable transition required by address exhaustion and routing problems to incorporate new features, there were a number of concerns about bundling too much additional complexity into an already difficult problem. It wasn't even clear that the technology yet existed to handle some of the new features that had been touted for IPng. IPng should appear simply like a new release of IPv4; although this would not necessarily bring new features, people would still transition through enlightened self-interest—to avoid disconnection from the global Internet in the future. There was no consensus about how to resolve this dilemma, since both smooth transition and multimedia support are musts.

Various parties were identified as needing to assist in the evaluation process:

- Operators, who need to understand deployment costs and scenarios.
- Vendors, who understand the implementation consequences.

The Decision Process

There is an IETF process for making a decision on protocol standards: working groups can be given deadlines to submit papers to the IESG which then decides which to progress as standards. It was suggested that this process has only broken down in that the deadlines had not been applied.

Other suggestions included:

- Urging coalitions between the different working groups.
- Forming an “IPng” working group either to make recommendations or to draw together the different proposals.
- Asking the IESG or even the IAB to drive the decision process.

On the basis of a straw poll, there was strong consensus that the decision should be made on technical grounds alone (subject to reasonable costs of implementation, deployment and transition).

It was repeatedly stated that an obvious requirement was that the proposed solution should work. There were at least two components to this: interoperability and scaling. This would be difficult to establish without large-scale piloting. There was no consensus on who might reasonably be expected to participate in such an exercise.

The following day, at the Thursday open plenary session, a proposal that the IESG should take the responsibility of recommending an IPng choice to the IETF met with strong consensus. This proposal included a series of steps that the IESG should take to develop a progressive decision with community involvement.

Attendees

| | |
|-------------------|-----------------------------------|
| George Abe | abe@infonet.com |
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Nick Alfano | alfano@mpr.ca |
| James Allard | jallard@microsoft.com |
| Bernt Allonen | bal@tip.net |
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Frederik Andersen | fha@dde.dk |
| Per Andersson | pa@cdg.chalmers.se |
| Toshiya Asaba | asaba@iij.ad.jp |
| Josee Auber | Josee_Auber@hpgnd.grenoble.hp.com |
| Anders Baardsgaad | anders@cc.uit.no |
| Dennis Baker | dbaker@wellfleet.com |
| Jim Barnes | barnes@xylogics.com |
| Tony Bates | tony@ripe.net |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Axel Belinfante | Axel.Belinfante@cs.utwente.nl |
| Vincent Berkhout | berkhout@cs.utwente.nl |
| Per Bilse | bilse@ic.dk |
| Jim Binkley | jrb@ibeam.intel.com |
| Robert Blokzijl | K13@nikhef.nl |
| Rebecca Bostwick | bostwick@es.net |
| Jim Bound | bound@zk3.dec.com |
| Robert Braden | braden@isi.edu |
| Stefan Braun | smb@cs.tu-berlin.de |
| Thomas Brisco | brisco@pilot.njin.net |
| Ronald Broersma | ron@nosc.mil |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Steve Buchko | stevebu@newbridge.com |
| Ross Callon | rcallon@wellfleet.com |
| Peter Cameron | cameron@xylint.co.uk |
| C. Allan Cargille | allan.cargille@cs.wisc.edu |
| Brian Carpenter | brian@dxcern.cern.ch |
| Vinton Cerf | vcerf@cnri.reston.va.us |
| George Chang | gkc@cct.bellcore.com |
| A. Lyman Chapin | lyman@bbn.com |
| Chris Chiotasso | chris@andr.ub.com |
| Henry Clark | henryc@oar.net |

| | |
|---------------------|--|
| Richard Colella | colella@nist.gov |
| David Conrad | davidc@iij.ad.jp |
| Al Costanzo | al@akc.com |
| Stephen Crocker | crocker@tis.com |
| Dave Cullerot | cullerot@ctron.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Stephen Deering | deering@parc.xerox.com |
| Steve DeJarnett | steve@ibmpa.awdpa.ibm.com |
| Kurt Dobbins | dobbins@ctron.com |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Francis Dupont | francis.dupont@inria.fr |
| Toerless Eckert | Toerless.Eckert@informatik.uni-erlangen.de |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Robert Enger | enger@reston.ans.net |
| Hans Eriksson | hans@sics.se |
| Deborah Estrin | estrin@usc.edu |
| Dino Farinacci | dino@cisco.com |
| Stefan Fassbender | stf@easi.net |
| Eric Fleischman | ericf@act.boeing.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Osten Franberg | euaokf@eua.ericsson.se |
| Paul Francis | Francis@thumper.bellcore.com |
| Dan Frommer | dan@jeremy.enet.dec.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Vince Fuller | vaf@stanford.edu |
| Peter Furniss | p.furniss@ulcc.ac.uk |
| Eugene Geer | ewg@cc.bellcore.com |
| Robert Gilligan | Bob.Gilligan@Eng.Sun.Com |
| Joseph Godsil | jgodsil@ncsa.uiuc.edu |
| Tim Goodwin | tim@pipex.net |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Marcel Graf | graf%dhdibm1.bitnet@vm.gmd.de |
| Terry Gray | gray@cac.washington.edu |
| Ron Greve | rgreve@cs.utwente.nl |
| Phillip Gross | pgross@ans.net |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Joel Halpern | jmh@network.com |
| Susan Hares | skh@merit.edu |
| Denise Heagerty | denise@dxcoms.cern.ch |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Robert Hinden | hinden@eng.sun.com |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Marc Horowitz | marc@mit.edu |

| | |
|-------------------------|--------------------------------------|
| Chris Howard | chris_howard@inmarsat.org |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Geoff Huston | g.huston@aarnet.edu.au |
| Phil Irely | pirey@relay.nswc.navy.mil |
| Ole Jacobsen | ole@interop.com |
| David Jacobson | dnjake@vnet.ibm.com |
| Ronald Jacoby | rj@sgi.com |
| Ola Johansson | ojn@tip.net |
| David Johnson | dbj@cs.cmu.edu |
| Laurent Joncheray | lpj@merit.edu |
| Philip Jones | p.jones@jnt.ac.uk |
| Cyndi Jung | cmj@3com.com |
| Thomas Kaepfner | kaepfner%heidelbg.vnet@ibmpa.ibm.com |
| Tomaz Kalin | kalin@rare.nl |
| Scott Kaplan | scott@wco.ftp.com |
| Anders Karlsson | sak@cdg.chalmers.se |
| Daniel Karrenberg | daniel@ripe.net |
| Frank Kastenholz | kasten@ftp.com |
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Sean Kennedy | liam@nic.near.net |
| Stephen Kent | kent@bbn.com |
| Zbigniew Kielczewski | zbig@eicon.qc.ca |
| John Klensin | Klensin@infoods.unu.edu |
| Mark Knopper | mak@merit.edu |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Ton Koelman | koelman@stc.nato.int |
| Mark Kusters | markk@internic.net |
| Glenn Kowack | Glenn@eu.net |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Arnold Krechel | krechel@gmd.de |
| John Larson | jl Larson@parc.xerox.com |
| Eliot Lear | lear@sgi.com |
| Jose Legatheaux Martins | jalm@fct.unl.pt |
| Tony Li | tli@cisco.com |
| Susan Lin | suelin@vnet.ibm.com |
| John Lindsay | lindsay@kingston.ac.uk |
| Robin Littlefield | robin@wellfleet.com |
| Anne Lord | anne@ripe.net |
| Peter Lothberg | roll@stupi.se |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| Paolo Malara | malara@crs4.it |
| Allison Mankin | mankin@cmf.nrl.navy.mil |
| Bill Manning | bmanning@rice.edu |

| | |
|----------------------|----------------------------------|
| David Marlow | dmarrow@relay.nswc.navy.mil |
| Cynthia Martin | martin@spica.disa.mil |
| Ignacio Martinez | martinez@rediris.es |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Keith McCloghrie | kzm@hls.com |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Greg Minshall | minshall@wc.novell.com |
| Keith Mitchell | keith@pipex.net |
| Pushpendra Mohta | pushp@cerf.net |
| Keith Moore | moore@cs.utk.edu |
| Kees Neggers | neggers@surfnet.nl |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Erik Nordmark | nordmark@eng.sun.com |
| David O'Leary | doleary@cisco.com |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Jorg Ott | jo@cs.tu-berlin.de |
| Christian Panigl | christian.panigl@cc.univie.ac.at |
| Andrew Partan | asp@uunet.uu.net |
| Michael Patton | map@bbn.com |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Charles Perkins | perk@watson.ibm.com |
| Drew Perkins | ddp@fore.com |
| David Piscitello | dave@mail.bellcore.com |
| Mel Pleasant | pleasant@hardees.rutgers.edu |
| Willi Porten | porten@gmd.de |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Juergen Rauschenbach | jrau@dfn.de |
| Alex Reijnierse | a.a.reijnierse@research.ptt.nl |
| Victor Reijs | reijs@surfnet.nl |
| Robert Reschly | reschly@brl.mil |
| Georg Richter | richter@uni-muenster.de |
| Dan Romascanu | dan@lannet.com |
| Luc Rooijackers | lwj@cs.kun.nl |
| Marjo Rottschaefer | |
| Hal Sandick | sandick@vnet.ibm.com |
| Miguel Sanz | miguel.sanz@rediris.es |
| Jon Saperia | saperia@tay.dec.com |
| Eve Schooler | schooler@isi.edu |
| John Scudder | jgs@merit.edu |
| Tim Seaver | tas@concert.net |
| Kitty Shih | kmshih@novell.com |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| W. David Sincoskie | sincos@thumper.bellcore.com |
| Simon Spero | simon_spero@unc.edu |
| Vladimir Sukonnik | sukonnik@process.com |

| | |
|-----------------------|---------------------------------|
| Fumio Teraoka | tera@csl.sony.co.jp |
| Marten Terpstra | marten@ripe.net |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Susan Thomson | set@bellcore.com |
| Martin Toet | toet@cs.utwente.nl |
| Antoine Trannoy | trannoy@crs4.it |
| Robert Ullmann | ariel@world.std.com |
| Willem van der Scheun | scheun@sara.nl |
| Guido van Rossum | guido@cwi.nl |
| Werner Vogels | werner@inesc.pt |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Steven Waldbusser | waldbusser@andrew.cmu.edu |
| Sandro Wallach | sandro@elf.com |
| Abel Weinrib | abel@bellcore.com |
| Douglas Williams | dougw@ralvmg.vnet.ibm.com |
| Kirk Williams | kirk@sbctri.sbc.com |
| Steven Willis | steve@wellfleet.com |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Jessica Yu | jyy@merit.edu |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

IPng Decision Process

Goal of BOF:

Consensus on decision processes

for next generation IP

We will NOT discuss

- **technical criteria**
- **details and merits of proposals**
- **CIDR**

We WILL discuss

- **Who decides?**
Is there a decision point?
Will the market decide?
- **Evaluation procedures (analysis, bake-off, ...)**
- **Do we attempt to combine strong points?**
- **How to get multiple implementations for bake-offs if the market is to decide?**
- **How to get a core of pilot operators?**
- **Open up the discussion to Internet users?**
(WAN and LAN network operators as well as end users)
- **How do we measure market approval?**

- **What timescale is realistic for the decision process?**
If Internet runs out of addresses in year N,
 - **when (year N-x) should the decision be made?**
 - **when (year N-y) should the software be ready for field test?**
 - **what is the latest time (N-z) to start the transition?**
- **Does NSF solicitation affect timescale?**
- **Can we avoid an irrational choice?**
- **Does decision process need financial or manpower support?**

CURRENT MEETING REPORT

Reported by Paul Francis/Bellcore

Minutes of the Network Address Translators BOF (NAT)

Purpose

The purpose of the meeting was to:

- Describe Kjeld Egevang's implementation of a simple NAT box.
- Determine what benefits might come from NAT.
- Determine what problems exist with NAT.
- Determine how we might use Kjeld's implementation to learn more about NAT.

Kjeld's Implementation

Kjeld's NAT implementation is described in the NAT Internet-Draft. The scheme in that document is not dynamic in that the addresses used for translation are statically assigned to single hosts for long periods of time. It is possible, however, to re-assign them to other hosts. Another aspect of the scheme described is that the addresses on the backbone side of the translator must be globally unique. It was pointed out that other NAT schemes do not have these characteristics (for instance, one proposed by Van Jacobson).

NAT Benefits

Some of the potential benefits of NAT discussed during the meeting were:

1. Make number administration of IP addresses generally easier by limiting that administration to border routers and DNS, particularly the renumbering of IP domains.
2. Using NAT to aid in address re-use by allowing a small number of hosts inside a domain, which have re-used addresses, to be able to talk outside through NAT.
3. Learn more about address translation in general so that we can better do translation for IPng (or, so that we can decide not to try translation for IPng).

There was some opinion that benefit 2 could much better be accomplished by simply giving the hosts that can talk outside multiple addresses: a re-used one for intra-domain use and a globally unique one for inter-domain use. There was some opinion that application level gateways might be a better approach in general.

NAT Problems

A number of NAT problems were discussed. Some were already known and described in Kjeld's talk. For instance, it is necessary for the router to have to dig into application headers to modify carriage of IP addresses. In the case of FTP, this requires that packet lengths be changed, and that sequence numbers in all subsequent TCP packets be changed. This is a heavy processing burden on routers, and requires router state, with the resulting scaling and reliability problems.

Any encryption of higher layer protocols that rely on IP information, such as TCP and FTP, will break with NAT. This also breaks Kerberos authentication. Any application that depends on carriage of an IP address that NAT does not account for will break with NAT. There does not exist a complete list of what applications those are, but it is clear that a number of things do work with NAT, such as telnet and mail.

It was mentioned that RFC 1006 applications break with NAT, but it is not clear why and the reasons were not discussed.

Conclusion

It was generally felt that it would be useful to the IP community to have more knowledge of the pitfalls of NAT. This is particularly true because anybody can install a NAT box independent of anybody else, and in the absence of any NAT standard.

Paul Francis was given an action item to find the list of applications that work over NAT that was generated when he experimented with NAT a couple of years ago. It was decided that there should be experimentation with NAT, with a goal of producing a document describing completely the characteristics of NAT. Kjeld was given the action item of coordinating these experiments. Nobody felt a need to follow up this BOF near-term with another meeting. It might be useful to meet once again after results are obtained, but this was left open until that time.

Attendees

| | |
|---------------------|-------------------------|
| Nick Alfano | alfano@mpr.ca |
| Frederik Andersen | fha@dde.dk |
| Robert Blokzijl | K13@nikhef.nl |
| Jim Bound | bound@zk3.dec.com |
| Ross Callon | rcallon@wellfleet.com |
| Richard Colella | colella@nist.gov |
| Francis Dupont | francis.dupont@inria.fr |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Eric Fleischman | ericf@act.boeing.com |

| | |
|-------------------|-----------------------------------|
| Peter Ford | peter@goshawk.lanl.gov |
| Paul Francis | Francis@thumper.bellcore.com |
| Terry Gray | gray@cac.washington.edu |
| Chris Howard | chris_howard@inmarsat.org |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| Geoff Huston | g.huston@aarnet.edu.au |
| Bill Manning | bmanning@rice.edu |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Greg Minshall | minshall@wc.novell.com |
| David Piscitello | dave@mail.bellcore.com |
| Robert Reschly | reschly@brl.mil |
| Luc Rooijackers | lwj@cs.kun.nl |
| Keith Sklower | sklower@cs.berkeley.edu |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Richard Thomas | rjthomas@bnr.ca |
| Claudio Topolcic | topolcic@cnri.reston.va.us |
| Noriko Yokokawa | norinori@wide.ad.jp |

2.2.1 Dynamic Host Configuration (DHC)

Charter

Chair(s):

Ralph Droms, droms@bucknell.edu

Mailing Lists:

General Discussion: host-conf@sol.bucknell.edu

To Subscribe: host-conf-request@sol.bucknell.edu

Archive: sol.bucknell.edu:~/dhcwg

Description of Working Group:

The purpose of this working group is to investigate network configuration and reconfiguration management, and determine those configuration functions that can be automated, such as Internet address assignment, gateway discovery and resource location, and those which cannot be automated (i.e., those that must be managed by network administrators).

Goals and Milestones:

- Done Write a BOOTP extensions document.
- Done Identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to: exchange packets with other hosts, obtain packet routing information, access the Domain Name System, and access other local and remote services.
- Done Summarize those mechanisms already in place for managing the information identified by objective 1.
- Done Suggest new mechanisms to manage the information identified by objective 1.
- Done Having established what information and mechanisms are required for host operation, examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.

Internet-Drafts:

“Clarifications and Extensions for the Bootstrap Protocol”, 05/03/1991, Walt Wimer <draft-ietf-dhc-bootp-02.txt>

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

“DHCP Options and BOOTP Vendor Extensions”, 06/30/1992, S. Alexander, R. Droms <draft-ietf-dhc-options-04.txt>

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

2.2.2 IP Over AppleTalk (APPLEIP)

Charter

Chair(s):

John Veizades, veizades@wco.ftp.com

Mailing Lists:

General Discussion: apple-ip@apple.com

To Subscribe: apple-ip-request@apple.com

Archive:

Description of Working Group:

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

Goals and Milestones:

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

Internet-Drafts:

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

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

Request For Comments:

RFC 1243 “AppleTalk Management Information Base”

2.2.3 IP Over Asynchronous Transfer Mode (ATM)

Charter

Chair(s):

Mark Laubach, laubach@hpl.hp.com

Mailing Lists:

General Discussion: atm@sun.com

To Subscribe: atm-request@sun.com

Archive: Send message to atm-request@sun.com

Description of Working Group:

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

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

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

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

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

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

Goals and Milestones:

- Done First Meeting. Establish detailed goals and milestones for Working Group.
- Done Post an Internet-Draft for a mechanism for IP over ATM. (Multi-Protocol Interconnect over ATM AAL5)
- Done Submit the Multi-Protocol Interconnect over ATM AAL5 to the IESG as a Proposed Standard.
- Mar 1993 Post Internet-Draft for "Internet Requirements for ATM Signaling."
- Jul 1993 Submit "Internet Requirements for ATM Signaling" to the IESG for consideration as an Informational Document.

Internet-Drafts:

- "Partial Address Resolution in ATM Networks", 03/03/1993, S. Subramaniam <draft-ietf-atm-address-resolve-00.txt>
- "IP over ATM : architecture, address translation, and call control", 03/22/1993, F. Liaw <draft-ietf-atm-address-translation-00.txt>
- "Default IP MTU for use over ATM AAL5 Services", 06/11/1993, R. Atkinson <draft-ietf-atm-mtu-01.txt>
- "Classical IP and ARP over ATM", 06/14/1993, M. Laubach <draft-ietf-atm-classic-ip-02.txt>

Request For Comments:

- RFC 1483 "Multiprotocol Encapsulation over ATM Adaptation Layer 5"

CURRENT MEETING REPORT

Reported by Mark Laubach/Hewlett-Packard

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

Monday

The first session opened with a formal announcement by Robert Hinden that he has stepped down as the ATM Working Group chair and that Mark Laubach has assumed the responsibility.

The agenda was presented and approved.

A review of recent ATM Forum activities was presented by Steve Willis. He reported that the User Network Interface (UNI) Specification Version 3.0 document is expected to be ratified in August.

An overview of the European ATM pilot project was presented by Juha Heinanen.

The topic of "routing IP over the switched virtual cloud" was presented by Joel Halpern, and he volunteered to write a proposal. Consensus is that the ATM Working Group will host the proposal, but actual work will be moved to another group that deals with routing over large public networks.

A general discussion was held to collect comments on Randall Atkinson's Internet-Draft, "Default IP MTU for use over ATM AAL5 Services." The author was not in attendance.

The last order of business was discussion of Mark Laubach's "Classical IP and ARP over ATM" Internet-Draft (henceforth called "Classical"). Discussion and consensus building continued over the next two meetings.

Tuesday

The second session opened with discussion of a timetable of ATM activities for the rest of 1993.

Both the Bellcore and Naval Research Laboratory (NRL) reference signaling codes will become available in late August or early September. Both implementations will be ATM Forum UNI 3.0 compliant, with the exception of point-to-multipoint.

An IP over UNI 3.0 document is expected to be completed and have implementation experience by the November IETF meeting.

The rest of the session was spent on discussion of Classical. During the discussion, the Internet Area Director, Stev Knowles, made it perfectly clear that Classical was not complete until ARP and IP multicast were fully addressed. (The position that area directors may delay an Internet-Draft from being submitted into the standards process was supported by the IAB in an open meeting later that evening.) Document review continued with a renewed sense of focus. LLC/SNAP was adopted by consensus as the default (the minimum required that implementors must support) IP encapsulation method. The IP MTU default size of 9180 octets was also adopted by consensus.

Wednesday

The last session opened with congratulations to Juha Heinanen for the publication of RFC 1483, "Multiprotocol Encapsulation over ATM Adaptation Layer 5."

Work then continued on Classical with the discussion of PVC support. A section on PVC support was generated for the document by an ad hoc team, and the text was approved by the group. An edited version of the text will be included in the document.

Further discussion on Classical took place following a presentation by Mark Laubach on a solution for ARP using an APR server. The group eventually reached consensus on the solution. Mark also presented solutions for the treatment of IP broadcast and IP multicast in ATM. These were also approved.

Having reached consensus on all issues, discussion on Classical was closed. Mark will produce a rewrite within the next two weeks.

Juha Heinanen led a discussion on his "NBMA Address Resolution Protocol (NBMA ARP)" Internet-Draft. Much discussion was generated on this topic, but unfortunately not enough time was available to conclude all issues. Juha will meet with others in the working group to resolve outstanding issues.

Editor's Note: Detailed summaries of discussions and consensus decisions by the working group are available via FTP or mail server from the remote directories as /ietf/atm/atm-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Attendees

| | |
|------------------|----------------------------|
| George Abe | abe@infonet.com |
| Roland Acra | acra@cisco.com |
| Masuma Ahmed | mxasabre@bellcore.com |
| Kannan Alagappan | kannan@DSMAIL.ENET.DEC.COM |
| Arun Arunkumar | nak@3com.com |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |

| | |
|-------------------|-----------------------------------|
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Lou Berger | lberger@bbn.com |
| Vincent Berkhout | berkhout@cs.utwente.nl |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Michael Brescia | |
| Caralyn Brown | cbrown@wellfleet.com |
| Tracy Brown | tacox@mail.bellcore.com |
| Theodore Brunner | tob@thumper.bellcore.com |
| Steve Buchko | stevebu@newbridge.com |
| John Burnett | jlb@adaptive.com |
| Ramon Caceres | ramon@mitl.research.panasonic.com |
| Brian Carpenter | brian@dxcern.cern.ch |
| Les Clyne | l.clyne@jnt.ac.uk |
| Jonathan Davar | jdavar@synoptics.com |
| Kurt Dobbins | dobbins@ctron.com |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Tom Easterday | tom@cic.net |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Robert Enger | enger@reston.ans.net |
| Julio Escobar | jescobar@bbn.com |
| Mark Fedor | fedor@psi.com |
| Dennis Ferguson | dennis@ans.net |
| James Forster | forster@cisco.com |
| Osten Franberg | euaokf@eua.ericsson.se |
| David Fresquez | fresquez@vnet.ibm.com |
| Dan Frommer | dan@jeremy.enet.dec.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Eugene Geer | ewg@cc.bellcore.com |
| David Ginsburg | ginsb@us-es.sel.de |
| Mike Goguen | goguen@synoptics.com |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Marcel Graf | graf%dhdibm1.bitnet@vm.gmd.de |
| Ron Greve | rgreve@cs.utwente.nl |
| Joel Halpern | jmh@network.com |
| Patrick Hanel | hanel@yoyodyne.trs.ntc.nokia.com |
| Ken Hayward | crm57d@bnr.ca |
| Geert Heijenk | heijenk@cs.utwente.nl |
| Juha Heinanen | juha.heinanen@datanet.tele.fi |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Jeff Hughes | jeff@col.hp.com |
| Sascha Ignjatovic | sascha@veda.co.at |
| Phil Irely | pirey@relay.nswc.navy.mil |
| Ronald Jacoby | rj@sgi.com |
| David Johnson | dbj@cs.cmu.edu |
| John Johnston | john@berlioz.nsc.com |

| | |
|--------------------|--------------------------------|
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Lothar Klein | lothar.klein@gmd.de |
| Mark Laubach | laubach@hpl.hp.com |
| Mark Lewis | Mark.S.Lewis@telebit.com |
| Carl Madison | carl@startek.com |
| Andrew Malis | malis_a@timeplex.com |
| Allison Mankin | mankin@cmf.nrl.navy.mil |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Keith McCloghrie | kzm@hls.com |
| Donald Merritt | don@arl.army.mil |
| Topi Miettinen | tm86214@cs.tut.fi |
| William Miskovetz | misko@cisco.com |
| Daniel Myers | dan@nsd.3com.com |
| David O'Leary | doleary@cisco.com |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Zbigniew Opalka | zopalka@agile.com |
| Charles Perkins | perk@watson.ibm.com |
| Drew Perkins | ddp@fore.com |
| Roy Perry | rperry@advtech.uswest.com |
| Philip Prindeville | philipp@res.enst.fr |
| J. Mark Pullen | mpullen@cs.gmu.edu |
| James Reeves | jreeves@synoptics.com |
| Tony Richards | richards@icm1.icp.net |
| Benny Rodrig | brodrig@rnd-gate.rad.co.il |
| Hal Sandick | sandick@vnet.ibm.com |
| Tim Seaver | tas@concert.net |
| Henk Sennema | sennema@sara.nl |
| W. David Sincoskie | sincos@thumper.bellcore.com |
| Timon Sloane | timon@timon.com |
| Kenneth Smith | kensmith@bnr.ca |
| Michael St. Johns | stjohns@darpa.mil |
| Antoine Trannoy | trannoy@crs4.it |
| Catherine Treca | Catherine.Treca@dione.urec.fr |
| Hisao Uose | uose@tnlab.ntt.jp |
| Dono van-Mierop | dono_van_mierop@3mail.3com.com |
| Werner Vogels | werner@inesc.pt |
| Scott Wasson | sgwasson@eng.xyplex.com |
| James Watt | james@newbridge.com |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |
| Marcel Wiget | wiget@switch.ch |
| Kirk Williams | kirk@sbctri.sbc.com |
| Steven Willis | steve@wellfleet.com |
| Rachel Willmer | rachelw@spider.co.uk |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

2.2.4 IP Over Large Public Data Networks (IPLPDN)

Charter

Chair(s):

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

Mailing Lists:

General Discussion: iplpdn@cnri.reston.va.us

To Subscribe: iplpdn-request@cnri.reston.va.us

Archive: ietf.cnri.reston.va.us:~/ietf-mail-archive/iplpdn/*

Description of Working Group:

The IP over Large Public Data Networks Working Group will specify the operation of the TCP/IP protocol suite over Public Data Networks (PDNs) such as SMDS, ISDN, X.25 PDNs, and Frame Relay. The working group will develop and define algorithms for the resolution of IP addresses and for the routing of IP datagrams over large, potentially global, public data networks.

The IP over SMDS Working Group has defined the operation of the Internet protocols when SMDS is used to support relatively small virtual private networks, or Logical IP Subnets (LISs). Issues arising from public and global connectivity were delegated to the IPLPDN Working Group.

The IPLPDN Working Group will also continue the work of the Private Data Network Routing Working Group (PDNROUT) on X.25 PDNs. This work will be extended to include call management and the use of the ISDN B channels for the transport of IP datagrams.

Address resolution and routing over Frame Relay will also be discussed.

Goals and Milestones:

- | | |
|------|--|
| TBD | Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents. |
| TBD | Routing of IP datagrams across very large public data networks such as SMDS and Frame Relay. |
| Done | Establish priorities and dates of completion for documents. |

Internet-Drafts:

“Determination of Encapsulation of Multi-protocol Datagrams in Circuit-switched Environments”, 02/17/1993, K. Sklower <draft-ietf-iplpdn-multi-isdn-02.txt>

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

“Management Information Base for Frame Relay DTEs”, 03/24/1993, C. Brown, F. Baker, C. Carvalho <draft-ietf-iplpdm-frmib-dte-00.txt>

“A Multilink Protocol for Synchronizing the Transmission of Multi-protocol Datagrams.”, 07/08/1993, K. Sklower <draft-ietf-iplpdm-simple-multi-01.txt>

Request For Comments:

RFC 1293 “Inverse Address Resolution Protocol”

RFC 1294 “Multiprotocol Interconnect over Frame Relay”

RFC 1315 “Management Information Base for Frame Relay DTEs”

RFC 1356 “Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode”

RFC 1433 “Directed ARP”

RFC 1490 “Multiprotocol Interconnect over Frame Relay”

CURRENT MEETING REPORT**Minutes of the IP Over Large Public Data Networks Working Group (IPLPDN)**

Report not submitted. Please refer to the Internet Area Report for a summary. The minutes of the joint PPPEXT/IPLPDN session follow the PPPEXT minutes. The attendee list below is from the joint session.

Attendees

| | |
|---------------------|--|
| George Abe | abe@infonet.com |
| Nick Alfano | alfano@mpr.ca |
| Bernt Allonen | bal@tip.net |
| Frederik Andersen | fha@dde.dk |
| Arun Arunkumar | nak@3com.com |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Fred Baker | fbaker@acc.com |
| Per Bilse | bilse@ic.dk |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Rich Bowen | rkb@ralvm11.vnet.ibm.com |
| Robert Braden | braden@isi.edu |
| Caralyn Brown | cbrown@wellfleet.com |
| Steve Buchko | stevebu@newbridge.com |
| Les Clyne | l.clyne@jnt.ac.uk |
| Thomas Cordetti | tomc@digibd.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Marty Del Vecchio | marty@shiva.com |
| Bob Downs | bdowns@combinet.com |
| Ian Duncan | id@cc.mcgill.ca |
| Toerless Eckert | Toerless.Eckert@informatik.uni-erlangen.de |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Eugene Geer | ewg@cc.bellcore.com |
| David Ginsburg | ginsb@us-es.sel.de |
| Marcel Graf | graf@dhdibm1.bitnet@vm.gmd.de |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Joel Halpern | jmh@network.com |
| Jari Hamalainen | jah@rcrtre.nokia.com |
| Patrick Hanel | hanel@yoyodyne.trs.ntc.nokia.com |
| Ken Hayward | crm57d@bnr.ca |
| Gerd Holzhauer | Holzhauer1@applelink.apple.com |
| John Hopkins | J_Hopkins@icrf.icnet.uk |

| | |
|-----------------------|-------------------------------|
| Chris Howard | chris_howard@inmarsat.org |
| David Jacobson | dnjake@vnet.ibm.com |
| Philip Jones | p.jones@jnt.ac.uk |
| Frank Kastenholz | kasten@ftp.com |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Dave Langley | davel@hprnd.hp.com |
| Eliot Lear | lear@sgi.com |
| Paolo Malara | malara@crs4.it |
| Andrew Malis | malis@maelstrom.timeplex.com |
| Shehzad Merchant | merchant@erg.sri.com |
| Gerry Meyer | gerry@spider.co.uk |
| William Miskovetz | misko@cisco.com |
| Keith Mitchell | keith@pipex.net |
| Daniel Myers | dan@nsd.3com.com |
| Drew Perkins | ddp@fore.com |
| David Piscitello | dave@mail.bellcore.com |
| Lars Poulsen | lars@cmc.com |
| Dave Rand | dave_rand@novell.com |
| Juergen Rauschenbach | jrau@dfn.de |
| Tony Richards | richards@icm1.icp.net |
| Benny Rodrig | brodrig@rnd-gate.rad.co.il |
| Miguel Sanz | miguel.sanz@rediris.es |
| Henk Sennema | sennema@sara.nl |
| Keith Sklower | sklower@cs.berkeley.edu |
| Timon Sloane | timon@timon.com |
| Kenneth Smith | kensmith@bnr.ca |
| Henk Steenman | henk@sara.nl |
| Richard Sweatt | rsweatt@synoptics.com |
| Antoine Trannoy | trannoy@crs4.it |
| Hisao Uose | uose@tnlab.ntt.jp |
| Rene van der Hauw | rene@geveke.nl |
| Willem van der Scheun | scheun@sara.nl |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Scott Wasson | sgwasson@eng.xyplex.com |
| Kirk Williams | kirk@sbctri.sbc.com |
| Rachel Willmer | rachelw@spider.co.uk |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

2.2.5 P. Internet Protocol (PIP)

Charter

Chair(s):

Paul Francis, Francis@thumper.bellcore.com

Mailing Lists:

General Discussion: pip@thumper.bellcore.com

To Subscribe: pip-request@thumper.bellcore.com

Archive: thumper.bellcore.com:~/pub/tsuchiya/pip-archive

Description of Working Group:

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

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

While the PIP overview demonstrates a number of powerful mechanisms, much work remains to be done to bring PIP to a full specification. This work includes, but is not limited to, specifying the header format; specifying a basic set of error messages (PCMP messages); specifying the PIP forwarding rules; specifying host interface messages (particularly the directory service query response); specifying rules for host PIP header construction; specifying modifications to existing protocols for use with PIP (BGP-4, OSPF, ARP, DNS, etc.); specifying PIP MTU discovery techniques; and specifying a transition strategy for PIP.

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

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

Goals and Milestones:

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

- Done Post as an Internet-Draft a description of the PIP Packet Format and Forwarding Engine, the PIP Control Message Protocol (PCMP), the PIP Host Interface Message Protocol, and the PIP MTU Discovery Protocol.
- Oct 1992 Post as an Internet-Draft a description of the modifications to BGP-4 for PIP, the Modifications to OSPF for PIP, and the modifications to ARP for PIP.
- Done Presentation and review of the PIP specification by the IESG. If acceptable, the first Working Group meeting will be held.
- Done Post as an Internet-Draft the modifications to DNS for PIP, the Address assignment in PIP, and the PIP transition strategy.

Internet-Drafts:

“Pip Header Processing”, 10/30/1992, P. Francis <draft-ietf-pip-processing-02.txt>

“Pip Identifiers”, 11/03/1992, P. Francis <draft-ietf-pip-identifiers-02.txt>

“Use of DNS with Pip”, 01/29/1993, P. Francis, S. Thomson <draft-ietf-pip-dns-01.txt>

“Pip Near-term Architecture”, 02/22/1993, P. Francis <draft-ietf-pip-architecture-01.txt>

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

“Pip Address Conventions”, 06/11/1993, P. Francis <draft-ietf-pip-address-conv-00.txt>

“Pip Host Operation”, 06/11/1993, P. Francis <draft-ietf-pip-host-operation-00.txt>

“PCMP: Pip Control Message Protocol”, 06/11/1993, P. Francis <draft-ietf-pip-control-msg-00.txt>

“IP Independent Transition (IPIT) for Pip”, 07/06/1993, P. Francis <draft-ietf-pip-ipit-transition-00.txt>

CURRENT MEETING REPORT

Reported by Geoff Huston/Australian Academic and Research Network

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

Overview

A specification overview was presented to the attendees. The specification of forwarding has remained unchanged for the past 3 months. The DNS architecture to support PIP has been revised. The PIP identifier structure has been revised. IDRP routing support for PIP has revisions in progress. The host operations specifications has been revised. The PIP Control Message Protocol is new, and is currently incomplete. The PIP transition specification is new. Missing from the specification is a MIB definition. Routing still requires further definition.

PIP Progress

- PIP DNS

The use of the DNS as a support tool for PIP transition is still under review. The major new area of support functionality required is that of timestamped queries, as described in the PIP DNS specification. In addition, the use of the DNS in PIP transition is described in the PIP transition specification.

- PIP IDS

The hierarchical structure of PIP identifiers has been weakened, and a flat ID structure is considered sufficient while allowing simple integration of auto-configuration mechanisms. The ID structure is that of a 2-byte identifier prefix and a 6-byte static host identifier. It was noted that there were questionable returns for a richer identifier structuring. It was noted that within the current specification of PIP there was no visible requirement for reverse lookups based on PIP IDs to discover PIP addresses, on the basis that PIP IDs and PIP addresses are intended to be passed together. Further structuring of the PIP host identifiers was left as an open issue.

- PIP Routing

Routing is based on a multilevel path vector, coupled with IDRP as the routing framework. The basic algorithms for PIP routing are essentially complete, but any-cast, tunnelling and Quality of Service attributes have yet to be implemented. IDRP is used as a mechanism to support neighbour reachability and sequencing.

- PIP Transition

Evaluation of transition arrangements using IPAE and an IP Independent Transition structure have been undertaken. The meeting focussed on this topic in further detail.

- PIP Host Operations

The host will be required to perform a choice of multiple PIP addresses, within the context of two hosts performing an address choice which allows optimal end-to-end reachability. The host operations include heuristics for host address selection and the use of PCMP messages in order to instruct the host to select an alternative address.

- PCMP

Currently PCMP has support for “packet not delivered” with 12 reasons. Other PCMP types, including router discovery mechanisms, are to be specified.

IP to PIP Transition

Concerns were expressed with the IPAE approach as an answer to the transition problem. The meeting reviewed an alternative approach to transition using a translating boundary architecture, the IP Independent Transition (IPIT) approach.

In evaluating the usefulness of IPAE it was noted that the use of IP addresses within an IPng packet allowed packet header translation in the direction of IPng to IP to be relatively straightforward. The packet header translation in the direction of IP to IPng does require an inverse lookup in order to generate the IPng address from the destination IP address. The static nature of this lookup does have negative implications where support for auto-configuration and mobility is desired within the transitioning environment.

The IPIT approach uses a translational approach where the binding of an IP address to an IPng host is dynamic, and the binding is undertaken by the boundary translating router. The nature of the binding (static/dynamic reuse) is reliant of the relative size of the pool of bindable IP addresses and the number of IPng hosts. The participants noted that this approach did have application layer implications where applications included explicit description of network layer addresses. The participants also noted that there was a requirement for the host to regularly inform the translating router that the IP address is in use, and also explicitly inform the router when the address can be returned to the pool for subsequent rebinding to another IPng host. The meeting explored various scenarios of pool allocation, as they related to packet header translation. The meeting noted that various operational practices, such as support of end-to-end traceroute will imply extensive use of the pool with a requirement for careful management of binding structures of IP addresses within the IPng domain. SNMP management from the IP domain of IPng resources was also discussed, with the outcome that management within an IPng domain would be from within the domain.

The objective of IPIT is to use dynamic binding of IP addresses to IPng hosts in order to ensure that transition can use a smaller set of IP addresses than a static binding would imply. Pool size can be further reduced by using the IPng/translation IP address pair as the translation table index, allowing different IPng hosts be assigned the same translational IP address (under a set of specific conditions).

Experimentation with IPIT was proposed, on the basis that if major operational flaws were exposed through this approach, the IPAE structure could be used as a fallback.

The participants discussed the topic of whether early or late partitioning was considered desirable, and the dinosaur argument was proposed, where the view was expressed that extensive transitional structures designed to provide an unnatural extension of life for retrograde hosts were considered to be a unnatural practice.

The event sequence for the binding of an IP address to an IPng host was examined. It was noted that the use of the DNS in the process of choice of a translating router implied that initial IP address binding from the pool was performed without explicit knowledge of the IP domain end host, and that the state requirements within the translating router, coupled with the requirement for DNS sequences, did imply fate-sharing on the basis of a requirement for synchronisation of the operation of the DNS and the translating routers. The translating routers also form a critical single point of failure within the IPIT structure.

The participants also discussed the bootstrap phase for the setup of the DNS forwarding across the IP/IPng domain boundary, and it was noted that IPng DNS servers would require a permanent IP address binding which was known to all boundary routers. The role and configuration of IPng DNS servers within this context was discussed.

PIP support for provider selection as a component of the transitional environment was discussed, and the use of reversal of an IP source route was considered, with the overall conclusion that provider selection would not map across the IP/IPng boundary within the transition environment.

DNS Operations

DNS operations within the PIP environment were presented at the meeting. The DNS operation requires the introduction of a new PIP class. The PIP ID is to be stored as an A RR, and the PIP address as ADDR RRs. The function of IP inverse lookup domain is supported within the PIP DNS environment as reverse domains for ID and address to map to domain names, and a third domain to map from ID to address.

The role of the DNS within IPIT was discussed, and it was noted that there was a requirement during transition for the PIP domain to be supported within an incomplete domain space within the PIP class. This implies that recursive resolvers must determine whether NSs are defined within the PIP class, which also implies that stub resolvers within the transitional environment will be inefficient.

The inclusion of support for timestamped queries was discussed, with a motivation that PIP addresses are more likely to change in response to provider changes, and a mechanism for effectively specifying a request for more recent information from the DNS was required.

It was noted that timestamp queries are more widely applicable, and that this function is on the DNS Working Group agenda for consideration. This is documented in the pip-dns Internet-Draft.

Deployment

The parts of PIP deployment which have been completed are the host code, the forwarding engine, PIP to IP translation and IP to PIP translation, encapsulation, P-ARP and PCMP. In addition pconf has been written as a configuration generator, which takes a network specification and generates specific configuration descriptions.

An experimental deployment on the PIP Backbone on 20 hosts across the Internet has been completed.

Future plans focus on deployment across further hosts and routers.

Attendees

| | |
|--------------------|--------------------------------------|
| Nick Alfano | alfano@mpr.ca |
| Bernt Allonen | bal@tip.net |
| Frederik Andersen | fha@dde.dk |
| Per Andersson | pa@cdg.chalmers.se |
| Michael Anello | mike@xlnt.com |
| Anders Baardsgaad | anders@cc.uit.no |
| John Ballard | jballard@microsoft.com |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Per Bilse | bilse@ic.dk |
| Jim Binkley | jrb@ibeam.intel.com |
| Robert Blokzijl | K13@nikhef.nl |
| Ronald Broersma | ron@nosc.mil |
| Steve Buchko | stevebu@newbridge.com |
| John Burnett | jlb@adaptive.com |
| Ross Callon | rcallon@wellfleet.com |
| Brian Carpenter | brian@dxcern.cern.ch |
| George Clapp | clapp@ameris.center.il.ameritech.com |
| David Conrad | davidc@iij.ad.jp |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Stephen Deering | deering@parc.xerox.com |
| Kurt Dobbins | dobbins@ctron.com |

| | |
|----------------------|-----------------------------------|
| Ian Duncan | id@cc.mcgill.ca |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Dino Farinacci | dino@cisco.com |
| Stefan Fassbender | stf@easi.net |
| Dennis Ferguson | dennis@ans.net |
| Eric Fleischman | ericf@act.boeing.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Osten Franberg | euaokf@eua.ericsson.se |
| Paul Francis | Francis@thumper.bellcore.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Eugene Geer | ewg@cc.bellcore.com |
| Robert Gilligan | Bob.Gilligan@Eng.Sun.Com |
| Joseph Godsil | jgodsil@cnsa.uiuc.edu |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Joel Halpern | jmh@network.com |
| Jari Hamalainen | jah@rcrtre.nokia.com |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Chris Howard | chris_howard@inmarsat.org |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| Geoff Huston | g.huston@aarnet.edu.au |
| David Jacobson | dnjake@vnet.ibm.com |
| Cyndi Jung | cmj@3com.com |
| Scott Kaplan | scott@wco.ftp.com |
| Frank Kastenholz | kasten@ftp.com |
| Dave Katz | dkatz@cisco.com |
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Ton Koelman | koelman@stc.nato.int |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Tony Li | tli@cisco.com |
| Robin Littlefield | robin@wellfleet.com |
| Greg Minshall | minshall@wc.novell.com |
| Keith Mitchell | keith@pipex.net |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Erik Nordmark | nordmark@eng.sun.com |
| Michael O'Dell | mo@uunet.uu.net |
| Petri Ojala | ojala@eunet.fi |
| Michael Patton | map@bbn.com |
| David Piscitello | dave@mail.bellcore.com |
| Luc Rooijackers | lwj@cs.kun.nl |
| Henry Sanders | henrysa@microsoft.com |
| W. David Sincoskie | sincos@thumper.bellcore.com |
| Henk Steenman | henk@sara.nl |
| Vladimir Sukonnik | sukonnik@process.com |
| Richard Thomas | rjthomas@bnr.ca |
| Susan Thomson | set@bellcore.com |

| | |
|-----------------------|----------------------------|
| Hisao Uose | uose@tnlab.ntt.jp |
| Willem van der Scheun | scheun@sara.nl |
| Werner Vogels | werner@inesc.pt |
| Scott Wasson | sgwasson@eng.xyplex.com |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |
| Kirk Williams | kirk@sbctri.sbc.com |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Noriko Yokokawa | norinori@wide.ad.jp |
| Jessica Yu | jyy@merit.edu |
| Romeo Zwart | romeo@sara.nl |

2.2.6 Point-to-Point Protocol Extensions (PPPEXT)

Charter

Chair(s):

Fred Baker, fbaker@acc.com

Mailing Lists:

General Discussion: ietf-ppp@ucdavis.edu

To Subscribe: ietf-ppp-request@ucdavis.edu

Archive:

Description of Working Group:

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

Goals and Milestones:

None specified

Internet-Drafts:

“The PPP Internetwork Packet Exchange Control Protocol (IPXCP)”, 06/10/1992, W. Simpson <draft-ietf-pppext-ipxcp-04.txt>

“Compressing IPX Headers Over WAN Media (CIPX)”, 12/08/1992, S. Mathur, M. Lewis <draft-ietf-pppext-cipx-04.txt>

“PPP LCP Extensions”, 01/08/1993, W. Simpson <draft-ietf-pppext-lcpext-04.txt>

“PPP over ISDN”, 03/10/1993, W. Simpson <draft-ietf-pppext-isdn-02.txt>

“PPP in Frame Relay”, 03/10/1993, W. Simpson <draft-ietf-pppext-frame-relay-01.txt>

“PPP over SONET/SDH”, 03/10/1993, W. Simpson <draft-ietf-pppext-sonet-01.txt>

“PPP in X.25”, 03/10/1993, W. Simpson <draft-ietf-pppext-x25-01.txt>

“PPP in HDLC Framing”, 07/02/1993, W. Simpson <draft-ietf-pppext-hdlc-framing-02.txt>

“The Point-to-Point Protocol (PPP)”, 07/06/1993, W. Simpson <draft-ietf-pppext-lcp-main-02.txt>

“Point-to-Point Protocol Extensions for Bridging”, 07/08/1993, F. Baker, R. Bowen <draft-ietf-pppext-for-bridging-00.txt>

“A Multilink Protocol for Synchronizing the Transmission of Multi-protocol Datagrams.”, 09/02/1993, K. Sklower <draft-ietf-pppext-multilink-00.txt>

“Requirements for an Internet Standard Point-to-Point Protocol”, 09/16/1993, D. Perkins <draft-ietf-pppext-requirements-00.txt>

Request For Comments:

RFC 1220 “Point-to-Point Protocol Extensions for Bridging”

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

RFC 1332 “The PPP Internet Protocol Control Protocol (IPCP)”

RFC 1333 “PPP Link Quality Monitoring”

RFC 1334 “PPP Authentication Protocols”

RFC 1376 “The PPP DECnet Phase IV Control Protocol (DNCP)”

RFC 1377 “The PPP OSI Network Layer Control Protocol (OSINLCP)”

RFC 1378 “The PPP AppleTalk Control Protocol (ATCP)”

RFC 1471 “The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol”

RFC 1472 “The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol”

RFC 1473 “The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol”

RFC 1474 “The Definitions of Managed Objects for the Bridge Network Control Protocol of the Point-to-Point Protocol”

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

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

PPP Extensions for Bridging

Editor's Note: A list of new features and editorial changes to the document is available via FTP or mail server from the remote directories as /ietf/pppext/pppext-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Other Documents

The group discussed several other documents in addition to PPP Extensions for Bridging.

Editor's Note: A list of documents and recommendations is available via FTP or mail server from the remote directories as /ietf/pppext/pppext-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

PPP Compression

A separate breakout meeting was held for the bulk of the work, and the slides from the two presentations that were given follow these minutes. They contain a lot of information.

Algorithms Under Consideration

Five candidate protocols are under active consideration:

1. Predictor – Free, but poor compression ratio - implement with CRC
2. Gandalf FZA – \$20K without patent protection
3. V.42bis – \$20K one time
4. HP PPC – About \$20 one time with patent protection
5. STAC – \$5 per, royalty on software with patent protection, \$40 on chip

Although we wanted to, the PPPEXT Working Group does not recommend one of them for universal implementation. The reason is that the group cannot, under IETF rules and marketplace sense, require everyone to license code or silicon from a single vendor, and the one unencumbered algorithm we have found has significant (64K per link) memory requirements. We therefore only provide the means to negotiate them.

Packet format for Predictor is:

Address
Control
PPP Compression Data Protocol ID
Original Frame Length (not compressed)
Compressed Frame
Frame CRC-16 (not compressed)

The reason for the CRC-16 is to help detect frame loss (and resultant dictionary desynchronization) in the case where a reliable link is not in use.

Reliable Link Negotiation

How to implement without a reliable link: decompress. If a frame fails to correctly decompress, send a Compression Control Protocol Configure request on the link.

- Reasons *not* to use a reliable link:
 - Would like to use the same algorithm on all WAN code
 - Links are generally reliable anyway
 - Unreliable links are perceived to be simpler
- Reasons *to* use a reliable link:
 - Loss of buffers introduced problems
 - More graceful degradation in the presence of errors

LAPB Negotiation Option

LAPB will be negotiated, but the minimum configuration will not support LAPB. The LAPB LCP Negotiation Option will have the following format:

LCP Option
Length
Window

Compression Control Protocol Negotiation Option

There will be one option number per compression algorithm, with a special one for proprietary algorithms. They will be listed in the order of preference, and the sender's preferences will be respected in each direction, as the most effort is in the compression of the frame.

The general format of these is:

```
COMPRESSION CONTROL PROTOCOL Option
Length
Parameters as required by the algorithm
```

The proprietary protocol option will have the vendors IEEE 802 Organizational Unit Identifier as the first three octets of the parameter field. It is recommended that vendors use the fourth octet as a version number. This allows a vendor to use a proprietary algorithm among its own equipment without revealing its intellectual property to the IANA. Note that this option may occur more than once—a vendor may support multiple versions of its own algorithm, or may support several vendors algorithms. The procedures defined in the PPP LCP for handling multiple instances of the same option apply in this case.

Attendees

| | |
|-------------------|--------------------------------------|
| Steve Alexander | stevea@lachman.com |
| Arun Arunkumar | nak@3com.com |
| Fred Baker | fbaker@acc.com |
| Per Bilse | bilse@ic.dk |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Caralyn Brown | cbrown@wellfleet.com |
| Steve Buchko | stevebu@newbridge.com |
| George Clapp | clapp@ameris.center.il.ameritech.com |
| Thomas Cordetti | tomc@digibd.com |
| Bob Downs | bdowns@combinet.com |
| Ian Duncan | id@cc.mcgill.ca |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Jari Hamalainen | jah@rctre.nokia.com |
| Dave Langley | davel@hprnd.hp.com |
| Gerry Meyer | gerry@spider.co.uk |
| William Miskovetz | misko@cisco.com |
| Drew Perkins | ddp@fore.com |
| Lars Poulsen | lars@cmc.com |
| Dave Rand | dave_rand@novell.com |
| Benny Rodrig | brodrig@rnd-gate.rad.co.il |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| Keith Sklower | sklower@cs.berkeley.edu |
| Scott Wasson | sgwasson@eng.xyplex.com |
| James Watt | james@newbridge.com |

PPP Data Compression

Dave Rand - Novell, Inc.

- dave_rand@Novell.com
- dir@bungie.com

ppp-comp mailing list archives are available at sgi.com/other/ppp-comp.

Requests for addition/deletion to the mailing list should go to ppp-comp-request@bungie.com.

Why use data compression?

- Better utilization of link
- Faster response time
- Higher reliability
- Cost savings
- Customers want it

When NOT to use data compression

- Encrypted data
- Already-compressed data
- Highly variable-latency-sensitive protocols
- When compression time exceeds transport time

Which compression method is best?

- Depends on link speed
- Depends on processor/compression speed
- Depends on type of traffic
- Depends on amount of memory available
- Depends on amount of money available (license costs)
- Who knows?!!

Algorithm Selection

- Traffic is not constant
- Traffic is not one file
- Traffic is not one protocol
- Individual files don't make a good test case
- Every case is different - real traffic is too hard to compare different algorithms against.

My solution: simulated traffic.

So - which algorithm is best?

- There is no "BEST" algorithm.
- Time
- Space
- Compression ratio
- RGF
- Reliable data links

Candidates

- Gandalf's FZA algorithm is a good choice for low-to-medium speed links. dcarr@gandalf.ca has pricing information.
- V.42bis has stable characteristics, licensing is reasonable (about \$20,000 USD). Other LZW variants, such as UNIX Compress may be a good choice.
- Prediction coding is free, and simple to implement. Compression ratio is poor.
- HP's PPC follows the PPP spirit much better, in that no reliable link is required. davel@hprnd.rose.hp.com has pricing information.
- Others?

Data Link Issues

- Almost all compression algorithms must have sequential, ordered delivery of data, in order to keep the compressor and decompressor in sync.
 - HP's PPC DOES NOT require a reliable link.
 - LAPB is a good choice, and is not too complex to implement.
- Negotiation of a reliable link may occur before or after negotiation of compression. If the algorithm selected requires a reliable link, and one is not negotiated, LCP must re-negotiate compression off prior to opening the link.

Negotiation of Compression

- Negotiation of compression will occur via a standardized mechanism.

| | | | | |
|---------|------|----------------|----------------|-----|
| LCPtype | Size | Preferred type | Supported type | ... |
|---------|------|----------------|----------------|-----|

Compression types will be ordered as to the sender's preference. Each compression algorithm will be assigned a different number. Each side offers all compression types it considers appropriate for the current link. If none match, no compression is performed.

Each compression type will use a single, generic LCP option for negotiation of algorithm-specific parameters. Algorithm specific parameters may be negotiated only after an algorithm is selected.

| | | |
|---------|------|---|
| LCPtype | Size | Free format/defined by compression algorithm. |
|---------|------|---|

Reliable link negotiation

| | | |
|---------|------|-------------|
| LCPtype | Size | Window size |
|---------|------|-------------|

This indicates the number of mru-sized blocks of data we can have outstanding on the link. If the window size is less than or equal to 7, LAPB modulo 8 should be chosen. If the window size is 8 or larger, LAPB modulo 128 should be chosen.

Data format

| | | | |
|-------------|------------------|-----------------|------|
| LAPB header | packet size high | packet size low | data |
|-------------|------------------|-----------------|------|

This format allows the link layer to alter the physical size of the frame to fit the underlying LAPB transport. We will only utilize the number of bytes specified by the two bytes preceding the data block.

Recovery

- No one is perfect.
- No algorithm is perfect.
- A method of detecting errors on a compressed data link is essential.

| |
|-------------------------|
| LAPB header |
| Packet size |
| Frame size |
| Frame data (compressed) |
| Frame CRC |
| Frame size... |

With the original CRC of the frame available, we have a second level of detecting errors in the compressed link.

Packet by Packet Compression (PPC) for Networking

For IETF, July 1993
Dave Langley, Hewlett Packard

Roseville Networks Division



Agenda


- Introduction to HP compression architecture
- HP PPC
- PPC compression performance
- Combining PPC with header compression
- H/W implementations
- Patents
- Licensing
- Summary

HP's Involvement in Compression

- HP has been working on compression technology since 1984
- HP developed its first compression ASIC IC in 1987
- HP is a leader in compression technology for DAT drives with both algorithms and H/W ASIC's
- HP is developing advanced video compression algorithms for multimedia
- HP maintains compression expertise in many fields across several divisions (Palo Alto Labs, Bristol Labs, ASIC divisions, product divisions)
- HP employs Abraham Lempel (of Lempel-Ziv)

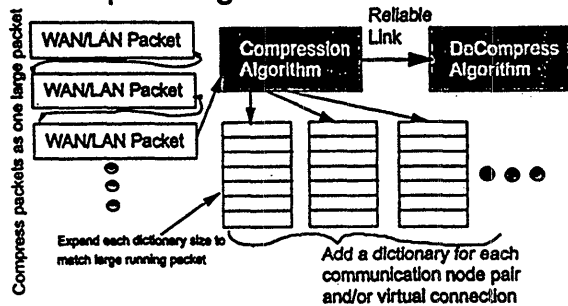
Compression & Networks

Why is LAN/WAN Compression difficult?



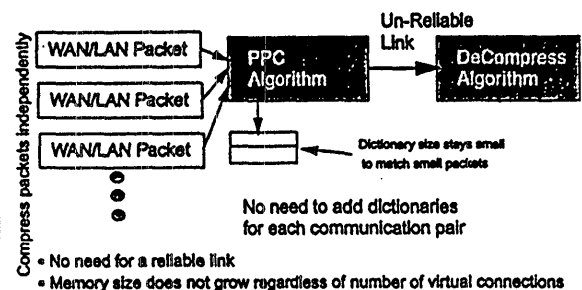
- Networking splits data into small packets
- Intermixed node traffic yields non-homogeneous data
- Standard compression performance is low on packets

Compressing small WAN Packets



- Traditional approach is to maintain a running dictionary
- Expensive as running packets, node pairs, and connections grow
- Must have a reliable link


HP Packet by Packet (PPC) Alg.



- No need for a reliable link
- Memory size does not grow regardless of number of virtual connections
- Intermixed communications node pairs have no effect on compression
- Any protocol (X.25, Frame Relay, ATM, etc.) can utilize this architecture
- Transparent to rest of network system

Dictionary cost comparison

Memory Size vs. Number of Virtual Connections
Memory Size, Kbytes




HP Packet by Packet, PPC, compression concept

- A low cost method to implement compression
- A compression system requiring low memory regardless of intermixed traffic type or number virtual connections
- A compression method that can work in a wide variety of network systems (ATM, X.25, ISDN, LAN's etc)
- A compression scheme that does not require a reliable link
- A packet by packet system that works well with header compression
- In 1991 began working on effective compression of small packets

Small LAN Packet compression

Compare running dictionary with HP pkt by pkt



Header Requirements

- Need appropriate headers on the protocols to handle the compression (max 1 byte)
- Header information is minimal
 - compressed yes/no (for expansion protection)
 - type of compression used on this packet
 - future needs

H/W ASIC implementation

- HP is investigating ASIC to implement this compression technology
- Contacting multiple vendors for ASIC sourcing
- Several possible implementations
 - Very low latency "on the fly" compression
 - Look aside uP bus type
 - Pass through compression

Patents

- HP has patents pending on this PPC technology
- PPC is built upon LZ2 and therefore has patent issues
- HP has worked through LZ compression patents before on previous standards and is doing so again for PPC
- HP is negotiating with holders of patents in the appropriate areas
- HP expects to obtain a reasonably priced agreement to utilize PPC in WAN products

Licensing

- HP is negotiating with patent holders to obtain "one-stop" licensing service
- Purchase of PPC license will cover outstanding patent licensing as well as licensing for HP PPC patent licensing (HP PPC charges minimal)
- Licensing is intended to be negotiated for either a one time up front OR a per unit royalty
- A license will provide right to use in WAN applications and C source code
- Use of compression ASIC would be independent of S/W agreement (all licensing would be included in ASIC price)
- HP has licensed compression patents before

Summary

- PPC offers implementation simplicity with good performance
- PPC is a low cost method to implement compression
- PPC requires small memory regardless of intermixed traffic type or number of virtual connections
- PPC combined with header compression provides even better compression on small packets
- PPC has applications in many WAN protocols where a running dictionary would be difficult/impossible
- HP intends to go through all the patent/licensing issues to offer an easy license service
- HP intends to facilitate the offering of a PPC ASIC from multiple sources

CURRENT MEETING REPORT

Reported by Fred Baker/ACC and George Clapp/Ameritech

Minutes of the joint session of IPLPDN and PPPEXT Working Groups

RFC 1356 X.25

RFC 1356 will be recommended as a Draft Standard. There have been six to seven implementations with no interoperability problems.

RFC 1294 has already been recommended for advancement to Draft Standard.

Protocol Discrimination

A PPP NLPID has been requested by the PPPEXT Working Group for use in NLPID-encapsulated protocols. The request has unfortunately gotten lost in the mail. Bill Simpson will resend the request to Lyman Chapin, who has agreed to make it happen. There is a separate issue with the ISDN Lower Layer Compatibility Information Element; George Clapp will pursue obtaining a value indicating PPP.

- **IP/Circuit Switched Service**

The question was seriously discussed whether we in fact need a default way to send IP over circuit switched services such as ISDN B channel. It was observed that the question is malformed; we do not need a default way to send IP over a V.35 or V.11 interface, for example. We need a way to speak to a peer system at the data link layer, which might be a Frame Relay or X.25 switch, or a peer host or router.

We already have standards for PPP, Frame Relay, and X.25. In different contexts, we are willing to run any of the three standards.

This approach is recommended for circuit switched services:

- Systems must implement PPP, on the assumption that circuit switched communications are generally [host or router] to [host or router].
- Systems may implement other protocols such as Frame Relay or X.25

The implication here is not that all calls will be initiated with PPP signaling and encapsulation, but that PPP signaling and encapsulation will be a universally implemented option.

Multi-link Protocol

The header will be changed to one of the following:

```
+-----+
|M|P|0|0| Sequence Number   |
+-----+
```

- M - More - 1 if a non-terminal fragment, 0 if the last fragment
- P - Phase - has the same value on each fragment of a message, inverts from message to message
- 0 - Reserved, must be zero
- Sequence Number - 0 to 4095 fragment sequence number

```
+-----+
|F|L|0|0| Sequence Number   |
+-----+
```

- F - First - 1 if first fragment in a message
- L - Last - 1 if last fragment in a message
- 0 - Reserved, must be zero
- Sequence Number - 0 to 4095 fragment sequence number

Including a link in the multi-link group is done by authenticating inclusion in the multi-link group and negotiation of the Fragmentation Protocol Control Protocol (FPCP).

Removing a link from the multi-link group is done by terminating the FPCP on that link.

In the worst case, receiver recovery from a sequence error (fragment loss) is done by sending an FPCP Configure Request in the OPEN state on all links; in most cases, one of the following two conditions is sufficient to detect and step past the loss of a sequenced fragment:

1. Receipt of a frame on each link with a successor to the omitted sequence number.
2. Expiration of an implementation-specific receipt timer; this should be long enough to handle the relevant timing issues.

There is a separate LCP negotiation, authentication step, and set of Control Protocol negotiations for each link in a multi-link group.

Several other options were considered, including the use of the RFC 1294 fragmentation header, which was agreed to in the March meeting; RFC 1294 provides the same essential features as this but requires four octets, and additionally provides only compatibility with RFC 1294.

PPP on Frame Relay

We need to have an Applicability Statement for PPP over Frame Relay, in view of the existence of RFC 1294. The default encapsulation is as described in the minutes of the March IETF. Various edits were recommended, which will be included in an updated draft, including collapsing of Keith Sklower's parameter negotiation document (with attribution as author) into this document.

LQM should not be used on a Frame Relay DLCI.

PPP on X.25

We need to have an Applicability Statement for PPP over X.25 in view of RFCs 877 and 1356. Various edits were recommended, which will be included in an updated draft. Primary attention should be given to reducing the size of the X.25 frame.

LQM should not be used in this environment.

The PPP NLPID SHOULD be placed in the call user data rather than being carried in each frame.

PPP/ISDN

Bill Simpson presented his paper on PPP over ISDN.

PPP must have the same default MRU (and any other defaults) on ISDN as in other environments. Keith Sklower will publish his IPLPDN document, "Determination of Encapsulation of Multi-Protocol Datagrams in Circuit Switched Environment," and Bill indicates that he would like to copy some of the technical material from it into this document. It was decided that he would reference Keith's document.

Parameter Negotiation

Keith and Bill will merge their documents. This document should be separate from the PPP over *foo* documents, as it is desired to be placed on the standards track, and the PPP over *foo* documents may not be placed on that track.

Attendees

| | |
|---------------------|--|
| George Abe | abe@infonet.com |
| Nick Alfano | alfano@mpr.ca |
| Bernt Allonen | bal@tip.net |
| Frederik Andersen | fha@dde.dk |
| Arun Arunkumar | nak@3com.com |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Fred Baker | fbaker@acc.com |
| Per Bilse | bilse@ic.dk |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Rich Bowen | rkb@ralvm11.vnet.ibm.com |
| Robert Braden | braden@isi.edu |
| Caralyn Brown | cbrown@wellfleet.com |
| Steve Buchko | stevebu@newbridge.com |
| Les Clyne | l.clyne@jnt.ac.uk |
| Thomas Cordetti | tomc@digibd.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Marty Del Vecchio | marty@shiva.com |
| Bob Downs | bdowns@combinet.com |
| Ian Duncan | id@cc.mcgill.ca |
| Toerless Eckert | Toerless.Eckert@informatik.uni-erlangen.de |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Eugene Geer | ewg@cc.bellcore.com |
| David Ginsburg | ginsb@us-es.sel.de |
| Marcel Graf | graf@dhdibm1.bitnet@vm.gmd.de |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Joel Halpern | jmh@network.com |
| Jari Hamalainen | jah@rctre.nokia.com |
| Patrick Hanel | hanel@yoyodyne.trs.ntc.nokia.com |
| Ken Hayward | crm57d@bnr.ca |
| Gerd Holzhauser | Holzhauser1@applelink.apple.com |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Chris Howard | chris_howard@inmarsat.org |
| David Jacobson | dnjake@vnet.ibm.com |
| Philip Jones | p.jones@jnt.ac.uk |
| Frank Kastenholz | kasten@ftp.com |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Dave Langley | davel@hprnd.hp.com |
| Eliot Lear | lear@sgi.com |
| Paolo Malara | malara@crs4.it |
| Andrew Malis | malis_a@timeplex.com |
| Shehzad Merchant | merchant@erg.sri.com |

| | |
|-----------------------|--|
| Gerry Meyer | <code>gerry@spider.co.uk</code> |
| William Miskovetz | <code>misko@cisco.com</code> |
| Keith Mitchell | <code>keith@pipex.net</code> |
| Daniel Myers | <code>dan@nsd.3com.com</code> |
| Drew Perkins | <code>ddp@fore.com</code> |
| David Piscitello | <code>dave@mail.bellcore.com</code> |
| Lars Poulsen | <code>lars@cmc.com</code> |
| Dave Rand | <code>dave_rand@novell.com</code> |
| Juergen Rauschenbach | <code>jrau@dfn.de</code> |
| Tony Richards | <code>richards@icm1.icp.net</code> |
| Benny Rodrig | <code>brodrig@rnd-gate.rad.co.il</code> |
| Miguel Sanz | <code>miguel.sanz@rediris.es</code> |
| Henk Sennema | <code>sennema@sara.nl</code> |
| Keith Sklower | <code>sklower@cs.berkeley.edu</code> |
| Timon Sloane | <code>timon@timon.com</code> |
| Kenneth Smith | <code>kensmith@bnr.ca</code> |
| Henk Steenman | <code>henk@sara.nl</code> |
| Richard Sweatt | <code>rsweatt@synoptics.com</code> |
| Antoine Trannoy | <code>trannoy@crs4.it</code> |
| Hisao Uose | <code>uose@tnlab.ntt.jp</code> |
| Rene van der Hauw | <code>rene@geveke.nl</code> |
| Willem van der Scheun | <code>scheun@sara.nl</code> |
| Ruediger Volk | <code>rv@informatik.uni-dortmund.de</code> |
| Scott Wasson | <code>sgwasson@eng.xyplex.com</code> |
| Kirk Williams | <code>kirk@sbctri.sbc.com</code> |
| Rachel Willmer | <code>rachelw@spider.co.uk</code> |
| Sam Wilson | <code>sam.wilson@ed.ac.uk</code> |
| Paul Zawada | <code>Zawada@ncsa.uiuc.edu</code> |

2.2.7 Router Requirements (RREQ)

Charter

Chair(s):

Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:

General Discussion: ietf-rreq@Jessica.Stanford.edu

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

Archive:

Description of Working Group:

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

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

The purposes of this project include:

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

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

Goals and Milestones:

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

Internet-Drafts:

“Requirements for IP Routers Volume 1: Introduction”, 09/17/1990, Philip
Almquist <draft-ietf-rreq-iprouters-04.txt>

Request For Comments:

RFC 1349 “Type of Service in the Internet Protocol Suite”

RFC 1354 “IP Forwarding Table MIB”

2.2.8 Simple Internet Protocol (SIP)

Charter

Chair(s):

Steve Deering, deering@parc.xerox.com

Robert Hinden, hinden@eng.sun.com

Mailing Lists:

General Discussion: sip@caldera.usc.edu

To Subscribe: sip-request@caldera.usc.edu

Archive:

Description of Working Group:

SIP is a candidate for IPng. The purpose of the working group is to finalize the SIP family of protocols, and to foster the early development and experimentation of this protocol.

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

Following the IP model, SIP uses globally-unique addresses, hierarchically structured for efficient routing. SIP addresses are 64 bits long, which is believed to be adequate to scale the Internet up to, say, thousands of internet-addressable devices in every office, every residence, and every vehicle in the world.

The quest of simplicity in SIP has been described as parallel to the RISC philosophy. The minimal SIP header contains only those fields which are necessary to achieve our goal: routing packets efficiently in a very large internet. As a result of this design philosophy, the SIP header is much simpler than the IP header. Simplicity facilitates high-performance implementation and increases the likelihood of correct implementation.

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

Goals and Milestones:

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

- Done Post an Internet-Draft specifying the SIP addressing and routing architecture. Include discussion of multicast and mobile host support as well as a discussion of how policy routing can be supported. Detail the changes required to OSPF, BGP, and RIP.
- Jan 1993 Post as an Internet-Draft a specification for the SIP MIB. Detail the operation of SNMP over SIP.
- Jan 1993 Make available a public domain implementation of SIP for the UNIX-BSD socket environment.
- Jan 1993 Make available a public domain version of modified TCP and UDP for the UNIX-BSD socket environment.
- Mar 1993 Post as an Internet-Draft a report on the initial implementation and experience with SIP.
- Jun 1993 Incorporate security into SIP.
- Done Post an Internet-Draft specifying changes to RIP needed for SIP.

Internet-Drafts:

“SIP-RIP”, 03/11/1993, G. Malkin, C. Huitema <draft-ietf-sip-rip-01.txt>

“SIP Program Interfaces for BSD Systems”, 04/05/1993, R. Gilligan <draft-ietf-sip-bsd-api-00.txt>

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

“SIP System Discovery”, 04/21/1993, W. Simpson <draft-ietf-sip-discovery-02.txt>

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

CURRENT MEETING REPORT

Reported by Robert Hinden/Sun Microsystems

Minutes of the Simple Internet Protocol Working Group (SIP)

These minutes are based on notes taken by Christian Huitema.

The SIP Working Group held two sessions and a demonstration at the Amsterdam IETF. The first session was 12 July at 4:00 p.m. The second session was 15 July at 1:30 p.m. Both sessions were audio/video multicast on the Internet. The demonstration was held on 14, 15, and 16 July.

Agenda

- Administrivia
- Review of Action Items
- Implementation Status Reports
- Demonstration Plans
- SIP Source Routing
- Review of Recent Work
- Assign Action Items

Administrivia

Bob Hinden introduced the agenda. Ross Callon mentioned his desire to add transition plans as a discussion item. The item was added, but due to a lack of time in the second session, was not discussed.

Review of Action Items

Editor's Note: A list of action items and their status is available via FTP or mail server from the remote directories as /ietf/atm/atm-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Implementation Status Reports

• Public Domain BSD

A presentation was given by Werner Vogel. Three BSD implementations have been done: Initial INRIA, full 64 bits, x-kernel. The target of INESC is BSD, specifically Mach and x-kernel.

Werner presented the architecture of the INRIA implementation:

- SIP processor in the kernel
- Interface configuration and route set up
- 64-bit ping
- 32-bit TCP and UDP without fragmentation

Other features are implemented but not yet tested.

Performance of the loop-back interface is faster than straight IP. Performance over Ethernet is equivalent to IP (same figure). NFS (block of 1K) and AFS work over SIP.

Next steps: more debugging, real 64-bit TCP, transport level support, integration of routing, use real interfaces, checksums, etc.

• Sun Solaris Implementation

Erik Nordmark gave the presentation. Sun included the “border router” code. SIP Multicast is implemented. VAT and NV work over SIP using multicast address translation. They are working on getting “traceroute” to work over the encapsulation, and avoiding the “lost ICMP” problem.

For solving the lost ICMP problem, the SIP process has to keep track of the tunnel’s MTU, and also of the “unreachable” status of tunnels. The TTL exceeded problem is harder to solve. This can be delegated to the routing process for “inter-router” tunnels, but cannot easily be used for “tail” tunnels. Tony Li mentioned that SDR is using “tunnel IDs” (64-bit encapsulation header) in order to solve this problem. He suggested we look at the SDR IDs.

• SIP IDR P Status

Sue Hares described the status of IDR P for SIP. She said that IDR P is part of “gated” which is already multiprotocol. She needs a SunOS 4.1 implementation (INRIA/INESC) to test the relaying of the packets over SIP, and for installing SIP routes. She believes the code is modular enough to install routes without problems. Yakov Rekhter mentioned the possibility of having extra attributes, for automatically installing tunnels. Sue also mentioned

extensions for multicast, for example, using the next hop information for memorizing the “broadcast tree” from a given source. A base level support could be ready for test within a month given a kernel. The link between IDRPs and IGPs other than IS-IS is neither done nor funded. She suggested finding volunteers within the working group. Code is public domain and can now be provided to “co-developers.”

She also mentioned that the ISO IDRPs specification will soon be published as an Informational RFC.

Demonstration Plans

Bob Gilligan presented the IETF demonstration set up. There were 6 sites participating:

- IETF at Amsterdam
- Xerox PARC
- TGV
- Sun
- Intercon Macintosh
- Beame & Whiteside (PC with DOS)

The first 4 sites run a SIP border router; at PARC and TGV, an IP host points to the SIP border router. In the last two sites, PCs and Macintoshes are isolated SIP hosts, connected to the routers in their domain space. Metro addressing is used. Werner volunteered the addition of a BSD SIP host in Portugal to the demonstration.

The demonstration featured Telnet, FTP, Ping, Traceroute, and VAT. FTP “third party” connections are limited to using the same prefix as the control connection.

SIP Source Routing

Charlie Perkins presented the use of source routing for solving the “mobile routing” problem. The classical problem is router efficiency: the forwarding of IPv4 packets with source routes was slow, which led to the use of IP encapsulation. Source Routing (SR) also has a bad reputation for security, though encapsulation has the same inherent problem. SR has slightly less overhead than encapsulation (16 vs 24 octets). ICMP messages are delivered to the source with SR, and to the encapsulator with IP encapsulation. There are also slight differences with fragmentation (reassembly at end of tunnel for encapsulation is less efficient), and MTU discovery in which tunnels are transparent.

The decision is in fact linked to “who does what.” The source itself should do SR, but intermediate hops should use encapsulation.

On the lesson of the mobile IP experience: The SIP specification should be clearer about “reversal of source routed packets.” It is not very clear, but it appears that “layer 4” solutions are generally inadequate. This design should really be studied inside the MOBILEIP Working Group. Tony Li mentioned that it is also being addressed by the SDR Working Group.

Review of Recent Works

• SIP RIP

Gary Malkin presented comments received from Garcia Luna Aceves on SIP RIP. The loop detection algorithm is not described precisely, and needs to be corrected. However, this is a major improvement over previous version of RIP, with the cost of more CPU. A consequence is that the maximum number of hops has been raised to 32.

Paul Francis asserted that loops are better than black holes, as you do not miss packets. He suggested that we look at using a path vector algorithm. Tony Li rejected the idea of accepting routing loops, as they are traffic multipliers that generate congestion; he also said that path tracing is a significant modification of the protocol. He said that cisco found that path tracing breaks when “route filtering” is in operation. He suggests that DUAL, which includes incremental updates, and guarantees loop freedom, is looked at. Toni also mentioned that some networks are larger than 32 hops, and that we should use path metrics, but that would make the whole thing much more complex.

Tony Li then offered to provide SIP-IGRP, giving change control to the IETF for SIP-specific extensions! After considerable discussion, the working group agreed that this should be pursued, given the usual caveats about licensing agreements and change control.

A proposal was made that SIP RIP should be limited to be used in “small networks.” This raises the question of how should the current SIP RIP draft be progressed. The working group decided to continue with a basic version of SIP RIP (without the loop control) and to ask the RIPv2 Working Group to take on the issue of loop control. The current version of SIP RIP (without loop control) will be called SRIP.

• System Discovery

Bill Simpson led the discussion on the system discovery draft.

Not a lot of implementation was done of the current version of the Router Discovery ICMP message type. It was nice, but it lacked extensibility. The current draft proposes a “single block with extensions” format:

| |
|--------------------|
| SIP + ICMP headers |
| IFACE |
| MAC |
| Services |
| Security label |
| Changed prefix |
| QOS |
| Authentication |

This format is similar to Novell's SAP. There are two messages: solicitation and response. The message operates similarly to the router's advertisement ICMP. "Changed prefix" is intended to enable dynamic address reconfiguration, which should have similar effects on TCP as on the current IP mobility solutions, i.e. require some form of source routing to retain the existing address.

The security label is really informational. QOS is "claiming to be the router for a particular QOS." This, as the security label, is equivalent to similar fields in the OSPF and IS-IS "hello" packets.

The service field is used to advertise the location of particular servers, e.g. "DNS" or "bootp."


Tony Li suggested having both a length and an AFI for the "iface" parameter. He also suggested making both "MAC" and "service" optional. Greg Minshall suggest MAC should, on the contrary, be present all the time in order to facilitate parsing. Greg also suggested that the experience acquired by Novell suggests that "service" is not a very good idea—he would prefer to use multicast queries. Steve Deering observes that there are more clients than servers, and that having servers advertise themselves is preferable (less traffic). Geert Jan de Groot questioned this assertion, as the "keep polling with backup" is more stable and easier to diagnose (the repeated packet pops in link analyzers, etc.). Bill Simpson mentioned that the algorithm which he described is exactly that of "IP router discovery," i.e., tested and true.

Paul Francis questioned the utility of the QOS field: there is no such thing as a QOS per router, but rather per router/destination tuple. The group agreed that redirection is a better solution. Paul also suggested that a strictly router-to-host protocol is much simpler than router-to-router hellos, and that the two groups do not have the same frequency and complexity requirements.

In order to do this for mobile systems, one also needs to carry a "list of routers heard by mobile" in the solicitation messages send by the mobiles. This needs to be discussed on the SIP mailing list.

• Host Auto Configuration

Bob Gilligan presented a set of “preliminary ideas” that he proposed to the mailing list on auto configuration. He proposes to represent the address as a combination of:



The suffix part is allocated by the system administrator. The prefix is heard from the router advertisement. At boot time, the system obtains (by various means) the “local suffix” (e.g. 32-bit IP address); then it obtains the “prefix” from the router advertisement and combines it to form a complete address.

Christian Huitema suggested that this is a very dangerous scheme as one can inadvertently boot the system in a new environment where the suffix is not unique. Bill Simpson suggested using a combination of IEEE 802 and directory names.

Paul Francis suggested the use of a two hop source route: the IEEE 802 unique SIP address of the host, and the router address obtained from the advertisement.

Conclusion and Assignment of Action Items

Steve Deering mentioned the need for more implementations, and also the need to start deployment. Members were encouraged to go see the demonstration in the terminal room, with border routers, VAT over SIP, Internet Talk Radio acquired over SIP, etc.

Attendees

| | |
|-------------------|---------------------------------|
| Kannan Alagappan | kannan@DSMAIL.ENET.DEC.COM |
| Steve Alexander | stevea@lachman.com |
| James Allard | jallard@microsoft.com |
| Frederik Andersen | fha@dde.dk |
| Michael Anello | mike@xlnt.com |
| Anders Baardsgaad | anders@cc.uit.no |
| Dennis Baker | dbaker@wellfleet.com |
| John Ballard | jballard@microsoft.com |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |

| | |
|---------------------|-----------------------------------|
| Per Bilse | bilse@ic.dk |
| Jim Binkley | jrb@ibeam.intel.com |
| David Borman | dab@cray.com |
| Jim Bound | bound@zk3.dec.com |
| Robert Braden | braden@isi.edu |
| Ronald Broersma | ron@nosc.mil |
| Suzy Brown | suzy_brown@genmagic.com |
| John Burnett | jlb@adaptive.com |
| Ross Callon | rcallon@wellfleet.com |
| Peter Cameron | cameron@xylint.co.uk |
| Henry Clark | henryc@oar.net |
| Dave Cullerot | cullerot@ctron.com |
| Walid Dabbous | Walid.Dabbous@sophia.inria.fr |
| James Davin | davin@thumper.bellcore.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Stephen Deering | deering@parc.xerox.com |
| Tim Dixon | dixon@rare.nl |
| Kurt Dobbins | dobbins@ctron.com |
| Pierre Dupont | dupont@mdd.comm.mot.com |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Mark Fedor | fedor@psi.com |
| Eric Fleischman | ericf@act.boeing.com |
| Paul Francis | Francis@thumper.bellcore.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Robert Gilligan | Bob.Gilligan@Eng.Sun.Com |
| Joseph Godsil | jgodsil@ncsa.uiuc.edu |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Marcel Graf | graf%dhdibm1.bitnet@vm.gmd.de |
| Terry Gray | gray@cac.washington.edu |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Robert Hinden | hinden@eng.sun.com |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| Gerd Holzhauser | Holzhauser1@applelink.apple.com |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| Ronald Jacoby | rj@sgi.com |
| David Johnson | dbj@cs.cmu.edu |
| Marijke Kaat | marijke@sara.nl |
| Tomaz Kalin | kalin@rare.nl |
| Neil Katin | katin@eng.sun.com |
| Ton Koelman | koelman@stc.nato.int |
| John Larson | jl Larson@parc.xerox.com |
| Tony Li | tli@cisco.com |
| Marjo Mercado | marjo@cup.hp.com |

| | |
|-----------------------|----------------------------------|
| Donald Merritt | don@arl.army.mil |
| Greg Minshall | minshall@wc.novell.com |
| Keith Mitchell | keith@pipex.net |
| Daniel Myers | dan@nsd.3com.com |
| Erik Nordmark | nordmark@eng.sun.com |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Zbigniew Opalka | zopalka@agile.com |
| Jorg Ott | jo@cs.tu-berlin.de |
| Christian Panigl | christian.panigl@cc.univie.ac.at |
| Charles Perkins | perk@watson.ibm.com |
| David Piscitello | dave@mail.bellcore.com |
| Aiko Pras | pras@cs.utwente.nl |
| James Reeves | jreeves@synoptics.com |
| Alex Reijnierse | a.a.reijnierse@research.ptt.nl |
| Yakov Rekhter | yakov@watson.ibm.com |
| Dan Romascanu | dan@lannet.com |
| Marjo Rottschaefer | |
| Henry Sanders | henrysa@microsoft.com |
| Kitty Shih | kmsih@novell.com |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| John Stewart | jstewart@cnri.reston.va.us |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Richard Thomas | rjthomas@bnr.ca |
| Susan Thomson | set@bellcore.com |
| Thierry Turletti | turletti@sophia.inria.fr |
| Hisao Uose | uose@tnlab.ntt.jp |
| Willem van der Scheun | scheun@sara.nl |
| Werner Vogels | werner@inesc.pt |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Jessica Yu | jyy@merit.edu |
| Romeo Zwart | romeo@sara.nl |

SIP WORKING GROUP MEETING

July 12, 1993

Internet Engineering Task Force Meeting
Amsterdam

AGENDA

- Administrivia
- Review Actions Items
- Implementation Status Reports
 - BSD 4.3 Implementation
 - Sun Implementation
 - Other Implementations
- Demonstration Plans
- SIP Source Routing / Charlie Perkins

SIP Working Group

July 12, 1993

AGENDA (CONT.)

- Review Recent Work
 - System Discovery
 - Host Auto Configuration
 - SIP RIP
 - SIP DNS
- Conclude and Assign Action Items

SIP Working Group

July 12, 1993

ACTION ITEMS

- ACTION:** Everyone read Auto-Configuration proposal and reply to list and put on agenda for next meeting.
- ACTION:** Simpson and Deering resolve differences and come up with one addressing allocation/assignment scheme.
- ACTION:** Crocker define and plan Amsterdam demo.
- ACTION:** Hinden/Deering agenda for Amsterdam meeting.
- ACTION:** Deering to send message to list outlining IPv4 ID generation choices and propose a solution.
- ACTION:** Gilligan/Mulligan to define and write up.

SIP Working Group

July 12, 1993

ACTION ITEMS (CONT.)

- ACTION:** Hinden to send Erik F. a message stating that the w.g. will develop a formal response.
- ACTION:** Gilligan to post his auto configuration proposal to list.
- ACTION:** Jim Bound to coordinate a w.g reply.
- ACTION:** Deering to contact other IPng chairs about coordinating IESG submissions.
- ACTION:** Gilligan to revise and reissue BSD API for SIP document.
- ACTION:** Deering to work on separate Flow ID document.
- ACTION:** Deering to talk to Ran Atkinson about status of SIP security proposal.

SIP Working Group

July 12, 1993

ACTION ITEMS (CONT.)

- ACTION:** Christian Huitema: To post as an Internet Draft of DNS changes for SIP (if not already posted).
- ACTION:** Simpson to get status of IDRP work and report to list.
- ACTION:** Deering/Hinden to ask John Moy to do revision of OSPF for SIP document.
- ACTION:** Deering to write ICMP for SIP document.
- ACTION:** Deering will also include IGMP changes to ICMP document.
- ACTION:** Christian Huitema: Will produce a new version of SIP for RIP document or get Gary Malkin to do it.
- ACTION:** Deering to look at SIP RIP to make sure it includes multicast

SIP Working Group

July 12, 1993

ACTION ITEMS (CONT.)

ACTION: Simpson to get status of IDRIP work and report to list.

ACTION: Deering/Hinden to ask John Moy to do revision of OSPF for SIP document.

ACTION: Deering to write ICMP for SIP document.

ACTION: Deering will also include IGMP changes to ICMP document.

ACTION: Christian Huitema: Will produce a new version of SIP for RIP document or get Gary Malkin to do it.

ACTION: Deering to look at SIP RIP to make sure it includes multicast support.

ACTION: Deering to get first version of SIP addressing document out before Amsterdam IETF.

SIP Working Group July 12, 1993

ACTION ITEMS (CONT.)

ACTION: Deering to send SIP list contents to Hinden.

ACTION: Hinden to revise charter and submit to Internet AD's

ACTION: Gilligan post a message to SIP list asking for volunteers to deploy and test Sun border router implementation.

ACTION: Mulligan send KA9Q code to Simpson.

ACTION: Deering to update SIP specification. Small amount of changes.

ACTION: Gilligan/Nordmark to provide updates to IPAE Specification by June 18.

ACTION: Crocker to do revision of IPAE specification by June 25.

SIP Working Group July 12, 1993

ACTION ITEMS (CONT.)

ACTION: Hinden to update and submit criteria as Informational RFC.

ACTION: Crocker to ask Marshall Rose to develop SIP MIBs.

SIP Working Group July 12, 1993

SIP DEMONSTRATION

July 14-16, 1993

Internet Engineering Task Force Meeting
Amsterdam

SIP OVERVIEW

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

SIP Working Group

July 14, 1993

WHAT IS BEING DEMONSTRATED


- SIP Communication
- SIP <-> IPv4 Translation
- SIP Encapsulation (IPAE)
- SIP Border Router **
- SIP Multicast **
- SIP Traceroute
- SIP Packet Monitoring

** New from Columbus IETF Demonstration

SIP Working Group

July 14, 1993

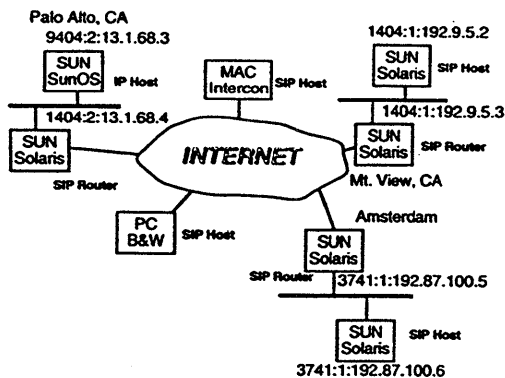
SIP DEMONSTRATION DOMAINS



SIP Working Group

July 14, 1993

DEMONSTRATION CONFIGURATION



SIP Working Group

July 14, 1993

SIP to SIP COMMUNICATION

- On Amsterdam SIP Host, enter:
 - % telnet 1404:1:192.9.5.2
- Traffic flows
 - Pure SIP from across Amsterdam subnet to Amsterdam SIP Border Router
 - IPAE across the Internet to Mt. View border router
 - Pure SIP across Mt. View subnet to Mt. View SIP Host.
- Monitor SIP and IPAE Traffic in *Snoop SIP* and *Snoop IPAE* windows
- Demonstrates SIP Forwarding between SIP Border Routers across an IP Backbone

SIP Working Group

July 14, 1993

SIP to IP COMMUNICATION

- On Amsterdam SIP Host, enter:
% telnet 9404:2:13.1.68.3
- Traffic flows
 - Pure SIP from across Amsterdam subnet to Amsterdam SIP Border Router
 - IPAE across the Internet to Palo Alto Border Router
 - IP across Palo Alto subnet to Palo Alto IP Host
 - Return IP Traffic is Mapped to SIP in Palo Alto Border Router
- Monitor SIP and IPAE Traffic in *Snoop SIP* and *Snoop IPAE* windows
- Demonstrates SIP Forwarding between SIP Border Routers across an IP Backbone and Translation to/from IP

SIP Working Group

July 14, 1993

IP to IP COMMUNICATION

- On Amsterdam SIP Host, enter:
% telnet 13.1.68.3
- Traffic flows
 - Pure IP from across Amsterdam subnet to Amsterdam SIP Border Router
 - Amsterdam Border Router Maps IP Traffic to SIP
 - IPAE across the Internet to Palo Alto Border Router
 - IP across Palo Alto subnet to Palo Alto IP Host
 - Return IP Traffic is Mapped to SIP in Palo Alto Border Router
- Monitor SIP and IPAE Traffic in *Snoop SIP* and *Snoop IPAE* windows
- Demonstrates IP Forwarding between SIP Border Routers across an IP Backbone and two Mapping to SIP

SIP Working Group

July 14, 1993

SIP MULTICAST

- VAT is running on SIP Host and SIP Border Router
 - Turn on Microphone and talk
- Traffic flows
 - Pure SIP Multicast across Amsterdam subnet
- Monitor SIP Multicast Traffic in *Snoop SIP* windows
- Demonstrates SIP Multicast using Unmodified VAT between two SIP Hosts on one subnet

SIP Working Group

July 14, 1993

OTHER APPLICATIONS

- FTP, TFTP, PING, TELNET
 - All Work in all Modes (SIP, IPAE, IP)
 - e.g. ping 9404:2:13.1.68.3
ping 13.1.68.3
 - Internet Talk Radio Play over SIP TCP Connection
 - % itr_play -v 1404:1:192.9.5.2
- Note: Release Audio Device by Clicking Lower Name Stripe in VAT

SIP Working Group

July 14, 1993

SIP TRACEROUTE

- On Amsterdam SIP Host, enter:
% /usr/sbin/traceroute 1404:1:192.9.5.2
- Observe
 - SIP Hops Along Path
 - Note: Currently don't show IPAE Hops
- Demonstrates Tracing of SIP/IPAE Routes

SIP Working Group

July 14, 1993

ACTIVE PARTICIPANTS

- | | |
|---------------------|---|
| • Beame & Whiteside | Implementation (PC) |
| • DEC | Implementation (Alpha/OSF) |
| • INRIA | Implementation (BSD, BIND), DNS & OSPF Specifications |
| • INESC | Implementation (BSD/Mach/x-kernel) |
| • Intercon | Implementation (MAC) |
| • MCI | Phone Conferences |
| • Merit | IDRP for SIP Specification |
| • Network General | Implementation (Sniffer) |
| • SGI | Implementation (IRIX, NetVisualizer) |
| • Sun | Implementation (Solaris 2.x, Snoop) |
| • TGV | Implementation (VMS) |
| • Xerox PARC | Steve Deering |
| • Bill Simpson | Implementation (KA9Q) |

SIP Working Group

July 14, 1993

IMPLEMENTATIONS

| OS | Organization | Status |
|---------------|-------------------|---------------------------------------|
| BSD/Mach | INESC | Completed (telnet, NFS, AFS, UDP) |
| DOS & Windows | Beame & Whiteside | Completed (telnet, ftp, tftp, ping) |
| IRIX | Silicon Graphics | In progress (ping) |
| KA9Q | Simpson | In progress (ping) |
| Mac OS | Intercon | Completed (telnet, ftp, finger, ping) |
| OSF | DEC | In Progress |
| Solaris | Sun | Completed (telnet, ftp, tftp, ping) |
| VMS | TGV | Completed (telnet, ftp) |

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

| MISC | Organization | Status |
|------|--------------|-----------|
| Bind | INRIA | Code done |

SIP Working Group

July 14, 1993

DOCUMENTS

- *SIP: A Simple Internet Protocol*, S. Deering, May '93 IEEE Network.
- *Simple Internet Protocol (SIP) Specification*, Internet Draft, draft-deering-sip-00.txt
- *SIP-RIP*, G. Malkin, C. Huitema, Internet Draft, draft-ietf-sip-rip-01.txt
- *IDRP for SIP*, S. Hares, Internet Draft, draft-ietf-ipidr-sip-00.txt
- *OSPF for SIP*, C. Huitema, Internet Draft, draft-ietf-sip-ospf-00.txt
- *SIP Addresses in the Domain Name Service Specifications*, C. Huitema, Internet Draft, draft-ietf-sip-dnss-00.txt
- *SIP Program Interfaces for BSD Systems*, R. Gilligan, Internet Draft, draft-ietf-sip-bsdapi-00.txt
- *SIP System Discovery*, W. Simpson, Internet Draft, draft-ietf-sip-discovery-02.txt
- *Administrative Allocation of the 64-bit Number Space*, W. Simpson, Internet Draft, draft-ietf-sip-64bit-plan-00.txt
- *IP Address Encapsulation (IPAE): A Mechanism for Introducing a New IP*, D. Crocker, Internet Draft, draft-ietf-ipae-new-ip-00.txt
- *IPv7 Criteria Analysis for IP Address Encapsulation (IPAE) and the Simple Internet Protocol (SIP)*, R. Hinden, Internet Draft, draft-ietf-ipae-ipv7-criteria-00.txt

SIP Working Group

July 14, 1993

WORKING GROUP INFORMATION

- SIP Working Group Mailing List
sip-request@caldera.usc.edu
- SIP Archive
parcftp.xerox.com /pub/sip/

SIP Working Group

July 14, 1993

2.2.9 TCP/UDP Over CLNP-Addressed Networks (TUBA)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

In addition to protocol layering issues and "proof of concept" work, the TUBA approach will place significant emphasis on the engineering and operational requirements of a large, global, multilateral public data network. TUBA will work to maximize interoperability with the routing and addressing architecture of the global CLNP infrastructure. The TUBA Working Group will work closely with the IETF NOOP and OSI IDRIP for IP Over IP Working Groups to coordinate a viable CLNP-based Internet which supports the applications which Internet users depend on such as TELNET, FTP, SMTP, NFS, X, etc. The TUBA Working Group will also work collaboratively with communities which are also using CLNP, and will consider issues such as interoperability, applications coexisting on top of multiple transports, and the evolution of global public connectionless datagram networks, network management and instrumentation using CLNP and TUBA, and impact on routing architecture and protocols given the TUBA transition.

The TUBA Working Group will consider how the TUBA scheme will support transition from the current IP address space to the future NSAP address

space without discontinuity of service, although different manufacturers, service providers, and sites will make the transition at different times. In particular, the way in which implementations relying on current 32-bit IP addresses will migrate must be considered. TUBA will ensure that IP addresses can be assigned, for as long as they are used, independently of geographical and routing considerations. One option is to embed IP addresses in NSAP addresses, possibly as the NSAP end-system identifier. Whatever scheme is chosen must run in a majority of *-GOSIPs and other NSAP spaces. The TUBA strategy will require a new mapping in the DNS from NAMEs to NSAP addresses.

The rationale RFC (RFC 1347) documents issues of transition and coexistence, among unmodified "IP" hosts and hosts which support "TUBA" hosts. Hosts wishing full Internet connectivity will need to support TUBA.

Goals and Milestones:

- Done Post Initial TUBA rational and discussion as an RFC. (RFC 1347)
- Done Post the Initial TUBA DNS specification. (RFC 1348)
- Done Review and approve the Charter.
- Done Post the TUBA CLNP profile as an Internet-Draft.
- Done Post a Routing and Addressing specification as an Internet-Draft, coordinated with the Network OSI Operations Working Group and the IDRP for IP Working Group.
- Nov 1992 Post a summary report on TUBA deployment in the Internet.
- Done Present the results of Working Group deliberations at the November IETF meeting.
- Nov 1992 Post an Internet-Draft on the changes required to Internet applications affected by the deployment of TUBA.
- Nov 1992 Post an Internet-Draft covering the methodologies, instrumentation, address administration, routing coordination and related topics.
- Done Post as an Internet-Draft a revision to RFC1347 reflecting lessons learned in the Working Group deliberation.

Internet-Drafts:

"Use of ISO CLNP in TUBA Environments", 09/04/1992, David Piscitello
<draft-ietf-tuba-clnp-04.txt>

"Assignment of System Identifiers for TUBA/CLNP Hosts", 04/30/1993, D. Piscitello
<draft-ietf-tuba-sysids-03.txt>

CURRENT MEETING REPORT

Reported by John Scudder/Merit

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

Summary

- Tasks
- Documents to be moved to Proposed Standard or Informational
- To-do list
- Presentations

Tasks

Editor's Note: A list of tasks is available via FTP or mail server from the remote directories as /ietf/tuba/tuba-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Documents to be Moved to Proposed Standard or Informational

- CLNP for TUBA [draft-ietf-tuba-clnp-03.txt]

Will be presented to the area director to be moved to Proposed Standard.

- Sysids [draft-ietf-tuba-sysids-01.txt]

Will be presented to the area director to be moved to Proposed Standard. This is already how OSI hosts at Merit are addressed.

It was suggested to present this to the ATM Forum—David Piscitello and Brian Carpenter will pursue this off-line.

- NSAP Allocation Guidelines [RFC 1237]

This document is currently a Proposed Standard. Ross Callon suggests that it needs editing (and volunteers to do it, too). Ross will edit it, place it in the tuba-docs directory on merit.edu, and send a notice to the mailing list (maybe to the NOOP list too).

RFC 1237 will be recommended to be moved to Draft Standard after editing is complete (before the next IETF).

- FUBAR (FTP and UDP with Bigger Addresses) [draft-piscitello-ftp-bigports-01.txt, tuba-only version]

TUBA and TP/IX implementations of FUBAR supposedly exist.

There was quite a bit of discussion about problems with FUBAR and TUBA translating gateways.

Some editing is needed on the document: five-letter commands need to be changed to four-letter, and various frivolities need to be elided. An appendix is to be written specifying use of FUBAR for TUBA.

In the spirit of compromise between problems with FUBAR over translating gateways and the need for some specification for big address FTP, there was agreement to move for Experimental status now, to be reviewed at the next IETF and then moved to Proposed Standard.

- DNS forward lookup (name → NSAP lookup)

There is a document for forward lookup only, no inverse lookup.

RFC 1238 needs to be moved to Historical, since reverse lookup is “broken.” Inverse lookup has been implemented, but is very slow. There is a new Internet-Draft that does not include reverse lookup. Richard Colella and Bill Manning will edit the Internet-Draft for next time. RFC 1238 will be left in place for now.

A DNS guru volunteer is needed. Richard is interested in working with this guru.

This will be discussed in Wednesday’s DNS meeting.

- Routing and Addressing Architecture

There is already such an architecture published in ISO 957x (David Piscitello or Dave Katz may know the real number).

957x will be translated to ASCII text. (Mark Knopper will work on doing this, Lyman Chapin, Yakov Rekhter and David Piscitello will provide a raw document.)

David Piscitello is changing 9542 to ASCII, Lyman is changing 8473, and Kunzinger has changed 10747. IS-IS and 957x need to be done. All will be recommended as Proposed Standards and made available both in ASCII and PostScript.

Relevant ISO documents are available (in PostScript) for anonymous FTP from merit.edu.

- EON

Will be recommended as a Proposed Standard.

To-do List

- DNS inverse lookup (mentioned above)
- Transition plan

To be discussed at the next meeting. Some anxiety was expressed that the plan needs to be finished well before the next IETF.

Peter Ford and John Curran are working on a transition plan.

A rough transition outline is:

```
Dual-stacked hosts
CLNP in routers
CLNP over IP infrastructure
IP over CLNP infrastructure
```

This segued into a discussion of the existing infrastructure, which led to discussion of EON: the EON RFC (RFC 1070) is still in Experimental status. There was some discussion about whether changes to EON are needed and worthwhile. Dave Katz volunteered to edit it and recommend it as a Proposed Standard. John Scudder will try out EON in BSD/386.

It might also be useful to have an IP in CLNP tunneling documents.

- Mobile hosts: Yakov Rekhter commented that TUBA will adopt whatever the Internet community decides on for IP.
- Formulate RFC 1380 responses.
- Working groups we have/want liaison with: DNS, FTP, ATM, RARE, NOOP, and any working groups arising from the OSIEXTND BOF.

Presentations**Autoconfiguration “a la” DEC (Chris Gunner)**

NSAP structure:

```
| -Area Address-----| -ID-----| -SEL-|
<-----n octets-----> <-6 oct--> <-1->
```

- Routers are configured (by hand) with area addresses
- End-systems “know” their IDs (e.g. MAC address) and “know” SEL(s)
- Routers send IS hellos (ISO 9542) with NET (NSAP)

- End-systems receive IS Hello and:
 - Extract area address
 - Create NSAP(s) (area address + ID + SEL(s))
 - Send ES Hello(s) with NSAP(s)

The migration to new area addresses is said to be pretty easy since an end-system can have both an “old area” and a “new area” NSAP.

Named objects, e.g. “node” (system), may have protocol “stack” attribute information, e.g.: (in DEC DNS)

```

+-----+      +-----+
| Upper Layers | ==> | SNMP      |
| CLNP, NSAP(s) |    | UDP, Port # |
+-----+      | CLNP, NSAPs |
                  +-----+

```

When an end system’s NSAP(s) change:

- Update naming service entry for objects for that system
- Requires name service protocol to do update
- System needs to have write access to these objects

This is basically a way for end-systems to update the DNS automatically when their addresses change. There was some concern of how to do this in the current DNS—Yakov commented that when standard IP DNS knows how to do this, TUBA will adopt it unchanged.

Issues:

- Frequency of updates
- Update failure – e.g. no write access – requires manual DNS override ability
- System state information about interaction with name service (transient failures)

Multicast (Dave Katz)

- Group NSAP addresses hack

Parallel AFI space (10-99 → A0-F9) (since AFI is in BCD)

- Synactically distinct but parallel space
- Hierarchy possible (unlike IP multicast space)

- CLNP
 - Multicast Data (MD) PDU
Distinct from DT PDU
 - Scope control options? (“I want this packet to go only this many administrative hops.”)
- ES-IS
 - NSAP → SNPA dynamic binding
 - Group membership announcement
 - Extra unicast hop – if you want to send multicast, you unicast your packet to an IS which then forwards it appropriately. You never get a redirect to start multicasting on the LAN.
- IS-IS
 - Could be changed to be MOSPF-like
 - No active work
- IDRP

No work yet
- For more information see OSI Extensions for use in the Internet BOF

ES-IS Address Administration (Dave Katz)

See ES-IS second edition. PostScript file on merit.edu.

```

      ES                                IS
      -----                          -----
      "who am I?"      ---->
      (to ask for an
      address)

                                <---  "You are foo" (for 18 hours)
                                <---  "You are bar"
                                (offers some addresses, guaranteed
                                to be reserved for ES for holding
                                timer duration)

      "I am bar" (ESH) ---->
      (to notify IS of
      who ES has decided
      to be, incl holding
      time of up to 18
      hours)
  
```

Issues:

- May not really want automatic assignment (security concerns)
- IS does not know some host information (e.g., IP address)—it might be nice to provide this input to construct the NSAP (or MAC address, other host-specific info)
- How can we deny service to undesired hosts? (e.g., send an end-system a bogus address to “shut him up”)

Attendees

| | |
|-------------------|--------------------------------------|
| Nick Alfano | alfano@mpr.ca |
| James Allard | jallard@microsoft.com |
| Bernt Allonen | bal@tip.net |
| Josee Auber | Josee_Auber@hpgnd.grenoble.hp.com |
| Anders Baardsgaad | anders@cc.uit.no |
| John Ballard | jballard@microsoft.com |
| Rebecca Bostwick | bostwick@es.net |
| Jim Bound | bound@zk3.dec.com |
| Thomas Brunner | brunner@switch.ch |
| Ross Callon | rcallon@wellfleet.com |
| Brian Carpenter | brian@dxcern.cern.ch |
| George Chang | gkc@cct.bellcore.com |
| Richard Colella | colella@nist.gov |
| David Conrad | davidc@ij.ad.jp |
| Tim Dixon | dixon@rare.nl |
| Kurt Dobbins | dobbins@ctron.com |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Francis Dupont | francis.dupont@inria.fr |
| Dino Farinacci | dino@cisco.com |
| Stefan Fassbender | stf@easi.net |
| Eric Fleischman | ericf@act.boeing.com |
| Osten Franberg | euaokf@eua.ericsson.se |
| Peter Furniss | p.furniss@ulcc.ac.uk |
| Eugene Geer | ewg@cc.bellcore.com |
| David Ginsburg | ginsb@us-es.sel.de |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Patrick Hanel | hanel@yoyodyne.trs.ntc.nokia.com |
| Susan Hares | skh@merit.edu |
| Denise Heagerty | denise@dxcoms.cern.ch |
| Jack Houldsworth | J.Houldsworth@ste0906.wins.icl.co.uk |
| Chris Howard | chris_howard@inmarsat.org |
| Geoff Huston | g.huston@aarnet.edu.au |

| | |
|----------------------|--------------------------------|
| Phil Irey | pirey@relay.nswc.navy.mil |
| David Jacobson | dnjake@vnet.ibm.com |
| J. Jensen | jensen@ic.dk |
| Thomas Johannsen | thomas@ebzaw1.et.tu-dresden.de |
| Dale Johnson | dsj@merit.edu |
| Philip Jones | p.jones@jnt.ac.uk |
| Cyndi Jung | cmj@3com.com |
| Anders Karlsson | sak@cdg.chalmers.se |
| Daniel Karrenberg | daniel@ripe.net |
| Frank Kastenholz | kasten@ftp.com |
| Dave Katz | dkatz@cisco.com |
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Mark Knopper | mak@merit.edu |
| Ton Koelman | koelman@stc.nato.int |
| Tony Li | tli@cisco.com |
| Susan Lin | suelin@vnet.ibm.com |
| Robin Littlefield | robin@wellfleet.com |
| Bill Manning | bmanning@rice.edu |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Cynthia Martin | martin@spica.disa.mil |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Jun Murai | jun@wide.ad.jp |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Andrew Partan | asp@uunet.uu.net |
| David Piscitello | dave@mail.bellcore.com |
| Willi Porten | porten@gmd.de |
| Aiko Pras | pras@cs.utwente.nl |
| Juergen Rauschenbach | jrau@dfn.de |
| Alex Reijnierse | a.a.reijnierse@research.ptt.nl |
| Victor Reijs | reijs@surfnet.nl |
| Yakov Rekhter | yakov@watson.ibm.com |
| Robert Reschly | reschly@brl.mil |
| Georg Richter | richter@uni-muenster.de |
| Luc Rooijakkers | lwj@cs.kun.nl |
| Henry Sanders | henrysa@microsoft.com |
| John Scudder | jgs@merit.edu |
| Keith Sklower | sklower@cs.berkeley.edu |
| Michael St. Johns | stjohns@darpa.mil |
| Henk Steenman | henk@sara.nl |
| Vladimir Sukonnik | sukonnik@process.com |
| Fumio Teraoka | tera@cs1.sony.co.jp |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Paul Traina | pst@cisco.com |

| | |
|-----------------------|---------------------|
| Antoine Trannoy | trannoy@crs4.it |
| Willem van der Scheun | scheun@sara.nl |
| Marcel Wiget | wiget@switch.ch |
| Kirk Williams | kirk@sbctri.sbc.com |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Noriko Yokokawa | norinori@wide.ad.jp |
| Jessica Yu | jyy@merit.edu |
| James Zmuda | zmuda@mls.hac.com |
| Romeo Zwart | romeo@sara.nl |

CURRENT MEETING REPORT

Reported by Peter Ford/LANL and John Scudder/Merit

Minutes of the joint session of TUBA, NOOP and RARE CLNS Working Groups

Marcel Wiget reported on the RARE CLNS Pilot. CLNS was run over a variety of media across Europe. The objectives of this project were routing coordination and testing of CLNS implementations to feedback to vendors and implementors. Systems were deployed on a wide variety of platforms including cisco, DEC, Sun, BSD systems, Novell and others. Applications tested included X.400, X.500, FTAM, DECnet, VT, TUBA, OSI Ping, and traceroute.

RARE developed databases for the distribution of addresses (NSAP, NSEL, TSEL, etc.) and routing topology. Tools used for managing the network included clnp-traceroute, tcpdump (INRIA), bis_ping, and proxy managers for sunnet manager. They also have extensions to the Internet rover. There are also tools for managing consistency between DNS and host files. All of these are available via anonymous FTP from ftp.rare.nl or anonymous FTAM/FTP from nic.switch.cs:/network/clns. To be added to the discussion list, send a request to wg-llt-clns-request@rare.nl.

Current activities include a task force for CLNS routing coordination. They are discussing how to take advantage of the RIPE routing registry which will require developing a schema for representing this information. There is also discussion of forming TUBA and IDRIP task forces but there is a sense that perhaps it is best to do that sort of work in the IETF.

In discussing the scale of the CLNS pilot, it was interesting to note that the UK has now turned on Phase V DECnet, and that the Italian high energy physics lab has over 100 CLNS routers at this time. It is estimated that there are over 200 CLNS routers in Europe at this time.

Dave Katz reported that the IDRIP specification is on its way to Geneva, and that it will become a full International Standard at the October ISO meeting in Seoul.

A spirited discussion was held on the use of IS-IS for routing the global CLNS network. Many felt that the size of the current CLNS system was too large to be managed with IS-IS. There are also issues of whether or not certain networks (e.g. CERN) would want to absorb the routing of the entire world into their own world. Dave Katz suggest building a "virtual core" for doing some of this. Further discussion of this will be done on the NOOP mailing list.

There are now many TUBA implementations:

- PC/MS-DOS by Richard Colella. Source code is available via anonymous FTP from [osi.ncsl.nist.gov](ftp://osi.ncsl.nist.gov).

- SunOS - kernel mode TUBA by Francis DuPont. Source not available unless you have a BSD 4.3 license.
- Francis Dupont reported on the user mode network layer translator implemented in user space on SunOS that currently runs on top of Sun's OSI implementation. The sources are available; send mail to Francis.
- Peter Ford reported on further work on top of Keith Sklower's implementation of TUBA on the BSDI platform. If you have a BSD 4.3 source code license you can get these sources for BSDI (send mail to Peter). There are now implementations of TELNET, FTP (using Dave Piscitello's FUBAR specification), finger, and there is a version of inetd which manages CLNP sockets for Unix implementations.
- Cyndi Jung reported that 3Com routers continue to support TUBA.
- Dino Faranacci reported cisco routers also support TUBA. Images can be obtained from `ftp.cisco.com:/beta921-dir/*.931_O.14.Z` where * is replaced with the cisco model you have. You need to send mail to `clns-beta-request@cisco.com` to get on the mailing list if you pick up one of these images.

Cyndi Jung reported on CLNS efforts for the next INTEROP in San Francisco (August 21 to August 26). She noted that this would be a good opportunity to demonstrate TUBA; George Chang is working on getting some spots around the show floor to demonstrate TUBA. The mailing list for discussion is `osiig@spitfire.interop.com`.

The meeting ended with a discussion of where to discuss TUBA deployment. It was decided that this should happen on the NOOP mailing list.

2.2.10 TP/IX (TPIX)

Charter

Chair(s):

Vladimir Sukonnik, sukonnik@process.com

Mailing Lists:

General Discussion: tpix@world.std.com

To Subscribe: tpix-request@world.std.com

Archive: [world.std.com:~/pub/tpix/](http://world.std.com/~pub/tpix/)*

Description of Working Group:

TP/IX is a new version of the IP, TCP, and UDP protocols, to advance the Internet technology to the scale and performance of the next generation of internetwork technology. TP/IX has been assigned the IP version number 7.

The working group is chartered to review the TP/IX and RAP protocols, evaluate issues arising during product development and deployment planning, and to document problems and explanations for any parts of the coexistence with IPv4 not covered directly in the TP/IX-IPv4 interoperation design.

The group will also be the initial forum for development of the RAP protocol while it is experimental; this work will need to be moved to the Routing Area when it is to be advanced.

Goals and Milestones:

- | | |
|----------|--|
| Done | Present the TP/IX (formerly IPv7) and the RAP protocols to the IETF Plenary. |
| Done | Post the TP/IX Protocol and the RAP protocol as Experimental RFCs. |
| Done | Hold Working Group meeting to discuss additional definitions. Prepare criteria to be met prior to standardization. |
| Nov 1993 | Hold Working Group meeting to evaluate the TP/IX and RAP protocols for Proposed Standard. |
| Dec 1993 | Submit the TP/IX and RAP Protocols to the IESG for consideration as a Proposed Standard. |

Internet-Drafts:

“Initial AD Assignment Plan”, 06/07/1993, R. Ullmann <draft-ietf-tpix-adplan-01.txt>

“Transit Policy Routing in TP/IX”, 06/15/1993, R. Ullmann <draft-ietf-tpix-transit-01.txt>

“TCP version 7 options”, 06/30/1993, R. Ullmann <draft-ietf-tpix-tcpopt-00.txt>

Request For Comments:

RFC 1475 “TP/IX: The Next Internet”

RFC 1476 “RAP: Internet Route Access Protocol”

CURRENT MEETING REPORT

Reported by Vladimir Sukonnik/Process Software Corporation

Minutes of the TP/IX Working Group (TPIX)

First Session Agenda

- AD assignment plan
- Transit network policy

Status of the TP/IX Proposal

The session started with Vladimir Sukonnik's short presentation on the status of TP/IX and RAP projects. RAP version 1 has been released as a commercial product by Process Software Corporation. The work has begun on implementing TP/IX in future releases of the product. Two Experimental RFCs have been published describing RAP and TP/IX, thus setting the stage for vendor prototype implementation. Vladimir has also outlined the main features of TP/IX and how they compare to other IPng proposals.

Transit Network Selection

Robert Ullmann described the Transit Network Selection Internet-Draft. The document outlines an approach to allow network users to select the carrier the same way the telephone customers in the US can select a long-distance provider. The idea is that the border router within the customer's management must be able to acquire knowledge in real time of the availability and costs of the transit networks, and be able to select one for each datagram forwarded to the external router.

Administrative Domain Assignment Plan

Robert presented an idea on how to assign Administrative Domain numbers for the version 7 Internet. The objective is to use a very small amount of space in the numbering system, while providing the necessary distribution of authority. AD numbers are assigned out of the same numbering plan as version 4 network numbers. This helps prevent confusion when the first part of an IPv7 8-byte address is erroneously used as an IPv4 address. It also may be useful in routing ADs with existing routing protocols. The AD 192.0.0 is assigned to the present version 4 numbering plan. This AD has a specific plan for assignment within it:

- The first 24 bits are the AD (192.0.0).
- The next 8 to 24 bits are a network number, each assigned to a specific organization.
- The remaining 16 to 40 bits are assigned to subnets and hosts by authority reserved to a specific organization.

Transition

Tim Dixon asked Robert and Vladimir to elaborate on the transition plan for TP/IX. As noted in RFC 1475, it is possible to provide a mostly-transparent bridge between IPv7 and IPv4. Most of the translation should consist of copying various fields, verifying fixed values in the datagram being translated, and setting fixed values in the datagram being produced. The objective of the conversion is to be able to upgrade systems, both hosts and routers, in whatever order desired by the owners. Organizations must be able to upgrade any given system without reconfiguration or modification of any other system; IPv4 hosts must also be able to interoperate essentially forever.

Future Plans

Robert was asked to elaborate on the future plans for TP/IX:

- RAP version 1 is done and shipping.
- Prototype TP/IX is planned to be shipped in the next release of the software.
- Design is ready for vendor prototype.

Second Session Agenda

- TCP large window/performance options
- Record Marking option

In the second day of the working group meeting, Robert described the TCP version 7 options Internet-Draft. By enlarging the TCP window and sequence number fields to 64 bits, we can avoid the problem that TCP v4 is having now. Mainly, the wrap-around time with the current TCP version is relatively short for fast networks.

Selective Acknowledgement Option

There is a new option to allow the receiver to indicate that some block of data, not “connected” to the left (start) edge of the TCP window, has been received. This option will allow unnecessary retransmissions to be avoided. Only lost segments will be retransmitted, not the whole window. This option is useful on connections with large RTTs and large bandwidths.

Timestamp Option

There is a new option to accurately measure the round-trip delay of the network path being used for a TCP connection. It contains a timestamp value selected by the sending TCP, and a copy of the most recently-received timestamp from the other TCP.

Record Mark Option

This option indicates the boundary of an application record. The record mark is constructed by the TCP service interface at the sender, and passed to the receiver's service interface. It is not used directly by the TCP, except that the TCP may use record marks as hints for where segments might be divided for maximum performance.

Large Port Number Field

Another proposal is to increase the TCP/UDP port number fields to 32 bits. The current version is suffering from "port burn-out." The current field size of 16 bits will max out at 16K connections in four minutes. Port numbers are divided into several ranges:

| | |
|-------------|--|
| 0 | Reserved |
| 1-32767 | Internet registered (well-known) protocols |
| 32768-98303 | Reserved to allow TCPv7-TCPv4 conversion |
| 98304 up | Dynamic assignment |

Attendees

| | |
|---------------------|---------------------------|
| Frederik Andersen | fha@dde.dk |
| Anders Baardsgaad | anders@cc.uit.no |
| Fred Baker | fbaker@acc.com |
| Jim Bound | bound@zk3.dec.com |
| Thomas Cordetti | tomc@digibd.com |
| Al Costanzo | al@akc.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Tim Dixon | dixon@rare.nl |
| Kurt Dobbins | dobbins@ctron.com |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Eric Fleischman | ericf@act.boeing.com |
| Phil Irey | pirey@relay.nswc.navy.mil |

William Simpson
Vladimir Sukonnik
Robert Ullmann

Bill.Simpson@um.cc.umich.edu
sukonnik@process.com
ariel@world.std.com

2.3 Network Management Area

Director(s):

- Marshall Rose: mrose.iesg@dbc.mtview.ca.us

Area Summary reported by Marshall Rose/Dover Beach Consulting

IFIP Electronic Mail Management BOF (EMAILMGT)

Harald Alvestrand presented his document defining an e-mail management model. There is an obvious need for a common language to be used between negotiating parties, e.g., Internet mail manager or GO-MHS managers who talk to PTTs or commercial partners. It is also necessary to have management tools for monitoring, controlling, and planning changes in a Mail Responsibility area. By having a precise model, relationships between the MTA manager, user and e-mail system customer will be more clear. Harald will finalize his paper by adding functions that need monitoring and control, as well as examples for the sendmail-based systems and circulate his next version of the document in the IFIP-EMAILMGT distribution list. After comments from the list members the document will be submitted as an RFC for the Houston meeting.

ATM MIB Working Group (ATOMMIB)

The ATOMMIB Working Group met to discuss the status of the work on ATM and SONET MIBs. The ATM work took up most of the meeting time. Compatibility with the ILMI MIB of the ATM Forum was an important issue. In addition, the working group did a detailed review of the proposed ATM MIB. As a result, a new version will be posted as an Internet-Draft. Some issues were deferred to the mailing list. The SONET MIB work is rather mature. Some minor changes were agreed to. One issue was deferred to the mailing list. A new version will be posted.

Frame Relay Service MIB Working Group (FRNETMIB)

The FRNETMIB Working group met twice at the 27th IETF. Items discussed were Kenneth Rodemann's SMA draft (which received general support), Kenneth's proposed VC-Table contained in the SMA draft (which, after being converted to a "flow" table received general support) and the current draft of the MIB (which received general support and a variety of changes). The FRNETMIB Working Group will continue work on both the FRS MIB and the SMA document with the goal of completing both by the December 1993 date of the original charter. The "flow" table work will be moved to the IFMIB Working Group.

A revised draft of the FRS MIB (in SNMPv2 format) will be posted to the Internet-Draft directory prior to INTEROP in August.

Interfaces MIB Working Group (IFMIB)

The IFMIB Working Group met twice. The interfaces evolution draft was discussed in detail. Issues in the June 1 version were considered first—none were contentious.

1. Consideration was given to replacing certain objects in the old interfaces group: ifOutQlen, ifType and ifSpecific. Proposals will be made to the mailing list.
2. 64-bit counters were discussed. It was resolved that only packet and octet counters would be 64 bits, and the conformance groups were changed.
3. The RFC 1229 additions were discussed. Some objects were deleted, some were not. Retained were ifPromiscuous, a modified ifTestTable, and the ifRcvTable.

Modem Management Working Group (MODEMMGT)

The meeting was attended by some 20 individuals representing SNMP network management vendors, modem manufacturers, the IETF Network Management Area Directorate, and a few users. We spent the bulk of the time discussing 5 significant architectural issues.

Good progress was made in resolving these points. We are hoping the group can produce an Internet-Draft near the end of July. The long-term goal of the group is to agree on the details of a written draft by the end of the year.

SNA NAU Services MIB Working Group (SNANAU)

The group reviewed the issues related to the current draft of the SNANAU MIB. Following is the list of specific results:

- A number of decisions have been reached concerning modifications of the MIB draft in preparation to release it as an Internet-Draft.
- There are some points which have to be discussed and agreed upon with the SNADLC Working Group. They will be contacted via e-mail after the IETF.
- Dave Perkins has been asked to clarify two problems related to SNMPv1-v2 migration (row creation) and the use of the new “if” group.

- We decided to suggest a supplemental working group meeting during one of the popular SNA implementors workshops (sometime in September). The idea is to provide another opportunity for people who already declared their interest or who are potentially interested in this effort to review the MIB draft and become involved in further work.

Token Ring Remote Monitoring Working Group (TRMON)

A final call was made for comments on the TRMON document. The comments received were incorporated into the next draft, which is available as an Internet-Draft. This document was also forwarded to the Network Management Area Director with a recommendation that it become the TRMON Proposed Standard. The Network Management Area Directorate (NMDIR) has started looking at the document and one of their action items is to ensure that there are no conflicts with the Token Ring MIB. Once NMDIR is satisfied with the document, it will be forward to the IESG.

The IEEE 802.5 committee is working on a new draft. Draft four should be available in September of 1993, with a conclusion of the process scheduled for March of 1994.

Once TRMON is released as a proposed standard, the TRMON Working Group will be disbanded and the RMON Working Group reformed for consideration of RFC 1271 advancement. A strawman charter was presented. It was decided that the date for working group recommendation should be March of 1994.

CURRENT MEETING REPORT

Reported by Suzanne Smith/Lawrence Livermore National Laboratory

Minutes of the IFIP Electronic Mail Management BOF (EMAILMGT)

Maria Dimou-Zacharova began the EMAILMGT meeting by requesting Urs Eppenberger to present a brief statement of purpose and meeting schedule for the group.

Harald Alvestrand presented his document defining an electronic mail management model. He explained the need for a common language and precise definitions to promote communication between working groups and commercial partners. Particular care should be taken to ensure that his model document and the requirements document (produced by Working Groups 6.5 and 6.6 of the International Federation for Information Processing (IFIP), and edited by Emily McCoy and Ray Freiworth) use the same terms and definitions.

Harald reviewed each section of the document and discussion followed. Urs Eppenberger pointed out that the user agent title "Users" should be titled "User Info" because users are outside of the mail management system. This change would keep the document consistent. Further attention was also given to the "Customer Services" section. Under this section, aliases and mail exploders should be modelled because many electronic mail customers demand this service. Harald would like to add an appendix to the document, with examples from popular e-mail routers such as sendmail.

Harald plans to finish the model document by the 28th IETF in Houston, and possibly publish it later as an Informational RFC.

The newest e-mail management document is dated 2 July 1993. Copies of the document and the figures in PostScript can be obtained from nic.switch.ch as the following file names:

```
/e-mail/ifip-emailmgt/docs/emgt-93-model-text.txt
/e-mail/ifip-emailmgt/docs/emgt-93-model-drawing.ps
```

Attendees

| | |
|-----------------------|-----------------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| Harald Alvestrand | Harald.Alvestrand@delab.sintef.no |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Robert Enger | enger@reston.ans.net |
| Urs Eppenberger | eppenberger@switch.ch |
| Tim Goodwin | tim@pipex.net |
| Christoph Graf | graf@switch.ch |
| Jeroen Houttuin | houttuin@rare.nl |
| Marko Kaittola | Marko.Kaittola@funet.fi |

| | |
|--------------------|--|
| Erik Lawaetz | <code>erik.lawaetz@uni-c.dk</code> |
| Robert Reschly | <code>reschly@brl.mil</code> |
| Marjo Rottschaefer | <code>Marjo.Rottschaefer@surfnet.nl</code> |
| Hugh Smith | <code>h.smith@cs.nott.ac.uk</code> |
| Suzanne Smith | <code>smith@es.net</code> |
| Aad van der Voort | <code>a.t.h.vandervoort@research.ptt.nl</code> |
| Paul Vetter | |
| Lea Viljanen | <code>Lea.Viljanen@helsinki.fi</code> |

CURRENT MEETING REPORT

Reported by Marshall Rose/Dover Beach Consulting

Minutes of the NM Area Directorate BOF (NMAREA)

The area director announced that an open meeting for the Network Management Area would be a regular event at each meeting of the IETF. These meetings will consist of an area status report and presentations/discussion on a specific topic of interest to the community.

The area status report is attached. However, two items require specific exposition:

1. Due to a lack of senior technical resources in the Network Management Area, the area director is imposing a moratorium on new working groups for the remainder of calendar year 1993. At the beginning of next year, the area director will evaluate whether the resources exist to allow the formation of new working groups. Although there is much interest in the Network Management Area, the number of volunteers with senior experience in SNMP philosophy and design is quite limited. Many of these volunteers are overloaded in terms of working group chair or working group consultant positions. As such, the area director is unable to allocate resources for new working groups—indeed, some existing working groups are suffering as a result.
2. E-mail to the area director should be addressed as `mrose.iesg@dbc.mtview.ca.us` in order to distinguish between the non-area director and area director roles held by the area director. In other words, people wanting to talk to the area director should send e-mail to `mrose.iesg`; people wanting to talk to Marshall Rose (Theorist, Implementor, and Agent Provocateur) should send mail to `mrose`.

Area Status Report

The topic for this open meeting of the Network Management Area was “SNMPv2 implementation and deployment issues.”

Keith McCloghrie made a presentation on his Internet-Draft “Algorithms for Automating Administration in SNMPv2 Managers,” which reports on one possible implementation strategy.

Steve Waldbusser discussed the current work on the SNMP implementation: an implementation strategy for an ultra-fast MIB compiler, and a TCL-based API for SNMP management applications. He also asked if there would be interest in a CMU-hosted SNMPv2 interoperability testing event during either September or October. Some interest was expressed.

Jeff Case reported on SNMPv2 and RMON demonstrations at ISINM 1993, held in April. Seven vendors demonstrated interoperability of a variety of SNMPv2 features including

manager-to-manager. Of particular interest was the performance boost when using get-bulk. He also discussed implementation experience with the automatic clock synchronization facilities for multi-threaded management applications.

Dave Perkins reported informal research experience with SNMPv2 with respect to size and performance.

Finally, there was general discussion on get-bulk performance and caching strategy.

The area director concluded the meeting by asking the attendees to be ever mindful of SNMPv2 implementation and deployment experiences so that they can contribute to the SNMPv2 evaluation process.

**NETWORK MANAGEMENT AREA:
OPEN MEETING**

Marshall T. Rose
Dover Beach Consulting, Inc.

July 13, 1993

AGENDA

- o NM Area Report
- o SNMPv2 implementation and deployment issues:
presentations and discussion

SNMPv1 FRAMEWORK

- o Core documents (full standards):

| RFC | Name |
|------|-------------------------------------|
| 1155 | Structure of Management Information |
| 1212 | Concise MIB Definitions |
| 1157 | Simple Network Management Protocol |
| 1213 | Management Information Base II |

- o Transport mappings (proposed standards):

| RFC | Name |
|------|---------------------|
| 1418 | SNMP over OSI |
| 1419 | SNMP over AppleTalk |
| 1420 | SNMP over IPX |

SNMPv2 FRAMEWORK

- o Core documents (proposed standards):

| RFC | Name |
|------|---------------------------------------|
| 1441 | Introduction to SNMPv2 |
| 1442 | SMI for SNMPv2 |
| 1443 | Textual Conventions for SNMPv2 |
| 1444 | Conformance Statements for SNMPv2 |
| 1445 | Administrative Model for SNMPv2 |
| 1446 | Security Protocols for SNMPv2 |
| 1447 | Party MIB for SNMPv2 |
| 1448 | Protocol Operations for SNMPv2 |
| 1449 | Transport Mappings for SNMPv2 |
| 1450 | MIB for SNMPv2 |
| 1451 | Manager-to-Manager MIB |
| 1452 | Coexistence between SNMPv1 and SNMPv2 |

Internet MIB Modules (as of 17 June 93)

| Area | Standard | | | Work In Progress |
|---------|----------|------------|---|--|
| | Full | Draft | Proposed | |
| General | MIB-II | | Appletalk DECnet Phase-IV User Identity | DNS Host resources Mail & Directory |
| Media | | Ether-like | Interface extensions RMON 802 bridges, 802.3 repeaters Token Ring, FDDI SIP, DS1, DS3, FR-DTE Character (serial, parallel, ports) PPP (LCP, Security, IP, Bridging) X.25 (PLP, LAPB, MIOX) | Interface layering TRMON MAU 802 source routing ATM FR-Service Modems SNA DLC/NAU Chassis, UPS |
| Routing | | | IP forwarding BGP v3, OSPF v2, RIP v2 | IS-IS IDPR, IDRP |

Acronyms

| | |
|-------|---|
| ATM | Asynchronous Transfer Mode |
| BGP | Border Gateway Protocol |
| CLNS | OSI Connectionless-mode Network Service |
| DNS | Domain Name System |
| FDDI | Fiber Distributed Data Interface |
| IDPR | Inter-Domain Policy Routing Protocol |
| IDRP | OSI Inter-Domain Routing Protocol |
| IS-IS | OSI Intermediate-System to Intermediate-System Protocol |
| LAPB | X.25 link layer |
| LCP | Link Control Protocol |
| MAU | Media Access Unit |
| MIOX | Multiprotocol Interconnect over X.25 |
| OSPF | Open Shortest Path First |
| PLP | X.25 Packet Layer Protocol |
| PPP | Point-to-Point Protocol |
| RIP | Routing Information Protocol |
| RMON | Remote Network Monitoring |
| SIP | SMDS Interface Protocol |
| TRMON | Token Ring RMON |
| UPS | Uninterruptible Power Supply |

- o There are now over 2000 standardized objects
- o There are also over 500 enterprise (proprietary) subtrees registered

NM DIRECTORATE

- Consider strategic evolution of the framework
- Provide architectural/engineering guidance, at the earliest stages of a WG
- Review submitted I-Ds for standards-track evaluation
- No standards-setting power, but is consulted whenever a WG (in any area) defines a MIB module
- Membership:
Fred Baker, Ted Brunner, Jeff Case, Keith McCloghrie, Dave Perkins, Bob Stewart, and Steve Waldbusser

MIB MODULES CURRENTLY UNDER REVIEW

draft-ietf-bridge-sr-objects-02.txt
draft-ietf-dns-resolver-mib-00.txt
draft-ietf-dns-server-mib-00.txt
draft-ietf-hostmib-resources-02.txt
draft-ietf-hubmib-man-02.txt
draft-ietf-hubmib-objects-00.txt
draft-ietf-rmonmib-trmib-01.txt

SNMPv2 AND MIB MODULES

- Now: use SNMPv1 SMI, but pay attention to SNMPv2 issues
- August—draft-standard: use SNMPv2 SMI, but without new object syntaxes
- Thereafter: use SNMPv2 SMI, without restrictions

WORKING GROUPS

- Active:
AToM MIB, Bridge MIB, Chassis MIB, FDDI MIB, FRS MIB, HR MIB, HUB MIB, Interfaces MIB, MADMAN, Modem, SNA DLC, SNA NAU, TRMON, UPS MIB, Character MIB, DECnet Phase IV MIB
- Inactive:
RMON, SNMPv2, Trunk MIB, X.25 MIB
- General SNMP list: snmp@ua.psi.com
- Moratorium on new WGs for CY1993:
lack of senior technical resources!

NM AD's STATEMENT OF DISCLOSURE

- Principal of a consultancy:
50% time for clients, 50% for community service; and, clients neither fund nor direct any community service

- Client list:

| Client | Market Area |
|--------------------------------|-----------------------------|
| North American Directory Forum | Directory services |
| SoftoSwitch | E-Mail & Directory products |
| AT&T Bell Laboratories | Network management services |
| Interop Company | US Program Committee |

- A small number of shares in PSI, otherwise no financial interest in any computer-communications company
- E-mail: mrose.iesg@dbc.atviev.ca.us

2. Implementation Model

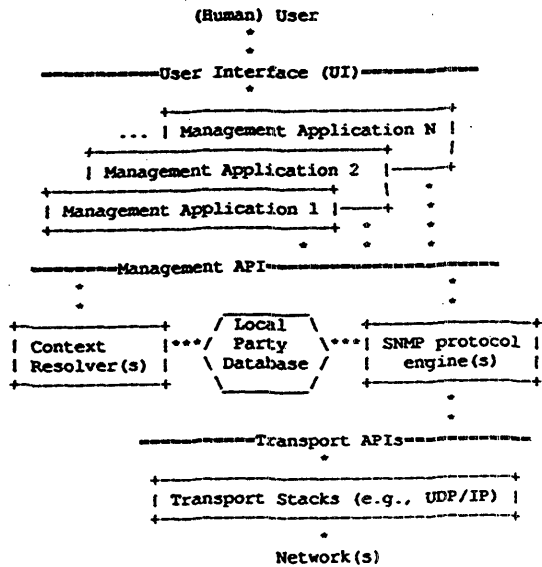


Figure 2.1 SNMPv2 Manager Implementation Model

- Users and applications do NOT know about SNMPv2 parties
- NMS has (logically) one Local Party Database
- Management stack (not applications) deals with:
 - clock synchronization,
 - secret update
 - max message size
- Special application deals with Local Party Database admin

- Application supplies "textual string"
- "Context Resolver" converts into "context handle"
- Application specifies "context handle" and security requirements with operation
- Management stack determines parties and context to be used (or SNMPv1 community, if necessary)
- Conventions for creating noAuth/noPriv entries automatically

- Management stack can dynamically (e.g., at initialization-time) determine parties/contexts for an md5Auth party, to:
 - clock synch,
 - update secret
- Management stack can dynamically determine the need for clock synch or secret update, and perform them without application having to know

- I-D suggests:
 - "context handle" points to context,
 - set of access control entries for a context-handle can be cached
- Alternative suggestion:
 - "context handle" points to set of access control entries.

2.3.1 ATM MIB (ATOMMIB)

Charter

Chair(s):

Kaj Tesink, kaj@cc.bellcore.com

Mailing Lists:

General Discussion: atommib@thumper.bellcore.com

To Subscribe: atommib-request@thumper.bellcore.com

Archive:

Description of Working Group:

The ATOM MIB Working Group is chartered to define sets of managed objects which will be useful in the management of ATM and SONET equipment, interfaces, networks, and/or services that conform to the relevant ATM and SONET specifications. The initial sets defined will be:

- An interface-specific MIB for ATM interfaces, which is aligned with the managed objects for interface layering being defined by the Interfaces MIB Working Group. The working group should consider the ATM Forum's ILMI MIB for its suitability in this respect, plus any extensions necessary to instrument the layers between the ATM layer and the IP layer (e.g., AAL5). The latter should take into account the work of the IP over ATM Working Group (e.g., the "Multi-Protocol over AAL5" specification).
- Managed objects for the monitoring and control of ATM PVCs and SVCs, both in ATM end-points and in ATM switches or networks. (Objects for ATM SVCs will be considered after completion of the work on ATM PVCs.)
- Managed objects that instrument devices with SONET interfaces that conform with the relevant SONET specifications. This work should closely align to other trunk MIBs (DS1/E1 MIB, DS3/E3 MIB). The working group should consider the existing Internet-Draft SONET MIB for its suitability in this respect.

Goals and Milestones:

- | | |
|----------|--|
| Done | Post an Internet-Draft of the ATM and SONET MIB. |
| Dec 1993 | Submit the ATM and SONET MIB to the IESG for consideration as a Proposed Standard. |

Internet-Drafts:

“Definitions of Managed Objects for the SONET/SDH Interface Type”, 06/30/1992,
Tracy Brown, Kaj Tesink <draft-ietf-atommib-sonet-00.txt>

“Definitions of Managed Objects for ATM Management”, 08/09/1993, M. Ahmed,
K. Tesink <draft-ietf-atommib-atm-00.txt>

CURRENT MEETING REPORT

Reported by Kaj Tesink/Bellcore

Minutes of the ATM MIB Working Group (ATOMMIB)

Ted Brunner kindly volunteered to take notes for these minutes.

ATM MIB(s) Work**Status**

A draft has been posted and discussed on the mailing list for some time. The scope is beyond that of ATM Forum's ILMI MIB, which is limited on local interface. A new version of the ILMI MIB will have address registration, and some minor polishing.

It has been suggested to maintain, as much as possible, ILMI compatibility and semantics. However, given the larger scope of the IETF ATM MIB, some differences will be unavoidable.

Discussion followed on Bellcore's proposed draft MIB.

Editor's Note: Not all points of discussion on the ATM MIB(s) are included in this version of the minutes. A complete version is available via FTP or mail server from the remote directories as /ietf/atommib/atommib-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The Need for ATM Local Interface Statistics Per VCC/VPC

The draft MIB proposes for each VC and VP interface a counter for the number of received and transmitted cells, and the number of discards due to traffic policing and shaping. These statistics could, for example, be used to detect congestion and configuration problems. Other suggestions were not to include these statistics at all for VC or VP interfaces, suggesting that hardware costs outweigh the benefits. Still another suggestion was to have a different conformance statement for public and private interfaces (i.e., public interfaces do have these statistics, and private interfaces do not). Keith McCloghrie suggested a compromise, i.e., to define a test capability that can measure these statistics for specific VP and VC interfaces for a short amount of time. Kaj Tesink will produce a draft, which will be discussed on the mailing list.

The Need for Physical Level Convergence Level Management

The current draft MIB specifies managed objects for the DS3 PLCP, and for the SONET TC. The contents of the corresponding tables were reviewed and agreed to. Discussion focused on whether these tables belonged in the ATM MIB or in the DS3 and SONET MIBs respectively. For practical reasons it was decided to keep them in the ATM MIB. Convergence layers for other types of physical facilities were not identified but could be added as needed.

The Need for ATM Connection Management

A more lengthy discussion took place on this subject. In addition, Ken Rodemann gave a presentation on a generic approach to the management of virtual connections, suggesting a common approach for Frame Relay, X.25, and ATM. The generic approach would serve as a sort of umbrella over connection tables that are specific to the X.25, Frame Relay, or ATM. The contents of the specific tables would not be affected by adoption of the generic approach. Rather, the specific approach would simplify the overall management of connections. Discussion of this topic was, due to lack of time, deferred to the FRNETMIB and IFMIB Working Group meetings.

On the specifics of the connection table, the following points were discussed:

- It was agreed that a connection table should work for both end-systems and intermediate-systems.
- The desire to maintain commonality with the ILMI MIB was expressed. It was also observed that the ILMI MIB does not have a connection table (not in its scope), and that the draft table caused only a difference in OID values for traffic parameters.
- The proposed draft makes the point of allowing the creation of a new association between an ingress and egress with a single table row.
- A connection table would benefit considerably from a rowStatus column as defined in the SNMPv2 TC.

Keith McCloghrie and Ted Brunner were tasked to review connection table alternatives, and post their findings to the mailing list for discussion.

The Need for SVC Management

Due to lack of time, network management needs for SVCs were not discussed. However, a discussion took place on the scope of the connection table. In general, the observation was supported that the connection table should not state whether it applies to SVCs and/or PVCs, leaving it to implementation as to how the table is applied.

SONET MIB Work

The issues discussed had been previously raised on the mailing list. In addition, the particular use of ifTable was discussed.

Status

The existing Internet-Draft has been kept highly compatible with other trunk MIBs. Only minor comments have been made on the mailing list.

Editor's Note: Not all points of discussion on the SONET MIB are included in this version of the minutes. A complete version is available via FTP or mail server from the remote directories as /ietf/atommib/atommib-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

The Need for Interval and Total Tables

The mailing list has pointed out that strictly speaking, these tables are redundant, since they can be deduced from the Current and first Interval tables. It was suggested to delete the Total tables, and leave the number of supported Interval tables as implementation specific. One suggestion was made to support at least an hour's worth of Interval tables. Another value that was suggested was eight hours. Given that some implementations of these tables may already exist, confirmation of this approach will be sought on the mailing list.

Attendees

| | |
|-------------------|---------------------------------|
| Masuma Ahmed | mx@sabre.bellcore.com |
| David Arneson | arneson@ctron.com |
| Anders Baardsgaad | anders@cc.uit.no |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Tracy Brown | tacox@mail.bellcore.com |
| Theodore Brunner | tob@thumper.bellcore.com |
| Jeff Case | case@cs.utk.edu |
| Chris Chiotasso | chris@andr.ub.com |
| Jonathan Davar | jdavar@synoptics.comm |
| David Engel | david@ods.com |
| Michael Erlinger | mike@jarthur.claremont.edu |
| David Fresquez | fresquez@vnet.ibm.com |
| Mike Goguen | goguen@synoptics.com |
| Juha Heinanen | juha.heinanen@datanet.tele.fi |

| | |
|----------------------|--------------------------------|
| Steven Horowitz | witz@chipcom.com |
| Jeff Hughes | jeff@col.hp.com |
| Carl Madison | carl@startek.com |
| Andrew Malis | malis@maelstrom.timeplex.com |
| Keith McCloghrie | kzm@hls.com |
| George Mouradian | gvm@arch3.att.com |
| Zbigniew Opalka | zopalka@agile.com |
| David Perkins | dperkins@synoptics.com |
| Drew Perkins | ddp@fore.com |
| James Reeves | jreeves@synoptics.com |
| Kenneth Rodemann | krr@qsun.att.com |
| Dan Romascanu | dan@lannet.com |
| Hal Sandick | sandick@vnet.ibm.com |
| Jon Saperia | saperia@tay.dec.com |
| Jean-Bernard Schmitt | jbs@vnet.ibm.com |
| Dono van-Mierop | dono_van_mierop@3mail.3com.com |
| James Watt | james@newbridge.com |
| Steven Willis | steve@wellfleet.com |

2.3.2 Bridge MIB (BRIDGE)

Charter

Chair(s):

Fred Baker, fbaker@acc.com

Mailing Lists:

General Discussion: bridge-mib@pa.dec.com

To Subscribe: bridge-mib-request@pa.dec.com

Archive:

Description of Working Group:

The Bridge MIB Working Group is chartered to define a set of managed objects that instrument devices that conform to the IEEE 802.1 standard for MAC-layer bridges.

This set of objects should be largely compliant with (and even draw from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

The MIB object definitions produced will be for use by SNMP and will be consistent with other SNMP objects, standards, and conventions.

Goals and Milestones:

- | | |
|------|--|
| Done | Publish initial proposal. |
| Done | Submit an Internet-Draft. |
| Done | Submit draft for RFC publication. |
| Done | Publish a draft revision to RFC 1286 that reflects implementation experience and the result of alignments with IEEE work as an Internet-Draft. |
| Done | Publish a draft SNMP MIB that instruments functions specific to source routed bridges as an Internet-Draft. |
| Done | Submit a draft MIB for source routing bridge functions to the IESG for consideration as a Proposed Standard. |

Internet-Drafts:

“Definitions of Managed Objects for Source Routing Bridges”, 05/05/1993, E. Decker, K. McCloghrie, P. Langille <draft-ietf-bridge-sr-objects-03.txt>

Request For Comments:

RFC 1286 “Definitions of Managed Objects for Bridges”

RFC 1493 “Definitions of Managed Objects for Bridges”

2.3.3 Character MIB (CHARMIB)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

General Discussion: char-mib@decwrl.dec.com

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

Archive:

Description of Working Group:

The Character MIB Working Group is chartered to prepare a recommendation to the IESG evaluating RFCs 1316-1318 (the Character MIBs) with respect to the standards track.

The recommendation will document implementation, interoperability, and deployment experience. If these experiences suggest that changes should be made to the documents, new drafts may be prepared. The recommendation will report one of four outcomes for each RFC:

- That the RFC should be advanced from Proposed to Draft status, without changes (if no problems are found);
- That a draft prepared by the working group should replace the RFC, and be designated a Draft Standard (if only minor changes are made);
- That a draft prepared by the working group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or,
- That the RFC should be designated as Historic (if this technology is problematic).

Goals and Milestones:

- | | |
|------|---|
| Done | Mailing list discussion of Charter and collection of concerns. |
| Done | Discussion and final approval of Charter; discussion on models and terminology. Make writing assignments. |
| Done | First draft document, discussion, additional drafts, special meeting? |
| Done | Review latest draft and if OK, give to IESG for publication as RFC. |
| Done | Reactivation of Working Group to prepare the Character MIBs for Draft Standard. |

Jun 1993 Post an Internet-Draft with the results of the survey of implementation and operational experiences with the Character MIBs. Post revised MIB documents if necessary.

Aug 1993 Submit the Character MIBs to the IESG for consideration as Draft Standards.

Request For Comments:

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

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

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

2.3.4 DECnet Phase IV MIB (DECNETIV)

Charter

Chair(s):

Jonathan Saperia, saperia@tay.dec.com

Mailing Lists:

General Discussion: phiv-mib@pa.dec.com

To Subscribe: phiv-mib-request@pa.dec.com

Archive:

Description of Working Group:

The DECnet Phase IV MIB Working Group is chartered to prepare a recommendation to the IESG evaluating the standards track status of RFC 1289 (the DECnet Phase IV MIB).

The recommendation will document implementation, interoperability, and deployment experience. If this experience suggests that changes should be made to the document, a new draft may be prepared. The recommendation will report one of four outcomes:

- That RFC 1289 should be advanced from Proposed to Draft status, without changes (if no problems are found).
- That a draft prepared by the working group, should replace RFC 1289, and be designated a Draft Standard (if only minor changes are made).
- That a draft prepared by the working group, should replace RFC 1289, and be designated a Proposed Standard (if major changes or feature enhancements are made).
- That RFC 1289 should be designated as Historic (if this technology is problematic).

Goals and Milestones:

- | | |
|----------|---|
| Done | Re-activate the Working Group to advance the Decnet Phase IV MIB to Draft Standard. |
| Done | Post an Internet-Draft of the results from the survey of implementation and operational experience. Post a revised version of the MIB if necessary. |
| Oct 1993 | Submit the DECnet Phase IV MIB to the IESG for consideration as a Draft Standard. |

Internet-Drafts:

“DECnet Phase IV MIB - Implementation Report”, 06/21/1993, J. Saperia
<draft-ietf-decnetiv-mib-implement-00.txt>

“DECnet Phase IV MIB Extensions”, 06/25/1993, J. Saperia <draft-ietf-decnetiv-
mibext-00.txt>

Request For Comments:

RFC 1289 “DECnet Phase IV MIB Extensions”

2.3.5 FDDI MIB (FDDIMIB)

Charter

Chair(s):

Jeffrey Case, case@cs.utk.edu

Mailing Lists:

General Discussion: fddi-mib@CS.UTK.EDU

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

Archive:

Description of Working Group:

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

Goals and Milestones:

- Done "Final" initial draft of required get/set variables.
- Done Initial implementations of required get/set variables.
- Done Revised "final" draft of required get/set variables.
- Done Adoption of draft of required get/set variables.
- Mar 1992 Submit the FDDI MIB to the IESG for consideration as a Proposed or Draft Standard depending on the magnitude of changes to RFC1285.
- Done Hold a meeting at the November IETF Plenary.
- Dec 1992 Post an Internet-Draft aligned with current the current ANSI document factoring in implementation experience with RFC 1285.

Request For Comments:

RFC 1285 "FDDI Management Information Base"

RFC 1512 "FDDI Management Information Base"

CURRENT MEETING REPORT

Reported by Jeff Case/University of Tennessee

Minutes of the FDDI MIB Working Group (FDDIMIB)

The status of the current MIB Internet-Draft was discussed. It is progressing through the Last Call, Directorate review, and IESG review with few comments, none of which are substantive; all are editorial in nature. If all continues as anticipated, the document will soon become an RFC with Proposed Standard status.

RFC 1285 will continue to be a Proposed Standard for management of ANSI 6.2 based FDDI network devices; the new RFC will be for the management of ANSI 7.x based FDDI network devices. RFC 1285 will expire at some point in the future but does not need to be dealt with at this time.

The majority of the (short) meeting time was devoted to consideration of the trap document. In a repetition of history (Atlanta IETF, July 30, 1991), the net result was that there is no need for a trap document. One reason for this is the observation that the Alarms group of the RMON MIB can be implemented in the RMON device and thereby provide all of the functions of all of the event notifications contained in the ANSI document. The working group decided at the Columbus IETF meeting that there would need to be standard mechanisms for enabling, disabling, and rate controlling any traps to be generated. Since the RMON MIB provides all of these capabilities and is already progressing on the standards track, the group concluded that there is no need for additional work on traps.

Consequently, having concluded its charter, the work of the group is finished. Members should anticipate the working group going dormant in the near future, to be re-chartered when there are new work items.

Attendees

| | |
|----------------------|---------------------------|
| David Arneson | arneson@ctron.com |
| Jeff Case | case@cs.utk.edu |
| Chris Chiotasso | chris@andr.ub.com |
| David Engel | david@ods.com |
| Phil Irely | pirey@relay.nswc.navy.mil |
| Ron Mackey | rem@dsiinc.com |
| Carl Madison | carl@startek.com |
| David Perkins | dperkins@synoptics.com |
| James Reeves | jreeves@synoptics.com |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Jean-Bernard Schmitt | jbs@vnet.ibm.com |

2.3.6 Frame Relay Service MIB (FRNETMIB)

Charter

Chair(s):

James Watt, james@newbridge.com

Mailing Lists:

General Discussion: frftc@nsco.network.com

To Subscribe: frftc-request@nsco.network.com

Archive:

Description of Working Group:

The Frame Relay Service MIB Working Group is chartered to define an initial set of managed objects which will be useful for customer network management of a provider's Frame Relay Service. The working group will consider existing definitions, including the Frame Relay Forum's work in this area. The objects defined by the working group will be consistent with the SNMP framework.

The working group will coordinate with both the Frame Relay Forum and the ATM MIB Working Group.

Goals and Milestones:

- | | |
|----------|--|
| Done | Post the initial Internet-Draft for discussion. |
| Dec 1993 | Submit the Frame Relay Service MIB to the IESG for consideration as a Proposed Standard. |

Internet-Drafts:

"Definitions of Managed Objects for Frame Relay Service", 05/13/1993, T. Brown <draft-ietf-frnetmib-fr-02.txt>

"Service Management Architecture for Virtual Connection Services", 06/15/1993, K. Rodemann <draft-ietf-frnetmib-virtual-sma-01.txt>

CURRENT MEETING REPORT

Reported by James W. Watt/Newbridge Networks Corporation

Minutes of the Frame Relay Service MIB Working Group (FRNETMIB)

Documents Discussed

- Service Management Architecture for Virtual Connection Services (1 July)
- Definitions of Managed Objects for Frame Relay Services (2 July)

Issues Discussed

- Frame Relay Forum (FRF)

Andy Malis, Chair of the FRF, presented the work plan for the FRF effort. The start and end dates of the FRF effort match those of the FRNETMIB Working Group and there are some intermediate meetings.

The goal shared by both groups is to have a standard MIB by December of 1993.

- Service Management Architecture

Ken Rodemann gave a presentation based on the 2 July draft of his document. The general reaction was positive, and in particular the “(virtual-)circuit table” was identified as a work item to be discussed with the ATOMMIB and IFMIB Working Groups.

Discussion of the rest of the Service Management Architecture document will continue in the FRNETMIB Working Group.

- Draft MIB

Based on discussions of the 2 July draft, consensus was reached on the following:

- The PVC table was split into a “flow” (or connection) table and an end-point table.
- The per-PVC InUnits and OutUnits counters were renamed to InOctets and OutOctets for consistency with other MIBs and frPVCUnitSize was renamed to (something like) OctetResolution (see action item #2).
- The LPortNetAddress object was dropped as a duplicate of ifPhysAddr.
- A reference to ifPhysAddress was placed in the description of LPortNumPlan.

- The frPVConnectStatus trap needs to have words to prevent flooding a manager with traps when a whole link fails. There is an added complexity with respect to the modeling of NNI links (see action item #3).
- The next draft of the MIB will be in SNMPv2.
- The focus of the MIB is primarily Customer Network Management (CNM) of PVCs for Frame Relay; however, it may be useful for other purposes.

Action Items

1. Tracy Brown to issue a revised MIB draft (in SNMP V2 format) prior to the August FRF meeting.
2. Tracy Brown to find consensus on where "unit size" should be in the MIB (end-point table or logical port table).
3. Tracy Brown, Ken Rodemann and James Watt to clarify wording on trap suppression in MIB draft.

Attendees

| | |
|----------------------|------------------------------|
| Masuma Ahmed | mxasabre@bellcore.com |
| David Arneson | arneson@ctron.com |
| Fred Baker | fbaker@acc.com |
| Caralyn Brown | cbrown@wellfleet.com |
| Tracy Brown | tacox@mail.bellcore.com |
| Theodore Brunner | tob@thumper.bellcore.com |
| Steve Buchko | stevebu@newbridge.com |
| Chris Chiotasso | chris@andr.ub.com |
| David Fresquez | fresquez@vnet.ibm.com |
| Andrew Malis | malis@maelstrom.timeplex.com |
| Keith McCloghrie | kzm@hls.com |
| George Mouradian | gvm@arch3.att.com |
| Daniel Myers | dan@nsd.3com.com |
| Zbigniew Opalka | zopalka@agile.com |
| David Perkins | dperkins@synoptics.com |
| Aiko Pras | pras@cs.utwente.nl |
| Dan Romascanu | dan@lannet.com |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Jean-Bernard Schmitt | jbs@vnet.ibm.com |
| Tibor Schonfeld | tibor@qsun.att.com |
| Kaj Tesink | kaj@cc.bellcore.com |

Dono van-Mierop
James Watt
Peter Wilson

dono_van_mierop@3mail.3com.com
james@newbridge.com
peter_wilson@3mail.3com.com

2.3.7 Interfaces MIB (IFMIB)

Charter

Chair(s):

Ted Brunner, tob@thumper.bellcore.com

Mailing Lists:

General Discussion: if-mib@thumper.bellcore.com

To Subscribe: if-mib-request@thumper.bellcore.com

Archive:

Description of Working Group:

The Interfaces MIB Working Group is chartered to accomplish two tasks.

First, to develop a collection of managed objects which model the relation between different entities in the data link and physical layers. The working group will explore different modeling approaches in order to develop a collection of objects which is both correct in the modeling sense and has an acceptable impact (if any) on the interfaces table from MIB-II and all media MIB modules on the standards track or under development by a working group. The objects defined by the working group will be consistent with the SNMP framework.

Second, to prepare a recommendation to the IESG evaluating RFC 1229 (the interface-extensions MIB), RFC 1231 (the token-ring MIB), RFC 1304 (the SMDS MIB), and RFC 1398 (the ethernet-like MIB) with respect to the standards track.

The recommendation will document implementation, interoperability, and deployment experience. If these experiences suggest that changes should be made to the documents, new drafts may be prepared.

For RFCs 1229, 1231, and 1304, the recommendation will report one of four outcomes for each RFC:

- that the RFC should be advanced from Proposed to Draft status, without changes (if no problems are found);
- that a draft prepared by the working group should replace the RFC, and be designated a Draft Standard (if only minor changes are made);
- that a draft prepared by the working group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or,
- that the RFC should be designated as Historic (if this technology is problematic).

For RFC 1398, the recommendation will report one of five outcomes:

- that the RFC should be advanced from Draft to Full status, without changes (if no problems are found);
- that a draft prepared by the working group should replace the RFC, and be designated a Standard (if only editorial changes are made);
- that a draft prepared by the working group should replace the RFCs, and be designated a Draft Standard (if only minor changes are made);
- that a draft prepared by the working group should replace the RFC, and be designated a Proposed Standard (if major changes or feature enhancements are made); or,
- that the RFC should be designated as Historic (if this technology is problematic).

Goals and Milestones:

- | | |
|----------|---|
| Done | Post the interface layering document as an Internet-Draft. |
| Sep 1993 | Submit the interface layering document to the IESG for consideration as a Proposed Standard. |
| Sep 1993 | Issue a call for implementation and operations experience with RFCs 1229, 1231, 1304, and 1398. |
| Oct 1993 | Evaluate experience and if necessary post revised MIBs as Internet-Drafts. |
| Dec 1993 | Submit recommendations on the various MIBs to the IESG. |

Internet-Drafts:

“Evolution of the Interfaces Group of MIB-II”, 06/04/1993, K. McCloghrie, F. Kastenholz <draft-ietf-ifmib-evolution-02.txt>

CURRENT MEETING REPORT

Reported by Theodore Brunner/Bellcore

Minutes of the Interfaces MIB Working Group (IFMIB)

The group's discussions were based on the "Evolution of the Interfaces Group of MIB-II" Internet-Draft. The version dated 1 June was discussed first, followed by discussion of additional issues in the 28 June version. The discussion was a continuation of issues raised on the mailing list.

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

Although only one session had originally been scheduled for IFMIB, a second session was scheduled ad hoc. The only other group from the Network Management Area whose schedule conflicted with the second session was UPSMIB.

The group did not have time to discuss issues such as a generic approach to virtual connections. However, this particular subject was also expected to be addressed at the meetings for the FRNETMIB and ATOMMIB Working Groups.

Attendees

| | |
|------------------|-------------------------------|
| Masuma Ahmed | mxsa@sabre.bellcore.com |
| David Arneson | arneson@ctron.com |
| Fred Baker | fbaker@acc.com |
| Jim Barnes | barnes@xylogics.com |
| Caralyn Brown | cbrown@wellfleet.com |
| Tracy Brown | tacox@mail.bellcore.com |
| Theodore Brunner | tob@thumper.bellcore.com |
| Jeff Case | case@cs.utk.edu |
| Chris Chiotasso | chris@andr.ub.com |
| Dave Cullerot | cullerot@ctron.com |
| David Engel | david@ods.com |
| David Fresquez | fresquez@vnet.ibm.com |
| Steven Horowitz | witz@chipcom.com |
| Frank Kastenholz | kasten@ftp.com |
| Andrew Knutsen | andrewk@sco.com |
| Deirdre Kostick | dck2@mail.bellcore.com |
| Kent Malave | kent@bach.austin.ibm.com |
| Keith McCloghrie | kzm@hls.com |
| Henri Moelard | henri.moelard@utrecht.ncr.com |

| | |
|----------------------|--|
| George Mouradian | <code>gvm@arch3.att.com</code> |
| Kenneth Rodemann | <code>krr@qsun.att.com</code> |
| Dan Romascanu | <code>dan@lannet.com</code> |
| Rick Royston | <code>rroyston@usr.com</code> |
| Jon Saperia | <code>saperia@tay.dec.com</code> |
| Jean-Bernard Schmitt | <code>jbs@vnet.ibm.com</code> |
| Kitty Shih | <code>kmshih@novell.com</code> |
| Timon Sloane | <code>timon@timon.com</code> |
| Suzanne Smith | <code>smith@es.net</code> |
| Kaj Tesink | <code>kaj@cc.bellcore.com</code> |
| Mario Vecchi | <code>mpv@thumper.bellcore.com</code> |
| Steven Waldbusser | <code>waldbusser@andrew.cmu.edu</code> |
| Scott Wasson | <code>sgwasson@eng.xyplex.com</code> |
| James Watt | <code>james@newbridge.com</code> |

2.3.8 Mail and Directory Management (MADMAN)

Charter

Chair(s):

Steve Kille, S.Kille@isode.com

Mailing Lists:

General Discussion: madman@innosoft.com

To Subscribe: mailserv@innosoft.com

In Body: subscribe ietf-madman <email address>

Archive: innosoft.com:~/ietf-madman/archive.txt

Description of Working Group:

The Mail and Directory Management Working Group is chartered to define four MIB modules: one for generic application monitoring, one for message relays (either SMTP or X.400 based), one for OSI Directory service (X.500), and a fourth for message stores. The MIB modules will provide basic monitoring capabilities, and will be consistent with the SNMP framework and existing SNMP standards.

Goals and Milestones:

- Done Post an Internet-Draft of the generic application monitoring MIB.
- Done Post an Internet-Draft of the message relay monitoring MIB.
- Done Post an Internet-Draft of the OSI X.500 Directory Service MIB.
- Jul 1993 Submit the directory monitoring, message relay, and generic application MIBs to the IESG for consideration as Proposed Standards.
- Jul 1993 Post an Internet-Draft of the message store monitoring MIB.
- Nov 1993 Submit the message store monitoring MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“Network Services Monitoring MIB”, 06/01/1993, N. Freed, S. Kille <draft-ietf-madman-networkmib-05.txt>

“Mail Monitoring MIB”, 06/21/1993, N. Freed, S. Kille <draft-ietf-madman-
mtamib-05.txt>

“Directory Monitoring MIB”, 06/28/1993, G. Mansfield, S. Kille <draft-ietf-
madman-dsa-mib-05.txt>

CURRENT MEETING REPORT

Reported by Urs Eppenberger/SWITCH

Minutes of the Mail and Directory Management Working Group (MADMAN)

The minutes of the MADMAN BOF, which met during the March 1993 IETF in Columbus were accepted nem con. The group then moved on to document review.

Network Services MIB

An application uptime counter will be added since it does not match with the system uptime due to application restarts. Identifying TCP/IP applications will be factored out to a separate document. The MIB will be submitted as a standards-track RFC.

Mail Monitoring MIB

The volume measured is user data above the message transfer protocol. The MIB will be submitted as a standards-track RFC.

Directory Monitoring MIB

Traps have been removed in the newest draft. The Navy is concerned about this since they wanted to get warned by the system directly without watching it all the time. The MIB will be submitted as a standards-track RFC, following online discussion.

Message Store MIB

It should be discussed on the list whether the group wants to continue this activity. The reasons given in Columbus for taking this up are still valid but there have been no volunteers to work on it. It will be dropped unless a document editor is found.

Attendees

| | |
|-----------------|--|
| Robert Cooney | cooney@wnyose.nctsw.navy.mil |
| Urs Eppenberger | eppenberger@switch.ch |
| Christoph Graf | graf@switch.ch |
| Jan Hansen | Jan.Hansen@teknologi.agderforskning.no |

| | |
|---------------------|-----------------------------|
| Anders Karlsson | sak@cdg.chalmers.se |
| Steve Kille | S.Kille@isode.com |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Arnold Krechel | krechel@gmd.de |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Brian May | Brian.May@mel.dit.csiro.au |
| Linda Millington | l.millington@noc.ulcc.ac.uk |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Kenneth Rossen | kenr@shl.com |
| Marjo Rottschaefer | |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Heiner Schorn | heiner.schorn@umdac.umu.se |
| Suzanne Smith | smith@es.net |
| Paul Vetter | |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Peter Yee | yee@atlas.arc.nasa.gov |

2.3.9 Modem Management (MODEMMGT)

Charter

Chair(s):

Mark Lewis, Mark.S.Lewis@telebit.com

Mailing Lists:

General Discussion: modemmgmt@Telebit.com

To Subscribe: majordomo@Telebit.com

In Body: subscribe modemmgmt <email address>

Archive: ftp.telebit.com:~/pub/modemmgmt

Description of Working Group:

The Modem Management Working Group is chartered to define a MIB module for dial-up modems and similar dial-up devices. This MIB module will provide a set of objects that are the minimum necessary to provide the ability to monitor and control those devices, and will be consistent with the SNMP framework and existing SNMP standards.

The working group will consider existing specifications including the RS-232-like, Character, PPP and other related MIB modules. It will consider enterprise-specific MIB modules which support modem-like devices. The working group will also consider the TSB Study Group 14's work on an OSI CMIS/CMIP object definition for V series DCEs entitled "Managed Object Template for V-Series DCE's."

Goals and Milestones:

Jun 1993 Post an Internet-Draft of the Modem Management MIB.

Oct 1993 Submit the Modem Management MIB to the IESG for consideration as a Proposed Standard.

INTERIM MEETING REPORT

Reported by Thomas Holodnik/Carnegie Mellon University

Minutes of the Modem Management Working Group (MODEMMGT)

A meeting of the Modem Management Working Group was held in Baltimore, June 28-29.

Agenda

- Outline
- Introductions
- Broad issues
- Structure of the MIB
- Detailed review
- Next meeting

Editor's Note: Not all agenda item summaries are included in this version of the minutes. A complete version is available via FTP or mail server from the remote directories as /ietf/modemmgmt/modemmgmt-minutes-93jun.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Action Items and Other Significant Points of Discussion

The attendees spent considerable time reducing the number of MIB objects deemed to be of little use, while some conveyed additional information that many felt was omitted in the initial MIB.

- A general statement is needed about vendors that may not support all values in the range specified. While many vendors may not support all values in the range specified, they will still be considered compliant with the MIB. RFC 1444 (SNMPv2 SMI) specifies that the range of the SYNTAX clause specifies the range that the variable may take which makes sense within the protocol.
- A general statement is needed indicating that many settings are country-specific (i.e., that many are set according to national standards). It may be permissible to change certain settings in one country but not another. Certain features may only be useful in one country but not another. Transmitter-level setting changes are illegal in some countries. This needs to be noted in the MIB. Regulatory agencies take precedence over what is allowed in the MIB.

- Is it desirable to manage remote modems via SNMP and modem proxy agents? That is, in addition to managing the chassis modem, you may also need to manage the stand-alone modem (via SNMP). The mechanics for this were left unresolved.
- Progress reports on V.id and V.Fast developments are needed. V.Fast will require additional objects or changes to existing objects. A list of these should be kept to give V.Fast adequate treatment. Power level adjustments are permitted under V.Fast.
- The connect failure reasons will need to be edited (reduced to a manageable size).
- The last call statistics group needs to be split into call statistics (to be renamed) and signal converter statistics. It was decided that many statistics kept over the last call were not terribly useful, but that some information should be kept for summary and reporting purposes.
- The list of MIB variables that need to be included in the call records should be developed further. The architecture for call records needs to be clearly developed.
- The modem MIB is to be experimental subtree 49.
- A good way to measure throughput is needed. Offered load is a strong factor in determining this. There is no agreement on this yet.
- The number of MIB variables and the time to construct the list of MIB variables, for the leased-sync area and the dial-asynch area, need to be carefully managed. If either of the two areas bog down or blow up, it should be jettisoned into another MIB.

Attendees

| | |
|-------------------|---------------------------|
| Jay Bain | jbain@uds.mot.com |
| Les Brown | brown_1@msm.cdx.mot.com |
| Ted Brown | ewb@telebit.com |
| Alan Clark | aclark@hayes.com |
| Thomas Holodnik | tjh@andrew.cmu.edu |
| Mark Lewis | Mark.S.Lewis@telebit.com |
| James Logan | logan@penril.com |
| Chris Payson | c_payson@telebit.com |
| Gregory Pearson | |
| Bill Richards | richards@gdc.com |
| Rick Royston | rroyston@usr.com |
| Erik Snoek | sdrierik@diamond.sara.nl |
| Richard Stuart | dickstuart@attmail.com |
| Steven Waldbusser | waldbusser@andrew.cmu.edu |

CURRENT MEETING REPORT

Reported by Mark Lewis/Telebit Corporation

Minutes of the Modem Management Working Group (MODEMMGT)

Agenda

- Introductions
- Working group goals and schedule
- Status of draft modem mib
- Comments on draft modem mib
- The next step

Introductions

There were 18 people present, several of whom said they were interested in implementing a modem MIB in their products. Also present were a few modem users, a couple of modem manufacturers, and members of the Network Management Directorate (NMDIR).


Working Group Goals and Schedule

It was discussed that this is the group's first meeting as a working group at an IETF meeting. A special two-day meeting had been held June 28-29 in Baltimore.

The major goal of this working group is to agree on a standard MIB for managing modems. The group would like to have an Internet-Draft written and accepted by the group by the end of 1993.

Architecture of the MIB

We discussed the idea of representing dial-up (switched), leased-line, and network modems as separate but related modem instances (model 2 below). In the diagram below, base objects refer to those which are the same regardless of mode (e.g. modem manufacturer). Common objects are those which may have different values for each mode (e.g. transmit output level). Switched and leased refer to groups of objects that are only relevant to that mode.



Model 1 presents the modem as a single instance with multiple instances of variables which might have different values for different modes (common). Model 2 presents multiple modem instances, one for each mode that the modem supports (e.g. switched, leased-line). For both models, there would be a single instance of the objects related to the specific mode (e.g. switched, leased-line).

We discussed the trade-offs between the two models above. Model 1 seemed better if the number of common objects were relatively small compared to the number of base objects. If the number were relatively large, model 2 seemed preferable.

We were unclear how many common objects might have different values for each mode. It was roughly estimated that somewhere between 10-80 objects would fall into this common category. It was noted that many implementations probably don't care to have different values for each mode.

Since the number of such objects is potentially large, there was general agreement to use model 2. This means management stations would deal with multiple instances, one for each mode the modem supports (e.g. switched, leased-line). (Note that after the meeting more detailed analysis was done which may indicate model 1 is more appropriate.)

```

MdmLineCapabilitiesEntry ::= SEQUENCE {
    mdmLineCapabilitiesIndex      INTEGER,
    mdmLineCapabilitiesID        OBJECT IDENTIFIER,
    mdmLineCapabilitiesEnableRequested  INTEGER,
    mdmLineCapabilitiesEnableGranted   INTEGER
}

```

We reviewed the capabilities table. There was general agreement that this fit the situation well. It provided a flexible method to let the agent describe the capabilities of the modem, as well as provide a way to enable and disable them.

Someone voiced a request that there be a linkage between an interface and a modem. There was agreement that this would be valuable where the modem was being used for packet connections using SLIP or PPP. It was also agreed that this would not be possible in cases

where the modem was being used in character only mode. The group resolved to coordinate with the working group designing the interface MIB to provide such a linkage.

There was a lengthy discussion of the idea of providing a record of calls for accounting and trouble-shooting. The advisability of using SNMP for accounting was considered. Since virtually all modem vendors provide such capabilities, it was decided to implement some method of tracking calls. Traps were deemed not suitable for this purpose. It was agreed that some type of a history of calls would be kept in the agent.

Several possible implementations of a call history were considered:

- Option 1

Rely on the management station to poll the agent often enough to get all call records. No traps were necessary using this approach.

- Option 2

Have the management station poll the agent, but also receive traps at a predefined point. For example, the agent would send traps after writing some 25 out of 100 call records. Note the predefined point could be configurable as a percentage.

- Option 3

Have the agent keep track of the last call record read by each management station. It would then be possible to send a trap to a particular management station when it is in danger of missing a call record. This assumes the same polling by the management of the agent.

Some of the trade-offs of these options were considered. Option 1 doesn't use traps but requires a high enough poll frequency (or a large amount of memory in the agent) to minimize the loss of call records. Options 2 and 3 improve reliability, and differ in their complexity.

We considered the 30 some events which were included in the draft. There was strong objection to defining 30 traps. The group thought the call history record provided an adequate record of the event. It was decided not to define traps for events.

The Next Step

The group hopes to produce the Internet-Draft by the end of July. At that point, it will be subject to review on the working group mailing list.

Attendees

| | |
|-------------------|------------------------------|
| Toshiya Asaba | asaba@iiij.ad.jp |
| John Ballard | jballard@microsoft.com |
| Jim Barnes | barnes@xylogics.com |
| Jeff Case | case@cs.utk.edu |
| Michael Erlinger | mike@jarthur.claremont.edu |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Mark Lewis | Mark.S.Lewis@telebit.com |
| Kim Chr. Madsen | kimcm@ic.dk |
| George Mouradian | gvm@arch3.att.com |
| David Perkins | dperkins@synoptics.com |
| John Plate | plate@ic.dk |
| Lars Poulsen | lars@cmc.com |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Rick Royston | rroyston@usr.com |
| Jon Saperia | saperia@tay.dec.com |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| Timon Sloane | timon@timon.com |
| Steven Waldbusser | waldbusser@andrew.cmu.edu |

2.3.10 Remote LAN Monitoring (RMONMIB)

Charter

Chair(s):

Mike Erlinger, mike@jarthur.claremont.edu

Mailing Lists:

General Discussion: rmonmib@jarthur.claremont.edu

To Subscribe: rmonmib-request@jarthur.claremont.edu

Archive: [jarthur.claremont.edu:/pub/rmon](http://jarthur.claremont.edu/pub/rmon)

Description of Working Group:

The RMON Working Group is chartered to prepare a recommendation to the IESG evaluating RFC 1271 (the RMON MIB) with respect to the standards track.

The recommendation will document implementation, interoperability, and deployment experience. If this experience suggests that changes should be made to the document, a new draft may be prepared. The recommendation will report one of four outcomes:

- that RFC 1271 should be advanced from proposed to draft status, without changes (if no problems are found);
- that a draft prepared by the working group, should replace RFC 1271, and be designated a draft standard (if only minor changes are made);
- that a draft prepared by the working group, should replace RFC 1271, and be designated a proposed standard (if major changes or feature enhancements are made); or,
- that RFC 1271 should be designated as historic (if this technology is problematic).

Goals and Milestones:

Sep 1993 Re-activation of WG, call for discussion of experiences.

Nov 1993 Meet at IETF to classify and evaluate experiences.

Apr 1994 Submit recommendation, possibly with new draft, to IESG.

Request For Comments:

RFC 1271 "Remote Network Monitoring Management Information Base"

2.3.11 SNA DLC Services MIB (SNADLC)

Charter

Chair(s):

Jeff Hilgeman, jeffh@apertus.com

Mailing Lists:

General Discussion: snadlcmib@apertus.com

To Subscribe: snadlcmib-request@apertus.com

Archive:

Description of Working Group:

The SNA DLC Working Group is chartered to define a set of managed objects for the SDLC and LLC-2 data link controls for SNA networks. These objects will be the minimum necessary to provide the ability to monitor and control those devices, providing fault, configuration, and performance management, and will be consistent with the SNMP framework and existing SNMP standards.

The working group will consider existing enterprise-specific MIB modules that define objects which support management of these devices. The group may choose to consider any work done by the IEEE in the area of managed object definition for LLC-2. It will also make sure that its work is aligned with the SNA NAU Services MIB Working Group, due to the close relationship between the devices being worked on by the two groups.

The working group recognizes that managed objects for other SNA data link controls and related components (e.g., QLLC, System/370 Channel, Data Link Switching, and ESCON) may need to be identified in the future. These objects are out of scope for the current charter; however, once the Group completes its charter, a new charter identifying some or all of these components may be considered.

Goals and Milestones:

- | | |
|----------|--|
| Done | Mailing List discussion of vendor proprietary MIBs. |
| Jul 1993 | Post an Internet-Draft of the SNA DLC MIB. |
| Dec 1993 | Submit the SNA DLC MIB to the IESG for consideration as a Proposed Standard. |

2.3.12 SNA NAU Services MIB (SNANAU)

Charter

Chair(s):

Zbigniew Kielczewski, zbig@eicon.qc.ca
Deirdre Kostick, dck2@mail.bellcore.com

Mailing Lists:

General Discussion: snanaumib@thumper.bellcore.com
To Subscribe: snanaumib-request@thumper.bellcore.com
Archive: thumper.bellcore.com:pub/tob/snanaumib

Description of Working Group:

The SNA NAU Services MIB Working Group is chartered to define a set of managed objects for PU type 2.0, and LU type 1, 2, and 3 devices for SNA networks. These objects will be the minimum necessary to provide the ability to monitor and control those devices, providing fault, configuration, and performance management, and will be consistent with the SNMP framework and existing SNMP standards.

The working group will consider existing enterprise-specific MIB modules that define objects which support management of these devices. It will also make sure that its work is aligned with the SNA DLC Services MIB Working Group, due to the close relationship between the devices being worked on by the two groups.

The working group recognizes that managed objects for other components (e.g., PU Type 4, PU Type 5, LU Types 1, 3, 4, 6.2 (APPC), APPN EN, APPN NN and APPI) may need to be identified in the future. These objects are out of scope for the current charter; however, once the group completes its charter, a new charter identifying some or all of these components may be considered.

Goals and Milestones:

- Jul 1993 Begin discussion of proprietary MIBS and develop a single proposal.
- Done Post an Internet-Draft of the SNA NAU Services MIB.
- Dec 1993 Submit the SNA NAU Services MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“Definitions of Managed Objects for SNA NAUs”, 08/02/1993, Z. Kielczewski, K. Shih <draft-ietf-snanau-snamib-00.txt>

CURRENT MEETING REPORT

Reported by Deirdre Kostick/Bellcore

Minutes of the SNA NAU Services MIB Working Group (SNANAU)

The SNANAU Working Group met on July 14 to review the proposed SNA NAU MIB draft distributed to the mailing list on June 30.

Changes To MIB Draft

There was a lack of feedback on the MIB draft via the mailing list. This, as well as the need to get feedback, was discussed.

Editor's Note: A list of changes to the document is available via FTP or mail server from the remote directories as /ietf/snanau/snanau-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Interim Meeting

An interim meeting may be scheduled to coincide with the APPN Implementor's Workshop (AIW) in September. The purpose of the interim meeting will be to review the planned Internet-Draft and to increase exposure of the MIB to encourage more feedback from the SNA community.

Action Item Summary

- The working group will be polled to determine if there is interest in having an interim meeting.
- The working group will be polled to gather feedback on the items listed in these minutes.
- The draft will be updated to reflect changes discussed during the meeting.
- Questions on mandatory/optional for the write capability and the experimental branch assignment will be investigated, and results reported back to the working group.
- An Internet-Draft will be posted by July 30 based on the MIB draft plus the changes discussed at the IETF meeting.

Attendees

| | |
|----------------------|------------------------|
| Zbigniew Kielczewski | zbig@eicon.qc.ca |
| Deirdre Kostick | dck2@mail.bellcore.com |
| David O'Leary | doleary@cisco.com |
| David Perkins | dperkins@synoptics.com |
| Kitty Shih | kmshih@novell.com |
| Ed Stern | els@proteon.com |

2.3.13 Token Ring Remote Monitoring (TRMON)

Charter

Chair(s):

Michael Erlinger, mike@jarthur.claremont.edu

Mailing Lists:

General Discussion: rmonmib@lexcel.com

To Subscribe: rmonmib-request@lexcel.com

Archive:

Description of Working Group:

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

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

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

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

Goals and Milestones:

- | | |
|----------|---|
| Done | Discussion and agreement on models and terminology. Comparison of RMON architecture and Token Ring requirements. Assign author and editor responsibilities. |
| Done | Working Group meeting at San Diego IETF. |
| Mar 1992 | Post Internet-Draft of the Token Ring Monitoring MIB. |
| Done | Working Group meeting at Cambridge IETF. |
| Nov 1992 | Submit the Token Ring MIB to the IESG as a Proposed Standard. |

Request For Comments:

RFC 1513 "Token Ring Extensions to the Remote Network Monitoring MIB"

CURRENT MEETING REPORT

Reported by Michael Erlinger/Harvey Mudd College

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

Introductions

A small group of individuals was present, so we went around the room giving short personal introductions.

TRMON Status

A final call was made for comments. The comments received were incorporated into the next draft, which is available as an Internet-Draft (draft-ietf-rmonmib-trmib-01.txt). This document was also forwarded to the Network Management Area Director with a recommendation that it become the TRMON Proposed Standard. At the working group meeting, it was pointed out that the Network Management Area Directorate has started looking at the document and that one action item is to ensure that there are no conflicts with the Token Ring MIB (RFC 1231). Once the Network Management Area Directorate is satisfied with the document, it will be forwarded to the IESG.

Token Ring IEEE Efforts and Relationship to TR RMON

The IEEE 802.5 committee is working on a new draft. The effort is directed towards clarification and explanation of the existing standard (in particular, clarification efforts will center on: MAC state tables, MAC insertion, and various timers). The 802.5 committee is meeting the week of 12 July 93 to discuss comments on draft three. Draft four should be available in September 93, with a conclusion of the process scheduled for March 94. Anyone interested in being on the IEEE 802.5 mailing list as an observer needs to contact Mike Erlinger. Whatever changes are made in 802.5 would be considered in a Draft Standard TRMON.

RFC 1271 Advancement Process

Once TRMON is released as a Proposed Standard, the TRMON Working Group will be disbanded and the RMONMIB Working Group reformed for consideration of RFC 1271 advancement. A strawman charter was presented. *Editor's Note: A draft of the RMONMIB Charter is available via FTP or mail server from the remote directories as /ietf/trmon/trmon-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.* Three

discussions ensued; the first concerned the milestone dates. It was decided that the date for the working group recommendation should be March 94, since November 93 is not realistic.

In the second discussion, it was agreed that the group should attempt to forward a recommendation that follows the second possibility in the charter, e.g., a new draft with minor modifications.

The third discussion area was a first pass list of RFC 1271 concerns. The following list presents only the areas for working group discussion (while the small group had various “solutions,” the effort was directed only at listing areas of concern). Once the RMONMIB Working Group is reconstituted, there will be an expanded discussion of RFC 1271 concerns. Thus, do not consider the list as final—just pet peeves of the attendees.

- Row Status Versus Entry Status

A long discussion ensued concerning pay now, pay later concerns for making this change. There was loose consensus in the group that since there is a desire for an RMON II, and since RMON II will be a distinct MIB from RMON I, that RMON II is the place for this major change (easier to market/sell this change with new functionality).

- History Group Break Up

Currently, Token Ring assumes a break up in the history group. This is only a reorganization and adds no new functionality.

- Packet Distribution for History

Packet distribution occurs in the TR history tables. Numerous people would like it to also be present in RMONMIB. The major reason for not including it is space consideration (agent). Depending on one’s view, this change may be considered as either minor or major.

- Clarification of Status States

Providing a state diagram and/or more text would clear up concerns with the row status field.

- Incorporation of SNMPv2 Macros

The Area Directorate’s schedule indicates that a new draft should use the new SMI, but must not use any type that will break SNMPv1 systems.

- Events in SNMPv1 versus those in SNMPv2—another issue for RMON II.

- Clarification of Filters “Not” Mask

It was agreed that strict reading of the current text provides a correct interpretation, but that more clarification would reduce questions and concerns.

- Order Applied to Dependent Tables

Nowhere in RFC 1271 is it explicitly stated that all orders of setting variables in dependent tables are acceptable, i.e., the specification does not specify a particular order.

- Clarify indexing in each table.

- The RFC 1212 reference to IP address representation seems to be missed by numerous RMON readers, and the text should be repeated in RMON.

- Log Entry Standardization

It is probably impossible to specify fields for log entries, but the RMON document should provided a list of recommended entries.

- RMON Trap Document

Should the document be recreated?

- Packet Match Event Removal

Concern about the possibility of generating large numbers of events.

- Clarify whether a probe “hears” its own packets.

Attendees

| | |
|------------------|----------------------------|
| David Arneson | arneson@ctron.com |
| Jeff Case | case@cs.utk.edu |
| David Engel | david@ods.com |
| Michael Erlinger | mike@jarthur.claremont.edu |
| Steven Horowitz | witz@chipcom.com |
| Jeff Hughes | jeff@col.hp.com |
| Carl Madison | carl@startek.com |
| Richard McBride | richiem@axon.com |
| Keith McCloghrie | kzm@hls.com |
| James Reeves | jreeves@synoptics.com |

| | |
|----------------------|--------------------------------|
| Dan Romascanu | dan@lannet.com |
| Jean-Bernard Schmitt | jbs@vnet.ibm.com |
| Richard Sweatt | rsweatt@synoptics.com |
| Steven Waldbusser | waldbusser@andrew.cmu.edu |
| Peter Wilson | peter_wilson@3mail.3com.com |
| Dirk Wisse | dirk.wisse@dnpap.et.tudelft.nl |

2.3.14 Uninterruptible Power Supply (UPSMIB)

Charter

Chair(s):

Jeff Case, case@cs.utk.edu

Mailing Lists:

General Discussion: ups-mib@cs.utk.edu

To Subscribe: ups-mib-request@cs.utk.edu

Archive: ucs.utk.edu:~/pub/ups-mib/mail-archive

Description of Working Group:

This working group will produce a document that defines MIB objects for use in monitoring and (possibly) controlling both high-end and low-end UPSs and related systems (e.g., power distribution systems or power conditioning systems). Related devices may be addressed in this effort to the extent that the primary focus on UPSs is not compromised.

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

At its discretion, the working group may fulfill its charter by the development of distinct MIB definitions for UPS systems of differing capabilities, but the number of MIB definitions produced by the working group will not exceed two.

At its discretion, the working group may produce an additional document defining traps that support the management of UPSs.

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

Because the structure of UPS implementations varies widely, the working group shall take special care that its definitions reflect a generic and consistent architectural model of UPS management rather than the structure of particular UPS implementations.

Goals and Milestones:

Done Hold Interim Working Group meeting to review draft.

Done Post initial draft MIB to Internet-Drafts.

Done Meet at March IETF meeting to reach closure on MIB document.

Apr 1993 Submit the UPS MIB to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“UPS Management Information Base”, 07/30/1993, J. Case <draft-ietf-upsmib-00.txt>

CURRENT MEETING REPORT

Reported by Jeff Case/University of Tennessee

Minutes of the Uninterruptible Power Supply Working Group (UPSMIB)

The UPSMIB Working Group met on Monday, July 12, at the 27th meeting of the Internet Engineering Task Force, in Amsterdam, The Netherlands. The meeting was a small one.

A new draft of the result of the strawman effort ongoing since the Columbus meeting was presented and reviewed by those in attendance. A few errors were identified during the review. These will be corrected and the results will be posted to the mailing list and as an Internet-Draft to achieve the widest possible circulation. The meeting was helpful in that it caused the generation of a new draft, bringing some closure to the strawman effort. This new draft will serve as the basis for future discussions, coalescing the work to date into a single document.

It was noted that the group is making progress more rapidly now that the level of activity on the mailing list has increased. It was also noted that this is essential, because the amount of time invested to date has exceeded the chartered time for completion, and the group must complete its work in the near term.

Attendees

| | |
|-----------------|------------------|
| Jeff Case | case@cs.utk.edu |
| Steven Horowitz | witz@chipcom.com |
| Jack Stiekema | jack@victron.nl |

2.4 Operational Requirements Area

Director(s):

- Scott Bradner: sob@harvard.edu

Area Summary reported by Scott Bradner/Harvard University

The one BOF held under the Operational Requirements Area was:

- Generic Internet Service Specification (GISS)

The working groups currently open in the Operational Requirements Area are:

- Benchmarking Methodology Working Group (BMWG)
- BGP Deployment and Application Working Group (BGPDEPL)
- Network Joint Management Working Group (NJM)
- Network OSI Operations Working Group (NOOP)
- Operational Statistics Working Group (OPSTAT)
- User Connectivity (UCP)

BGPDEPL and OPSTAT met in Amsterdam.

Generic Internet Service Specification BOF (GISS)

A presentation was given on the current GISS work. The document, aimed at service providers, has undergone many changes since the previous BOF.

The consensus of the group was that a document of this type is badly needed by service providers, but the proper location for the work is in an “operators forum,” rather than an IETF working group. Until a global operators forum comes into existence, the IETF seems to target the correct audience. The chair will talk to Scott Bradner, the Area Director of the Operational Requirements Area, about forming a working group. A prospective charter was discussed.

The group is changing its name from “GISS” to “GISD” for Generic Internet Service Description.

BGP Deployment and Application Working Group (BGPDEPL)

A summary of CIDR deployment, including both route aggregation and IP address assignment, was presented. The current plan for CIDR deployment is:

1. Deploy BGP4 without aggregation
2. Advertise test aggregated route
3. Aggregate at the site level or single policy level, whichever is a smaller block
4. Understand more
5. Aggregate more

Steps four and five will be repeated, one after the other, until CIDR is fully implemented.

The group agreed that IBGP doesn't scale to very large numbers. However, it is currently tractable, and will be supported for a while, but the group should consider options for the future.

It was suggested that another aggregation rule should be added saying that no network should aggregate routes without informing other networks; a route aggregate registry could provide a means for communicating this information.

It was also suggested that de-aggregation should not be done in the initial stage. This should not cause a problem since initial aggregation will only occur at the site level or single policy level.

ANS, 3com, cisco, Proteon and Wellfleet were all asked about the status of their CIDR implementations.

Feedback was requested on BGP4 interoperability tests so that the interoperability matrix could be updated and sent to the list.

A syntax for registering route aggregates was presented.

Several people volunteered to write a paper analyzing the effectiveness of CIDR to be ready by September 1993.

There was a discussion on both the effectiveness and complexity of renumbering. Further analysis will be done.

Operational Statistics Working Group (OPSTAT)

The chair of the working group was unable to make the meeting, and the agenda was shortened as a result.

The ideas about the client/server strawman developed at, and since, the last meeting need to be incorporated into the strawman. Henry Clark volunteered to do this and send the result to the list.

CURRENT MEETING REPORT

Reported by Tony Bates/RIPE

Minutes of the Generic Internet Service Specification BOF (GISS)

Agenda

- Short presentation of GISS work so far
- Discussion of whether a working group is needed
- Review of proposed working charter

It was noted that the attendance at this meeting was quite low. However, this can, in part, be attributed to an unfortunate timing clash with the most popular session of the IETF - the IPng Decision Process BOF (IPDECIDE).

GISS Work Presentation

Tony Bates gave a presentation of the current GISS work. The document has changed quite dramatically since the last BOF meeting. A second draft is available from:

`ftp.ripe.net:ripe/docs/ripe-drafts/giss.[txt,ps]`

The major changes are in the focus and structure of the document. The document is aimed at service providers. Six areas of GISS have been highlighted:

1. Access
2. Generic Services
3. Connectivity
4. Operations
5. Information Provision and Coordination
6. Security

Currently 38 aspects have been highlighted in the first cut of GISS. Many are still to be completed. The intention is to have aspects contributed by relevant experts rather than the authors themselves. The structure of the document is very much like a UNIX manual page, making it easy to index only the GISS aspects of direct interest. It was agreed that the current areas and aspects were more than enough for a first GISS document.

An issue was raised as to whether GISS could in fact be renamed to GISD "Generic Internet Service Description" as is intimated in the document itself. It was agreed that the next draft and all future work will be referred to as GISD.

Working Group Discussion

The issue of whether a working group is needed was discussed. The clear consensus was that such a document is badly needed by service providers. It was felt the correct forum should be some sort of “operators” forum rather than the IETF itself. However, such a group doesn’t really exist, at least in a global context. In Europe, RIPE is such a group and a GISS-based working group will continue within RIPE. Andrew Partan also mentioned the possibility of such an operators forum being formed in North America. However, until such time, the IETF seems to contain the correct target audience, or to be more precise, many service providers also attend the IETF.

It was agreed that the chair would request that a GISD Working Group be formed within the Operations Area once the charter has been accepted by the group. A request along with the charter would be sent to Scott Bradner.

Proposed Working Group Charter Review

The charter was circulated. Apart from typographical changes and changes in wording (to make it clear that providers meant all forms of service providers, not just commercial), the charter was accepted.

- Name: Generic Internet Service Description (GISD)
- Chairs: Tony Bates and Daniel Karrenberg
- Mailing Lists:
 - General Discussion: `giss-wg@ripe.net`
 - To Subscribe: `giss-wg-request@ripe.net`
- Description:

GISD collects short descriptions of Internet service aspects. Internet service in GISD means the interaction of Internet service providers among themselves and with their users. GISD aims to provide a common frame of reference and vocabulary to talk about an Internet service. For each aspect of the Internet service, it describes different options for service provision in use in the current Internet. GISD is merely descriptive and does not proscribe or mandate. GISD is intended to be a living document collecting work of many contributors.

The GISD Working Group will update and revise the GISD document to assist network service providers in a better understanding and description of what Internet service means. It will:

- Update and revise the GISD document that lists the areas and aspects of interest to TCP/IP network service providers.
 - Identify additional GISD areas and aspects appropriate to GISD.
 - Identify areas of overlap with other IETF working groups.
 - Create a reference document of GISD terms.
 - Establish procedures to ensure the ongoing maintenance of the document and identify an organisation willing to do it.
- Goals and Milestones:
 - Review current GISD draft and add any additional areas and aspects deemed to be essential.
 - Draft of GISD will be prepared, reviewed and modified. Initiate IETF Internet-Draft review process by submission of the GISD document to the IETF Secretariat.
 - Follow-up with final amendments to the document and the submission of the document to the RFC Editor for publication as an FYI RFC.

Attendees

| | |
|-------------------|---------------------------------|
| Tony Bates | tony@ripe.net |
| David Conrad | davidc@iij.ad.jp |
| Ola Johansson | ojn@se.ubn.net |
| Daniel Karrenberg | daniel@ripe.net |
| Anne Lord | anne@ripe.net |
| Andrew Partan | asp@uunet.uu.net |
| Marten Terpstra | marten@ripe.net |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |

2.4.1 BGP Deployment and Application (BGPDEPL)

Charter

Chair(s):

Jessica Yu, jyy@merit.edu

Mailing Lists:

General Discussion: bgpd@merit.edu

To Subscribe: bgpd-request@merit.edu

Archive: [merit.edu:~/pub/bgpd-archive](http://merit.edu/~pub/bgpd-archive)

Description of Working Group:

The major purpose of this group is to coordinate BGP deployment and application in the current Internet.

It intends to create a forum for BGP users to share BGP deployment experiences and also provide a channel for users to communicate with router vendors who implemented or who are implementing BGP. It also intends to discuss BGP policy application and coordinate policy implementation in the current Internet routing environment which includes defining the usage of policy, defining a mechanism to share policy information, etc.

Goals and Milestones:

- | | |
|---------|--|
| Ongoing | Facilitate the deployment of BGP as widely as possible. |
| TBD | Define the issues and the needs of policy routing in the current Internet architecture. Discuss how BGP policy routing capability applies to Internet policy routing needs. A document may be generated on this topic. |
| Done | Post as an Internet-Draft, a report of BGP deployment status. |
| Done | Post an Internet-Draft, defining a mechanism to share policy information between Administrative Domains. |

Internet-Drafts:

“Notes of BGP-4/CIDR Coordination Meeting of 11 March 93”, 03/24/1993,
C. Topolcic <draft-ietf-bgpdepl-minutes-93feb-00.txt>

Request For Comments:

RFC 1482 “Aggregation Support in the NSFNET Policy Routing Database”

CURRENT MEETING REPORT

Reported by Jessica Yu/Merit

Minutes of the BGP Deployment Working Group (BGPDEPL)

Thanks to Mark Knopper and Bill Manning for taking the notes from which these minutes were derived.

Overview

Jessica Yu presented a summary of CIDR deployment status. Several different efforts are working on CIDR deployment:

- BGPDEPL Working Group meetings at IETFs
- BGPDEPL subgroup meeting at INTEROP (March) by Claudio Topolcic
- NSFNET regional-techs meetings by Merit
- RIPE Routing Working Group meeting by Tony Bates

CIDR has two major components: the route aggregation strategy, and the IP address assignment strategy as described in draft-fuller-cidr-strategy-03.txt (an update to RFC 1338).

The IP address assignment strategy described in RFC 1466 has been implemented since the fall of 1992 by IR. Jeff Huston from AARnet mentioned that non-CIDR compatible IP address assignments have been handed to Pacific region networks. Jeff will document the situation and inform the NIC.

The status of route aggregation implementation is as follows:

- The CIDR-capable routing protocol specification is done.
- Software development is underway.
- An initial deployment plan has been developed.

Editor's Note: A list of CIDR-related documents is available via FTP or mail server from the remote directories as /ietf/bgpdepl/bgpdepl-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

CIDR Deployment

The Current Deployment Plan: (initial)

- Step 1 - Deploy BGP4 without aggregation
- Step 2 - Advertise test aggregated route
- Step 3 - Aggregate at site level or single policy level, whichever is a smaller block
- Step 4 - Understand more
- Step 5 - Aggregate more

Steps 4 and 5 are recursive until CIDR is fully implemented. As soon as the CIDR software is ready, step 1 could be executed.

The rules for aggregation at initial deployment stage:

- Aggregate based on manual configuration
- Proxy aggregation allowed (with agreement of the advertiser)
- Holes in aggregates allowed
- IGP/IBGP carry aggregation within a domain
- Coordination: bi-lateral and overall
- Aggregate routing registry

A concern was raised about the merit of IBGP and whether it would be continuously supported by vendor software.

Dennis Ferguson clarified that IBGP doesn't scale to very large numbers since it requires a full mesh peering session between all the border routers within a domain. On ANS's network, each external router has over 90 IBGP neighbors. If an external network is lost, 90 announcements go out. This results in large overhead when routes flap. If your network grows big enough you should have something other than IBGP. Dennis is willing to write a short paper on the usefulness and limits of IBGP.

Tony Li of cisco stated that running IBGP is currently tractable and will be supported for a while, though we may consider alternatives for the future.

It is generally agreed that IBGP is usable in the size of the current network. It works on NSF/ANSnet, which has about 90 nodes participating in IBGP. A typical network has much fewer nodes participating in IBGP.

It was suggested to add another rule of aggregation, i.e. no network should aggregate routes without informing other networks. It was agreed that the aggregated routes should be registered. Implementing a route aggregate registry will provide a means of sharing such information.

It was suggested that de-aggregation should not be done at the initial stage. Since initially the plan is to only aggregate at site level or a single policy level, it (hopefully) will not cause too much inconvenience. But if it does, the issue would be revisited.

It was also agreed that if a network de-aggregates, those de-aggregated networks should not be propagated outside the network. With more experience gained on aggregation (or de-aggregation), these issues will be discussed further.

CIDR Capable Software Implementation

Router vendors are asked to respond to the following questions about the status of software implementation:

1. Is a CIDR implementation available?
2. What features are included and what are not?
3. Which version and where to get them?
4. Aggregation configuration syntax?
5. Interoperability test plan?

- Gated (by Dennis Ferguson)

ANS gated's BGP4 and aggregation implementation is being debugged and will be a beta-release soon. It is being tested on the NSFNET research network as well as in the ANS labs. It is ready for interoperability testing and Dennis will test against all the other implementations he can find.

The code is able to form route aggregates. Aggregation by proxy is supported. Each route can only contribute to a single aggregate, though the aggregate can contribute to a larger aggregate. Null routes to non-existent networks can be installed but needs kernel to support it. No controlled de-aggregation exists in this implementation.

The code is expected to be available in about two weeks. Dennis will create a distribution when it looks like things are working.

Dale Johnson mentioned that the syntax for the Merit "Network Announcement Change Request" will allow aggregates. The syntax for aggregates is `x.x.x/len` where `len` is the prefix length.

- 3com (by Arun Arunkumar)

BGP4 is being tested in the lab, talking to itself and ready to do interoperability test with other vendor's code. The code accepts, forwards and manages aggregated routes properly, but does not form route aggregates yet. The current implementation does not support controlled de-aggregation but will support it in the future if necessary. This will be released as part of version 6.2 in the September-October time frame. Aggregates could be carried by BGP4 and OSPF as well. 3com is working on implementing OSPF-BGP interaction. One month's testing is still needed.

- cisco (by Paul Traina)

Pre-beta code is available via anonymous FTP from cisco. Send mail to Paul Traina if you intend to use this code, and you will be added to the `bgp-beta` mailing list

and told how to get the code. BGP4 is currently based on version 9.21 of the router software, which is not yet in beta and thus is very “experimental.”

The implementation currently carries, advertises, and redistributes aggregates, but aggregate generation is still a few weeks off. Aggregate configuration uses the new route map feature, for which the user interface is not yet stable. BGP4 will redistribute aggregate routes with any routing protocol that carries mask information (EIGRP, OSPF, and ISIS). BGP3- and BGP4-OSPF interaction and automatic tag stuff is supported. Controlled de-aggregation is not currently supported (and may never be). Automatic negotiation is supported for BGP versions 2 through 4. The code will be deployed in a few test routers on regional networks and is currently going through early field testing. It is ready for interoperability testing (and probably has been tested against gated by the time you read this).

- Proteon (by Ed Stern)

They are currently testing BGP4, but are not ready for interoperability testing. They are not able to aggregate for the first release, but can pass aggregated routes and forward on the longest match. BGP can exchange routes between all other protocols. BGP and EGP can run simultaneously.

- Wellfleet (by John Kraczyk)

Release 7.60 is going into beta next month. This version contains BGP3, which also implements OSPF-BGP interaction.

A beta version of CIDR/BGP4 will hopefully be available sometime in early 1994. Plans include accepting, forwarding, and forming aggregates (also proxy), OSPF-BGP4 interaction, and possibly OSPF LSA type 8. Some form of controlled de-aggregation will also be included. Interoperability testing will be done when the implementation is closer.

- EuropaNET (by Peder Chr)

EuropaNET is working on implementation of BGP4. The Megaswitch is used, which is a custom router.

Interoperability Test

Tony Li requested feedback on BGP4 interoperability tests that so he can update his interoperability matrix for BGP and send it to the list.

Yakov suggested running a virtual DMZ for BGP testing, i.e. to establish remote BGP4 sessions between BGP4 test boxes at different locations on the net.

Route Aggregation Registry

Mark Knopper presented a syntax to register route aggregates in the NSFNET policy routing database. It is written in RFC 1482, and comments and suggestions are welcome.

Mark Knopper also presented a summary of the discussion at the NSFNET regional-techs meeting held June 10-11, 1993. The transcript of Mark's presentation, and other presentations, given at the Merit meeting can be obtained via FTP from `merit.edu:/pub/nsfnet/regional-techs`.

CIDR Analysis

It was suggested to do a CIDR analysis to evaluate CIDR's impact on the lifetime of IPv4. The IAB chartered this working group to write a paper to include such an analysis. The group suggested that the following areas be included in the analysis.

- CIDR impact on the routing table growth
- CIDR impact on the rate of IP address space depletion
- The rate of use of IP address space
- Impact of policy (AUP) on CIDR efficiency

The following people volunteered to work together to produce this analysis paper: Peter Ford, Dale Johnson, Tony Li, Bill Manning and Yakov Rekhter. Peter Ford agreed to take a lead on this. The paper should be ready by September 1993, before the IAB meeting takes place.

There was also a discussion about renumbering hosts to make better use of the assigned IP address space and increase the efficiency of aggregation. Peter Ford observed that lots of assigned Class B addresses have only 50 or so hosts on it, leaving the rest of the space unused. It was agreed that autoconfiguration could be of great help to renumbering. It was also suggested that it does not hurt to study the renumbering process with currently available technology. John Kraczyk mentioned that Wellfleet manually renumbered its network recently. He will document the process as a case study.

Attendees

| | |
|------------------|--------------------------------|
| Arun Arunkumar | <code>nak@3com.com</code> |
| Toshiya Asaba | <code>asaba@ij.ad.jp</code> |
| Tony Bates | <code>tony@ripe.net</code> |
| Rebecca Bostwick | <code>bostwick@es.net</code> |
| Ronald Broersma | <code>ron@nosc.mil</code> |
| Thomas Brunner | <code>brunner@switch.ch</code> |
| Henry Clark | <code>henryc@oar.net</code> |

| | |
|----------------------|----------------------------|
| David Conrad | davidc@iij.ad.jp |
| Tom Easterday | tom@cic.net |
| Deborah Estrin | estrin@usc.edu |
| Stefan Fassbender | stf@easi.net |
| Mark Fedor | fedor@psi.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Vince Fuller | vaf@stanford.edu |
| Craig Haney | craig@icp.net |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| Geoff Huston | g.huston@aarnet.edu.au |
| David Jacobson | dnjake@vnet.ibm.com |
| Dale Johnson | dsj@merit.edu |
| Anders Karlsson | sak@cdg.chalmers.se |
| Daniel Karrenberg | daniel@ripe.net |
| Sean Kennedy | liam@nic.near.net |
| Lothar Klein | lothar.klein@gmd.de |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Mark Kosters | markk@internic.net |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Tony Li | tli@cisco.com |
| Susan Lin | suelin@vnet.ibm.com |
| Robin Littlefield | robin@wellfleet.com |
| Peter Lothberg | roll@stupi.se |
| Bill Manning | bmanning@rice.edu |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Keith Mitchell | keith@pipex.net |
| Jun Murai | jun@wide.ad.jp |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Michael O'Dell | mo@uunet.uu.net |
| David O'Leary | doleary@cisco.com |
| Andrew Partan | asp@uunet.uu.net |
| Michael Patton | map@bbn.com |
| Juergen Rauschenbach | jrau@dfn.de |
| Yakov Rekhter | yakov@watson.ibm.com |
| Robert Reschly | reschly@brl.mil |
| Duncan Rogerson | d.rogerson@nosc.ja.net |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Miguel Sanz | miguel.sanz@rediris.es |
| John Scudder | jgs@merit.edu |
| Tim Seaver | tas@concert.net |
| Henk Steenman | henk@sara.nl |
| Ed Stern | els@proteon.com |
| John Stewart | jstewart@cnri.reston.va.us |
| Marten Terpstra | marten@ripe.net |
| Claudio Topolcic | topolcic@cnri.reston.va.us |

| | |
|-----------------------|---------------------------------|
| Paul Traina | pst@cisco.com |
| Willem van der Scheun | scheun@sara.nl |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Jessica Yu | jyy@merit.edu |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

2.4.2 Benchmarking Methodology (BMWG)

Charter

Chair(s):

Scott Bradner, sob@harvard.edu

Mailing Lists:

General Discussion: bmg@harvard.edu

To Subscribe: bmg-request@harvard.edu

Archive:

Description of Working Group:

The major goal of the Benchmarking Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

- | | |
|------|--|
| TBD | Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices. |
| Done | The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class. |
| Done | Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput. |

Request For Comments:

RFC 1242 “Benchmarking Terminology for Network Interconnection Devices”

2.4.3 Network Joint Management (NJM)

Charter

Chair(s):

Gene Hastings, hastings@psc.edu

Mailing Lists:

General Discussion: njm@merit.edu

To Subscribe: njm-request@merit.edu

Archive:

Description of Working Group:

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the Internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the agenda of this group will include the IAB and its task forces, and groups within FARNET. In particular FARNET has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the group will represent organizations with production responsibilities. Most work will be carried on via email or teleconferencing.

Goals and Milestones:

None specified

2.4.4 Network OSI Operations (NOOP)

Charter

Chair(s):

Susan Hares, skh@merit.edu
Cathy Wittbrodt, cjw@barrnet.net

Mailing Lists:

General Discussion: noop@merit.edu
To Subscribe: noop-request@merit.edu
Archive: merit.edu:~/pub/noop-archive

Description of Working Group:

The working group is chartered to work on issues related to the deployment of CLNP in the Internet. The first area of this group's work has been the learning necessary to start deploying OSI in Internet networks. This phase includes planning for OSI deployment by creating routing plans for regional networks and education on using OSI routing protocols.

This first area of the group's work will be on-going as we continue to deploy OSI in the Internet. This step has lead to people deploying OSI for pilot projects and demonstrations of OSI.

The second step of deploying OSI will be the transition of OSI from a pilot service to a production service. During this phase we will work on specifying the network debugging tools and test beds. We will need to track the level of OSI support in the Internet. We will need to provide documentation for new users of OSI on the Internet.

Goals and Milestones:

- | | |
|----------|--|
| Ongoing | Provide a forum to discuss OSI routing plans by email or in group discussions. |
| Jan 1992 | Post as an Internet-Draft, a tutorial for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP. |
| Apr 1992 | Post as an Internet-Draft, a requirements document specifying what OSI network tools are needed on every host and router. |
| Jul 1992 | Post as an Internet-Draft, a collection of regional Routing and Addressing plans. |
| Done | Post as an Internet-Draft, a list of OSI Network Utilities available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication. |
| Jul 1992 | Post as an Internet-Draft, a description of OSI network layer debugging methods. |

- Done Post as an Internet-Draft, a list of OSI Network Layer NOC tools available in the public domain and from vendors. This list will be passed over to the NOC tools Group effort for joint publication.
- Jul 1992 Submit to the IESG for Proposed Standard, a requirements document specifying what network tools are needed on every OSI host and router.
- Aug 1992 Submit to the IESG as an Informational RFC, a description of OSI network layer debugging methods.

Internet-Drafts:

“An Echo Function for ISO 8473”, 11/10/1992, S. Hares, C. Wittbrodt <draft-ietf-noop-echo-02.txt>

“Essential Tools for the OSI Internet”, 11/10/1992, S. Hares, C. Wittbrodt <draft-ietf-noop-tools-03.txt>

2.4.5 Operational Statistics (OPSTAT)

Charter

Chair(s):

Bernhard Stockman, boss@ebone.net

Phillip Gross, pgross@ans.net

Mailing Lists:

General Discussion: oswg-1@wugate.wustl.edu

To Subscribe: oswg-1-request@wugate.wustl.edu

Archive: wuarchive.wustl.edu:~doc/mailling-lists/oswg-1

Description of Working Group:

Today there exists a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The working group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

- Done Agreement on a model.
- Done Survey for most useful and popular metrics.
- Done Survey for most useful and popular presentation formats.
- Done Identify similar efforts being performed by other groups.
- Done Define a common minimal set of metrics.
- Done Propose a MIB for metrics not already there.
- Done Define a common storage format to facilitate data sharing.
- Done Define common presentation formats to make data comparable.
- Done Develop outline, and make writing assignments for paper (Opstat1) documenting March 1991 milestones.
- Done Complete paper Opstat1.

- Done Possible mid-term meeting to review Opstat1.
- Done Submit Opstat1 as Internet-Draft.
- Done Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
- Done Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
- Done Propose old tools to be retrofitted.
- Done Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools.
- Sep 1992 Submit Internet-Draft of new and retrofitted tools.
- Dec 1992 Submit new and old tools Internet-Draft to IESG as an Informational RFC.
- Mar 1993 Post an Internet-Draft defining the client/Server Opstat protocol.
- Mar 1993 Post the Opstat Statistical MIB as an Internet-Draft.
- Jul 1993 Submit the Client/Server Opstat Protocol to the IESG for consideration as a Proposed Standard.
- Jul 1993 Submit the Statistical Opstat MIB to the IESG for consideration as a Proposed Standard.

Request For Comments:

RFC 1404 "A Model for Common Operational Statistics"

CURRENT MEETING REPORT

Reported by Daniel Karrenberg/RIPE NCC

Minutes of the Operational Statistics Working Group (OPSTAT)

Introduction

Due to unforeseen circumstances, Bernhard Stockmann could not make it to the meeting. Daniel Karrenberg volunteered to chair a short session to proceed pending business and take in new initiatives. The agenda was shortened accordingly. No other regrets were received.

Client/Server Strawman Progress

The ideas developed at, and since, the last meeting need to be incorporated into the strawman. Henry Clark volunteered to do this and to circulate it on the list as soon as it is ready. He is also still working on an implementation. The group asked him to pragmatically document what will work.

Conclusion

There was no other business, so the meeting was closed.

Attendees

| | |
|-------------------|---------------------------|
| Anders Baardsgaad | anders@cc.uit.no |
| Robert Blokzijl | K13@nikhef.nl |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Kursat Cagiltay | kursat@vm.cc.metu.edu.tr |
| Henry Clark | henryc@oar.net |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Daniel Karrenberg | daniel@ripe.net |
| Cynthia Martin | martin@spica.disa.mil |
| Michael Patton | map@bbn.com |
| Erdal Taner | erdal@vm.cc.metu.edu.tr |

2.4.6 User Connectivity (UCP)

Charter

Chair(s):

Dan Long, long@nic.near.net

Mailing Lists:

General Discussion: ucp@nic.near.net

To Subscribe: ucp-request@nic.near.net

Archive:

Description of Working Group:

The User Connectivity Working Group will study the problem of how to solve network users' end-to-end connectivity problems.

Goals and Milestones:

- | | |
|------|--|
| Done | Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems. |
| TBD | Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate. |

Request For Comments:

RFC 1297 "NOC Internal Integrated Trouble Ticket System Functional Specification Wishlist ("NOC TT REQUIREMENTS")"

2.5 Routing Area

Director(s):

- Bob Hinden: hinden@eng.sun.com

Area Summary reported by Bob Hinden/Sun Microsystems

Border Gateway Protocol Working Group (BGP) and OSI IDRIP for IP Over IP Working Group (IPIDRP)

The BGP and IPIDRP Working Groups met jointly. BGP and IPIDRP will be writing a joint usage document. Implementors' experiences were solicited for writing the Proposed Standard report by September for both protocols.

BGP and IDRIP will be forwarding final documents, plus the Proposed Standard report, to the Routing Area Director so that BGP4 and IDRIP can go forward. Both IPIDRP and BGP will be going into "hiatus" once the standard requests are granted.

Inter-Domain Multicast Routing Working Group (IDMR)

The Amsterdam IETF meeting was the first official meeting of the IDMR Working Group. The working group met for two 2-hour sessions.

During the first session, Deborah Estrin gave a presentation on ESL, one of the new proposals for inter-domain multicast routing. This was the result of a collaboration with Steve Deering, Dino Farinacci, and Van Jacobson. The motivation behind the design of ESL was, for groups with a relatively small number of senders (sources), to allow receivers to receive data from those sources either over a shared tree, or over a shortest-path tree rooted at the source. The latter is useful for applications requiring minimal delay between senders and receivers. It was agreed that, because ESL is in its early stages of development, there remain specification and engineering details that need to be resolved.

The second session was mostly dedicated to discussing the IDMR charter. It was unanimously agreed that the current charter is lacking with respect to many aspects of inter-domain multicasting, and it should be a goal of the working group to try to resolve many of these, for example, user group management and interoperability.

The conclusion of this discussion was that the charter should be re-worked and re-submitted to the area director after the items to be worked on have been enumerated in order of priority.

IP Routing for Wireless/Mobile Hosts Working Group (MOBILEIP)

The MOBILEIP Working Group met twice at the Amsterdam IETF, with only one of the previously most active contributors unable to attend. Outside of the working group meetings themselves considerable time was spent over coffee tables, meals, and trains discussing the major issues. There seems to be movement towards some common mechanisms (the question of “encapsulation” versus “source routing,” for example, *seems* to have been settled in favor of encapsulation).

There were reports on a user requirements document, as well as on liaison activities with IEEE 802.11. There were substantial discussions about common terminology, beaconing, and how the location of a host is discovered. The creation of an “IP encapsulation working group” within the IETF was suggested.

RIP Version II (RIPV2)

The use of the Routing Domain in RIP-2 was discussed. Its use is still unclear. It was determined that the use of the field could not be sufficiently well defined to meet the varying needs of those few people who would like to use it. The field also poses difficult MIB problems (discussed below). Therefore, it has been decided to remove the field from the protocol and leave a Must Be Zero field in its place.

There were two proposed changes to the MIB. The first was to deprecate the Routing Domain object. It has been pointed out that the tables cannot be indexed correctly unless the Routing Domain object was used as part of the index. Given that the Routing Domain field is not well defined, this change would result in an overall simplification of the MIB. The second proposal dealt with handling unnumbered interfaces. While the RIP-2 protocol does not expressly address them, their existence does require consideration since the MIB tables cannot be indexed properly with unnumbered interfaces. The proposal is to use a network number of zero and a host number of ifIndex to create a suitable IP address for use in indexing tables.

There are currently two independent implementations of RIP-2: gated and Xylogics's routed. The MIB has been implemented for gated. ACC has a partial implementation of RIP-2 and is planning to implement the remainder.

Gerry Meyer's Demand Routing proposal was discussed at length. It was agreed that it performed a useful function. It was pointed out that it simulated many of the functions of TCP and that other routing protocols, such as RAP, used TCP.

Source Demand Routing (SDR)

Following a brief overview of the SDR forwarding protocol, Deborah Estrin described successful experiments completed on small-scale network testbeds including DARTnet. Plans were made for continued experimentation in conjunction with MERIT and others. No changes have been made to the specification since the last IETF; however a few very minor changes are planned.

Tony Li presented a language for describing SDRP policies, and a simple request-response protocol for exchanging this information. The group also reviewed the draft specification for optional-setup mode in SDRP. The implementation of this functionality will be finished at the end of the summer. Drafts of the policy language and setup specification are available now, and will be submitted as Internet-Drafts in the coming month or two. In addition, a draft usage document and MIB will be submitted as Internet-Drafts before the next IETF. At the next IETF Tony Li will lead a detailed walk through of the SDRP specification.

2.5.1 Border Gateway Protocol (BGP)

Charter

Chair(s):

Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:

General Discussion: bgp@ans.net

To Subscribe: bgp-request@ans.net

Archive:

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

- | | |
|----------|---|
| Ongoing | Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community. |
| Done | Complete development of Version 2 of the Border Gateway Protocol (BGP). |
| Done | Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol. |
| Done | Develop a MIB for BGP Version 3. |
| Done | Work with the Security Area to enhance the provision for security in BGP. |
| Done | Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy. |
| Done | Post an Internet-Draft specifying multicast extensions to BGP. |
| Done | Post the specification of BGP 4 as an Internet-Draft. |
| Done | Post an Internet-Draft specifying a MIB for BGP Version 4. |
| Jan 1993 | Submit the multicast extensions to BGP to the IESG as a Proposed Standard. |
| Jan 1993 | Submit the specification for BGP Version 4 to the IESG for consideration as a Proposed Standard. |
| Jan 1993 | Submit the BGP Version 4 MIB to the IESG for consideration as a Proposed Standard. |

Internet-Drafts:

“A Border Gateway Protocol 4 (BGP-4)”, 05/05/1992, Y. Rekhter, T. Li
<draft-ietf-bgp-bgp4-06.txt>

“Definitions of Managed Objects for the Border Gateway Protocol (Version 4)”,
09/01/1992, S. Willis, J. Burruss, J. Chu <draft-ietf-bgp-mibv4-03.txt>

“BGP4/IDRP for IP—OSPF Interaction”, 09/15/1992, K. Varadhan, S. Hares,
Y. Rekhter <draft-ietf-bgp-bgp4ospf-interact-01.txt>

Request For Comments:

RFC 1105 “Border Gateway Protocol BGP”

RFC 1163 “A Border Gateway Protocol (BGP)”

RFC 1164 “Application of the Border Gateway Protocol in the Internet”

RFC 1265 “BGP Protocol Analysis”

RFC 1266 “Experience with the BGP Protocol”

RFC 1267 “A Border Gateway Protocol 3 (BGP-3)”

RFC 1268 “Application of the Border Gateway Protocol in the Internet”

RFC 1269 “Definitions of Managed Objects for the Border Gateway Protocol (Version
3)”

RFC 1364 “BGP OSPF Interaction”

RFC 1397 “Default Route Advertisement In BGP2 And BGP3 Versions Of The Border
Gateway Protocol”

RFC 1403 “BGP OSPF Interaction”

CURRENT MEETING REPORT

Reported by Dave Katz/cisco Systems

Minutes of the joint session of BGP and IPIDRP Working Groups

Call to Order

Susan Hares, Chair of the OSI IDRPs for IP Over IP Working Group (IPIDRP), and Yakov Rekhter, Chair of the Border Gateway Protocol Working Group (BGP), called the meeting to order.

Dave Katz agreed to act as recording secretary.

The following agenda was presented:

- BGP4 progression to Proposed Standard
- BGP4 MIB progression to Proposed Standard
- BGP/IDRP policy MIB
- BGP/IDRP over X.25
- IDRPs status
- IDRPs for IP next step
- IDRPs for IPng

BGP4 Progression to Proposed Standard

Yakov described the requirements for progressing a document to Proposed Standard status. One of these requirements is a report to the Routing Area Director describing, among other things, implementation experience. Yakov solicited an implementor to write this report and Paul Traina graciously volunteered.

BGP4 MIB Progression to Proposed Standard

Yakov reported that no comments were received on the Internet-Draft MIB, and that no implementations are currently known, although it is reported that several implementations are underway. Once implementations appear, the required report will be written, and the document will be progressed at that point.

BGP/IDRP Policy MIB

John Krawczyk was asked by Yakov shortly before the meeting to report on Wellfleet's implementation of BGP policy control via SNMP. John described Wellfleet's proprietary policy MIB and its elements.

Editor's Note: A description of these elements is available via FTP or mail server from the remote directories as /ietf/ipidrp/ipidrp-bgp-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

A discussion ensued about the practicality of standardizing a policy MIB. There is a tension between trying to provide a mechanism by which a router can be completely configured using SNMP, and the fact that policy mechanisms are product differentiation features that are likely to differ from vendor to vendor. A standardized MIB would be an advantage operationally, but it is likely to be difficult to describe a canonical MIB that would reflect the breadth of functionality available in all implementations. One possibility would be to define a standard canonical subset of functionality and then use proprietary MIB branches for features unique to particular implementations.

As there are very few policy MIB implementations at this point in time, it was agreed to table this topic until more implementations are available, at which time the issue will be revisited.

BGP/IDRP over X.25

Gerry Meyer described a scheme for use of BGP over tariffed switched circuits, including X.25. There is a general problem with protocols that send periodic messages, in that they tend to hold switched circuits open forever, and generate lots of (potentially costly) traffic. They also may require full mesh connectivity, which requires many circuits.

An Internet-Draft has been published that describes a mechanism for running RIP in such environments. Gerry described an adaptation of this mechanism to BGP.

Editor's Note: Gerry's description is available via FTP or mail server from the remote directories as /ietf/ipidrp/ipidrp-bgp-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Susan pointed out that a similar scheme for IDRP has been published by the Aeronautical Telecommunications Network forum.

A discussion ensued about how this type of functionality could be accomplished with minimal changes to BGP itself. It was decided that the only necessary functional change to the BGP protocol would be to specify that a negotiated holding timer value of zero would indicate that periodic keepalive packets are not in use. Other functionality would be part of a connection manager which would not be specified by the BGP Working Group, although

the functionality expected by BGP would be. Gerry agreed to modify the BGP usage document to specify the expected connection manager functionality. Yakov agreed to make the holding time modification to the BGP4 document.

BGP Working Group Status

Tony Li moved that the BGP Working Group status be changed to "hiatus." Yakov pointed out that this does not mean that the group is disbanding, but rather that it will lie dormant until it is necessary to resurrect it (in particular, to advance documents through the standards track). The group agreed by consensus.

IDRP Standard Status

Dave Katz reported that the IDRP specification, ISO 10747, had passed the DIS ballot and the document editor had accommodated all ballot comments. The final IS text was sent to the ISO Secretariat in Geneva this week, and ISO will formally ratify the document during the next ISO SC6 meeting in Seoul, South Korea in October.

It was announced that an ASCII version of the IDRP document, to be published as an RFC, had been created by the document editor and was available on the ISO document archive on merit.edu.

IDRP Usage Document

The question was raised as to whether IDRP requires a usage document, distinct from the BGP usage document. It was noted that IDRP has a number of features that BGP does not, in particular, Routing Domain Confederations and Distribution Lists, and that it would be useful for a usage document to discuss the use of these features. It was decided that there would be a unified BGP/IDRP usage document based on the current BGP usage document. Susan and John Scudder agreed to write this document; Tony agreed to review it.

IDRP MIB

A MIB document exists, with local MIB extensions. The question was raised as to whether it would be worthwhile to create a unified MIB document. It was agreed that the BGP and IDRP MIB documents would remain distinct.

IDRP Implementation Status

The Aeronautical Telecommunications Network standards specify IDRP over a fully connected SMDS-like network running without an intra-domain routing protocol. One European router vendor expects an implementation within “eight to ten weeks.” Two ports of gated are being worked on (Sun and Hewlett-Packard). CSC is doing a scratch implementation for the Federal Aviation Administration. IBM is also doing a scratch implementation, but stressed that this was not a product. The implementation is expected to become available (end of 1993).

No implementation provides any significant policy functionality at this time. The IBM implementation is intended to support the full policy syntax specified in an appendix in the IDRP specification. No implementation, except for IBM, supports Routing Domain Confederations at this time, though all implementors plan to do this (especially since it is a mandatory part of the protocol). The IBM implementation will support Routing Domain Confederations as part of its initial release (end of 1993).

IDRP is currently active in Europanet, connecting two parts of Europanet together.

Susan asked for volunteers to provide information on implementation experiences to be provided in support of progression of IDRP for IP to Proposed Standard status. The group agreed to pursue this progression.

IDRP for IPng

Susan reported that IDRP has been offered to all of the IPng working groups as a candidate inter-domain routing protocol. The TUBA group plans to use IDRP; the SIP and PIP groups are evaluating the protocol. Susan is preparing an IDRP for SIP document.

IDRP for IP Working Group Status

Susan moved that the IDRP for IP Working Group be moved to “hiatus” status until the documents are ready for progression to Proposed Standard status. The group agreed to this proposal.

The meeting was then adjourned.

Attendees

| | |
|----------------------|------------------------------|
| Bernt Allonen | bal@tip.net |
| Arun Arunkumar | nak@3com.com |
| Toshiya Asaba | asaba@iiij.ad.jp |
| Dennis Baker | dbaker@wellfleet.com |
| Tony Bates | tony@ripe.net |
| Ronald Broersma | ron@nosc.mil |
| Thomas Brunner | brunner@switch.ch |
| Henry Clark | henryc@oar.net |
| David Conrad | davidc@iiij.ad.jp |
| Tom Easterday | tom@cic.net |
| Hans Eriksson | hans@sics.se |
| Stefan Fassbender | stf@easi.net |
| Mark Fedor | fedor@psi.com |
| Dennis Ferguson | dennis@ans.net |
| Peter Ford | peter@goshawk.lanl.gov |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Vince Fuller | vaf@stanford.edu |
| Craig Haney | craig@icp.net |
| Susan Hares | skh@merit.edu |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| David Jacobson | dnjake@vnet.ibm.com |
| J. Jensen | jensen@ic.dk |
| Cyndi Jung | cmj@3com.com |
| Marijke Kaat | marijke@sara.nl |
| Dave Katz | dkatz@cisco.com |
| Sean Kennedy | liam@nic.near.net |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Ton Koelman | koelman@stc.nato.int |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Tony Li | tli@cisco.com |
| Robin Littlefield | robin@wellfleet.com |
| Peter Lothberg | roll@stupi.se |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Gerry Meyer | gerry@spider.co.uk |
| Keith Mitchell | keith@pipex.net |
| Ramin Najmabadi | najmabadi@helios.iihe.rtt.be |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Michael O'Dell | mo@uunet.uu.net |
| Andrew Partan | asp@uunet.uu.net |
| Willi Porten | porten@gmd.de |
| Juergen Rauschenbach | jrau@dfn.de |
| Yakov Rekhter | yakov@watson.ibm.com |
| Georg Richter | richter@uni-muenster.de |

| | |
|-------------------|----------------------------|
| Duncan Rogerson | d.rogerson@nosc.ja.net |
| Miguel Sanz | miguel.sanz@rediris.es |
| John Scudder | jgs@merit.edu |
| Henk Steenman | henk@sara.nl |
| Ed Stern | els@proteon.com |
| John Stewart | jstewart@cnri.reston.va.us |
| John Stewart | |
| Marten Terpstra | marten@ripe.net |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Paul Traina | pst@cisco.com |
| Rene van der Hauw | rene@geveke.nl |
| Rachel Willmer | rachelw@spider.co.uk |
| Jessica Yu | jyy@merit.edu |
| Romeo Zwart | romeo@sara.nl |

2.5.2 IP Routing for Wireless/Mobile Hosts (MOBILEIP)

Charter

Chair(s):

Steve Deering, deering@parc.xerox.com
Greg Minshall, minshall@wc.novell.com

Mailing Lists:

General Discussion: mobile-ip@parc.xerox.com
To Subscribe: mobile-ip-request@parc.xerox.com
Archive: parcftp.xerox.com:~/pub/mobile-ip/mail-archive

Description of Working Group:

The Mobile IP Working Group is chartered to develop or adopt architectures and protocols to support mobility within the Internet. In the near-term, protocols for supporting transparent host “roaming” among different subnetworks and different media (e.g., LANs, dial-up links, and wireless communication channels) shall be developed and entered into the Internet standards track. The work is expected to consist mainly of new and/or revised protocols at the (inter)network layer, but may also include proposed modifications to higher-layer protocols (e.g., transport or directory). However, it shall be a requirement that the proposed solutions allow mobile hosts to interoperate with existing Internet systems.

Longer term, the group may address, to the extent not covered by the mobile host solutions, other types of internet mobility, such as mobile subnets (e.g., a local network within a vehicle), or mobile clusters of subnets (e.g., a collection of hosts, routers, and subnets within a large vehicle, like a ship or spacecraft, or a collection of wireless, mobile routers that provide a dynamically changing internet topology).

Goals and Milestones:

- | | |
|----------|---|
| Done | Review and approve the Charter, making any changes deemed necessary. |
| Nov 1992 | Post an Internet-Draft documenting the Mobile Hosts protocol. |
| Mar 1993 | Review the Charter of the Mobile IP Working Group for additional work required to facilitate non-host mobility. |
| Mar 1993 | Submit the Mobile Host Protocol to the IESG as a Proposed Standard. |

CURRENT MEETING REPORT

Reported by John Johnston/National Semiconductor

Minutes of the IP Routing for Wireless/Mobile Hosts Working Group
(MOBILEIP)

Agenda

- Review of mobility model
- Liaison reports
- Document status
- Subcommittee reports
- Short presentations
- Interim Meeting

Mobility Model

Greg Minshall reviewed the mobility model for the first time attendees in the session. Basically, the problem was stated as finding a methodology (architecture, protocol, etc.) to support the routing of mobile hosts (MH). In most of the models presented to date, each MH has both a home address and a forwarding or care of address. Packets are sent to a MH via its home address. These packets are directed via normal routing to a base station which serves as the home base of the MH. This base station must know where the MH is at all times by maintaining the IP address of the base station currently serving the MH. Assuming the MH is “not at home” the base station then forwards the packet to its peer currently serving the MH, who in turn delivers the packet directly to the MH.

The base station to base station delivery mechanism is called tunneling which can result in inefficient (dogleg) routing. An extreme example is a US-based Amsterdam IETF attendee trying to connect to a local Amsterdam host. Packets would be first routed over the Atlantic to the home base station in the US, then routed back over the Atlantic to the MH’s current base station in Amsterdam, and finally delivered to the local Amsterdam host. This problem will be handled through the use of address caching.

Finally, Greg clarified the scope of the working group as supporting media independent mobility. One solution must handle wireless IR, wireless RF, ethernet, etc.

Liaison Reports

Charlie Perkins reported that 802.11 is standardizing IEEE MAC protocols for wireless media. This body is meeting during the same week as the IETF. At their last session,

Charlie proposed that the IETF working group inform 802.11 about all of the network layer events and indications that will be necessary to support mobile IP. Charlie indicated that 802.11 still has many open issues including MAC address selection (48-bit?).

Steve Alexander did not report on the Dynamic Host Configuration Working Group (DHC) because the group has not met since the last IETF.

Scott Kaplan, the liaison for the Domain Name System Working Group (DNS), did not report, because he was unable to attend the last DNS session.

There was no report given for the Internet Protocol Security Protocol Working Group (IPSEC) because John Ioannidis was unable to attend this MOBILEIP session.

Document Status

- Fumio Teraoka from Sony will distribute a new version of his document when he returns from the IETF.
- Dave Johnson has an updated version of his document available.
- The working group has not heard from Columbia University.
- Charlie Perkins will complete a new version of his document in several months.
- Matsushita's draft work is continuing. A version will be released in several weeks.

Subcommittee Reports

Editor's Note: The subcommittee reports are available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Short Presentations

Editor's Note: The presentations are available via FTP or mail server from the remote directories as /ietf/mobileip/mobileip-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Interim Meeting

There was consensus to meet somewhere on the east coast sometime before the November IETF. Possible dates and locations will be discussed via the mailing list.

Further Discussion

Yakov Rehkter's subcommittee report on user requirements created a discussion regarding CDPD. This technology was described by Mark Knopper:

- Consortium of 9 large US and 1 Canadian voice carriers.
- Data services over cellular infrastructure.
- Mobile End System makes itself known to a Mobile Intermediary System.
- Packets routed first to Intermediary System which forwards them to End System.
- Billing through X.400.

Suzy Brown expressed the desire for the IETF to press ahead quickly to avoid the potential for deployment of technology-specific solutions that will not interoperate with the Internet. Other infrastructure-based solutions are being developed (Ram Mobile Data, Mobitex, GSM, etc.).

Along the same lines, John Penners' review of mobile services spawned discussions centering on the relationship between mobile IP and the many specialized services and providers. Steve Deering presented a model that emphasized the logical separation from wireless service providers and the Internet. This led to several observations:

- We should view the technologies as physical medium below IP (service provider is at lower level).
- Successful mobile IP deployment could leverage incorporation into technology-specific switches.
- A goal is to avoid "doglegs."
- Single hop at IP level (not single physical hop) in service provider.
- Understanding providers' rules might help the dogleg problem.
- Service providers need mobile IP to connect their clouds.

Greg Minshall presented a similar model emphasizing how CDPD could create a coast-to-coast dogleg because it is not care of address-aware. This led to a discussion over whether it would be beneficial to take proactive measures to influence CDPD.

During Charlie Perkins' presentation on beaconing procedures, Steve Deering emphasized the desirability that mobile hosts transmit new base station updates (as opposed to IAPs). Also, Steve stated that he would like to use multicast addressing over broadcast whenever possible (addressing must be consistent within a cell), and Greg indicated that we should request a "well known" address for this purpose.

Open Issues

- Can we use existing router discovery methods to support beaconing?
- Does Dave Johnson's encapsulation technique break traceroute?
- Does Dave Johnson's encapsulation technique break MTU discovery?

Action Items

- Steve Deering will contact Columbia University for an update on their work.
- Determine if a working group should be formed within the IETF to deal with the issues of encapsulation.
- Obtain a well known multicast address.
- Obtain new ICMP number.

Attendees

| | |
|-------------------|-----------------------------------|
| Kannan Alagappan | kannan@DSMAIL.ENET.DEC.COM |
| Steve Alexander | stevea@lachman.com |
| Nick Alfano | alfano@mpr.ca |
| James Allard | jallard@microsoft.com |
| Bernt Allonen | bal@tip.net |
| Michael Anello | mike@xlnt.com |
| Anders Baardsgaad | anders@cc.uit.no |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Dennis Baker | dbaker@wellfleet.com |
| John Ballard | jballard@microsoft.com |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Per Bilse | bilse@ic.dk |
| Jim Binkley | jrb@ibeam.intel.com |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Michael Brescia | |
| Ronald Broersma | ron@nosc.mil |
| Ramon Caceres | ramon@mitl.research.panasonic.com |

| | |
|---------------------|---------------------------------|
| Thomas Cordetti | tomc@digibd.com |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Stephen Deering | deering@parc.xerox.com |
| Pierre Dupont | dupont@mdd.comm.mot.com |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Julio Escobar | jescobar@bbn.com |
| Dan Frommer | dan@jeremy.enet.dec.com |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Robert Gilligan | Bob.Gilligan@Eng.Sun.Com |
| Jari Hamalainen | jah@rctre.nokia.com |
| Mark Handley | mhandley@cs.ucl.ac.uk |
| Gerd Holzhauser | Holzhauser1@applelink.apple.com |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Steven Horowitz | witz@chipcom.com |
| Chris Howard | chris_howard@inmarsat.org |
| Geoff Huston | g.huston@aarnet.edu.au |
| David Johnson | dbj@cs.cmu.edu |
| John Johnston | john@berlioz.nsc.com |
| Philip Jones | p.jones@jnt.ac.uk |
| Marijke Kaat | marijke@sara.nl |
| Scott Kaplan | scott@wco.ftp.com |
| Ton Koelman | koelman@stc.nato.int |
| Mark Kosters | markk@internic.net |
| Mark Laubach | laubach@hpl.hp.com |
| Tony Li | tli@cisco.com |
| Susan Lin | suelin@vnet.ibm.com |
| Cynthia Martin | martin@spica.disa.mil |
| Donald Merritt | don@arl.army.mil |
| Paul Milazzo | milazzo@bbn.com |
| Greg Minshall | minshall@wc.novell.com |
| William Miskovetz | misko@cisco.com |
| Keith Mitchell | keith@pipex.net |
| Henri Moelard | henri.moelard@utrecht.ncr.com |
| Jun Murai | jun@wide.ad.jp |
| Ronny Nilsen | Ronny.Nilsen@usit.uio.no |
| Petri Ojala | ojala@eunet.fi |
| Zbigniew Opalka | zopalka@agile.com |
| John Penners | jpenners@advtech.uswest.com |
| Charles Perkins | perk@watson.ibm.com |
| Jim Rees | Jim.Rees@umich.edu |
| Yakov Rekhter | yakov@watson.ibm.com |
| Henry Sanders | henrysa@microsoft.com |
| Hal Sandick | sandick@vnet.ibm.com |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| Fumio Teraoka | tera@csl.sony.co.jp |

| | |
|------------------|---------------------------------|
| Antoine Trannoy | trannoy@crs4.it |
| Thierry Turletti | turletti@sophia.inria.fr |
| Werner Vogels | werner@inesc.pt |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |
| Kirk Williams | kirk@sbctri.sbc.com |
| Rachel Willmer | rachelw@spider.co.uk |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Jessica Yu | jyy@merit.edu |
| Romeo Zwart | romeo@sara.nl |

A Review of Existing Mobile Services and Technologies

John Penners
7/8/93
27th IETF

Purpose

Provide:

- Technical insight into alternative existing approaches
- A benchmark for evaluating mobile solutions

Services and Technologies (Not Exhaustive)

Mobile Data Services

RAM
ARDIS

Mobile Data Technologies

CDPD
PRMA

Cellular Technologies

AMPS
GSM
DECT

RAM (RADio/Modem)

- Partially owned by BellSouth
- Based on MOBITEK protocol which is openly available without fee to any vendor desiring to produce MOBITEK compatible equipment and/or applications.
- Hierarchical System has a Network Control Center, Main Switch, Regional Switch, Local Switch, and a Base Station.
- Messages are routed to destinations by the lowest network level for fast response times.
- 4.8 kbps effective throughput over a 8kbps data rate.
- Data only packet switched network
- 1-512 character Packet length

ARDIS

- A Partnership of IBM and Motorola provides Real-Time Data Access
- 19.2 / 4.8 Kbps
- 1 to 2048 bytes of user data in a message (RF packets are 240 bytes of user data)
- Host Protocols Include
SNA LU6.2, X.25, SNA 3270 and Bisynchronous 3720,
Asynchronous Communications
- 4.8 kbps operating in 400 major metropolitan areas
- Five year expansion plan includes 900 base stations with an estimated investment of \$150,000,000.

CDPD (Cellular Digital Packet Data)

- Developed by a consortium of cellular service providers and IBM
- 9.2 kbps over 30kHz channels
- Specification is extensive (approx 800pgs) and open.
- Data is sent on idle voice channels - voice always has priority
- Reuses cell sites and frequencies but requires additional base station components that are not expected to be cheap.
- "I believe that data transmission via cellular, given our economies with the use of CDPD, is the greatest opportunity that our business has before it today" - Craig McCaw

PRMA (Packet Reservation Multiple Access)

- Initiated by David Goodman at the Wireless Information Network Laboratory (WINLABS) at Rutgers University
- Primarily a data link and physical layer protocol that is based on Slotted ALOHA and TDMA, which provides a reservation mechanism
- Uses a packet header to guide packets through network switches and base stations.
- Information contained in packets are categorized as "Periodic" or "Random" Speech packets are periodic
- System being prototyped by Bellcore and WINLABS

AMPS (Advanced Mobile Phone System)

- First Generation Wireless Communication System (No Longer considered advanced)
- Hybrid of Digital and Analog Technology - Digital used for control and analog used for voice.
- Four Channels: Reverse Control Channel (RECC), Reverse Voice Channel (RVC), Forward Control Channel (FOCC), Forward Voice Channel (FVC)
- Inband Control Signaling called Blank and Burst
- The Voice Channels are dedicated, the Forward Control Channel is broadcast, and the Reverse Control Channel is random access with collisions.

GSM (Groupe Special Mobile)

- European Digital Cellular
- Six Basic Elements: Mobile Stations (MS), Base Stations (BS), Mobile Service Switching Centers (MSC), Home Location Registers (HLR), Visitor Location Registers (VLR), Equipment Identity Registers (EIR)
- Uses Mobile Assisted Handoff
- Six Signalling Protocols: Call Management (CM), Mobility Management (MM), Radio Resource Management (RRM), Signal Connection Control Part (SCCP), Message Transfer Part (MTP), Link Access Protocol - D Channel (LAPD)

DECT (Digital European Cordless Telecommunications)

- Intended for small systems (Similar to a wireless PBX with handoff)
- Single cell or Multicell capabilities
- Transmission rates up to 1.152 Mb/s
- High Quality Voice
- Five Control Channels: Call Management Special Features (C), Physical Layer Control (M), Handshaking (N), Paging (P), System Information (Q).
- Quick Handoffs controlled by the handset

A Mobile Host Routing Protocol

David B. Johnson

School of Computer Science
Carnegie Mellon University

dbj@cs.cmu.edu



The Entities

The Mobile Host

The Sending Host

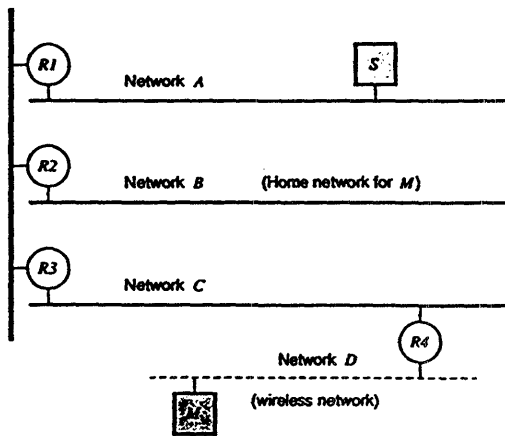
Support functions:

- The Base Station
- The Location Server
- Location Caches

The Base Station

Keeps list of visiting mobile hosts currently on network

Transmits arriving packets locally to mobile host

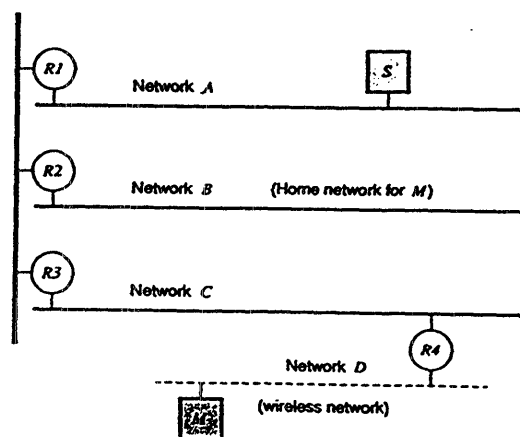


The Location Server

On home network, reachable by mobile host's IP address

Maintains database of location of network's mobile hosts (Base Station IP address)

Tunnels mobile host's packets to the Base Station

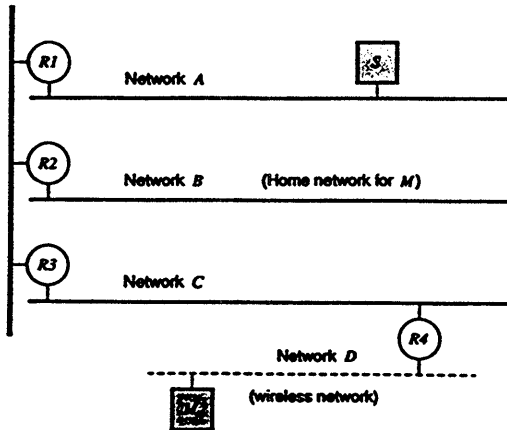


Location Caches

Any host or router that caches the location of a mobile host

Usually the sending host or first-hop router

Tunnels mobile host's packets to the Base Station



Moving a Mobile Host

Disconnecting from the network:

- Mobile host notifies old Base Station and Location Server

Reconnecting to the network:

- Mobile host notifies new Base Station and Location Server
- Can disconnect from old at the same time

Old Base Station may cache "forwarding pointer" to new Base Station (Location Cache)

New Base Station adds mobile host to list

Location Server:

- Maintains database of host's current location
- Arranges to intercept arriving packets for mobile host (e.g., "proxy" ARP)

Locating a Mobile Host

Sender never knows (does not need to know) that destination host is mobile

A mobile host's IP address will route packets for it to its home network

The Location Server on the home network:


- Intercepts packet for mobile host
- Forwards packet to Base Station
- Returns ICMP "mobile host redirect" with Base Station address to sender

Sender or any router can cache Base Station address (Location Cache)

Any out-of-date Location Caches are automatically corrected when (and if) needed

ICMP Mobile Host Redirect

Similar to standard ICMP redirect message:



Sent to original sender of packet

Base Station IP Address = 0 causes cache entry to be deleted

Redirect is handled similar to existing host-specific ICMP redirect:

- Store in same table (new type field on records)
- Lookup already required for host-specific redirect
- Sender can find Base Station address for mobile host in table with little or no extra cost

Tunneling to a Mobile Host

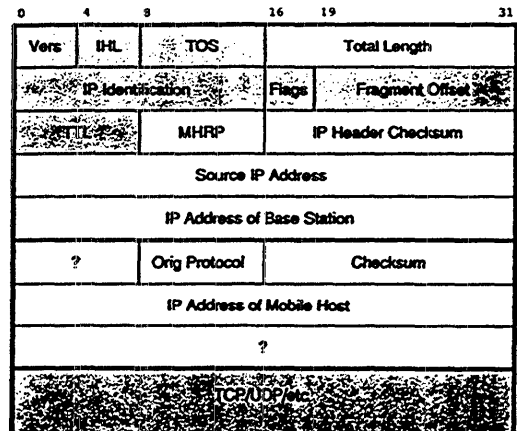
Tunneling to Base Station is initiated by:

- The sending host (Location Cache), or
- Any intermediate router (Location Cache), or
- The Location Server on the mobile host's home network

Why no longer using IP Loose Source Routing?

- Slows down packet forwarding
- Lots of buggy implementations
- Politically incorrect
- Can do only what it's already defined to do


Mobile Host Routing Protocol



Transform original packet in new IP packet addressed to Base Station

Can carry along extra information specific to mobile host routing

MHRP Protocol Description




IP source address is always "entry point" of tunnel

IP destination address is always "end point" of tunnel

"Forwarding pointer" Location Cache starts new tunnel:

- Previous IP destination is IP source for this tunnel
- Next IP destination is new Base Station address
- Keep list of previous IP source addresses in MHRP header

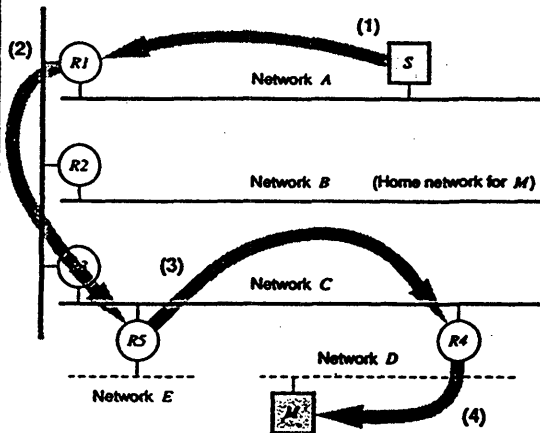
Example: Sender is a Location Cache



(1) IP = (MHRP, S, R4) MHRP = (TCP, M, ())

(2) IP = (TCP, S, M)

Example: Use of a Forwarding Pointer



- (1) $IP = \langle TCP, S, M \rangle$
- (2) $IP = \langle MHRP, R1, R5 \rangle$ $MHRP = \langle TCP, M, (S) \rangle$
- (3) $IP = \langle MHRP, R5, R4 \rangle$ $MHRP = \langle TCP, M, (S, R1) \rangle$
- (4) $IP = \langle TCP, S, M \rangle$

Location Cache Maintenance

Once packet received by correct Base Station:

- Send ICMP mobile host redirect to MHRP list of previous IP source addresses

If packet received by Location Server:

- Send ICMP mobile host redirect to MHRP list of previous IP source addresses
- Also send redirect to current IP source address

All out-of-date or missing cache entries used for this packet now point to correct Base Station

Base Station State Recovery

If Base Station "forgets" a mobile host, Base Station will tunnel arriving packets to the Location Server

Location Server sends normal ICMP mobile host redirects

Current Base Station will receive redirect to itself:

- Base Station could just add mobile host back to list of locally visiting mobile hosts
- Base Station could instead send query message on local network to see if mobile host is really there

Location Server discards original packet

When rebooting, Base Station broadcasts query on local network to find any mobile hosts still there

Robustness Against Routing Loops

Loops of Location Caches are not possible with a correct implementation of the protocol

Many different implementations of each protocol in the Internet

A buggy implementation might create loops and a lot of congestion

The protocol can easily be made robust against loops:

- Check for duplicate in list of previous IP source addresses in MHRP header when adding to list
- If duplicate, send ICMP mobile host redirect to all to delete cache entries
- Either discard original packet or tunnel to Location Server

Summary

- Each organization manages its own mobile hosts
- Any host running the right software can become mobile at any time
- No overhead when mobile host is "at home"
- Location Caches automatically corrected if needed
- Robust against routing loops
- No temporary IP address assignment needed
- No broadcast or multicast
- Scalable
- Low overhead
- Simple protocol

2.5.3 IS-IS for IP Internets (ISIS)

Charter

Chair(s):

Ross Callon, rcallon@wellfleet.com
Chris Gunner, gunner@dsmail.lkg.dec.com

Mailing Lists:

General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu
Archive:

Description of Working Group:

The ISIS Working Group will develop additions to the existing OSI IS-IS routing protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

- Done Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
- Done Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.
- Done Post a revision of the IS-IS as an Internet-Draft.
- Mar 1993 Submit the revised IS-IS to the IESG as a Draft Standard.
- Mar 1993 Submit the IS-IS MIB to the IESG as a Proposed Standard.

Internet-Drafts:

- “Integrated IS-IS Management Information Base”, 11/05/1991, Chris Gunner <draft-ietf-isis-mib-02.txt>
- “Further Integration of IS-IS; Appletalk, IPX, and Other Protocols”, 06/25/1993, R. Perlman, C. Gunner <draft-ietf-isis-atipx-00.txt>
- “Routing over Nonbroadcast Multiaccess Links”, 07/07/1993, R. Perlman, C. Gunner <draft-ietf-isis-nbma-00.txt>
- “Multiple Levels of Hierarchy with IS-IS”, 08/09/1993, R. Perlman, C. Gunner <draft-ietf-isis-multilevel-routing-00.txt>

Request For Comments:

RFC 1195 “Use of OSI IS-IS for Routing in TCP/IP and Dual Environments”

CURRENT MEETING REPORT

Reported by Chris Gunner/Digital Equipment Corporation

Minutes of the IS-IS for IP Internets Working Group (ISIS)

The ISIS Working Group met for one session on Thursday morning. The meeting started with a review of what was discussed at the previous meeting, followed by status reports of work items from that meeting.

RFC 1195

Little progress has been made in the effort to advance the standardization status of Integrated IS-IS (RFC 1195) from Proposed Standard to Draft Standard. Examples of operational experience, especially interworking between different implementations, and implementations of the MIB were gratefully received. Other examples should be reported to Chris Gunner.

Multicasting

At the last meeting, Radia Perlman briefly discussed a simple idea for supporting multicast with Integrated IS-IS. Little progress has been made on this. There was some discussion about how this relates to the multicast work being done for CLNP (ISO 8473) and ES-IS (ISO 9542) by ISO (although ISO is not working on multicast support in routing protocols). The group felt that there should be some liaison between the IETF group and the ISO group working on this topic. It is possible that both multicast schemes might be useful and used in the same network, for example, using the ISO scheme at level 2 while using Radia's scheme at level 1.

Multi-Area Router

No progress has been made on the multi-area router idea discussed at the last meeting. The group still thinks this is worthwhile and should be written down as an Internet-Draft.

Multiple Levels of Hierarchy

No progress has been made on the multiple levels of hierarchy idea proposed by Radia Perlman at the last meeting. There was some discussion about whether this is useful if networks use IDRPs as well as IS-IS, since then two levels of hierarchy in IS-IS may be enough.

Integrated IS-IS

There were a number of modifications to Integrated IS-IS discussed at the last meeting. No progress has been made in defining these further:

- Designated router “tenure.”
- Increasing the LSP number space to greater than 256.
- Modifying the metric range to 16 bits for each link and 32 bits for external metrics.
- Increasing the number of pseudonode IDs to greater than 256 (this can be done by applying the multi-area router idea).

Internet-Drafts

The other two subjects discussed at the last meeting, integration of IPX and Appletalk and routing over non-broadcast multi-access (NBMA) circuits, have made progress. There is an Internet-Draft for each. The rest of the meeting was mostly spent discussing these documents.

The group felt that there was sufficient interest to continue working on both documents.

The routing over NBMA Internet-Draft covers operation of Integrated IS-IS, ES-IS, CLNP and IP. It could be applied to any NBMA-style circuit such as SMDS, X.25, Frame Relay or ATM. There was some discussion about how this overlapped with work going on in other IETF working groups and ISO. The IP and CLNP aspects overlap with work done or being done by the IPLPDN and ATM Working Groups. The group felt that we should liaise with the other working groups where there was overlap. The ES-IS aspects would require additions to the ES-IS protocol such as the addition of an IP address option. Like the Integrated IS-IS protocol itself, this modifies a protocol existing as an ISO standard. The group’s modus operandi so far has been, and will continue to be, to carry on extending these protocols in the absence of any formal relationship between the IETF and ISO.

Attendees

| | |
|---------------------|--|
| Ross Callon | <code>rcallon@wellfleet.com</code> |
| Richard Colella | <code>colella@nist.gov</code> |
| Jonathan Davar | <code>jdavar@synoptics.com</code> |
| Kjeld Borch Egevang | <code>kbe@craycom.dk</code> |
| Dino Farinacci | <code>dino@cisco.com</code> |
| Eugene Geer | <code>ewg@cc.bellcore.com</code> |
| Mike Goguen | <code>goguen@synoptics.com</code> |
| Chris Gunner | <code>gunner@dsmail.lkg.dec.com</code> |

| | |
|------------------|----------------------------------|
| Patrick Hanel | hanel@yoyodyne.trs.ntc.nokia.com |
| Susan Hares | skh@merit.edu |
| Denise Heagerty | denise@dxcoms.cern.ch |
| Ton Koelman | koelman@stc.nato.int |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Christian Panigl | christian.panigl@cc.univie.ac.at |
| Georg Richter | richter@uni-muenster.de |
| Rachel Willmer | rachelw@spider.co.uk |

2.5.4 Inter-Domain Multicast Routing (IDMR)

Charter

Chair(s):

Tony Ballardie, A.Ballardie@cs.ucl.ac.uk
Paul Francis, Francis@thumper.bellcore.com

Mailing Lists:

General Discussion: idmr@cs.ucl.ac.uk
To Subscribe: idmr-request@cs.ucl.ac.uk
Archive: cs.ucl.ac.uk:/darpa/idmr-archive.Z

Description of Working Group:

Existing inter-domain multicast routing protocols are not scalable to a large internetwork containing very large numbers of active wide-area groups. The purpose of the IDMR Working Group, therefore, is to discuss proposed inter-domain multicast routing protocols, and put forward one (or a hybrid of several/all) as a Proposed Standard protocol to the IESG.

Several proposals have been made to date, including Core-Based Tree (CBT) multicasting, Core-Based Join (CBJ) multicasting, and Scalable Reverse Path Multicasting (SRPM). Some of the above have yet to be reviewed.

Goals and Milestones:

- | | |
|----------|---|
| Done | Post the Core Based Trees architecture as an Internet-Draft. |
| Nov 1993 | Meet at IETF. All proposals must be submitted by this date. Discuss all proposals which have been submitted. |
| Dec 1993 | Submit the Core Based Trees architecture Internet-Draft to the IESG to be published as an Informational RFC. |
| Jul 1994 | Meet at IETF. Discuss security issues with respect to the proposed protocol(s). |
| Aug 1994 | Post an Internet-Draft for a single protocol (which may be one of the proposals, or a combination of proposals), and an Internet-Draft serving as a protocol analysis document for that protocol (as required by RFC 1264). |
| Jan 1995 | Submit the single protocol to the IESG as a Proposed Standard. |
| Mar 1995 | Post an Internet-Draft for an IDMR MIB. |
| Jul 1995 | Submit the IDMR MIB Internet-Draft to the IESG as a Proposed Standard. |

Internet-Drafts:

“Core Based Trees (CBT) An Architecture for Scalable Inter-Domain Multicast Routing”, 08/03/1992, A. Ballardie, P. Tsuchiya, J. Crowcroft <draft-ballardie-cbt-02.txt>

CURRENT MEETING REPORT

Reported by Tony Ballardie/University College London

Minutes of the Inter-Domain Multicast Routing Working Group (IDMR)

The 27th IETF meeting in Amsterdam was the first official meeting of the IDMR Working Group. The working group met over two 2-hour sessions.

During the first session, Deborah Estrin gave a presentation on ESL (work done in collaboration with Steve Deering, Dino Farinacci, and Van Jacobson), one of the new proposals for inter-domain multicast routing. The motivation behind the design of ESL was, for groups with a relatively small number of senders (sources), to allow receivers to receive data from those sources either over a shared tree, or over a shortest-path tree rooted at the source. The latter is useful for applications requiring minimal delay between senders and receivers. It was agreed that, because ESL is in its early stages of development, there remains specification and engineering details that need to be resolved.

The second session was mostly dedicated to discussing the IDMR charter. It was unanimously agreed that the current charter is lacking with respect to many aspects of inter-domain multicasting, and it should be a goal of the working group to try to resolve many of these, for example, user group management and interoperability.

The conclusion of this discussion was that the charter should be re-worked and re-submitted to the Routing Area Director after the items to be worked on have been enumerated in order of priority.

Attendees

| | |
|-----------------|-------------------------------|
| Toshiya Asaba | asaba@iij.ad.jp |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Dennis Baker | dbaker@wellfleet.com |
| Tony Ballardie | A.Ballardie@cs.ucl.ac.uk |
| Stefan Braun | smb@cs.tu-berlin.de |
| Michael Brescia | brescia@bbn.com |
| John Burnett | jlb@adaptive.com |
| Walid Dabbous | Walid.Dabbous@sophia.inria.fr |
| Stephen Deering | deering@parc.xerox.com |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Hans Eriksson | hans@sics.se |
| Julio Escobar | jescobar@bbn.com |
| Deborah Estrin | estrin@usc.edu |
| Dino Farinacci | dino@cisco.com |
| Dennis Ferguson | dennis@ans.net |

| | |
|-------------------|-----------------------------------|
| Osten Franberg | euaokf@eua.ericsson.se |
| Paul Francis | Francis@thumper.bellcore.com |
| Susan Hares | skh@merit.edu |
| Robert Hinden | hinden@eng.sun.com |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| John Larson | jlarson@parc.xerox.com |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Gerry Meyer | gerry@spider.co.uk |
| Ronny Nilsen | Ronny.Nilsen@usit.uio.no |
| David O'Leary | doleary@cisco.com |
| Jorg Ott | jo@cs.tu-berlin.de |
| James Reeves | jreeves@synoptics.com |
| Benny Rodrig | brodrig@rnd-gate.rad.co.il |
| Tim Seaver | tas@concert.net |
| Ed Stern | els@proteon.com |
| John Stewart | jstewart@cnri.reston.va.us |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Thierry Turletti | turletti@sophia.inria.fr |
| Werner Vogels | werner@inesc.pt |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |

2.5.5 Inter-Domain Policy Routing (IDPR)

Charter

Chair(s):

Martha Steenstrup, msteenst@bbn.com

Mailing Lists:

General Discussion: idpr-wg@bbn.com

To Subscribe: idpr-wg-request@bbn.com

Archive:

Description of Working Group:

The Inter-Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

- Done Write an architecture document.
- Done Draft Protocol Specification of key elements of the protocol.
- Done Develop a prototype implementation of the protocols.
- Done Submit the IDPR Specification to the IESG as a Proposed Standard.

Internet-Drafts:

“Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)”, 07/22/1991, R.A. Woodburn <draft-ietf-idpr-mib-02.txt>

Request For Comments:

- RFC 1126 “Goals and functional requirements for inter-autonomous system routing”
- RFC 1477 “IDPR as a Proposed Standard”
- RFC 1478 “An Architecture for Inter-Domain Policy Routing”
- RFC 1479 “Inter-Domain Policy Routing Protocol Specification: Version 1”

2.5.6 Multicast Extensions to OSPF (MOSPF)

Charter

Chair(s):

John Moy, jmoy@proteon.com

Mailing Lists:

General Discussion: mospf@comet.cit.cornell.edu

To Subscribe: mospf-request@comet.cit.cornell.edu

Archive:

Description of Working Group:

This working group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

- Done Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.
- Done Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.
- Done The Group should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using as an example, the existing BSD OSPF code, written by Rob Coltun of the University of Maryland.
- Done Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations.
- Done Submit the MOSPF Specification to the IESG as a Proposed Standard.

Internet-Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-04.txt, .ps>

"MOSPF: Analysis and Experience", 04/16/1993, J. Moy <draft-ietf-mospf-analysis-02.txt>

Request For Comments:

RFC 1469 "IP Multicast over Token-Ring Local Area Networks"

2.5.7 OSI IDRP for IP Over IP (IPIDRP)

Charter

Chair(s):

Sue Hares, skh@merit.edu

Mailing Lists:

General Discussion: idrp-for-ip@merit.edu

To Subscribe: idrp-for-ip-request@merit.edu

Archive: [merit.edu:~/pub/archive/idrp](http://merit.edu/~pub/archive/idrp)

Description of Working Group:

The IDRP for IP over IP Working Group is chartered to standardize and promote the use of IDRP (ISO Inter-Domain Routing Protocol) as a scalable inter-autonomous system routing protocol capable of supporting policy-based routing for TCP/IP internets. The objective is to take IDRP, as it is defined by ISO standards, and define backward compatible extensions and/or network adaptation layers to enable this protocol to be used in the TCP/IP internets. If any ISO standardization efforts overlap with this area of work, it is intended that the ISO work will supersede the standards proposed by this group.

1) IDRP for IP over IP document (standards track)

This document contains the appropriate adaptations of the IDRP protocol definition that enables it to be used as a protocol for exchange of "inter-autonomous system information" among routers to support forwarding of IP packets across multiple autonomous systems.

2) IDRP MIB document (standards track)

This document contains the MIB definitions for IDRP. These MIB definitions are in two parts; IDRP General MIB, and IDRP for IP MIB. An appendix is planned: IDRP For IP GDMO

3) IDRP - OSPF Interactions (standards track)

This document will specify the interactions between IDRP and OSPF. This document will be based on a combination of the BGP-OSPF interactions document and IDRP - ISIS interactions document.

4) IDRP for IP Usage document (standards track)

Most of the IDRP for IP Usage document will reference the CIDR (supernetting document) Internet-Draft. Any additional terms or protocol definitions needed for IDRP for IP will also be specified here.

Goals and Milestones:

- Done IDRP for IP submitted for Internet-Draft.
- Jun 1992 IDRP MIB document submitted for Internet-Draft.
- Jun 1992 IDRP - OSPF Interactions document submitted for Internet-Draft.
- Jun 1992 IDRP Usage document submitted for Internet-Draft.
- Nov 1992 IDRP for IP submitted to the IESG for Proposed Standard.
- Nov 1992 IDRP Usage document submitted to the IESG for Proposed Standard.
- Nov 1992 IDRP MIB Submitted to the IESG for Proposed Standard.
- Nov 1992 IDRP - OSPF Interactions document submitted to the IESG for Proposed Standard.

Internet-Drafts:

“IDRP for SIP”, 03/22/1993, S. Hares <draft-ietf-ipidrp-sip-00.txt>

CURRENT MEETING REPORT**Minutes of the OSI IDRP for IP Over IP Working Group (IPIDRP)**

The minutes of the joint BGP/IPIDRP sessions follow the BGP charter. The attendee list below is from the joint sessions.

Attendees

| | |
|-------------------|------------------------------|
| Bernt Allonen | bal@tip.net |
| Arun Arunkumar | nak@3com.com |
| Toshiya Asaba | asaba@iij.ad.jp |
| Dennis Baker | dbaker@wellfleet.com |
| Tony Bates | tony@ripe.net |
| Ronald Broersma | ron@nosc.mil |
| Thomas Brunner | brunner@switch.ch |
| Henry Clark | henryc@oar.net |
| David Conrad | davidc@iij.ad.jp |
| Tom Easterday | tom@cic.net |
| Hans Eriksson | hans@sics.se |
| Stefan Fassbender | stf@easi.net |
| Mark Fedor | fedor@psi.com |
| Dennis Ferguson | dennis@ans.net |
| Peter Ford | peter@goshawk.lanl.gov |
| Shoji Fukutomi | fuku@furukawa.co.jp |
| Vince Fuller | vaf@stanford.edu |
| Craig Haney | craig@icp.net |
| Susan Hares | skh@merit.edu |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| David Jacobson | dnjake@vnet.ibm.com |
| J. Jensen | jensen@ic.dk |
| Cyndi Jung | cmj@3com.com |
| Marijke Kaat | marijke@sara.nl |
| Dave Katz | dkatz@cisco.com |
| Sean Kennedy | liam@nic.near.net |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Ton Koelman | koelman@stc.nato.int |
| John Krawczyk | jkrawczy@wellfleet.com |
| Tony Li | tli@cisco.com |
| Robin Littlefield | robin@wellfleet.com |
| Peter Lothberg | roll@stupi.se |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Gerry Meyer | gerry@spider.co.uk |
| Keith Mitchell | keith@pipex.net |

| | |
|----------------------|------------------------------|
| Ramin Najmabadi | najmabadi@helios.iihe.rtt.be |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Michael O'Dell | mo@uunet.uu.net |
| Andrew Partan | asp@uunet.uu.net |
| Willi Porten | porten@gmd.de |
| Juergen Rauschenbach | jrau@dfn.de |
| Yakov Rekhter | yakov@watson.ibm.com |
| Georg Richter | richter@uni-muenster.de |
| Duncan Rogerson | d.rogerson@nosc.ja.net |
| Miguel Sanz | miguel.sanz@rediris.es |
| John Scudder | jgs@merit.edu |
| Henk Steenman | henk@sara.nl |
| Ed Stern | els@proteon.com |
| John Stewart | jstewart@cnri.reston.va.us |
| John Stewart | john@bunter.fdc.iaf.nl |
| Marten Terpstra | marten@ripe.net |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Paul Traina | pst@cisco.com |
| Rene van der Hauw | rene@geveke.nl |
| Rachel Willmer | rachelw@spider.co.uk |
| Jessica Yu | jyy@merit.edu |
| Romeo Zwart | romeo@sara.nl |

2.5.8 Open Shortest Path First IGP (OSPF)

Charter

Chair(s):

John Moy, jmoy@proteon.com

Mailing Lists:

General Discussion: ospfigp@trantor.umd.edu

To Subscribe: ospfigp-request@trantor.umd.edu

Archive:

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

- | | |
|------|--|
| Done | Design the routing protocol, and write its specification. |
| Done | Develop multiple implementations, and test against each other. |
| Done | Obtain performance data for the protocol. |
| Done | Make changes to the specification (if necessary) and publish the protocol as a Draft Standard RFC. |
| TBD | Gather operational experience with the OSPF protocol and submit the document as a Standard. |

Internet-Drafts:

“OSPF Version 2”, 11/11/1992, J. Moy <[draft-ietf-ospf-version2-04.txt](#), .ps>

“The OSPF External Attributes LSA”, 03/23/1993, D. Ferguson <[draft-ietf-ospf-extattr-00.txt](#)>

“Guidelines for Running OSPF Over Frame Relay Networks”, 05/03/1993, O. deSouza, M. Rodrigues <[draft-ietf-ospf-guidelines-frn-00.txt](#)>

Request For Comments:

RFC 1131 “OSPF specification”

RFC 1245 “OSPF Protocol Analysis”

- RFC 1246 “Experience with the OSPF Protocol”
- RFC 1247 “OSPF Version 2”
- RFC 1248 “OSPF Version 2 Management Information Base”
- RFC 1252 “OSPF Version 2 Management Information Base”
- RFC 1253 “OSPF Version 2 Management Information Base”

2.5.9 RIP Version II (RIPv2)

Charter

Chair(s):

Gary Malkin, gmalkin@xylogics.com

Mailing Lists:

General Discussion: ietf-rip@xylogics.com

To Subscribe: ietf-rip-request@xylogics.com

Archive: xylogics.com:gmalkin/rip/rip-arc

Description of Working Group:

RIP Version 2 and the Version 2 MIB was approved as a Proposed Standard in January 1993. They were published as RFC 1388 and RFC 1389. Since the minimum required period has elapsed for a protocol to remain as a Proposed Standard, RIP V2 can now be considered for advancement to Draft Standard.

The RIP Version 2 Working Group will prepare a recommendation to the IESG evaluating the standards track status of RIP Version 2 and the RIP Version 2 MIB. The recommendation will document implementation, interoperability and deployment experience as required by RFC 1264 "Routing Protocol Criteria."

This group is chartered to prepare revisions of RFC 1388, RIP Version 2, RFC 1389, the RIP Version 2 MIB, and RFC 1387, analysis of the protocol if necessary.

The RIP Version 2 Working Group is further chartered to evaluate the proposal for "Routing over Demand Circuits using RIP" for standards track consideration.

Goals and Milestones:

- | | |
|----------|--|
| Done | Review of RIP-II Internet-Draft to ensure the additions are useful and backwards compatible. Also ensure that the additions cannot cause routing problems. |
| Done | Final review of RIP-II Internet-Draft and submission into the standards track. First review of RIP-II MIB. |
| Done | Review of implementations. Final review of MIB. |
| Jul 1993 | Hold working group meetings to review RIP Version 2 implementations and make any changes needed to the specifications. |
| Jul 1993 | Review the RIP over Demand Circuits Internet-Draft. |
| Aug 1993 | Submit the RIP over Demand Circuits to the IESG for consideration as a Proposed Standard. |

- Nov 1993 Post as an Internet-Draft a report describing the implementation and operational experience of the RIP v2 protocol in accordance with the RFC 1264 "Routing Protocol Criteria."
- Mar 1994 Submit the RIP Version 2 protocol to the IESG for consideration as a Draft Standard.

Request For Comments:

- RFC 1387 "RIP Version 2 Protocol Analysis"
- RFC 1388 "RIP Version 2 Carrying Additional Information"
- RFC 1389 "RIP Version 2 MIB Extension"

CURRENT MEETING REPORT

Reported by Gary Malkin/Xylogics

Minutes of the RIP Version II Working Group (RIPV2)

Agenda

- Review charter
- Review RFC 1388 (protocol spec)
- Review RFC 1389 (MIB)
- Review implementation experience
- Review the Demand Routing specification
- Summary of decisions and actions

Summary

The charter was approved as written.

The use of the Routing Domain in RIP-2 is still unclear. It was determined that the use of the field could not be sufficiently well defined to meet the varying needs of those few people who would like to use it. The field also poses difficult MIB problems (discussed below). Therefore, it has been decided to remove the field from the protocol and leave a must-be-zero field in its place. Presumably, a motivated person could propose a third version of RIP which would define a use for this field. This change does not, to the knowledge of those attending the meeting, invalidate any existing implementations and may therefore be made without requiring the specification to remain at the Proposed Standard level.

There were two proposed changes to the MIB. The first was to deprecate the Routing Domain object. It has been pointed out that the tables cannot be indexed correctly unless the Routing Domain object was used as part of the index. Given that the Routing Domain field is not well defined, this change will result in an overall simplification of the MIB. The second proposal dealt with handling unnumbered interfaces. While the RIP-2 protocol does not expressly address them, their existence does require consideration since the MIB tables cannot be indexed properly with unnumbered interfaces. The proposal is to use a network number of zero and a host number of if_index to create a suitable IP address for use in indexing tables. These changes do not, to the knowledge of those attending the meeting, invalidate any existing implementations and may therefore be made without requiring the specification to remain at the Proposed Standard level.

There are currently two independent implementations of RIP-2: gated and Xylogics's routed. The MIB has been implemented for gated. ACC has a partial implementation of RIP-2 and is planning to implement the remainder.

Gerry Meyer's Demand Routing proposal was discussed at length. It was agreed that it performed a useful function. However, Robert Ullman pointed out that it simulated many of the functions of TCP and that RAP used TCP. Robert and Gerry will continue the discussion after Gerry has had an opportunity to read the RAP RFC. The slides from Gerry's presentation follow these minutes.

Attendees

| | |
|----------------------|------------------------------|
| Steve Alexander | stevea@lachman.com |
| Michael Anello | mike@xlnt.com |
| Dennis Baker | dbaker@wellfleet.com |
| Fred Baker | fbaker@acc.com |
| Dave Cullerot | cullerot@ctron.com |
| Bob Downs | bdowns@combinet.com |
| Dino Farinacci | dino@cisco.com |
| Vince Fuller | vaf@stanford.edu |
| David Jacobson | dnjake@vnet.ibm.com |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Tony Li | tli@cisco.com |
| Susan Lin | suelin@vnet.ibm.com |
| Paolo Malara | malara@crs4.it |
| Gary Malkin | gmalkin@xylogics.com |
| Gerry Meyer | gerry@spider.co.uk |
| Keith Mitchell | keith@pipex.net |
| Ramin Najmabadi | najmabadi@helios.iihe.rtt.be |
| Willi Porten | porten@gmd.de |
| Juergen Rauschenbach | jrau@dfn.de |
| Benny Rodrig | brodrig@rnd-gate.rad.co.il |
| A. Velu Sinha | avsinha@attmail.com |
| Paul Traina | pst@cisco.com |
| Robert Ullmann | ariel@world.std.com |
| Douglas Williams | dougw@ralvmg.vnet.ibm.com |

Routing Over Demand Circuits - RIP

- The problem with routing protocols.
- Routing information integrity.
- Intelligent circuit manager.
- Presumption of reachability.
- Implementations.

Routing Over Demand Circuits - RIP

The problem with routing protocols:

- Periodic updates (, or hellos or flooding).
- SVC connection open ~ permanently.
- Cost of a connection (duration).
- Cost of data (X.25).
- Mesh of connections ~ $O(N^{**2})$.

Routing Over Demand Circuits - RIP

Routing information integrity:

- Triggered updates of 'complete' database
- Guaranteed delivery:
 - New packet types with sequence and fragment number.
 - Individual fragments acknowledged.
 - Fragments re-transmitted until acknowledged.
 - Routing database not updated until all fragments are received.
- Retention of alternative routes.

Routing Over Demand Circuits - RIP

Intelligent circuit manager:

- Able to determine unreachability in a predictable finite period.
- Internal circuit DOWN and circuit UP messages
 - DOWN indicates no SVC, and unable to get one for ANY datagram.
 - UP (sent after a down) indicates SVC established.
- After DOWN, circuit manager is responsible for re-establishing connection.

Routing Over Demand Circuits - RIP

Presumption of reachability:

- Unless evidence to the contrary:
 - Don't time out routes learned from the WAN.
- Evidence to the contrary:
 - Routing entry missing from an update.
 - Routing entry received with infinite metric.
 - Circuit DOWN indication from circuit manager.
 - (Eventual) failure to receive ACK of update.

Routing Over Demand Circuits - RIP

Implementations:

- Spider: IP RIP 1, IPX RIP and SAP over ISDN-PPP and X.25 (gerry@spider.co.uk).
- XtcN: IPX RIP and SAP over ISDN-HDLC (lamb@xtcn.com).
- Others??

2.5.10 Source Demand Routing (SDR)

Charter

Chair(s):

Deborah Estrin, estrin@usc.edu
Tony Li, tli@cisco.com

Mailing Lists:

General Discussion: sdrp@caldera.usc.edu
To Subscribe: sdrp-request@caldera.usc.edu
Archive: [jerico.usc.edu:~/pub/sdrp](http://jerico.usc.edu/~pub/sdrp)

Description of Working Group:

The SDR Working Group is chartered to specify and promote the use of SDRP (Source Demand Routing Protocol) as an inter-domain routing protocol capability in conjunction with IDRP and BGP inter-domain routing protocols. The purpose of SDR is to support source-initiated selection of inter-domain routes, to complement the intermediate node selection provided by BGP/IDRP.

The goal of the SDR Working Group is to release the components of SDR as IETF Prototypes and to obtain operational experience with SDR in the Internet. Once there is enough experience with SDR, the working group will submit the SDR components to the IESG for standardization.

SDR has four components: packet formats for protocol control messages and encapsulation of user datagrams, processing and forwarding of user data and control messages, routing information distribution/collection and route computation, and configuration and usage.

The group's strategy is to:

1. Define the format, processing and forwarding of user datagram and control messages so that SDR can be used very early on as an efficient means of supporting "configured" inter-domain routes. User packets are encapsulated along with the source route and forwarded along the "configured" route. Routes are static at the inter-domain level, but are not static in terms of the intra-domain paths that packets will take between specified points in the SDR route. The impact of encapsulation on MTU, ICMP, performance, etc., are among the issues that must be evaluated before deployment.
2. Develop simple schemes for a) collecting dynamic domain-level connectivity information, and b) route construction based on this information, so that those domains that want to can make use of a richer, and dynamic set of SDR routes.
3. In parallel with 1 and 2, develop usage and configuration documents and prototypes that demonstrate the utility of static-SDR and simple-dynamic-SDR.

4. After gaining some experience with the simple schemes for distribution, develop a second generation of information distribution and route construction schemes. The Group hopes to benefit from discussions with IDPR and NIMROD developers at this future stage because the issues faced are similar.
5. The Group will also investigate the addition of security options into the SDRP forwarding and packet format specifications.

Goals and Milestones:

- Mar 1993 Post an Internet-Draft of packet forwarding and control message format and protocol for IP.
- Jun 1993 Post as an Internet-Draft the SDR MIB.
- Jun 1993 Post as an Internet-Draft the SDR Usage and Configuration document. This is the highest priority after the draft specification in order to demonstrate how even static-SDR can be used to achieve concrete objectives.
- Sep 1993 Post as an Internet-Draft the BGP/IDRP Extensions Specification. As mentioned in the Internet Draft there are a few extensions to BGP/IDRP needed to support SDR. These must be detailed and documented.
- Done Submit as an Internet-Draft a specification for Route Setup.
- Nov 1993 Post as an Internet-Draft a SDR Deployment Plan.
- Dec 1993 Post as an Internet-Draft a document describing the distribution/acquisition of Information to construct richer SDR routes. The initial versions of SDR will use only configured information (some of which may be derived from BGP/IDRP) as the basis for constructing source routes.
- Dec 1993 Post as an Internet-Draft a specification for SDR Multicast.
- Mar 1994 Submit the set of SDR specifications to the IESG for consideration as a Proposed Standard.
- Mar 1994 Submit the set of SDR specifications to the IESG for consideration as a Prototype protocol.

Internet-Drafts:

“Source Demand Routing Policy Language”, 06/21/1993, T. Li <draft-ietf-sdr-pl-00.txt>

“Source Demand Routing: Route Setup”, 06/23/1993, D. Estrin, D. Zappala, T. Li <draft-ietf-sdr-route-setup-00.txt>

“BGP SDRP_SPEAKERS Attribute”, 09/13/1993, K. Varadhan <draft-ietf-sdr-speakers-attribute-00.txt>

CURRENT MEETING REPORT

Reported by Peter Ford/Los Alamos National Laboratory

Minutes of the Source Demand Routing Working Group (SDR)

Tony Li opened the meeting and bashed the agenda into shape. It was subsequently dynamically reordered.

Deborah Estrin gave an overview of SDRP, noting the specification for SDRP has not changed since the last meeting. She requested that people read the specification and provide more comments.

Deborah reported on a prototype implementation by Daniel Zapalla at USC based on the SunOS DARTnet kernel. Tests were conducted on a small 4-node testbed at USC, and on DARTnet. There is a kernel interface establishing source routes, filtering and encapsulation. There is a routing socket interface for the D-FIB.

USC is interested in seeing more people pick up the code and build experimental testbed islands, and then interconnecting them for later interdomain experimentation.

Christian Huitema asked about anycast, and Yakov noted that the AS number is a group address. Further work remains to be done in defining the requirements for anycasting, but there appears to be nothing in the SDRP specification that would be a substantial limitation.

Tony Li reported on his work on BGP/IDRP interaction with SDRP. The original idea was to route SDRP control messages to AS representatives (e.g. to query AS for policy). This functionality is possible by changing routing of control messages, with the last address being an AS address. Thus the BGP attribute is no longer needed for this functionality. Tony will fix the specification to reflect this change. The IDRP attribute is still useful for tunneling. Yakov has a draft of this attribute and will post it as an Internet-Draft shortly.

Tony reported on the policy language he has been working on. It is C-like. Tony would like to get more comments on the specification. The evaluation of policy is boolean. The functionality will be familiar to people who are familiar with cisco access lists. The OR operator (distinct from ||) can be used to ignore a term which cannot be evaluated (bottom). Christian Huitema suggested that the access control aspects of this language should be checked against the work on IP security. Future work will include policy based on source and destination AS number (using the DNS?), source and destination communities, and there may be some work on richer pattern matching on the entire SDRP route if there is a need. There could be some work on time varying characteristics, such as load or delay. Steve Hotz sent mail suggesting that Tony think about evaluating policy terms to preference continuous values instead of boolean values.

Tony reported on his work on Information Distribution. The plan is to get the current AS topology from some static source. It was suggested to see RIPE-81 for an expression of this information. Christian noted that we probably need to consider higher levels of aggregation of this information above ASs.

Deborah stated that she needs to do a rework of the futures document which deals with scaling issues. She will update the document.

Deborah reported on SDRP setup work. Setup is done via an explicitly source routed packet with the probe bit set. The motivation for these setups is to reduce the header size. It is not a requirement for a router to participate. It can strip the probe bit, send it ahead, and send a setup rejected message back to the originator.

There was a question if the routers have to maintain information on source routes that it is currently setup for: the answer was yes. It was again noted that this is soft state.

There is a draft specification of SDRP Setup that will be distributed later in the summer.

Christian stated that setup needs to be investigated to see what its interactions are with regard to load splitting.

Deborah noted additional work that will be considered, including multicast (and its relationship to ESL). This would be used to establish branches on multicast trees. There also needs to be better tools for building SDRP routes.

Editor's Note: A list of assigned tasks is available via FTP or mail server from the remote directories as /ietf/sdr/sdr-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Attendees

| | |
|-----------------|----------------------------|
| Michael Anello | mike@xlnt.com |
| Toshiya Asaba | asaba@ij.ad.jp |
| Cynthia Bagwell | cbagwell@gateway.mitre.org |
| Dennis Baker | dbaker@wellfleet.com |
| John Ballard | jballard@microsoft.com |
| John Burnett | jlb@adaptive.com |
| Ross Callon | rcallon@wellfleet.com |
| Henry Clark | henryc@oar.net |
| Francis Dupont | francis.dupont@inria.fr |
| Deborah Estrin | estrin@usc.edu |
| Dino Farinacci | dino@cisco.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Vince Fuller | vaf@stanford.edu |

| | |
|-----------------------|------------------------------|
| Susan Hares | skh@merit.edu |
| Denise Heagerty | denise@dxcoms.cern.ch |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| Nandor Horvath | horvath@sztaki.hu |
| David Jacobson | dnjake@vnet.ibm.com |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Tony Li | tli@cisco.com |
| Robin Littlefield | robin@wellfleet.com |
| Jun Matsukata | jm@eng.isas.ac.jp |
| Keith Mitchell | keith@pipex.net |
| Ramin Najmabadi | najmabadi@helios.iihe.rtt.be |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Erik Nordmark | nordmark@eng.sun.com |
| Petri Ojala | ojala@eunet.fi |
| Juergen Rauschenbach | jrau@dfn.de |
| James Reeves | jreeves@synoptics.com |
| Yakov Rekhter | yakov@watson.ibm.com |
| Duncan Rogerson | d.rogerson@nosc.ja.net |
| Shawn Routhier | sar@epilogue.com |
| John Scudder | jgs@merit.edu |
| John Stewart | jstewart@cnri.reston.va.us |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Paul Traina | pst@cisco.com |
| Rene van der Hauw | rene@geveke.nl |
| Willem van der Scheun | scheun@sara.nl |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |
| Jessica Yu | jyy@merit.edu |

2.6 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker/TIS and Jim Galvin/TIS

The Security Area within the IETF is responsible for development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet. The appendix to this report defines what is meant by “security” on the Internet.

Much of the work of the Security Area is performed in coordination with working groups in other areas. The Security Area Advisory Group (SAAG) is a group of security experts which provides both consulting help to other areas and direct management of working groups within the Security Area.

The main bulk of the work for the SAAG consists of a set of formal work items. These work items correspond to working groups within the IETF Security Area, security-relevant developments within working groups in areas other than security, and internal SAAG work items which do not merit the creation of formal working groups but which do need some level of attention.

Below is the status of each of the working groups officially chartered or initiated within the Security Area. Immediately following those reports is an update on other security issues as well as security related work in other IETF areas.

Authorization and Access Control Working Group (AAC)

The AAC Working Group met on Wednesday afternoon. The need for demonstrations of the technology was discussed; possible applications for an authorization API include Prospero, remote execution, database access, and file transfer. The representation of authorization information for access control lists and distributed authorization credentials was discussed. There was some contention on this topic.

At the SAAG meeting it was suggested that some of the work arising from this working group may be best addressed initially by the PSRG. Cliff will follow-up on this.

Common Authentication Technology Working Group (CAT)

The CAT Working Group met for two sessions at the Amsterdam IETF, discussing (in about equal proportion) general CAT issues and FTP security integration. We reviewed the status of implementation and specification activities, identified items requiring follow-up work, and managed to associate individuals and subset groups with several of the desired items.

The GSS-API base specification, GSS-API C Language Bindings, and Kerberos Version 5 documents have been submitted, adopted, and published by the IESG as Proposed Standards.

The DASS document has been submitted, adopted, and published by the IESG as an Experimental Protocol.

Commercial Internet Protocol Security Option Working Group (CIPSO)

There is a new draft of the CIPSO specification; Steve Crocker will make sure it gets published as an Internet-Draft.

There are several folks who are unsatisfied with this document, including both SAAG and PSRG members. There has been some difficulty getting the issues communicated effectively to Ron Sharp. Steve Crocker has been tasked with resolving the conflicts.

Internet Protocol Security Protocol Working Group (IPSEC)

The meeting was opened by Al Hoover with an introduction to the IPSEC Working Group for first time attendees, a review of the approved IPSEC charter, a review of liaisons between IPSEC and IPng, IEEE, CAT, and other working groups. A brief discussion of preliminary implementations related to IPSP was discussed. Absences of IPSEC implementors limited the scope to a review of the various approaches (SwIPe, NLSP, SP3).

IPSEC would like to target demonstrations of preliminary implementations (non-interoperable) for the Houston (November) IETF. Demonstrations of preliminary interoperable implementations is targeted for the March IETF.

Network Access Server Requirements Working Group (NASREQ)

The NASREQ Working Group meeting was sparsely attended on Thursday morning. Lacking critical mass, there was little discussion about any of the outstanding issues. The modified charter and schedule was reviewed. The chair requested volunteers to help write

sections of the draft. (One person expressed willingness to help, provided he received approval from his management.) Representatives from the ACCT Working Group and the AAC Working Group requested more contact and information flow between the groups.

Privacy-Enhanced Electronic Mail Working Group (PEM)

The PEM Working Group met for one, 2-hour session on Wednesday afternoon. The meeting opened with reports on implementation status from seven PEM developers, representing commercial, research, and academic communities.

Next, the meeting addressed the continuing topic of MIME-PEM integration. No substantive progress was made on this topic, due to a lack of written submissions for review prior to the meeting. Nonetheless, there was a discussion of this topic, based on a very recent discussion between two of the authors of the relevant Internet-Draft. A presentation on the use of Distinguished Names versus Domain Name System names in certificates followed, but was truncated because of time limitations.

A presentation on the rationale for the current certification system design was skipped, also due to time constraints. Full text of the slides for both of these presentations will be included as an appendix to the meeting minutes. The meeting concluded with a discussion of triple-DES modes of use for PEM, and a paper exploring this issue was distributed at the meeting. While there was substantial sentiment for one of the modes, it was agreed that further analysis is called for.

SNMP Security Working Group (SNMPSEC)

In conjunction with the SNMPv2 Working Group, twelve documents have been completed and published as Proposed Standards. This work item was officially closed at the Amsterdam IETF.

TELNET Working Group (TELNET) - Applications Area

There is no security-relevant progress to report from the Amsterdam IETF.

Router Requirements Working Group (RREQ) - Internet Area

The previous single document has been split into four documents and a number auxiliary documents. Phil Almquist has responsibility for finishing the documents and submitting them to the IESG for publication.

Domain Name System Working Group (DNS) - Service Applications Area

A mailing list and subcommittee of the DNS Working Group has been created. Work on DNS security is expected to begin on the mailing list.

Trusted Network File System Working Group (TNFS) - Service Applications

The TNFS Working Group meets principally under the auspices of the Trusted Systems Interoperability Group.

No progress to report.

Audio/Video Transport Working Group (AVT) - Transport Area

This activity was to be reviewed to identify the security issues for the Amsterdam meeting.

No progress to report.

Integrated Directory Services Working Group (IDS) - User Services Area

Privacy constraints exist for the data, but there was no substantial discussion at this time.

Export Control Issues

Vint Cerf and Steve Crocker need to press forward on drafting a document.

IP: The Next Generation

A plan for processing a security review of the competing next generation proposals were to be drafted for the Amsterdam meeting.

No progress to report.

IRC

A document in RFC format exists that purports to document this protocol. At this time, this work item exists to track this activity.

ITAR Publication

An on-line version of the US International Traffic in Arms Regulations (ITAR) will be created as soon as it has been published in the Federal Register, probably as an informational RFC.

Key Management Strategies

A review of key management strategies and activities was to be drafted for the Amsterdam meeting.

No progress to report. At the SAAG meeting it was asked why this work item is called out separately from the IP security work item. John will be asked to address this question in his draft.

Mobile IP Security

John Ioannidis was to report on the relationship between this work item and the IP security work item at the Amsterdam SAAG meeting; he was not present at this meeting. Al Hoover will follow-up with John on the status of this work item.

Random Number Generation Issues

A revised Internet-Draft is overdue. Jeff Schiller will follow-up with Don Eastlake to arrange a new draft.

Routing Security Plan

No progress to report. Steve Crocker will follow-up with Radia Perlman.

Security Area Architecture

Editor's Note: A summary description of the Security Area is available via FTP or mail server from the remote directories as /ietf/93jul/area.security.93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Working Group Liaison Checklist

A checklist was prepared and distributed at the SAAG meeting. A copy will be distributed to the SAAG Interest mailing list for discussion.

2.6.1 Authorization and Access Control (AAC)

Charter

Chair(s):

Clifford Neuman, bcn@isi.edu

Mailing Lists:

General Discussion: ietf-aac@isi.edu

To Subscribe: ietf-aac-request@isi.edu

Archive: prospero.isi.edu:~/pub/aac/*

Description of Working Group:

The goal of the Authorization and Access Control Working Group is to develop guidelines and an Application Program Interface (API) through which network accessible applications can uniformly specify access control information. This API will allow applications to make access control decisions when clients are not local users, might not be members of a common organization, and often not known to the service or application in advance.

Several authentication mechanisms are in place on the Internet, but most applications are written with local applications in mind and no guidelines exist for supporting authorization and access control based on the output of such authentication mechanisms. The CAT working group developed the GSS-API, a common API to support authentication. The AAC Working Group will develop a common API that accepts the identity of a client (perhaps the output of the GSS-API), a reference to an object to be accessed, and optionally an indication of the operation to be performed. The API will return a list of authorized operations or a yes/no answer that can be easily used by the application.

A second, longer term purpose of the working group will be to examine evolving mechanisms and architectures for authorization in distributed systems and to establish criteria which enable interworking of confidence and trust across systems. The working group will develop additional goals and milestones related to this purpose and will submit a revised charter once the appropriate goals and milestones are determined. To the extent possible this additional work will encourage evolution toward credential formats that more readily allow support for or translation across multiple mechanisms.

Goals and Milestones:

- | | |
|------|--|
| Done | Submit charter and milestones for approval. |
| Done | Meet at the Columbus IETF to identify common characteristics of evolving distributed authorization mechanisms and begin discussion of approaches for interoperability across mechanisms. |

- Jun 1993 Post draft API as an Internet-Draft.
- Jun 1993 Post an Internet-Draft of the guidelines for authorization and access control for network accessible applications.
- Aug 1993 Submit the AAC guidelines document for approval as an Informational RFC.
- Jan 1994 Submit the AAC API for consideration as an Experimental RFC.

CURRENT MEETING REPORT

Reported by B. Clifford Neuman/Information Sciences Institute

Minutes of the Authorization and Access Control WG (AAC)

The Authorization and Access Control Working Group met at the July IETF for the first time since the approval of its charter and official inception as a working group. The preceding three meetings were BOF sessions.

The charter, past minutes, mailing list discussions, and other documents mentioned in these minutes are available by anonymous FTP from prospero.isi.edu in the directory /pub/aac.

Agenda

- Report on approval of the AAC charter.
- Presentation of a list of restrictions and privilege attributes needed by applications and existing security systems, and a proposed method for representing them.
- Discussion of the intended use of these restrictions by applications, and the presentation of an Application Program Interface (API) to provide a simple interface for application developers.
- Discussion of the information maintained in the security context. The security context maintains information about the user that is used to make authorization decisions.

Report on the AAC Charter

It was reported that the working group charter was approved by the IESG. Steve Crocker brought up several desires that were raised in the discussion by the IESG. Among these desires is the need for some kind of demonstration of the technology, in particular, integration with possible applications.

Discussion of Possible Applications (Digression)

Possible applications were discussed. An early test will be the Prospero Directory Service which already has support for access control lists, and an access control list type reserved for the mechanism developed by the working group.

Another possible application is in support of cross-site, remote execution. In particular, Tom Hutton is looking for a simple way to specify access controls for data and processing resources distributed across several sites.

File transfer provides a third set of applications. Steve Crocker pointed out the need for secure file transfer to and between large diverse groups. This is related to the FTP extension work in the Common Authentication Technology Working Group (CAT) in that those extensions make available to the application the authenticated network identity of the client, and that identity might be used as a basis for authorization decisions. Some of the Washington University FTP daemon extensions are also of interest here.

A final application that was discussed was network databases. Daisy Rose mentioned that the Network Database Working Group (NETDATA) has a need for an authorization mechanism that will allow them to determine which remote principals are authorized to access a database, and which local user ID is to apply to such remote accesses.

How It Will Be Used by Applications

Throughout the discussion of possible applications, the issue of how authorization information would be specified by applications was raised. There seemed to be two classes: applications that are aware of network identities, and those that are not.

Applications that are not aware of network identities rely on local authorization using local user identities. A separate mechanism is used to map network identities to local identities. For such applications, authorization is confined to initially establishing who is authorized to assume a particular identity at the time a connection is initiated. It is not clear if this is an authentication issue or an authorization issue.

Applications that are aware of network identities make a call to the authorization API for each operation that is to be mediated. The authorization API will return a yes or no answer, or a list of what the principal can do, based on the principal's network identity.

Access control list entries could identify the type of authentication required, in addition to the name of the principal authorized by an entry. Sam Sjogren suggested allowing the specification of weaker authentication methods including regular expression matches on network address or hostname and usernames in addition to stronger methods. This would allow the authorization API to be used with an existing application that does not have support for strong authentication, and would allow easier transition to stronger mechanisms if they are later integrated into the application.

There was a brief discussion about whether an administrative interface to maintain access control lists needs to be defined. This issue was deferred until it is decided what access control lists and the API for checking authorization will look like. The definition of an external representation for an access control list should be enough to get started.

Presentation of Restrictions and Privilege Attributes

A draft list of restrictions was distributed at the meeting. The list defines some common restrictions that are useful for representing privilege attributes and constraints on the use of credentials in distributed systems.

Several restrictions were discussed. Sam Sjogren suggested that it might be useful to think of these in terms of the questions who, what, when, where, and how (why is more appropriate for audit than authorization). With this taxonomy, the restrictions discussed were:

- who - `for_use_by_principal`, `for_use_by_group`;
- what - `local_uid`, `group_membership`, `dce_pac`, `authorized`, `quota`, `netmask`;
- when - `accept_n_times`, `authorized_times`;
- where - `for_use_on_server`, `limit_restriction`, `limit_application`; and
- how - `connection_type` (dial-in, hard-wired from a secure area, etc), `application_name`.

Even with this breakdown, there was a great deal of confusion about the difference between the “who” restrictions which limit who may exercise the proxy, and the “what” restrictions that seem to assert local user IDs and group membership, instead of restrict them. It is clear from the discussion that the model needs to be refined so that this distinction is more understandable, or replaced so that positive and negative attributes are considered separately.

During discussion after the meeting, some ideas for addressing this confusion were generated. A revised specification incorporating one of these ideas will be distributed to the mailing list by the third week of August, and it will be decided at that time if the concerns have been addressed.

Discussion of the Security Context

In the few minutes that remained, Piers McMahon discussed possible information to be included in the security context, a structure that stores information about a principal and is passed as input to the authorization API which uses it, to decide which access control list entries are applicable. The presentation outlined the security-relevant information about a session maintained by, exported by, or used by several systems.

The Generic Security Service Application Program Interface (GSS-API) supports authentication and message protection. Separate authorization mechanisms provide access mediation and enforcement. The network user identity authenticated by the GSS-API is part of the security context and can be used by the authorization API.

In the OSF Distributed Computing Environment, a set of privileges are added to the security context. These privileges are securely transmitted in privilege attribute certificates signed using Kerberos. These privileges become part of the security context once validated by the end-server.

The security context for Sesame includes privilege attributes and control attributes that can limit delegation and permissible targets. Max Six includes labels and audit IDs in the security contexts.

Representation of attributes is likely to be needed in a security context used for access control. It is recommended that the GSS-API security context be extended to include privilege attributes. John Linn pointed out that if this is done, a set of widely accepted attributes will be needed.

Thanks to Richard Graveman for his notes which were helpful in the preparation of these minutes.

Attendees

| | |
|-----------------------|--------------------------------|
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Mohammad Alaghebandan | mohammad_alaghebandan@3com.com |
| Luc Boulianne | lucb@cs.mcgill.ca |
| Stefan Braun | smb@cs.tu-berlin.de |
| Stephen Crocker | crocker@tis.com |
| Kurt Dobbins | dobbins@ctron.com |
| Richard Graveman | rfg@ctt.bellcore.com |
| Neil Haller | nmh@thumper.bellcore.com |
| Alton Hoover | hoover@ans.net |
| Thomas Hutton | hutton@opus.sdsc.edu |
| Dale Johnson | dsj@merit.edu |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Kim Chr. Madsen | kimcm@ic.dk |
| Clifford Neuman | bcn@isi.edu |
| Mel Pleasant | pleasant@hardees.rutgers.edu |
| Lars Poulsen | lars@cmc.com |
| Robert Reschly | reschly@brl.mil |
| Michael Roe | mrr@cl.cam.ac.uk |
| Daisy Rose | daisy@watson.ibm.com |
| Wolfgang Schneider | schneiw@darmstadt.gmd.de |
| Heiner Schorn | heiner.schorn@umdac.umu.se |
| Sam Sjogren | sjogren@tgv.com |
| James Zmuda | zmuda@mls.hac.com |

2.6.2 Commercial Internet Protocol Security Option (CIPSO)

Charter

Chair(s):

Ron Sharp, rls@neptune.att.com

Mailing Lists:

General Discussion: cipso@wdl1.wdl.loral.com

To Subscribe: cipso-request@wdl1.wdl.loral.com

Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Commercial Internet Protocol Security Option Working Group is chartered to define an IP security option that can be used to pass security information within and between security domains. This new security option will be modular in design to provide developers with a single software environment which can support multiple security domains.

The CIPSO protocol will support a large number of security domains. New security domains will be registered with the Internet Assigned Numbers Authority (IANA) and will be available with minimal difficulty to all parties.

There is currently in progress another IP security option referred to as IPSO (RFC 1108). IPSO is designed to support the security labels used by the US Department of Defense. CIPSO will be designed to provide labeling for the commercial, US civilian and non-US communities.

The Trusted Systems Interoperability Group (TSIG) has developed a document which defines a structure for the proposed CIPSO option. The working group will use this document as a foundation for developing an IETF CIPSO specification.

Goals and Milestones:

- | | |
|----------|--|
| Ongoing | Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard. |
| Done | Review and approve the charter for the IETF CIPSO Working Group. Review revised TSIG CIPSO Specification. |
| Done | Review outstanding comments/issues from mailing list. Continue work on specification and prepare it for submission as an Internet-Draft by the end of May. |
| Jul 1991 | Review outstanding comments/issues from mailing list. The specification will be submitted to the IESG for consideration as a Proposed Standard. |

Mar 1992 Submit specification to the IESG for consideration as a Draft Standard. There must be at least two interoperable implementations by this time.

Internet-Drafts:

“COMMON IP SECURITY OPTION”, 03/10/1993, R. Sharp <draft-ietf-cipso-ipsec-option-00.txt>

2.6.3 Common Authentication Technology (CAT)

Charter

Chair(s):

John Linn, linn@gza.com

Mailing Lists:

General Discussion: cat-ietf@mit.edu

To Subscribe: cat-ietf-request@mit.edu

Archive: [bitsy.mit.edu:~/cat-ietf/archive](http://bitsy.mit.edu/~cat-ietf/archive)

Description of Working Group:

The goal of the Common Authentication Technology Working Group is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security-oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accommodating evolution of underlying technologies.

In support of these goals, the working group will pursue several interrelated tasks. We will work towards agreement on a common service interface allowing callers to invoke security services, and towards agreement on a common authentication token format, incorporating means to identify the mechanism type in conjunction with which authentication data elements should be interpreted. The CAT Working Group will also work towards agreements on suitable underlying mechanisms to implement security functions; two candidate architectures (Kerberos V5, based on secret-key technology and contributed by MIT, and X.509-based public-key Distributed Authentication Services being prepared for contribution by DEC) are under current consideration. The CAT Working Group will consult with other IETF working groups responsible for candidate caller protocols, pursuing and supporting design refinements as appropriate.

Goals and Milestones:

| | |
|------|--|
| Done | Progress Internet-Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms. |
|------|--|

- Done Preliminary BOF session at IETF meeting, discussions with TELNET and Network Printing Working Groups.
- Done Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet-Draft process.
- Done First IETF meeting as full working group: review charter distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.
- Done Update mechanism-independent Internet-Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.
- Done Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.
- Done Submit service interface specification to to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“FTP Security Extensions”, 04/06/1993, S. Lunt <draft-ietf-cat-ftpsec-02.txt>

Request For Comments:

- RFC 1507 “Distributed Authentication Security Service”
- RFC 1508 “Generic Security Service Application Program Interface”
- RFC 1509 “Generic Security Service API : C-bindings”
- RFC 1510 “The Kerberos Network Authentication Service (V5)”
- RFC 1511 “Common Authentication Technology Overview”

CURRENT MEETING REPORT

Reported by John Linn/Geer Zolot Associates and Sam Sjogren/TGV

Minutes of the Common Authentication Technology Working Group (CAT)

The CAT Working Group met for two sessions at the Amsterdam IETF, discussing (in about equal proportion) general CAT issues and FTP security integration.

Review of CAT Activities

We reviewed the status of CAT-related Internet-Drafts: “Generic Security Service Application Program Interface” (GSS-API) and “Generic Security Service API : C-bindings” are in the hands of the RFC Editor pending advancement to Proposed Standard status, as is “The Kerberos Network Authentication Service (V5).” The Kerberos V5 GSS-API implementation has not received recent development effort, and is not currently compliant, but a plan to make volunteer resources available is being explored.

Chuck McManis discussed CAT-related activities ongoing at Sun Microsystems. Sun currently supports Kerberos V4, and plans to migrate to V5. Kerberos is invoked (using its native interface, rather than GSS-API) from RPC. Separate work on layering RPC atop GSS-API had been ongoing at Sun, but has not yet yielded conclusive results. One of the US National Laboratories had ported beta 2 of Kerberos V5 to Solaris, and Sun is working with the resultant code base.

CAT Technical Discussion

Two proposals for incremental changes to the GSS-API documents were considered:

1. A terminology change in response to a request from X/Open, renaming the per-message protection primitives from GSS_Sign to GSS_GetMIC, GSS_Seal to GSS_Wrap, GSS_Verify to GSS_VerifyMIC, and GSS_Unseal to GSS_Unwrap (to avoid conflict with other usage, without change to function, preserving (though deprecating use of) existing names in existing code for backward compatibility) was tentatively accepted pending e-mail review.

In evaluating the request and considering alternatives, it was observed that ISO’s usage of the term “seal” echoes the notion of applying a wax seal to a document. It was also observed that the current Kerberos V5 implementation of GSS_Sign emits a token containing the entirety of the input message rather than just a signature.

It was also observed that no GSS-API per-message protection interface currently exists to provide confidentiality without integrity, and post-meeting review (GSS-API specification, Section 1.2.2) confirmed the related point that mechanisms indicating the availability of per-message confidentiality services are also expected to indicate and offer per-message integrity. No definitive conclusion was reached about the level of demand for confidentiality without integrity.

2. A proposal to add `GSS_set_default_cred` and `GSS_lookup_default_cred` routines was rejected for reasons of semantics which were considered to be environment-specific (though considered as a likely initial entry in a set of extensions for POSIX and like environments). Much of the motivation for this feature derives from a desire to control the set of credentials which will be transferred by inheritance across the UNIX fork operation. It was observed that it would be difficult to implement the `set_default_cred` function within the current Kerberos V5 code base, and that different implementors could implement the proposal as defined with conflicting semantics which would not support application portability. Given an observation that credentials structures are ephemeral, use of `acquire_cred` with (non-ephemeral) principal names as arguments was recommended as an alternative approach which would survive UNIX forks.

Chuck McManis expressed interest in using `set_default_cred` as a means to spawn threads using different mechanisms for different threads, and saw this as a more critical priority than use of different identities within a single mechanism; he also expressed a desire that credentials be “lightweight” structures.

CAT Follow-On Tasks and Action Items

Follow-on tasks identified were:

- Kerberos V5 GSS-API mechanism specification and code enhancement;
- Kerberos V4 GSS-API implementation;
- “negotiated” mechanism definition (a task to which a framework discussion authored by Bob Blakely and forwarded to the list was considered relevant);
- CATS stream-oriented overlay definition;
- documentation of mechanism implementor’s guidance/agreements; and
- environment-specific specifications and extensions (e.g., credential inheritance semantics).

Individuals and subset groups were associated with several of these items. Activity on the “negotiated” mechanism’s design was argued as not being critical at this time; it will assume greater importance once multiple mechanisms are actively supported.

FTP Security

The discussion on FTP security was moderated by, and this section of the minutes was reported by, Sam Sjogren.

Review of FTP Activities

Discussions on the CAT mailing list as well as at the Columbus IETF meeting in March resulted in changes to the specification for security in FTP. Steve Lunt revised the “FTP Security Extensions” Internet-Draft and submitted it to the IETF Secretariat for announcement. A list of changes made to the document was also produced and was sent to the CAT mailing list. Since Steve was not able to physically attend the group’s meeting in Amsterdam, arrangements were made to allow Steve to participate via speakerphone. The list of changes was the focus of most of the group’s discussions.

FTP Technical Discussion

One of the additions to the FTP security document is a specification using the GSS-API authentication type. This specification needs to be reviewed in detail and any problems corrected. John Linn will communicate his observations to Steve on this.

Although the interaction of the AUTH, PASS, and USER commands have been clarified somewhat, it was agreed that the various possible cases (including those involving users for whom passwords (which should be protected) may be required in addition to other forms of authentication) should be more rigorous.

The form of Base-64 encoding used by FTP security has been brought into line with that used by PEM. One concern is that the length of a Base-64 string is currently unbounded, and that may cause problems for small-machine implementations. This will be addressed in the small-machine discussion on the mailing list.

For the time being, proxy file transfers are deferred. One of the effects of this is that the requirements for negotiating session keys are eased. However, the negotiating of session keys with the various possible mechanisms should still be investigated to make sure that in the future we will not be precluded from supporting this feature.

It is necessary for a server to indicate to a client, somehow, what levels of security are supported (e.g., integrity but not encryption). Although this has been left purely to the

particular mechanism, there is a feeling that the protocol itself should provide some support for determination of this when mechanisms themselves do not support it. So, a 402 reply code is defined which indicates to a client that ENC and/or MIC commands are not accepted, thereby allowing a client to probe a server to determine the levels of security it supports. Note that this even allows a server to force the use of privacy but not allow mere integrity assurance. This method is authentication-mechanism independent.

In the case of a server allowing integrity but not privacy, implementors are encouraged to warn the user that the level of security available is less than they have requested.

Another potential small-machine issue surfaced in the specification of buffer size and length for protected file transfers. Although the length field in the specification has been reduced from 4 bytes to 2 bytes, thereby reducing the buffer size from 4 gigabytes to 64 kilobytes, even a 64-kilobyte buffer may prove to be a problem for some small machines. This issue will be discussed further on the mailing list.

Instead of the commonly used 'rcmd' principal that is usually used with Kerberized TELNET and R-Utilities, the principal name 'ftp' has been specified for use with FTP security. There was a feeling that a number of sites may wish to avoid the additional overhead of creating another principal for each machine, so there should be some capability to fallback to use of the 'rcmd' principal. This would appear to be an issue left to particular implementations and site policy. Perhaps it should be mentioned in the Internet-Draft that clients are recommended to first try the 'ftp' principal and if the 'ftp' principal does not exist or the FTP server will not accept the 'ftp' principal, then try the 'rcmd' principal.

Various other small changes were made to the Internet-Draft that were either corrections or clarifications and are not worth mentioning in the minutes.

FTP Follow-On Tasks and Action Items

There will be discussion on the CAT Working Group mailing list regarding the buffer size issues and how a small-system implementation may be affected by large buffers or buffers of indefinite size.

Steve will incorporate the changes that arose from the group's discussions into the Internet-Draft and produce a new revision of the document and a list of changes.

The most important thing to do at this stage is to gain more implementation experience. Sam will solicit implementors through various e-mail lists and other channels.

Attendees

| | |
|--------------------|-----------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| David Arneson | arneson@ctron.com |
| Stefan Braun | smb@cs.tu-berlin.de |
| Richard Graveman | rfg@cctt.bellcore.com |
| Neil Haller | nmh@thumper.bellcore.com |
| Thomas Hutton | hutton@opus.sdsc.edu |
| John Johnston | john@berlioz.nsc.com |
| Nada Kapidzic | nadak@dsv.su.se |
| John Larson | jl Larson@parc.xerox.com |
| Arjen Lenstra | lenstra@bellcore.com |
| John Linn | linn@gza.com |
| Kim Chr. Madsen | kimcm@ic.dk |
| Chuck McManis | chuck.mcmanis@eng.sun.com |
| Clifford Neuman | bcn@isi.edu |
| John Penners | jpenners@advtech.uswest.com |
| Robert Reschly | reschly@brl.mil |
| Jeffrey Schiller | jis@mit.edu |
| Wolfgang Schneider | schneiw@darmstadt.gmd.de |
| Sam Sjogren | sjogren@tgv.com |
| Theodore Ts'o | tytso@mit.edu |
| Erik Zegwaart | zegwaart@surfnet.nl |
| James Zmuda | zmuda@mls.hac.com |

2.6.4 Internet Protocol Security Protocol (IPSEC)

Charter

Chair(s):

Al Hoover, hoover@ans.net

Paul Lambert, paul_lambert@email.mot.com

Mailing Lists:

General Discussion: ipsec@ans.net

To Subscribe: ipsec-request@ans.net

Archive: ftp.ans.net:~/pub/archive/ipsec

Description of Working Group:

Rapid advances in communication technology have accentuated the need for security in the Internet. The IP Security Protocol Working Group (IPSEC) will develop mechanisms to protect client protocols of IP. A security protocol in the network layer will be developed to provide cryptographic security services that will flexibly support combinations of authentication, integrity, access control, and confidentiality. The protocol formats for the IP Security Protocol (IPSP) will be independent of the cryptographic algorithm. The preliminary goals will specifically pursue host-to-host security followed by subnet-to-subnet and host-to-subnet topologies.

Protocol and cryptographic techniques will also be developed to support the key management requirements of the network layer security. The key management will be specified as an application layer protocol that is independent of the lower layer security protocol. The protocol will initially support public key based techniques. Flexibility in the protocol will allow eventual support of Key Distribution Center (KDC - such as Kerberos) and manual distribution approaches.

Goals and Milestones:

- Jun 1993 Post as an Internet-Draft the IP Security Protocol.
- Jul 1993 Post as an Internet-Draft the specification for Internet key management.
- Nov 1993 Report on pilot implementation of the IP Security Protocol. Update Protocol as needed.
- Mar 1994 Report on pilot implementation of the Internet Key Management Protocol. Update Internet-Draft as needed.
- Jul 1994 Submit the IP Security Protocol to the IESG for consideration as a Proposed Standard.

Jul 1994 Submit the Internet Key Management Protocol to the IESG for consideration as a Proposed Standard.

CURRENT MEETING REPORT**Minutes of the Internet Protocol Security Protocol Working Group (IPSEC)**

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

Attendees

| | |
|--------------------|------------------------------|
| Richard Colella | colella@nist.gov |
| Al Costanzo | al@akc.com |
| Stephen Crocker | crocker@tis.com |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Pierre Dupont | dupont@mdd.comm.mot.com |
| Dan Frommer | dan@jeremy.enet.dec.com |
| Richard Graveman | rfg@cct.bellcore.com |
| Neil Haller | nmh@thumper.bellcore.com |
| Alton Hoover | hoover@ans.net |
| Stephen Kent | kent@bbn.com |
| Arjen Lenstra | lenstra@bellcore.com |
| John Linn | linn@gza.com |
| Kim Chr. Madsen | kimcm@ic.dk |
| Cynthia Martin | martin@spica.disa.mil |
| David O'Leary | doleary@cisco.com |
| John Penners | jpenners@advtech.uswest.com |
| Charles Perkins | perk@watson.ibm.com |
| Philip Prindeville | philipp@res.enst.fr |
| Robert Shirey | shirey@mitre.org |
| William Simpson | Bill.Simpson@um.cc.umich.edu |
| Michael St. Johns | stjohns@darpa.mil |
| Vladimir Sukonnik | sukonnik@process.com |
| Robert Ullmann | ariel@world.std.com |
| Peter Yee | yee@atlas.arc.nasa.gov |
| James Zmuda | zmuda@mls.hac.com |

2.6.5 Network Access Server Requirements (NASREQ)

Charter

Chair(s):

Allan Rubens, acr@merit.edu

John Vollbrecht, jrv@merit.edu

Mailing Lists:

General Discussion: nas-req@merit.edu

To Subscribe: nas-req-request@merit.edu

Archive:

Description of Working Group:

The Network Access Server Requirements Working Group has as its primary goal, to identify functions and services that should be present in IP Network Access Servers (NASs) and to specify the standards that provide for these functions and services. The term "Network Access Server" is used instead of the more conventional term "Terminal Server" as it more accurately describes the functions of interest to this group. A "Network Access Server" is a device that provides for the attachment of both traditional "dumb terminals" and terminal emulators as well as workstations, PCs or routers utilizing a serial line framing protocol such as PPP or SLIP. A NAS is viewed as a device that sits on the boundary of an IP network, providing serial line points of attachment to the network. A NAS is not necessarily a separate physical entity; for example, a host system supporting serial line attachments is viewed as providing NAS functionality and should abide by NAS requirements.

This group will adopt (or define, if need be) a set of standard protocols to meet the needs of organizations providing network access. The immediate needs to be addressed by the group are in the areas of authentication, authorization, and accounting (AAA). In general, this group will select a set of existing standards as requirements for a NAS. If necessary, the group will identify areas of need where Internet standards don't already exist and new standardization efforts may be required.

Initially the group will independently investigate the two cases of character and frame-oriented access to the NAS. This investigation will be aimed at determining what work is being done, or needs to be done, in this and other working groups in order to be able to define the set of NAS requirements. While the ultimate goal of this group is to produce a NAS Requirements document, it may be necessary to define standards as well. This initial investigation will help determine what the goals of this group need to be. The group will also work with appropriate working groups to define required NAS standards that fall into the areas of these other groups.

Goals and Milestones:

- Done NAS Requirements Document posted as an Internet-Draft.
- Nov 1992 Post an Internet-Draft on Character oriented Authentication, Authorization, and Accounting(AAA).
- Nov 1992 Post an Internet-Draft on frame oriented AAA requirements.
- Nov 1993 Submit the NAS Requirements document to the IESG as a Proposed Standard.

Internet-Drafts:

“Network Access Server Proposed Requirements Document”, 10/01/1992, J. Vollbrecht, A. Rubens, G. McGregor <draft-ietf-nasreq-nasrequirements-01.txt>

CURRENT MEETING REPORT

Reported by Jim Barnes/Xylogics

Minutes of the Network Access Server Requirements Working Group
(NASREQ)

Agenda

- Introduction and brief review
- Discussion of revised charter
- Review of goals and schedule
- NAS “helper” subgroup report
- Discussion of latest nasreq draft
- Request for volunteer writers

The NASREQ Working Group met in Amsterdam with a small number of attendees. Lacking a critical mass of participants, the discussion was limited and the meeting brief.

The latest revision of the nasreq charter was briefly discussed. One question was whether the NAS is only being viewed as the entry point into a network or whether the requirements for a NAS should also reflect that the NAS could be a gateway out of a network to the outside world. Along this same theme, there was a later question about whether there were any NAS dialout requirements (for example to provide a PPP dialout link on demand).

During the discussion of accounting, two relevant points were raised:

1. The accounting group is in the process of being restarted. Setting up a subgroup with representatives from both NASREQ and the accounting group to exchange information was viewed as being a good idea.
2. It was noted that there is now a Modem Management Working Group (MODEM-MGT). One of the areas they are investigating is what accounting statistics would be useful for modems. Perhaps there is some common area that is interesting to both the NASREQ Working Group and the MODEMMGT Working Group.

In the authentication area, it was noted that the GSSAPI and Kerberos V5 documents have been recommended for Proposed Standard.

The authentication “helper” subgroup had little news to report. A draft document is still targeted for the November IETF meeting.

Cliff Neuman relayed a request from the Authorization and Access Control Working Group (AAC) for the NASREQ group to provide AAC with information on the NASREQ authorization requirements.

A request for volunteers to help write some sections of the NASREQ document was made. One person expressed interest in helping to do some of the writing.

Attendees

| | |
|-------------------------|------------------------------------|
| Mohammad Alaghebandan | mohammad_alaghebandan@3com.com |
| Jim Barnes | barnes@xylogics.com |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Peter Cameron | cameron@xylint.co.uk |
| Richard Graveman | rfg@ctt.bellcore.com |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Stephen Kent | kent@bbn.com |
| Andrew Knutsen | andrewk@sco.com |
| Klaus-Peter Kossakowski | Kossakow@informatik.uni-hamburg.de |
| John Linn | linn@gza.com |
| Kim Chr. Madsen | kimcm@ic.dk |
| Clifford Neuman | bcn@isi.edu |
| Timon Sloane | timon@timon.com |

2.6.6 Privacy-Enhanced Electronic Mail (PEM)

Charter

Chair(s):

Stephen Kent, kent@bbn.com

Mailing Lists:

General Discussion: pem-dev@tis.com

To Subscribe: pem-dev-request@tis.com

Archive: pem-dev-request@tis.com

Description of Working Group:

PEM is the outgrowth of work by the Privacy and Security Research Group (PSRG) of the IRTF. At the heart of PEM is a set of procedures for transforming RFC 822 messages in such a fashion as to provide integrity, data origin authenticity, and, optionally, confidentiality. PEM may be employed with either symmetric or asymmetric cryptographic key distribution mechanisms. Because the asymmetric (public-key) mechanisms are better suited to the large scale, heterogeneously administered environment characteristic of the Internet, to date only those mechanisms have been standardized. The standard form adopted by PEM is largely a profile of the CCITT X.509 (Directory Authentication Framework) recommendation.

PEM is defined by a series of documents. The first in the series defines the message processing procedures. The second defines the public-key certification system adopted for use with PEM. The third provides definitions and identifiers for various algorithms used by PEM. The fourth defines message formats and conventions for user registration, Certificate Revocation List (CRL) distribution, etc. (The first three of these were previously issued as RFCs 1113, 1114 and 1115. All documents have been revised and are being issued first as Internet-Drafts.)

Goals and Milestones:

- | | |
|---------|---|
| Done | Submit first, third, and fourth documents as Internet-Drafts. |
| Ongoing | Revise Proposed Standards and submit to IESG for consideration as a Draft Standard, and repeat for consideration as an Internet Standard. |
| Done | Submit second document as an Internet-Draft. |
| Done | First IETF working group meeting to review Internet-Drafts. |
| Done | Submit revised Internet-Drafts based on comments received during working group meeting, from pem-dev mailing list, etc. |

- Done Submit Internet-Drafts to IESG for consideration as Proposed Standards.
- Done Post an Internet-Draft of the MIME/PEM Interaction specification.
- Apr 1993 Submit the PEM/MIME specification to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“MIME-PEM Interaction”, 11/23/1992, S. Crocker, N. Freed, M. Rose <draft-ietf-pem-mime-02.txt>

Request For Comments:

- RFC 1319 “The MD2 Message-Digest Algorithm”
- RFC 1320 “The MD4 Message-Digest Algorithm”
- RFC 1321 “The MD5 Message-Digest Algorithm”
- RFC 1421 “Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures”
- RFC 1422 “Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management”
- RFC 1423 “Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers”
- RFC 1424 “Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services”

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

Minutes of the Privacy Enhanced Mail Working Group (PEM)

PEM Implementation Status Report

- TIS

TIS/PEM is now available via anonymous FTP and the distribution includes CA software. The TIS UK office is working on a commercial PEM implementation to be made available in Europe.

- MIT

TechMail with PEM, a Mac PEM implementation integrated with a Mac e-mail system, will be available very soon via anonymous FTP. TechMail requires a POP3 server and has been shown to interoperate with TIS/PEM. Ray Lau, a well-known Mac software developer, has developed a stand-alone version of PEM software for the Mac which may be available soon.

- PASSWORD Project

Three Unix PEM implementations are now available, and a fourth is on the way as a result of this EC-sponsored project. The implementations are from UCL, Cambridge University, INRIA, and GMD. All can make use of X.500 directories and all include CA software. All the implementations interoperate with one another and there have been some interworking tests with TIS/PEM.

The UCL version is "somewhat" integrated with MH and runs on a SPARC now.

The GMD version is part of a larger security tool kit, using a variety of algorithms, and secure (strong authentication) X.500 operations. The GMD implementation is a standalone version that is not integrated into an e-mail system, although it is integrated with X.500 directory access code to fetch certificates and CRLs. It will be available via anonymous FTP (subject to COCOM restrictions). There are plans to establish a PCA for academic users within Germany.

The Cambridge version supports symmetric key management as well as asymmetric key management (using RSA). The Cambridge version is the subject of experimentation with other algorithm suites, e.g., DSA and SHA. It also works without access to a directory server.

- COST (Sweden)

This commercial PEM implementation runs on DEC and Sun workstations, includes CA software, and includes a centralized certificate server for all COST/PEM users. COST intends to run its own PCA. COST/PEM is now undergoing beta testing at universities in Dublin and Stockholm. They are performing interoperability tests with TIS. A smart card version is being developed.

- RSADSI

A new TIPEM (commercial product) version is in beta now and should be available in eight to twelve weeks. The Certificate Issuing Systems (a commercial product) software and hardware will be available in six to eight weeks. A description of the RSADSI plans for operating a low-assurance PCA with two CAs will be available very soon. One CA (a free service) under this PCA will be a PERSONA CA and it is now operating on a trial basis; a second (non-PERSONA) CA will become operational soon along with a policy statement available via FTP. A commercial PCA, using the CIS noted above, will be available soon, and the policy for that PCA also will be available soon.

MIME-PEM Discussion

It was the intent of this PEM Working Group meeting to review the latest technical proposal (Internet-Draft) for MIME-PEM integration, as announced in the agenda distributed several weeks prior to the meeting. John Linn and Jeff Schiller were asked to review this Internet-Draft and prepare comments for this meeting. Unfortunately, the Internet-Draft available immediately prior to the meeting did not reflect changes discussed at the last working group meeting and in subsequent pem-dev mailing list discussion. Thus no substantive progress was made on this topic.

Both John and Jeff expressed concern about the continued presence of inline “optional” fields to indicate the PEM processing to be applied, or to indicate the PEM processing that has been applied to a message. They also expressed concern about “distinguished encoding” problems that may arise in complex MIME messages where a signature might encompass MIME headers embedded in these complex messages. Both observed that one possible means of simplifying the MIME-PEM “integration problem” is to treat an RFC 1421 “message” as a body part to be carried by MIME. This might not offer all of the flexibility of the recent proposals, but it would significantly simplify both processing and backward compatibility for RFC 822-PEM implementations.

Steve Crocker presented several technical points related to the structure of MIME-PEM messages, based on a very recent meeting with Marshall Rose. There was agreement that signatures on complex MIME objects cannot be effected as previously proposed, and that new application type MIME objects were required to “protect” signed data from possible manipulation by MIME relays.

He agreed that representing PEM messages as application context types would allow RFC 1421 data to be a body part, as suggested above. An analysis by Steve Crocker and Marshall suggested that a (MIME) message reader can easily distinguish among a vanilla text message, a MIME message, or a PEM-MIME message. Disambiguation would require introduction of an additional blank line into the RFC 1421 message format. However, if an RFC 1421 message were to pass through an intermediate MIME relay, it might be transformed in a way that would make it ambiguous as to whether this was initially an RFC 1421 message or initially a MIME-PEM message.

There was no resolution of these issues. There was agreement that a new MIME-PEM Internet-Draft must be written to reflect the recent improved understanding of these problems, and to reflect the comments of the previous PEM Working Group meetings. A proposal for a simple, MIME encapsulation of RFC 1421 messages may be developed, but no authors were identified.

Certificate Name Discussion

At the previous PEM Working Group meeting, Steve Crocker initiated a discussion of the form of names that should appear in certificates used in PEM. He advocated the use of domain name system (DNS) mailbox names in lieu of distinguished names (DNs). This discussion took the form of several hand-written slides that were reported upon in the meeting minutes, but there was no written follow-up on the pem-dev mailing list. In response to this previous presentation, Steve Kent assembled material to compare the use of DNS names and DN's. Due to time constraints, only some of this material was presented during the meeting. This material argues that many (though not all) of the objections previously raised to the use of DN's can be addressed through good PEM implementation practice and that the use of DN's offers a number of advantages relative to the use of DNS.

Editor's Note: The material is available via FTP or mail server from the remote directories as /ietf/pem/pem-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Certification System Discussion

At the previous PEM Working Group meeting, Steve Crocker also initiated a discussion of the concept of relaxing some of the constraints imposed by the PEM certificate management system (RFC 1422). This discussion took the form of several hand-written slides that were reported upon in the meeting minutes, but there was no written follow-up on the pem-dev mailing list. In response to this previous presentation, Steve Kent assembled material to review some of the rationale that underlies the current certification system. Due to time constraints, none of this material was presented.

Editor's Note: The material is available via FTP or mail server from the remote directories as /ietf/pem/pem-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Algorithm Discussion

There was extensive discussion of ways to use "triple-DES" on the pem-dev mailing list prior to this meeting. Mike Roe made a brief presentation on triple-DES options and distributed to attendees an extensive analysis he had prepared. There is agreement that use of EDE as a codebook is intrinsically slower than approaches that perform multiple chaining passes, if parallel hardware is employed. However, in the PEM (e-mail) context it is not clear that this performance improvement is a significant factor, especially if software DES (not executed in parallel on multiple processors) is the dominant implementation mode. In this context, the various triple-DES options are essentially equivalent in performance.

It was suggested that the primary motivation for using triple-DES is security, not performance. Although there is significant literature supporting the security of EDE as a codebook, there is little if any analysis of the security of the other proposed modes. It was announced that RSA Labs will perform a study of the security of the different proposed modes and make a report available to the PEM Working Group.

No resolution of this issue was reached at this meeting.

Attendees

| | |
|-----------------------|------------------------------|
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Frederik Andersen | fha@dde.dk |
| James Conklin | jbc@bitnic.educom.edu |
| Stephen Crocker | crocker@tis.com |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Kishan Dudkikar | kishan@icm1.icp.net |
| Steve Dusse | spock@rsa.com |
| Barbara Fraser | byf@cert.org |
| James Galvin | galvin@tis.com |
| Tim Goodwin | tim@pipex.net |
| Richard Graveman | rfg@cct.bellcore.com |
| Neil Haller | nmh@thumper.bellcore.com |
| Alton Hoover | hoover@ans.net |
| Nada Kapidzic | nadak@dsv.su.se |
| Stephen Kent | kent@bbn.com |
| Peter Kirstein | P.Kirstein@cs.ucl.ac.uk |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| Arjen Lenstra | lenstra@bellcore.com |

| | |
|--------------------|--------------------------|
| John Linn | linn@gza.com |
| Kim Chr. Madsen | kimcm@ic.dk |
| Keith Moore | moore@cs.utk.edu |
| Michael Roe | mrr@cl.cam.ac.uk |
| Jeffrey Schiller | jis@mit.edu |
| Wolfgang Schneider | schneiw@darmstadt.gmd.de |
| Robert Shirey | shirey@mitre.org |
| Theodore Ts'o | tytso@mit.edu |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Erik Zegwaart | zegwaart@surfnet.nl |
| James Zmuda | zmuda@mls.hac.com |

2.7 Service Applications Area

Director(s):

- David Crocker: dcrocker@mordor.stanford.edu

Area Summary reported by David Crocker/Silicon Graphics

The Service Applications Area was formed at the Columbus IETF meeting, splitting off from Transport Services, to provide focus on the “middleware” range of end-system support services, including file access, time synchronization, directory lookup and canonical access and representation procedures.

NFS and ONC IETF Standards Effort BOF (ONC)

ONC is a suite of protocols developed by Sun Microsystems. These protocols include:

- XDR for canonical data representation
- RPC for remote procedure call
- NFS and NFS+ for file access
- LOCKD for resource access coordination
- NIS and NIS+ for resource location

Sun is offering these protocols to the IETF for standardization, and the Amsterdam IETF meeting included a technical presentation followed by a discussion BOF. The presentation covered the nature and state of the various protocols. The BOF discussed the protocols in greater detail and further discussed the IETF’s interest in pursuing their standardization. There was clearly sufficient interest to warrant pursuing the matter further, including formation of a working group to consider immediate standardization of some of the protocols, and to perform whatever modifications are necessary to then standardize the others.

Standardization will require that Sun formally assign “change control” ownership to the IETF. Development of the necessary paperwork will be pursued with ISOC and its counsel.

Domain Name System (DNS)

The DNS Working Group covers a wide range of development and maintenance activities for the Domain Name System. Rather than dividing into multiple working groups, it is currently operating with a series of sub-groups. The load balancing subgroup is interested in using the DNS to spread users across multiple machines/interfaces. The security subgroup is concerned with authentication and integrity of DNS data. The big zones subgroups is

attending to the question of very large “flat” portions of the DNS, with “.com” providing the major impetus.

The load balancing subgroup is basically done, having to write the informational paper and let people comment on it; there are no proposed protocol changes. The security subgroup has not done much until now, but will soon start doing the cryptographic signature work that has been discussed. The big zones subgroup has done some exploration but has not reached any kind of real closure. The subgroup will keep trying, with a few more people promising to help on this.

The RFC Editor has asked the DNS Working Group to review a paper on “Service Advertisement using the DNS.” Marshall Rose asked for advice about some technical points of using DNS wild cards. There was an updated summary of the PIP (IPng) DNS design work by Sue Thomson; this sparked a resurrection of the old debate about the usefulness of the DNS class mechanism, which debate was stopped by the chair when it started looping. There was a discussion on some timestamp-related mechanisms that have been proposed both as part of the PIP work and as part of an incremental zone transfer protocol proposed by Anant Kumar. The general feeling was that the DNS Working Group should look into this but they do not yet understand exactly what they want. The working group agreed to take on the draft “Common DNS Errors and Suggested Fixes” submitted by Jon Postel, et al. The chair announced the existence of several new DNS-related Internet-Drafts, and asked other members of the working group to please review them.

MHS-DS (MHSDS)

The MHSDS Working Group is seeking to integrate use of the X.500 directory service into Internet X.400 operation, including e-mail routing.

The MHSDS Working Group decided to publish three Internet-Drafts as Proposed Standards. An additional Internet-Draft will be published as Experimental. Minor editorial changes will be made to these documents, a final call for comments will be made to the working group, and then the documents will be progressed. In addition to document review, the working group reviewed progress of its pilot project, Project Long Bud. The Internet-Draft which describes this project will be updated to reflect comments made at this IETF meeting, the project’s FTP archive will be reorganized and updated, and actions were assigned to begin investigation of ways to improve the quality of Internet X.500 service related to support of X.400 routing and address mapping. Finally, the working group held a tutorial session to help some of its membership better understand the technical details of its X.400 routing and mapping algorithms.

Service Location Protocol (SVRLOC)

Sun's NIS+ proposal was discussed, to understand what part of the solution space it covers and whether the current service location proposal will cover the needs of NIS+ clients. The working group went through what is thought to be the list of items that are remaining to complete for the proposal to begin its travels down the IETF standards track. A document ready to submit to the standards track should be completed by the next IETF meeting.

Minimal OSI Upper-Layers (THINOSI)

The THINOSI Working Group is specifying a subset of the OSI upper-layer infrastructure protocols, to facilitate implementation and operational efficiency.

Various issues were discussed in the review of the cookbook. A point that came up more than once was how the cookbook should relate to the parallel work in the OSI regional workshops (the Common Upper-Layer Requirements Part 3: Minimal OSI Profile (CULR3)) and in X/Open (specification of use of the XTI interface for minimal OSI (XTI/mOSI)). The possibility of the cookbook having a formal statement of compliance to CULR-3 was discussed. The eventual status of the cookbook was discussed, and it was believed it should be targetted for the standards track, as the specification of the supporting protocol layers for the relevant applications. Since the charter was written (following the BOF held in Washington, DC), the coverage of the cookbook has changed to more than just "byte-stream" (although the amount of new text is small).

Trusted Network File Systems (TNFS)

TNFS did not meet in Amsterdam. The working group has submitted a draft specification to the standards process. Final details are being resolved, prior to formal IESG review.

CURRENT MEETING REPORT

Reported by Dave Crocker/Silicon Graphics

Minutes of the NFS and ONC IETF Standards Effort BOF (ONC)

Creation of the Service Applications Area suggests increased interest in the “middleware” category of support services, above transport and below specific application semantics. Sun Microsystems’ Open Network Computing (ONC) suite of protocols provides a number of services in this arena and has an established user base. Sun has expressed an interest in pursuing IETF standardization of the ONC suite.

At the Amsterdam IETF, two events were held to consider this possibility. The first was a Monday morning plenary presentation and the second was a BOF. The plenary presentation was to announce the discussions with Sun and to provide a basic introduction to the technology. The BOF continued the technical exposition, with detailed review of the ONC components, and then discussed the IETF’s interest in standardizing the suite. The protocols that would be candidates include:

- XDR, data representation
- RPC, remote procedure call
- NFS Version 2 & Version 3, file access
- LOCKD, resource access coordination
- NIS & NIS+, resource location

XDR, RPC and NFS V2 have already been published as Informational RFCs. Further, they have been quite stable for a long time. Version 3 is quite new and has been implemented only on a few platforms; it is just being introduced to the NFS development community. NIS+ also is quite new and not fully documented.

The BOF attendees seemed quite interested in ALL of the modules, though there was no clear agreement to standardize all of them. That is, there is interest in proceeding, taking things one step at a time. It was agreed to form a working group which would serially process the protocols, in whatever manner the working group felt best. It was expected that that would mean direct standardization of some of the protocols, without modification, and possibly changing some of the others prior to considering their standardization.

Quite understandably, the BOF discussion also probed Sun’s willingness to turn over control of the protocols to the IETF. (This is a formal requirement, documented in RFC 1310.) Sun has repeatedly offered its assurances that a) it understands the requirement, b) it is comfortable with it, and c) it wishes to work with the IETF on further growth of the suite.

Attendees

| | |
|--------------------|-----------------------------------|
| Steve Alexander | stevea@lachman.com |
| James Allard | jallard@microsoft.com |
| Harald Alvestrand | Harald.Alvestrand@delab.sintef.no |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| David Borman | dab@cray.com |
| Piet Bovenga | p.bovenga@uci.kun.nl |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Brian Carpenter | brian@dxcern.cern.ch |
| David Conrad | davidc@iij.ad.jp |
| David Crocker | dcrocker@mordor.stanford.edu |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Dennis Ferguson | dennis@ans.net |
| Osten Franberg | euaokf@eua.ericsson.se |
| Robert Gilligan | Bob.Gilligan@Eng.Sun.Com |
| Terry Gray | gray@cac.washington.edu |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Gerd Holzhauser | Holzhauser1@applelink.apple.com |
| Phil Irely | pirey@relay.nswc.navy.mil |
| Scott Kaplan | scott@wco.ftp.com |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Sylvain Langlois | Sylvain.Langlois@exp.edf.fr |
| John Larson | jl Larson@parc.xerox.com |
| Chuck McManis | chuck.mcmanis@eng.sun.com |
| Keith Moore | moore@cs.utk.edu |
| Jun Murai | jun@wide.ad.jp |
| Clifford Neuman | bcn@isi.edu |
| Michael O'Dell | mo@uunet.uu.net |
| Drew Perkins | ddp@fore.com |
| Jim Rees | Jim.Rees@umich.edu |
| Robert Reschly | reschly@brl.mil |
| Daisy Rose | daisy@watson.ibm.com |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Henry Sanders | henrysa@microsoft.com |
| Jon Saperia | saperia@tay.dec.com |
| John Stewart | |
| Antoine Trannoy | trannoy@crs4.it |
| John Veizades | veizades@wco.ftp.com |
| Noriko Yokokawa | norinori@wide.ad.jp |

IETF Standardization of ONC/NFS?

D. Crocker
Silicon Graphics, Inc.
dcrocker@sgi.com

The dream home

- **WHAT FEATURES & FUNCTIONS ARE NEEDED?**
 - Support for distributed applications
 - Support for networked hosts
 - + ??

IETF Standardization of ONC/NFS

The properties: A set of technologies

- XDR
- RPC
- PORTMAPPER
- NFS (2, 3)
- NIS (+)

IETF Standardization of ONC/NFS

Willing buyer? What does the IETF want

- **NOTABLE HOLES IN IETF APPLICATIONS INFRASTRUCTURE**
 - Poor "subroutine" facilities
 - No application-level request-response protocol
 - No standard file access
- **MAKE/BUY ALTERNATIVES?**

IETF Standardization of ONC/NFS

Willing seller? What does Sun want?

- **SUN MICROSYSTEMS CURRENTLY HOLDS REAL CHANGE CONTROL**
- **IETF REQUIRES THAT IT BE FORMALLY HANDED CHANGE CONTROL**
- **WHAT IS SUN WILLING TO DISCUSS GIVING TO IETF?**

IETF Standardization of ONC/NFS

Letter of intent?

- **COMPONENTS OF A CANDIDATE AGREEMENT**
- **ROUGH CONSENSUS?**
- **HOW SHALL WE PROCEED?**

IETF Standardization of ONC/NFS

2.7.1 Domain Name System (DNS)

Charter

Chair(s):

Rob Austein, sra@epilogue.com

Mailing Lists:

General Discussion: namedroppers@nic.ddn.mil

To Subscribe: namedroppers-request@nic.ddn.mil

Archive: nicfs.nic.ddn.mil:~/namedroppers/*.Z

Description of Working Group:

The DNS Working Group is concerned with the design, operation, and evolution of the Domain Name System within the Internet. As the Internet continues to grow, we expect to serve as a focal point for work on scaling problems within the current framework, work on protocol evolution as new mechanisms become necessary, and documentation of current practice for DNS implementors and administrators. We are also responsible for oversight of DNS activities by other groups within the IETF to the extent that we review the impact such work will have on the DNS and make recommendations to the working groups and IESG as necessary. Since some of these are ongoing tasks, we do not expect the working group to disband anytime soon.

Several issues are of particular concern at this time:

Scaling: The DNS is the victim of its own success. The global DNS namespace has grown to the point where administering the top levels of the tree is nearly as much work as the old NIC host table used to be. We need to work on ways to distribute the load. Some of the solutions are likely to be technical, some political or economic; we still treat the top-level DNS service the way we did when DARPA was footing the bill, and the funding for that service is in the process of going away.

Security: The DNS is a zero-security system; it is not even as strong as the IP layer above which it operates. As a result, accidental spoofing (cache pollution) is an all-too-frequent occurrence. We need to make the DNS more robust against accidental corruption, and must provide at least an optional authentication mechanism for that portion of the community that wants one. At the same time, we must not cripple the existing system by drastically increasing its bandwidth consumption or by mandating use of cryptographic techniques that would preclude worldwide distribution of DNS software. The global DNS database is exactly that, an existing world-wide database representing hosts on six continents and (at least) forty-five countries. A solution that does not take this into account is not acceptable.

Management: The group has a draft document describing MIB extensions to manage the DNS. It also needs to specify a standard way to dynamically create and destroy DNS records; SNMP may be an appropriate tool for this task, but we haven't yet specified enough of the details to know for certain. The impact that a dynamic update mechanism will have on the DNS needs to be examined, with particular attention given to security and scaling issues.

IPng/Routing: As the fur starts flying in the battle between the IPng proponents and the new-routing-architecture proponents, it is expected that groups on both sides will need some amount of support from the DNS. Such support is likely to be minimal and straightforward, but these proposals are likely to need "rush service" for whatever support they require. So the working group needs to monitor these activities, stay involved, and generally do what it can to make sure that DNS support is not a bottleneck.

The DNS Working Group also needs to examine the impact that any proposed IPng system would have on the DNS, since the DNS database and protocols have special provision for IP addresses.

Goals and Milestones:

- | | |
|----------|---|
| Done | Post as an Internet-Draft a description of the Responsible Person Record. |
| Done | Post an Internet-Draft specifying the addition of network naming capability to the DNS. |
| TBD | Submit to the IESG the document for load balancing in the DNS as an Informational document. |
| Done | Submit the Responsible Person Record to the IESG for consideration as a Proposed Standard. |
| Ongoing | Monitor and offer technical support to the various groups working on the next version of IP. |
| TBD | Post an Internet-Draft of the "Big Zone" policy recommendations for root and first-level zone administration. |
| TBD | Submit the "Big Zone" policy document to the IESG for consideration as a policy statement. |
| Done | Submit the specification for network naming to the IESG for consideration as a Proposed Standard. |
| Done | Post the DNS MIB as an Internet-Draft. |
| Feb 1993 | Submit the DNS MIB to the IESG for consideration as a Proposed Standard. |
| Mar 1993 | Post an Internet-Draft specifying the dynamic resource record creation and deletion. |

- Mar 1993 Submit to the IESG the incremental zone transfer mechanism as a Proposed Standard.
- Mar 1993 List and prioritize the Working Group's goals, and pick a subset that is appropriate to pursue at the present time.
- Jun 1993 Post an Internet-Draft for adding load balancing capability to the DNS.
- Nov 1993 Submit the proposal for dynamic resource record creation/deletion to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

- "DNS Support for IDPR", 03/22/1993, R. Austein <draft-ietf-dns-idpr-01.txt>
- "DNS Server MIB Extensions", 06/11/1993, R. Austein, J. Saperia <draft-ietf-dns-server-mib-01.txt>
- "DNS Resolver MIB Extensions", 06/11/1993, R. Austein, J. Saperia <draft-ietf-dns-resolver-mib-01.txt>
- "Common DNS errors and suggested fixes.", 07/06/1993, A. Kumar, J. Postel, C. Neuman <draft-ietf-dns-common-errors-01.txt>
- "Common DNS Data File Configuration Errors", 08/11/1993, P. Beertema <draft-ietf-dns-config-errors-00.txt>

Request For Comments:

RFC 1480 "The US Domain"

CURRENT MEETING REPORT

Reported by Rob Austein/Epilogue Technology

Minutes of the Domain Name System Working Group (DNS)

Thanks to Bill Manning for providing the notes on which these minutes are based.

The first part of the meeting consisted of status reports from the chair of the working group and the leaders of several subgroups that have undertaken specific tasks assigned at previous meetings.

The first report was from James Gavin, leader of the subgroup working on DNS security (please see the end of these minutes for subgroup mailing list information). Per recent discussions on the DNS Working Group mailing list, the security subgroup believes that an IP-level security mechanism does not provide the service security needed by the DNS, and that the right model for the DNS is a digital signature providing end-to-end authentication of RR data. The exact digital signature mechanism to be used is still under discussion. The subgroup expects to begin serious work in the near future (that is, before the 28th IETF in Houston).

The working group explicitly absolved James's subgroup from responsibility for the so-called "just as good as IP security" issues, some of which have already been addressed by code contributed to BIND version 9.1 by USC-ISI.

The DNS MIB has been split into two separate MIBs (one for resolvers, one for name servers), per advice from the Network Management Directorate (NMAREA). The latest revisions of the MIB documents (`draft-ietf-dns-resolver-mib-01.txt` and `draft-ietf-dns-server-mib-01.txt`) have been submitted to the IESG for approval as Proposed Standards. Calls for objections were issued both to the DNS Working Group mailing list and verbally at the working group meeting; the authors of the MIB documents feel that they have successfully defended the current documents against the one objection that was raised (to the authors' last-minute decision to remove the variable `dnsServCounterNonAuthNoNames` from the server MIB), and that the documents are (finally!) ready for promotion to Proposed Standard status. We expect a decision from the IESG in the near future, certainly before the 28th IETF.

Liaison work with the X.400 Operations Working Group (X400OPS) has been proceeding in fits and starts, but we believe that we are making progress. As of the X400OPS meeting on the morning of 14 July, we believe we have an understanding with X400OPS on how their DNS work should proceed, and we expect to receive a copy of the next draft of the X400OPS "mapping table" paper from Claudio Allocchio, our liaison within X400OPS, as soon as he has a chance to write it.

On 1 July, the RFC Editor asked the DNS Working Group to review a short document entitled "Service Advertisement Using the DNS." This document had been submitted directly

to the RFC Editor without starting life as an Internet-Draft. The DNS Working Group chair reviewed the document, solicited other reviewers from the working group and sent comments to the RFC Editor.

The report for the Load Balancing subgroup was given by Thomas Brisco. Based on commentary from the DNS Working Group Chair and the Service Applications Area Director, the load balancing subgroup believes that their problem would be best solved by implementation hacks, without attempting to extend the DNS protocol by adding new magic RR types. Accordingly, the subgroup will now write a document describing the kinds of implementation hacks that best address their problem, put said document up for review and publication as an Informational RFC, and terminate the subgroup after a suitable review period. The document will include text warning about known implementation problems (e.g., zero TTLs) and required sanity checking.

Next, the working group heard a short presentation by Marshall Rose, outlining some technical details of how Marshall's "experiment in remote printing" uses DNS MX RRs with wildcard owner names to map international telephone numbers to SMTP servers. In brief, an international phone number like +1-415-123-4567 would be mapped to the DNS name 7.6.5.4.3.2.1.5.1.4.1.TPC.INT, thus allowing all of the San Francisco area to be covered by a wildcard name such as *.5.1.4.1.TPC.INT. We concluded that Marshall's proposal was technically feasible, but warned him that his scheme could be construed as duplication of the global authority tree, and that he might encounter administrative or political problems similar to the ones encountered by X400OPS. See RFC 1486, "An Experiment in Remote Printing," for more details on this topic.

A brief discussion followed on adding timestamps to the DNS protocols. Several proposals currently under discussion (the P. Internet Protocol Working Group (PIP) DNS work and Anant Kumar's proposed incremental zone transfer protocol) involve use of a timestamp mechanism to detect out-of-date RRs. One way of retrofitting a timestamp mechanism into the DNS protocols would be to define a new DNS class; all the RR types in this class would have a timestamp as the first part of their RDATA portions. We would also need to allocate new RR type codes for timestamped versions of all the "class-invariant" RR types. This is ugly, but would retain backwards compatibility with existing DNS code that thinks it knows how to parse any RR. Several members of the working group suggested using a new DNS opcode instead of a new DNS class; this avoids all the delegation problems associated with a new class, but doesn't preserve strict backwards compatibility with the existing protocol. This is still a research topic.

During the timestamp discussion, Masataka Ohta pointed out that the timestamp-based incremental zone transfer protocol as circulated, does not provide any way to delete RRs, only to add them. Fixing this shouldn't be hard, it just requires some kind of deletion pseudo-type as in Paul Mockapetris's original proposal (the DNS2 BOF held at the 25th IETF).

Next, Sue Thomson presented the most recent DNS design work done by the PIP Working Group. The details of this work are described in the current Internet-Draft "draft-ietf-

pip-dns-01.txt.” Briefly, the document proposes to allocate a new DNS class for PIP; this solves several of the problems discussed at the Columbus (26th IETF) DNS Working Group meeting, but introduces all the known difficulties associated with use of multiple DNS classes. The document also suggests using a timestamp mechanism. This is still a snapshot of a work in progress.

Last, the working group agreed to take on responsibility for the Internet-Draft, “Common DNS Errors and Suggested Fixes” submitted to the working group by Jon Postel. There was not enough time to discuss the document itself. Please read the Internet-Draft and send comments to Anant Kumar, anant@isi.edu, or to the DNS Working Group mailing list. Anant will coordinate changes.

Subgroup Mailing Lists

DNS Security

- General Discussion: dns-security@tis.com
- To Subscribe: dns-security-request@tis.com

Load Balancing

- General Discussion: dns-wg-lb@ns1.rutgers.edu
- To Subscribe: dns-wg-lb-request@ns1.rutgers.edu

Attendees

| | |
|--------------------|--|
| Robert Austein | sra@epilogue.com |
| Anders Baardsgaad | anders@cc.uit.no |
| Tony Bates | tony@ripe.net |
| David Borman | dab@cray.com |
| Erik-Jan Bos | erik-jan.bos@surfnet.nl |
| Thomas Brisco | brisco@pilot.njin.net |
| Henry Clark | henryc@oar.net |
| Geert Jan de Groot | geertj@ica.philips.nl |
| Francis Dupont | francis.dupont@inria.fr |
| Osten Franberg | euaokf@eua.ericsson.se |
| John Hopkins | J_Hopkins@icrf.icnet.uk |
| Marc Horowitz | marc@mit.edu |
| Steven Horowitz | witz@chipcom.com |
| Phil Irey | pirey@relay.nswc.navy.mil |
| Thomas Johannsen | thomas@ebzaw1.et.tu-dresden.de |
| Dale Johnson | dsj@merit.edu |
| Marijke Kaat | marijke@sara.nl |

| | |
|-------------------------|---------------------------------|
| Frank Kastenholz | kasten@ftp.com |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Mark Kusters | markk@internic.net |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Eliot Lear | lear@sgi.com |
| Jose Legatheaux Martins | jalm@fct.unl.pt |
| Carl Malamud | carl@malamud.com |
| Bill Manning | bmanning@rice.edu |
| Greg Minshall | minshall@wc.novell.com |
| Keith Mitchell | keith@pipex.net |
| Clifford Neuman | bcn@isi.edu |
| Peder Chr. Noergaard | pcn@tbit.dk |
| Masataka Ohta | mohta@cc.titech.ac.jp |
| Petri Ojala | ojala@eunet.fi |
| Michael Patton | map@bbn.com |
| Charles Perkins | perk@watson.ibm.com |
| Lars Poulsen | lars@cmc.com |
| Juergen Rauschenbach | jrau@dfn.de |
| Robert Reschly | reschly@brl.mil |
| John Romkey | romkey@elf.com |
| Luc Rooijackers | lwj@cs.kun.nl |
| Marshall Rose | mrose@dbc.mtview.ca.us |
| Miguel Sanz | miguel.sanz@rediris.es |
| Jon Saperia | saperia@tay.dec.com |
| Tim Seaver | tas@concert.net |
| John Stewart | jstewart@cnri.reston.va.us |
| Erdal Taner | erdal@vm.cc.metu.edu.tr |
| Marten Terpstra | marten@ripe.net |
| Susan Thomson | set@bellcore.com |
| Gregory Vaudreuil | gvaudre@cnri.reston.va.us |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Jost Weinmiller | jost@prz.tu-berlin.d400.de |
| Sam Wilson | sam.wilson@ed.ac.uk |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Romeo Zwart | romeo@sara.nl |

2.7.2 MHS-DS (MHSDS)

Charter

Chair(s):

Kevin Jordan, Kevin.E.Jordan@cdc.com
Harald Alvestrand, Harald.Alvestrand@uninett.no

Mailing Lists:

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

Description of Working Group:

The MHS-DS Working Group works on issues relating to Message Handling Services use of Directory Services. The Message Handling Services are primarily X.400, but issues relating to RFC 822 use of Directory and Directory support for RFC 822 and X.400 interworking are in the scope of the group. Directory and Directory Services refer to the services based upon the CCITT X.500 recommendations and additional ISO standards, stable implementation agreements, and RFCs, as specified by the OSI-DS Working Group. The major aims of the MHS-DS Working Group are:

1. Define a set of specifications to enable effective, large-scale deployment of X.400.
2. Study issues associated with supporting X.400 communities which lack access to X.500 Directory, and define requirements for tools which: a) extract information from the X.500 Directory for use by non-X.500 applications, b) upload information into the X.500 Directory.
3. Coordinate a pilot project which deploys MHS information into the X.500 Directory and uses it to facilitate mail routing and address mapping. The results of this pilot will be documented, and experience gained from the project will be fed back into the Internet specifications created by the working group.

Goals and Milestones:

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

- Jan 1993 Post an overview of MHS use of Directory as an Internet-Draft.
- Done Post a document on representing tables and subtrees in the directory as an Internet-Draft.
- Done Post an Internet-Draft on representing the O/R Address hierarchy in the Directory Information Tree.
- Done Post an Internet-Draft on MHS use of Directory to support MHS Routing.
- Done Post as an Internet-Draft a document on the use of the directory to support mapping between X.400 and RFC822 addresses.
- Done Post as an Internet-Draft a document describing a simple profile for MHS use of Directory.
- Done Post as an Internet-Draft a document on the use of the Directory to support routing for RFC822 and related protocols.
- Done Submit as an Internet-Draft a document on MHS use of Directory to support MHS Context Conversion.
- Done Post as an Internet-Draft a document describing the use of the Directory to support distribution lists.
- Aug 1993 Submit the set of MHS-DS documents to the IESG for consideration as Experimental and Informational documents.

Internet-Drafts:

- “A simple profile for MHS use of Directory”, 04/09/1992, S. Kille <draft-ietf-mhsds-mhsprofile-03.txt, .ps>
- “Representing Tables and Subtrees in the Directory”, 04/09/1992, S. Kille <draft-ietf-mhsds-subtrees-03.txt, .ps>
- “Representing the O/R Address hierarchy in the Directory Information Tree”, 04/09/1992, S. Kille <draft-ietf-mhsds-infotree-03.txt, .ps>
- “Use of the Directory to support mapping between X.400 and RFC 822 Addresses”, 04/09/1992, S. Kille <draft-ietf-mhsds-supmapping-03.txt, .ps>
- “Use of the Directory to support routing for RFC 822 and related protocols”, 04/09/1992, S. Kille <draft-ietf-mhsds-822dir-03.txt, .ps>
- “MHS use of Directory to support MHS Routing”, 04/17/1992, Steve Kille <draft-ietf-mhsds-routdirectory-03.txt>
- “MHS use of Directory to support MHS Content Conversion”, 11/10/1992, S. Kille <draft-ietf-mhsds-convert-01.txt, .ps>
- “Introducing Project Long Bud Internet Pilot Project for the Deployment of X.500 Directory Information in Support of X.400 Routing”, 06/21/1993, H. Alvestrand, K. Jordan, S. Langlois <draft-ietf-mhsds-long-bud-intro-00.txt>

CURRENT MEETING REPORT

Reported by Keven Jordan/Control Data Systems

Minutes of MHS-DS Working Group (MHSDS)

Introductions and Administrivia

The meeting opened with participant introductions and statements of interests in MHS-DS and Project Long Bud. The minutes from the Columbus IETF were then approved. It was pointed out that the MHSDS Working Group is now under the new Service Applications Area with Dave Crocker as the SAP Area Director.

Action Items

- Eight Internet-Drafts were updated by Steve Kille.
- Three of the Internet-Drafts were recommended to be progressed as Proposed Standards, after some minor editorial changes:

“Representing the O/R Address hierarchy in the Directory Information Tree”

“Representing Tables and Subtrees in the Directory”

“Use of the Directory to support mapping between X.400 and RFC 822 Addresses”

- The Internet-Draft, “MHS use of Directory to support MHS Routing,” was recommended to be published as an Experimental RFC after minor editorial changes.

Although there are some minor editorial and technical errors, no substantive errors exist in this document. This document now needs some higher degree of visibility. Two independent implementations already exist and provide proof of the basic concepts. Project Long Bud will show if the concepts scale well by deploying MHS-DS in the Internet. The working draft is stable enough to be published as an RFC. The Experimental status is initially appropriate since Long Bud feedback might result in important changes. The final goal is to place the routing document onto the standards track.

- The “Simple Profile” document is a facilitating document and it should be dropped once the pilot is further along and implementations become more mature. Most of the content of this document can be merged into the routing document, specifying some of the functionality as mandatory and some as optional. This has been done for other Internet protocol specifications.

- It was suggested that an acronym be created for the document set (e.g. MIME), as RFC numbers change, but the acronym would continue to be meaningful.

Editor's Note: A summary of documents and the working group's recommendations for progression is available via FTP or mail server from the remote directories as /ietf/mhsds/mhsds-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Summary of Actions from the Last Meeting

- Kevin Jordan to write an Internet-Draft providing an overview of the main set of MHS-DS RFCs. Status: not done, but some progress has finally been made.
- Harald Alvestrand to write pseudo code for the routing document. Status: not done.
- Steve Kille to update the document set and repost as Internet-Drafts. Status: done.
- Jim Romaguera to coordinate the writing of Project Long Bud definition and participation document(s). Status: done. The "Introduction to Project Long Bud" document was written and submitted as an Internet-Draft. Sylvain Langlois is the principal editor.
- Urs Eppenberger to write specifications for a tool which can be used to browse and verify X.500 routing and address mapping information. Status: not done. This tool is desirable but Urs can not commit the time to write its specification.
- Panos Tsigaridas to write specifications for, and begin implementation of, tools for synchronizing X.500 directory information with GO-MHS routing and mapping tables. Status: done. A beta version of a tool for reading information from the directory and generating GO-MHS tables can be found on the MHS-DS file server.
- Kevin Jordan and Long Bud Design Team to prepare an informal MHS-DS demonstration at IETF in Amsterdam. Status: done. Kevin demonstrated live MHS-DS technology and tools in the public terminal room of the RAI Conference Center.

Long Bud Report

- The meeting participants reviewed the extent of MHS-DS information available in the Internet DIT.

- All of the US MTA and organizational information provided in the routing documents available from the University of Wisconsin has been added under c=US. However, it seems that most of the MTAs registered under PRMD=XNREN no longer exist; most do not respond to connection requests. Perhaps the overall state of PRMD=XNREN needs to be reviewed.
- Countries in the DIT having MHS-DS routing entries include: the United States (US), Great Britain (GB), Denmark (DK), Germany (DE), Switzerland (CH), Spain (ES), Portugal (PT) and France (FR).

Review of "Introduction to Project Long Bud"

Editor's Note: A detailed review is available via FTP or mail server from the remote directories as /ietf/mhsds/mhsds-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Implementations and Tools

The current status of MHS-DS implementations and tools was reviewed. The Long Bud Status Report will provide details. The Status Report will also be updated periodically to reflect the current status.

DIT QOS

Reliability of the DSAs:

- MHS-DS requires the Internet DSA network to provide a good quality of service.
- The current QOS provided by the Internet DSA network is marginal at best. The US root-level DSAs seem to be particularly problematic, especially the DSAs which hold the top-level US domains under O=Internet.
- These problems might be solved by moving responsibility for top-level information to an organization which is funded well enough to provide good QOS. InterNIC probably qualifies, and has expressed an interest in providing this service.

Updated MHS-DS Internet-Drafts

Steve Kille briefly described what changes had been made to the MHS-DS Internet-Drafts since the previous revisions. Most of the changes were editorial in nature. A very few were more substantial. For example:

- Diagram change in the routing document.
- The representation for personal name was changed. An RDN with multiple AVAs is now used instead of the RFC 1327 representation used previously.
- O/R address syntax has been aligned to ISO syntax.

Some issues were raised on the routing document:

- The need for shared bilateral tables was introduced as a new concept. It was recommended that the `bilateralTable` attribute be changed to a sequence of DNs. This would allow a community of MTAs, e.g. the GO-MHS community, to share a potentially large table of information about MTAs. This could be used, for example, to establish a basis for deciding whether or not a connection request should be accepted or rejected. If an MTA outside of the community attempts to create a connection to an MTA within the community, the internal MTA could reject the connection after discovering that the caller is not registered in the shared bilateral table.

In addition to using the shared table, some MTAs might also have a need for maintaining a bilateral table which records agreements which are truly bilateral.

Thus, there appears to be a legitimate need for defining the `bilateralTable` attribute as a sequence of DNs. It was decided that this should be discussed further on the mailing list.

- ‘Next Tree First’ routing failure action when the top of a private routing tree is reached needs further discussion. This change needs to be discussed off-line. Steve, Harald, Kevin and Julian (in absentia) will discuss this.
- It was pointed out that an ‘Initiator Calling Address’ attribute may be needed. This will be discussed further on the mailing list.

Tutorial BOF

A tutorial BOF was scheduled for late in the afternoon. Thanks to Kevin for giving the tutorial on such short notice. It was well received.

Planning for the Next Meeting

- MHS-DS will schedule two time slots at the Houston IETF meeting.
- Four of the Internet-Drafts should have been progressed as RFCs by then.
- Progress on Long Bud will be reviewed.
- A new revision of the “Introduction to Project Long Bud” Internet-Draft will have been distributed, and its disposition will be discussed.
- The remaining four MHS-DS Internet-Drafts will be discussed.

Attendees

| | |
|--------------------------|--|
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Per Andersson | pa@cdg.chalmers.se |
| Piet Bovenga | p.bovenga@uci.kun.nl |
| C. Allan Cargille | allan.cargille@cs.wisc.edu |
| Robert Cooney | cooney@wnyose.nctsw.navy.mil |
| David Crocker | dcrocker@mordor.stanford.edu |
| Urs Eppenberger | eppenberger@switch.ch |
| Tony Genovese | genovese@es.net |
| Jan Hansen | Jan.Hansen@teknologi.agderforskning.no |
| Steven Horowitz | witz@chipcom.com |
| Jeroen Houttuin | houttuin@rare.nl |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Ola Johansson | ojn@tip.net |
| Kevin Jordan | Kevin.E.Jordan@cdc.com |
| Peter Jurg | jurg@surfnet.nl |
| Anders Karlsson | sak@cdg.chalmers.se |
| Steve Kille | S.Kille@isode.com |
| Paul Klarenberg | klarenberg@netconsult.ch |
| Bruno Koechlin | Bruno.Koechlin@inria.fr |
| Arnold Krechel | krechel@gmd.de |
| Sylvain Langlois | Sylvain.Langlois@exp.edf.fr |
| Erik Lawaetz | erik.lawaetz@uni-c.dk |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| Ignacio Martinez | martinez@rediris.es |
| Linda Millington | l.millington@noc.ulcc.ac.uk |
| Paul-Andre Pays | pays@faugeres.inria.fr |
| Catherine Pierre-Radenac | catherine.pierre-radenac@inria.fr |
| Jim Romaguera | romaguera@netconsult.ch |
| Marjo Rottschaefer | |

| | |
|--------------------------|----------------------------|
| Panos-Gavriil Tsigaridas | Tsigaridas@fokus.gmd.de |
| Eftimios Tsigros | tsigros@helios.iihe.rtt.be |
| Paul Vetter | |
| Peter Yee | yee@atlas.arc.nasa.gov |
| Steve Zeber | zeber@stc.nato.int |

2.7.3 Minimal OSI Upper-Layers (THINOSI)

Charter

Chair(s):

Peter Furniss, p.furniss@ulcc.ac.uk

Mailing Lists:

General Discussion: thinosi@ulcc.ac.uk

To Subscribe: thinosi-request@ulcc.ac.uk

Archive: pluto.ulcc.ac.uk:/ulcc/thinosi/thinosi-mail-archive.txt

Description of Working Group:

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

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

The working group is chartered to produce two documents:

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

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

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

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

Goals and Milestones:

- May 1993 Post an Internet-Draft for “Skinny bits for Byte-Stream.”
- Aug 1993 Post an Internet-Draft for “Skinny Bits for Directory.”
- Dec 1993 Submit the “Skinny Bits for Byte-Stream” specification to the IESG for consideration as a Proposed Standard.
- Mar 1994 Submit the “Skinny Bits for Directory” specification to the IESG for consideration as a Proposed Standard.

Internet-Drafts:

“Octet sequences for upper-layer OSI to support basic communications applications”, 08/09/1993, P. Furniss <draft-ietf-thinosi-cookbook-00.txt>

CURRENT MEETING REPORT

Reported by Peter Furniss/Consultant

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

The THINOSI Working Group met on 15th July at the Amsterdam IETF. Most of the time was devoted to reviewing the upper-layer cookbook.

Upper-Layer Cookbook

Various issues were discussed in the review of the cookbook. A point that came up more than once was how the cookbook should relate to the parallel work in the OSI regional workshops (the Common Upper-Layer Requirements Part 3: Minimal OSI profile (CULR-3)) and in X/Open (specification of use of the XTI interface for minimal OSI (XTI/mOSI)).

The three (or four) categories of application that the cookbook can support need to be further clarified, making the distinctions purely on which OSI facilities are used and not the use of ASN.1. Examples of application protocols in each category would be useful. If possible, the categories should be linked to those in CULR-3 and XTI/mOSI. This may not be straightforward due to the different approaches of the three documents: OSI-style profile, API definition and implementor's profile/respecification. Josee Auber will attempt to compare and contrast the three approaches in a message to the THINOSI list.

The possibility of the cookbook having a formal statement of compliance to CULR-3 was discussed. CULR-3 defines compliance statements by which another specification (e.g. the cookbook) can state that its use and support of the OSI facilities complies with CULR-3. However, CULR-3 is being developed in the OSI regional workshops (OIW leading, EWOS involved) and will probably have to follow the conformance requirements of the base OSI standards. Some of the base standards have over-enthusiastic conformance requirements, which go beyond the cookbook target, which is "interworking with conformant implementations." This is especially true for presentation. Moves are afoot in ISO/IEC to get this sorted out, but for the time being at least the cookbook should not be committed to comply with CULR-3.

The cookbook specifies the use of indefinite lengths where possible for sending the Presentation PCI (this is equivalent to the "Canonical Encoding Rules," a newly-defined subset of BER). There had been a suggestion that the opposite choice (Distinguished Encoding Rules—definite-length throughout) would be a preferable simplification. Which is found to be the easier to encode or decode is considerably a matter of the coding approach taken. Peter claimed that for these supporting layers (as distinct from the encoding of application protocols) the canonical choice was best.

Peter will attempt to get a new draft of the cookbook out by the end of July and submit it for posting as an Internet-Draft.

The eventual status of the cookbook was discussed, and it was believed that it should be targeted for the standards track, as the specification of the supporting protocol layers for the relevant applications.

Revision of Charter

Since the charter was written (following the Washington BOF), the coverage of the cookbook has changed to more than just “byte-stream” (although the amount of new text is small), and the charter needs to be changed to reflect this. Other possible changes were discussed. It was concluded that the intended thinDAP document should stay in the plan, but the dates may need revision. Given the nature of the cookbook, it will be worth considering the production of some very brief “mapping specifications” that would state precisely how particular applications (Z39.50 for example) would use the cookbook.

Peter Furniss will work up a draft revision and post it to the list.

Over the Fence: Activities in Other Arenas

CULR-3 has been revised again (June 1993) and it is intended to be submitted for ballot to become an ISP early next year.

The XTI/mOSI specification will shortly be published by X/Open as a “Preliminary Specification.” A Preliminary Specification is valid for a year, and may change subsequently.

Implementation Plans

Peter Furniss was expecting to start extending the X/osi code to a more general THINOSI implementation, with XTI/mOSI as the upper interface. Terry Sullivan is very keen to start something to support Z39.50.

Attendees

| | |
|-----------------|-----------------------------------|
| Josee Auber | Josee_Auber@hpgnd.grenoble.hp.com |
| Michael Brescia | |
| David Crocker | dcrocker@mordor.stanford.edu |
| Walid Dabbous | Walid.Dabbous@sophia.inria.fr |
| Peter Furniss | p.furniss@ulcc.ac.uk |

Brian May

Mark Needleman

Kamlesh Tewani

Brian.May@mel.dit.csiro.au

mhn@stubbs.ucop.edu

ktt@arch2.att.com

2.7.4 Network Database (NETDATA)

Charter

Chair(s):

Daisy Rose, daisy@watson.ibm.com

Mailing Lists:

General Discussion: ietf-ndb@ucdavis.edu

To Subscribe: ietf-ndb-request@ucdavis.edu

Archive:

Description of Working Group:

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

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

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

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

Goals and Milestones:

- | | |
|------|---|
| Done | Review and approve the charter, making any changes necessary. Examine needs, resources for this network database protocol and define the scope of work. Begin work on a framework for the solution. Assign writing assignments for first draft of the document. |
| Done | First draft to be completed. |
| Done | Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting. |
| Done | Continue revisions based on comments received at meeting and e-mail. Start making document an Internet-Draft. |

- Mar 1992 Review final draft. If it is OK, give it to IESG for publication as an RFC.
- Jun 1992 Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

2.7.5 Network Printing Protocol (NPP)

Charter

Chair(s):

Glenn Trewitt, trewitt@pa.dec.com

Mailing Lists:

General Discussion: print-wg@pa.dec.com

To Subscribe: print-wg-request@pa.dec.com

Archive:

Description of Working Group:

The Network Printing Working Group has the goal of pursuing those issues which will facilitate the use of printers in an internetworking environment. In pursuit of this goal it is expected that we will present one or more printing protocols for consideration for standards status in the Internet community.

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

Goals and Milestones:

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

Request For Comments:

RFC 1179 "Line Printer Daemon Protocol"

2.7.6 Service Location Protocol (SVRLOC)

Charter

Chair(s):

John Veizades, veizades@wco.ftp.com
Scott Kaplan, scott@wco.ftp.com

Mailing Lists:

General Discussion: srv-location@apple.com
To Subscribe: srv-location-request@apple.com
Archive: [apple.com:~/pub/srv-location/svr-loc-archive](ftp://apple.com/~pub/srv-location/svr-loc-archive)

Description of Working Group:

The Service Location Working Group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based service location protocols. Protocols that would be designed by this group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the service location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this group to understand the operation of service location and then determine the correct action in their view whether it be to use current protocols to suggest a service location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

- | | |
|----------|--|
| Done | Open discussion and determine if a working group should be formed. |
| Done | Continue discussion trying to refine the problem statement and possible resolutions. |
| Jul 1991 | Do we take the RFC track or do we write a report on our conclusion and leave it at that? |

Internet-Drafts:

“Resource Location Protocol”, 03/12/1993, S. Kaplan <draft-ietf-svrloc-resloc-00.txt, .ps>

CURRENT MEETING REPORT

Reported by John Veizades/FTP

Minutes of the Service Location Protocol Working Group (SVRLOC)

The meeting began by opening the floor to questions on the current status of SVRLOC's work.

Sun made a presentation on NIS+ which is being offered to the IETF along with the rest of the Sun ONC work (RPC, XDR and NFS) for standardization under the IETF umbrella. The presentation was requested to understand the NIS+ work and to see if the current service location proposal will solve the issues addressed by NIS+.

The following is a list of issues that need to be resolved before the document can go down the standards track:

- An architectural overview needs to be added to the document.
- Security considerations for authentication, privacy and spoofing—some sort of awareness of these issues needs to be added to the document.
- Addresses—to be able to run over multiple network protocols, a standard for address encoding needs to be put in place. Suggestions included taking the defined address specifications in the sockets.h file and registering them through the IANA.
- A length field should be in the packet.
- Language and character sets—the locale should be sent using the ISO standard locale encoding, and character sets would be specified for every string. The suggestion was made that services may want to register one service entity for each language instances that is available. For instance, if a particular service supports French, English and Spanish, one service would be registered for each language, and user agents requesting a particular language would be able to filter on the language type to acquire the appropriate service for their language needs.
- Rendezvous mechanism for specifying the end point of the answering service (address, port and other information)—the rendezvous information is used by the particular user agent service stub to make the connection to the appropriate service endpoint on the service agent. This will also allow directory agents to respond for service agents, and for service agents to return service-specific rendezvous information to the upper layer protocol. For example:

address type=IP; address: 90.1.0.12; port: 98; service info: ATS3=0

would be a string that may be returned from a modem pool to be used by the serial line service to send configuration information to the modem pool server to get the particular type of service specified by the user agent.

- Examples for several common services (e.g. printing, FTP, mail server, name server, and network management trap).
- Multicast addresses should be acquired from the IANA.

A technical presentation was given in Thursday's open plenary, outlining the service location protocol and giving status information on the service location protocol proposal. The work was well received by the audience.

The latest version of the documents can be found on: wco.ftp.com/resloc.

Attendees

| | |
|-------------------|--|
| Steve Alexander | stevea@lachman.com |
| Stefan Braun | smb@cs.tu-berlin.de |
| Eric Fleischman | ericf@act.boeing.com |
| Thomas Kaepfner | kaepfner%heidelbg.vnet@ibmpa.ibm.com |
| Scott Kaplan | scott@wco.ftp.com |
| Andrew Knutsen | andrewk@sco.com |
| John Larson | jl Larson@parc.xerox.com |
| Tony Li | tli@cisco.com |
| Paolo Malara | malara@crs4.it |
| Chuck McManis | chuck.mcmanis@eng.sun.com |
| John Veizades | veizades@wco.ftp.com |
| Steven Waldbusser | waldbusser@andrew.cmu.edu |

2.7.7 Trusted Network File Systems (TNFS)

Charter

Chair(s):

Fred Glover, fglover@zk3.dec.com

Mailing Lists:

General Discussion: tnfs@wdl1.wdl.loral.com

To Subscribe: tnfs-request@wdl1.wdl.loral.com

Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Trusted Network File System Working Group is chartered to define protocol extensions to the Network File System (NFS) Version 2 protocol which support network file access in a Multilevel Secure (MLS) Internet environment. MLS functionality includes Mandatory Access Control (MAC), Discretionary Access Control (DAC), authentication, auditing, documentation, and other items as identified in the Trusted Computer System Evaluation Criteria (TC-SEC) and Compartmented Mode Workstation (CMW) documents.

The primary objective of this working group is to specify extensions to the NFS V2 protocol which support network file access between MLS systems. It is intended that these extensions should introduce only a minimal impact on the existing NFS V2 environment, and that unmodified NFS V2 clients and servers will continue to be fully supported.

Transferring information between MLS systems requires exchanging additional security information along with the file data. The general approach to be used in extending the NFS V2 protocol is to transport additional user context in the form of an extended NFS UNIX style credential between a Trusted NFS (TNFS) client and server, and to map that context into the appropriate server security policies which address file access. In addition, file security attributes are to be returned with each TNFS procedure call. Otherwise, the NFS V2 protocol remains essentially unchanged.

The Trusted System Interoperability Group (TSIG) has already developed a specification which defines a set of MLS extensions for NFS V2, and has also planned for the future integration of Kerberos as the authentication mechanism. The TNFS Working Group should be able to use the TSIG Trusted NFS document as a foundation, and to complete the IETF TNFS specification within the next 3-6 months.

Goals and Milestones:

- Mar 1991 Verify the interoperability of TNFS implementations at the 1992 NFS Connectionathon.
- Done Review and approve the TNFS Working Group Charter, review revised TSIG TNFS Specification, and publish a proposed standard following the July meeting.
- Jul 1991 Review revised TSIG TNFS specification.
- Oct 1991 Review outstanding comments/issues from mailing list.
- Oct 1991 Make any final revisions to TNFS document based on comments, issues, and interoperability testing.
- Nov 1991 Publish a Proposed Standard following the July meeting.
- Mar 1992 Request IESG to make the revised document a Draft Standard.

Internet-Drafts:

“A Specification of Trusted NFS (TNFS) Protocol Extensions”, 07/23/1991,
Fred Glover <draft-ietf-tnfs-spec-03.txt>

2.8 Transport Area

Director(s):

- Allison Mankin: mankin@cmf.nrl.navy.mil

Area Summary reported by Allison Mankin/Naval Research Laboratory

Transport Area Directorate

Dave Borman dab@cray.com
Sally Floyd floyd@ee.lbl.gov
Jim Hughes hughes@hughes.network.com
Matt Mathis mathis@pele.psc.edu

The Transport Area

The Transport Area and the Service Application Area were separated from each other at the 26th IETF meeting in Columbus. The Transport Area deals with protocols and algorithms that provide end-to-end transmission services in the Internet. We maintain the notion of transport services, not just transport “protocols,” because of the increasing variety of end-to-end requirements that the Internet is meeting or will be expected to meet in the near future. TCP in itself supports applications of many different characteristics, from brief transaction-like exchanges in X-windows on through hypertext support (HTTP running over TCP) and many more. Then there is the range of distributed file systems, over transport services such as the combination of RPC and UDP. Then we come to audio and video, receiving their service in the MBONE from the new transport protocol under development in the AVT Working Group, supported by UDP over multicast IP.

So far the Internet is not only holding its own with these major transport services, but they are continuing to spread and to multiply. The Transport Area includes AVT and the Multiparty Multimedia Session Control Working Group, which will lead to new ways in which multiparty services can be provided. The TCP Large Windows Working Group continues to refine the extensions that allow TCP to work effectively in gigabit networks. As to future Transport work areas, the Transport Area will be watching the progress of research work such as RSVP and will support such efforts when they are ready for engineering and deployment. We have recently received a proposal for work to begin on a mobile transport protocol.

We hosted one BOF in Amsterdam, following email review of a proposal and draft specification. The TMUX BOF, chaired by Jim Barnes, gathered to work on decreasing the streams of small packets between host pairs that are generated by terminal servers. These

small packet streams cause there to be a perception that proprietary protocols can perform more efficiently than TCP/IP. The TMUX BOF is summarized further below.

Transport Area Working Groups

The Transport Area currently has three working groups:

- Audio/Video Transport (AVT)

This group did not meet in Amsterdam. Protocol specifications and experimental implementations are nearing completion. May have an interim meeting over the Internet.

- Multiparty Multimedia Session Control (MMUSIC)

Summarized below.

- TCP Large Windows (TCPLW)

Summarized below.

TCP Multiplexing BOF (TMUX)

One of the problems with the use of terminal servers is the large number of small packets they can generate. Frequently, most of these packets are destined for only one or two hosts. The Connection Multiplexing Protocol (CMP), an approach to decreasing the number of these by multiplexing between the application and TCP (draft-*cameron-cmp-01.txt*), was described by Peter Cameron, along with a brief description of some implementation results. After a general discussion of the proposal, Dave Crocker presented a counter proposal that did the multiplexing between IP and TCP instead. Discussion of the two proposals continued with advantages and disadvantages for each proposal. The consensus of those attending the BOF was that CMP addressed a valid problem, but at the wrong place. A request was made that the CMP developers try an implementation of the IP multiplexing proposal (TMUX) to determine whether that was a valid solution. Work continued soon after the Amsterdam meeting, with the result that a specification is now available (draft-*cameron-tmux-00.txt* by Pete Cameron and Dave Crocker).

Multiparty Multimedia Session Control Working Group (MMUSIC)

The MMUSIC Working Group met officially as a working group for the first time in Amsterdam. Two sessions were held that were multicast over the MBONE. The first meeting was used to set the context and to discuss the progress made since the BOFs held at the last

IETF. After review of the modified charter, we discussed proposals for a set of common terminology, an end-system architecture, the MMUSIC protocol requirements, implementation considerations and conference styles. To narrow the scope of the discussion, we emphasized the need to think in terms of a “version 0” negotiation protocol. During the second meeting, the foundations for the strawman protocol were discussed and included proposals for the definition of session state and for naming conference components. After thorough descriptions were given of the main protocol assumptions, we delved into the basic message types, examples of how they might be used, and default session policies.

TCP Large Windows Working Group (TCPLW)

The TCP Large Windows Working Group met for one session in Amsterdam. The major accomplishments were:

- Review of draft-ietf-tcplw-extensions-00.txt
- Consideration of advancing RFC 1323 to Draft Standard

The reviewed document is a compilation of a few bug fixes and clarifications that need to be made to RFC 1323. It was compiled by Bob Braden. It also includes a pseudo-code presentation of the RFC 1323 TCP extensions. Bob and David Borman led a walk through of the document, to help explain it and to find out if any other changes that are needed to RFC 1323 were missed. The contents of the draft will be folded into the text of RFC 1323 before that specification is submitted for Draft Standard status.

CURRENT MEETING REPORT

Reported by Jim Barnes/Xylogics

Minutes of the TCP Multiplexing BOF (TMUX)

Agenda

- Introduction
- CMP Presentation
- Discussion
- Alternative proposals
- Where do we go from here

Discussion

The TMUX BOF began with Peter Cameron's presentation on the Connection Multiplexing Protocol (CMP). The CMP protocol is defined in the Internet-Draft: draft-cameron-cmp-01.txt.

A couple of changes have been made since the last version was published. One change was the removal of the close reply message type. This message type is replaced by just sending the close message in response to a received close, just as TCP sends a FIN in response to a received FIN.

During the following discussion, a number of issues were raised:

- How would the implementation of CMP on top of TCP affect the TCP window dynamics?
- CMP may be fine for multiplexing a large number of small packets but if FTP connections are multiplexed, the FTP subconnections will fight each other for available window space.
- The idea of falling back to a normal TCP connection if a request to open a CMP connection fails was well received.
- There will necessarily be bandwidth reduction due to the multiplex protocol headers.
- A misbehaving CMP client may exceed the allowed credit and force the receiver to control the flood with the TCP window mechanism.
- Performance versus complexity of implementation was mentioned as a possible issue.

Kent Malave briefly described his experiences when multiplexing SPX packets. The reasons for doing the multiplexing, and the experiences in implementing the protocol, were similar to those behind CMP.

Dave Crocker gave a brief presentation of an alternate proposal. TMux is a multiplexing protocol between the IP and TCP layers (in contrast to CMP which is a layer on top of TCP). The apparent advantages of TMux over CMP were largely due to the simplicity of the design. Someone noted that there were no delay timers in the protocol.

When discussing the advantages of one proposal over another, it was noted that:

- Data loss in CMP will cause delays in all other TCP data over that connection until the lost packet was retransmitted.
- TMux will require that the multiplexed packets are padded so that each multiplexed packet starts on a word boundary.
- A single bit error in a CMP packet requires that the entire packet be retransmitted. A single bit error in a TMux packet will require only the erroneous packet to be retransmitted, but the other multiplexed packets in the IP packet can be delivered to the application.

Allison Mankin discussed the concern that the IESG has with changing the architecture. When describing what would be significant issues in the IESG approval process, steady-state performance was deemed to not be a deciding factor. The main issues would be the required architectural changes to the protocol stack, the ease of implementation of any solution, and the behavior of the protocol under aberrant conditions.

When the chair requested a consensus on what to do next, the following suggestions were made:

- A test implementation of TMux should be done and compared with CMP keeping the above decision criteria in mind.
- The TMux proposal should be published as an Internet-Draft.
- Greg Minshall will investigate the availability of documentation describing the NPI protocol between the IP and TCP layers in System V.4.

Attendees

| | |
|------------------|------------------------------|
| Jim Barnes | barnes@xylogics.com |
| Julian Bates | bates@xylogics.com |
| David Borman | dab@cray.com |
| Peter Cameron | cameron@xylint.co.uk |
| Les Clyne | l.clyne@jnt.ac.uk |
| David Crocker | dcrocker@mordor.stanford.edu |
| Sun-Kwan Kimn | sunkimn@hp.com |
| Andrew Knutsen | andrewk@sco.com |
| John Krawczyk | jkrawczyk@wellfleet.com |
| Kent Malave | kent@bach.austin.ibm.com |
| Gary Malkin | gmalkin@xylogics.com |
| Allison Mankin | mankin@cmf.nrl.navy.mil |
| Marjo Mercado | marjo@cup.hp.com |
| Greg Minshall | minshall@wc.novell.com |
| Douglas Williams | dougw@ralvmg.vnet.ibm.com |
| Gordon Young | young@xylogics.com |

CMP Overview

CMP TECHNICAL OVERVIEW

Pete Cameron
Xylogics International Ltd.
Email: cameron@xylint.co.uk
Phone: +44 908 222112

CMP Overview

Background

- About 1 year ago a small team at Xylogics International started looking at reasons why people used non-networked terminal connection solutions rather than networked solutions.
- One of the reasons we found, was the (perceived) problem of higher host load from networked solutions on high port counts.
- After looking at many solutions, we came up with the idea of multiplexing many ports into a single TCP pipe.
- We realised that this was a general solution, so we have proposed it as CMP to the Internet community.


CMP Overview

The Problem

- TELNET and Rlogin give high host overhead.
 - When using terminal servers, typically have large number of remote sessions.
- ⇒ Host is very heavily loaded when using terminal servers.
- Host cpu load caused by :
 - One (or more) context switch for each packet
 - Ethernet drivers overhead the same for all packet sizes
 - One interrupt for every packet in or out (at least)
- ⇒ Overheads are per packet, not per character

CMP Overview

Telnet Daemon Example (System V.4)



- This example shows the layout of the modules involved in transmitting a character on System V.4 based systems with streams.
- The problem here is due to the large overhead of passing a packet containing a single character through the stream.


CMP Overview

The Proposed solution: CMP

- CMP multiplexes all of the TELNET and Rlogin packets from a single system (eg. terminal server) to a single host.
- Could support other low volume protocols.
- CMP uses a short delay to allow all data (from many ports) sent in a short space of time to be placed in a single packet.
 - The number of packets the host has to deal with is greatly reduced, which lowers the host load.
- ✓ Preliminary performance results are very encouraging.

CMP Overview

Design overview




- Multiplexed connection can be initiated from either end.
- Once multiplexed connection is made, either end can start subconnections.
- When last subconnection is broken, multiplexed connection is killed.

CMP Overview

Brief details of the protocol

- Message based, fully symmetric protocol.
- Simulates full TCP connection on each multiplexed connection.
- Short 4 byte header on each message.
- Supports standard and urgent data.
- Per subconnection flow control.
- See *draft-cameron-cmp-01.txt* for full details.



CMP Overview

The messages

| Type | USE | Len field contains |
|------|--------------------------------|------------------------|
| 0 | Data transfer | # of octets of data |
| 1 | Pointer to (later) urgent data | 2 |
| 2 | Open new subconnection | 6 |
| 3 | Reply to OPEN message | 6 |
| 4 | Close one end of subconnection | 2 |
| 6 | Credit update | # octets of new credit |

Dest-ID field is a unique number specifying which remote session the packet is for. Typically a table index or socket number.

CMP Overview

Test results

Test program: Simulation of 20 typists at 5 cps with host echo.

Test system: 33MHz SPARC with System V.4

Results:

Without CMP - Uses 32% of Host CPU

With CMP - Uses 14% of Host CPU

CMP Overview

Conclusion

- CMP provides a mechanism for multiplexing many low volume connections into one higher volume connection.
- Because the network software has only to deal with one packet for several ports, it reduces the host overhead.
- Whilst this was designed primarily for use over a LAN, nothing in the design precludes use over other types of network (eg. WAN).

TMUX: MULTIPLEXING TRANSPORT SEGMENTS OVER IP

D. CROCKER
SILICON GRAPHICS, INC.
DCROCKER@SGL.COM

The problem

- MANY SMALL PACKETS BETWEEN HOST PAIRS
 - Lots of network packet handling
 - Lots of host interrupts

D. Crocker, TMux

2

Approach


- FUNCTIONAL REQUIREMENT
 - Fewer packets
That's all; anything else is extra
 - Minimal protocol modification
No redundant mechanisms
- TMux
 - IP is a host-to-host protocol
 - Transport protocols are between "processes"
 - Hence, multiplex transport segments in the same IP datagram

D. Crocker, TMux

3

Formats

|IP|TM hdr|Tport segment|TM hdr|Tport segment|...|



D. Crocker, TMux

4

2.8.1 Audio/Video Transport (AVT)

Charter

Chair(s):

Stephen Casner, casner@isi.edu

Mailing Lists:

General Discussion: rem-conf@es.net

To Subscribe: rem-conf-request@es.net

Archive: [nic.es.net:~/ietf/rem-conf/av-transport-archive](http://nic.es.net/~ietf/rem-conf/av-transport-archive)

Description of Working Group:

The Audio/Video Transport Working Group was formed to specify experimental protocols for real-time transmission of audio and video over UDP and IP multicast. The focus of this group is near-term and its purpose is to integrate and coordinate the current AV transport efforts of existing research activities. No standards-track protocols are expected to be produced because UDP transmission of audio and video is only sufficient for small-scale experiments over fast portions of the Internet. However, the transport protocols produced by this working group should be useful on a larger scale in the future in conjunction with additional protocols to access network-level resource management mechanisms. Those mechanisms, research efforts now, will provide low-delay service and guard against unfair consumption of bandwidth by audio/video traffic.

Similarly, initial experiments can work without any connection establishment procedure so long as a priori agreements on port numbers and coding types have been made. To go beyond that, we will need to address simple control protocols as well. Since IP multicast traffic may be received by anyone, the control protocols must handle authentication and key exchange so that the audio/video data can be encrypted. More sophisticated connection management is also the subject of current research. It is expected that standards-track protocols integrating transport, resource management, and connection management will be the result of later working group efforts.

The AVT Working Group may design independent protocols specific to each medium, or a common, lightweight, real-time transport protocol may be extracted. Sequencing of packets and synchronization among streams are important functions, so one issue is the form of timestamps and/or sequence numbers to be used. The working group will not focus on compression or coding algorithms which are domain of higher layers.

Goals and Milestones:

Done Define the scope of the working group, and who might contribute. The first step will be to solicit contributions of potential protocols from projects that

- have already developed packet audio and video. From these contributions the group will distill the appropriate protocol features.
- Done Conduct a teleconference working group meeting using a combination of packet audio and telephone. The topic will be a discussion of issues to be resolved in the process of synthesizing a new protocol.
- Done Review contributions of existing protocols, and discuss which features should be included and tradeoffs of different methods. Make writing assignments for first-draft documents.
- Done Post an Internet-Draft of the lightweight audio/video transport protocol.
- May 1993 Post a revision of the AVT protocol addressing new work and security options as an Internet-Draft.
- Jun 1993 Submit the AVT protocol to the IESG for consideration as an Experimental Protocol.

Internet-Drafts:

- “RTP: A Transport Protocol for Real-Time Applications”, 12/16/1992, H. Schulzrinne, S. Casner <draft-ietf-avt-rtp-03.txt, .ps>
- “Media Encodings”, 12/16/1992, H. Schulzrinne <draft-ietf-avt-encodings-02.txt>
- “Sample Profile for the Use of RTP for Audio and Video Conferences with Minimal Control”, 12/16/1992, H. Schulzrinne <draft-ietf-avt-profile-02.txt>
- “Packetization of H.261 video streams”, 03/11/1993, T. Turletti, C. Huitema <draft-ietf-avt-video-packet-01.txt>

2.8.2 Multiparty Multimedia Session Control (MMUSIC)

Charter

Chair(s):

Eve Schooler, schooler@isi.edu
Abel Weinrib, abel@bellcore.com

Mailing Lists:

General Discussion: confctrl@isi.edu
To Subscribe: confctrl-request@isi.edu
Archive: venera.isi.edu:~/confctrl/confctrl.mail

Description of Working Group:

The demand for Internet teleconferencing has arrived, yet an infrastructure to support this demand is barely in place. Multimedia session control, defined as the management and coordination of multiple sessions and their multiple users in multiple media (e.g., audio, video), is one component of the infrastructure. The Multiparty Multimedia Session Control Working Group is chartered to design and specify a protocol to perform these functions.

The protocol will provide negotiation for session membership, underlying communication topology and media configuration. In particular, the protocol will support a user initiating a multimedia multiparty session with other users ("calling" other users) over the Internet by allowing a teleconferencing application on one workstation to explicitly rendezvous with teleconferencing applications running on remote workstations. Defining a standard protocol will enable session-level interoperability between different teleconferencing implementations.

The focus of the working group is to design a session negotiation protocol that is tailored to support tightly-controlled conferences. The MBONE currently carries primarily loosely-controlled sessions, i.e., sessions with little to no interaction among members and with no arbitration facility, security, or coordination of quality-of-service options for time-critical media. Users may learn of available sessions using the "sd" utility or other out of band mechanisms (e.g., email). However, there is clearly also a need for tightly-controlled sessions that provide mechanisms for directly contacting other users to initiate a session and for negotiating conference parameters such as membership, media encodings and encryption keys. In addition, these sessions should support renegotiation during a session, for example to add or delete members or change the media encoding. It is possible that the protocol will, in the limiting case, also support loosely-controlled sessions.

The main goal of the working group will be to specify the session control protocol for use within teleconferencing software over the Internet. The working

group will focus on the aspects of the session control problem that are well understood, while keeping an eye on evolving research issues. Toward this end, the working group has made an inventory of existing conferencing systems and their session control protocols. The working group will document the requirements of the existing prototypes as a basis for the protocol development. The working group will iteratively refine the protocol based on implementation and operational experience.

Furthermore, the working group will coordinate with other efforts related to multimedia conferencing, such as directory services for cataloguing users and conferences, the RTP and RTCP protocols developed by the Audio/Video Transport Working Group, resource reservation and management at the network level, and schemes for multicast address allocation.

Goals and Milestones:

- May 1993 Hold an on-line working group meeting to discuss the conference control framework, the relevant terminology, a functional taxonomy and how different conversational styles place requirements on session protocols.
- Jun 1993 Submit the Conference Session Control Protocol to the IESG for consideration as an Experimental Protocol.
- Aug 1993 Post an Internet-Draft describing the Session Control Requirements.
- Nov 1993 Post an Internet-Draft of the Session Control Protocol.
- Mar 1994 Submit a revised Internet-Draft based on implementation experience.

CURRENT MEETING REPORT

Reported by Eve Schooler/ISI

Minutes of the Multiparty Multimedia Session Control Working Group (MMUSIC)

An on-line copy of the minutes and the accompanying slides may be found in the directory `venera.isi.edu:confctrl/minutes` as files `ietf.7.93` and `slides.7.93.ps`.

The MMUSIC Working Group met officially for the first time in Amsterdam. We held two sessions that were multicast over the MBONE. The first meeting was used to set the context and to discuss the progress made since the BOFs held at the last IETF. During the second session, we began to lay the groundwork for a strawman MMUSIC protocol.

First Session: Context and Progress

After review of the modified charter, we discussed proposals for a set of common terminology, an end-system architecture, the MMUSIC protocol requirements, implementation considerations and conference styles. To narrow the scope of the discussion, we emphasized the need to think in terms of a “version 0” negotiation protocol.

Terminology, Framework, Requirements

Highlights from Lakshman’s (`lakshman@ms.uky.edu`) proposed session control glossary were presented (slides 4-8). The key points were:

- The differentiation between a “session” (an association of members for control) and a “conference” (a logical abstraction among multiple participants for multimedia real-time communication that consists not only of a control session) but also of related media associations and conference policies.
- The identification of the main system components for an end-system teleconferencing architecture as being the conference session manager, media agents and a resource manager.

Since “media” is an overloaded word, we are open to suggestions for a better term than “media association,” which is currently defined as the encapsulation of the transport (point-to-point or multipoint) in a single medium.

Some clarification was needed for the term “reflector participant,” a participant who neither generates nor terminates data but acts as a go-between. It is one of several participant types

that arise out of policy choices. Julio Escobar commented that it is somewhat of a misnomer since a reflector implies the “reflecting back” of data, and a reflector participant may be used in a variety of fashions (e.g., it may translate or combine data). A reflector might be considered a service access point.

In a change since last time, we emphasized that the conference session manager is not necessarily the central system component, as shown in the slide “Framework 1,” but that we think of it more as in the “Framework 2” slide. However, to a large extent the relationship among the various components is immaterial and implementation-specific. For instance, a third approach, where the conference session manager is part of one monolithic application, is equally valid. The working group focus is somewhat separate from the specifics of the framework choices, since we are primarily interested in the interaction between the MMUSIC negotiation protocol and the conference session manager.

Since the last meeting, we refined the session control protocol requirements (slide 11). They encompass several functions: those for distributed session management, dynamic membership management, session policy management, and domain specific tasks (e.g., media associations and configurations). These groups of functions reflect the management of a “conference” itself, and the management of its control, policy and media elements.

In addition, we need to distinguish between the policies that are carried by the protocol and those understood by the protocol. Does the protocol simply carry policy in the same way that it simply carries media information, as a payload or “bag of bits”? While session policy is not meant as an optional characteristic, since it is what defines the session type, policy enforcement is probably outside the scope of the protocol. However, policy enforcement will impact the degree to which we can provide session privacy and security.

The idea of advance reservation was a recurring topic, and one which needs further scrutiny. We expect the protocol to provide hooks for pre-scheduling sessions, though the conference session manager has no direct effect on reservations in the network, nor reservation strategies (optimistic versus pessimistic). Ambiguities remain about the definition of resources, since they occur at a number of levels (e.g., people, rooms, hardware devices, workstation capabilities, network bandwidth), and about how different policies will cause different outcomes for resource scheduling and contention resolution. One suggestion was to create a proper session MIB to assist with end-system management of configuration, capabilities and policies.

There was also speculation about the interaction between the session manager and media agents, for instance when the transport for a media agent fails but the control path is still functioning. Although the end-system architecture is somewhat outside the venue of the MMUSIC protocol, we expect that some system component implementations will enable up-calls from (down-calls to) media agents and that are conveyed to (from) the session manager either directly or indirectly. The session manager (media agents) may issue modifications as a consequence. Similar mechanisms are needed for a session chair to be able to turn on and off media agent data flows.

Implementation for the Internet

What are the building blocks needed for implementation and operation of a MMUSIC protocol in the Internet (slide 12)? Perhaps the biggest questions were:

- What transport platform is needed to support a general purpose negotiation service?
- To what degree do we need reliable multicast?
- How reliable does reliable have to be?

A critical, near term action item is to find an individual or a collection of individuals who can advise us on our options and recommend a solution. It was noted that INRIA is planning to make a reliable multicast solution available to the public shortly. Regardless of the approach taken, it was agreed that the interface to MMUSIC should be designed so that the requirements for the underlying service are clearly stated and that the actual choice may vary. Different transport implementations might be differentiated on their port number. The requisite comment was made that the goals of reliable delivery and scalability of sessions conflict with each other.

In addition, it was suggested that we investigate the universal identifier naming infrastructure already used in the WorldWide Web (WWW). Another correlation was noted between synchronous and asynchronous group negotiation, for instance for mailing list coordination. Yet the timing characteristics for conferee interactions differ by an order of magnitude between real-time conference session control and mailing list management.

Conversation Styles

Pre-IETF, we posted to the mailing list some musings on the relationship between conversation styles and their requirements on the underlying communication infrastructure. It was agreed that a number of other dimensions, besides size, need to be considered to describe the list of session types more completely. The question was raised about whether it is easier to move from tight to loose or loose to tight schemes when devising an approach, since the complicated scenarios occur somewhere in the middle of the continuum (slide 13).

There also was debate about the existence of an upper bound on the number of conferees generating media (actively participating). The united-nations model and the distributed simulation model are good counterexamples to an upper bound. A criticism about the original assumption is that we need to be careful not to introduce artifacts of the capabilities of the current generations of tools into our interaction model. Although a preliminary document on the range of conversation styles has been drafted, further details are needed.

Second Session: Outline for a Protocol

The foundations for the strawman protocol were discussed (slides 14-20) and included proposals for the definition of session state and for naming conference components. After thorough descriptions were given of the main protocol assumptions, we delved into the basic message types, examples of how they might be used and default session policies. A lively discussion ensued that helped us compile a list of outstanding issues and action items (slides 22 and 23).

Session State and Naming

There were no strong opinions about whether or not naming should be opaque or structured; in other words whether or not names should be based on arbitrary identifiers, or structured around common identifiers already in use, such as login ids, host addresses, port numbers and timestamps. In any event, the identifiers must be unique. The inclusion of a sequence number was felt to be overkill. There are no side effects if the `session_id` is based on the initiator's `member_id` and the initiator drops out of the session; the incorporation of the initiator's `member_id` is simply a technique to make the `session_id` unique. Aliases were deemed useful from an application standpoint, but not necessary for the operation of the session control protocol itself. The idea behind the aliases were to provide RTCP-like support, though there are other textual pieces of information that RTCP carries in addition to conferee names and the session name.

However, there are broader privacy concerns if we tie the `member_id` and `session_id` to identifiable naming structures. Also, should naming be any different if the conference session manager acts on behalf of a an individual user, conference room of participants, reflector participant (the proxy), or an automated service (the virtual user)?

We also need to think about naming for mobility, both at session setup (to forward requests when you have multiple addresses) and during long-lived sessions (to allow users to move around), for example, were individuals equipped with locator badges? Can we leverage off of the location services for naming transparency being designed in the MOBILEIP Working Group?

Protocol Assumptions

Questions about looser styles of conferencing came up. How does someone simply tune-in in a loose control fashion, given that we're thinking about requiring the participation of at least one member for a session to persist? One idea was to create a virtual member to "own" the session. This virtual member would not only establish the session, but also take responsibility to terminate it. The idea of a virtual member also could be applied to pre-scheduling sessions (again, this is different from reserving the network or other resources), since the virtual member would establish the session ahead of time and only participate from

a control standpoint. Another suggestion was to view this as allowing empty membership lists, since the virtual member is not an active member.

Basic Message Types

The basic message types do not in and of themselves provide a negotiation service; they are meant as building blocks. Whether or not they are delivered reliably is a separate issue. We proposed a three-phase commit handshake (Propose-Reply-Announce) for proposals needing negotiations. Suggestions about the messages that are under advisement:

- Add an optional “reason” field to the Reply message to make informed decisions about initiating another round of proposals after receiving a reject. This is useful for handling error messages when not in the correct state to receive a Proposal. More generally, this optional payload field could be added to both Reply and Propose messages:
 - In the Reply message: to allow a reject or an accept response to include hints that could assist with renegotiation. This results in four types of Replies.
 - In a Propose message: to provide enough information up-front to reduce the number of negotiated rounds.
- Differentiate between an Announce message that:
 - Has been agreed on versus one that a proposer has decided on alone.
 - Includes the delta versus the entire state of the conference. If the state is large, one may want to send the hash of the state. Does an Announce send state, operations, or both?
- How to handle/avoid a questionable Announce message? To cut down on false Announces, one option is to multicast all messages to all members.

Examples were provided of how the messages might be used, from simple scenarios (using an Announce to leave a session or to produce keep-alive messages) to session initiation or modification. We assumed that a proposal may be comprised of multiple operations. Although many hooks were discussed, version 0 may defer using some of them.

A key open issue is if the protocol requires serializability, i.e., that all proposals are acted on in the same order at all conference session managers involved in a session. We maintained that a sufficient measure of tight control can be enforced without serializability, and without requiring absolute global state consistency. To reduce conflicts, multicast Propose messages to all others, even if not involved in the accept phase.

Policies

Slide 20 introduced the main policies we expect to associate with a session. Underlined options represent the default choices, should no policy be chosen. We clarified that members are always allowed to leave a session, regardless of the termination policy. In fact there was a motion to replace explicit session termination by implicit termination when the membership count drops to zero, or after some duration beyond when the membership goes to zero (this would avoid odd behavior caused by the non-serializability of proposals). On a similar note, we may want to support a policy where one member is able to delete all other members in order to terminate a session.

We discussed the rigidity of session policies and the need to determine if different session policies conflict with one another, especially with regard to “static” sessions (e.g., unchangeable sessions in terms of their members, policies and media components).

The concern was that if there are communication failures that prohibit approval for a session change (e.g., when the policy is that *all* must approve), that this would result in a deadlock, a malfunctioning session that does not terminate, or a scenario where resources are never returned. Clearly, in filling in the protocol details we will need to differentiate between the receipt of a reject reply versus no reply at all. We will also need to state how policies will be handled (possibly become more relaxed) in the event of a communication failure. The communication failure may be due to an intermittent lapse in connectivity, to a person leaving their workstation unattended and not being able to immediately reply to a query, or to the member having left the conference at a network failure point and consequently being in the wrong state. It was felt that a conference session manager should behave like TCP reset after a failure and retain no previous state.

Even though we proposed an initially small set of policy choices, the richness and completeness of this set needs further scrutiny. To decide where version 0 falls on the session style continuum, we solicit input on suggested policies and their default values. Additionally, to what level of granularity should we institute policies? Do we need to have global policies? Per-operation policies? Per-initiator policies? Will policies be shared with media agents and other system components?

A good deal of discussion centered around policies about who may make proposals, since non-members may be restricted from initiating proposals, and in some cases not all members are proposers. These policies apply to changes in general, and are separate from who is needed to approve proposals (e.g., coordinate a vote or approve an initiation request). The solution proposed was that a non-member find a sponsor for a proposal, ensuring the notion of trusted membership. However, is this approach too stringent? Because of the premium on being a member versus a non-member, there is also interest in making assurances that one person isn't impersonating another.

More specifically, the issue boiled down to the question of joining a session. How does one join? Who does one contact? Again, the idea is to assign a “doorman” or “doormen.” For version 0, to support an open session in the sd style, a doorman could simply say yes all the time. Similarly, is there a more straightforward approach out there?

Outstanding Issues

How does floor control interact with session control? Is the notion of floor control its own protocol? With its own messages? We are looking to other projects, such as MiCE, to advise us on this.

Could the negotiation protocol be used for other purposes, such as a calendar scheduler, booking service, or even to measure agreement over topics during an ongoing conference? We are interested in someone studying the range of related applications for the MMUSIC protocol.

Attendees

| | |
|-------------------|-------------------------------|
| George Abe | abe@infonet.com |
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Axel Belinfante | Axel.Belinfante@cs.utwente.nl |
| Lou Berger | lberger@bbn.com |
| Jim Binkley | jrb@ibeam.intel.com |
| Carsten Bormann | cabo@cs.tu-berlin.de |
| Robert Braden | braden@isi.edu |
| Stefan Braun | smb@cs.tu-berlin.de |
| Michael Brescia | |
| Les Clyne | l.clyne@jnt.ac.uk |
| Walid Dabbous | Walid.Dabbous@sophia.inria.fr |
| Steve DeJarnett | steve@ibmpa.awdpa.ibm.com |
| Ed Ellesson | ellesson@vnet.ibm.com |
| Hans Eriksson | hans@sics.se |
| Julio Escobar | jescobar@bbn.com |
| Deborah Estrin | estrin@usc.edu |
| Osten Franberg | euaokf@eua.ericsson.se |
| David Ginsburg | ginsb@us-es.sel.de |
| Ron Greve | rgreve@cs.utwente.nl |
| Mark Handley | mhandley@cs.ucl.ac.uk |
| Geert Heijen | heijen@cs.utwente.nl |
| Rune Hjelsvold | Rune.Hjelsvold@idt.unit.no |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| Xinli Hou | xinli@cs.utwente.nl |
| Sascha Ignjatovic | sascha@veda.co.at |

| | |
|------------------|--------------------------------------|
| Phil Irely | pirey@relay.nswc.navy.mil |
| John Johnston | john@berlioz.nsc.com |
| Thomas Kaepfner | kaepfner%heidelbg.vnet@ibmpa.ibm.com |
| Peter Kirstein | P.Kirstein@cs.ucl.ac.uk |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| John Larson | jl Larson@parc.xerox.com |
| Mark Laubach | laubach@hpl.hp.com |
| Allison Mankin | mankin@cmf.nrl.navy.mil |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Shehzad Merchant | merchant@erg.sri.com |
| Donald Merritt | don@arl.army.mil |
| Topi Miettinen | tm86214@cs.tut.fi |
| Paul Milazzo | milazzo@bbn.com |
| Ronny Nilsen | Ronny.Nilsen@usit.uic.no |
| Zbigniew Opalka | zopalka@agile.com |
| Jorg Ott | jo@cs.tu-berlin.de |
| Mark Prior | mrp@itd.adelaide.edu.au |
| J. Mark Pullen | mpullen@cs.gmu.edu |
| Eve Schooler | schooler@isi.edu |
| Henk Sennema | sennema@sara.nl |
| A. Velu Sinha | avsinha@attmail.com |
| Kenneth Smith | kensmith@bnr.ca |
| Kamlesh Tewani | ktt@arch2.att.com |
| Claudio Topolcic | topolcic@cnri.reston.va.us |
| Antoine Trannoy | trannoy@crs4.it |
| Thierry Turletti | turletti@sophia.inria.fr |
| Guido van Rossum | guido@cwi.nl |
| Dono van-Mierop | dono_van_mierop@3mail.3com.com |
| Mario Vecchi | mpv@thumper.bellcore.com |
| Abel Weinrib | abel@bellcore.com |

Multiparty Multimedia Session Control Working Group (mmusic)

Eve Schooler <schooler@isi.edu>
Abel Weinrib <abel@bellcore.com>

13, 14 July 1993

Agenda

MMusic overview
Terminology
Framework
Protocol Requirements
Conference Styles
Implementation Considerations
Strawman Protocol
Action Items
Other Topics?

2

MMUSIC Overview

Motivation for work

- history

Working group charter

Changes from BOF

- name
- focus
- area

Milestones

- 8/93 Post an Internet Draft describing the Session Control requirements
- 11/93 Post an Internet Draft of the Session Control Protocol
- 3/94 Submit a revised Internet Draft based on implementation experience
- 6/94 Submit the Session Control protocol to the IESG for consideration as an Experimental Protocol

3

Terminology

Conference

- A logical abstraction among multiple participants for multimedia real-time communication. A conference consists of a control session, related media associations (?), and conference policies.

Conferee

- A participant in a conference.

Session

- An association of members for control; for instance, to control a conference with multiple conferees.

Media Association (?)

- Encapsulation of the transport (point-to-point or multipoint) in a single medium.

Conference Policies

- Rules regarding the style of interaction for a conference.

4

Terminology—control

(Multiparty Multuser) Session Control Protocol

- The protocol used for session control, the management and coordination of multiple sessions, and their multiple members in multiple media.

Distributed Control

- A control model where control functions are distributed among session managers.

Centralized Control

- A control model where control functions are the responsibility of a centralized agent.

Tight Control

- A session style in which state is actively shared among participants and that aims to keep state consistent among participants.

Loose Control

- A session style in which state information is passively shared among participants. In the extreme, no state sharing is performed.

5

Terminology—system components

Conference Session Manager

- A software entity that resides at each conferee's end system to coordinate the initiation, maintenance and interaction of sessions. A communication endpoint for the session control protocol.

Media Agent

- A software entity that handles media-specific functions such as encoding, compression and transport packetization that are used by conferences. Media in a conference might include audio, video, graphics and text.

Resource Manager

- A software entity that manages the media agents on a workstation. It understands the static end-system descriptions (hardware and software capabilities), as well as the negotiated per-session preferences.

Conference Directory Service

- A directory that provides user network addresses, conference IDs and addresses, the conference begin time, conference topic, etc.

Conference Scheduling

- The advertisement of a conference's start time with a session directory service.

6

Terminology—policies

Initiator

- A conferee that initiates a conference.

Chair

- A designated conferee having more authority than other conferees in the conference. For example, the chair might decide the policy on late joins, media floor control, interaction style, etc.

Receiver

- A conferee that receives session data.

Sender

- A conferee that transmits/sends session data.

Passive Participant

- A conferee that only acts as a receiver.

Active Participant

- A conferee that acts both as a sender and receiver.

Reflector Participant

- An entity that relays data between conferees, acting as a go-between.

7

Terminology—policies

Floor Control

- Coordinated control over who may or may not send and/or receive data.

Connectivity Style

- The interconnectivity of conferees (e.g., 1-to-N, N-to-N, M-to-N) in either the control or data realm.

Access Control

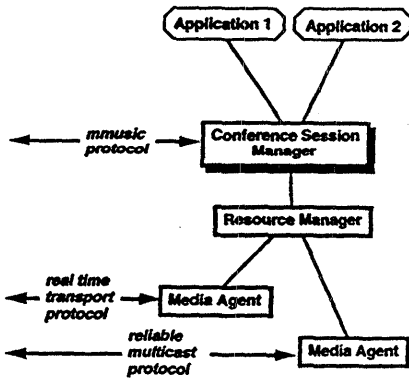
- The accessibility of a session to potential conferees.

Interaction Policies

- The model and rules used by conferees to interact with one another in a conference.

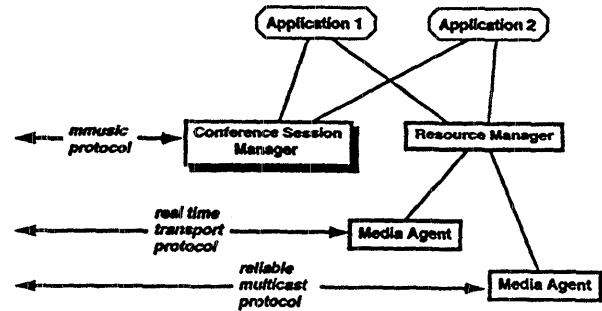
8

Framework 1



9

Framework 2



10

Session Control Protocol Requirements

Distributed session management

- establish, terminate, merge
 - once established, a session has a unique name (for example, to advertise the session in a directory service)
- schedule—negotiate the session in advance

Dynamic membership management

- add member, delete member

Session policy management

- set, modify, remove
 - permissions on: who may change policies, who may terminate, who may join, who may invite, etc.

Domain specific (a bag of bits)

- Specification of media associations and configurations
- Security: encryption key distribution
- Domain-specific policies: floor control, sender vs. receiver, etc.

11

Implementation for the Internet

Internet realities

- packet loss
- delay variations
- ever growing

Communication building blocks

- UDP vs. TCP
- unreliable vs. reliable "datagrams"
- point-to-point vs. multipoint
 - CCP reliable multipoint
- causal/atomic multicast
 - ISIS group services

Active vs. passive

- tightly coordinated request/replies
- individual status announcements

Adaptivity to delay variations

- adapt peer-to-peer timeouts to dynamic network conditions

RTCP considerations

12

More Basic Message Types

Status

- sent by a session manager to share the state of the session
- no reply is required

Query

- sent by a session manager to request the session state from another session manager
- the receiving session manager responds with an Announce message

19

Session Policies

Who may terminate a session *without* negotiation

- ALL, designated subset of members, NONE
 - ALL is all members
 - NONE means that a session only terminates when all members leave

Who may initiate a change to the session

- ALL, designated subset of members, NONE
 - NONE means that the session is "static"

Who must approve a change to the session (depending on the initiator of the proposal?)

- ALL, designated subset of members, any n of the members, NONE
- may be more complex, e.g., n of a designated subset, etc.
- at session establishment, this determines who must agree to participate


Should the latter two policies be specified for each type of operations, or apply to all possible changes?

Who may be told about session state

- ANYONE, ALL, designated set of users
 - ANYONE is all users


20

Session Establishment



21

Session Change



22

Outstanding Issues

Mobility at session establishment and during the session

Distributed system issues

- keep alive messages using Announce
- state synchronization
- join by contacting a *doorman* or *doormen*
- replace terminate session by using delete all members
- error handling
- serializability conflicts and fault tolerance

Transport Infrastructure

- UDP, TCP, Sun RPC, VMTP, TP4, ISIS, etc.?

Schedule and preschedule

Floor control

Richness of policies

Other uses?

23

Action Items

Documents

- requirements
 - framework
 - terminology
 - conversation styles
 - protocol requirements
- specification

24

2.8.3 TCP Large Windows (TCPLW)

Charter

Chair(s):

David Borman, dab@cray.com

Mailing Lists:

General Discussion: tcplw@cray.com

To Subscribe: tcplw-request@cray.com

Archive:

Description of Working Group:

The TCP Large Windows Working Group is chartered to produce a specification for the use of TCP on high delay, high bandwidth paths. To this end, this working group recommended RFC 1072 "TCP extensions for long-delay paths" and RFC 1185 "TCP Extension for High-Speed Paths" be published jointly as a Proposed Standard. Deficiencies in the technical details of the documents were identified by the End-to-End Research Group of the IRTF. Rather than progress the standard with known deficiencies, the IESG tasked the End-to-End Research Group to fix and merge these two documents into a single protocol specification document. This review was done on the e2e-interest@isi.edu mailing list.

The TCP Large Windows Working Group is being resurrected for a one time meeting, to review and if appropriate, approve this new document.

Goals and Milestones:

Done Review the TCP Extended Window Size proposal from the IRSG End to End Research Group and if acceptable, recommend it for standards status.

Internet-Drafts:

"TCP Extensions for High Performance: An Update", 06/23/1993, R. Braden
<draft-ietf-tcplw-extensions-00.txt>

Request For Comments:

RFC 1323 "TCP Extensions for High Performance"

CURRENT MEETING REPORT

Reported by David Borman/Cray Research

Minutes of the TCP Large Windows Working Group (TCPLW)

The TCP Large Windows Working Group met on Wednesday, July 14 at the 27th IETF meeting in Amsterdam. The agenda was:

- Review of draft-ietf-tcplw-extensions-00.txt
- Consideration of advancing RFC 1323 to Draft Standard
- Status of SACK option

The document draft-ietf-tcplw-extensions-00.txt is a compilation of bugs and clarification that need to be made to RFC 1323. It was compiled by Bob Braden. Bob and David Borman led a walk through of the document to help explain it and to find out if any other changes that are needed to RFC 1323 were missed.

RTTM: Relationship to Karn's Algorithm

One of the items that Karn's algorithm addresses is how to get valid RTT values in the presence of lost data. With current methods, the answer is that whenever data has to be retransmitted, no valid RTT estimate can be made in that window. The reason for this is that when the ACK is received, it cannot be determined whether that ACK came from the original packet or the retransmitted packet.

With the use of the timestamps option, this ambiguity is removed, since the timestamp echoed in the ACK will always be from the data packet that caused the ACK of new data to be generated.

The other part of Karn's Algorithm, that of doing exponential back-off on retransmissions, still needs to be done.

RTTM: Which TS to Echo

The discussion that took place during the session was a bit muddled. This summary provides a clearer explanation.

This is the one real bug in RFC 1323. As stated currently, the TSval is only copied to TS.Recent when a packet is received that advances the left edge of the window. The change allows the TSval to be copied when the left edge of the window is advanced, or if the packet is before the left edge of the window and it has a newer timestamp.

The scenario that this addresses is when a packet is sent, its ACK is lost, and so the packet is then resent. The second packet will not advance the left edge of the window, but will cause a duplicate ACK to be generated. To the sender, this will be the first ACK that is received, so it will want to use the timestamp in that packet to update the RTT estimate. If the TSval is not copied from the resent packet, the the sender will get an inflated RTT estimate.

By making this change, with the use of timestamps for doing RTT measurements, it can be guaranteed that in the worst case, at least one RTT measurement can be made per window. Without the timestamps option, in many TCP implementations one RTT estimate per window is the best case.

Additionally, this change fixes a problem with a long-lived unidirectional connection. Since all the ACKs will have the same sequence number, they would never cause TS.Recent to be updated. With this change, the string of ACK-only packets will cause TS.Recent to be kept up-to-date, and allow PAWS to work properly.

TCP Options and MSS

The discussion of how the TCP MSS option interacts with the presence of TCP and/or IP options needed some clarification. Since many TCP implementations use the MSS option to indicate how large of a packet can be received without fragmentation, and the MSS does not include the TCP/IP headers, the question is: should it also be adjusted for the length of the options that will probably be in each packet?

The answer is no, and when sending data the sender should always assume that the MSS received did not account for the TCP and IP options, so the MSS should be reduced by the length of the TCP and IP options when determining how much data can be sent. The following grid shows why:

| | MSS is adjusted to include options | MSS is not adjusted to include options |
|--|---------------------------------------|---|
| Sender adjusts length for options | Packets are too short | Packets are the correct length |
| Sender doesn't adjust length for options | Packets are the correct length | Packets are too long |

The goal is to not send IP datagrams that have to be fragmented by IP. Packets sent with the constraints in the lower right of this grid will cause IP fragmentation. The only way to

ensure that this doesn't happen is for the data sender to decrease the MSS by the length of the IP and TCP options.

Modification to TCP Event Processing Rules

The draft document contains a new set of rules for how to process the TCP extensions. Bob wanted feedback on whether or not the format was easy to read and understand. The general feeling was that it was easy to read and understand.

SACK

There was some discussion of the status of the SACK option. Right now, there is not a lot of visible activity in this area. There are some test implementations of the SACK option as originally defined in RFC 1072. The plan is once there is some hard data on how well it works, then a new draft document will be generated on the SACK option.

Action Items

- Bob Braden will update RFC 1323 with the changes that are described in draft-ietf-tcplw-extensions-00.txt. The new document will be sent to the mailing list for final review, and then sent to the IESG for consideration for advancement to Draft Standard status.
- Bob Braden will update draft-ietf-tcplw-extensions-00.txt to include an updated description of which TS to echo to explain the bug as described above, and then it will be published as an Informational RFC.
- David Borman will send out the current list of implementations, so that people can send in updates. The updated list will be given to the IESG at the same time as the updated version of RFC 1323.

Attendees

| | |
|----------------|-------------------------------|
| David Borman | dab@cray.com |
| Jim Bound | bound@zk3.dec.com |
| Robert Braden | braden@isi.edu |
| Walid Dabbous | Walid.Dabbous@sophia.inria.fr |
| Jeffrey Dunn | dunn@neptune.nrl.navy.mil |
| Francis Dupont | francis.dupont@inria.fr |
| Robert Enger | enger@reston.ans.net |

| | |
|-----------------------|---------------------------|
| Joseph Godsil | jgodsil@ncsa.uiuc.edu |
| Frank Hoffmann | hoffmann@dhdibm1.bitnet |
| Phil Irey | pirey@relay.nswc.navy.mil |
| Ronald Jacoby | rj@sgi.com |
| Paolo Malara | malara@crs4.it |
| Kent Malave | kent@bach.austin.ibm.com |
| Allison Mankin | mankin@cmf.nrl.navy.mil |
| Willi Porten | porten@gmd.de |
| Henry Sanders | henrysa@microsoft.com |
| Tim Seaver | tas@concert.net |
| Keith Sklower | sklower@cs.berkeley.edu |
| Robert Ullmann | ariel@world.std.com |
| Willem van der Scheun | scheun@sara.nl |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

2.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reynolds/Information Sciences Institute

Ten working groups in the User Services Area of the IETF met in Amsterdam.

Integrated Directory Services Working Group (IDS)

The IDS Working Group is chartered to facilitate the integration and interoperability of current and future directory services into a unified directory service. This work will unite directory services based on a heterogeneous set of directory service protocols (X.500, Whois++, etc.). In addition to specifying technical requirements for the integration, the IDS Working Group will also contribute to the administrative and maintenance issues of directory service offerings by publishing guidelines on directory data integrity, maintenance, security, and privacy and legal issues for users and administrators of directories.

The "Advanced Usages of X.500" document from the IDS Working Group has been sent to the RFC Editor for publication. The "Pilot Project" catalog is out as an Internet-Draft. The revision of FYI 11 has been released as an Internet-Draft and will be submitted to the RFC Editor for publication by the end of August 1993. A draft on "Legal & Privacy Issues in Directory Services" was submitted and will be published as an Internet-Draft. A draft on "A Guide to Available Directory Services" was circulated and will be advanced as an Internet-Draft for eventual publication as an Informational RFC. The "Directory Services Policy Handbook" will be split into pieces and will be released as a series of documents.

Integration of Internet Information Resources Working Group (IIIR)

IIIR is chartered to facilitate interoperability between Internet information services, and to develop, specify, and align protocols designed to integrate the plethora of Internet information services (WAIS, Archie, Prospero, etc.) into a single "virtually unified information service."

IIIR met twice at this IETF. At the first session, the HTML Internet-Draft from Tim Berners-Lee was discussed, and will be submitted as an Informational RFC. The HTTP protocol will also be released as an Internet-Draft before the next IETF. HTML+ may be brought in as a potential standard. At the second session, the "Vision of an Integrated Internet Information Architecture" document was discussed. The document is still an Internet-Draft, and comments have come in that need to be incorporated. The "Resource

Transponders” document was also discussed. The potential overlap with the SVRLOC Working Group’s work has still not been resolved.

The “Taxonomy of Information Services” document has been shelved as other papers have recently been published on this work. There are some minor problems with the Gopher RFC which will be addressed in an addendum. The Prospero protocol was suggested as a means of integrating information services. Discussion on this item will be taken to the IIIR Working Group e-mail list. The Gopher/IIIR data type specifications will be registered as MIME types. A document will be written on failure tracking and resolution between information services. “Z39.50 over TCP/IP” will be released as an Internet-Draft, to eventually become an Informational RFC. A paper on integrating data elements between the IAFA Working Group’s paper and the library community will be released as an Internet-Draft.

Internet School Networking Working Group (ISN)

The Internet School Networking Working Group is chartered to facilitate the connection of the United States’ K-12 (Kindergarten-12th grade) schools, public and private, to the Internet, and school networking in general.

After some background on the history of ISN, the group shared experiences in school networking and their own companies and projects. There were 30 or 35 people in attendance and almost all had more than a passing interest in school networking. Next, the group gave a lot of good feedback on the document currently under construction, the Internet-Draft on FAQs for the primary and secondary school community. Also discussed was a document on connectivity models for schools. It was decided that this document will be combined with the USERDOC2’s document on how to connect to the Internet, and will be put out as an Internet-Draft. The final agenda item was a review and revision of the group’s charter. After some discussion, three volunteers from three different countries agreed to draft the text for a new charter and post it to the list. From there, the group will determine new deliverables and milestones.

Network Information Services Infrastructure Working Group (NISI)

NISI is exploring the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services “infrastructure” for the Internet.

NISI discussed two main topics. First, comments on the current Internet-Draft were solicited and a discussion of the relationships between network information centers (NICs) ensued, resulting in suggested revisions to the document’s diagram. Secondly, the scope of recommendations for interactions between NICs when more than one are involved in helping a user were clarified. The group benefited greatly from the international input available in this forum, and both documents currently in progress should be much improved.

Network Training Materials Working Group (TRAINMAT)

The Network Training Materials Working Group is chartered to enable the research community to make better use of the networked services. Towards this end, the working group will work to provide a comprehensive package of "mix and match" training materials for the broad academic community which will: 1) enable user support staff to train users to use the networked services and 2) provide users with self-paced learning material. In the first instance, it will not deal with operational training. This working group is the IETF component of a joint RARE/IETF group working on network training materials.

The meeting opened with introductions and descriptions of training activities in which each person is involved. Jill Foster discussed the RARE ISUS Network Training Materials Task Force as well as the NISP/ITTI project at the University of Newcastle, UK. The group reviewed data elements for the training materials catalog, and agreed on the general categories which Jill presented with a few suggested changes. A small group will finalize the template and send it out to the list. Volunteers were recruited to work on the catalog. Another project the group will pursue is a subject resource guide. Issues in using the network to deliver interactive and multimedia training were discussed, with the possibility for developing multimedia pilots reviewed. Jill gave a demonstration of the materials developed at Newcastle.

Networked Information Retrieval Working Group (NIR)

NIR is chartered to increase the useful base of information about networked information retrieval tools, their developers, interested organizations, and other activities that relate to the production, dissemination, and support of NIR tools. NIR is a cooperative effort of the IETF, RARE, and CNI.

Jim Fullton gave an overview of CNIDR. Jill Foster gave a brief overview of the RARE ISUS work in the NIR area. Anders Gillner discussed his work on the Eurogher project. This project involves establishing a subject-based, as well as geographic-based, Gopher infrastructure in Europe. The group then spent most of the meeting editing the NIR report. Major changes included the merging of the WAIS and freeWAIS sections, NCSA's Mosaic was moved under the WWW section, and the NLM sections were removed. Various methods of publishing and maintaining this information on-line were discussed. Several suggestions for evaluating NIR tools was discussed. It was decided to put together a simple checklist as the basis of the evaluation. There will be further discussion on the mailing list. The charter was updated to remove the reference to documentation and training materials.

Uniform Resource Identifiers Working Group (URI)

URI is chartered to define a set of standards for the encoding of system independent resource location and identification information for the use of Internet information services.

The URI Working Group met in three sessions to discuss various issues related to URLs, URNs and URCs. The current URL proposal, along with an amendment to expand the allowable character set, was approved by a vote of twenty-nine to one. A report on the Info Mesh was made by Karen Sollins.

URNs were discussed in great detail. Agreement was reached on the basic format of URNs as well as their content. Sub-elements of the URN remain undefined and will be discussed at the next IETF.

A paper describing a proposed URM was discussed, as were proposals for the format of URCs. A call for papers describing URMs and URCs was made, with discussion to be carried out on the list. Next time: 1) discussion of URN sub-fields, and 2) discussion of URC/URM formats.

User Documents Revisions Working Group (USERDOC2)

The User Documents Revisions Working Group is preparing a revised bibliography of on-line and hard copy documents, reference materials, and training tools addressing general networking information and how to use the Internet. The target audience includes those individuals who provide services to end users, and end users themselves.

The USERDOC2 Working Group reviewed recent activities, including the release of FYI RFCs 19 and 20, since the last IETF meeting. The archive, "Introducing the Internet," has been updated, adding two additional documents and revising the access guide. Four sites currently house copies of the archive, and others are encouraged to maintain it. A discussion on documentation needs was held, starting with a talk by Bert Stals on activities by the ISUS Documentation group. Issues relating to documentation in languages other than English and covering a more world-wide perspective in RFCs was covered. The final topic was new activities, with two future FYI RFCs in the works. These are a revision of FYI 3 which is to be finalized at the next meeting, and a new document on how to get connected to the Internet which is just being started.

User Services Working Group (USWG)

USWG provides a regular forum for people interested in all user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet.

Joyce Reynolds reported on the IETF User Services Area activities including: working groups coming to closure, new working groups starting up, new publications, and current user services related Internet-Draft postings.

April Marine led a discussion on an FYI RFC update: FYI 4 "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions," (also RFC 1325), May 1992. Bill Manning brought up a discussion on how to "empower" users to utilize and document tools. There also seems to be missing from the FYI RFC publications a series of notes that tries to address the manners and morals of the collective body. Another topic that may need to be worked on in this forum is how to deal with the basics (e.g., how do "I" get attached?).

Jill Foster presented an update on RARE activities, including a report on the RARE Information Services/User Services (ISUS) activities.

David Sitman, EARN (European Academic and Research Network) representative, presented a discussion and review of EARN's "Guide to Network Resource Tools" in preparation for submission to the RFC Editor for FYI RFC publication.

Whois and Network Information Lookup Service Working Group (WNILS)

The purpose of WNILS is to expand and define the standard for Whois services, to resolve issues associated with the variations in access, and to promote a consistent and predictable service across the network.

Many attendees were new to the WNILS Working Group and requested an introduction to the Whois++ architecture. Peter Deutsch presented an overview of the Whois++ architecture and design philosophy. A lengthy discussion followed and additions to the protocol were requested. Peter will incorporate the additional specifications and submit an Internet-Draft by July 30.

Chris Weider also began his presentation with an overview of centroids. Two problem areas were identified: a) centroids don't scale for databases containing large quantities of unique data, and b) the potential exists for cycles or looping in queries because the directory is in a mesh configuration. Chris agreed further work is needed on indexing but this first implementation should provide a functional system for existing datasets and a testbed for the limitations of this model.

The recommended modifications to the Whois protocol have been submitted as an Internet-Draft. The host record section will be updated. The goals and milestones will be updated to reflect the following: 1) all draft papers will be submitted as Internet-Drafts by the end of July, and 2) two working implementations of clients, servers and centroids will be available by September 30.

2.9.1 Integrated Directory Services (IDS)

Charter

Chair(s):

Chris Weider, clw@merit.edu
 Tim Howes, tim@umich.edu

Mailing Lists:

General Discussion: ids@merit.edu
 To Subscribe: ids-request@merit.edu
 Archive: merit.edu:~/pub/ids-archive

Description of Working Group:

The Integrated Directory Services Working Group is chartered to facilitate the integration and interoperability of current and future directory services into a unified directory service. This work will unite directory services based on a heterogeneous set of directory services protocols (X.500, WHOIS++, etc.). In addition to specifying technical requirements for the integration, the IDS Working Group will also contribute to the administrative and maintenance issues of directory service offerings by publishing guidelines on directory data integrity, maintenance, security, and privacy and legal issues for users and administrators of directories.

IDS will also assume responsibility for the completion of the outstanding Directory Information Services Infrastructure (DISI) Internet-Drafts, which are all specific to X.500, and for the maintenance of FYI 11, "A catalog of available X.500 implementations".

IDS will need to liase with the groups working on development and deployment of the various directory service protocols.

The IDS Working Group is a combined effort of the Applications Area and the User Services Area of the IETF.

Goals and Milestones:

- | | |
|---------|--|
| Ongoing | Track emerging directory service protocols to specify standards for interoperation with existing protocols. |
| Ongoing | Liase with groups working on deployment and development of directory services to locate and fix interoperability problems. |
| Ongoing | Identify unfilled needs of directory service offerers, administrators, and users. |
| Done | Submit to the IESG the DISI "Advanced Usages of X.500" paper as an informational document. |

- Jun 1993 Submit to the IESG the 1993 revision of FYI 11, "A catalog of available X.500 implementations" as an informational document.
- Jul 1993 Submit as an Internet-Draft a "Specifications for interoperability between WHOIS++ and X.500".
- Jul 1993 Submit as an Internet-Draft a "Guide to administering a directory service", which covers data integrity, maintenance, privacy and legal issues, and security.
- Jul 1993 Submit as an Internet-Draft a "Catalog of available WHOIS++ implementations".
- Done Post the "X.500 Pilot Project Catalog" paper as an Internet-Draft.
- Nov 1993 Submit to the IESG the DISI "X.500 Pilot Project Catalog" paper as an informational document.
- Nov 1993 Submit to the IESG the "Specifications for interoperability between WHOIS++ and X.500" as a standards document.
- Nov 1993 Submit as an Internet-Draft a "User's guide to directory services on the Internet".
- Mar 1994 Submit to the IESG the "Guide to administering a directory service" as an informational document.
- Mar 1994 Submit to the IESG the 1994 revision of FYI 11.
- Mar 1994 Submit to the IESG the "Catalog of available WHOIS++ implementations" as an informational document.

Internet-Drafts:

"X.500 Pilot Projects", 06/15/1993, A. Marine <draft-ietf-ids-pilots-00.txt>

Request For Comments:

RFC 1491 "A Survey of Advanced Usages of X.500"

CURRENT MEETING REPORT

Reported by Chris Weider/Merit Network, Inc.

Minutes of the Integrated Directory Services Working Group (IDS)

Review of Previous Minutes

There were no modifications to previous minutes.

Liaison Reports

- There was no NADF liaison report.
- WG-NAP has an X.500 focus, and is working on data management issues, privacy and legal issues, and problems with distributed entry and multiple providers.
- Thomas Lenggenhager has been working on a Data Management for White Pages paper for organizations wishing to join the X.500 directory.
- OSI-DS has released RFCs 1274 and 1279 as its core work, and is continuing work on LDAP, IP and network management in the X.500 directory, and will be guiding new schema work in a new schema working group. They will also be working with the 1993 version of X.500.

Status Reports

- Whois++
Four independent groups are working on servers. There is already a Whois++ front-end to X.500 developed by Mark Prior of the University of Adelaide.
- LDAP
The LDAP Internet-Draft is moving on to Proposed Standard status, and the ISODE Consortium will be publishing the LDAP API.

InterNIC Directory Services

Sri Sataluri gave a talk on the directory services activity of the InterNIC. The InterNIC is running a number of directories on their X.500 service, and is also running NetFind and Whois.

X.500 Index DSA Work

Paul Barker presented his work on building special purpose DSAs to help alleviate the problems encountered when trying to search across multiple branches of the DIT. His technique is to replicate entries appropriate to a given search into a special purpose DSA, and restricting the search to just that DSA. His experimental DSA holds all UK computer science people.

Progress of Assigned Documents

- FYI 11: The X.500 Implementation Catalog

This paper is nearly finished and should be out as an Internet-Draft by the end of August 1993.

- Pilot Project Catalog

Although there was some discussion as to the necessity of this document, it was finally decided that it should go ahead and will be released as an Internet-Draft soon.

- Advanced Usages of X.500

This paper was released as RFC 1491 and FYI 21 in late July.

- Whois++ / X.500 Interoperability

This paper will be released as an Internet-Draft before Houston.

- Directory Service Policy Handbook

It was decided that rather than try to write one all-encompassing paper on this subject, that the topics covered in the Policy Handbook outline should be split up across a number of papers. Thomas Lenggenhager's WG-NAP paper on data management will be released as one of these papers; Erik Huizer's upcoming paper on legal issues and privacy will be released through the IETF process as another of these papers.

Discussion of Directory Services and Privacy Issues

Erik Huizer gave a talk on the results of the research he has done in the legal status of directory services. The US has essentially no relevant laws; the ones it does have are designed to protect the providers of the service rather than the people whose information is being held. In the European Community (EC), the situation is a patchwork; each country has different laws, with the strictest being Germany. The EC has developed a directory services/database policy which must be placed into law in all EC countries by 1995; this policy is stricter than most of the current national laws but weaker than the strongest laws.

In any case, export of personal data to countries with no privacy laws is strictly prohibited (for example, Germany has already stated that no data can be exported to the US). The laws seem perfect for centralized databases but place prohibitive restrictions on distributed global databases such as X.500. Erik's paper will be out by the end of the year.

New Document Discussion

- Whois++ Implementation Catalog

This will be released as an Internet-Draft before Houston, written by Chris Weider.

- Updates to FYI 13 and 14

These papers will be reviewed for any potential update by Chris Weider, and will be modified before Houston if required.

Other Business

Marco Hernandez has a draft of a Guide to Directory Services document which covers all the major directory service paradigms. This would be an ideal addition to the FYI series, and will be released as an Internet-Draft as soon as Marco feels that it is complete.

Attendees

| | |
|-----------------------|--------------------------------|
| Paul Barker | p.barker@cs.ucl.ac.uk |
| Luc Boulianne | lucb@cs.mcgill.ca |
| Piet Bovenga | p.bovenga@uci.kun.nl |
| James Conklin | jbc@bitnic.educom.edu |
| Robert Cooney | cooney@wnyose.nctsw.navy.mil |
| Peter Deutsch | peterd@bunyip.com |
| Thomas DeWitt | tdewitt@osi.ncsl.nist.gov |
| Maria Dimou-Zacharova | dimou@dxcern.cern.ch |
| Terry Gray | gray@cac.washington.edu |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| J. Jensen | jensen@ic.dk |
| Thomas Johannsen | thomas@ebzaw1.et.tu-dresden.de |
| Peter Jurg | jurg@surfnet.nl |
| Anders Karlsson | sak@cdg.chalmers.se |
| Michael Khalandovsky | mlk@ftp.com |
| Steve Kille | S.Kille@isode.com |
| Bruno Koechlin | |
| Mark Kosters | markk@internic.net |

| | |
|--------------------------|-----------------------------------|
| Arnold Krechel | krechel@gmd.de |
| John Kunze | jak@violet.berkeley.edu |
| Sylvain Langlois | Sylvain.Langlois@exp.edf.fr |
| Erik Lawaetz | erik.lawaetz@uni-c.dk |
| Eliot Lear | lear@sgi.com |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Steen Linden | steen.linden@uni-c.dk |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| John Martin | John.Martin@newcastle.ac.uk |
| Linda Millington | l.millington@noc.ulcc.ac.uk |
| Keith Moore | moore@cs.utk.edu |
| Clifford Neuman | bcn@isi.edu |
| Paul-Andre Pays | pays@faugeres.inria.fr |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Catherine Pierre-Radenac | catherine.pierre-radenac@inria.fr |
| Mel Pleasant | pleasant@hardees.rutgers.edu |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Kenneth Rossen | kenr@shl.com |
| Srinivas Sataluri | sri@internic.net |
| Rickard Schoultz | schoultz@admin.kth.se |
| Karen Sollins | sollins@lcs.mit.edu |
| Simon Spero | simon_spero@unc.edu |
| Catherine Treca | Catherine.Treca@dione.urec.fr |
| Ton Verschuren | Ton.Verschuren@SURFnet.nl |
| Paul Vetter | |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Chris Weider | clw@merit.edu |
| Russ Wright | wright@lbl.gov |
| Steve Zeber | zeber@stc.nato.int |

2.9.2 Integration of Internet Information Resources (IIIR)

Charter

Chair(s):

Chris Weider, clw@merit.edu
 Kevin gamiel, kgamiel@cnidr.org

Mailing Lists:

General Discussion: iiir@merit.edu
 To Subscribe: iiir-request@merit.edu
 Archive: [merit.edu:~/pub/iiir-archive](http://merit.edu/~pub/iiir-archive)

Description of Working Group:

The Integration of Internet Information Resources Working Group (IIIR) is chartered to facilitate interoperability between Internet information services, and to develop, specify, and align protocols designed to integrate the plethora of Internet information services (WAIS, ARCHIE, Prospero, etc.) into a single "virtually unified information service" (VUIS). Such protocols would include, but are not limited to, update protocols for distributed servers, a "query routing protocol" to pass queries between existing services, protocols for gateways between existing and future services, and standard exchange formats (perhaps based on Z39.50) for cross-listing specific information.

Also, where necessary, IIIR will create technical documentation for protocols used for information services in the Internet.

Goals and Milestones:

- | | |
|----------|--|
| Ongoing | Track emerging Internet information services in order to specify technical requirements for their integration into the VUIS. |
| Ongoing | Liaise with other groups working on deployment and integration of Internet information services: e.g., The Coalition for Networked Information, RARE Working Group 3, etc. |
| Ongoing | Create specifications for interoperability between Internet information systems. |
| Done | Post an Internet-Draft on 'A vision of integrated information resources.' |
| Jul 1993 | Post an Internet-Draft on 'Taxonomy of Internet Information Services.' |
| Jul 1993 | Submit final version of 'A vision of integrated information resources' to the IESG as an Informational RFC. |
| Jul 1993 | Submit final version of 'Taxonomy of Internet Information Services' to the IESG as an Informational RFC. |

- Nov 1993 Post an Internet-Draft defining common exchange formats.
- Nov 1993 Post an Internet-Draft defining a Query Routing Protocol.
- Mar 1994 Submit final version of common exchange format to the IESG as a Proposed Standard.
- Jul 1994 Submit final version of Query Routing Protocol to the IESG as a Proposed Standard.

Internet-Drafts:

“Resource Transponders”, 03/22/1993, C. Weider <draft-ietf-iiir-transponders-00.txt>

“A Vision of an Integrated Internet Information Service”, 03/26/1993, C. Weider, P. Deutsch <draft-ietf-iiir-vision-00.txt>

“Hypertext Markup Language (HTML): A Representation of Textual Information and MetaInformation for Retrieval and Interchange”, 06/23/1993, T. Berners-Lee, D. Connolly <draft-ietf-iiir-html-01.txt, .ps>

CURRENT MEETING REPORT

Reported by Chris Weider/Merit

Minutes of the Integration of Internet Information Resources Working Group (IIIR)

Acknowledgements

Many thanks to Kevin Gamiel for taking notes for these minutes during the second session.

Prologue

IIIR was originally scheduled for only one session, on the evening of the 14th. However, Tim Berners-Lee, author of the HTML Internet-Draft, would have been unable to attend, so an additional session was added on the 13th to discuss the document.

HTML Internet-Draft

The HTML Internet-Draft describes current practice for the Hypertext Markup Language used in the World Wide Web. As such, it will be submitted for publication as an Informational RFC as soon as possible. There are two other protocols used in the Web which will be brought into the Internet standards track: HTML-2, and the Hypertext Transport Protocol (HTTP).

Previous Minutes

No changes were suggested.

Vision Document

Chris Weider gave a brief overview of the vision paper. Jill Foster wanted an additional layer of integration on top of the architecture by having a single client that speaks all the information service protocols. Mitra indicated a need for a set of shared, common libraries for each server to use for URN → URL access. These suggestions will be incorporated into a future version of the paper.

Transponder Document

Chris Weider gave an overview of the transponder document. The final consensus of the working group was to consider the transponder document as a “call to arms” to start thinking about ways to make maintenance of the URN → URL infrastructure automatic.

Taxonomy Document

Chris Weider and Peter Deutsch were not able to complete the taxonomy document. However, Cliff Neuman pointed out that there were several recently published papers that cover most of the ground of the taxonomy document, and that they would serve as well. Cliff will post references to those papers to the mailing list and the taxonomy document will be removed from the list of goals and milestones.

Gopher

There were still concerns that RFC 1436 documenting the base Gopher protocol did not adequately reflect current operational practice. Chris Weider took an action item to talk to the authors and resolve the issue by the end of August.

Prospero

Cliff Neuman gave a talk on Prospero and outlined how it might be used to integrate disparate information services on the Internet.

IIIR Format Types

Mitra stated that there needs to be a set of data/format type specifiers for resources that is consistent across all the NIDR systems. Historically each system has created its own format specifiers, which impedes inter-operability. The MIME content type specification process is now “open” for additions, so a consistent set can now be registered. Mark McCahill has agreed to register the Gopher types in MIME, and that will serve as a basis for the IIIR set. An earlier BOF on datatypes resulted in the formation of a mailing list designed to discuss this issue; the list is `ietf-types@cs.utk.edu`. At this writing the activation of the list has not yet been announced.

Data Elements

Much concern was expressed by members of the working group that a paper on data elements for metainformation transfer had not been released. A small working group composed of librarians and IETFers are working on it, but there was also concern that the process of creating the document was not open. Peter Deutsch noted that in many cases current practice is being accepted over architecting a working solution. Chris Weider agreed that a draft of the paper would be released to the working group as soon as it was available.

Quality Assurance

Mitra gave a brief talk on his concerns about the quality (or lack thereof) of pointer maintenance, and noted that if a given retrieval request fails, that it is impossible to determine where it has failed. Chris Weider took an action item to submit a request for volunteers to write a paper on technical quality assurance.

Z39.50

Jim Fullton mentioned that a paper by Cliff Lynch on Z39.50 over TCP/IP will be submitted for consideration as an Informational RFC in time for Houston. Brewster Kahle, John Kunze, and Jim Fullton agreed to write a paper on the use of Z39.50 in information systems.

Gopher+

Time expired before Gopher+ could be discussed, but will be taken to the list. Of particular concern is that the Gopher+ extensions allow interoperability with other information systems deployed on the Internet.

Attendees

| | |
|----------------|-----------------------------|
| Luc Boulianne | lucb@cs.mcgill.ca |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| James Conklin | jbc@bitnic.educom.edu |
| Peter Deutsch | peterd@bunyip.com |
| Reinhard Doelz | embnet@comp.bioz.unibas.ch |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Jim Fullton | fullton@cnidr.org |
| Kevin Gamiel | kgamiel@cnidr.org |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Anders Gillner | awg@sunet.se |

| | |
|----------------------|--------------------------------|
| Rune Hjelsvold | Rune.Hjelsvold@idt.unit.no |
| Xander Jansen | xander.jansen@surfnet.nl |
| Olle Jarnefors | ojarnef@admin.kth.se |
| Michael Khalandovsky | mlk@ftp.com |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| April Marine | april@atlas.arc.nasa.gov |
| John Martin | John.Martin@newcastle.ac.uk |
| Mitra | mitra@pandora.sf.ca.us |
| Clifford Neuman | bcn@isi.edu |
| Lisa Nielsen | lisa@hawaii.edu |
| Lars-Gunnar Olsson | Lars-Gunnar.Olsson@data.slu.se |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Rickard Schoultz | schoultz@admin.kth.se |
| Patricia Smith | psmith@merit.edu |
| Karen Sollins | sollins@lcs.mit.edu |
| Vladimir Sukonnik | sukonnik@process.com |
| Peter Svanberg | psv@nada.kth.se |
| Guido van Rossum | guido@cw.nl |
| Ton Verschuren | Ton.Verschuren@SURFnet.nl |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Chris Weider | clw@merit.edu |

2.9.3 Internet Anonymous FTP Archives (IAFA)

Charter

Chair(s):

Peter Deutsch, peterd@bunyip.com
Alan Emtage, bajan@bunyip.com

Mailing Lists:

General Discussion: iaafa@cc.mcgill.ca
To Subscribe: iaafa-request@cc.mcgill.ca
Archive: archive.cc.mcgill.ca:~/pub/iaafa-archive

Description of Working Group:

The Internet Anonymous FTP Archives Working Group is chartered to define a set of recommended standard procedures for the access and administration of anonymous FTP archive sites on the Internet. Such a set of procedures will provide a framework for:

- (a) Allowing the inexperienced Internet user the ability to more easily navigate the hundreds of publically accessible archive sites.
- (b) Allowing users and network-based tools to retrieve specific site information such as access policies, contact information, possible areas of information specialization, archived package descriptions, etc., in a standardized manner.

Particular emphasis will be placed on the possible impact of these procedures on the FTP site administrators.

Attention will be paid to the impact of newer archive indexing and access tools on the operation of such archive sites. A set of suggestions will be offered to allow archive site administrators to better integrate their offerings with such tools as they are developed.

The security of the anonymous FTP site configuration will also be considered to be an integral part of this document. It is expected that remote management of the archives will be adequately handled by existing network management procedures.

Goals and Milestones:

- | | |
|----------|---|
| Done | First IETF Meeting: review and approve the Charter making any changes deemed necessary. Examine the scope of the recommended procedures and impact on site administrators. Assign writing assignments for the first draft of the documents. |
| Mar 1992 | Review first draft and determine necessary revisions. Follow up discussion will occur on mailing list. |

- Jun 1992 Make document an Internet-Draft. Continue revisions based on comments at IETF and on the mailing list.
- Nov 1992 Fourth IETF meeting. Review final drafts and if OK, give to IESG for publication as an RFC.

Internet-Drafts:

“How to Use Anonymous FTP”, 06/15/1993, P. Deutsch, A. Emtage, A. Marine
<draft-ietf-iafa-howftp-00.txt>

“Publishing Information on the Internet with Anonymous FTP”, 08/17/1993,
P. Deutsch, A. Emtage <draft-ietf-iafa-publishing-00.txt>

“Data Element Templates for Internet Information Objects”, 08/17/1993, P.
Deutsch, A. Emtage <draft-ietf-iafa-templates-00.txt>

2.9.4 Internet School Networking (ISN)

Charter

Chair(s):

Jennifer Sellers, sellers@nsipo.nasa.gov
 Arthur St. George, stgeorge@bootes.unm.edu

Mailing Lists:

General Discussion: isn-wg@unmvma.unm.edu
 To Subscribe: listserv@unmvma.unm.edu
 In Body: `subscribe isn-wg <first name> <last name>`
 Archive:

Description of Working Group:

The Internet School Networking Working Group is chartered to facilitate the connection of the United States' K-12 (Kindergarten-12th Grade) schools, public and private, to the Internet, and school networking in general.

It is critically important that national networking for K-12 education proceed along established lines of protocol, using existing network structures. The working group's first priority will be to establish guidelines for specialized user interfaces. K-12 networking will also require other support services, such as directories, online and hotline help, specialized training programs and collaborative projects with instructional and curriculum groups, disciplinary groups and post-secondary institutions.

While the initial focus is school networking in the US, the working group will coordinate its efforts with similar activities in other countries and regions of the world.

Goals and Milestones:

- | | |
|----------|--|
| Done | Meet for the first time at IETF and establish approval of Charter. Examine the status of projects in process when Working Group was created. Begin work on list of deliverables. |
| Jan 1992 | Release X.500 "K-12 People Directory" version in collaboration with Merit. Develop plans and milestones for K-12 Resources Directory. |
| Mar 1992 | First draft of information packet document for computing directors to assist them in connecting K-12 schools. First draft of user interface guideline statement. |
| May 1992 | Release X.500 K-12 Resource Directory version in collaboration with Merit. Present final draft guideline statement. |

Internet-Drafts:

“FYI on Questions and Answers: Answers to Commonly Asked ”Elementary and Secondary School Internet User” Questions”, 06/11/1993, J. Sellers, A. Marine <draft-ietf-isn-faq-01.txt>

CURRENT MEETING REPORT

Reported by Jennifer Sellers/NASA NREN

Minutes of the Internet School Networking Working Group (ISN)

Introductions

The meeting opened with a sharing of experiences as the participants introduced themselves and told of projects in school networking of which they were aware. Several projects were mentioned.

Editor's Note: A list of projects is available via FTP or mail server from the remote directories as /ietf/isn/isn-minutes-93jul.txt. Refer to Section 1.2 of the proceedings for retrieval instructions.

Educator Presentation

Klaus Fueller, a German educator, gave a presentation on his perspectives in school networking. He mentioned that some of the pedagogical potential of networks is that using them in instruction can encourage students to read, write, follow discourse and discussion, and practice foreign language as a means of communication rather than as a classroom exercise. For teachers, the use of networks can allow for collaboration. In teaching students about networks, they learn a new means of information retrieval, how to make (local and international) contacts, and how to use network resources. Klaus emphasized the use of e-mail and suggested that e-mail-only, rather than full Internet access, may serve many schools very well. He suggested that non-technical teacher training materials, developed by teachers on the net, are sorely needed. He will be giving a teacher training course on 20 September and is looking for teacher volunteers around the world to coordinate with him. If interested, please contact him.

Klaus then gave an overview of the domain naming of schools in Germany, which follows the scheme of school_name.county_name.state_name.schule.de.

Finally, Klaus mentioned the technical issues and barriers to elementary and secondary school networking, some of which are that there are no technical people at the local level, materials for technical training are scarce, software is needed to simplify use of the network, and schools without LANs and only stand-alone machines need to be considered. For further information, send mail to postmaster@schule.de.

Review of FAQ Document

The FAQ document, which is now an Internet-Draft, draft-ietf-isn-faq-01.txt, will be an RFC FYI document by the time of the Houston IETF in November. There will be information added on “low-low” level connectivity, that is, mail and news only. The document is currently extremely US-centric, and those outside the US are encouraged to give input that will help broaden it, both in choice of vocabulary and in scope. It is hoped that all members of the working group will not only review the document themselves but also disseminate it to the educators with whom they work, for review and comment. A number of comments for improvement were made during the meeting. Thanks to all who gave feedback, and please follow up by posting the feedback electronically to this list or directly to Jennifer Sellers by 6 August if possible.

Connectivity Models Document

The group discussed Gene Hastings' (hastings@psc.edu) document, “Connectivity Models for Internet Access” (available via anonymous FTP as [models.mar93/a.psc.edu](ftp://models.mar93/a.psc.edu)). Ellen Hoffman, co-Chair of the USERDOC2 Working Group, will coordinate with Gene and Jennifer to add to the beginning of the connectivity models document some general issues covered in the “How to Connect” document of the USERDOC2 Working Group. The ISN Working Group hopes to release the document, at a minimum, as an Internet-Draft, and possibly an FYI RFC, by the Houston IETF in November.

Review and Revise Charter

The milestone stating that the group will release a directory (X.500) of people involved in networks in elementary and secondary education was revised. It had been expected that Merit would receive funds to carry out much of this task, and those funds were not forthcoming. Additionally, it is outside the scope of a working group to maintain data, and in order to make this directory useful, data would need to be maintained. Some discussion of directory services offered by the InterNIC ensued, and April Marine will coordinate with InterNIC staff to discover what precisely is available and if it might serve the original intent of the group's milestone. April will also talk to the players in RIPE who might be able to coordinate the service in Europe.

The text of the charter was also discussed, and the group agreed that a few key concepts should be incorporated into the new charter:

- The group should be international in focus.
- The wording “primary and secondary” (education) should replace “K-12.”

- The group could be most useful by identifying barriers and/or problems in the area of school networking around the world and proposing solutions.
- The group will work to identify gaps where solutions are not readily available or don't seem to be on the horizon.

Klaus Fueller, Bruce Nelson, and Aad Nienhuis volunteered to craft a draft charter and post it to the ISN mailing list for comments.

The meeting lasted approximately two hours and was adjourned around 15:30 local time. Thanks to Ray Harder who took notes for the minutes during the meeting.

Attendees

| | |
|----------------------|-------------------------------|
| Rich Bowen | rkb@ralvm11.vnet.ibm.com |
| Thomas Brisco | brisco@pilot.njin.net |
| J. Nevil Brownlee | nevil@ccu1.aukuni.ac.nz |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Klaus Fueller | klausf@osgo.ks.he.schule.de |
| Raymond Harder | rharder@eis.calstate.edu |
| Maria Heijne | maria.heijne@surfnet.nl |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Ellen Hoffman | ellen@merit.edu |
| Andrew Knutsen | andrewk@sco.com |
| John Larson | jl Larson@parc.xerox.com |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Paul Lustgarten | Paul.Lustgarten@att.com |
| Paolo Malara | malara@crs4.it |
| April Marine | april@atlas.arc.nasa.gov |
| Cynthia Mills | cmills@bbn.com |
| Mitra | mitra@pandora.sf.ca.us |
| Bruce Nelson | bnelson@novell.com |
| David O'Leary | doleary@cisco.com |
| Mel Pleasant | pleasant@hardees.rutgers.edu |
| Joyce K. Reynolds | jkrey@isi.edu |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| Patricia Smith | psmith@merit.edu |
| Milan Sova | sova@feld.cvut.cz |
| Claudio Topolcic | topolcic@cnri.reston.va.us |
| Marc van der Noordaa | noordaa@rare.nl |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Sam Wilson | sam.wilson@ed.ac.uk |

2.9.5 Network Information Services Infrastructure (NISI)

Charter

Chair(s):

April Marine, april@atlas.arc.nasa.gov
Pat Smith, psmith@merit.edu

Mailing Lists:

General Discussion: nisi@merit.edu
To Subscribe: nisi-request@merit.edu
Archive:

Description of Working Group:

The NISI Working Group will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services “infrastructure” for the Internet. The work will begin with existing NIC functions and services and should build upon work already being done within the Internet community. A primary goal of the group is to facilitate the development of relationships between NICs that will result in the presentation of a seamless user support service. NISI will work with all NICs, including the InterNIC, to achieve the goal of a fully-functioning, cooperative mesh of worldwide NICs. In addition to creating policies for interaction, NISI will address areas such as common information formats, methods of access, user interface, and issues relating to security and privacy of Internet databases.

Goals and Milestones:

- | | |
|------|--|
| Done | Complete draft for phase 2 suggesting cooperative agreements for NICs. |
| Done | Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of ‘cooperative agreements’ which will allow NICs to work together more effectively to serve users. |
| Done | Revised draft document ready for working group review. Document defines NIC functions and suggests some standardizations for NIC services, as well as offers new mechanisms for exchanging information between NICs. |
| Done | Document submitted as Internet-Draft for comment from a wider Internet audience. |
| Done | Working group discussed current Internet-Draft and suggested minor revisions. Decision made to continue Working Group activity beyond this document. |
| Done | First document released as Informational RFC. Outline and discuss new NISI tasks at IETF meeting. |

- Done Write a document explaining the security issues of privacy and accuracy in Internet databases. Publish as an Informational RFC.
- Jun 1993 Post an Internet-Draft describing NIC interrelationships.
- Jun 1993 Post an Internet-Draft of a NIC user handoff procedure based on the UCP work.
- Done Post an Internet-Draft describing accessing the nic-profiles data in the X.500 database at Merit.
- Dec 1993 Submit the NIC Interrelationship document for consideration as an FYI RFC.
- Dec 1993 Submit the User Handoff procedures for consideration as an FYI RFC.
- Dec 1993 Submit the Nic-Profiles paper for consideration as an FYI RFC.

Internet-Drafts:

“Current NIC Interrelationships”, 06/28/1993, A. Marine <draft-ietf-nisi-nics-00.txt>

Request For Comments:

- RFC 1302 “Building a Network Information Services Infrastructure”
- RFC 1355 “Privacy and Accuracy Issues in Network Information Center Databases”

CURRENT MEETING REPORT

Reported by April Marine/NASA NAIC

Minutes of the Network Information Services Infrastructure Working Group (NISI)

Introductions

Taking advantage of our first meeting in Europe, we took some time to introduce ourselves. The chairs were very happy that the group included so many new faces. In fact, the whole meeting benefited from the new perspectives brought by our European and other non-American participants.

Document Discussions

NISI had one Internet-Draft and one document proposal to discuss, and these discussions took up the bulk of the meeting.

- “Current NIC Interrelationships”

The current Internet-Draft aims to capture how the many different types of NICs currently interrelate. A draft drawing that breezed by the earlier discussion in Columbus, Ohio, brought forth much more discussion in Amsterdam. It was finally decided to modify the draft in a couple of ways. First, the one “generic” picture currently included will be scrapped in favor of at least three pictures depicting actual examples, each representing a different region, e.g. Europe, the US, and the Pacific Rim. Other areas are welcome, but were not initially discussed. Plus, a deeper discussion of the various roles different organizations acting as NICs play will be included. Other minor editorial suggestions were also made.

- Paper on Coordinating Handoff of User Problems Between NICs

The next document proposed was for recommendations regarding how to coordinate interaction between NICs when a user query requires the services of more than one organization. This also engendered an interesting discussion, with some feeling that as the net becomes more commercialized, “customer service” segments of it will interact less and less. So, it was agreed that the thrust of the document would be toward making recommendations as to how NICs would interact with each other i.e., leave the NIC-to-end-user interaction to each NIC. NICs would be encouraged to follow such recommendations as they choose. NISI recognizes that we can’t tell the NICs what to do, only make suggestions.

Other Items

Pat announced that the NIC profile information is now available via gopher. While it is now much easier to view the NIC information, we need to consider other means of keeping the information current. Pat will talk further with the InterNIC about this. Pushpendra Mohta of the InterNIC agreed to follow up with Pat.

A suggestion was made to have a mailing list where NIC people could exchange e-mail on various problems. We have such a list, called `nic-forum@merit.edu`. We decided to try to jump-start discussion on the list. In addition, perhaps the InterNIC could mention the mailing list at their "NIC Fest" coming up.

The problem of getting people to register IN-ADDR information once they have a net number assigned was mentioned. Ellen Hoffman said that some discussion of that could fit into the "How to Connect" document that the USERDOC2 Working Group will be working on.

Wrap Up

The meeting concluded with the chairs promising to solicit input, especially drawings of NIC relationships, from the attendees.

Attendees

| | |
|----------------|--|
| Matti Aarnio | <code>mea@nic.nordu.net</code> |
| Jodi-Ann Chu | <code>jodi@uhunix.uhcc.hawaii.edu</code> |
| Henry Clark | <code>henryc@oar.net</code> |
| David Conrad | <code>davidc@iij.ad.jp</code> |
| Jill Foster | <code>Jill.Foster@newcastle.ac.uk</code> |
| Osten Franberg | <code>euaokf@eua.ericsson.se</code> |
| Ben Geerlings | <code>ben.geerlings@surfnet.nl</code> |
| Marcel Graf | <code>graf%dhdibm1.bitnet@vm.gmd.de</code> |
| Maria Heijne | <code>maria.heijne@surfnet.nl</code> |
| Ellen Hoffman | <code>ellen@merit.edu</code> |
| Per Jacobsen | <code>p.h.jacobsen@usit.uio.no</code> |
| Xander Jansen | <code>xander.jansen@surfnet.nl</code> |
| Brewster Kahle | <code>Brewster@wais.com</code> |
| Lothar Klein | <code>lothar.klein@gmd.de</code> |
| Steen Linden | <code>steen.linden@uni-c.dk</code> |
| John Lindsay | <code>lindsay@kingston.ac.uk</code> |
| Anne Lord | <code>anne@ripe.net</code> |
| Gary Malkin | <code>gmalkin@xylogics.com</code> |
| April Marine | <code>april@atlas.arc.nasa.gov</code> |

| | |
|----------------------|-------------------------------|
| Brian May | Brian.May@mel.dit.csiro.au |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Pushendra Mohta | pushp@cerf.net |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Juergen Rauschenbach | jrau@dfn.de |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Srinivas Sataluri | sri@internic.net |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| Patricia Smith | psmith@merit.edu |
| Marc van der Noordaa | noordaa@rare.nl |
| Ruediger Volk | rv@informatik.uni-dortmund.de |

2.9.6 Network Training Materials (TRAINMAT)

Charter

Chair(s):

Jill Foster, Jill.Foster@newcastle.ac.uk

Mailing Lists:

General Discussion: us-wg@nnsf.nsf.net

To Subscribe: us-wg-request@nnsf.nsf.net

Archive: nnsf.nsf.net:~/nsfnet/us-wg*

Description of Working Group:

Widespread familiarity with global network services and competence in using them brings benefit to individual users, enriches the information skills and resources of the community and optimises the return in investment in networked services.

The Network Training Materials Working Group is chartered to enable the research community to make better use of the networked services. Towards this end, the working group will work to provide a comprehensive package of "mix and match" training materials for the broad academic community which will: 1) enable user support staff to train users to use the networked services, and 2) provide users with self-paced learning material. In the first instance, it will not deal with operational training.

This working group is the IETF component of a joint RARE/IETF group working on network training materials.

The working group will create a catalogue of existing network training materials (using the TopNode cataloguing fields where appropriate), identify the gaps in network training materials and work to identify the problems associated with hands on training workshops using networked services providing a real service.

Goals and Milestones:

- | | |
|----------|---|
| Done | First working group meeting. Review and approve the charter with a review of documents and materials to be written. |
| Jul 1993 | Post the catalogue of training materials as an Internet-Draft. |
| Dec 1993 | Submit the catalogue of training materials for review and publication as an Informational RFC. |

CURRENT MEETING REPORT

Reported by Ellen Hoffman/Merit

Minutes of the Network Training Materials Working Group (TRAINMAT)

Ellen Hoffman announced that, due to a change in her role with Merit, she could no longer act as co-chair. At the end of the meeting, two volunteers offered to consider becoming co-chair. Jill Foster and Joyce Reynolds, User Services Area Director, will work out the details in the next few weeks.

Round Table Sharing of Experiences

The session started with introductions, and each person described the training activities at their organizations. This sharing of information included training activities around the world (Europe, Pacific Rim, North America and Australia). These included subject-based training, involvement of librarians in training, and training educators (in NASA, K-12 and German schools). The group would like to see network training materials used in a pilot on networked multimedia delivery.

RARE and IETF Activities

Jill briefly discussed the RARE Network Training Task Force. She noted that the group is studying what exists and seeks to provide a mix-and-match package of materials; it is more operations- than standards-oriented. She outlined a plan to develop a training video building on a COSINE video.

Jill also described the NISP/ITTI network training materials project at the University of Newcastle, UK. An initial stage of the project developed a catalog of training materials and work is ongoing to develop curriculum materials for Internet training. The generic training materials can be used by trainers as-is, or tailored to particular circumstances and/or subject groups. The first unit of this package, an overview entitled "The World from your Desktop," was produced by Margaret Isaacs, Margaret.Isaacs@newcastle.ac.uk. This is a presentation (Powerpoint), accompanying handout, speakers notes and suggestions for demonstrations. There are worksheets for an accompanying hands-on workshop. Jill ran a short demonstration of the disk-based presentation and handed out examples of the accompanying material. This is UK-biased, but should be easily adaptable by trainers. Margaret sent out messages about obtaining the files via FTP from tuda.ncl.ac.uk in `/pub/network-training`. README files have more details on what is available.

Catalogue of Network Training Materials

The catalogue of network training materials was discussed. Various volunteers will look at defining the template for collecting the information, and will each take about ten items from the NISP/ITTI catalogue and update the information on these items, as well as collect information on further items.

- Data Elements

Data elements for the catalog were reviewed and the group reached consensus on most areas. Jill will mail a revised set to the TRAINMAT mailing list. Issues discussed included audience type, keywords, media type and access information, cost as concept rather than specific number (free versus charge), and whether to include all fields for all types (FTP, telnet, etc.) in every entry, or just those that apply.

- Subject Headings

Volunteers will agree on a limited set of subject headings based on the actual material collected for the catalogue.

Subject Resource Guides

A discussion of subject-specific Internet resource guides was held. It was agreed that subject resource guides were a useful aid in making training relevant. Many exist and we should pool what we have and develop others. These should be along the lines of "the top 20 most interesting resources in subject area X" rather than an exhaustive list. Australian and UK librarians are willing to help on this. Other volunteers are encouraged.

Using the Network to Deliver Training.

Issues in using the network for training delivery were briefly reviewed in an open discussion. Included were recovery of costs for training development, self- or peer-based training through applications like IRC or interactive "games" like MUD or MUSE. Jill noted the multimedia group was looking for pilots, and that Internet training might be a possibility. This will be further discussed at the next meeting. There was also the recognition that training for the absolute novice did not always work on the Internet, as these folks are least likely to be able to use the Internet when they start. Jill requested information on any training projects others might know of.

Liaison Activities

Susan Harris has been assisting in posting a limited set of training information to the mailing list from net-train.

The RARE ISUS Working Group and the USWG Working Group mailing lists are currently being used for TRAINMAT discussion. The USWG addresses can be found in the TRAINMAT charter; the RARE addresses are:

General discussion: wg-isus@rare.nl
 To subscribe: mailserver@rare.nl
 In the body of the message type: subscribe wg-isus [your first name] [your last name]

Charter Revisions

Jill noted that the dates on the charter need revision, mainly because the data element issues have slowed production of a catalog. This will be done and posted to the list.

Attendees

| | |
|---------------------|--------------------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| George Clapp | clapp@ameris.center.il.ameritech.com |
| Reinhard Doelz | embnet@comp.bioz.unibas.ch |
| Michael Erlinger | mike@jarthur.claremont.edu |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Klaus Fueller | klausf@osgo.ks.he.schule.de |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Raymond Harder | rharder@eis.calstate.edu |
| Maria Heijne | maria.heijne@surfnet.nl |
| Ellen Hoffman | ellen@merit.edu |
| Nandor Horvath | horvath@sztaki.hu |
| Xander Jansen | xander.jansen@surfnet.nl |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| John Kunze | jak@violet.berkeley.edu |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| John Lindsay | lindsay@kingston.ac.uk |
| April Marine | april@atlas.arc.nasa.gov |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Joyce K. Reynolds | jkrey@isi.edu |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| Patricia Smith | psmith@merit.edu |

| | |
|----------------------|------------------------------------|
| Bert Stals | <code>bert.stals@surfnet.nl</code> |
| Hisao Uose | <code>uose@tnlab.ntt.jp</code> |
| Marc van der Noordaa | <code>noordaa@rare.nl</code> |
| Chris Weider | <code>clw@merit.edu</code> |

2.9.7 Networked Information Retrieval (NIR)

Charter

Chair(s):

Jill Foster, Jill.Foster@newcastle.ac.uk
George Brett, George.Brett@cnidr.org

Mailing Lists:

General Discussion: nir@mailbase.ac.uk
To Subscribe: mailbase@mailbase.ac.uk
In Body: subscribe nir <first name> <last name>
Archive: mailbase.ac.uk:~/pub/nir

Description of Working Group:

As the network has grown, along with it there has been an increase in the number of software tools and applications to navigate the network and make use of the many, varied resources which are part of the network. Within the past year and a half we have seen a wide spread adoption of tools such as the Archie servers, the Wide Area Information Servers (WAIS), the Internet Gopher, and the WorldWide Web (WWW). In addition to the acceptance of these tools there are also diverse efforts to enhance and customize these tools to meet the needs of particular network communities.

There are many organizations and associations that have recently begun to focus on the proliferating resources and tools for Networked Information Retrieval (NIR). The Networked Information Retrieval Working Group will be a cooperative effort of three major players in the field of NIR: IETF, RARE, and the Coalition for Networked Information (CNI) specifically tasked to collect and disseminate information about the tools and to discuss and encourage cooperative development of current and future tools.

The NIR Working Group intends to increase the useful base of information about NIR tools, their developers, interested organizations, and other activities that relate to the production, dissemination, and support of NIR tools, to produce documentation that will enable user services organizations to provide better support for NIR tools, to develop materials that will assist the support and training of end users and to evolve in the future as necessary to meet and anticipate changes in the field (i.e., NIR tools, protocols, network topology, etc.).

Goals and Milestones:

- Done Review and comment on proposed charter. Discuss applications template and organizational template.
- Sep 1992 Post an Internet-Draft containing the Applications and Organizational Templates.
- Oct 1992 Post an Internet-Draft of the “Consumer Report” with introductory material and completed templates.
- Dec 1992 Submit “Consumer Report” to the IESG for publication as an Informational RFC.

Internet-Drafts:

“A Status Report on Networked Information Retrieval: Tools and Groups”,
03/24/1993, J. Foster, G. Brett, P. Deutsch <draft-ietf-nir-status-report-00.txt>

CURRENT MEETING REPORT

Reported by Kevin Gamiel/MCNC - CNIDR

Minutes of the Networked Information Retrieval Working Group (NIR)

Jim Fullton gave a general overview of the MCNC Clearinghouse for Networked Information Discovery and Retrieval (CNIDR). CNIDR is funded by the National Science Foundation to promote the use and development of networked information tools and systems.

Anders Gillner from a RARE ISUS task force gave a presentation on the Eurogopher. The Eurogopher is an attempt to establish a geography-based as well as subject-based gopher infrastructure throughout Europe.

The NIR report was discussed in some detail and edited by the group.

- It was decided that NCSA's Mosaic for X should not be a stand-alone category, but rather indicated as a client under the World Wide Web category.
- The WAIS and freeWAIS categories will be merged as WAIS/freeWAIS. Brewster Kahle and Jim Fullton will write parallel descriptions of each and will resubmit those writings to Jill for the next update.
- The Grateful Med item was voted out of the report as it was generally considered not to be an open network resource. Also voted to be removed from the report was the group Lister Hill National Center for Biomedical Communications, National Library of Medicine.
- Section 3 of the report will be edited to expand all acronyms.
- The ZIT has been dissolved and therefore will be removed from the report.
- Mark Needleman completed a tool template for the Z39.50 protocol. Since it is a protocol and not a tool, it was decided to port the template to a group template on the Z39.50 Implementors Group (ZIG).
- comp.infosystems.www will be added to the newsgroup templates.
- CNIDR is considering publishing the report on-line. Access could be via FTP, gopher, Wais, and WWW. Any details will be posted to the NIR list as they become available (after the meeting, a representative from the InterNIC expressed interest in doing this as well).
- It was agreed that the full report should be updated once per year, although more regular updating of the individual sections held on-line should be possible.

Reinhard Doelz discussed his NIR tool checklist. The checklist was very extensive and ambitious. Several schemes for maintaining the list were discussed, including on-line telnet "holes" allowing the user to complete the checklist. As a first cut, however, April volunteered to prune the checklist into a fairly simple list. Reinhard's work was deemed a useful resource for users deciding on NIR tools and will be posted to the mailing list.

The group agreed to remove the sections concerning the provision of documentation and training materials for NIR tools from the charter. It was agreed that this is too ambitious and out of scope.

Attendees

| | |
|-------------------|-----------------------------|
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Toshiya Asaba | asaba@iij.ad.jp |
| Zavisa Bjelogrić | zavisa@roma.intecs.it |
| Luc Boulianne | lucb@cs.mcgill.ca |
| Kursat Cagiltay | kursat@vm.cc.metu.edu.tr |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| David Conrad | davidc@iij.ad.jp |
| Reinhard Doelz | embnet@comp.bioz.unibas.ch |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Jim Fullton | fullton@cnidr.org |
| Kevin Gamiel | kgamiel@cnidr.org |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Maria Heijne | maria.heijne@surfnet.nl |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Rune Hjelsvold | Rune.Hjelsvold@idt.unit.no |
| Per Jacobsen | p.h.jacobsen@usit.uio.no |
| Xander Jansen | xander.jansen@surfnet.nl |
| Brewster Kahle | Brewster@wais.com |
| John Kunze | jak@violet.berkeley.edu |
| John Lindsay | lindsay@kingston.ac.uk |
| Anne Lord | anne@ripe.net |
| April Marine | april@atlas.arc.nasa.gov |
| Linda Millington | l.millington@noc.ulcc.ac.uk |
| Pushpendra Mohta | pushp@cerf.net |
| Lisa Nielsen | lisa@hawaii.edu |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Rickard Schoultz | schoultz@admin.kth.se |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| A. Velu Sinha | avsinha@attmail.com |
| Patricia Smith | psmith@merit.edu |

| | |
|----------------------|-----------------------|
| Milan Sova | sova@feld.cvut.cz |
| Bert Stals | bert.stals@surfnet.nl |
| Marc van der Noordaa | noordaa@rare.nl |
| Guido van Rossum | guido@cwi.nl |
| Chris Weider | clw@merit.edu |

2.9.8 Uniform Resource Identifiers (URI)

Charter

Chair(s):

Jim Fullton, fullton@cnidr.org
Alan Emtage, bajan@bunyip.com

Mailing Lists:

General Discussion: uri@bunyip.com
To Subscribe: uri-request@bunyip.com
Archive: archives.cc.mcgill.ca:~/pub/uri-archive

Description of Working Group:

The Uniform Resource Identifiers Working Group is chartered to define a set of standards for the encoding of system independent resource location and identification information for the use of Internet information services.

This working group is expected to produce a set of documents that will specify standard representations of Uniform Resource Locators (URLs) for encoding location and access information across multiple information systems. Such standards are expected to build upon the document discussed at the UDI BOF session held during the 24th IETF meeting in Boston, Unique Resource Serial Numbers (URSNs) which specify a standardized method for encoding unique resource identification information for Internet resources, and Uniform Resource Identifiers (URIs) which specify a standardized method for encoding combined resource identification and location information systems to be used for resource discovery and access systems in an Internet environment.

Such a set of standards will provide a framework that allows the Internet user to specify the location and access information for files and other resources on the Internet, users and network-based tools to uniquely identify specific resources on the Internet, and the creation and operation of resource discovery and access systems for the Internet. The security of such resource discovery services will also be considered to be an integral part of the work of this group.

Goals and Milestones:

- | | |
|------|---|
| Done | Review and approve the charter making any changes deemed necessary. Examine the scope of the recommended documents. Review the first draft of a proposal for Uniform Resource Locators already available. |
| Done | Submit URL document as an Internet-Draft. Review additional draft documents and determine necessary revisions. Follow up discussion will occur on mailing list. |

Nov 1993 Submit the URL document to the IESG for publication as a Proposed Standard RFC.

Internet-Drafts:

“Uniform Resource Locators”, 04/26/1993, T. Berners-Lee <draft-ietf-uri-url-01.txt, .ps>

“Uniform Resource Names”, 05/17/1993, C. Weider, P. Deutsch <draft-ietf-uri-resource-names-00.txt>

CURRENT MEETING REPORT

Reported by Jim Fullton/MCNC - CNIDR

Minutes of the Uniform Resource Identifiers Working Group (URI)

Session I

The meeting began with introductions, followed by Tim Berners-Lee's presentation of the current URL draft. After discussion, the following changes were proposed:

- Remove the AFS type.
- Add the "mailto" type, `mailto:timbl@cern.ch`, designed to be a generalized mail transmission URL, and to help solve the problem of how to contact an author.
- Wrapping—the formal URL wrapper is defined as `<>`.

A significant portion of the first session was spent discussing the merits of creating a more generalized URL specification, with the current specification as a specific case. Peter Deutsch advocates the creation of this far more generalized specification and will examine the issue further before the next IETF. During these discussions, it became clear that the allowable character set needed refining.

The three general proposals were:

1. No restrictions on character sets.
2. Significant restrictions on the character set; i.e. the current document.
3. Looser restrictions on the character set, but stay with the current proposal.

Further discussion of URL issues was put off until the third session. Chris Weider closed out the first session with a tutorial on current URN issues.

Session II

The second session was spent on URN issues. Karen Sollins made a presentation describing a system under development at MIT.

After further discussion of URNs, the group agreed that several issues needed to be resolved before URNs could be properly specified:

- Detection of duplicates
- Fragment specifiers

- Version specifiers
- Type specifiers

After vast amounts of discussion, fragment, version, and typing information was moved to a Uniform Resource Citation object, which has not received a formal name. The URN was defined as a URN identifier (RURNS), a naming authority string and an opaque string defined by the naming authority as being unique within its name space. The whole URN is wrapped in <>, just like a URL:

<URN:NA:OpStr>

A URN should occupy the same character space as a URL, for ease of applications upgrades.

Session III

The discussion of URLs was continued in the third session. Mitra led a discussion about modifications to the URL character set, and the group agreed to modify the URL draft to allow the use of the following characters:

:
;
=
,
”

Other characters not included in the draft are excluded.

A vote on the modified draft was held, with the group recommending adoption of the modified draft by a vote of twenty-nine to one.

The URN format as described in the second session was informally agreed upon, which led into a discussion of naming authorities. It is assumed that many naming authorities will exist, including authorities within an organization whose existence will not be known outside the organization.

Next Time

At the next IETF, an agenda item will be specified for a “sub-ID” discussion to define the format of the naming authority string and opaque string.

Attendees

| | |
|----------------------|------------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| Chris Adie | C.J.Adie@edinburgh.ac.uk |
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Tim Berners-Lee | timbl@info.cern.ch |
| Zavisa Bjelogrić | zavisa@roma.intecs.it |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| James Conklin | jbc@bitnic.educom.edu |
| Glen Daniels | gub@elf.com |
| Peter Deutsch | peterd@bunyip.com |
| Alan Emtage | bajan@bunyip.com |
| Urs Eppenberger | eppenberger@switch.ch |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Jim Fullton | fullton@cnidr.org |
| Kevin Gamiel | kgamiel@cnidr.org |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Anders Gillner | awg@sunet.se |
| Judith Grass | grass@cnri.reston.va.us |
| Maria Heijne | maria.heijne@surfnet.nl |
| Marco Hernandez | marco@mh-slip.cren.edu |
| Rune Hjelsvold | Rune.Hjelsvold@idt.unit.no |
| Nandor Horvath | horvath@sztaki.hu |
| Per Jacobsen | p.h.jacobsen@usit.uio.no |
| Xander Jansen | xander.jansen@surfnet.nl |
| Olle Jarnefors | ojarnef@admin.kth.se |
| Borka Jerman-Blazic | jerman-blazic@ijs.si |
| Brewster Kahle | Brewster@wais.com |
| Michael Khalandovsky | mlk@ftp.com |
| Lothar Klein | lothar.klein@gmd.de |
| Jim Knowles | jknowles@binky.arc.nasa.gov |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| John Kunze | jak@violet.berkeley.edu |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Thomas Lenggenhager | lenggenhager@switch.ch |
| Steen Linden | steen.linden@uni-c.dk |
| John Lindsay | lindsay@kingston.ac.uk |
| Anne Lord | anne@ripe.net |
| April Marine | april@atlas.arc.nasa.gov |
| Chip Matthes | chip@delphi.com |
| Brian May | Brian.May@mel.dit.csiro.au |
| Mitra | mitra@pandora.sf.ca.us |
| Keith Moore | moore@cs.utk.edu |
| Mark Needleman | mhn@stubbs.ucop.edu |
| Clifford Neuman | bcn@isi.edu |

| | |
|------------------------|---|
| Lisa Nielsen | <code>lisa@hawaii.edu</code> |
| Lars-Gunnar Olsson | <code>Lars-Gunnar.Olsson@data.slu.se</code> |
| Jorg Ott | <code>jo@cs.tu-berlin.de</code> |
| Jim Rees | <code>Jim.Rees@umich.edu</code> |
| Ulla Sandberg | <code>ulla@kiera.ericsson.se</code> |
| Rickard Schoultz | <code>schoultz@admin.kth.se</code> |
| David Sitman | <code>a79vm.tau.ac.il</code> |
| Karen Sollins | <code>sollins@lcs.mit.edu</code> |
| Milan Sova | <code>sova@feld.cvut.cz</code> |
| Simon Spero | <code>simon_spero@unc.edu</code> |
| Arend van den Boogaart | <code>a.k.vandenboogaart@research.ptt.nl</code> |
| Marc van der Noordaa | <code>noordaa@rare.nl</code> |
| Guido van Rossum | <code>guido@cwi.nl</code> |
| Ton Verschuren | <code>Ton.Verschuren@SURFnet.nl</code> |
| Chris Weider | <code>clw@merit.edu</code> |

2.9.9 User Documents Revisions (USERDOC2)

Charter

Chair(s):

Ellen Hoffman, ellen@merit.edu

Lenore Jackson, jackson@nsipo.arc.nasa.gov

Mailing Lists:

General Discussion: user-doc@merit.edu

To Subscribe: user-doc-request@merit.edu

Archive:

Description of Working Group:

The focus of the USERDOC2 Working Group is on identifying and locating documentation about the Internet. A major activity is the revision of an existing bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "How to use the Internet" (RFC 1175, FYI 3). This effort will also be used to help locate documentation produced by other organizations and examine the means by which such documents are made available on the Internet. The target audience is those individuals who provide services to end users and end users themselves. The group is also developing a new FYI RFC document designed as a very short bibliography targeted at novice users.

The USERDOC2 Working Group will:

- (1) Identify and categorize useful documents, reference materials, training tools, and other publications about the Internet, particularly those available on-line.
- (2) Publish on-line and hard copies of the bibliography(s) produced and other reference material on documentation as needs are identified.
- (3) Develop and implement procedures to maintain and update the bibliography and investigate methods to provide the information in an on-line format.
- (4) As a part of the update process, identify new materials for inclusion into the active bibliography and identify additional needs which are required for locating documentation and other publications.
- (5) Review procedures for periodic review of the bibliography by the User Services Working Group.
- (6) Examine methods for delivering documentation and work with providers to improve the availability of basic Internet documentation.

Goals and Milestones:

Done Identify new "sources of information" (e.g., individuals, mailing lists, bulletins, etc.) Review existing document and obtain comments from others in USWG about needed revisions at the San Diego IETF.

- Done Publish an Internet-Draft of the short bibliography for novice users.
- Done Submit the revised FYI document to the IESG for publication as an RFC.
- Done Post a revised version of FYI3, "A bibliography of Internetworking Information" as an Internet-Draft.
- Apr 1993 Submit the revised FYI3 to the IESG for publication as an Informational RFC.

Request For Comments:

- RFC 1463 "FYI on Introducing the Internet—A Short Bibliography of Introductory Internetworking Readings for the Network Novice"

CURRENT MEETING REPORT

Reported by Ellen Hoffman/Merit

Minutes of the User Documents Revisions Working Group (USERDOC2)

Current Working Group Activities

The USERDOC2 Working Group has completed two FYI RFC efforts since the last IETF. The new FYIs are:

- FYI 19: "Introducing the Internet—A Short Bibliography of Introductory Internet-working Readings for the Network Novice" by Ellen Hoffman and Lenore Jackson.
- FYI 20: "What is the Internet?" by Ed Krol and Ellen Hoffman.

The on-line files associated with FYI 19 are archived in "Introducing the Internet," along with newer materials that are of interest to beginning users. These archives can be accessed by e-mail, FTP, Gopher, and WAIS, as well as dial-up for those lacking Internet connectivity. Four sites currently maintain copies of the archive, and volunteers were requested for other sites willing to duplicate the files. For additional information on the archive, send an e-mail message to nis-info@merit.edu with the text, "send access.guide" in the body of the message.

Open Forum on Documentation Activities

The discussion began with a presentation by Bert Stals. He described the work of the User Documentation Task Force, a sub-group of the RARE Information Services and User Support Working Group (ISUS). Their goal is to produce documentation for end users on an introductory and intermediate level, without duplicating other efforts such as User-Doc. Their first project is a set of one-pagers, each on a separate applications topic (gopher, e-mail, WAIS, WWW, archie, file transfer, Mailbase). Bert noted user support staff could use the text in their own documentation series (using their own house-style). Alternatively, simple formatted versions of the pamphlet will be available for those who wish to use them directly in several formats, including ASCII text and WordPerfect files. The working group maintains a mailing list for those interested.

To join, send to: mailbase@mailbase.ac.uk
with the text: `subscribe rare-userdoc your_firstname your_lastname`

Bert also described a cooperative project with EARN to update the first edition of EARN's "Guide to Network Resource Tools."

Ellen Hoffman asked for comments regarding whether the IETF should get documentation translated into other languages. The discussion raised several points, including concerns that introductory rather than technical information was the area most in need of translation as most technical people understand English. David Sitman (EARNInfo Officer) said that in his experience people were unwilling to translate documents but were more comfortable creating their own. He also suggested that documentation written in the US tended to be chatty which made it more difficult for non-native English speakers. Issues were also raised about when documentation is effective and the amount of effort required to produce it. Jill Foster noted that there are various levels of documentation, and that each has its own requirements.

On-Going Projects

The next project is a revision of RFC 1175, the long bibliography. The goal is to have a fairly complete draft by the next IETF. A revised charter will be posted (currently the file where the charter should be is empty—a situation the chairs will work on correcting after the meeting). Ellen noted that she would like to include non-English introductory texts in the bibliography. Suggestions for materials to include should be sent to her.

Attendees

| | |
|----------------------|-----------------------------|
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| Henry Clark | henryc@oar.net |
| Klaus Fueller | klausf@osgo.ks.he.schule.de |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Ellen Hoffman | ellen@merit.edu |
| Nandor Horvath | horvath@sztaki.hu |
| Per Jacobsen | p.h.jacobsen@usit.uio.no |
| Dennis Jennings | jennings@irlearn.ucd.ie |
| Ola Johansson | ojn@tip.net |
| April Marine | april@atlas.arc.nasa.gov |
| Chuck McManis | chuck.mcmanis@eng.sun.com |
| Pushpendra Mohta | pushp@cerf.net |
| Lisa Nielsen | lisa@hawaii.edu |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| David Sitman | a79vm.tau.ac.il |
| Patricia Smith | psmith@merit.edu |
| Milan Sova | sova@feld.cvut.cz |
| Bert Stals | bert.stals@surfnet.nl |
| Marc van der Noordaa | noordaa@rare.nl |

2.9.10 User Services (USWG)

Charter

Chair(s):

Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:

General Discussion: us-wg@nic.near.net

To Subscribe: us-wg-request@nic.near.net

Archive: [ftp.near.net:/mail-archives/us-wg*](ftp://ftp.near.net/mail-archives/us-wg*)

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF working groups created to perform certain projects, or outside organizations such as SIGUCCS.)

(1) Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should:

- Clearly address user assistance needs;
- Produce an end-result (e.g., a document, a program plan, etc.);
- Have a reasonably clear approach to achieving the end-result (with an estimated time for completion); and
- Not duplicate existing or previous efforts.

(2) Create working groups or other focus groups to carry out projects deemed worthy of pursuing.

(3) Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

None specified

Request For Comments:

RFC 1150 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. notes"

RFC 1177 "FYI on Questions and Answers - Answers to Commonly Asked "New Internet User" Questions"

- RFC 1206 “FYI on Questions and Answers - Answers to Commonly asked ”New Internet User” Questions”
- RFC 1207 “Answers to Commonly asked ”Experienced Internet User” Questions”
- RFC 1325 “FYI on Questions and Answers Answers to Commonly asked ”New Internet User” Questions”
- RFC 1462 “FYI on ”What is the Internet?””

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

Minutes of the User Services Working Group (USWG)

The meeting opened with introductions and affiliations of each of the User Services Working Group (USWG) participants. Since this was the first IETF outside of the United States, we had increased global attendance from people representing countries in Scandinavia and Europe (including Corsica), Turkey, Australia, New Zealand, Japan, Canada, and the US.

Joyce Reynolds reported on the IETF User Services Area activities including: working groups coming to closure and new working groups starting up, new publications, and current user services related Internet-Drafts postings.

Joyce also mentioned that the e-mail correspondence list for this working group has a new home. The "us-wg" e-mail lists have moved from nsc.nsf.net. Refer to the USWG charter for the new addresses.

Joyce thanked NEARnet, and in particular, Corinne Carroll, Nina Mecht, and Cyndi Mills for their continued assistance and support in maintaining this e-mail list on behalf of the USWG members.

One working group completed its task and came to closure: NOCTools2 (Robert Enger, Chair).

New FYI RFC publications since the last IETF:

| | | | |
|--------|----------|--|--------------|
| FYI 21 | RFC 1491 | "A Survey of Advanced Usages of X.500" | July 1993 |
| FYI 2 | RFC 1470 | "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices" | June 1993 |
| FYI 20 | RFC 1462 | "FYI on "What is the Internet?" | May 1993 |
| FYI 19 | RFC 1463 | "FYI on Introducing the Internet—A Short Bibliography of Introductory Internetworking Readings" | May 1993 |
| FYI 18 | RFC 1392 | "Internet Users' Glossary" | January 1993 |

Current User Services Area related Internet-Draft postings include:

draft-ietf-isn-faq-00.txt
draft-ietf-ids-pilots-00.txt
draft-ietf-iafa-howftp-00.txt
draft-ietf-ids-x500-survey-02.txt
draft-ietf-iiir-html-00.txt, ps
draft-ietf-nisi-nics-00.txt

Other User Services Area related draft postings include:

FYI 4, RFC 1325, "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions" has been updated and is now under review. It may be obtained via anonymous FTP from: naic.nasa.gov as files/fyi4-june93.txt.

April Marine led a review/discussion on the FYI RFC update; FYI 4 "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions", (Also RFC 1325), May 1992.

Bill Manning (Rice University) brought up a discussion on how to "empower" users to utilize and document tools. He also commented to the group that there seems to be missing from the FYI RFC publications a series of notes that tries to address the manners and morals of the collective body. Another topic that may need to be worked on in this forum is how to deal with the basics (e.g., How do "I" get attached?). After some discussion by the participants, Joyce Reynolds, Jodi Chu, and Bill volunteered to take this discussion off-line, review Bill's comments and come back to the USWG with proposed projects.

Cyndi Mills talked briefly about the shutdown of BBN's NNSC site. The services formerly provided by the NNSC have been transferred to a new Network Information Services Management team, collectively known as the InterNIC (the Internet Network Information Center). The NNSC Info-Server has been shut down, and anonymous FTP to nnsf.net has been discontinued.

David Sitman, EARN (European Academic and Research Network) representative, presented a discussion and review of EARN's "Guide to Network Resource Tools" in preparation for submission to the RFC Editor for FYI RFC publication.

The "Guide to Network Resource Tools" is available via Listserv in plain text and PostScript versions. Send mail to `LISTSERV@EARNCC.BITNET` and include the line, `GET NETTOOLS MEMO`, for the plain text version, or `GET NETTOOLS PS`, for the PostScript version. The guide should now be available from `ripe.net` for access via FTP, gopher, WAIS and WWW (at `ftp.ripe.net`, `gopher.ripe.net`, `wais.ripe.net`, etc.)

Jill Foster presented an update on RARE activities, including a report on the RARE Information Services/User Support Working Group (ISUS) activities. ISUS is broken down into several sub-areas: Network User Support, Asynchronous Group Communication, and Networked Information Retrieval and Services Liaison. ISUS Task Forces are being set up to cover the main tasks outlined in the workplan. This includes the following task forces:

- User Documentation
- Document Delivery
- Networked Information Retrieval (joint RARE/IETF/CNI endeavor)
- Coordination of Networked Information Retrieval Services
- Maintenance and Gathering of Information on Networks and Networked Resources
- UNITE “Total Solution” User Interface
- RARE Technical Report 1
- Training
- Publicity and Awareness (joint RARE/IETF working group)
- Support for Special Interest Communities
- Automatic Mailing List Servers
- Multimedia Information Services

Attendees

| | |
|-----------------|-----------------------------|
| Matti Aarnio | mea@nic.nordu.net |
| James Allard | jallard@microsoft.com |
| Kursat Cagiltay | kursat@vm.cc.metu.edu.tr |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| Henry Clark | henryc@oar.net |
| Peter Deutsch | peterd@bunyip.com |
| Reinhard Doelz | embnet@comp.bioz.unibas.ch |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Klaus Fueller | klausf@osgo.ks.he.schule.de |
| Joan Gargano | jcgargano@ucdavis.edu |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Anders Gillner | awg@sunet.se |
| Judith Grass | grass@cnri.reston.va.us |
| Maria Heijne | maria.heijne@surfnet.nl |
| Ellen Hoffman | ellen@merit.edu |
| Nandor Horvath | horvath@sztaki.hu |
| Per Jacobsen | p.h.jacobsen@usit.uio.no |
| Ola Johansson | ojn@tip.net |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| John Lindsay | lindsay@kingston.ac.uk |
| Anne Lord | anne@ripe.net |
| Gary Malkin | gmalkin@xylogics.com |
| Bill Manning | bmanning@rice.edu |

| | |
|----------------------|--------------------------------|
| April Marine | april@atlas.arc.nasa.gov |
| Chip Matthes | chip@delphi.com |
| Cynthia Mills | cmills@bbn.com |
| Pushpendra Mohta | pushp@cerf.net |
| Jun Murai | jun@wide.ad.jp |
| Lisa Nielsen | lisa@hawaii.edu |
| Lars-Gunnar Olsson | Lars-Gunnar.Olsson@data.slu.se |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Tibor Schonfeld | tibor@qsun.att.com |
| Jennifer Sellers | sellers@nsipo.nasa.gov |
| David Sitman | a79vm.tau.ac.il |
| Patricia Smith | psmith@merit.edu |
| Milan Sova | sova@feld.cvut.cz |
| Bert Stals | bert.stals@surfnet.nl |
| Erdal Taner | erdal@vm.cc.metu.edu.tr |
| Fumio Teraoka | tera@csl.sony.co.jp |
| Marc van der Noordaa | noordaa@rare.nl |

2.9.11 Whois and Network Information Lookup Service (WNILS)

Charter

Chair(s):

Joan Gargano, jcgargano@ucdavis.edu

Mailing Lists:

General Discussion: ietf-wnils@ucdavis.edu

To Subscribe: ietf-wnils-request@ucdavis.edu

Archive: [ucdavis.edu:~/archive/wnils](http://ucdavis.edu/~archive/wnils)

Description of Working Group:

The Network Information Center (NIC) maintains the central NICNAME database and server, defined in RFC 954, providing online look-up of individuals, network organizations, key nodes, and other information of interest to those who use the Internet. Other distributed directory information servers and information retrieval tools have been developed and it is anticipated more will be created. Many sites now maintain local directory servers with information about individuals, departments and services at that specific site. Typically these directory servers are network accessible. Because these servers are local, there are now wide variations in the type of data stored, access methods, search schemes, and user interfaces. The purpose of the Whois and Network Information Lookup Service (WNILS) Working Group is to expand and define the standard for WHOIS services, to resolve issues associated with the variations in access and to promote a consistent and predictable service across the network.

Goals and Milestones:

- | | |
|----------|--|
| Done | Review and approve the charter making any changes deemed necessary. Examine the particular functional needs for expanded whois directory service. Begin work on a framework for recommendations. Assign writing assignments for first draft of document. |
| Done | Post the Whois and Network Information Lookup Service Recommendations document as an Internet-Draft. |
| Apr 1993 | Post the revised WHOIS protocol and index service document to the IESG as an Internet-Draft. |
| Done | Post the "Architecture of the Whois++ Index Service" as an Internet-Draft. |
| Jun 1993 | Submit the Whois and Network Information Lookup Service Recommendations document as an Informational RFC. |
| Jun 1993 | Submit the "Architecture of the WHOIS++ Index Service" to the IESG for consideration as an Informational RFC. |

Sep 1993 Submit a revised WHOIS protocol specification and index service document to the IESG for consideration as a Draft Standard.

Internet-Drafts:

“Architecture of the Whois++ Index Service”, 11/23/1992, C. Weider, J. Fulton, S. Spero <draft-ietf-wnils-whois-01.txt>

“Whois and Network Information Lookup Service Whois++”, 07/06/1993, J. Gargano, K. Weiss <draft-ietf-wnils-whois-lookup-00.txt>

CURRENT MEETING REPORT

Reported by Joan Gargano/University of California, Davis

Minutes of the Whois and Network Information Lookup Service Working Group (WNILS)

Approval of the Minutes

The minutes of March 29, 1993 were unanimously approved without changes.

Status of Whois++ Architecture

Many attendees were new to the WNILS Working Group and requested an introduction to the Whois++ architecture. Peter Deutsch presented an overview of the Whois++ architecture and design philosophy. A lengthy discussion followed and the following additions to the protocol were requested.

- Provide a mechanism for supporting 16-bit Unicode.
- Specify the mechanism of handling multimedia data using the MIME standard.
- Provide an option for handling compressed data.

Peter will incorporate the additional specifications and submit an Internet-Draft by July 30, 1993.

Status of The Distributed Whois++ Model - Centroids

Chris Weider presented an overview of centroids. Two problem areas were identified.

1. Centroids don't scale for databases containing large quantities of unique data.
2. The potential exists for cycles or looping in queries because the directory is in a mesh configuration.

Chris agreed further work is needed on indexing but this first implementation should provide a functional system for existing data sets and a testbed for the limitations of this model.

Status of Recommended Modifications to the Whois Protocol

The status update was given by Joan Gargano. The recommended modifications to Whois have been submitted as an Internet-Draft. The host record section will be updated with the following information and resubmitted:

| | |
|------------------------------------|----------|
| Host records | required |
| Full domain name | required |
| IP address | optional |
| System administrator name | optional |
| System administrator telephone | optional |
| System administrator address | optional |
| System administrator email address | optional |
| Type of machine | optional |
| Operating system | optional |
| Mail exchanger | optional |
| Last update | required |
| Location of additional information | optional |

Update of Goals and Milestones

The goals and milestones will be updated to reflect the following:

1. All draft papers will be submitted as Internet-Drafts by the end of July.
2. Two working implementations of clients, servers and centroids will be available by September 30, 1993.

Attendees

| | |
|-----------------|-----------------------------|
| Paul Barker | p.barker@cs.ucl.ac.uk |
| Tim Berners-Lee | timbl@info.cern.ch |
| Kursat Cagiltay | kursat@vm.cc.metu.edu.tr |
| Jodi-Ann Chu | jodi@uhunix.uhcc.hawaii.edu |
| James Conklin | jbc@bitnic.educom.edu |
| Thomas DeWitt | tdewitt@osi.ncsl.nist.gov |
| Jill Foster | Jill.Foster@newcastle.ac.uk |
| Barbara Fraser | byf@cert.org |
| Kevin Gamiel | kgamiel@cnidr.org |
| Joan Gargano | jcgargano@ucdavis.edu |
| Ben Geerlings | ben.geerlings@surfnet.nl |
| Anders Gillner | awg@sunet.se |
| Judith Grass | grass@cnri.reston.va.us |

| | |
|----------------------|---------------------------------|
| Nandor Horvath | horvath@osztaki.hu |
| Olle Jarnefors | ojarnef@admin.kth.se |
| Borka Jerman-Blazic | jerman-blazic@ijs.si |
| Thomas Johannsen | thomas@ebzaw1.et.tu-dresden.de |
| Neil Katin | katin@eng.sun.com |
| Sean Kennedy | liam@nic.near.net |
| Michael Khalandovsky | mlk@ftp.com |
| Peter Koch | pk@techfak.uni-bielefeld.de |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Eliot Lear | lear@sgi.com |
| Steen Linden | steen.linden@uni-c.dk |
| John Lindsay | lindsay@kingston.ac.uk |
| April Marine | april@atlas.arc.nasa.gov |
| John Martin | John.Martin@newcastle.ac.uk |
| Chip Matthes | chip@delphi.com |
| Brian May | Brian.May@mel.dit.csiro.au |
| Mitra | mitra@pandora.sf.ca.us |
| Clifford Neuman | bcn@isi.edu |
| Lisa Nielsen | lisa@hawaii.edu |
| Ronny Nilsen | Ronny.Nilsen@usit.uio.no |
| Petri Ojala | ojala@eunet.fi |
| Lars-Gunnar Olsson | Lars-Gunnar.Olsson@data.slu.se |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Joyce K. Reynolds | jkrey@isi.edu |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Miguel Sanz | miguel.sanz@rediris.es |
| Heiner Schorn | heiner.schorn@umdac.umu.se |
| Rickard Schoultz | schoultz@admin.kth.se |
| Patricia Smith | psmith@merit.edu |
| Suzanne Smith | smith@es.net |
| Milan Sova | sova@feld.cvut.cz |
| Erdal Taner | erdal@vm.cc.metu.edu.tr |
| Marc van der Noordaa | noordaa@rare.nl |
| Guido van Rossum | guido@cwi.nl |
| Ton Verschuren | Ton.Verschuren@SURFnet.nl |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Chris Weider | clw@merit.edu |
| Wilfried Woeber | Wilfried.Woeber@CC.UniVie.ac.at |
| Russ Wright | wright@lbl.gov |

Chapter 3

IAB Open Meeting

3.1 IAB Open Meeting

Reported by Bob Braden/ISI

The IAB held an open meeting at the Amsterdam IETF, on Tuesday evening. About 120 observers attended, approximately 20% of the IETF meeting. The following is a brief summary of the meeting. A more complete summary is available by anonymous FTP from host ftp.isi.edu with pathname pub/IAB/IABmins.jul93.txt.

Standards Procedures Document

Another round of revisions will be made in the replacement for RFC 1310, and a new Internet-Draft will be circulated. However, the IAB feels this document should be published as an RFC as soon as possible. A key issue is the rules for intellectual property, particularly copyrights. The IAB will take steps to inform and involve the Internet community, as soon as ISOC lawyers have prepared new text.

Proposed ISOC Liaison Agreements with ISO and ITU

The IAB accepted a recommendation from Vint Cerf, President of the Internet Society, that a Memorandum of Understanding (MOU) between ISO and ISOC be drafted. This MOU, if accepted by both sides, would form the basis for a Category A liaison relationship with ISO. It would be framed to protect the successful IETF processes for standards making, while establishing the ground rules for interaction between IETF working groups and ISO subcommittees, and any other relations deemed helpful. Vint agreed to draft such a document, for presentation to the Internet community for comments and discussions.

Liaison with the ITU, delayed by their reorganization, is now under active consideration.

Projections of CIDR Effects

There was an extensive discussion of the existing projections of the effects of CIDR on preserving the IP address space and preventing a routing explosion. The uncertainties are still very large, and further studies, with their assumptions carefully documented, are needed.

Architecture

The IAB has initiated a study of modifications of the Internet architecture for shared media, like public data networks.

Steve Kent summarized the ongoing work in the IETF and IRTF towards a security architecture for the Internet.

Attendees

| | |
|-------------------|--|
| Bernt Allonen | bal@tip.net |
| Harald Alvestrand | Harald.Alvestrand@uninett.no |
| Per Andersson | pa@cdg.chalmers.se |
| Anders Baardsgaad | anders@cc.uit.no |
| Tony Bates | tony@ripe.net |
| Nutan Behki | Nutan_Behki@qmail.newbridge.com |
| Robert Blokzijl | K13@nikhef.nl |
| Rebecca Bostwick | bostwick@es.net |
| Robert Braden | braden@isi.edu |
| Ronald Broersma | ron@nosc.mil |
| J. Nevil Brownlee | nevil@ccui.aukuni.ac.nz |
| Ross Callon | rcallon@wellfleet.com |
| Vinton Cerf | vcerf@cnri.reston.va.us |
| George Chang | gkc@ctt.bellcore.com |
| A. Lyman Chapin | lyman@bbn.com |
| Henry Clark | henryc@oar.net |
| Les Clyne | l.clyne@jnt.ac.uk |
| Richard Colella | colella@nist.gov |
| Stephen Coya | scoya@cnri.reston.va.us |
| David Crocker | dcrocker@mordor.stanford.edu |
| Stephen Deering | deering@parc.xerox.com |
| Thomas DeWitt | tdewitt@osi.ncsl.nist.gov |
| Francis Dupont | francis.dupont@inria.fr |
| Tom Easterday | tom@cic.net |
| Toerless Eckert | Toerless.Eckert@informatik.uni-erlangen.de |
| Dino Farinacci | dino@cisco.com |
| Mark Fedor | fedor@psi.com |
| Dennis Ferguson | dennis@ans.net |
| Eric Fleischman | ericf@act.boeing.com |
| Peter Ford | peter@goshawk.lanl.gov |
| Osten Franberg | euaokf@eua.ericsson.se |
| Paul Francis | Francis@thumper.bellcore.com |
| Vince Fuller | vaf@stanford.edu |
| Peter Furniss | p.furniss@ulcc.ac.uk |

| | |
|-------------------------|--------------------------------------|
| Tony Genovese | genovese@es.net |
| David Ginsburg | ginsb@us-es.sel.de |
| Joseph Godsil | jgodsil@ncsa.uiuc.edu |
| Ramesh Govindan | rxg@thumper.bellcore.com |
| Phillip Gross | pgross@ans.net |
| Chris Gunner | gunner@dsmail.lkg.dec.com |
| Joel Halpern | jmh@network.com |
| Jari Hamalainen | jah@rcetre.nokia.com |
| Craig Haney | craig@icp.net |
| Susan Hares | skh@merit.edu |
| Robert Hinden | hinden@eng.sun.com |
| Jack Houldsworth | J.Houldsworth@ste0906.wins.icl.co.uk |
| Chris Howard | chris_howard@inmarsat.org |
| Christian Huitema | Christian.Huitema@sophia.inria.fr |
| Erik Huizer | Erik.Huizer@SURFnet.nl |
| Geoff Huston | g.huston@aarnet.edu.au |
| Sascha Ignjatovic | sascha@veda.co.at |
| Ole Jacobsen | ole@interop.com |
| Ronald Jacoby | rj@sgi.com |
| Dennis Jennings | jennings@irlearn.ucd.ie |
| Borka Jerman-Blazic | jerman-blazic@ijs.si |
| David Johnson | dbj@cs.cmu.edu |
| Philip Jones | p.jones@jnt.ac.uk |
| Cyndi Jung | cmj@3com.com |
| Scott Kaplan | scott@wco.ftp.com |
| Daniel Karrenberg | daniel@ripe.net |
| Frank Kastenholz | kasten@ftp.com |
| Dave Katz | dkatz@cisco.com |
| Peter Kaufmann | kaufmann@dfn.dbp.de |
| Sean Kennedy | liam@nic.near.net |
| Stephen Kent | kent@bbn.com |
| Lothar Klein | lothar.klein@gmd.de |
| John Klensin | Klensin@infofoods.unu.edu |
| Mark Knopper | mak@merit.edu |
| Stev Knowles | stev@ftp.com |
| Rajeev Kochhar | rajeev_kochhar@3com.com |
| Ton Koelman | koelman@stc.nato.int |
| Pekka Kytolaakso | pekka.kytolaakso@csc.fi |
| Sylvain Langlois | Sylvain.Langlois@exp.edf.fr |
| Mark Laubach | laubach@hpl.hp.com |
| Jose Legatheaux Martins | jalm@fct.unl.pt |
| Tony Li | tli@cisco.com |
| John Lindsay | lindsay@kingston.ac.uk |
| Peter Lothberg | roll@stupi.se |
| Paul Lustgarten | Paul.Lustgarten@att.com |

| | |
|----------------------|--------------------------------|
| Carl Malamud | carl@malamud.com |
| Paolo Malara | malara@crs4.it |
| Bill Manning | bmanning@rice.edu |
| David Marlow | dmarlow@relay.nswc.navy.mil |
| Ignacio Martinez | martinez@rediris.es |
| Chuck McManis | chuck.mcmanis@eng.sun.com |
| Peter Merdian | merdian@rus.uni-stuttgart.de |
| Keith Mitchell | keith@pipex.net |
| Jun Murai | jun@wide.ad.jp |
| Peder Chr. Noergaard | pcn@tbit.dk |
| David O'Leary | doleary@cisco.com |
| Petri Ojala | ojala@eunet.fi |
| Geir Pedersen | Geir.Pedersen@usit.uio.no |
| Charles Perkins | perk@watson.ibm.com |
| David Perkins | dperkins@synoptics.com |
| David Piscitello | dave@mail.bellcore.com |
| Willi Porten | porten@gmd.de |
| Lars Poulsen | lars@cmc.com |
| Mark Prior | mrp@itd.adelaide.edu.au |
| Juergen Rauschenbach | jrau@dfn.de |
| Alex Reijnierse | a.a.reijnierse@research.ptt.nl |
| Yakov Rekhter | yakov@watson.ibm.com |
| Jim Romaguera | romaguera@netconsult.ch |
| John Romkey | romkey@elf.com |
| Luc Rooijackers | lwj@cs.kun.nl |
| Ulla Sandberg | ulla@kiera.ericsson.se |
| Miguel Sanz | miguel.sanz@rediris.es |
| Jon Saperia | saperia@tay.dec.com |
| John Scudder | jgs@merit.edu |
| Robert Shirey | shirey@mitre.org |
| W. David Sincoskie | sincos@thumper.bellcore.com |
| Michael St. Johns | stjohns@darpa.mil |
| John Stewart | jstewart@cnri.reston.va.us |
| Marten Terpstra | marten@ripe.net |
| Kamlesh Tewani | ktt@arch2.att.com |
| Richard Thomas | rjthomas@bnr.ca |
| Susan Thomson | set@bellcore.com |
| Antoine Trannoy | trannoy@crs4.it |
| Lea Viljanen | Lea.Viljanen@helsinki.fi |
| Ruediger Volk | rv@informatik.uni-dortmund.de |
| Kirk Williams | kirk@sbcetri.sbc.com |
| Jessica Yu | jyy@merit.edu |
| Paul Zawada | Zawada@ncsa.uiuc.edu |

Chapter 4

IETF Protocol Presentations

4.1 IPng - SIP

Presented by Steve Deering/Xerox PARC

Bio: *Steve Deering is a member of the research staff at the Xerox Palo Alto Research Center (PARC). He has been an active participant in the IETF and IRTF since 1984, and has served as chair of several IETF working groups. His current interests include addressing and routing for very large internets, with support for multicast, mobility, and multi-media services.*


SIP STATUS REPORT

July 16, 1993

Steve Deering
Bob Hinden

Internet Engineering Task Force Meeting
Amsterdam

SIP OVERVIEW



24 bytes long, 4 more than IP

SIP Working Group

July 16 1993

SIP DEMONSTRATION

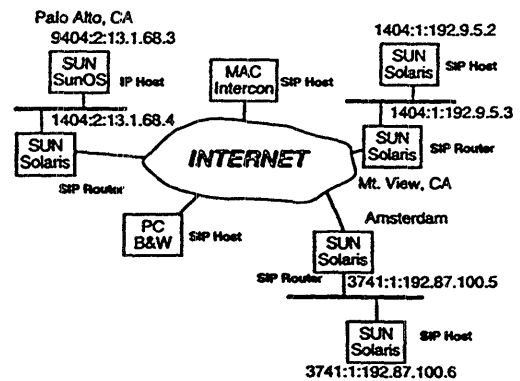
- SIP Inter-Domain Communication **
- SIP <-> IPv4 Translation
- SIP Encapsulation (IPAE)
- SIP Border Router **
- SIP Multicast **
- SIP Traceroute
- SIP Packet Monitoring

** New from Columbus IETF Demonstration

SIP Working Group

July 16 1993

DEMONSTRATION CONFIGURATION



SIP Working Group

July 16 1993

IMPLEMENTATIONS

| OS | Organization | Status |
|---------------|-------------------|---------------------------------------|
| BSD/Mach | INESC | Completed (telnet, NFS, AFS, UDP) |
| DOS & Windows | Beame & Whiteside | Completed (telnet, ftp, tftp, ping) |
| IRIX | Silicon Graphics | In progress (ping) |
| KA9Q | Simpson | In progress (ping) |
| Mac OS | Intercon | Completed (telnet, ftp, finger, ping) |
| OSF | DEC | In Progress |
| Solaris | Sun | Completed (telnet, ftp, tftp, ping) |
| VMS | TGV | Completed (telnet, ftp) |

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

| MISC | Organization | Status |
|------|--------------|-----------|
| Bind | INRIA | Code done |

SIP Working Group

July 16 1993

RECENT WORK

- Host to Router Protocol
 - System Discovery
 - Address Resolution
 - Redirects
 - Black Hole Detection
 - Mobile Host Support
- Auto Configuration
 - Plug & Play
 - Dynamic Address Prefix Reconfiguration
- SIP Numbering Plan
- SRIP - RIP for SIP
- Flow Label

SIP Working Group

July 16 1993

WHAT'S NEXT

- Expand SIP Infrastructure
- SIP Payload Security
 - Authentication, Integrity, and Privacy
 - Based on recent IPSEC work
- Test Advanced SIP Services on DARTnet
 - Flows
 - Mobility
 - Multicast

SIP Working Group

July 16 1993

DOCUMENTS

- *SIP: A Simple Internet Protocol*, S. Deering, May '93 IEEE Network.
- *Simple Internet Protocol (SIP) Specification*, Internet Draft, draft-deering-sip-00.txt
- *SIP-RIP*, G. Malkin, C. Huitema, Internet Draft, draft-ietf-sip-rip-01.txt
- *IDRP for SIP*, S. Hares, Internet Draft, draft-ietf-ipidr-sip-00.txt
- *OSPF for SIP*, C. Huitema, Internet Draft, draft-ietf-sip-ospf-00.txt
- *SIP Addresses in the Domain Name Service Specifications*, C. Huitema, Internet Draft, draft-ietf-sip-dnss-00.txt
- *SIP Program Interfaces for BSD Systems*, R. Gilligan, Internet Draft, draft-ietf-sip-bsdapi-00.txt
- *SIP System Discovery*, W. Simpson, Internet Draft, draft-ietf-sip-discovery-02.txt
- *Administrative Allocation of the 64-bit Number Space*, W. Simpson, Internet Draft, draft-ietf-sip-64bit-plan-00.txt
- *IP Address Encapsulation (IPAE): A Mechanism for Introducing a New IP*, D. Crocker, Internet Draft, draft-ietf-ipae-ncw-ip-00.txt
- *IPv7 Criteria Analysis for IP Address Encapsulation (IPAE) and the Simple Internet Protocol (SIP)*, R. Hinden, Internet Draft, draft-ietf-ipae-ipv7-criteria-00.txt

SIP Working Group

July 16 1993

WORKING GROUP INFORMATION

- SIP Working Group Mailing List
sip-request@caldera.usc.edu
- SIP Archive
parcftp.xerox.com /pub/sip/
- Working Group Chairs
Steve Deering deering@parc.xerox.com
Bob Hinden hinden@eng.sun.com

SIP Working Group

July 16 1993

4.2 IPng - PIP

Presented by Paul Francis/Bellcore

Bio: Paul Francis (formerly Tsuchiya) is currently a member of the technical staff in Bellcore's Information Networking Research Laboratory. He is also currently a Ph. D. student at the University College London. He has been researching routing and addressing problems for nearly a decade. He is the author of several RFCs and Internet-Drafts, covering, among other things, the topics of subnet number assignment, shortcut routing, network address translation, and PIP.

Pip Status

Amsterdam IETF

Paul Francis

Belcore

Belcore
Amsterdam IETF
Primary
Frame 2 of 13

Outline

- Brief Background
- Specification Progress
- Implementation Progress
- Installation Progress
- Future Plans

What is Pip?

- Pip header designed to be flexible and evolvable, but still fast
 - Variable length addressing
 - Multiple routing paradigms
 - Rich QoS handling
- Addressing and routing algorithms designed to provide new features
 - Provider selection
 - Auto-configuration and re-configuration
 - Flexible address assignment
- Working on header/algorithms for flows
 - Pipe = Pip Enhanced.....

Belcore
Amsterdam IETF
Primary
Frame 3 of 13

Specification Overview

- | | |
|---------------------------|----------------------|
| • Pip Forwarding | no change |
| • Pip DNS | revised |
| • Pip Identifiers | revised |
| • Pip Routing | revision in progress |
| • Pip Transition | new |
| • Pip Host Operation | new |
| • PCMP | new (incomplete) |
| • Pip Address Conventions | new |

Belcore
Amsterdam IETF
Primary
Frame 4 of 13

Specification Progress

- Pip DNS
 - Support for transition
 - Time-stamped queries
- Pip Identifiers
 - Weakened hierarchical structure
 - Flat IDs are simpler at Pip layer
- Pip Routing
 - Multi-Level Path Vector (MLPV)—coupled with IDRP
 - Basic algorithms complete
 - Still to do: anycast, tunneling, QoS

Belcore
Amsterdam IETF
Primary
Frame 5 of 13

Specification Progress

- Pip Transition
 - Previously planned IPAE-like transition
 - Found IPAE reliance on IP addresses too constraining
 - New scheme called IPIT (IP Independent Transition)
 - Allows Pip IDs and addresses to be IP independent
- Pip Host Operation
 - Describes how Pip host chooses among multiple Pip addresses
 - Provider selection

Belcore
Amsterdam IETF
Primary
Frame 6 of 13

Specification Progress

- PCMP (Pip's ICMP)
 - Have specified "Packet Not Delivered" messages
 - Analogous to ICMP's Destination Unreachable
 - Have specified Router Discovery messages
 - Other messages yet to be specified
- Pip Address Conventions
 - Details conventions for hierarchical unicast Pip addresses
 - "Class D" style multicast
 - CBT style multicast
 - Anycast

Comments on Implementation and Testing

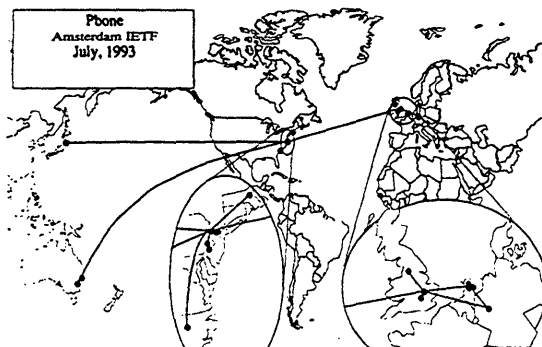
- Given:
 - Importance of IPng
 - Distributed nature and complexity of IPng
 - Difficulty and unknowns of transition
- It is important that we have significant installed-base experience with IPng before we commit
- Concentrate on one complete implementation
 - Rather than maximize number of different implementations

Bellcore Pip Implementation Progress

- Ramesh Govindan, Sue Thomson
 - PCMP
 - Ping
 - Traceroute
 - Router Discovery
 - Host algorithm for choosing among multiple providers
 - gethostbyname() application library replacement
 - Generally allow for address manipulation
 - DNS transition
 - General tightening up of things.....

Installation Progress

- Have already installed 18-node Pbone
 - Amsterdam demo included Ping and Traceroute of Pbone sites
- Purpose of current Pbone:
 - Test forwarding and PCMP on small scale
 - Test Pbone configuration software
 - Test Pip/IP interoperability



Pbone Participants

- AT&T
- Bellcore
- CSIRO (Australia)
- Imperial Cancer Research Fund (London)
- MIT
- NTT
- RIPE (Amsterdam)
- University College London
- University of Manchester (England)
- University of Pennsylvania
- University of Stuttgart

Future Plans

- Install "seed" Pbone
 - Initial infrastructure for Pip deployment
- To do this:
 - Fill out PCMP, transition, DNS specifications
 - Complete the routing specification
 - Complete the implementation
 - Install and configure inverse name lookup support
 - Install and configure Pbone routers
 - Start bringing up users
- Seed Pbone installation in 3 months

4.3 IPng - TUBA

Presented by Mark Knopper/Merit

Bio: Mark Knopper is the manager of Merit's Internet Engineering group which is responsible for NSFNET engineering. Before moving over to the IE group, he was the manager of Merit's Michnet engineering group. He is also active in the IETF OSI-DS Working Group and was co-investigator on the DARPA/NSF FOX (Field Operational X.500) project. Mark is currently co-Chair of the TCP/UDP Over CLNP-Addressed Networks Working Group (TUBA) of the IETF.

State of the TUBA
Working Group

27th IETF Netherlands

July 17, 1993

TUBA Agreements

- tuba-clnp-03.txt TUBA Spec proposed as Proposed Standard
- tuba-sysids-01.txt - 6-bit SysID to be Informational RFC
- RFC 1237 NSAP Allocation to be revised for standards track
- ftp-bigports-01.txt FOOBAR for TUBA to experimental RFC
- DNS: forward lookup from RFC 1238 stable.
Inverse lookup work done this week. New RFC is for standards track.
- Routing and addressing architecture docs from ISO to be used for TUBA
- EON tunneling RFC to be revised for standards track
- IS-IS autoconfiguration scheme for TUBA presented
- Multicast for CLNP being developed in ISO/ANSI
- IS-IS scheme for global CLNP inter-domain routing will be implemented by network providers.
- Interop CLNP demo to be held at Interop in August
- CLNP routing registration to be coordinated by RIPE

TUBA Operations and Deployment

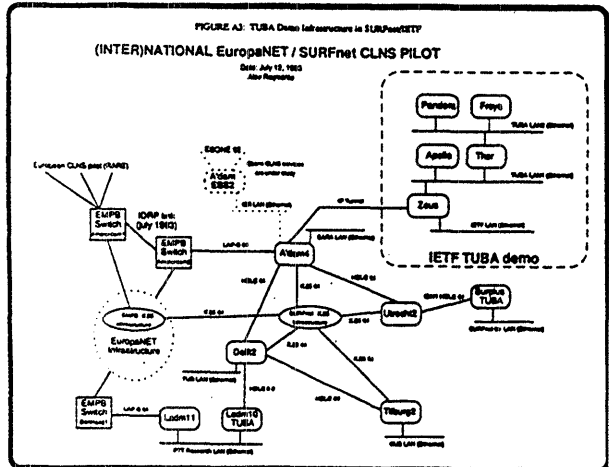
CLNP infrastructure is growing because of TUBA.

TUBA, NOOP groups met jointly with RARE CLNS WG

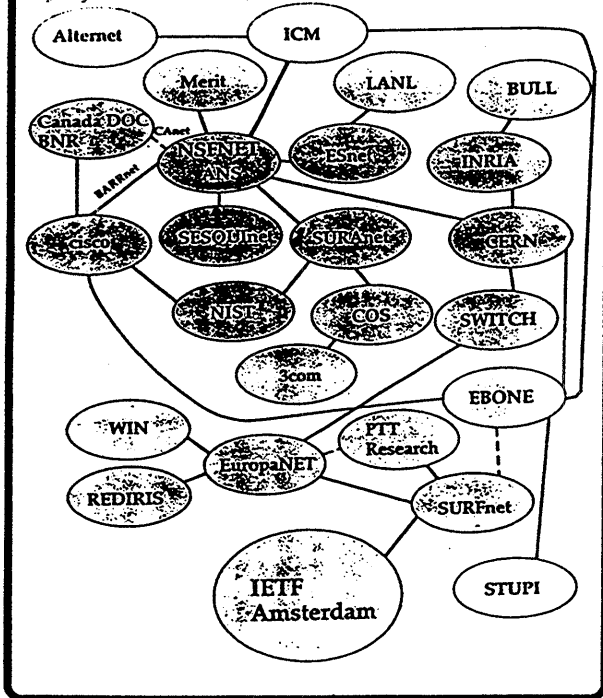
- Static routing is biggest problem.
- Agreed to design and implement dynamic routing using IS-IS in global CLNP infrastructure.
- HEPnet uses CLNP for DECnet V in production.
- New CLNP net management applications available:
Rover, ping alert tool, traceroute.
- RARE WG, RIPE and Merit to collaborate on CLNP object definition for RIPE database

TUBA group is now coordinating with these groups:

- DNS
- FTP
- ATM
- NOOP
- RARE CLNS WG and new RARE TUBA WG
- OSI-EXTEND (BOF)



July, 1993 TUBA Infrastructure and Participants



Next Actions and IPng

IPng chairs will meet with area directors regularly.

TUBA and CLNP work likely to continue, even if withdrawn from IPng consideration in event another approach is chosen.

Group work partitioning:

TUBA/NOOP joint group will continue to discuss TUBA deployment.

"Plenary" TUBA group will discuss protocol issues. OSI-EXTEND will work on next generation of CLNP.

4.4 IPng - TPIX

Presented by Vladimir Sukonnik/Process Software Corporation

Bio: Vladimir Sukonnik is a Senior Software Engineer at Process Software Corporation. He keeps busy by developing TCP/IP application software and worrying about the future of the Internet. In his spare time Vladimir teaches Computer Science at Boston University. Vladimir holds a Master of Science degree in Computer Science from Boston University and Bachelor of Science degrees in Mathematics and Computer Science from Northeastern University.

TP/IX : Internet v7

- Increases address fields and permits hierarchical allocation of addresses
- Explicitly designed to support high bandwidth applications through enlarging the window size, sequence (and ACK) space, port numbers, and PDU size.
- TP/IX datagrams have a forward route ID for extremely fast path, circuit, or flow-based forwarding

Where We Are

- TP/IX and RAP are ready for vendor prototype
- RAP version 1 product is available
- Published experimental RFCs 1475 and 1476
- Published 3 Internet drafts on options
 - AD assignment plan
 - Transit network selection
 - TCP options

Where We Are Going

- RAP version 1 is a product
- TP/IX product release beginning of 1994
- RAP and TP/IX are ready for vendor prototype
- Working group meeting in Houston

Version 4 to Version 7 transition

- Not an architecture change
- No constraints or dependencies in deployment
- No mandatory administration retraining cost
- No "flag day"
- Hosts and routers can be upgraded incrementally

Chapter 5

Technical Presentations

5.1 The IETF from a European Perspective

Presented by Erik Huizer/SURFnet

Bio: *Erik Huizer is Senior Network Consultant at SURFnet bv, the company which operates the academic and research network in The Netherlands. He spends most of his time managing various national and international projects on e-mail and directory services. Huizer lives with internetworking technologies such as TCP/IP, OSI, information services, directory services and user support, as a service developer, beta-tester and mediator. He is currently the Area Director for the Applications Area of the Internet Engineering Task Force. Huizer received his Ph. D. degree in Science and Technology from Delft University of Technology, The Netherlands, in 1987. His preference for chocolate milkshakes is well known.*

SURFnet

Mijn mening

A European's perspective on the IETF

Erik Huizer

SURFnet bv

Erik Huizer
IETF July 99 1

SURFnet

Welkom

This is a private perspective on:

- European networking
- Why IETF in Europe
- What did we do
- What is the cost
- What is the expected benefit
- What are the hurdles

Erik Huizer
IETF July 99 2

SURFnet

Langzaam maar zeker

- Political swamp
- Too much focus on
 - Infrastructure
 - Bandwidth
- Too little attention for
 - applications
 - services
 - users

Erik Huizer
IETF July 99 3

SURFnet

Netwerken voor Onderzoek en Hoger Onderwijs

- Mostly national research networks
 - e.g. SURFnet
- Some multinational
 - Nordunet, EARN
- Mostly multiprotocol
 - IP and TCP/IP suite of protocols
 - X.25 and OSI suite of protocols
- Mostly up to 2 Mbit/s

Erik Huizer
IETF July 99 4

SURFnet

De ruggesgraat

Two international backbones:

- Ebone
 - Ip-only 2 Mb/s infrastructure
 - participants pay
- Europanet
 - Multiprotocol backbone
 - 2 Mb/s still mostly pilot
 - CEC funding

Erik Huizer
IETF July 99 5

SURFnet

Haat en Liefde

RARE :

- network politics
- technical programme
- secretariat

Ripe

- IP-network coordination
- Network coordination center

Erik Huizer
IETF July 99 6

SURFnet

Soms roeien tegen de stroom op

RARE Technical work structure

- RARE Technical Committee (RTC)
- 7 WGs
 - IMM
 - NAP
 - MSG
 - ISUS
 - CHAR
 - LLT
 - Security

Erik Heester
IETF July 98 7

SURFnet

In het land der blinden

Commercial networks:

- Mostly X.25 based
 - PTOs are losing monopoly
 - X.400 services on top
 - Closed User group networks
- IP-based
 - growing
 - various national providers
 - EUnet

Erik Heester
IETF July 98 8

SURFnet

De schildpad en de haas

- TCP/IP
 - still mostly academic and research
 - but growing fast
 - not subsidised
 - not really accepted for doing bussiness (yet)
- OSI
 - PTO
 - Commercial industry
 - Governments
 - CEC

Erik Heester
IETF July 98 9

SURFnet

Schapev voor de dam

Why IETF to Europe?

- Part of larger plan

OSI has:

- Large commitment
- Large investment

Erik Heester
IETF July 98 10

SURFnet

Doelstellingsdefinitie

To get internet technology accepted:

- Internationalization of the Internet (Society)
- Stable, organization
- Not just US-funded
- Open and accepted standards process
- With equal non-US participation

Erik Heester
IETF July 98 11

SURFnet

Alle kleine beetjes

- ISO liaison
- International participation
- on all levels
 - Isoc
 - IAB, IESG
 - IETF
- of all kinds
 - academic & research
 - commercial
 - government

Erik Heester
IETF July 98 12

SURFnet

Kind in het badwater

What we do NOT want:

- Destroy the IETF thru bureaucracy
- Change bottom-up to top-down
- Focus on "rough code and running consensus"
- Introduce three-piece suits

Sorry Vint

6th Meeting
IETF July 93 13

SURFnet

Timmeren aan de weg

What is being done?

- Stimulate participation in IETF
- RARE charter member Isoc
- Involve government and CEC
- Build good internetworks
- Coordination (Ripe NCC)
- Contribute to Kids Computer World
- Volunteer as IETF local host

6th Meeting
IETF July 93 14

SURFnet

Zoals de waard is

Local host responsibilities:

- preparation
- information
- terminal room
- audio video multicast
- social event

6th Meeting
IETF July 93 15

SURFnet

Het halve werk

Last IETF meetings:

- 650 participants
- 40 Europeans

Guestimate for Amsterdam:

- 500 participants
- 200 Europeans
- Lots of AV-participants

6th Meeting
IETF July 93 16

SURFnet

Een gewaarschuwd mens

- Information server
- Tourist info
- Copyright

6th Meeting
IETF July 93 17

SURFnet

Computerzaal

Sponsors:

- SURFnet bv
- Unisource
- Sun Microsystems (The Netherlands)
- Digital Equipment Corporation (The Netherlands)
- Cisco Systems Europe
- Geveke Electronics BV
- NIKHEF
- Source Information Technology BV
- Tadpole Technology Plc.
- Foundation SURF
- ftp Software, Inc.

6th Meeting
IETF July 93 18

SURFnet

Bedankt!

- SARA (Academic Computer Center Amsterdam)
- RIPE NCC
- Computer Science, Twente University
- DNPAP, Fac. of Electrical Engineering, Delft Univ. of Technology

G.H. Heuser
1977 July 99 19

SURFnet

Waar, wanneer, wat en hoe.

- In diamond lounge
- from 07:00 - 24:00
- dial-in facilities
- SUN, DEC, MAC, Sparcbook
- Drop-in
 - ethemet
 - serial
 - appletalk
- No full time staff
- T-shirts

G.H. Heuser
1977 July 99 20

SURFnet

Geluids- en beeldrondstrooiing

- 2 AVT channels
- Rooms N & M
- Mbone restructured

G.H. Heuser
1977 July 99 21

SURFnet

Uitje

- Training through the Netherlands
- you need a ticket
- you have to pay
- Leaves at 18:10 sharp!

G.H. Heuser
1977 July 99 22

SURFnet

Voor niets gaat de zon op

- SURFnet is a small company
- First time Europe
- Equipment etc. \$ 120.000
- Manpower ca 100 days

G.H. Heuser
1977 July 99 23

SURFnet

Boter

- Lower barrier for European participation
- Involvement
- Acceptance

G.H. Heuser
1977 July 99 24

SURFnet

Adress onder het gras

- IPng decision process
- Focus on Applications and services
- Process formalization
- Acceptance of standards (process)
- RFC publications

Surfnet
RTT July 98 25

SURFnet

Veel plezier

Surfnet
RTT July 98 26

5.2 Advances on ONC

Presented by Chuck McManis/SunSoft

Bio: Chuck McManis is currently a senior staff engineer with SunSoft. His job title is Senior Technologist responsible for the ONC+ portion of Solaris. He graduated from the University of Southern California in 1983 with an Electrical Engineering degree and has worked in both hardware and software positions. He has been with Sun for over six years. Prior to working at Sun, he worked for Intel Corporation on '286 and '386 systems.

Service applications area (SAP)

- APPLICATION INFRASTRUCTURE
- SUPPORT DISTRIBUTED APPLICATIONS
 - Remote procedure call
 - Canonical data representation
 - File access
- DISTRIBUTED SUPPORT FOR HOSTS
 - Security
 - Time
 - Directory

D. Crocker, IETF Amsterdam Plenary

1

ONC & NFS

- HOLE IN THE INTERNET SUITE
 - Look to fill it
 - Consider the alternatives
- PROCESS
 - Consider a technology (Now)
 - Consider adopting it (BOF)

D. Crocker, IETF Amsterdam Plenary

2

Open Network Computing Current Protocols Future Directions

Chuck McManis


Distributed Systems Problem Space

- ❑ Construct Distributed Services
- ❑ Leverage existing knowledge
- ❑ Document Services Protocols
- ❑ Provide a mechanism for evolving protocols.

The Procedure Call Model

- ❑ Expresses a request/response relationship well
 - ➔ - Doesn't express all network semantics
 - ➔ - Doesn't include good model for parallelism
- ❑ Easy to understand, but has some problems.
 - ➔ Procedure calls don't care about calling other procedures
 - ➔ Global variables hold state (poor encapsulation)
 - ➔ Doesn't include a good model for partial failure.

The ONC Distributed System Services



Components

Evolution Of the System

- ❑ RPC Message Protocol - No Changes since rfc1057
- ❑ XDR Data Representation Protocol - No Changes since rfc1014
- ❑ Host Rendezvous protocol (Portmapper) - updated in 1989 to utilize non-IP specific data types for addresses. (no longer returns 'ports')
- ❑ Authentication Protocol - No change from the original published in scheme however additional flavors (Kerberos V4) have been added.
- ❑ Naming Protocols - In 1991 Sun replaced the existing NIS service with the NIS+ naming service.
- ❑ Distributed File Systems - V2 is most prevalent, Some slight changes to make V3.

Areas of Ongoing Research and Development

- ❑ New flavors of authentication.
 - ➔ Kerberos V5
 - ➔ GSS Layer
 - ➔ RSA Publickey system
- ❑ Support for non-procedural operations
 - ➔ One to Many, Many to One interactions.
 - ➔ Asynchronous and batched calls.
 - ➔ Multithreaded clients and servers.
- ❑ File system Issues
 - ➔ Caching and cache coherency
 - ➔ ACLs
 - ➔ non-linear file representation (forked files)

Interoperability present and future

- Connectathon Event - tests existing implementations.
- Successful unencumbered implementations have occurred in the past and are encouraged.
- Clear documentation on all protocols.
- Clear interrelationships among the pieces.

ONC and the IETF

- What we Want
 - Putting these useful protocols (RPC, XDR, and NFS) on to the standards track.
 - Enlisting the community in a collaborative effort to evolve the protocols.
 - Forum for discussing new protocols and their general applicability (NIS+, Authentication, etc)
 - Providing a means for the community to drive the direction of the evolution.

ONC and the IETF

- What we don't want:
 - Endless debates on whether or not these protocols are useful.
 - Requirements to solve a different set of problems. (transactions, pure procedure call semantics, etc)
 - Lack of progress.

Summary

- The ONC base protocols are mature.
 - The ONC protocols have been around for a significant period of time.
 - Implementations have been done on nearly all operating systems.
 - ONC implementations exhibit nearly 100% interoperability.
 - "Duty Free" implementations are available.
- We believe this makes the case for an existing practice standard to be developed.

ONC Network File System

V3 Protocol Revision

Chuck McManis

Why we are doing this revision

- Previous attempts at doing V3 stalled.
 - To Wide in scope (TOPS <-> NFS merge)
 - To Ambitious (NeFS)
 - To Radical
- New circumstances require change *NOW*
 - 64 bit chips and Operating Systems
 - POSIX (ACLs, pathconf)
 - Big window networks (ATM, FDDI)
 - High latency networks (ISDN, PPP)
- Incorporate good research into the "mainstream" product.

What it is

- Minimal revision to NFS (evolutionary not revolutionary)
 - Changes measured against real requirements.
 - Non-contentious among vendors.
 - Short timeframe for implementation.
- A Multivendor Effort
 - Recruited efforts from top ONC licensees.
 - Small group achieving consensus fairly quickly.
 - Avoid previous "kitchen sink" type efforts
- Real
 - Prototypes from vendors will be tested next week at Connectathon.
- Assumes that V4 will be undertaken

What it is specifically

- Current Specification
 - 64 bit sizes and offsets
 - ACCESS RPC for over the wire access control
 - Increased transfer size
 - Exclusive CREATE
 - Reduced Field overloading
 - Support for Weak Cache consistency
 - Fewer GETATTR calls (returned more often)
 - Support for POSIX pathconf
 - General bug fixes.

What it is NOT

- The Kitchen Sink
 - Previous three attempts tried to solve every problem known.
 - Tossed out of this effort: Forked file systems, file versions, generalized error returns, collapsing RMDIR/MKDIR into REMOVE, idempotency procedure, ZERO procedure, *huge* file handles, extended attributes, advisory close, and append mode writes.
- Cache consistency protocol deferred
- ACL support considered a separate issue
- Does not merge TSIG TNFS work
- Not a response to DCE DFS

Milestones

- Current Status
 - Specification and protocol description available @ Connectathon
 - Mount and locking changes prototyped and under review
 - SunSoft clients and servers running internally.
- Next Steps
 - Interoperability testing at Connectathon
 - Release of final specification to general public.
 - Put together plan for V4 development
 - Explore the possibility of doing V4 as an IETF working group.

5.3 Setting Up a Routing Registry in Europe

Presented by Daniel Karrenberg/RIPE

Bio: *Daniel Karrenberg has been active in European networking since 1981, first at Dortmund University, from where he received a graduate degree in Computer Science in 1987 and later at CWI in Amsterdam, The Netherlands. Daniel has helped to create the pan-European network EUnet which at first offered UUCP services but has since become a full Internet service provider. He is currently responsible for the RIPE Network Coordination Centre which coordinates the European part of the Internet and acts as the European regional Internet registry.*

Daniel Karrenberg


RIPE NCC

Topics

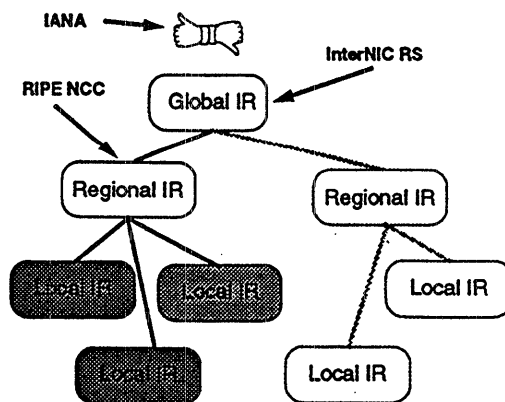
- RIPE Network Coordination Center
- European Routing Registry

The European Internet

- 20 - 25% of the Internet



Internet Registries



European Internet Registries

- EU registry system started a year ago
- today 61 local registries (21 non-provider)
- CIDR blocks from the start
- 193.0.0 - 193.255.255 are European
- ~1000 network numbers assigned / month
- ~4 Bs / month (recovered almost as many)
- bulk of the work done by local registries !!!!

RIPE Database

- vital information for operations
 - networks
 - domains
 - autonomous systems (routing)
 - persons
- operational since 1989
- accessible by WHOIS at whois.ripe.net
- tagged format (make your own reports)
- data (going to be) exchanged automatically with InterNIC

RIPE Database

person: Daniel Karrenberg
address: RIPE Network Coordination Centre
address: Kruislaan 409
address: NL-1098 SJ Amsterdam
address: Netherlands
phone: +31 20 592 5065
fax-no: +31 20 592 5090
e-mail: dfk@ripe.net
nic-hdl: DK58
changed: ripe-dbm@ripe.net 920826
source: RIPE

Other RIPE NCC Activities

- RIPE Support
- Joint Projects

- RIPE Document Store
 - interactive: telnet info.ripe.net
pad 204129004331
 - FTP: ftp.ripe.net
 - GOPHER: gopher.ripe.net
 - WAIS: wais.ripe.net
 - WWW: www.ripe.net
 - Mail: mail-server@ripe.net

RIPE NCC

- core activities
 - funded by (all) European Service Providers
 - reported in public Quarterly Reports
 - Anne Lord
 - Marten Terpstra
 - Daniel Karrenberg

- joint projects
 - funded per project
 - reported in Project Reports (RTRs, RFCs ...)
 - Tony Bates

Organisational Environment

- RIPE
 - Forum for European Internet Service Providers
 - does *NOT* operate a network
 - meets 3 times a year
 - defines RIPE NCC activities

- RARE
 - Association of Research Networks
 - does *NOT* operate a network
 - provides legal and financial framework for NCC
 - provides framework for technical projects

Operational Environment

- many Internet Service Providers (Operators)
- no single backbone
- mesh structure



- routing stability is a big concern!

Routing Registry


The Need:

Routing Stability

Routing Registry

- Strategy: help the SPs to manage external routing
- Tools for Service Providers
 - prconfig - generate router configurations
 - prcheck - check policies
 - prtraceroute - traceroute with policy information
 - prpath - extract possible paths between nets
 - prconn - extract connectivity of a network
- Tools use the Routing Registry

Routing Registry



Routing Registry

- Basic RR Concepts
 - describe routing exchanges between SPs
 - easy to explain
 - no creeping featureism
 - descriptions stay local (no paths!)
 - can make use of BGP (but does not require it)
 - leverages off existing RIPE database

Representation of AS

A simple example



```

as: AS1
<administrivia go here>
as-out: AS2 AS1          # announce all networks in AS1 to AS2
as-in:  AS2 100 AS2 AS1  # accept AS2 and AS1 networks from AS2

as: AS2
<administrivia go here>
as-out: AS1 AS2 AS3      # announce all networks in AS2 and AS3 to AS1
as-in:  AS1 100 AS1      # accept all AS1 networks from AS1
as-out: AS3 AS2 AS1      # announce all networks in AS3 to AS2
as-in:  AS3 100 AS3      # accept AS2 and AS1 networks from AS2

as: AS3
<administrivia go here>
as-out: AS2 AS3          # announce all networks in AS3 to AS2
as-in:  AS2 100 AS1 AS2  # accept AS2 and AS1 networks from AS2
    
```

prtraceroute

```
prtraceroute jolly.nis.garr.it
```

```

1 AS1104 hef-router.nikhef.nl (192.87.45.80) [I]
2 AS1755 Amsterdam-EBS1.Ebone.NET (192.87.4.17) [D1]
3 AS1755 Cern-EBS1.Ebone.NET (192.87.4.10) [I]
4 AS 513 chepl.cern.ch (192.65.185.2) [D1]
5 AS 137 GARR-gw.mi.infn.it (192.12.193.49) [E1]
6 AS 137 GARR-gw.cilea.it (192.12.193.41) [I]
7 AS 137 GARR-gw.cineca.it (192.12.193.21) [I]
8 AS 137 GARR-gw.cnr.it (192.12.193.14) [I]
9 AS 137 jolly.nis.garr.it (192.12.192.5) [I]
    
```

```
Path followed: 1104 1755 513 137
```


```

AS1104 = IXI RBS at NIKHEF
AS1755 = EBONE-INTERNAL
AS 513 = CERN-AS
AS 137 = GARR/INFN IT
    
```

RS Status

Breakdown of observed European ASes

In DB, no RIPE-81 Policy (6.7%) Unknown (1.1%)



Status

- prtraceroute: β
- prconfig: α
- other tools: design

- registry: needs populating

- coordinating with similar projects at MERIT and CIX

- document: RIPE-81
 - ripe/docs/ripe-docs/ripe-81.(txt,ps)
 - open

Plans

- PRIDE Project
 - Policy-based Routing Implementation and Deployment in Europe
 - open

- Route Server at MAE-East

- ----> operational service

5.4 Electronic Cash: Theory To Application

Presented by David Chaum/DigiCash

Bio: *David Chaum is head of the Cryptography Group at the Center for Mathematics and Computer Science (CWI) in Amsterdam. He is also a founder of DigiCash, which develops electronic payment systems. Dr. Chaum received his Ph. D. in Computer Science from the University of California, Berkeley, in 1982 and joined CWI in 1984. He helped to found the International Association for Cryptologic Research and remains active on its board; he also consults internationally on cryptology.*

Cryptography, and public key techniques in particular, are surprisingly powerful and general in theory. Practical schemes for their use have been developed for electronic money with at least the main advantages of paper bank notes. One scheme, not requiring any special hardware, seems well suited to an electronic mail environment. Possible extensions improve security and allow compatibility with systems using credit card-sized IC cards and their upgrade to PDAs.


Electronic Cash: Theory Towards Application

David Chaum
DigiCash/CWI/CAFE


Question

Can a number be a better form of cash than a piece of paper?


Bank-Proof Number-as-Coin System



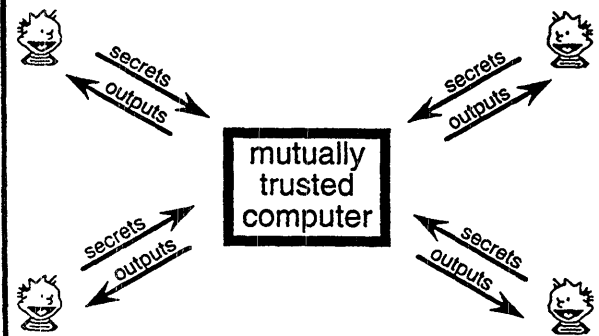
Straw-Man Number-as-Coin System




Theory—Public Key



Science Fiction — DES



Multiparty-Secure Number-as-Coin System




RSA Signatures

The bank can create a number system in which $(f(n)^{1/3})^3 = f(n)$


and in which anyone can raise to the 3 but only the bank can raise to the 1/3;

signature pairs $n, x=f(n)^{1/3}$ cannot be created without the bank—but anyone can verify a pair by testing if $f(n) = x^3$


Bank-Proof Signature-as-Coin System




Bank-Proof Unlinkable Coin System



Paper Cash




SmartCash



Offline Unlinkable Cash

Original paper co-authored with Amos Fiat & Moni Naor



Performance of Electronic Coin Systems

| | Withdrawal | | | | Payment | | |
|--------------|-------------------------|------|--------|-----------------|-------------------------|------|------------------|
| | Multiplications User | Bank | Trans. | Bytes Stored | Multiplications User | Shop | Bytes Transm. |
| Discrete Log | 1820 | 420 | 160 | 256 | 2 | 770 | 275 |
| RSA | 2304 | 2100 | 730 | 256 | 200 | 512 | 275 |

CAFE (Hybrid)



Basic Types of Electronic Money

| | Smart Cards | CAFE | Software Only |
|---------------|----------------|-------------|------------------|
| DES | ? | X | X |
| Public Key | ✓ | ☆ | ✓ |

5.5 JVTOS for Workstations

Presented by Bernhard Plattner/ETH Zurich

Bio: *Bernhard Plattner is a Professor of Computer Engineering at ETH Zurich, where he leads a communication systems research group. He received a diploma in electrical engineering from ETH in 1975 and a Ph. D. in Computer Science in 1983. His research currently focuses on applications of communication systems and higher layer protocols. He is also directing work on high-speed networking, protocol engineering, and multimedia applications, and is interested in real-time computing, process execution monitoring and debugging. He has been active in the design, implementation and operation of the Swiss National Network for Research and Education (SWITCH).*

JVTOS is an advanced teleservice to support cooperative work over distance which allows distributed users to work in a collaborative fashion with multimedia. JVTOS comprises facilities for session management, floor control, multimedia application sharing, telepointing, and audio/video communication. It provides generic support for cooperation-aware multimedia applications.

JVTOS offers services for multimedia collaboration across high-speed networks and is primarily aimed at running in heterogeneous workstation environments comprising different hardware platforms and also different operating and window systems.

This presentation describes the design of JVTOS as well as its implementation on different platforms which is currently under development.

**Sharing Multimedia Applications
Among Heterogeneous Workstations**

Bernhard Platner

Computer Engineering and Networks Laboratory (TIK)
ETH Zürich, Switzerland

RACE

European Communications Research Program

Phase I: Technology
(Fiber -optic Transmission, ATM)

Phase II:
Applications, Pilots,
Communication Experiments

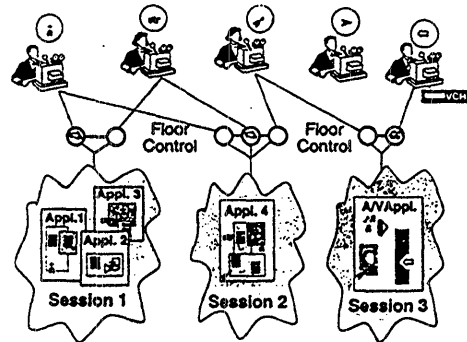
RACE II Project CIO (R2060)

Multimedia Mail Service
Joint Viewing and Tele-Operation Service


Focus:
**Telecooperation Environment for Multimedia
in a Heterogeneous Setup**

Project duration: 1.1.92 - 31.12.94

JVTOS Scenario



The Multimedia Application Sharing Service



Sharing *cooperation-unaware, single-user, multimedia* applications

Application output:
Simultaneous view through
shared windows

Application input:
Subject to *floor control*

**Application Sharing
versus
Cooperation-Aware Applications**

- Many application sharing scenarios
- Diversity of computer applications
- Off-the-shelf applications may be shared
- Users need not learn new applications

Aims

(1) Support and bridge heterogeneous platforms:


- Output on any platform Sun Sparcstation (SunOS/Solaris)
- Input from any platform Apple Macintosh (MacOS)
- Execution on any platform IBM PC (MS-DOS/Windows)

(2) Multimedia applications:


- text/graphics • audio • video (still and moving)

(3) Dynamic user participation:

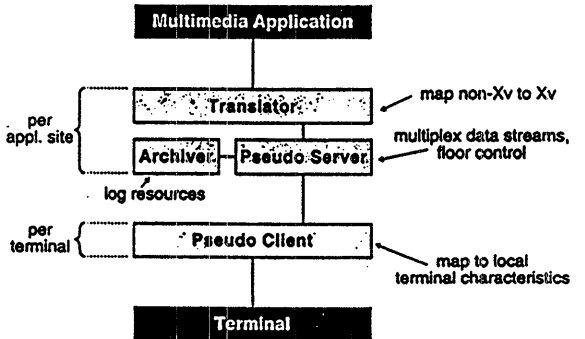
- join an ongoing session
- leave a session



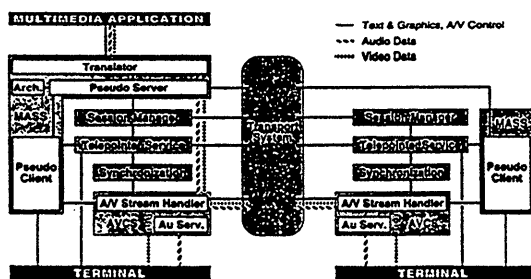
Sharing X Applications



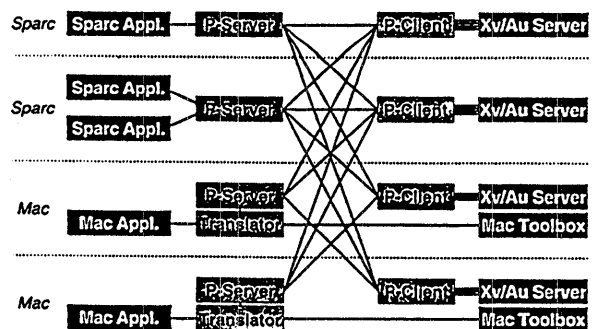
Service Components

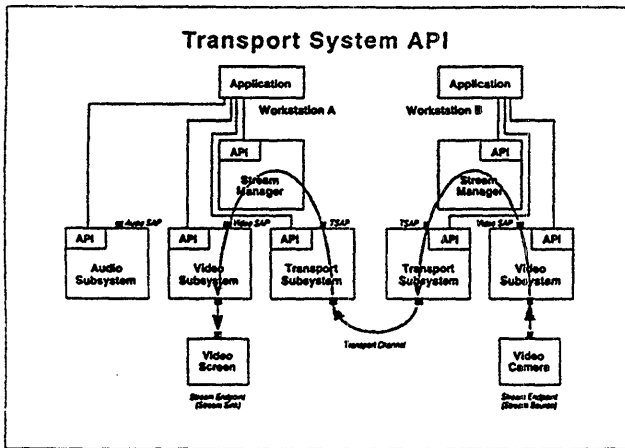


Service Architecture



Heterogeneous Setup





Implementation

- Basic system implemented on Sparc
- P-Client port to Macintosh in progress
- Integration of audio/video in progress
- Detailed design of translators

JVTOS prototype shown at CeBIT '93

Summary

Run Single-User Applications in a Group Context:

- (1) *Heterogeneous Workstation Environment*
- (2) *Support for Multimedia*
- (3) *Dynamic User Participation*

JVTOS = Application Sharing + Picturephone + Telepointing
 Tele-Cooperation Environment for Multimedia in a Heterogeneous Setup

5.6 Service Location

Presented by John Veizades/FTP Software and Scott Kaplan/FTP Software

Bio: John Veizades has been working in the IP community for over ten years and has worked on personal computer based TCP/IP implementations for Xerox and Apple. He has been co-Chair of the Service Location Protocol Working Group since its inception and has been working with the IETF to bring ease of use and configuration to TCP/IP based systems. He is now the engineering manager of the service location work at FTP Software's West Coast Operations.

Bio: Scott Kaplan is the Director of Engineering at FTP Software's West Coast Operations. He is co-Chair of the Service Location Protocol Working Group of the IETF. He is currently studying networking technologies which will make hosts easier to configure and maintain. The ultimate goal of this work is to dispell the myth that grandmothers cannot work computers.

John Veizades

Soon to be of FTP Software, Inc.
veizades@ftp.com

Scott Kaplan

FTP Software, Inc.
scott@ftp.com

Mailing list srv-location@ftp.com

Drafts available on wco.ftp.com

How do users find services in a network?

- Word of mouth
- Administrative procedures
- Wandering the network or the office space

Service Location's approach

- Allows network services to advertise their presence
- Allows users to browse for network services using common terms
- Allows users to broadcast their needs to services
- Allows users to find services by describing their needs to constrain responses
- Solves the dentist office scenario

Service Location Protocol requirements

- No End System configuration
- Multicast support at the network layer
- Not dependent on any network layer addressing scheme
- Scalable to a campus environment (i.e. < 100 nets)
- Scales to a global network with a scaleable directory service

Other related work

- Xerox's Clearinghouse and Network Binding Protocol
- Apple's Name Binding Protocol
- Legato's Resource Administration Platform

Current Solutions

- X.500
- DNS
- ICMP Gateway Discovery
- DHCP/BootP

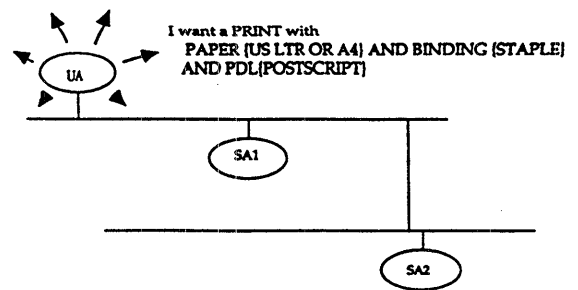
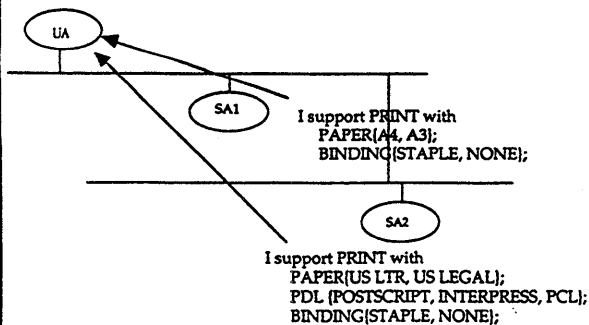

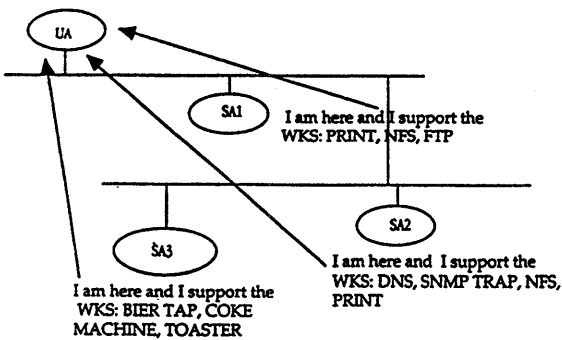

Service Location Comparison

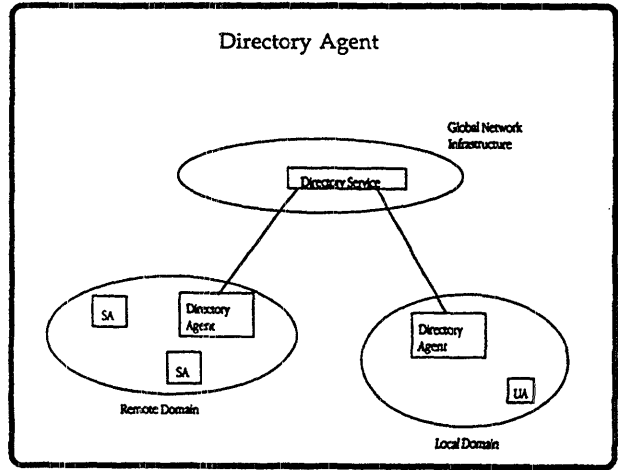
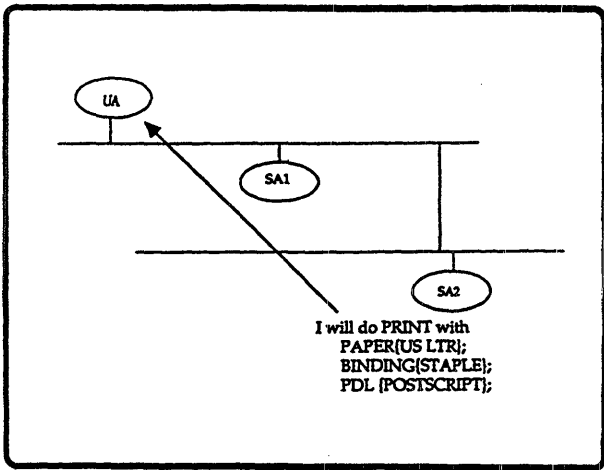
| | X.500 | DNS | ICMP Gateway Discovery | RARP | BootP | Service Location |
|-------------|-------|-----|------------------------------|------|-------|---------------------|
| Simplicity | F | D | A | B | B | B |
| Distributed | B | B | C | D | D | B |
| Timeliness | D | D | A | D | D | A |
| Generality | A | C | F | F | B | B |
| Salability | A | A | F | F | F | B/C |

Better- A B C D F - Worse

A Brief Protocol Description

- User agent multicasts to find services
- Service agent responds with well known service attributes
- User agent requests attribute dictionary from services it is interested in
- User agent makes a detailed service request using elements from the dictionary





5.7 An Experiment in Remote Printing

Presented by Marchall Rose/Dover Beach Consulting

Bio: Marshall T. Rose is Principal at Dover Beach Consulting, Inc., a California-based computer-communications consultancy. He spends half of his time working with clients, and the other half involved in self-supported, openly-available projects. Rose lives with internet-working technologies, such as TCP/IP, OSI, network management, and directory services, as a theorist, implementor, and agent provocateur. He is the author of four professional texts—on Open Systems Interconnection, internet Management, OSI Directory Services, and Electronic Mail. Rose received the Ph. D. degree in Information and Computer Science from the University of California, Irvine, in 1984. His subscriptions to The Atlantic and Rolling Stone Magazine are in good standing.

Although electronic mail is preferable as a means of third-party communication, in some cases it may be necessary to print information, in hard-copy form, at a remote location. The remote output device may consist of a standard line printer, a printer with multiple fonts and faces, a printer that can reproduce graphics, or a facsimile device. Remote output may be accompanied by information that identifies the intended recipient. This presentation describes a technique for “remote printing” using the Internet mail infrastructure. In particular, this focuses on the case in which remote printers are connected to the international telephone network. Furthermore, it describes an experiment in remote printing.

AN EXPERIMENT IN REMOTE PRINTING

Marshall T. Rose
Dover Beach Consulting, Inc.

July 16, 1993

AN ACKNOWLEDGEMENT

Carl Malamud of Internet Multicasting Service provided substantive comments on the design of the experiment.

AN EXPERIMENT IN REMOTE PRINTING

- Suppose you want to send a message to someone, but they don't use e-mail
- The next best thing is to send the message to their printer, i.e., their G3-facsimile device on the international telephone network
- Well, fax modems for computers have been around a long time, but...
why not use the mail infrastructure to route the message to the closest fax modem?

THE ADVANTAGE OF A GENERAL-PURPOSE INFRASTRUCTURE

- This project is about outreach:
integrating the e-mail and fax communities
- Enhancing enterprise-wide access, independent of:
institution: commercial, educational, or government
size: global, regional, or local
making it easier for the "outside world" to talk to an organization's personnel, e.g.,
the sales person;
the university registrar; or,
the (elected) official
who uses fax, but not e-mail

PROCEDURE

- UA constructs template
- User fills-in composition headers and body
- UA detects recipient uses a remote printer:
constructs logical address of remote printer server; and,
constructs a multipart/mixed content, containing:
the original content; and,
a application/remote-printing content
- UA adds submission headers,
and engages submission protocol
- When remote printer server finishes,
it replies to the originator with the "outcome"

WHAT CONTENTS CAN BE PRINTED

- text/plain:
not all character sets may be available
- message/rfc822:
processed recursively
- application/postscript:
evaluated in an "execution safe" environment
- image/tiff:
Tag Image File Format (class F)
- any multipart:
processed recursively,
either on new page (mixed) or same page (parallel)

NAMING, ADDRESSING, AND ROUTING

- o Printer is identified by it's telephone number, e.g., "+1 415 968 2510"
- o Straight-forward mapping to address of remote printer server, e.g., "remote-printer@0.1.5.2.8.6.9.5.1.4.1.tpc.int"
- o MX records used for routing (as usual), e.g.,
 - o.8.6.9.5.1.4.1.tpc.int. IN MX 10 dbc.mtview.ca.us.identifies a relay willing to access any printer with a prefix of "+1 415 968"
(of course, there could be multiple relays)
- o Note that finding an MX RR doesn't imply:
 - o corresponding telephone number is valid; or,
 - o a G3-facsimile device is attachedsimilarly, there may not be an server for every printer...

THE application/remote-printing CONTENT-TYPE

```
<content> ::= <recipient-info> CRLF
              <originator-info>
              [CRLF *(text CRLF)]

<recipient-info> ::= "Recipient" ":" <value> CRLF
<address-info> ::= "Address" ":" <value> CRLF

<address-info> ::= ["Title" ":" <value> CRLF]
                  ["Department" ":" <value> CRLF]
                  ["Organization" ":" <value> CRLF]
                  ["Mailstop" ":" <value> CRLF]
                  ["Address" ":" <value> CRLF]
                  ["Telephone" ":" <value> CRLF]
                  ["Facsimile" ":" <value> CRLF]
                  ["Email" ":" <value> CRLF]

<value> ::= *text
          [CRLF LVSP-char <value> ]
```

AN EXAMPLE

```
To: remote-printer@0.1.5.2.8.6.9.5.1.4.1.tpc.int
From: "John Q. Public" <jpublic@tpd.org>
Date: Sun, 11 Apr 1993 20:34:13 -0800
Subject: Comments on "An Experiment in Remote Printing"
Message-ID: <19930411203413000.456@tpd.org>
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="----- _aaaaaaaa0"

----- _aaaaaaaa0
```

AN EXAMPLE (cont.)

```
Content-Type: application/remote-printing

Recipient: Marshall Ross
Title: Principal
Organization: Dover Beach Consulting, Inc.
Address: 420 Whisman Court
         Mountain View, CA 94043-2186
         US
Telephone: +1 415 968 1052
Facsimile: +1 415 968 2510

Originator: John Q. Public
Organization: The Public Domain
Telephone: +1 801 555 1234
Facsimile: +1 801 555 6789
E-Mail: "John Q. Public" <jpublic@tpd.org>

Any text appearing here would go on the cover-sheet.
```

AN EXAMPLE (cont.)

```
----- _aaaaaaaa0
Content-Type: text/plain; charset="us-ascii"

Here are my comments on your draft.
...

----- _aaaaaaaa0--
```

UNIFORM ADDRESSING

- o Two problems:
 - o can't include e-mail and remote printing recipients in same message; and,
 - o UA has to know about remote printing recipients
- o The solution: give remote printing recipients real e-mail address, using an "rpxploder"
- o When E-mail recipients generate a reply, they can include remote printing recipients!

AN EXAMPLE

To: arose@rpxplode.tpd.org
cc: Arlington Hewes <tpcadmin@dbc.stviev.ca.us>
From: "John Q. Public" <jpublic@tpd.org>
Date: Sun, 11 Apr 1993 20:34:12 -0800
Subject: Comments on "An Experiment in Remote Printing"
Message-ID: <19930411203412000.1236@tpd.org>
MIME-Version: 1.0
Content-Type: text/plain; charset=us-ascii

Here are my comments on your draft.
...

INFRASTRUCTURE

- o "tpc.int." divided into zones,
for each IDDD country code
- o A site wishing to participate contacts zone administrator, e.g.,
1.3.tpc.int. IN SOA ns.ripe.net. hostmaster.ripe.net. (
10005 ;serial
43200 ;refresh every 12 hours
14400 ;retry every 4 hours
345600 ;expire after 4 days
86400);minimum ttl one day
says "hostmaster@ripe.net" is responsible for +31 (NL)

ACCOUNTING AND PRIVACY

- o Minimal auditing:
 - the date the message was received;
 - the From: and Message-ID: headers;
 - the size of the body;
 - the telephone number;
 - any telephony-related information (e.g., call duration); and,
 - any G3-related information (e.g., recipient ID)

FUTURE ISSUES

- o Determination of supported content-types and character sets
- o Introduction of:
 - authentication;
 - integrity;
 - privacy;
 - authorization; and,
 - accounting services
- o Preferential selection of a remote printer server
- o Aggregation of multiple recipients in a single message

INTERESTED IN PARTICIPATING?

- o Server site requirements:
 - a computer on the IP-connected Internet;
 - a fax modem;
 - a phone line;
 - fax spooling software; and,
 - glue software
- o Client site requirements:
 - Internet e-mail connectivity; and,
 - modified UAs and/or rpxploder software

FOR MORE INFORMATION

- o Discussion group: send a note to
tpc-rp-request@aar.net.edu.au
and ask to be added to the
tpc-rp@aar.net.edu.au
list
- o Prototype implementation:
 - site: ftp.ics.ucl.edu
 - area: arose/tpc
 - file: rp.tar.Z
 - mode: binary

ACKNOWLEDGEMENTS

- Sam Leffler: the openly-available FlexFAX package
- Ned Freed/Dan Newman: Innosoft's PMDF for OpenVMS

Chapter 6

Workshop for Working Group Chairs

6.1 Workshop for Working Group Chairs

Presented by David Crocker/Silicon Graphics

Attendees

| | |
|---------------------|------------------------------|
| David Crocker | dcrocker@mordor.stanford.edu |
| Steve DeJarnett | steve@ibmpa.awdpa.ibm.com |
| Luca Delgrossi | luca@ibmpa.awdpa.ibm.com |
| Kjeld Borch Egevang | kbe@craycom.dk |
| Borka Jerman-Blazic | jerman-blazic@ijs.si |
| Mark Laubach | laubach@hpl.hp.com |
| Tony Li | tli@cisco.com |
| Greg Minshall | minshall@wc.novell.com |
| John Stewart | jstewart@cnri.reston.va.us |
| Vladimir Sukonnik | sukonnik@process.com |
| James Watt | james@newbridge.com |

Workshop for Working Group Chairs

D. Crocker

Silicon Graphics
dcrocker@sgi.com / +1 415 390 1804

Hello

- INTRODUCTION
 - What are we doing here, at this hour?
 - There is only rough consensus about the process
- AGENDA
 - I. IETF structure
 - II. Formal process
 - III. The inner working group
 - IV. Conflict resolution

D. Crocker, Workshop for Working Group Chairs

2

The need for working group chair training

- IETF LARGE, DIVERSE
 - Process increasingly formal
 - No voting means (very) rough consensus
- DIFFICULTY MAKING PROGRESS AND BEING FAIR
 - Listen to all points of view
 - Keep working group focus
- CHAIRS OFTEN UNCLEAR ABOUT LIMITATIONS AND AUTHORITIES

D. Crocker, Workshop for Working Group Chairs

3

Documents

- THE INTERNET STANDARDS PROCESS (RFC 1310)
- IETF WORKING GROUP GUIDELINES AND PROCEDURES (IN PROGRESS)
- ["EVOLVING THE SYSTEM" IN INTERNET SYSTEM HANDBOOK, LYNCH & ROSE, EDS.]
- ["MAKING STANDARDS THE IETF WAY" IN ACM STANDARDSVIEW, SUMMER 1993; REPRINTED IN CONNEXIONS, AUGUST 1993.]

D. Crocker, Workshop for Working Group Chairs

4

I. IETF Structure

| | |
|------------------|--|
| ISOC | Internet Society legal cover |
| IAB | Internet Architecture Board Design cohesion, process appeals, IETF liaisons |
| IETF secretariat | Staff support |
| IESG | Internet Engineering Steering Group IETF oversight |
| AD | Area director Oversight for specific working groups |
| WG chair | Manage a working group to a productive end |
| Working group | The people who do the work |

D. Crocker, Workshop for Working Group Chairs

5

Working group roles

- CHAIR OVERSEES ENTIRE PROCESS, BUT:

| | |
|-------------|---|
| Facilitator | Process management, things fair, focused, on time |
| Judge | Evaluation of technical options and driver towards "right" choice |
| Scribe | Record-keeper and editor of documents |
- WORKING GROUP IS JURY, PROVIDING IDEAS, REVIEW, CONSENSUS
- DESIGN TEAM IS PRIMARY ADVOCATE AS SELF-SELECTING GROUP WITH COMMON VISION, PROVIDING CORE EFFORT

D. Crocker, Workshop for Working Group Chairs

6

Formal Process

- FORMAL LABELS FOR A SPECIFICATION
- DEVELOPMENTAL STEPS
- ACCEPTANCE CRITERIA

Formal labels

- Internet draft:** no official standing, fluid working document
- Proposed Standard:** stable spec, no known errors, might have implementation
- Draft Standard:** multiple, interoperable implementations testing all functionality
- Internet Standard:** field experience and clear community acceptance (and use)

IETF acceptance criteria

- Competence:** technically sound
- Constituency:** providers & users
- Coherence:** clear writing
- Consensus:** rough but clear

II. Developmental steps

0. BIRDS OF A FEATHER (BOF)
- "Market research" to determine interest and ability to pursue topic
 - Optional, one-shot meeting
1. CHARTER
- Role:** Public announcement & project management plan
- Scope:** What is to be pursued
- Approach:** How it will be pursued
- Product:** What will be delivered
- Checkpoints:** Milestones and dates

II. Developmental steps

2. DOCUMENT SPECIFICATION
- Clarity of purpose
 - Clarity of writing
 - Clarity of solution
3. WG CONSENSUS
- Clearly dominant agreement
 - Diversity of opinion about solution may be resolved by agreement to make *some* decision
 - Agreement about parts may permit eventual agreement about whole

II. Developmental steps

4. AREA DIRECTOR APPROVAL
- Technical review
 - Process review
 - Independent review when results of wg in question
5. SUBMISSION TO IESG
- Via secretariat & AD

II. Developmental steps

6. LAST CALL

- Request for final feedback from IETF
- Intended to detect major errors in process or content that might have slipped through cracks
- Not intended as formal, full review

7. IESG REVIEW (& APPROVAL)

- May conduct independent review

II. Developmental steps

(7.5) IAB CONFLICT RESOLUTION

- If formal challenge not resolved by IESG

8. RFC PUBLICATION

- RFC editor has publication criteria

III. The inner working group

- THE LIVES OF A CHAIR
- GROUP STYLE
- GROUP ROLES
- DEVELOPMENTAL PHASES (PROBLEM SOLVING 101)
- VENUES
- DEBATE
- CONFLICT MANAGEMENT

The lives of a chair

- HOW TO KEEP FROM BEING SAT ON
 - Agenda & schedule
 - Adequate debate, but not more than that
 - Maintain clear focus
 - Rehash *only* if constructive *and* working group desires
- PROACTIVE MANAGEMENT
 - Maintain pressure for forward progress
 - Escalate to IETF management when progress stalled

Working group style

- FREE-FLOWING
 - Cohesive group
 - Clear purpose
- TIGHTLY-MANAGED
 - Complex topic
 - Group diversity
 - Major differences in philosophy

WG management roles

- Facilitator : ensuring fairness and a thorough airing of views and alternatives
- Judge : evaluation of choices and movement towards choice
- Scribe: keeping track of things
- Design team: Primary advocates for the core effort, when wg diverse & topic complex; **must** work to keep wg consensus
- Working group: Jury & other contributors

Problem solving 101

- **PROBLEM STATEMENT**
- **SOLUTION EXPLORATION**
- **SOLUTION ADOPTION**
- **SPECIFICATION REFINEMENT**

Discussion & decision venues

- **EMAIL**
 - International participation
 - Inefficient, but extensive
 - The real place for consensus
 - Can be run as "meetings"
- **FACE-TO-FACE MEETINGS**
 - Well-advertised ahead of time
 - Inherently restricted attendance
 - Limited time
 - Need for clear agenda and crisp management

Email vs. Meetings (One person's perspective)

- **WG RESULTS MUST SHOW APPROVAL BASED ON ENTIRE WORKING GROUP**
- **MEETINGS HAVE LIMITED ATTENDANCE**
- **TREAT MEETINGS AS "STRONG INDICATOR" PRIMA FACIE BASIS FOR DECISIONS**
- **ENSURE VERIFICATION THROUGH EMAIL**

Debate

- **CAN CLARIFY PURPOSE, IMPLICATIONS, ALTERNATIVES**
- **CAN TEAR THE GROUP APART**
- **MUST BE TOLERATED AND EVEN ENCOURAGED, UNTIL RESOLUTION OR IMPASSE**

IV. Conflict Resolution

- **PREFERABLE TO SOLVE WITHIN WORKING GROUP**
 1. Conflict types
 2. Timing of objections
- **OFTEN CAN'T**
 3. Chain of appeal

1. Conflict types

- **TECHNICAL**
 - Specific detail: minor vs. show-stopper
 - Basic philosophies: rarely resolved
- **PROCESS**
 - Unfair practice: usually claim against wg chair
 - Topic missed: oops. (showstopper?)

2. Timing of objections

- **TECHNICAL SHOWSTOPPERS WELCOME ANYTIME**
- **SMALL DETAILS WELCOME ONLY AT TIME WG COVERS THE SUBJECT**
- **PHILOSOPHICAL DEBATE WELCOME ONLY AT TIME WG MAKING DECISIONS ABOUT APPROACH**
- **UNFAIR PRACTICE COMPLAINTS ALLOWED WHENEVER INFRACTION FELT**
- **WG MAY ALLOW TOPIC TO BE RE-OPENED IF WG FEELS ISSUE COMPELLING OR NEW ALTERNATIVE INTRIGUING.**

D. Crocker, Workshop for Working Group Chairs

25

3. Chain of appeal

- **WG CHAIR**
- **AREA DIRECTOR**
- **IESG (PLENARY)**
- **IAB**

D. Crocker, Workshop for Working Group Chairs

26

If you can keep your head when those around you...

- **Most IETF members are remarkably well-intentioned**
- **Differences happen**
 - **Tempers often flare, but then settle down**
 - **Not all differences can be settled**
 - **When minority view clearly will not sway working group, respect opinion, but move on**
- **Ask questions**
- **Make it happen!**

D. Crocker, Workshop for Working Group Chairs

27

Appendix A

Attendees

Matti Aarnio
NORDUnet
Phone: +358 21 6338668
Fax: +358 21 329059
mea@nic.nordu.net

George Abe
Infonet Services Corporation
Phone: 310-335-2867
Fax: 310-335-2876
abe@infonet.com

Roland Acra
cisco Systems, Europe
Phone: +33 1 60 92 20 00
Fax: +33 1 69 28 83 26
acra@cisco.com

Chris Adie
Edinburgh University
Phone: +44 31 650 3363
Fax: +44 31 662 4809
C.J.Adie@edinburgh.ac.uk

Masuma Ahmed
Bell Communications Research
Phone: 908-758-2515
Fax: 908-758-4192
mxa@sabre.bellcore.com

Kannan Alagappan
Digital Equipment Corporation
Phone: 508-486-5953
Fax: 508-486-5279
kannan@DSMAIL.ENET.DEC.COM

Mohammad Alaghebandan
3Com Corporation
Phone: 408-764-5984
mohammad_alaghebandan@3com.com

Steve Alexander
Lachman Technology, Inc.
Phone: 708-505-9555
Fax: 708-505-9574
stevea@lachman.com

Nick Alfano
MPR Teltech, Ltd.
Phone: 604-293-5342
Fax: 604-293-5787
alfano@mpr.ca

James Allard
Microsoft Corporation
Phone: 206-882-8080
Fax: 206-936-7329
jallard@microsoft.com

Claudio Allocchio
National Institute for
Nuclear Physics - Italy
Phone: +39 40 3758523
Fax: +39 40 226338
Claudio.Allocchio@elettra.trieste.it

Bernt Allonen
Unisource Business Networks
Phone: +46 8 688 2181
Fax: +46 8 688 2006
bal@tip.net

Harald Alvestrand
SINTEF DELAB
Phone: +47 7 59 70 94
Fax: +47 7 53 25 86
Harald.Alvestrand@uninett.no

Frederik Andersen
Dansk Data Elektronik A/S
Phone: +45 42 84 50 11
Fax: +45 42 84 52 20
fha@dde.dk

Per Andersson
Chalmers University of
Technology
Phone: +46 31 772 17 81
Fax: +46 31 772 36 63
pa@cdg.chalmers.se

Michael Anello
XLNT Designs, Inc.
Phone: 619-487-9320
Fax: 619-487-9768
mike@xlnt.com

David Arneson
Cabletron Systems, Inc.
Phone: 603-337-5038
Fax: 603-424-8121
arneson@ctron.com

Arun Arunkumar
3Com Corporation
Phone: 408-764-5104
Fax: 408-764-5002
nak@3com.com

Toshiya Asaba
Internet Initiative Japan
Phone: +81 3 3580 3781
Fax: +81 3 3580 3782
asaba@iiij.ad.jp

Josee Auber
Hewlett-Packard
Phone: +33 76 62 11 93
Fax: +33 76 62 12 86
Josee_Auber@hpgnd.grenoble.hp.com

Robert Austein
Epilogue Technology
Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sra@epilogue.com

Anders Baardsgaad
University of Tromso
Phone: +47 834 4105
Fax: +47 834 4100
anders@cc.uit.no

Cynthia Bagwell
The MITRE Corporation
Phone: 703-883-5438
Fax: 703-883-7142
cbagwell@gateway.mitre.org

Dennis Baker
Wellfleet Communications, Inc.
Phone: 508-436-3849
Fax: 508-670-8760
dbaker@wellfleet.com

Fred Baker
Advanced Computer
Communications
Phone: 805-685-4455
Fax: 805-685-4465
fbaker@acc.com

John Ballard
Microsoft Corporation
Phone: 206-936-4688
Fax: 206-936-7329
jballard@microsoft.com

Tony Ballardie
University College London
Phone: +44 71 387 7050
Fax: +44 71 387 1397
A.Ballardie@cs.ucl.ac.uk

Paul Barker
University College London
Phone: +44 71 380 7366
Fax: +44 71 387 1397
p.barker@cs.ucl.ac.uk

Jim Barnes
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-3159
barnes@xylogics.com

Julian Bates
Xylogics International, Ltd.
Phone: +44 908 222112
Fax: +44 908 222115
bates@xylogics.com

Tony Bates
RIPE Network Coordination
Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
tony@ripe.net

Nutan Behki
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
Nutan_Behki@qmail.newbridge.com

Axel Belinfante
University of Twente
Phone: +31 53 893774
Fax: +31 53 333815
Axel.Belinfante@cs.utwente.nl

Lou Berger
Bolt Beranek and Newman
Phone: 703-284-4651
Fax: 703-284-2766
lberger@bbn.com

Vincent Berkhout
University of Twente
Phone: +31 53 893742
Fax: +31 53 333815
berkhout@cs.utwente.nl

Tim Berners-Lee
CERN
Phone: +41 22 767 3755
Fax: +41 22 767 7155
timbl@info.cern.ch

Per Bilse
Internet Consult
Phone: +45 3117 1233
Fax: +45 3116 1607
bilse@ic.dk

Jim Binkley
Intel
Phone: 503-696-5054
jrb@ibeam.intel.com

Zavisa Bjelogric
Intecs Sistemi SpA
Phone: +39 6 41886244
Fax: +39 6 4191667
zavisa@roma.intecs.it

Robert Blokzijl
RIPE/NIKHEF
Phone: +31 20 5925102
Fax: +31 20 5925155
K13@nikhef.nl

David Borman
Cray Research, Inc.
Phone: 612-683-5571
Fax: 612-683-5599
dab@cray.com

Carsten Bormann
Technische Universitaet Berlin
Phone: +49 30 314 24315
Fax: +49 30 314 25156
cabo@cs.tu-berlin.de

Erik-Jan Bos
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
erik-jan.bos@surfnet.nl

Rebecca Bostwick
Lawrence Livermore National
Laboratory
Phone: 510-422-4017
Fax: 510-422-1482
bostwick@es.net

Luc Boulianne
McGill University
Phone: 514-398-6698
Fax: 514-398-3883
lucb@cs.mcgill.ca

Jim Bound
Digital Equipment Corporation
Phone: 508-486-5180
Fax: 508-486-7417
bound@zk3.dec.com

Piet Bovenga
 University of Nijmegen
 Phone: +31 80 617945
 Fax: +31 80 617979
 p.bovenga@uci.kun.nl

Rich Bowen
 IBM Corporation
 Phone: 919-543-9851
 rkb@ralvm11.vnet.ibm.com

Robert Braden
 Information Sciences Institute
 Phone: 310-822-1511
 Fax: 310-823-6714
 braden@isi.edu

Stefan Braun
 TELES GmbH
 Phone: +49 30 88473370
 Fax: +49 30 8852741
 smb@cs.tu-berlin.de

Michael Brescia
 Bolt Beranek and Newman
 Phone: 617-873-3662
 Fax: 617-873-3315
 brescia@bbn.com

Thomas Brisco
 Rutgers University
 Phone: 908-932-5020
 Fax: 908-932-2968
 brisco@pilot.njin.net

Ronald Broersma
 Naval Command, Control and
 Ocean Surveillance Center
 Phone: 619-553-2245
 Fax: 619-553-6811
 ron@nosc.mil

Caralyn Brown
 Wellfleet Communications, Inc.
 Phone: 508-436-3835
 Fax: 508-670-8760
 cbrown@wellfleet.com

Suzy Brown
 General Magic
 Phone: 415-966-6286
 Fax: 415-965-1830
 suzy_brown@genmagic.com

Tracy Brown
 Bell Communications Research
 Phone: 908-758-2107
 Fax: 908-758-4177
 tacox@mail.bellcore.com

J. Nevil Brownlee
 University of Auckland
 Phone: +64 9 373 7599
 Fax: +64 9 373 7425
 nevil@ccu1.aukuni.ac.nz

Theodore Brunner
 Bell Communications Research
 Phone: 201-829-4578
 Fax: 201-829-2504
 tob@thumper.bellcore.com

Thomas Brunner
 SWITCH
 Phone: +41 12618179
 Fax: +41 12618133
 brunner@switch.ch

Steve Buchko
 Newbridge Networks Corporation
 Phone: 603-591-3600
 Fax: 603-591-3680
 stevebu@newbridge.com

John Burnett
 Adaptive Corporation
 Phone: 415-780-7768
 Fax: 415-780-4323
 jlb@adaptive.com

Ramon Caceres
 Matsushita Information
 Technology Laboratory
 Phone: 609-734-0800
 Fax: 609-987-8827
 ramon@mitl.research.panasonic.com

Kursat Cagiltay
 Middle East Technical
 University
 Phone: +90 4 210 1120
 Fax: +90 4 210 1120
 kursat@vm.cc.metu.edu.tr

Ross Callon
 Wellfleet Communications, Inc.
 Phone: 508-436-3936
 Fax: 508-670-8760
 rcallon@wellfleet.com

Peter Cameron
 Xylogics International, Ltd.
 Phone: +44 908 222112
 Fax: +44 908 222115
 cameron@xylint.co.uk

C. Allan Cargille
 University of Wisconsin
 Phone: 608-262-5084
 Fax: 608-262-9777
 allan.cargille@cs.wisc.edu

Brian Carpenter
 CERN
 Phone: +41 22 767 4967
 Fax: +41 22 767 7155
 brian@dxcern.cern.ch

Jeff Case
 University of Tennessee
 Phone: 615-573-1434
 Fax: 615-573-9197
 case@cs.utk.edu

Vinton Cerf
 Corporation for National
 Research Initiatives
 Phone: 703-620-8990
 Fax: 703-620-0913
 vcerf@cnri.reston.va.us

George Chang
 Bell Communications Research
 Phone: 908-699-3879
 Fax: 908-336-2969
 gkc@ctt.bellcore.com

A. Lyman Chapin
 Bolt Beranek and Newman
 Phone: 617-873-3133
 Fax: 617-873-3766
 lyman@bbn.com

Chris Chiotasso
 Ungermann-Bass, Inc.
 Phone: 508-691-6301
 Fax: 508-687-4442
 chris@andr.ub.com

Jodi-Ann Chu
 University of Hawaii
 Phone: 808-956-2400
 Fax: 808-956-2412
 jodi@uhunix.uhcc.hawaii.edu

George Clapp
 Ameritech Advanced Data
 Services
 Phone: 708-248-6507
 Fax: 708-248-6038
 clapp@ameris.center.il.ameritech.com

Cynthia Clark
 Corporation for National
 Research Initiatives
 Phone: 703-620-8990
 Fax: 703-620-0913
 cclark@cnri.reston.va.us

Henry Clark
 OARnet
 Phone: 614-292-6483
 Fax: 614-292-7168
 henryc@oar.net

Stuart Clayman
 University College London
 Phone: +44 71 387 7050
 sclayman@cs.ucl.ac.uk

Les Clyne
Rutherford Appleton Laboratory
Phone: +44 235 446459
Fax: +44 235 446251
l.clyne@jnt.ac.uk

Richard Colella
National Institute of
Standards and Technology
Phone: 301-975-3627
Fax: 301-590-0932
colella@nist.gov

James Conklin
EDUCOM
Phone: 202-872-4200
Fax: 202-872-4318
jbc@bitnic.educom.edu

David Conrad
Internet Initiative Japan
Phone: +81 3 3580 3781
Fax: +81 3 3580 3782
davidc@iiij.ad.jp

Robert Cooney
Naval Computer and
Telecommunications Station
Phone: 202-685-1050
Fax: 202-433-3188
cooney@wnyose.nctsw.navy.mil

Thomas Cordetti
DigiBoard, Inc.
Phone: 612-943-0405
Fax: 612-943-0803
tomc@digibd.com

Al Costanzo
AKC Consulting
Phone: 908-298-9000
al@akc.com

Stephen Coya
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
scoya@cnri.reston.va.us

David Crocker
Silicon Graphics, Inc.
Phone: 415-390-1804
Fax: 415-962-8404
dcrocker@mordor.stanford.edu

Stephen Crocker
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
crocker@tis.com

Dave Cullerot
Cabletron Systems, Inc.
Phone: 603-337-5003
Fax: 603-424-8121
cullerot@ctrtron.com

Walid Dabbous
INRIA
Phone: +33 93 65 77 18
Fax: +33 93 65 77 65
Walid.Dabbous@sophia.inria.fr

Glen Daniels
ELF Communications
Phone: 617-576-2856
Fax: 617-576-6715
gub@elf.com

Jonathan Davar
SynOptics Communications, Inc.
Phone: 408-764-1445
Fax: 408-764-1669
jdavar@synoptics.comm

Megan Davies Walnut
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
mwalnut@cnri.reston.va.us

James Davin
Bell Communications Research
Phone: 201-829-4798
Fax: 201-829-2504
davin@thumper.bellcore.com

Geert Jan de Groot
Philips Consumer Electronics
Phone: +31 40 797408
Fax: +31 40 732340
geertj@ica.philips.nl

Stephen Deering
Xerox Corporation
Phone: 415-812-4839
Fax: 415-812-4471
deering@parc.xerox.com

Steve DeJarnett
IBM Corporation
Phone: 415-694-3896
Fax: 415-694-3073
steve@ibmpa.awdpa.ibm.com

Marty Del Vecchio
Shiva Corporation
Phone: 617-270-8300
Fax: 617-270-8599
marty@shiva.com

Luca Delgrossi
IBM European Networking Center
Phone: +49 6221 594330
Fax: +49 6221 593300
luca@ibmpa.awdpa.ibm.com

Peter Deutsch
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8134
peterd@bunyip.com

Taso Devetzis
Bell Communications Research
Phone: 201-829-4764
Fax: 201-829-2504
devetzis@bellcore.com

Thomas DeWitt
General Services
Administration
Phone: 202-501-0139
Fax: 202-208-5555
tdewitt@osi.ncsl.nist.gov

Maria Dimou-Zacharova
CERN
Phone: +41 22 767 3356
Fax: +41 22 767 7155
dimou@dxcern.cern.ch

Tim Dixon
RARE
Phone: +31 20 639 1131
Fax: +31 20 639 3289
dixon@rare.nl

Kurt Dobbins
Cabletron Systems, Inc.
Phone: 603-337-5003
Fax: 603-424-8121
dobbins@ctrn.com

Reinhard Doelz
Universitaet Basel
Phone: +41 61 264 2247
Fax: +41 61 267 2078
embnet@comp.bioz.unibas.ch

Bob Downs
Combinet
Phone: 408-522-9020
Fax: 408-732-5479
bdowns@combinet.com

Kishan Dudkikar
SKD Systems
Phone: 301-984-7657
Fax: 301-816-2982
kishan@icm1.icp.net

Ian Duncan
McGill University
Phone: 514-398-3710
Fax: 514-398-4658
id@cc.mcgill.ca

Jeffrey Dunn
Naval Research Laboratory
Phone: 202-767-3870
Fax: 202-404-7453
dunn@neptune.nrl.navy.mil

Francis Dupont
INRIA
Phone: +33 1 39 635213
Fax: +33 1 39 635330
francis.dupont@inria.fr

Pierre Dupont
Motorola, Inc.
Phone: 604-241-6243
Fax: 604-241-6243
dupont@mdd.comm.mot.com

Steve Dusse
RSA Data Security, Inc.
Phone: 415-595-8782
Fax: 415-595-1873
spock@rsa.com

Tom Easterday
CICNet, Inc.
Phone: 313-998-6285
Fax: 313-998-6105
tom@cic.net

Toerless Eckert
University Erlangen-Nuernberg
Informatik 4
Phone: +49 9131 85 7278
Fax: +49 9131 39388
Eckert@informatik.uni-erlangen.de

Bill Edwards
US Sprint
Phone: 913-967-5334
Fax: 913-967-2526
william.x.edwards@sprint.sprint.com

Kjeld Borch Egevang
Cray Communications
Phone: +45 44 53 01 00
kbe@craycom.dk

Ed Ellesson
IBM Corporation
Phone: 919-254-4115
Fax: 919-254-5410
ellesson@vnet.ibm.com

Alan Emtage
Bunyip Information Systems
Phone: 514-875-8611
Fax: 514-875-8134
bajan@bunyip.com

David Engel
Optical Data Systems, Inc.
Phone: 214-234-6400
Fax: 214-234-4059
david@ods.com

Robert Enger
Advanced Network and
Services, Inc.
Phone: 703-758-7722
Fax: 703-758-7717
enger@reston.ans.net

Urs Eppenberger
SWITCH
Phone: +41 1 261 8112
Fax: +41 1 261 8133
eppenberger@switch.ch

Hans Eriksson
Swedish Institute of
Computer Science
Phone: +46 8 752 1527
Fax: +46 8 751 7230
hans@sics.se

Michael Erlinger
Harvey Mudd College
Phone: 714-693-2617
Fax: 714-621-8465
mike@jarthur.claremont.edu

Julio Escobar
Bolt Beranek and Newman
Phone: 617-873-4579
Fax: 617-873-3776
jescobar@bbn.com

Deborah Estrin
University of Southern
California
Phone: 310-822-1511
Fax: 310-823-6714
estrin@usc.edu

Dino Farinacci
cisco Systems
Phone: 415-688-4696
Fax: 415-688-4575
dino@cisco.com

Stefan Fassbender
IBM Corporation
Phone: +49 6221 594 446
Fax: +49 6221 593 300
stf@easi.net

Mark Fedor
Performance Systems
International, Inc.
Phone: 518-283-8860
Fax: 518-283-8904
fedor@psi.com

Dennis Ferguson
Advanced Network and
Services, Inc.
Phone: 914-789-5367
Fax: 914-789-5310
dennis@ans.net

Eric Fleischman
Boeing Computer Services
Phone: 206-957-5334
Fax: 206-957-5050
ericf@act.boeing.com

Peter Ford
Los Alamos National Laboratory
Phone: 505-665-0058
Fax: 505-665-4939
peter@goshawk.lanl.gov

James Forster
cisco Systems
Phone: 415-688-7854
Fax: 415-688-8282
forster@cisco.com

Jill Foster
University of
Newcastle upon Tyne
Phone: +44 91 222 8250
Fax: +44 91 222 8765
Jill.Foster@newcastle.ac.uk

Osten Franberg
Ellementeal AB
Phone: +46 8 7274121
Fax: +46 8 6478276
euaokf@eua.ericsson.se

Paul Francis
Bell Communications Research
Phone: 201-829-4484
Fax: 201-829-2504
Francis@thumper.bellcore.com

Barbara Fraser
CERT Coordination Center
Phone: 412-268-5010
Fax: 412-268-6989
byf@cert.org

David Fresquez
IBM Corporation
Phone: 919-543-3657
Fax: 919-543-4517
fresquez@vnet.ibm.com

Dan Frommer
Digital Equipment Corporation
Phone: +972 2 782 551
Fax: +972 2 782 373
dan@jeremy.enet.dec.com

Klaus Fueller
Lichtenberg-Schule
Phone: +49 561 405366
klausf@osgo.ks.he.schule.de

Shoji Fukutomi
Furukawa Electoric
Phone: +81 463 23 1045
Fax: +81 463 23 0416
fuku@furukawa.co.jp

Vince Fuller
BARRNet
Phone: 415-723-6860
Fax: 415-723-0010
vaf@stanford.edu

Jim Fullton
MCNC - CNIDR
Phone: 919-248-1499
fullton@cnidr.org

Peter Furniss
Peter Furniss Consultants
Phone: +44 81 313 1833
Fax: +44 81 313 1833
p.furniss@ulcc.ac.uk

James Galvin
Trusted Information Systems
Phone: 301-854-6889
Fax: 301-854-5363
galvin@tis.com

Kevin Gamiel
MCNC - CNIDR
Phone: 919-248-1499
Fax: 919-248-1405
kgamiel@cnidr.org

Joan Gargano
University of California,
Davis
Phone: 916-752-2591
Fax: 916-752-9017
jcgargano@ucdavis.edu

Eugene Geer
Bell Communications Research
Phone: 908-758-2149
Fax: 908-758-4184
ewg@cc.bellcore.com

Ben Geerlings
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
ben.geerlings@surfnet.nl

Tony Genovese
Lawrence Livermore National
Laboratory
Phone: 415-422-4017
Fax: 415-422-1482
genovese@es.net

Robert Gilligan
Sun Microsystems, Inc.
Phone: 415-336-1012
Fax: 415-336-6015
Bob.Gilligan@Eng.Sun.Com

Anders Gillner
KTHnoc/SUNET
Phone: +46 8 7906518
Fax: +46 8 241179
awg@sUNET.se

David Ginsburg
Alcatel SEL
Phone: +49 711 8213146
Fax: +49 711 8212470
ginsb@us-es.sel.de

Joseph Godsil
National Center for
Supercomputing Applications
Phone: 217-333-1989
Fax: 217-244-1987
jgodsil@ncsa.uiuc.edu

Mike Goguen
SynOptics Communications, Inc.
Phone: 408-764-1289
Fax: 408-764-1299
goguen@synoptics.com

Tim Goodwin
PIPEX, Ltd.
Phone: +44 223 250100
Fax: +44 223 250101
tim@pipex.net

Ramesh Govindan
Bell Communications Research
Phone: 201-829-4406
Fax: 201-829-4504
rxg@thumper.bellcore.com

Christoph Graf
SWITCH
Phone: +41 1 268 1520
Fax: +41 1 268 1568
graf@switch.ch

Marcel Graf
IBM European Networking Center
Phone: +49 6221 59 4532
Fax: +49 6221 59 3300
graf@dhdibm1.bitnet.vm.gmd.de

Judith Grass
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
grass@cnri.reston.va.us

Richard Graveman
Bell Communications Research
Phone: 908-699-4611
Fax: 908-336-2943
rfg@ctt.bellcore.com

Terry Gray
University of Washington
Phone: 206-543-3880
Fax: 206-543-3909
gray@cac.washington.edu

Ron Greve
University of Twente
Phone: +31 8373 15436
Fax: +31 53 333815
rgreve@cs.utwente.nl

Phillip Gross
Advanced Network and
Services, Inc.
Phone: 914-789-5335
Fax: 914-789-5310
pgross@ans.net

Chris Gunner
Digital Equipment Corporation
Phone: 508-486-7792
Fax: 508-486-5279
gunner@dsmail.lkg.dec.com

Neil Haller
Bell Communications Research
Phone: 201-829-4478
Fax: 201-829-2504
nmh@thumper.bellcore.com

Joel Halpern
Network Systems Corporation
Phone: 612-424-1606
Fax: 612-424-1736
jmh@network.com

Jari Hamalainen
Nokia Mobile Phones
Phone: +358 31 165896
Fax: +358 31 165888
jah@rctre.nokia.com

Mark Handley
University College London
Phone: +44 71 387 7050
Fax: +44 71 380 1397
mhandley@cs.ucl.ac.uk

Patrick Hanel
Nokia Telecommunications
Phone: +358 51044927
Fax: +358 51044952
hanel@yoyodyne.trs.ntc.nokia.com

Craig Haney
 ICMNET
 Phone: 703-904-2084
 Fax: 703-904-2292
 craig@icp.net

Alf Hansen
 SINTEF DELAB
 Phone: +47 7 592982
 Fax: +47 7 532586
 Alf.Hansen@delab.sintef.no

Jan Hansen
 Agder Research Foundation
 Phone: +47 370 42555
 Fax: +47 370 40696
 Hansen@teknologi.agderforskning.no

Raymond Harder
 Consultant
 Phone: 909-983-4713
 rharder@eis.calstate.edu

Susan Hares
 Merit Network, Inc.
 Phone: 313-936-2095
 Fax: 313-747-3745
 skh@merit.edu

Ken Hayward
 Bell Northern Research
 Phone: 613-763-4042
 Fax: 613-763-2626
 crm57d@bnr.ca

Denise Heagerty
 CERN
 Phone: +41 22 767 49 75
 Fax: +41 22 767 71 55
 denise@dxcoms.cern.ch

Geert Heijenck
 University of Twente
 Phone: +31 53 892834
 Fax: +31 53 333815
 heijenck@cs.utwente.nl

Maria Heijne
 SURFnet bv
 Phone: +31 30 310290
 Fax: +31 30 340903
 maria.heijne@surfnet.nl

Juha Heinanen
 Telecom Finland
 Phone: +358 49 500958
 Fax: +358 31 2432211
 juha.heinanen@datanet.tele.fi

Marco Hernandez
 CREN
 Phone: 202-872-4200
 Fax: 202-872-4318
 marco@mh-slip.cren.edu

Robert Hinden
 Sun Microsystems, Inc.
 Phone: 415-336-2082
 Fax: 415-336-6015
 hinden@eng.sun.com

Rune Hjelsvold
 UNINETT
 Phone: +47 7 59 3686
 Fax: +47 7 59 4466
 Rune.Hjelsvold@idt.unit.no

Ellen Hoffman
 Merit Network, Inc.
 Phone: 313-936-3000
 Fax: 313-747-3745
 ellen@merit.edu

Frank Hoffmann
 IBM Deutschland Information
 System GmbH
 Phone: +49 6221 594427
 Fax: +49 6221 593300
 hoffmann@dhdibm1.bitnet

Gerd Holzhauer
 Apple Computer, Inc.
 Phone: +33 1 4901 4704
 Fax: +33 1 4090 0954
 Holzhauer1@applelink.apple.com

Alton Hoover
Advanced Network and
Services, Inc.
Phone: 914-789-5341
Fax: 914-789-5310
hoover@ans.net

John Hopkins
Imperial Cancer Research Fund
Phone: +44 71 269 3631
Fax: +44 71 269 3624
J_Hopkins@icrf.icnet.uk

Marc Horowitz
OpenVision
Phone: 617-374-3700
Fax: 617-374-3715
marc@mit.edu

Steven Horowitz
Chipcom Corporation
Phone: 508-460-8922
Fax: 508-460-8952
witz@chipcom.com

Nandor Horvath
Hungarian Academy of Sciences
Phone: +36 1 1497986
Fax: +36 1 1297866
horvath@sztaki.hu

Xinli Hou
University of Twente
Phone: +31 53 893747
Fax: +31 53 333815
xinli@cs.utwente.nl

Jack Houldsworth
ICL
Phone: +44 438 738 941
Fax: +44 438 738 904
J.Houldsworth@ste0906.wins.icl.co.uk

Jeroen Houttuin
RARE
Phone: +31 20 6391131
Fax: +31 20 6393289
houttuin@rare.nl

Chris Howard
INMARSAT
Phone: +44 71 728 1456
Fax: +44 71 728 1625
chris_howard@inmarsat.org

Jeff Hughes
Hewlett-Packard
Phone: 719-531-4777
Fax: 719-531-4505
jeff@col.hp.com

Christian Huitema
INRIA
Phone: +33 93 65 77 15
Fax: +33 93 65 77 65
Christian.Huitema@sophia.inria.fr

Erik Huizer
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
Erik.Huizer@SURFnet.nl

Geoff Huston
Australian Academic and
Research Network
Phone: +61 62493385
Fax: +61 62473425
g.huston@aarnet.edu.au

Thomas Hutton
San Diego Supercomputer Center
Phone: 619-534-5136
hutton@opus.sdsc.edu

Sascha Ignjatovic
VEDA Data Systems, Inc.
Phone: +43 1 876 3677
Fax: +43 1 587 0216
sascha@veda.co.at

Phil Irely
NSWC
Phone: 703-663-1582
Fax: 703-663-1675
pirey@relay.nswc.navy.mil

Ole Jacobsen
ConneXions
Phone: 415-962-2515
Fax: 415-949-1779
ole@interop.com

Per Jacobsen
University of Oslo
Phone: +47 22852470
Fax: +47 22852730
p.h.jacobsen@usit.uio.no

David Jacobson
IBM Corporation
Phone: 919-254-5590
Fax: 919-254-9515
dnjake@vnet.ibm.com

Ronald Jacoby
Silicon Graphics, Inc.
Phone: 415-390-1371
Fax: 415-969-2314
rj@sgi.com

Xander Jansen
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
xander.jansen@surfnet.nl

Olle Jarnefors
Royal Institute of Technology (KTH)
Phone: +46 8 790 71 26
Fax: +46 8 10 25 10
ojarnef@admin.kth.se

Dennis Jennings
Network & Computer Associates
Phone: +353 1 904826
Fax: +353 1 904825
jennings@irlearn.ucd.ie

J. Jensen
Internet Consult
Phone: +45 3117 1233
Fax: +45 3116 1607
jensen@ic.dk

Borka Jerman-Blazic
Jozef Stefan Institute
Phone: +38 61 159 199
Fax: +38 61 161 029
jerman-blazic@ijs.si

Thomas Johannsen
Dresden University of
Technology
Phone: +49 351 463 4621
Fax: +49 351 463 7163
thomas@ebzaw1.et.tu-dresden.de

Ola Johansson
Unisource Business Networks
Phone: +46 8 6882156
Fax: +46 8 6882007
ojn@tip.net

David Johnson
Carnegie Mellon University
Phone: 412-268-7399
Fax: 412-681-5739
dbj@cs.cmu.edu

Dale Johnson
Merit Network, Inc.
Phone: 313-936-3335
Fax: 313-747-3745
dsj@merit.edu

John Johnston
National Semiconductor
Phone: 408-721-3925
Fax: 408-739-6204
john@berlioz.nsc.com

Laurent Joncheray
Merit Network, Inc.
Phone: 313-936-3335
lpj@merit.edu

Philip Jones
JNT
Phone: +44 265 446 618
Fax: +44 235 446 251
p.jones@jnt.ac.uk

Kevin Jordan
Control Data Systems, Inc.
Phone: 612-482-6835
Fax: 612-482-4455
Kevin.E.Jordan@cdc.com

Cyndi Jung
3Com Corporation
Phone: 408-764-5173
Fax: 408-764-5002
cmj@3com.com

Peter Jurg
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
jurg@surfnet.nl

Marijke Kaat
SARA
Phone: +31 20 5928038
marijke@sara.nl

Thomas Kaepfner
IBM European Networking Center
Phone: +49 6221 59 4424
Fax: +49 6221 59 3300
kaepfner%heidelbg.vnet@ibmpa.ibm.com

Brewster Kahle
Wide Area Information Services
Phone: 415-617-0444
Fax: 415-327-6513
Brewster@wais.com

Marko Kaittola
FUNET
Phone: +358 31 162 933
Fax: +358 31 162 913
Marko.Kaittola@funet.fi

Tomaz Kalin
RARE
Phone: +31 20 639 1131
Fax: +31 20 639 3289
kalin@rare.nl

Nada Kapidzic
Stockholm University and Royal
Institute of Technology
Phone: +46 8 16 16 20
Fax: +46 8 703 90 25
nadak@dsv.su.se

Scott Kaplan
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
scott@wco.ftp.com

Anders Karlsson
Chalmers University of
Technology
Phone: +46 31 772 1684
Fax: +46 31 772 3663
sak@cdg.chalmers.se

Daniel Karrenberg
RIPE Network Coordination
Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
daniel@ripe.net

Frank Kastenzholz
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4488
kasten@ftp.com

Neil Katin
SunSoft
Phone: 415-336-3618
Fax: 415-336-6776
katin@eng.sun.com

Dave Katz
cisco Systems
Phone: 415-668-8284
Fax: 415-326-1989
dkatz@cisco.com

Peter Kaufmann
Deutsches Forschungsnetz
Phone: +49 30 8842 9932
Fax: +49 30 8842 9970
kaufmann@dfn.dbp.de

Sean Kennedy
Bolt Beranek and Newman
Phone: 617-873-3354
Fax: 617-873-5620
liam@nic.near.net

Stephen Kent
Bolt Beranek and Newman
Phone: 617-873-3988
Fax: 617-873-4086
kent@bbn.com

Michael Khalandovsky
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
mlk@ftp.com

Zbigniew Kielczewski
Eicon Technology Corporation
Phone: 514-631-2592
Fax: 514-631-3092
zbig@eicon.qc.ca

Steve Kille
ISODE Consortium
Phone: +44 81 332 9091
Fax: +44 81 332 9019
S.Kille@isode.com

Sun-Kwan Kimn
Hewlett-Packard
Phone: 408-447-3664
Fax: 408-447-3660
sunkimn@hp.com

Peter Kirstein
University College London
Phone: +44 71 380 7286
Fax: +44 71 387 1397
P.Kirstein@cs.ucl.ac.uk

Paul Klarenberg
NetConsult AG
Phone: +41 31 235750
Fax: +41 31 235792
klarenberg@netconsult.ch

Lothar Klein
German National Research
Center for Computer Science
Phone: +49 2241 14 3179
lothar.klein@gmd.de

John Klensin
United Nations University
Phone: 617-864-3247
Fax: 617-491-6266
Klensin@infofoods.unu.edu

Mark Knopper
Merit Network, Inc.
Phone: 313-763-6061
Fax: 313-747-3745
mak@merit.edu

Jim Knowles
Sterling Software
Federal Systems
Phone: 415-604-3056
Fax: 415-604-7300
jknowles@binky.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
Phone: 508-685-4000
Fax: 508-794-4484
stev@ftp.com

Andrew Knutsen
The Santa Cruz Operation, Inc.
Phone: 408-427-7538
Fax: 408-458-0811
andrewk@sco.com

Peter Koch
University of Bielefeld
Phone: +49 521 106 2924
Fax: +49 521 106 2962
pk@techfak.uni-bielefeld.de

Rajeev Kochhar
3Com Corporation
Phone: 408-764-5176
Fax: 408-764-5003
rajeev_kochhar@3com.com

Bruno Koechlin
INRIA
Phone: +33 39 63 57 36
Bruno.Koechlin@inria.fr

Ton Koelman
SHAPE Technical Centre
Phone: +31 70 3142429
Fax: +31 70 3142111
koelman@stc.nato.int

Klaus-Peter Kossakowski
Universitaet Hamburg,
FB Informatik
Phone: +49 40 54715 262
Fax: +49 40 54715 206
Kossakow@informatik.uni-hamburg.de

Mark Kosters
Network Solutions, Inc.
Phone: 703-742-4795
Fax: 703-742-4811
markk@internic.net

Deirdre Kostick
Bell Communications Research
Phone: 908-758-2642
Fax: 908-758-4177
dck2@mail.bellcore.com

Glenn Kowack
EUnet
Phone: +31 20 592 51 09
Fax: +31 20 592 51 55
Glenn@eu.net

John Krawczyk
Wellfleet Communications, Inc.
Phone: 508-436-3811
Fax: 508-670-8760
jkrawczyk@wellfleet.com

Arnold Krechel
Gesellschaft fuer Mathematik
und Datenverarbeitung
Phone: +49 2241 14 3188
Fax: +49 2241 14 3002
krechel@gmd.de

John Kunze
University of California,
Berkeley
Phone: 510-642-1530
Fax: 510-643-5385
jak@violet.berkeley.edu

Pekka Kytolaakso
Centre for Scientific
Computing
Phone: +358 4572246
Fax: +358 4572302
pekka.kytolaakso@csc.fi

Dave Langley
Hewlett-Packard
Phone: 916-758-4973
Fax: 916-786-9185
davel@hprnd.hp.com

Sylvain Langlois
Electricite de France
Phone: +33 1 47 65 44 02
Fax: +33 1 47 65 35 23
Sylvain.Langlois@exp.edf.fr

John Larson
Xerox Corporation
Phone: 408-662-9755
Fax: 408-662-9756
jlarson@parc.xerox.com

Mark Laubach
Hewlett-Packard Laboratories
Phone: 415-857-3513
Fax: 415-857-8526
laubach@hpl.hp.com

Erik Lawaetz
UNI-C
Phone: +45 45 93 83 55
Fax: +45 45 93 02 20
erik.lawaetz@uni-c.dk

Eliot Lear
Silicon Graphics, Inc.
Phone: 415-390-2414
Fax: 415-961-9584
lear@sgi.com

Debra Legare
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
dlegare@cnri.reston.va.us

Jose Legatheaux Martins
Faculdade de Ciencias e
Tecnologia da UNL
Phone: +351 1 295 4464
Fax: +351 1 295 4461
jalm@fct.unl.pt

Thomas Lenggenhager
SWITCH
Phone: +41 1 261 81 78
Fax: +41 1 261 81 33
lenggenhager@switch.ch

Arjen Lenstra
Bell Communications Research
Phone: 201-829-4878
Fax: 201-829-2645
lenstra@bellcore.com

Mark Lewis
Telebit Corporation
Phone: 408-745-3232
Fax: 408-745-3810
Mark.S.Lewis@telebit.com

Tony Li
cisco Systems
Phone: 415-688-8186
Fax: 415-688-4575
tli@cisco.com

Susan Lin
IBM Corporation
Phone: 203-795-1023
Fax: 203-783-7026
suelin@vnet.ibm.com

Steen Linden
UNI-C
Phone: +45 45938355
Fax: +45 45930220
steen.linden@uni-c.dk

John Lindsay
School of Information Systems
Phone: +44 81 547 2000
Fax: +44 81 547 1457
lindsay@kingston.ac.uk

John Linn
Geer Zolot Associates
Phone: 617-374-3700
linn@gza.com

Robin Littlefield
Wellfleet Communications, Inc.
Phone: 508-436-3812
Fax: 508-670-8760
robin@wellfleet.com

Anne Lord
RIPE Network Coordination
Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
anne@ripe.net

Peter Lothberg
EBONE
Phone: +46 8 6699720
Fax: +46 8 102510
roll@stupi.se

Paul Lukacs
Advanced Research Projects
Agency
Phone: 703-527-1970
Fax: 703-527-1973

Paul Lustgarten
AT&T Bell Laboratories
Phone: 908-576-2928
Fax: 908-576-6406
Paul.Lustgarten@att.com

Ron Mackey
Distributed Systems
International, Inc.
Phone: 708-665-4693
Fax: 708-665-4706
rem@dsiinc.com

Carl Madison
Star-Tek, Inc.
Phone: 508-393-9393
Fax: 508-365-6934
carl@startek.com

Kim Chr. Madsen
Internet Consult
Phone: +45 3117 1233
Fax: +45 3116 1607
kimcm@ic.dk

Carl Malamud
Internet Multicasting Service
Phone: 202-628-2044
Fax: 202-628-2042
carl@malamud.com

Paolo Malara
CRS4
Phone: +39 70 27 96 257
Fax: +39 70 27 96 245
malara@crs4.it

Kent Malave
IBM Corporation
Phone: 512-838-3486
Fax: 512-838-3509
kent@bach.austin.ibm.com

Andrew Malis
Ascom Timeplex, Inc.
Phone: 508-470-2565
Fax: 508-264-4999
malis@maelstrom.timeplex.com

Gary Malkin
Xylogics, Inc.
Phone: 617-272-8140
Fax: 617-272-3159
gmalkin@xylogics.com

Allison Mankin
Naval Research Laboratory
Phone: 202-404-7030
Fax: 202-767-1683
mankin@cmf.nrl.navy.mil

Bill Manning
Rice University
Phone: 713-285-5415
Fax: 713-527-6099
bmanning@rice.edu

April Marine
NASA NAIC
april@atlas.arc.nasa.gov

David Marlow
NSWC
Phone: 703-663-1675
dmarlow@relay.nswc.navy.mil

Cynthia Martin
Defense Information Systems
Agency
Phone: 703-487-3038
Fax: 703-487-3019
martin@spica.disa.mil

John Martin
University of
Newcastle upon Tyne
Phone: +44 91 222 8087
Fax: +44 91 222 8765
John.Martin@newcastle.ac.uk

Ignacio Martinez
RedIRIS
Phone: +34 1 4351214
Fax: +34 1 5781773
martinez@rediris.es

Jun Matsukata
Institute of Space and
Astronautical Science
Phone: +81 427 51 3911
Fax: +81 427 59 4252
jm@eng.isas.ac.jp

Chip Matthes
General Videotex Corporation
Phone: 617-491-3342
Fax: 617-491-6642
chip@delphi.com

Brian May
CSIRO
Phone: +61 3 282 2613
Fax: +61 3 282 2600
Brian.May@mel.dit.csiro.au

Richard McBride
Axon Networks, Inc.
Phone: +44 31 331 7816
richiem@axon.com

Keith McCloghrie
Hughes LAN Systems
Phone: 415-966-7934
Fax: 415-966-7980
kzm@hls.com

Chuck McManis
SunSoft
Phone: 415 336-1001
Fax: 415-336-6015
chuck.mcmanis@eng.sun.com

Marjo Mercado
Hewlett-Packard
Phone: 408-447-2826
Fax: 408-447-3660
marjo@cup.hp.com

Shehzad Merchant
SRI International
Phone: 415-859-3594
Fax: 415-859-4812
merchant@erg.sri.com

Peter Merdian
BelWue
Phone: +49 711 1319129
Fax: +49 711 682357
merdian@rus.uni-stuttgart.de

Donald Merritt
Army Research Laboratory
Phone: 410-278-6808
Fax: 410-278-5075
don@arl.army.mil

Gerry Meyer
Spider Systems Limited
Phone: +44 31 554 9424
Fax: +44 31 554 0649
gerry@spider.co.uk

Topi Miettinen
Tampere University of
Technology
Phone: +358 31 161 927
Fax: +358 31 161 857
tm86214@cs.tut.fi

Paul Milazzo
Bolt Beranek and Newman
Phone: 617-873-4325
Fax: 617-873-3776
milazzo@bbn.com

Linda Millington
University of London Computer
Centre
Phone: +44 71 405 8400
Fax: +44 71 242 1845
l.millington@noc.ulcc.ac.uk

Cynthia Mills
Bolt Beranek and Newman
Phone: 617-873-4143
Fax: 617-873-5620
cmills@bbn.com

Greg Minshall
Novell, Inc.
Phone: 510-975-4507
Fax: 510-947-1238
minshall@wc.novell.com

William Miskovetz
cisco Systems
Phone: 415-688-4682
Fax: 415-688-8282
misko@cisco.com

Keith Mitchell
PIPEX, Ltd.
Phone: +44 223 250100
Fax: +44 223 250101
keith@pipex.net

Mitra
Pandora Systems
Phone: 415-488-0944
Fax: 415-931-7931
mitra@pandora.sf.ca.us

Henri Moelard
NCR
Phone: +31 3402 76551
Fax: +31 3402 39125
henri.moelard@utrecht.ncr.com

Pushpendra Mohta
CERFnet
Phone: 619-455-3908
Fax: 619-455-3990
pushp@cerf.net

Keith Moore
University of Tennessee
Phone: 615-974-3735
Fax: 615-974-8296
moore@cs.utk.edu

George Mouradian
AT&T Bell Laboratories
Phone: 908-949-7671
Fax: 908-949-8569
gvm@arch3.att.com

Jun Murai
Keio University
Phone: +81 466 49 1100
Fax: +81 466 49 1101
jun@wide.ad.jp

Daniel Myers
3Com Corporation
Phone: 408-764-6865
Fax: 408-764-6134
dan@nsd.3com.com

Ramin Najmabadi
Universite Libre de Bruxelles
Phone: +32 2 650 5044
Fax: +32 2 641 3816
najmabadi@helios.iihe.rtt.be

Mark Needleman
University of California
Phone: 510-987-0530
Fax: 510-839-3573
mhn@stubbs.ucop.edu

Kees Neggers
SURFnet
Phone: +31 30 310 290
Fax: +31 30 340 903
neggers@surfnet.nl

Bruce Nelson
Novell, Inc.
Phone: 801-429-3875
Fax: 801-429-3086
bnelson@novell.com

Clifford Neuman
Information Sciences Institute
Phone: 310-822-1511
bcn@isi.edu

Lisa Nielsen
 University of Hawaii
 Phone: 808-956-5896
 Fax: 808-956-3548
 lisa@hawaii.edu

Torben Nielsen
 University of Hawaii
 Phone: 808-956-3499
 Fax: 808-956-3548
 torben@hawaii.edu

Ronny Nilsen
 University of Oslo
 Phone: +47 22 85 27 31
 Fax: +47 22 85 27 30
 Ronny.Nilsen@usit.uio.no

Peder Chr. Noergaard
 Telebit Communications ApS
 Phone: +45 86 28 81 76
 Fax: +45 86 28 81 86
 pcn@tbit.dk

Erik Nordmark
 Sun Microsystems, Inc.
 Phone: 415-336-2788
 Fax: 415-336-6015
 nordmark@eng.sun.com

Michael O'Dell
 UUNET Technologies, Inc.
 Phone: 703-204-8000
 Fax: 703-204-8001
 mo@uunet.uu.net

David O'Leary
 cisco Systems
 Phone: 714-261-5681
 Fax: 714-261-5684
 doleary@cisco.com

Masataka Ohta
 Tokyo Institute of Technology
 Phone: +81 3 5499 7084
 Fax: +81 3 3729 1940
 mohta@cc.titech.ac.jp

Petri Ojala
 EUnet Finland, Ltd.
 Phone: +358 49 425722
 Fax: +358 0 4572302
 ojala@eunet.fi

Lars-Gunnar Olsson
 Swedish University of
 Agriculture
 Phone: +46 18 67 17 76
 Fax: +46 18 30 15 33
 Lars-Gunnar.Olsson@data.slu.se

Zbigniew Opalka
 Agile Networks, Inc.
 Phone: 508-287-0700
 Fax: 508-287-0606
 zopalka@agile.com

Jorg Ott
 Technische Universitaet Berlin
 Phone: +49 30 314 73389
 Fax: +49 30 314 25156
 jo@cs.tu-berlin.de

Christian Panigl
 Vienna University Computer
 Center - AConet
 Phone: +43 1 436111 383
 Fax: +43 1 436111 170
 christian.panigl@cc.univie.ac.at

Andrew Partan
 UUNET Technologies, Inc.
 Phone: 703-204-8000
 Fax: 703-204-8001
 asp@uunet.uu.net

Michael Patton
 Bolt Beranek and Newman
 Phone: 617-873-2737
 map@bbn.com

Paul-Andre Pays
 INRIA
 Phone: +33 1 39 63 54 58
 pays@faugeres.inria.fr

Geir Pedersen
University of Oslo UNINETT
Phone: +47 22 85 24 78
Fax: +47 22 85 27 30
Geir.Pedersen@usit.uio.no

John Penners
US West Advanced Technologies
Phone: 303-541-6106
Fax: 303-541-8239
jpenners@advtech.uswest.com

Charles Perkins
IBM Corporation
Phone: 914-784-7350
Fax: 914-784-7007
perk@watson.ibm.com

David Perkins
SynOptics Communications, Inc.
Phone: 408-764-1516
Fax: 408-988-5525
dperkins@synoptics.com

Drew Perkins
Fore Systems, Inc.
Phone: 412-967-4040
Fax: 412-967-4044
ddp@fore.com

Roy Perry
US West Advanced Technologies
Phone: 303-541-6245
Fax: 303-541-6276
rperry@advtech.uswest.com

Catherine Pierre-Radenac
INRIA
Phone: +33 39 63 5261
catherine.pierre-radenac@inria.fr

David Piscitello
dave@mail.bellcore.com

John Plate
Internet Consult
Phone: +45 3117 1233
Fax: +45 3116 1607
plate@ic.dk

Mel Pleasant
Rutgers University
Phone: 908-932-2023
Fax: 908-932-2968
pleasant@hardees.rutgers.edu

Willi Porten
Gesellschaft fuer Mathematik
und Datenverarbeitung
Phone: +49 2241 14 3190
Fax: +49 2241 14 3002
porten@gmd.de

Lars Poulsen
Rockwell International
Corporation
Phone: 805-968-4262
Fax: 805-968-8256
lars@cmc.com

Aiko Pras
University of Twente
Phone: +31 53 893778
Fax: +31 53 333815
pras@cs.utwente.nl

Philip Prindeville
Ecole Nationale Supérieure des
Telecommunications
Phone: +33 1 45 81 73 14
Fax: +33 1 45 89 16 64
philipp@res.enst.fr

Mark Prior
The University of Adelaide
Phone: +61 8 228 5680
Fax: +61 8 223 6245
mrp@itd.adelaide.edu.au

J. Mark Pullen
George Mason University
Phone: 703-993-1538
Fax: 703-993-2738
mpullen@cs.gmu.edu

Dave Rand
Novell, Inc.
Phone: 408-321-1259
Fax: 408-946-9761
dave_rand@novell.com

Juergen Rauschenbach
Deutsches Forschungsnetz
Phone: +49 30 88429946
Fax: +49 30 88429970
jrau@dfn.de

Jim Rees
University of Michigan
Phone: 313-763-4174
Fax: 313-763-4434
Jim.Rees@umich.edu

James Reeves
SynOptics Communications, Inc.
Phone: 408-764-1440
Fax: 408-988-5008
jreeves@synoptics.com

Alex Reijnierse
PTT Research
Phone: +31 70 3325779
Fax: +31 70 3326477
a.a.reijnierse@research.ptt.nl

Victor Reijs
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
reijs@surfnet.nl

Yakov Rekhter
IBM Corporation
Phone: 914-945-3896
Fax: 914-945-2141
yakov@watson.ibm.com

Robert Reschly
US Army Research Laboratory
Phone: 410-278-6808
Fax: 410-278-5075
reschly@brl.mil

Joyce K. Reynolds
Information Sciences Institute
Phone: 310-822-1511
Fax: 310-823-6714
jkrey@isi.edu

Tony Richards
Sprint Government Systems
Division
Phone: 703-904-2460
Fax: 703-904-2463
richards@icm1.icp.net

Georg Richter
Universitotsrechenzentrum
Phone: +49 251 832476
Fax: +49 251 832090
richter@uni-muenster.de

Kenneth Rodemann
AT&T Bell Laboratories
Phone: 908-949-6229
Fax: 908-949-1726
krr@qsun.att.com

Benny Rodrig
RAD Network Devices, Ltd.
Phone: +972 3 6458568
Fax: +972 3 5447368
brodrig@rnd-gate.rad.co.il

Michael Roe
Cambidge University
Phone: +44 223 334718
Fax: +44 223 334678
mrr@cl.cam.ac.uk

Duncan Rogerson
University of London Computer
Centre
Phone: +44 71 405 8400
Fax: +44 71 242 1845
d.rogerson@nosc.ja.net

Jim Romaguera
NetConsult AG
Phone: +41 31 235750
Fax: +41 31 235792
romaguera@netconsult.ch

Dan Romascanu
Lannet Data Communications
Phone: +972 3 6458414
Fax: +972 3 5447146
dan@lannet.com

John Romkey
ELF Communications
Phone: 617-576-2856
Fax: 617-576-6715
romkey@elf.com

Luc Rooijackers
SPC Group
Phone: +31 412 043 888
Fax: +31 412 025 225
lwj@cs.kun.nl

Daisy Rose
IBM Corporation
Phone: 914-945-3632
Fax: 914-945-4439
daisy@watson.ibm.com

Marshall Rose
Dover Beach Consulting, Inc.
Phone: 415-968-1052
Fax: 415-968-2510
mrose@dbc.mtview.ca.us

Kenneth Rossen
SHL Systemhouse
Phone: 617-661-7474
Fax: 617-661-3038
kenr@shl.com

Marjo Rottschaefer
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
Marjo.Rottschaefer@surfnet.nl

Shawn Routhier
Epilogue Technology
Corporation
Phone: 617-942-0915
Fax: 617-942-7846
sar@epilogue.com

Rick Royston
US Robotics, Inc.
Phone: 708-933-5430
Fax: 708-982-1348
rroyston@usr.com

Ulla Sandberg
Ericsson Radio Systems AB
Phone: +46 8 757 03 81
Fax: +46 8 751 71 41
ulla@kiera.ericsson.se

Henry Sanders
Microsoft Corporation
Phone: 206-882-8080
henrysa@microsoft.com

Hal Sandick
IBM Corporation
Phone: 919-254-4614
Fax: 919-254-4975
sandick@vnet.ibm.com

Miguel Sanz
RedIRIS - Fundesco
Phone: +34 1 4351214
Fax: +34 1 5781773
miguel.sanz@rediris.es

Jon Saperia
Digital Equipment Corporation
Phone: 508-952-3171
Fax: 508-496-9929
saperia@tay.dec.com

Martina-Angela Sasse
 University College London
 Phone: +44 71 380 7212
 Fax: +44 71 387 1397
 a.sasse@cs.ucl.ac.uk

Srinivas Sataluri
 AT&T Bell Laboratories
 Phone: 908-949-7782
 Fax: 908-949-1726
 sri@internic.net

Jeffrey Schiller
 Massachusetts Institute of
 Technology
 Phone: 617-253-0161
 Fax: 617-258-8736
 jis@mit.edu

Jean-Bernard Schmitt
 IBM Corporation
 Phone: +33 92 11 56 60
 Fax: +33 93 24 78 09
 jbs@vnet.ibm.com

Wolfgang Schneider
 Gesellschaft fuer Mathematik
 und Datenverarbeitung
 Phone: +49 6151 869 719
 Fax: +49 6151 869 785
 schneiw@darmstadt.gmd.de

Hans Scholtes
 University of Twente
 Phone: +31 53 894287
 Fax: +31 53 333815
 scholtes@cs.utwente.nl

Tibor Schonfeld
 AT&T Bell Laboratories
 Phone: 908-949-3152
 Fax: 908-992-9893
 tibor@qsun.att.com

Eve Schooler
 Information Sciences Institute
 Phone: 310-822-1511
 Fax: 310-823-6714
 schooler@isi.edu

Heiner Schorn
 UMDAC
 Phone: +46 90 167767
 Fax: +46 90 166762
 heiner.schorn@umdac.umu.se

Rickard Schoultz
 KTHnoc/SUNET
 Phone: +46 8 790 9088
 Fax: +46 8 102 510
 schoultz@admin.kth.se

John Scudder
 Merit Network, Inc.
 Phone: 313-936-3335
 jgs@merit.edu

Tim Seaver
 MCNC - CONCERT Network
 Phone: 919-248-1973
 Fax: 919-248-1405
 tas@concert.net

Jennifer Sellers
 NASA NREN
 Phone: 202-434-8954
 Fax: 202-434-4599
 sellers@nsipo.nasa.gov

Henk Sennema
 SARA
 Phone: +31 20 592 3075
 Fax: +31 20 668 3167
 sennema@sara.nl

Kitty Shih
 Novell, Inc.
 Phone: 408-747-4305
 Fax: 408-747-4900
 kmshih@novell.com

Robert Shirey
The MITRE Corporation
Phone: 703-883-7210
Fax: 703-883-1397
shirey@mitre.org

Keld Simonsen
DKnet
Phone: +45 39 179900
Fax: +45 39 179899
keld@dkuug.dk

William Simpson
Computer Systems
Consulting Services
Phone: 313-623-7979
Bill.Simpson@um.cc.umich.edu

W. David Sincoskie
Bell Communications Research
Phone: 201-829-4426
Fax: 201-829-2504
sincos@thumper.bellcore.com

A. Velu Sinha
Delphi/AGS-NYNEX
Phone: 301-983-3899
avsinha@attmail.com

David Sitman
EARN Association
Phone: +972 3 640 8586
Fax: +972 3 640 9118
a79vm.tau.ac.il

Sam Sjogren
TGV, Inc.
Phone: 408-427-4366
Fax: 408-427-4365
sjogren@tgv.com

Keith Sklower
University of California,
Berkeley
Phone: 510-642-9587
Fax: 510-642-5775
sklower@cs.berkeley.edu

Timon Sloane
Timon Sloane Consulting
Services
Phone: 415-515-0741
Fax: 415-941-9726
timon@timon.com

Hugh Smith
Nottingham University/NEXOR
Phone: +44 602 514218
h.smith@cs.nott.ac.uk

Kenneth Smith
Northern Telecom Limited
Phone: 613-763-7720
Fax: 613-765-2186
kensmith@bnr.ca

Patricia Smith
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3745
psmith@merit.edu

Suzanne Smith
Lawrence Livermore National
Laboratory
Phone: 510-422-4017
Fax: 510-422-1482
smith@es.net

Karen Sollins
Massachusetts Institute of
Technology
Phone: 617-253-6002
Fax: 617-253-2673
sollins@lcs.mit.edu

Milan Sova
Czech Technical University
Phone: +42 2 3323925
Fax: +42 2 3111786
sova@feld.cvut.cz

Simon Spero
University of North Carolina
at Chapel Hill
Phone: 919-962-9107
Fax: 919-962-5604
simon_spero@unc.edu

Michael St. Johns
Advanced Research Projects
Agency
Phone: 703-696-2271
Fax: 703-696-2202
stjohns@darpa.mil

Bert Stals
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
bert.stals@surfnet.nl

Henk Steenman
SARA
Phone: +31 20 5928038
henk@sara.nl

Ed Stern
Proteon, Inc.
Phone: 508-898-2800
els@proteon.com

John Stewart
Phone: +31 17 51 41762
john@bunter.fdc.iaf.nl

John Stewart
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
jstewart@cnri.reston.va.us

Jack Stiekema
Victron bv
Phone: +31 50 44 62 22
Fax: +31 50 42 41 07
jack@victron.nl

Don Stikvoort
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
don.stikvoort@surfnet.nl

Vladimir Sukonnik
Process Software Corporation
Phone: 508-879-6994
sukonnik@process.com

Peter Svanberg
Royal Institute of Technology
Phone: +46 8 790 71 46
Fax: +46 8 790 09 30
psv@nada.kth.se

Richard Sweatt
SynOptics Communications, Inc.
Phone: 408-764-1260
Fax: 408-764-1299
rsweatt@synoptics.com

Erdal Taner
Middle East Technical
University
Phone: +90 4 210 1120
Fax: +90 4 210 1120
erdal@vm.cc.metu.edu.tr

Fumio Teraoka
Sony Computer Science
Laboratory, Inc.
Phone: +81 3 3448 4380
Fax: +81 3 3448 4273
tera@csl.sony.co.jp

Marten Terpstra
RIPE Network Coordination
Centre
Phone: +31 20 592 5065
Fax: +31 20 592 5090
marten@ripe.net

Kaj Tesink
 Bell Communications Research
 Phone: 908-758-5254
 Fax: 908-758-4196
 kaj@cc.bellcore.com

Kamlesh Tewani
 AT&T Bell Laboratories
 Phone: 908-949-5369
 Fax: 908-949-1196
 ktt@arch2.att.com

Richard Thomas
 Bell Northern Research
 Phone: 613-763-7025
 Fax: 613-765-4920
 rjthomas@bnr.ca

Susan Thomson
 Bell Communications Research
 Phone: 201-829-4514
 Fax: 201-829-2504
 set@bellcore.com

Dick Thoolen
 News Breakers
 Phone: +31 23 36 13 04
 Fax: +31 23 36 33 55

Martin Toet
 University of Twente
 Phone: +31 53 894287
 Fax: +31 53 333815
 toet@cs.utwente.nl

Jacques Togtema
 University of Twente
 Phone: +31 53 894287
 Fax: +31 53 333815
 togtema@cs.utwente.nl

Claudio Topolcic
 Corporation for National
 Research Initiatives
 Phone: 703-620-8990
 Fax: 703-620-0913
 topolcic@cnri.reston.va.us

Paul Traina
 cisco Systems
 Phone: 415-325-7943
 Fax: 415-325-4328
 pst@cisco.com

Antoine Trannoy
 CRS4
 Phone: +39 70 27 96 260
 Fax: +39 70 27 96 245
 trannoy@crs4.it

Catherine Treca
 CNRS/UREC
 Phone: +33 1 44 27 42 60
 Fax: +33 1 44 27 42 61
 Catherine.Treca@dione.urec.fr

Theodore Ts'o
 Massachusetts Institute of
 Technology
 Phone: 617-253-8091
 Fax: 617-248-8736
 tytso@mit.edu

Panos-Gavriil Tsigaridas
 Research Center for Open
 Communication Systems FOKUS
 Phone: +49 30 25499232
 Fax: +49 30 25499202
 Tsigaridas@fokus.gmd.de

Eftimios Tsigros
 Universite Libre de Bruxelles
 Phone: +32 2 650 57 02
 Fax: +32 2 641 38 16
 tsigros@helios.iihe.rtt.be

Thierry Turetletti
 INRIA
 Phone: +33 93 65 77 18
 Fax: +33 93 65 77 66
 turetletti@sophia.inria.fr

Robert Ullmann
 Process Software Corporation
 Phone: 617-247-7459
 ariel@world.std.com

Hisao Uose
NTT Telecommunication
Networks Laboratories
Phone: +81 422 59 3476
Fax: +81 422 59 2518
uose@tnlab.ntt.jp

Arend van den Boogaart
Royal PTT Nederland N.V.
PTT Research
Phone: +31 70 3325242
Fax: +31 70 3326477
a.k.vandenboogaart@research.ptt.nl

Rene van der Hauw
Geveke Electronics
Phone: +31 20 5869815
Fax: +31 20 5869777
rene@geveke.nl

Marc van der Noordaa
RARE
Phone: +31 20 639 1131
Fax: +31 20 639 3289
noordaa@rare.nl

Willem van der Scheun
SARA
Phone: +31 20 5923079
scheun@sara.nl

Aad van der Voort
PTT Research
Phone: +31 70 332 5391
Fax: +31 70 332 6477
a.t.h.vandervoort@research.ptt.nl

Jan P.M. van Oorschot
Delft University of Technology
Phone: +31 15 786179
Fax: +31 15 784898
j.p.m.voorschot@et.tudelft.nl

Guido van Rossum
CWI
Phone: +31 20 592 4127
Fax: +31 20 592 4199
guido@cw.nl

Jaap van Till
Consultancy Group STRATIX BV
Phone: +31 20 6400026
Fax: +31 20 6400609
vantill@dutetvd.et.tudelft.nl

Dono van-Mierop
3Com Corporation
Phone: 408-764-5930
Fax: 408-764-5002
dono_van_mierop@3mail.3com.com

Gregory Vaudreuil
Tigon
Phone: 214-733-2722
Fax: 214-733-8678
gvaudre@cnri.reston.va.us

Mario Vecchi
Bell Communications Research
Phone: 201-829-4374
Fax: 201-829-2504
mpv@thumper.bellcore.com

John Veizades
FTP Software, Inc.
Phone: 415-543-9001
veizades@wco.ftp.com

Ton Verschuren
SURFnet bv
Phone: +31 30 310290
Fax: +31 30 340903
Ton.Verschuren@SURFnet.nl

Paul Vetter
RCC
Phone: +31 55 778822
Fax: +31 55 215960

Lea Viljanen
University of Helsinki
Phone: +358 708 4347
Fax: +358 708 4441
Lea.Viljanen@helsinki.fi

Werner Vogels
INESC
Phone: +351 1 3100316
Fax: +351 1 525843
werner@inesc.pt

Ruediger Volk
University of Dortmund
Phone: +49 231 755 4760
Fax: +49 231 755 2386
rv@informatik.uni-dortmund.de

Steven Waldbusser
Carnegie Mellon University
Phone: 412-268-6628
Fax: 412-268-4987
waldbusser@andrew.cmu.edu

Sandro Wallach
ELF Communications
Phone: 617-576-2856
Fax: 617-576-6715
sandro@elf.com

Scott Wasson
Xyplex, Inc.
Phone: 508-952-4746
sgwasson@eng.xyplex.com

James Watt
Newbridge Networks Corporation
Phone: 613-591-3600
Fax: 613-591-3680
james@newbridge.com

Chris Weider
Merit Network, Inc.
Phone: 313-936-3000
Fax: 313-747-3185
clw@merit.edu

Terry Weigler
Corporation for National
Research Initiatives
Phone: 703-620-8990
Fax: 703-620-0913
tweigler@cnri.reston.va.us

Jost Weinmiller
Berlin Technical University
Phone: +49 30 314 26820
Fax: +49 30 314 21114
jost@prz.tu-berlin.d400.de

Abel Weinrib
Bell Communications Research
Phone: 201-829-4454
Fax: 201-829-5889
abel@bellcore.com

Marcel Wiget
SWITCH
wiget@switch.ch

Douglas Williams
IBM Corporation
Phone: 919-254-4801
dougw@ralvmg.vnet.ibm.com

Kirk Williams
Southwestern Bell Technology
Resources
Phone: 314-529-7586
Fax: 314-529-7674
kirk@sbctri.sbc.com

Steven Willis
Wellfleet Communications, Inc.
Phone: 508-436-3927
Fax: 508-670-8760
steve@wellfleet.com

Rachel Willmer
Spider Systems Limited
Phone: +44 31 554 9424
Fax: +44 31 554 0649
rachelw@spider.co.uk

Peter Wilson
3Com Corporation
Phone: +44 442 278102
peter_wilson@3mail.3com.com

Sam Wilson
The University of Edinburgh
Phone: +44 31 650 4979
Fax: +44 31 662 4712
sam.wilson@ed.ac.uk

Dirk Wisse
Delft University of Technology
Phone: +31 15 786209
Fax: +31 15 784898
dirk.wisse@dnpap.et.tudelft.nl

Wilfried Woeber
Vienna University Computer
Center - AConet
Phone: +43 1 436111 355
Fax: +43 1 436111 170
Wilfried.Woeber@CC.UniVie.ac.at

Russ Wright
Lawrence Berkeley Laboratory
Phone: 510-486-6965
Fax: 510-486-6363
wright@lbl.gov

Peter Yee
NASA Ames Research Center
Phone: 415-604-3812
Fax: 415-604-6999
yee@atlas.arc.nasa.gov

Noriko Yokokawa
Keio University
Phone: +81 466 49 1100
Fax: +81 466 49 1101
norinori@wide.ad.jp

Gordon Young
Xylogics International, Ltd.
Phone: +44 908 222112
Fax: +44 908 222115
young@xylogics.com

Jessica Yu
Merit Network, Inc.
Phone: 313-936-2655
Fax: 313-747-3745
jyy@merit.edu

Paul Zawada
National Center for
Supercomputing Applications
Phone: 217-333-1189
Fax: 217-244-1987
Zawada@ncsa.uiuc.edu


Steve Zeber
SHAPE Technical Centre
Phone: +31 70 314 2427
Fax: +31 70 314 2111
zeber@stc.nato.int

Erik Zegwaard
SURFnet
Phone: +31 30 310290
Fax: +31 30 340903
zegwaard@surfnet.nl

James Zmuda
Hughes Aircraft Company
Phone: 714-732-7240
Fax: 714-732-2000
zmuda@mls.hac.com

Romeo Zwart
SARA
Phone: +31 20 5928038
romeo@sara.nl

Corporation
for
National
Research
Initiatives



*Research
and
Development
for
the
National
Information
Infrastructure*

1895 Preston White Drive
Suite 100
Reston, VA 22091
USA

Tel: (703) 620-8990
Fax: (703) 620-0913