Proceedings of the Twenty-First Internet Engineering Task Force

BellSouth – Atlanta
July 29 - August 2, 1991

Corporation
for
National
Research
Initiatives

PROCEEDINGS OF THE TWENTY-FIRST INTERNET ENGINEERING TASK FORCE

BELLSOUTH
ATLANTA
July 29 - August 2, 1991

Compiled and Edited by
Megan M. Davies

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

This document, and the IETF plenary it reports, have been jointly sponsored by the National Science Foundation, the Defense Advanced Research Projects Agency, the Department of Energy, and the National Aeronautics and Space Administration through a grant from the the National Science Foundation (NCR-8820945).

Contents

Chairman's Message

Thanks to hosts BellSouth and Caroline Cranfill

I'd like to extend my appreciation to BellSouth and Caroline Cranfill for hosting the Atlanta IETF meeting. The facilities and connectivity were excellent. Our hosts deserve special thanks for the Tuesday evening social activity, which gave us a fine introduction to Atlanta hospitality. Thanks to our hosts for helping to make the IETF in Atlanta a successful one.

IETF and IESG Status Report

At the start of the Atlanta IETF meeting, there were 63 working groups in 9 areas (one area, Standards Practices, does not have working groups). For Atlanta, 9 working groups submitted "final " Internet-Drafts. Some of these documents represent standards actions, for which the IESG will be forwarding recommendations to the IAB. Others are informational documents. Most of these working groups will retire when their documents are published as RFCs (either as standards or informationl RFCs).

These numbers, both the total number of working groups and the number of working groups approaching conclusion at one meeting, represent high-water marks (yet-again) for IETF activity. Detailed listings and status of working groups will be given elsewhere in these Proceedings.

There has been some new additions to the IESG, and some minor restructuring, since the March IETF meeting.

First, we have added a new "Transport and Services Area", led by David Borman of Cray Research. Long time IETF attendees will recognize David from his numerous reports in the past regarding his work implementing high performance TCP for the Cray operating system, based Van Jacobson's enhancements. This new area will also include activities that don't quite fall properly into either the Internet or Application Areas (from a "layerist" perspective). The DNS activities are an example. We used to have an area called "Host and User Services", led by Craig Partridge. When Craig left for his postdoctoral position at the Swedish Institute of Computer Science (SICS) last fall, we divided the transport and other host issues between the Internet and Application Areas and elevated User Services as a distinct area (led by Joyce Reynolds, USC-ISI). However, that move added additional topics to two of the largest IETF areas (i.e., Internet and Applications). Plus, we soon realized that we had need for specific transport level expertise on the IESG. Therefore, we decided to form the new "Transport and Services Area". We are very pleased that David has been able to join the IESG to fill this new position.

Next, I am very happy to announce that the Operational Requirements Area is now fully staffed. Bernhard Stockman (Nordunet) and Susan Estrada (Cerfnet) have joined the IESG as co-Area Directors for this important area. Bernhard brings a special understanding of international networking to this task, which will help us as the Internet, and the surrounding technical issues, expand globally. Susan brings a perspective based on two important views – her long time key association with FARNET and her role as Executive Director of a midlevel network. I had been acting as an interim Director for this Area, and I will continue to work with Susan and Bernhard for the near future. One of the first objectives will be to form an Operational Requirements Area Directorate. We held an open session in Atlanta which we hope will act as an organizational session for the ORAD. You can read Susan's report on that meeting later in these Proceedings.

I am also pleased to welcome Philip Almquist onto the IESG, joining Noel Chiappa as co-Director of the Internet Area. The Internet Area is one of the largest, and historically one of the more important, in the IETF and it will profit from the additional attention that Noel and Philip together will now be able to bring to bear. Philip will bring an important perspective as the current chair of the Router Requirements working group.

I am also sorry to have to report that Rob Hagens (University of Wisconsin) has had to leave the IESG due to other pressing time commitments. We will miss Rob as the co-Director for the OSI Integration Area, with Ross Callon (DEC). Fortunately, we will not lose Rob completely because he will remain as co-chair of the X.400 Operational working group. The OSI Integration Area is a large area, requiring a special set of skills, and we already are looking for someone to join Ross in leading this area. Rob, we will miss you on the IESG, but we look forward to your continued contributions to the IETF as a working group chair.

With the changes above, there are now 10 IETF Areas and 13 members of the IESG. A complete listing of the areas and members is given in Chapter 1 of these Proceedings.

New IETF Secretariat Staffing

Over the past couple years, the IETF has greatly expanded its scope and activities. Most of the technical activities of the IETF are handled by volunteers from either academia or industry (i.e., the working group chairs, the IESG members, and/or the many individual working group contributors). However, with its current schedule of 3 yearly meetings of 350-400 attendees (and with a 400 page Proceedings, and numerous online Internet-Drafts and reports, documenting each meeting), the administration and logistics of the IETF can no longer be handled solely by volunteers.

To handle the growing administrative and logistics necessities, a professional staff has grown up at CNRI over the last few years to support the largely volunteer technical contributors. With the hiring this summer of Steve Coya (CNRI) as IETF Executive

Director, perhaps it is now time to recognize this "IETF Secretariat" for the important role it now plays.

This is not unlike other professional associations (like ACM or IEEE) which make heavy use of volunteer contributors. In those organizations, the principle technical contributing positions (analogous to our WG chairs or IESG) are filled by volunteers, but there is a core professional staff which carries out the inevitable but important administrative and logistics details that make the organization run smoothly.

It is clear, at least in my mind, that the IETF, with its current size and scope, could no longer function as it now does without this important support.

The current full-time professional staff comprising the "IETF Secretariat" at CNRI is:

Steve Coya	IETF Executive Director
Greg Vaudreuil	IESG Secretary
Megan Davies	Proceedings and Meeting Coordinator
Cynthia Clark	IETF Mailing List Coordinator
Debra Legare	Internet Drafts and Registration

Steve Coya is still fairly new to the IETF, and I would like to take this opportunity to welcome him to this new important position. While I was at CNRI, I performed various functions for the IETF which could be characterized under the two separate headings of "executive director" responsibilities and "overall chair" responsibilities. Those duties that comprised the executive director responsibilites now reside with Steve Coya. This includes, but is not limited to, overall responsibility for insuring that the IETF Secretariat efficiently and effectively discharges its collective duties. Steve also attends all IESG meetings, and will be working with me to complete an "IETF Handbook", which will serve as introduction and guideline for operations of the IETF. (This handbook will be be submitted as an Internet-Draft for general review and comment, when complete.) As another example, Steve will serve as the primary point of contact for submitting the monthly IETF report to the "Internet Monthly Report". The IETF chair and any on the Area Directors may also contribute to the report, but Steve will have the responsibility for reporting on the overall IETF status and technical administration (e.g. new WGs formed, old WGs completed, Internet-Drafts submitted, protocols actions taken by the IESG, etc). Plus he will act to compile all other IESG contributions into the overall report, and submit it to ISI. Again, let me welcome Steve to the IETF in this important new position. We will be seeing a lot of Steve in the future!

Please join me in thanking these folks and CNRI for the wonderful support they have provided in the past and will continue to provide at future meetings.

If there are any questions or suggestions regarding the Secretariat, please feel free to correspond directly with Steve Coya (scoya@nri.reston.va.us). Either Vint Cerf (CNRI) or myself can also answer questions, if Steve is not available.

Atlanta's Open Plenary

The Atlanta meeting represents the largest number of working groups the IETF has ever sponsored in one place at one time. As a measure of the growth and activity of the IETF, this is an exciting statistic. However, it also raises important questions about how large the IETF can grow and still enjoy its present success and utilize its present format.

These and other issues were very much on the minds of attendees during the open plenary session on Thursday afternoon.

Specifically, this most recent growth surge has re-surfaced some problems that we have seen, and dealt with, in the past. For example, there is the simple logistics hassle of scheduling many parallel WG meetings so that important groups do not overlap. There is also the problem of bringing new players into the WG activities without completely rehashing all the old technical issues. One answer is to increase our effective usage of electronic mail. However, then we must be careful to balance the "group consensus" reached by email versus the consensus reached by onsite attendees. And, of course, there is the general problem of reaching closure on technical issues in a large group, whether that group is meeting face-to-face or by email.

These are all problems that we have seen in various degrees, and attempted to deal with in various ways, in the past. These are all problems that seem to re-surface at the leading edge of any new IETF growth surge.

What's less clear, at least to me, is whether we have finally hit a "wall", in which the basic procedures that the IETF has developed to conduct business are no longer valid, or whether we merely need to adapt to the new environemnt (yet-again) by evolving our operating procedures.

One of the strengths of the IETF has always been its ability to evolve and to adapt to new conditions. In many ways, we are still learning lessons about how best to conduct our business. "Making it up as we go" has been both a blessing and a curse. At times we have seen interactions between the IAB and IESG, or between the IESG and the working groups, suffer from lack of clarity of "procedures". On the other hand, when conditions change around us, we have always had the flexibilty to adapt quickly.

My personal feeling is that our ability to evolve has been one of the major keys to our success. I think we must all recognize that the networking world has changed in recent years in ways that very few could have predicted, and the IETF has had the ability to learn from any missteps and change along with it. We are still evolving, and I'd like to think that we have a ways to go (and grow) before we hit a "wall".

Some excellent suggestions came out at Thursday's open plenary sesssion. Acting on these suggestions, we will add an extra afternoon session on Friday at the next IETF meeting. This will allow us to add one additional WG session to the current generic schedule and to devote both all of Monday and Tuesday to WG sessions. We will also try to make good use of the new/old idea of "Birds of a feather" sessions, to explore subjects before commiting expensive (and scarce) personnel resources to starting up a full WG. (In Atlanta, we had 5 BOFs on important subjects.) We will make an extra effort to schedule "overview" sessions on topics of interest, particularly as status reports on ongoing WG activities. This will help keep general attendees and new players abreast of activities on other areas.

Finally, we will enhance our current set of guidelines for WG chairs to provide suggestions on how to handle new attendees who need to be brought up to speed and how to achieve closure in the face of questions from these new players. For example, acting on a suggestion from the open plenary, we will ask WG chairs to do a better job of recording the rationale for all technical decisions. In this way, new players will have a better record of why specific decision were made. Some suggestions were as simple as to make better use of agendas and to invoke a different set of meeting guidelines as a WG nears conclusion. For example, in Atlanta the Router Requirements WG had a strict agenda of still-open topics, and the chair announced beforehand that previous issues would not be re-opened in Atlanta without very good cause. It's also important for WG chairs to understand that they are bound to reach consensus, not unanimity. That is, in the face of unresolved technical opinions, it is perfectly valid for the chair to adapt the consensus view and then move forward.

It is important to understand that the open plenary sessions have always played an important role in guiding the IETF. (In fact, in *very* beginning, there were no WGs, so the entire meeting was an open plenary!) The open plenaries have included IESG reports since the IESG formation at Hawaii meeting (Fall 1989).

Of course, there are other ways to provide comments on IETF activity. The main IETF mailing list is available for that purpose (ietf@isi.edu). The IESG can be reached individually or collectively (iesg@nri.reston.va.us). I am always available either by phone or email (pgross@nis.ans.net, 914-789-5335), and the IETF Secretariat can be reached through Steve Coys, IETF Excutive Director, (scoya@nri.reston.va.us).

I believe it is very important to have direct interaction with attendees and others interested in IETF activities. This makes the IETF very much a self-guiding body, and I think that has contributed as much to the IETF's success as any other factor.

Internet Society

This meeting saw the announcement of the Internet Society by Vint Cerf (CNRI). (See the Technical Presentations Chapter for a summary of Vint's presentation.)

The Internet Society will be a professional society which we hope will grow to encompass the activities of the IAB and IETF in a positive way. The IAB and IESG are very supportive of the Internet Society, and we look forward to working with Vint on developing the specific method by which the IAB and IETF both become part of this new important group.

Focus on International Growth

Another topic that received considerable attention at this meeting was the growing focus in the IETF on global networking issues. This emphasis is almost unavoidable as the Internet grows internationally. (MERIT's network status report again showed that the non-US portion continues to be the fastest growing segment of the Internet.)

We had the largest showing of non-US attendees in Atlanta, and we now have non-US representation on both the IAB and IESG.

This raised the possibility of holding an IETF meeting outside of North America (we have already held an IETF in British Columbia). I think this is a very natural and desirable outcome. However, we have to approach it cautiously. Many US employers and US federal agencies all treat non-US travel quite differently from domestic travel, and we need to make sure that the environment has matured properly so that we can expect to get representative participation at any IETF meeting outside North America. As important as it is to begin holding IETF meetings in non-US venues, it could be damaging to our very goals of internationalization if we held a non-US meeting that was poorly attended by US participants.

We normally schedule IETF meetings 12-18 months in advance, so the earliest spot not yet scheduled is Fall 1992 or Winter 1993. Fortunately, this provides time to prepare and plan for our first meeting outside North America.

One way to help prepare the climate for such a meeting is make the Internet Society a strong and successful organization. One of the principal goals of the Internet Society will be to focus on international networking issues. So please consider joining the Internet Society. For my part, I will work directly with Vint Cerf and incoming IAB chair Lyman Chapin (BBN) to see that the IAB and IETF are incorporated into the Internet Society in a positive and natural fashion.

Final Agenda of the Twenty-First IETF

(July 29 - August 2, 1991)

MONDAY, July 29

8:00-9:00 am	IETF Registration and Continental Breakfast	
9:00-9:30 am	Introductions and Local Arrangements	
9:30-12:00 noon	Morning Sessions	
	APP	Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	MGT	Remote LAN Monitoring WG (Mike Erlinger/Micro Technology)
	OPS	Network Status Reports WG (Phill Gross/ANS)
	OSI	OSI Directory Services WG (Steve Hardcastle-Kille/UCL)
	OSI	OSI General WG (FTP-FTAM Gateway Specification Review) (Rob Hagens/UWisc)
	RTG	Inter-Domain Policy Routing WG (Martha Steenstrup/BBN)
	SEC	Security Area Advisory Group (observers welcome) (Stephen Crocker/TIS)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	Network News Transport Protocol WG (Eliot Lear/Intelligenetics)
	APP	Network Printing Protocol WG (Glenn Trewitt/DEC)
	BOF	Subnets (Philip Almquist/Consultant)
	MGT	IEEE 802.3 Hub MIB WG (Keith McCloghrie/Hughes and Donna McMaster/SynOptics)
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	OSI	Network OSI Operations WG (Sue Hares/Merit)
	OSI	OSI Directory Services WG (Steve Hardcastle-Kille/UCL)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	SEC	Privacy-Enhanced Electronic Mail I WG (Steve Kent/BBN)

	USV	User Services WG (Joyce Reynolds/ISI)telnet)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Afternoon Sessions II	
	APP	Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	OPS	Operational Statistics WG (Phill Gross/ANS and Bernhard Stockman/NORDUnet)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	Open Shortest Path First IGP WG (John Moy/Proteon)
	SEC	Privacy-Enhanced Electronic Mail I WG (Steve Kent/BBN)
	SEC	SNMP Security WG (James Galvin/TIS and Keith McCloghrie/Hughes)
	USV	Directory Information Services Infrastructure WG (Chris Weider/Merit)
7:00-10:00 pm	Evening Sessions	
	BOF	Automated Internet Mailing List Services (Dave Lippke/UTexas)
	BOF	Conditioning of By-Request Network Resources (Andy Nicholson/Cray Research)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	TSV	Dynamic Host Configuration WG (Ralph Droms/Bucknell)
	USV	NOC-Tools Catalogue Revisions WG (Robert Enger/Contel and Gary Malkin/FTP Software)

TUESDAY, July 30, 1991

8:30-9:00 am	Continental Breakfast, No Morning Plenary	
9:00-12:00 noon	Morning Sessions	
	APP	Network Database WG (Daisy Shen/IBM)
	APP	Telnet WG (Dave Borman/Cray Research)
	INT	Point-to-Point Protocol Extensions WG (Noel Chiappa/Consultant)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	MGT	Internet Accounting WG (Cyndi Mills/BBN)
	MGT	X.25 Management Information Base WG (Dean Throop/Data General)
	OSI	X.400 Operations WG (Alf Hansen/UWisc)
	RTG	Inter-Domain Policy Routing WG (Martha Steenstrup/BBN)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	USV	Network Information Services Infrastructure WG (Dana Sitzler and Pat Smith/Merit)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	APP	Internet Message Extensions WG (Greg Vaudreuil/CNRI)
	APP	Network Database WG (Daisy Shen/IBM)
	MGT	FDDI MIB WG (Jeff Case/UTenn)
	OPS	Benchmarking WG (Scott Bradner/Harvard)
	OPS	Operational Area Directorate (observers welcome) (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
	OSI	X.400 Operations WG (Alf Hansen/UWisc)
	RTG	IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG	ISIS for IP Internets WG (Ross Callon/DEC)
	SEC	Commercial Internet Protocol Security Option WG (Ron Sharp/ATT)

	SEC	Privacy-Enhanced Electronic Mail II WG (Steve Kent/BBN)
	SEC	Site Security Handbook WG (Paul Holbrook/CICNet and Joyce Reynolds/ISI)
	TSV	Service Location Protocol WG (John Veizades/Apple)
3:30-4:00 pm	Break (Refreshments provided)	
4:00-6:00 pm	Technical Presentations	

- Security Issues and Directions (Steve Crocker/TIS)
- Site Security Handbook Report (Paul Holbrook/CICNet)
- Guidelines for the Secure Operation of the Internet (Steve Crocker/TIS and Barbara Fraser/CERT)
- Trusted Systems Interoperability Group (Paul Cummings/DEC)

WEDNESDAY, July 31, 1991

8:30-9:00 am	Continental Breakfast	
9:00-9:30 am	Technical Presentations	
	•	NREN Architecture and Goals (St. Louis follow-up) (William Johnston/LBL and Peter Ford/LANL)
9:30-12:00 noon	Morning Sessions	
	APP	Automated Internet Mailing List Services WG (Dave Lippke/UTexas)
	BOF	NREN Architecture and Goals (Peter Ford/LANL)
	INT	IP over Appletalk WG (John Veizades/Apple)
	INT	Router Requirements WG (Philip Almquist/Consultant)
	MGT	Internet Accounting WG (Cyndi Mills/BBN)
	MGT	Simple Network Management Protocol (Marshall Rose)
	OSI	X.400 Operations WG (Alf Hansen/UWisc)
	RTG	Border Gateway Protocol WG* (Yakov Rekhter/IBM)
	RTG	IP over Large Public Data Networks WG* (George Clapp/Ameritech)
	SEC	Commercial Internet Protocol Security Option WG (Ron Sharp/ATT)
	SEC	Common Authentication Technology WG (John Linn/DEC)
Breaks	Coffee available throughout morning.	
1:30-3:30 pm	Afternoon Sessions I	
	BOF	IP Logical Networks (Philip Almquist/Consultant)
	INT	IP over Appletalk WG (John Veizades/Apple)
	MGT	SNMP Network Management Directorate (Chuck Davin/MIT)
	OPS	Network Joint Management WG (Gene Hastings/PSC)
	OSI	Office Document Architecture WG (Peter Kirstein/UCL)
	OSI	X.400 Operations WG (Alf Hansen/UWisc)
	RTG	Border Gateway Protocol WG* (Yakov Rekhter/IBM)
	RTG	IP over Large Public Data Networks WG* (George Clapp/Ameritech)

	SEC	Commercial Internet Protocol Security Option WG (Ron Sharp/ATT)
	SEC	Common Authentication Technology WG (John Linn/DEC)
	TSV	Domain Name System WG (Mike Reilly/DEC)
	TSV	Dynamic Host Configuration WG (Ralph Droms/Bucknell)

3:30-3:45 pm — Break (Refreshments provided)

3:45-6:30 pm — Technical Presentations

- NREN Legislative Update (Mike Roberts/EDUCOM)
- IP over Frame Relay Report (Caralyn Brown/Wellfleet)
- BellSouth Telecommunications (Caroline Cranfill/BellSouth)
- Introduction to the Internet Society (Vint Cerf/CNRI)
- Toward a New Routing Architecture (Noel Chiappa/Consultant)

7:00-10:00pm — Evening Session

	APP	Network News Transport Protocol WG (Eliot Lear/Intelligenetics)
	BOF	Internet Support for Mobile Hosts (Steve Deering/Xerox)
	BOF	IP Address Enhancements (Noel Chiappa/Consultant)
	BOF	RFC 1148 Bis Editing (Steve Hardcastle-Kille/UCL)
	MGT	Management Services Interface WG (Oscar Newkerk/DEC)

* BGP and IPLPDN will meet jointly to discuss "Discovery and Routing over SMDS"

THURSDAY, August 1, 1991

8:30-9:00 am	Continental Breakfast
9:00-9:45 am	Technical Presentations • NSFNET T3 Deployment (Elise Gerich/Merit and Jordan Becker/ANS)
10:00-12:00 noon	Morning Sessions
	APP Automated Internet Mailing List Services WG (Dave Lippke/UTexas)
	APP Network Fax WG (Mark Needleman/UC)
	INT Router Requirements WG (Philip Almquist/Consultant)
	MGT Simple Network Management Protocol (Marshall Rose)
	MGT Internet Accounting WG (Cyndi Mills/BBN)
	OSI OSI General (CO/CL Interworking Review) (Rob Hagens/UWisc)
	RTG IP over Large Public Data Networks WG (George Clapp/Ameritech)
	RTG Multicast Extensions to OSPF WG (Steve Deering/Xerox PARC)
	SEC Security Area Advisory Group (observers welcome) (Stephen Crocker/TIS - Meets in Hilton)
Breaks	Coffee available throughout the morning.
1:30-3:30 pm	Technical Presentations • Introduction to Coalition for Networked Information (Clifford Lynch/UCAL) • IAB Workshop Report (Bob Braden/ISI) • IESG Evolution Plan (Phill Gross/ANS)
3:30-4:00 pm	Break (Refreshments provided)
4:00-6:30 pm	Open Plenary and IESG • IESG Evolution Plan (Phill Gross/ANS) • Protocol Standards Actions

FRIDAY, August 2, 1991

8:30-9:00 am	Continental Breakfast	
9:00-11:30 am	Working Group Area and Selected Working Group Presentations	
		Introduction to Archie (Peter Deutsch/McGill Univ)
		Recent Results from the DDN NIC Host Count and DNS Count Program (April Marine/SRI)
	APP	Applications Area (Russ Hobby/UC Davis)
	INT	Internet Area (Noel Chiappa/Consultant and Philip Almquist/Consultant)
	MGT	Network Management Area (Chuck Davin/MIT)
	OPS	Operations Area (Susan Estrada/CERFnet, Phill Gross/ANS, Bernhard Stockman/NORDUnet)
	OSI	OSI Integration Area (Ross Callon/DEC and Rob Hagens/UWisc)
	RTG	Routing Area (Bob Hinden/BBN)
	SEC	Security Area (Steve Crocker/TIS)
	TSV	Transport and Services Area (Dave Borman/Cray Research)
	USV	User Services Area (Joyce K. Reynolds/ISI)
11:30-12:00 noon	Concluding Remarks (Phill Gross/ANS)	
12:00 pm	Adjourn	

Key to Abbreviations

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

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) has grown into a large open 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 began in January 1986 as a forum for technical coordination by contractors working on the ARPANET, DDN, and the Internet core gateway system.

The IETF mission includes:

- Specifying the short and mid-term Internet protocols and architecture for the Internet,
- Making recommendations regarding Internet protocol standards for IAB approval,
- Identifying and proposing solutions to pressing operational and technical problems in the Internet,
- Facilitating technology transfer from the Internet Research Task Force, and
- Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

Technical activity on any specific topic in the IETF is addressed within Working Groups. All Working Groups are organized roughly by function into eight technical areas. Each is led by an Area Director who has primary responsibility for that one area of IETF activity. Together with the Chair of the IETF, these eight technical Directors compose the Internet Engineering Steering Group (IESG).

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

IETF and IESG Chair:	Phill Gross/ANS
Applications:	Russ Hobby/UC-Davis
Internet:	Noel Chiappa
	Philip Almquist/Consultant
Network Management:	James Davin/ MIT
OSI Integration:	Ross Callon/DEC
Operational Requirements:	Phill Gross/ANS
	Bernard Stockman/Nordunet
	Susan Estrada/CERFnet
Routing:	Robert Hinden/BBN
Security:	Steve Crocker/TIS
Transport and Services	Dave Borman/Cray
User Services	Joyce Reynolds/ISI
Standards Management:	Dave Crocker/DEC
IESG Secretary:	Greg Vaudreuil/CNRI

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 plenary sessions three times a year composed of Working Group Sessions, Technical Presentations and Network Status Briefings. The meetings are currently four and one half days long and include an open IESG meeting.

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

Mailing Lists

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

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

Information and logistics about upcoming meetings of the IETF are distributed on the general IETF mailing list. For general inquiries about the IETF, send a request to `ietf-request@isi.edu`. An archive of mail sent to the IETF list is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 Future IETF Meeting Sites

Fall 1991

Los Alamos National Laboratory
Host: Dale Land and John Morrison
November 11-22, 1991

Spring 1992

San Diego Supercomputer Center
Host: E. Paul Love, Jr. and Hans-Werner Braun
March 16-20, 1992 (tentative)

1.2 On Line IETF Information

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

Directory Locations

Information pertaining to the IETF, its Working Groups and Internet Drafts can be found in either the "IETF" Directory or the "Internet-Drafts" Directory. (For a more detailed description of these Directories, please see Section 1.2.1 and 1.2.2). To retrieve this information via FTP, establish a connection, then Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

cd ietf
cd internet-drafts

Individual files can then be retrieved using the GET command:

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

East Coast (US) Address: nnsc.nsf.net (192.31.103.6)

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

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

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

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

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

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

Europe Address: nic.nordu.net (192.36.148.17)

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

1.2.1 The IETF Directory

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

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

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

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

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

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

1.2.2 The Internet-Drafts Directory

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

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

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

where <std wg abrev> is the Working Group acronym, <docname> is a very short name, and <rev> is the revision number.

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

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

where <author> is the author's name.

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

1.3 Guidelines to Authors of Internet Drafts

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

Internet Drafts are not an archival document series. These documents should not be cited or quoted from in any formal document. Unrevised documents placed in the Internet Drafts directories have a maximum life of six months. After that time, they will either be submitted to the RFC editor or 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.

Following the practice of the RFCs, submissions are to be sent in ASCII. Postscript is also acceptable, however, we still require the submission of a matching ascii version (even if figures must be deleted) for readers without postscript printers and for on-line searches.

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

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

The Internet Draft should neither state nor imply that it is a proposed standard. To do so conflicts with the role of the IAB, the RFC Editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the Internet Draft. These are common words in the "Status of the Memo" section and may cause confusion if placed in the title. If the Internet Draft becomes an RFC, the Status of the Memo section will be filled in by the RFC editor with a status assigned by the IAB. As an Internet Draft, that section should contain a statement approximating one of the following statements:

1. This draft document will be submitted to the Internet Activities Board as a standards document. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

2. This document will be submitted to the Internet Activities Board as an proposed standard. This document defines an experimental extension to the SNMP MIB. Upon publication as a proposed standard, a new MIB number will be assigned. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

3. This draft document will be submitted to the RFC editor as an informational document. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

4. This draft document will be submitted to the RFC editor as an experimental protocol. This is a working document only, it should neither be cited nor quoted in any formal document. This document will expire before <Date, six months from current date>. Distribution of this memo is unlimited. Please send comments to <working group mailing list>

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

Chapter 2

Steering Group Report

2.1 Standards Progress Report

Between the March meeting hosted by Washington University in St. Louis and the July meeting hosted by Bell South in Atlanta there have been many IETF originating protocols and informational documents published as RFC's.

RFC1214	OSI Internet Management: Management Information Base
RFC1220	Point-to-Point Protocol Extensions for Bridging It is a product of the Point to Point Protocol Extensions Working Group group.
RFC1224	Techniques for Managing Asynchronously Generated Alerts It is a product of the Alert Management Working Group.
RFC1229	Extensions to the Generic-Interface MIB It is a product of the SNMP Working Group.
RFC1230	IEEE 802.4 Token Bus MIB It is a product of the SNMP Working Group.
RFC1231	IEEE 802.5 Token Ring MIB It is a product of the SNMP Working Group.
RFC1232	Definitions of Managed Objects for the DS1 Interface Type It is a product of the SNMP Working Group.
RFC1233	Definitions of Managed Objects for the DS3 Interface Type It is a product of the SNMP Working Group.
RFC1237	Guidelines for OSI NSAP Allocation in the Internet It is a product of the OSI NSAP Guideline Working Group.
RFC1238	CLNS MIB - for use with Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542) It is a product of the SNMP Working Group.
RFC1242	Benchmarking Terminology for Network Interconnection Devices It is a product of the Benchmarking Methodology Working Group.
RFC1243	AppleTalk Management Information Base It is a product of the IP-Appletalk Working Group.

RFC1244	Site Security Handbook It is a product of the Site Security Handbook Working Group of the IETF.
RFC1245	OSPF Protocol Analysis It is a product of the OSPF Working Group.
RFC1246	Experience with the OSPF Protocol It is a product of the OSPF Working Group.
RFC1247	OSPF Version 2 It is a product of the OSPF Working Group.
RFC1248, RFC1252	OSPF Version 2 Management Information Base It is a product of the OSPF Working Group.

2.2 Minutes of the Open Plenary and IESG

Agenda:

- Intro to the IETF
- IESG Evolution Plan
- IGP Statement
- Open Plenary

2.2.1 Introduction to the IETF

Phill Gross opened the IESG and Open Plenary session with a review of the structure of the IETF and IESG. The IESG has added several new members: Susan Estrada and Bernhard Stockman have joined Phill Gross as Directors of the Operations Area; Dave Borman has joined the IESG as the Director of the recreated Transport and Services Area, formerly Craig Partridge's Host and User Services; and Philip Almquist has joined Noel Chiappa as a Director for the Internet Area. For a complete listing of Areas and Directors, please see the Introduction to the IETF, in Section 1.

The IETF Secretariat has also grown and now numbers five, including a new Executive Director, Steve Coya. Greg Vaudreuil continues work as the IESG Secretary, tracking and managing the IETF standards process. Megan Davies plans and coordinates the IETF Plenary sessions, including logistics, meeting arrangements, and the scheduling of working groups slots, and is the editor of the IETF Proceedings. Debra Legare is the name behind the Internet drafts process, as well as handling IETF meeting registrations. Cynthia Clark, a new addition to the Secretariat, is working to maintain the IETF mailing list and also provides administrative support as needed

The IETF now has over 63 working groups. A surge of MIB activity has produced much work and has spawned many working groups. Considering the groups that have finished their work, and those likely to finish in the near future, the number should again be in the comfortable under-50 range.

2.2.2 IESG Evolution Plan

The IETF is responsible for the short to mid-term evolution of the Internet Protocol Suite. To lead the IETF in this mission, the Chairman of the IETF created the IESG to coordinate and plan the work of the IETF. This activity begins in one of three ways:

1. A person comes to an IETF meeting with an idea, and that idea becomes the basis of working group,

2. The RFC editor receives an RFC submission that he feels would benefit from the broader exposure the IETF can give, and

3. The IESG Area Directors may have a plan and commence efforts to achieve a goal.

The IESG evolution plan was created to facilitate the planning of the work by the IETF. The current draft of the document is available in the Internet Drafts directories. During this session, each Area Director gave a brief overview of their plans. Please see the Internet Draft of the five year plan for more detail <draft-ietf-iesg-evolutionplan-00.txt>.

2.2.3 Protocol Actions

IGP Statement

The IESG recognized that as far back as the February 1990 IETF meeting in Tallahassee, Florida, multi-vendor interoperability of routers with a modern Interior Gateway Protocol requires the implementation of a common IGP between all platforms. By a common IGP, the IESG means that IGP vendors are expected to implement, not that this is the preferred IGP for any particular environment. The IESG has chosen OSPF as the common IGP for the Internet. This recommendation is available for comment in the Internet Drafts directory, and will be sent to the IAB as an Applicability Statement.

2.3 Open Plenary

2.3.1 IETF Finances

Vint Cerf provided a summary of the allocation of fees paid to the Corporation for National Research Initiatives in support of the IETF meetings. These figures are presented in slides which are included later in this report. The presentation was well received by the IETF membership who indicated that the questions which had been raised were answered.

2.3.2 Ethernet MIB

A statement by the SNMP Working Group meeting was read stating the group's concerns over the manner in which the IAB and the Working Group interacted on the issue of the Ethernet MIB. The discussion which followed focused both on the technical and policy issues involved with that particular document, and the more general concern with the nature of the interaction between the IAB, the IESG, and the IETF working groups.

MIB Discussions

Discussion focused on the fact that the Ethernet MIB contains several required variables which are not in the required set of the IEEE 802.3 MIB. Two concerns were raised by the IAB: interactions with IEEE and the acceptable actions of a working group, and the feasibility of implementing the variables given the current hardware availability. The Working Group felt they had addressed the concern of hardware availability, stating that there are several implementations which use the contested variables, often retaining counters in the driver software rather than using hardware registers.

The IAB noted that the IETF is not the only body responsible for the standardization of management instrumentation for ethernet-like media. There was concern that the IAB not overstep its authority by making mandatory new instrumentation above and beyond that already required by other standards. This policy issue was not one considered by the working group. The working group did point out, however, that in the two years the group was publicly working n the MIB, they never received guidance from the IAB.

Policy Discussions

While there were specific technical disagreements between the IAB and the SNMP Working Group, the underlying cause of the friction was perceived by most to be the process by which the IAB and the IETF working groups interact. There was a strong feeling that the work of the IETF is done in working groups, and the ensuing discussion centered on the appropriate time for substantive technical input by IAB members. It was understood that the IAB has the right and obligation to conduct a final technical review, but there was a uniform sense that the final review of the IAB is not the appropriate nor reasonable time for giving policy guidance, and any last minute changes in broad direction points to a deficiency of the review process.

The Ethernet MIB situation was exacerbated by an ill-defined process for resolving disagreements between the IAB and the Working Group. There was a feeling in the Working Group that the IAB was mandating changes to the group in an authoritative and closed manner. Upon further discussion, this turned out not to be the case. The IAB took a long time to reach an understanding of the issues, but communicated its view in an ad-hoc manner which gave the impression of a mandate. This issue is still open for discussion.

A tentative solution discussed in this meeting involved the clear statement of process for resolution. As understood in the plenary, the IAB has the authority to reject a specification for technical reasons. If it does so, it needs to send the rejection to the IESG, and the relevant Working Group Chair(s) with a technical description of the problem.

Working group minutes are generally not complete, and mailing list archives are unwieldy to search. When the IAB has questions, a dialogue between the IAB and the Working Group Chair will often result in an understanding and resolution of the issues. If the problem is serious, a dialogue should ensue resulting in resolution. This mechanism is the expected process understood by the IETF.

2.3.3 Open Meetings

There was concern expressed in the IETF plenary over what is viewed as a proliferation of "closed" meetings, including IESG and Area Directorate sessions. With little discussion, it was recognized that the primary motivation for most closed meetings was to insure that progress is made. During the week of IETF meetings, a few new mechanisms were tried, including limiting participation to only the "core" group, but allowing silent observers who were free to comment during breaks. This was successful and was held up as a better means of doing business.

The IESG meets weekly by teleconference, and in executive session at IETF plenary meetings. Meeting minutes and other documentation (listing of individual action items and status, crafting of positions, formulating recommendations, etc.) have been distributed only among IESG members and eventually sent to the IAB. In the spirit of openness, the IESG committed itself to releasing timely minutes which will be distributed to the IETF mailing list.

2.3.4 Plenary Meeting Discussions

Terminal Room Hours

The terminal room at IETF meetings is a wonderful service, and attendees wish it could stay open forever. It was noted that there are often long lines waiting for terminal access. Many comments and suggestions were made ranging from buying a dozen vt100 like terminals, 24 hour access, use of rent-a-cops, to offering dial up service in the hotel.

The terminal room is a service provided by the local host, but is coordinated by the IETF Secretariat. Suggestions and comments may be sent to the Executive Director, Steve Coya <scoya@nri.reston.va.us>.

Meeting Duration and Interval

In addition to having the highest number of attendees, the number of working group and BOF meetings held during the week reached an all time high, as did the complaints that it was becoming increasing difficult for individuals to attend all the meetings desired. Several suggestions were offered and discussed during the plenary: conducting more meetings each year, more days per meetings (starting on Sunday or full days on Friday), fewer technical presentations, more evening sessions, etc. Other suggestions included limiting the number of times any one group could meet, limiting the number of working groups that would meet, and restricting BOF meetings to evening sessions.

After a significant period of discussion, the attendees expressed a preference for making Friday a full day, thereby adding another working group session. It was decided that the number of IETF meetings will remain at three per year, and the Secretariat will initially extend the duration on Friday to allow the addition of an extra working group session during the week.

International Meetings

The IETF is increasingly an international group, and as such there is increasing pressure to hold meetings internationally. There was a great deal of interest in holding a meeting in Europe, especially for the OSI Area groups.

A large percentage of participants in this plenary session stated they would be willing to travel to Europe, while there were many notable "core" participants who stated they would not be able to travel. It was noted that some small companies simply cannot afford international travel, and many U.S. employers perceive international travel as a perk or luxury. The IETF is currently difficult to characterize to employers, with a most general description as "a working meeting of a lot of people." Many individuals stated that moving the meeting to an international site would make it that much more difficult to gain travel authorization.

No decision was reached in the discussions, though all reiterated the desire that meetings be held internationally.

OPERATING THE IETF

- o CNRI has five full-time staff devoted to IETF Functions

  Steve Coya - Executive Director

  Greg Vaudreuil - IESG Secretary

  Megan Davies - Proceedings and IETF Meeting Logistics

  Cynthia Clark - IETF Mailing Lists and assistance to MD

  Debra Legare - Secretary
  Internet Drafts
  IETF Meeting Registration

OPERATING THE IETF (2)

- o Sources of support:

  - FNC via NSF

  - IETF Attendance Fees

  - Internal CNRI resources

- o All income from IETF-related activity is reported to NSF and offsets audited costs for IETF support.

OPERATING THE IETF (3)

- o CNRI does not keep separate account of staff costs associated with IETF meetings.

- o Out-of-pocket costs include:

  Food costs
  Audio-visual rentals
  Local copying
  Room rental
  Proceedings Production
  Teleconferences

OPERATING THE IETF (4)

- o Typical Cost Ranges

Food:	\$13,000 - \$20,000
AV:	\$ 1,400 - \$ 2,400
Copy:	\$ 1,300 - \$ 1,500
GRT[1]:	\$ 1,300 - \$ 1,400
Proc:	\$10,000 - \$14,000

	\$27,000 - \$39,300

Note 1: GRT:Gross Receipts Tax

- o IESG Teleconference Calls

  Nov 90 - Mar 91 ~ \$400/month

  Mar 91 - Jun 91 ~ \$1300/month

  Annual est: \$15K - \$20K/year
  (Note: int'l calls now)

INCOME

ST. Louis	40,000	(approx)
ATLANTA	48,000	(est.)
	88,000	

COSTS

ST. LOUIS	27,000	(approx)
ATLANTA	39,000	(est.)
Teleconf.	1,600	NOV 90 - MAR 91
Teleconf	5,200	MAR 91 - JUNE 91
	72,800	

Net	15,200	applied to other IETF costs

OPERATING THE IETF (5)

SUMMARY

- o All income is allocated to IETF Secretariat Operation
- o Any net after out-of-pocket offsets reduces charges against NSF/CNRI coop. agreement. The FNC is not able to justify subsidizing the entire cost of IETF operation
- o CNRI is a non-profit with income solely from grants and fees
- o Ultimately, it is hoped that Internet Society income can help to support IETF operating costs

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UCDavis

Protocols in Support of Personal Communications

At the Atlanta meeting there were four Working Groups, Internet Message Extensions Working Group (822EXT), Internet Mail Extensions Working Group (SMTPEXT), Network News Transport Protocol Working Group (NNTP), and the Automated Internet Mailing List Services Working Group (LIST), all of which had a common goal: A better system for communications between people and groups of people. Currently electronic mail is widely used for personal communication on the Internet. Network News has also become a very useful tool for information exchange. However, these systems need updating to provide the capabilities that people want.

Ideally, a user would use a single User Agent to interact with both email and network news. That agent would combine the strong points from both of these types of information exchange. With the implementation of that User Agent in mind, there are three main areas in which the protocols need work.

1. Message Format. ASCII text is no longer sufficient for the kind of information that we want to exchange. As we develop multimedia information, we need to devise methods of including it in the messages sent over the Internet. The 822EXT Working Group is defining the method to provide this function.

2. Message Delivery. Improved methods are needed for the delivery of messages that allow the new formats, and improve security and efficiency. The SMTPEXT Working Group is working on this for email. The biggest question right now is how to go beyond the current seven bit characters specified for SMTP. The NNTP Working Group is updating the transport for network news. Both of these groups will be using the format developed in the 822EXT Working Group.

3. EMAIL Management. Currently the use and management of email lists is not very easy on the Internet. The LIST Working Group wants to define a syntax for the maintenance of email lists by the users themselves. This would allow for a user to be added, deleted as well as for other operations to be carried out simply by sending a message to the automated list maintainer. The LIST Working Group will be providing requirements to the other groups for any features needed in the message format or transport.

Other Working Groups Meeting in Atlanta

The Telnet Working Group had discussions on the best ways to allow different methods of authentication.

The Network Database Working Group continued work on SQL over TCP/IP.

The Network Printing Protocol Working Group worked on the LPR/LPD document and the Printer Access Protocol document.

The Network FAX Working Group finalized an image format to be used on the Internet for FAX images.

Future Direction of the Applications Area

The number one item that was stressed at the IAB retreat in San Diego was the need for common formats of information for the Internet community. We may develop great protocols for transporting information over the network, but it does not do us any good if we cannot understand that information. We need to develop a common multimedia "language" that we can speak on the Internet. As the email groups have discovered, text formats need some work, but we also need to agree on image, audio and other formats to create a multimedia Internet.

One problem is that there are already so many standards to choose from in this area. For the most part we don't need to invent new formats, but rather agree to use a set of existing ones that best suit the Internet. I would like to solicit ideas on how best to select these formats for the Internet community.

CURRENT MEETING REPORT

Reported by David Lippke/UTEXAS

LISTBOF Minutes

The Automated Internet Mailing List Services BOF was held Monday night, before the regular Working Group meetings on Wednesday and Thursday mornings. The purpose in holding the BOF was simply to gather interested folks in order to conduct a rather open and free-wheeling discussion about the list server problem and the Working Group's Agenda.

A rather wide-ranging discussion did ensue. Topics ranged from "How is this any different than netnews?", through the enumeration of specific features that the "phase 2" list server world should offer, and on to the presentation of interconnection models and the cures for certain problems that will arise in the "phase 2" work.

In spite of the meeting's unstructured format, the group reached three conclusions:

1. Simple statement of why the Working Group exists: Dealing with Internet mailing lists is a pain for everyone involved – users, list owners, and postmasters alike. Internet mailing lists lack fundamental features such as access control and standardized archive maintenance. In short, the Internet mailing list world is a very primitive one... one which is in serious need of improvement.

2. Netnews groups and list server lists are closely related methods of group communication. While each has its own area of most appropriate application, their services are more alike than different and, in particular, they are nearly identical (in principle) at the user level. Consequently, eventual integration of their services is expected — at least at the user interface level if not also in the very provisioning of the services themselves.

3. Implementation point: If the phase 2 world seems to require a new (transport-level) protocol definition, the burden of proof (that the protocol is actually necessary) rests with the Working Group.

Attendees

Thomas Brisco	brisco@rutgers.edu
James Conklin	conklin@bitnic.educom.edu
John Curran	jcurran@bbn.com
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.arc.nasa.gov
Russ Hobby	rdhobby@ucdavis.edu

Neil Katin	katin@eng.sun.com
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@turbo.bio.net
Louis Leon	osll@emuvm1.cc.emory.edu
David Lippke	lippke@utdallas.edu
Daniel Long	long@nic.near.net
Joseph Malcolm	jmalcolm@sura.net
Keith Moore	moore@cs.utk.edu
Chris Myers	chris@wugate.wustl.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Harri Salminen	hks@funet.fi
Theodore Tso	
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.1 Automated Internet Mailing List Services (list)

Charter

Chair(s):

David Lippke, lippke@utdallas.edu

Mailing Lists:

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

Description of Working Group:

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

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

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

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

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by David Lippke/UTEXAS

LIST Minutes

The Automated Internet Mailing List Services Working Group had two separate meetings, one on Wednesday morning and another the following morning. The second meeting was simply a continuation of the first and, consequently, these notes do not distinguish between the two. Further, the order of presentation here doesn't necessarily reflect the order of discussion at the meetings (where each topic was generally visited more than once).

To start things off, we reviewed why the Working Group exists and what we are, and are not, trying to do. The view expressed was the following:

> We are here because dealing with Internet mail lists is painful for everyone involved — users, list owners, and postmasters alike. Current Internet mailing list services generally lack fundamental features such as access control and standardized archive maintenance. In short, the Internet mailing list world is a very primitive one... and one which is in serious need of improvement. The Working Group exists to address this need.
>
> However, we are NOT here to create the ultimate group communication system. Although list-style group communication should eventually become part of an integrated group communication system, our goals are more focused and short term. The feeling is that we have to learn to walk before we can run.

After establishing this, the group went on to discuss the agenda for the meetings. Two major potential directions were identified. Either we could define a baseline user interface or we could spend the time trying to develop a picture of the phase 2 list server world. The following list of pros and cons was reviewed:

Reasons to define a baseline interface:

1. Damage control in the name of minimizing user confusion. Alternative view: we need to define the first few articles of a "user's bill of rights" (e.g., users have the right to receive confirmation of all transactions, users have the right to see whether or not they are subscribed to a given list, etc).

2. Enable implementors to begin work now.

3. We can define the baseline quickly (assumption).

4. It's a fail-safe strategy for the Working Group (i.e., recognize that there's a significant chance that the phase 2 work will fail. If so, the Working Group will have at least accomplished something).

Reasons to NOT define a baseline interface:

1. Perhaps we would codify something that won't fit future models well.
2. Perhaps the subset we define now will be too limited to be useful.

After some discussion, we decided to attempt a definition of the baseline's contents and use that process to learn more about the problem and to see if we hit any show stoppers (which would indicate that defining a baseline user interface at this time is not the proper thing to do).

Thus, we began discussing a long list of functions and design issues. We rated each of the functions as "in" or "out" of the baseline interface and discussed each of the design issues long enough to develop views on how each should be treated.

Each of these items is given separate treatment below, but to jump straight to the end, the final conclusion was that we were satisfied with our efforts to define a baseline user interface and that enough functionality was contained within it to warrant its publication.

BASELINE COMPONENTS

- INCLUDED: Subscribe/Unsubscribe capability

  Discussion: An obvious conclusion. Also concluded that any subscription policy was allowable (e.g., open, closed, by service area, etc), but that the user is always owed a confirmation, explanation, or denial. See more general comments in the ISSUES section.

- INCLUDED: List parameter review capability

  Discussion: If users can see that a list exists, then they should be able to review its operational definition (e.g., see who the owner is, see what the subscription policy is, etc). Also, there was a general consensus expressed that a list's definition should include a "keywords" parameter which could be used as an aid in searching. The expression of the various parameters is not to be specified.

- INCLUDED: List subscriber review capability

  Discussion: If users can see that a list exists, then they should be able to obtain a list of its subscribers, UNLESS list policy dictates otherwise. In all cases, requesting users are owed either the list or an explanation of why they cannot retrieve it.

- INCLUDED: List of lists capability

  Discussion: Users should be able to obtain a list of all lists a given list server knows about (and they're allowed to know exist). We agreed that list servers needed to somehow identify for what "domain" they spoke, but tabled the implementation details for discussion on the Working Group list.

- INCLUDED: Various minor commands (HELP, POLICY, STOP, etc.)

  Discussion: We resolved that this wasn't worth spending time on and that the details should be worked out on the Working Group mailing list.

- EXCLUDED: Per-user options

  Discussion: This was tabled as being too demanding of implementations and because we predicted that there would be no quick agreement on what a set of baseline options should be. After the initial conclusion, it was later countered that users had a fundamental right to conceal their membership on a list and that the implementation of this was not overly complex even with simple-minded sendmail alias implementations. The ensuing discussion revealed that while the Mailbase implementors currently allowed per-user concealment, they will soon remove that capability since their users had raised the opposite argument (i.e., that they had a right to see who they were talking to when they posted a message to a list). This counter-counterpoint showed that the issue was a debatable one. Since our razor was that if something was debatable, it was not baseline, we returned to our initial conclusion.

- EXCLUDED: Archive Searching and Archive Retrieval

  Discussion: Although it was universally accepted that archival services are important, exploration of this topic revealed a number of sticky issues which we felt could be not quickly resolved. Difficulties ranged from problems related to the (previously agreed upon) need for program- interpretable list server responses to the quagmires of search method specification. Thus, the whole area of archive services was booted. The interim suggestion is that the output of a list parameter review mention how the archives are to be obtained.

- EXCLUDED: File services

  Discussion: This died for reasons similar to those that killed the inclusion of archive services.

- EXCLUDED (with proviso): Authentication

  Discussion: All cookie approaches do significant damage to the current pattern of user interaction. We have no experience with these approaches nor have we spent time looking for alternatives. Consequently, the introduction of such a facility in the baseline was deemed a real bad idea. HOWEVER, the baseline definition will mandate that all list server transactions be logged for X (TBD) period of time in a way that allows listmasters to reverse transactions, should the need arise. Also, any transaction

which affects a user (mail address) should result in a confirmation/informational message being sent to that user (mail address). The feeling was that this is similar to what goes on now and at least offers some degree of reactive protection.
It was also noted that PEM does not address the question of whether or not a person can speak authoritatively for a given mail address (although it may diminish the exposure since one at least knows *who* caused the trouble).

- EXCLUDED: Proxy Operations

  Discussion: Proxy operations are desirable, but we uncovered a complex set of problems and possible approaches once we dug into the issue. Also riding against their inclusion was the lack of solid authentication (these two issues seem to feed on each other ...).

DESIGN ISSUES / PHILOSOPHIES

- LISTSERV Compatibility

  The LISTSERV command set and interaction methods/patterns are the de facto standard. We should not be afraid to vary from that standard, but we should only do so when there is ample cause.

- Where should mail commands be sent?

  Directly to the list server agent address for the most part, but mail to listname-request and listname-owner should do a reasonable thing (which, even on BITNET, could be simple aliases to the list owner).

- How should the results of commands be returned?

  By default, they should be returned via the mechanism the commands were received. Command results should also be machine-interpretable. The intent was that we should define how this is done, but the issue was tabled for Working Group list discussions. In any case, the view is that both humans and GUI tools need to be able to make requests and understand the response(s).

- General syntax rules

  Tabled for discussion on the list.

- Channels and other provisions for upwards compatibility.

  Part of the above and likewise tabled.

- General note on command interaction

Users are always owed a message confirming (directly or indirectly) the reception and disposition of their requests.

- Identity of list servers

  A minor issue, but list servers should identify themselves (general type, version number, etc.) in some appropriate way during most transactions. (where "appropriate" and "most" is TBD on the list).

- Header handling

  Although further debate is assured, the group came up with the following guidelines in regard to how list mail should be sent to subscribers.

  1. Steps should be taken to ensure that rejections are never routed back to a list.
  2. "Sender:" and SMTP return path should never be set to the list address.
  3. Header trace information should not be stripped.
  4. The list equivalent of a "Received:" line is needed (e.g., Exploded:). Resolved to work with the 822 Extensions folks on this.
  5. Messages from a list should be unambiguously identifiable as coming from that list. Header extensions may be required for this as well.
  6. "Reply-To:" should not be modified.

Attendees

James Conklin	conklin@bitnic.educom.edu
Peter Deutsch	peterd@cc.mcgillica
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Jill Foster	jill.foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
James Galvin	galvin@tis.com
Neil Katin	katin@eng.sun.com
Darren Kinley	kinley@crim.ca
Mark Knopper	mak@merit.edu
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@turbo.bio.net
Louis Leon	osll@emuvm1.cc.emory.edu
David Lippke	lippke@utdallas.edu
Joseph Malcolm	jmalcolm@sura.net
Chris Myers	chris@wugate.wustl.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Jan Michael Rynning	jmr@nada.kth.se
Harri Salminen	hks@funet.fi
Larry Snodgrass	snodgras@bitnic.educom.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Chris Weider	clw@merit.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Russ Wright	wright@lbl.gov

3.1.2 Distributed Scheduling Protocol (chronos)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

Jan 1991 Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. Prototype implementations.

Feb 1991 Make document an Internet Draft. Continue revisions based on comments received over e-mail.

Mar 1991 Spring IETF meeting. Review final draft and if OK, give to IESG for publication as RFC. Begin implementations.

Jul 1991 Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

3.1.3 Internet Mail Extensions (smtpext)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done — Review the Charter of the group. Determine if changes to SMTP are necessary. Discuss the needs for backward compatability, and interoperabiliy. This discussion will be held by email.

Aug 1991 — Discuss the elimination of the 7 bit restrictions in SMTP, and the implications of removing this restriction in terms of interoperation.

Aug 1991 — Discuss the issues involved with binary transmission. Determine whether a "binary" mode should be pursued, and whether the SMTP line length restriction should be eliminated.

Dec 1991 — Write a document specifying the changes to SMTP agreed to by the group. Post as an Internet Draft.

Mar 1992 — Review and finalize the SMTP Extensions document.

Mar 1992 — Submit the SMTP Extensions document as a proposed standard.

Internet Drafts:

"SMTP Extensions for Transport of Text-Based Messages Containing 8-bit Characters", 07/10/1991, John Klensin, R. Kankkunen, G. Vaudreuil <draft-ietf-smtpext-8bittransport-00.txt>

INTERIM MEETING REPORT

Reported by Greg Vaudreuil/CNRI

SMTPEXT Minutes

Agenda

- 8 Bit Transport
- Overview of Current Status
- Review of Current Proposals
  - Negotiated 8 Bit Support
  - Unnegotiated 8 Bit Support
  - Use of 7 Bit Transport with Encoding
- Discussion: Which is the Preferred Proposal?
- Mail Enclave Issues
- Local use of Non-standard practices: Real Problem or Not?
  - Non-standard or Non-support of transfer encodings
  - Local Use of Non-standard Character sets
- Define Enclaves
  - Administratively Limited?
  - Universal Mesh of Capabilities?

Minutes

8 Bit Transport Issues

Much discussion has occured both on the mailing list and in private with the Chair calling into question the conclusions the SMPTEXT Working Group reached during the St. Louis IETF in March '91. This issue was raised at the Copenhagen meeting to reconsider or reaffirm the conclusions the Working Group at that earlier meeting. There continue to be two credible proposals for transition to 8 bit transport. The first is a proposal to redefine the SMTP protocol to standardize the existing practice sending 8 bit mail over standard SMTP channels. The second is a proposal to send 8 bit textual data after negotiating that capability.

The Working Group reviewed the two proposals and came to the following understandings about the proposals. A third "proposal", do nothing, was evaluated as well.

- Redefine SMTP to pass 8 bit data.

  The proposal stems from the existing practice of sending 8 bit data between SMTP implementations without negotiation or confirmation of the capabilities of the receiver.

Benefits

- It works (mostly.)
- Easy modification to existing code to gain functionality.
- Currently deployed by several vendors, and tested extensively in current mail environments.

Costs

- There is no assurance that the message was delivered as intended.
- Use of C1 codespace may be compromised.

Discussion

- The extensions have been extensively tested in a "friendly" environment where character sets sent have not used C1 codespace for graphic characters, nor have multi-byte characters been sent.
- Some unpredictable behavior has been noted.
- The costs are continuing, they never go away.

• Negotiate the sending of 8 Bit Data

  Benefits

  - Backwards compatible with current conformant implementations.
  - Failure is detectable, and recovery by encoding and resending is possible.

  Costs

  - Not compatible with some (much?) current deployed software.
  - Failure recovery after negative negotiation potentially complex.
  - Code changes are more complex.

  Discussion

  - The costs of the transition are one time, and will fade away with time.

• Send no 8 bit data

  Benefits

  - The hassle of upgrading current transport is unneeded.
  - All functionality is supported through encoding.

  Costs

  - Encoding required additional resources including computer time as well as communications bandwidth.
  - Local users may use 8 bit transport anyway.

  Discussion

- The technical analysis of this issue is but a small part in the problem. There is a strong feeling, almost religion, among site administrators and many "users" that encoding data that is easily transportable over the network infrastructure is wasteful, inelegant, and just plain wrong.

Conclusions

The attendees of this meeting reaffirmed the Working Group consensus that standards for the transmission of 8 bit characters without negotiation have costs which would be too in terms of expected mail performance to be acceptable. The main points underscoring this conclusion were the inability to "know" the transaction was successful, and the effective loss of the ability to use C1 codespace in future character sets intended to be transportable over SMTP. While it was noted that much experience has been gathered with current implementations using un-negotiated 8 bit mail, it was understood that this experience was gathered in a relatively homogeneous environment with friendly character sets. Problems were expected by the Working Groups in general application in the Internet and in sending characters sets like IBM PC codepages which use C1 codespace.

Enclave Issues

The Working Group felt that the concept of enclaves was not something that had to be defined. Specifically, the idea that enhanced capabilities should be confined to an administrative or geographical region was seen as being too restrictive. The attendees preferred to maintain the end-to-end model of electronic mail, rather than formalize the concept of autonomous mail domains.

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

SMTPEXT Minutes

Agenda

- Where are we, and where are we going?
  - Just send 8 bits
  - Negotiate 8 bits
  - Do nothing
- If negotiated, how to do transport conversion?
  - Encapsulation
  - Message Munging
- Defining the "New" SMTP

Where are We, and Where are We Going?

The Chair began this meeting by reviewing the history of this Working Group and the goals as they have evolved. This meeting was called in part to affirm the progress on the mailing list in a room where true give-and-take could be had. In a nutshell, the SMTP extensions were first motivated by those who want to be able to send 8 bit textual data via SMTP. This is already being done in practice. The group discussed the goals and in light of current deployment of non-standard systems, refined the goals to include a more general extension to the SMTP protocol.

There was a general feeling among many participants that a simple extension to support only 8 bit textual data was not worth the transition costs involved in upgrading the system. There are however many reasons to update the mail transport protocol. Among these needs are arbitrary options negotiation, binary transport, maximum message length restrictions, and "real" authentication. A sampling of opinions from the meeting:

- The Europeans REALLY REALLY want to send their stuff without encoding it. They REALLY REALLY want to do this via a negotiated option so they could have an assurance that the mail was delivered as intended.

- Existing software vendors, Prime, Sun, and others not so visible, do not feel that 8 bit textual data transmission is worth the costs of modification. This was strongly asserted at the St. Louis IETF, while the mailing list (led in part by the Chair) went off and wrote an SMTP extensions specification for 8 bit mail anyway.

- Even the multi-part multi-media mail people could agree with the assertion that the world would be a better place if binary data could be sent.

After a bit of soul searching, the group agreed to work on a complete change to SMTP which would allow future new features to be added via negotiation, and would allow binary and 8 bit transport.

Interworking of 7 bit, 8 bit and Binary Transport.

Now that the Working Group decided to move ahead on new functionality, the next question to be solved was the definition of an interworking strategy. Fortunately for this group, the Message Format Extensions group decided to keep nested transport encodings in their proposed standard document. While this feature is tentative and subject to the results of implementation experience, it provides a mechanism for initial implementations. After a short amount of discussion, the group decided to write a specific, well defined conversion algorithm which specifies that messages which need to be converted between transport environments, MUST be encapsulated into a new message of the form defined in the message format extensions document. This encapsulation will result in a message with a single body part MESSAGE with an appropriate transport encoding. If the message format document is changed to make illegal nested transport encodings, this issue will have to be revisited.

The strict definition of the transport encoding to be used was discussed, and the consensus of the group was that a strict specification of which transport encoding to use could not easily be made to work. The best approach for an implementor is to scan the document and determine statistically whether it would be better to encode the entire message in a Base 64 encoding or escape the few offending characters via a quoted printable encoding.

Defining the "New" SMTP

The Working Group began work on the new SMTP version. It was argued that the greatest change necessary is to define a negotiation mechanism for new capabilities. Some of these capabilities are:

- 8 bit Text
- Binary Transport
- Authentication
- Delivery Notification
- Message Size Negotiation
- Explicit Batch Mode

Several modifications to the protocol were suggested that were feature-independent. Among the suggestions were:

A Second TCP Connection for Data

A second data connection would make it possible to do data checkpointing, and would reduce the cost of sending binary data. Drawbacks include the overhead of opening and tearing down a second channel, and running SMTP over non-tcp single-channel protocols such as X.25. The Working Group decided not to pursue this approach. The cost of sending binary data over the existing channel either by escaping or byte counting was found to be preferable over the cost of opening a new TCP connection. Checkpointing in FTP is still not widely used, and is considered by this group to be of dubious value.

Asynchronous Operation

Currently SMTP commands are batched by several implementations and sent in a single packet to save round-trips. This has been demonstrated to work with known SMTP implementations. An extension to tag the data and the commands to allow full asynchronous operation was proposed. This offers very significant improvements in throughput by reducing packet per verb to control packet per session in the best case. The Working Group debated this point and concluded that full asynchronous operation would push SMTP into a not-so-simple MTP.

A Negotiation Infrastructure

The group agreed that a mechanism needs to be defined to allow the extension of SMTP. The current approach of this Working Group has been to add functionality via the addition of new verbs. While this approach is seen by some to be the strait forward answer, using new verbs can cost significant time in round-trip delay while playing a network version of the old card game "go-fish". Other suggestions included a telnet like negotiation.

The Working Group began exploring features of a new negotiation mechanism for the SMTP protocol. Among the possible goals are:

- Symmetry – should the receiver and the sender both request an option?
- Batchable – should more than one option be negotiated at a time?
- Duration – per-session, per message, or per-recipient?
- Default behavior - should the default be better than current SMTP service?

Symmetry: Symmetry was suggested as a means to allow authentication of the sender by the receiver. At this time there is no formal authentication mechanism, and the negotiated use of CAT or Kerberos was seen as a good thing. After lengthy debate, the group decided that authentication of the sending SMTP in a store and forward network was of dubious value and was not worth the added complexity a symmetric negotiation entails.

Batchable: Batching negotiated parameters offers great savings in round-trip times. It is not clear how this would work in practice, but the group felt that this was a good goal.

Duration: This was a tricky subject. Currently SMTP does not provide any information about the users environment. Any use of per-recipient or per-message requires the keeping of more knowledge about the end-user than the system has now. It was not clear to the group that any per-recipient options exist that could not be duplicated by a local delivery agent.

Default: This turned into a no-brainer. The group unanimously felt that the new SMTP needed to be backward compatible, and in the case of complete failure of any negotiation, the mail would continue to go through as specified in RFC 821 and HR.

The meeting concluded with the discussion of several specific negotiation strategies. Several attendees volunteered to write up proposals for negotiation mechanisms. This discussion will be continued on the mailing list.

Attendees

Peter Boos	peterb@bnr.ca
James Conklin	conklin@bitnic.educom.edu
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Ned Freed	ned@innosoft.com
Olafur Gudmundsson	ogud@cs.umd.edu
Russ Hobby	rdhobby@ucdavis.edu
Neil Katin	katin@eng.sun.com
Darren Kinley	kinley@crim.ca
Jim Knowles	jknowles@trident.arc.nasa.gov
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@turbo.bio.net
Jack Liu	liu@koala.enet.dec.com
Joseph Malcom	jmalcom@sura.net
Leo McLaughlin	ljm@wco.ftp.com
Keith Moore	moore@cs.utk.edu
Michael Patton	map@lcs.mit.edu
Jan Michael Rynning	jmr@nada.kth.se
Mark Saake	saake@llnl.gov
Harri Salminen	hks@funet.fi
Keld Simonsen	keld.simonsen@dkuug.dk
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.4 Internet Message Extensions (822ext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

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

Internet Drafts:

"Mechanisms for Specifying and Describing the Format of Internet Message Bodies", 06/18/1991, Nathaniel Borenstein, Ned Freed <draft-ietf-822ext-messagebodies-00.txt, .ps>

"Mnemonic Text Format", 07/08/1991, Philippe-Andre Prindeville, Keld Simonsen <draft-ietf-822ext-qreadable-02.txt>

"Character Mnemonics and Character Sets", 07/08/1991, Keld Simonsen <draft-ietf-822ext-char-00.txt>

INTERIM MEETING REPORT

Reported by Greg Vaudreuil/CNRI

822EXT Minutes

Agenda

1. Character Set Selection
   - Status and Input to the ISO 10646 process.
     - Unicode <=> ISO 10646 Union?
     - Use of CO and C1 codespace.
   - Selection of "Common" character sets or schemes.
     - ISO 8859-1, ISO 8859-n, Profiles for the use of ISO 2022?
     - Specifying "requiredness".
   - Specification of 8 bit character sets in headers.

Minutes

Unified Character Set

1. Administrative

   At last word, the ISO DIS 10646 received 9 YES votes and 14 NO votes, and work is proceeding to resolve the remaining issues. An unofficial but promising effort is the work underway to unify ISO DIS 10646 and Unicode, another scheme for a global character set. This effort is being conducted outside the normal ISO process. This Working Group was asked to discuss this effort and endorse it if possible. The Working Group discussed this effort, and agreed that the efforts to combine Unicode and 10646 were in fact positive.

2. Technical

   The unification of ISO DIS 10646 and Unicode requires the resolution of several technical issues. The primary issue, tentatively resolved involves "Han unification" a scheme that re-uses many of the graphics of the various Kanji character sets. Other issues involve the use of CO and C1 codespace. The use of C0 and C1 codespace involves transport issues and this Working Group was asked for its input.

   C0 codespace consists of the spaces between 0 and 31 and 127, traditionally used for control characters. There is a proposal to use this space in the second octet of a multi-byte character for graphic characters. The Working Group discussed this and rejected the use of this space. A graphic character in the C0 space will likely be interpreted by a transport protocol as a control character. Many transport protocols which interpret in-band data such as SMTP may behave unpredictably in this situation. One example is where the sequence of graphics legally sent by a 8 bit sender may be mis-interpreted by a 7 bit receiver after bit stripping as a 13-10-46-13-10 sequence terminating the SMTP session prematurely. Other related anomalies were envisioned. Unless all transport protocols are made aware of the multi-byte nature of the data, an unlikely occurrence any time soon, reuse of C0 space is not recommended.

C1 codespace consists of the spaces between 128-150, space that may be interpreted as control characters if the high order bit is stripped. ISO 8859-n character sets, and the current 10646 proposal reserve this space for control characters only, with an eye toward backward compatibility with 7 bit systems. The Working Group discussed this and concluded that use of C1 codespace could be used for graphics if transport protocols could be relied upon to never strip the high order bit and interpret the resulting character as control sequences. The Working Group did not make a specific recommendation, only that the use of C1 space to compact a character set was a positive thing, and future evolution transport protocols should support the use of this space for graphics.

Common Character Sets

In the absence of a single international standard character set,the Working Group needs to profile the use of a limited number of the 200+ character sets in use worldwide to facilitate interoperation. Keld Simonsen gave an overview of the current character sets in usage.

ISO 7 bit family:
- ASCII
- National Versions
  - 10 National use
  - 2 Alternate rep **# $**
- ECMA registry
  - 7, 8, 16 bit
  - ISO 2022 shifts

ISO 8 bit 8859 family:
- 1 char = 1 octet
- ASCII in pos 0-127
- Pos 160-255
  - Latin sets (5)
  - Cyrillic
  - Greek
  - Arabic
  - Hebrew

ISO 6937-2 family 8/16 bit:
- 6937-2, T.61
- Non-Spacing accents

- 1 char = 1 or 2 bytes
- about 330 graphical chars

Vendor 8 bit sets
- DEC-MCS
- HP Roman8
- IBM PC codepages (5)
  - Uses also 128-159 (C1)
- IBM EBCDIC]
  - Many versions
  - Not ASCII Compatible

16 bit char sets
- Japanese: JIS 0208, 0212
- Chinese: GB 1980
- Korean:
- Japanese 8/16 bit: Shift JIS
- Unicode: New vendor charset unifies CN, JP, KO sets
  - Incompatible with ISO

Multi-byte:
- EUC: Extended UNIX code
  - ISO 2022 shiftingS
  - SS1 SS2 SS3
  - 4 char sets
  - 8/16/24 bits

32 Bits:
- ISO 10646
  - Also usable in 8, 16, or 24 bit compaction methods
  - Proper encoding subsets: ASCII and ISO 8859-1

Control Character Sets:
- ISO 646: 0-31, 127
- ISO 6429: 0-31, 127-159
- EBCDIC: as ISO 646

Several ideas were batted around, including strict use of ISO2022, profiling language to character set mapping, and the use of "preferred" character sets. The Working Group felt that the best approach was to codify existing practice in the interim, pending adoption of an "international" character set. This existing practice was reduced to the following.

If possible, use ISO 8859, with the lowest version number possible, (i.e., use 8859-1 (Latin 1) over 8859-10? (Latin 5?). If the characters needed are not in the 8859 sets (i.e., Kanji) use the 2022 character switching standard, declaring 2022 in the header of the document. While this may lead to the use of any of the many characters in the ECMA registry, the Working Group felt that in practice, only the current Oriental mail systems will use the ISO 2022 system and only with limited character sets.

Use of Non-ASCII Character Sets in Headers.

What a mess! The attendees of this meeting spent over an hour working on various schemes for indicating character sets in the headers of a message other than ASCII. It was identified as a requirement that the fields defined as TEXT be able to have variable character sets. While this goal was stated, no mechanism for the implementation was agreed upon.

A modification of the BNF notation was suggested by Keld Simonsen.

```
CHAR-EIGHT     = <any Eight-bit character>; (0-377, 0.-255)

qtext          = <any CHAR-EIGHT excepting <">,"\" & CR, and
               including linear-white-space>

quoted-pair    = "\" CHAR-EIGHT

text           = <any CHAR-EIGHT, including bare CR & bare LF but
               NOT including CRLF>
```

This notation was accepted by the attendees of the meeting, however several problems were identified and not resolved.

- Identification of the header character set and the need for conversion, and
- Encoding the header character sets in 7 bit transport format.

It was not clear how a conversion gateway would know that the header was 8 bit and needed encoding. A suggestion accepted by the group was that the use of the new BNF requires the use of a header-charset header line. This additional header adds complexity to user agents and conversion gateways by requiring two passes of the header to determine and convert the header into a passable or readable form. It was felt that this was in-elegant but do-able.

Several proposals were discussed for encoding the 8 bit text strings when 7 bit transport was required. It was accepted that this was a hard requirement.

1. Variable Substitution

   One proposal for the insertion of 8 bit text was to substitute a variable name in the header for each text string needing 8 bit characters. The variable could then be defined elsewhere in the header, including the encoded actual string and a token indicating the character set. This was rejected as messy and difficult to implement in current user agents.

2. Message Encapsulation

   Encapsulate the mail message using the message type body part and a suitable transport encoding, preferable quoted-printable. This proposal is controversial among at least one implementor of the message format standard as having excessive complexity for the user agent. It is not clear the encapsulated message will be permitted to have a transport encoding.

3. Encoded Text Fields

   This proposal would specify a standard encoding for the header fields, possibly quoted-readable or quoted-printable and identify this fact in a header-transport-encoding header or the header-character-set header.

Conclusions

While no one was happy, the group tentatively agreed not to permit 8 bit text in the headers. The only reasonable way to encode 7 bit text was to encode the text fields, and insert a new header line. With this overhead the group agreed that while not ideal, a requirement that extended character sets should always be encoded, eliminating the need for intermediate gateways to parse and convert the headers.

Attendees

Byungnam Chung	bnchung@sokri.etri.re.kr
Johnny Eriksson	bygg@sunet.se
Phillip Gross	pgross@nis.ans.net
Robert Kummerfeld	Bob@cs.su.oz.au
Jan Michael Rynning	jmr@nada.kth.se
Keld Simonsen	keld.simonsen@dkuug.dk
Peter Svanberg	psu@nada.kth.se
Gregory Vaudreuil	gvaudre@nri.reston.va.us

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

822EXT Minutes

Agenda

- Review of Implementations.
- Review of the Message Format Document.
- Type/ Subtype Clarifications.
- Resolution of the Encapsulated Message Format.
- State and Status of the Mnemonic Encoding and Character Set Document.

The meeting began with a review of the work of implementors on the Message format documents. Four attendees had implemented from the document, with at least two others not in attendance. Three of the known implementations were mail readers and two were gateway products.

A review of the Message format document was begun. Due to the limited time the Working Group had in face to face session, it was agreed to only discuss those points which were substantive and potentially controversial.

Issue 1: Character set designation in a new header line. There was dissatisfaction with indicating the character set only as a subtype of text. One implementor found it useful to have a character set is non-text objects. A review of the reasons for putting character sets in a sub-type resulted in no objections to moving this information into a new header.

Issue 2: The character designation discussion opened up a issue regarding the syntax of optional and required fields in the type designation. An objection, with request for explanation, was made to the split between the type and the subtype field. The original rational for this hierarchy, to aid gateways and mail readers in "doing the right thing" with unrecognized content-types is less compelling in light of the realization that the content-type is little more than a hint as to which transfer encoding should be used, and there are many cases where selection of a transfer-encoding will result in a less efficient choice than could be made. Other participants argued that the type field offers a valuable help to mail readers which try to do something with unrecognized subtypes. Resolution was reached with the observation that the type/subtype notation could be interpreted by a mail reader as a single level content-type. The proposed standard version will use the two level hierarchy.

Issue 3: The syntax of type/subtype is not clean. Some subtypes have mandatory fields, such as text, and others have an attribute pair notation for options. Much of this notation is a holdover from the RFC 934 multi-part specification. The Working Group re-affirmed the preference for simplicity and elegance over compatibility with that previous specification. After discussion the following was proposed and accepted overwhelmingly: required parameters for a type or subtype should be included as part of that content-type header line, and optional parameters should be put in a new header line per option. It was noted that several options may be used by many body-types and so there is a reduction of complexity. Among the new optional parameters suggested were content-filename and content-conversion. Other fields were left up to the document editors to define as needed to clean up the type/subtype syntax.

After this warmup the Working Group moved on to the issue of nested transfer encodings. After some initial implementation experience, it has become clear that allowing arbitrary nested body parts each with a transfer encoding, causes a significant increase in the complexity of mail readers. No one disagreed that nested encodings are possible on almost all know platforms. People realized that some of this complexity could be pushed off onto gateways, but after exchanging sendmail horror stories for 30 minutes, it became clear that having gateways mung messages was really "sickening" to many in attendance.

In return for this complexity, two key advantages are realized. The first is the ability to allow 8 bit text in the headers of messages, provided the message is encoded as a whole a transfer encoding. Without the ability to nest the encodings, including headers in this fashion would only be possible for simple messages with no encoded body parts. The second advantage is the ability to specify a simple encapsulation for gateways between diverse transport environments as well as non-smtp based environments. By allowing the encapsulation of a binary or 8bit message without requiring part by part examination and conversion, the need for a gateway to parse complex multi-part messages and understand the content-types is significantly reduced.

After two hours of talking, a strawman poll was taken in which the group was fairly evenly split between those interested in preserving the nested encodings and those who did not. In the interest of making progress with an issue that has defied consensus, the Chair proposed a compromise position. Because the group could agreed that it is far easier to drop the nested encodings in a future version than to add it the following was stated.

POSITION: The Proposed Standard version of the message format document will allow nested encodings. If in implementing this specification, it is determined by the group that it is either technologically unfeasible or excessively cumbersome, it will be dropped at the Draft Standard stage.

Beginning the second session, the Working Group discussed the two documents by Keld Simonson. The first of these documents describes many character sets, both ISO standards and others that are of interest to the Internet Community. Furthermore, this document defines naming conventions for both the characters and character sets. This naming functionality is not duplicated in any other registry, although it is expected that a similar ISO registry will be set up at some time. This document uniquely names the character sets intended to be used in the Message Format document and other IETF work. With the addition of a provision that future character sets will be registered with the IANA, the Working Group endorsed it's publication as an informational document.

The second of Simonsen's documents, the mnemonic encoding document was discussed in terms of the message format document. This document uses the character names in the character set document to define a readable escape sequence for characters which cannot be represented in ASCII. This has been proposed as an alternative to the use of a native character set and transport encoding. The Working Group thought this was a wonderful idea, and endorsed it's publication as an experimental protocol. The Message Format document will reference this as a mechanism for sending 8 bit text where it is known the receiver is only capable of reading text on an ASCII or invariant 646 display.

The Working Group discussed the need to resolve the problem with the growing anarchy in email error message, both in terms of the interpretation of RFC822 headers for the designation of error recipients, and the format of those messages. It was felt that this work should be begun in two areas, a revision to RFC 822, to clarify ambiguous sections, and defining a standard machine-parsable error message using the message format standard. This effort began with a call for ideas and strawman proposals on the Working Group. Due to lack of time, this was not discussed further in this meeting.

Attendees

Philip Budne	phil@shiva.com
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Ned Freed	ned@innosoft.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Russ Hobby	rdhobby@ucdavis.edu
P. Allen Jensen	allen@audfax.audiofax.com
Neil Katin	katin@eng.sun.com
Douglas Kerr	dougk@mtxinu.com
Darren Kinley	kinley@crim.ca
Jim Knowles	jknowles@trident.arc.nasa.gov
Vincent Lau	vincent.lau@eng.sun.com
Eliot Lear	lear@turbo.bio.net
Jack Liu	liu@koala.enet.dec.com
Joseph Malcom	jmalcom@sura.net
Louis Mamakos	louie@ni.umd.edu
Leo McLaughlin	ljm@wco.ftp.com
Keith Moore	moore@cs.utk.edu
Michael Patton	map@lcs.mit.edu
Jan Michael Rynning	jmr@nada.kth.se
Harri Salminen	hks@funet.fi
Keld Simonsen	keld.simonsen@dkuug.dk
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Wobus	jmwobus@suvm.acs.syr.edu

3.1.5 Network Database (netdata)

Charter

Chair(s):

Daisy Shen, daisy@watson.ibm.com

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Review and approve the Charter, making any changes necessary. Examine needs, resources for this network database protocol and define the scope of work. Begin work on a framework for the solution. Assign writing assignments for first draft of the document.
Done	First draft to be completed.
Aug 1991	Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting.
Dec 1991	Continue revisions based on comments received at meeting and e-mail. Start making document an Internet Draft.

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

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

Internet Drafts:

"Network Database", 06/26/1991, Daisy Shen <draft-ietf-netdata-netdata-00.txt, .ps>

CURRENT MEETING REPORT

Reported by Daisy Shen/IBM

NetData Minutes

This was the second meeting of the Working Group Chaired by Daisy Shen. The meeting Agenda is shown below:

- Review the Charter.
- Give the History of the SQL Access Group.
- Presentation of the Draft.
- Discussion of the Draft and Problems That are Related to the Subject.
- Discussion of the Effort of Other Vendors and OSF Related to the Subject.
- Future Work.

Reviewed the Charter

All the members felt that we met the Charter and the Milestones on schedule as of June 1991.

History of the SQL Access Group

All the members agreed that the SQL Access group doesn't seem to be willing to work with us. The SQL Access group gave us a two-page document which is useless. Although, our group and the SQL Access group seem to have some similar interests, we are solving the problem via TCP/IP while they use the ISO standard. Actually those two have different goals.

Presentation and Discussion of the Draft and Related Problems

The first draft was presented for those who had not read it or had questions. There were some issues discussed and solved.

- SQL statements

  The protocol works with both SQL static and dynamic statements.

- Transport means

  Our protocol will apply to both TCP and UDP. It is one of the parameters that users have to specify. The default value is TCP.

- Data conversion

  The header information along with data will be treated as a string by RPC/XDR and the protocol will use ASN format which is an IOS standard to represent the reply data.

- Security

  Kerberos is not required. Users can have options to use or not to use Kerberos. For those machines who don't have Kerberos, the protocol will provide security which relies on uid and gid. Authentication will be done on the RPC level.

- RPC

  Before there is a standard RPC, the protocol will work with all versions of RPC. However, all members of NETDATA believe that we should start an RPC Working Group standardizing RPC, especially a secured RPC.

The Network Management Operational Statistics Working Group had some discussion on database issues. We should let them know that our group is working on the problem. It would be nice if they could tell us their requirements. Perhaps, we can help them solve their problem.

Discussion of the Effort of Other Vendors and OSF Related to the Subject

We could not find any more information on the effort of other vendors and OSF. We will continue to search.

Future Work

- Update the first draft and make it a protocol only.
- Create a new draft for the structure and implementation issues.
- Get a better understanding on ASN.1.
- Get more familiar with ISO standard.
- Look into SQL statement's additional negotiation options.
- Look into starting an RPC Working Group.
- Connect the Network Management Operational Statistics Working Group.
- Find volunteers to do the second version of the implementation.

Attendees

Chi Chu	chi@sparta.com
John Cook	cook@chipcom.com
Bill Fardy	fardy@ctron.com
Shari Galitzer	shari@gateway.mitre.org
Russ Hobby	rdhobby@ucdavis.edu
Mike Janson	mjanson@mot.com
Vincent Lau	vincent.lau@eng.sun.com

Louis Leon	osll@emuvm1.cc.emory.edu
Leo McLaughlin	ljm@wco.ftp.com
Michael Patton	map@lcs.mit.edu
Geir Pedersen	geir.pedersen@use.uio.no
Daisy Shen	daisy@watson.ibm.com
Erik Sherk	sherk@nmc.cit.cornell.edu
Anil Singhal	
David Waitzman	djw@bbn.com

3.1.6 Network Fax (netfax)

Charter

Chair(s):

Mark Needleman, mhn@stubbs.ucop.edu

Mailing Lists:

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

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Review and approve Charter making any changes deemed necessary. Refine definition of scope of work to be accomplished and initial set of RFC's to be developed. Begin working on framework for solution.
Done	Continue work on definition of issues and protocols. Work to be conducted on mailing list.
Aug 1991	First draft of RFC to be completed. To be discussed at IETF meeting and revised as necessary.

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

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

CURRENT MEETING REPORT

Reported by Mark Needleman/U California

NETFAX Minutes

The Netfax Working Group met on August 1, 1991 at the IETF meeting in Atlanta. The primary goal of the meeting was to attempt to get consensus on the draft paper written by ISI on a common image file format for fax in the internet making use of TIFF encoding.

Alan Katz of ISI gave a short presentation about the paper, highlighting its major points. Discussion was then held. Consensus was achieved among those present that this was the proper way to go in the short-run with the intention of more closely examining the possibility of using ODA as it became more prevalent in the internet.

Mark Needleman and Alan Katz agreed to work together to turn the draft paper into a format suitable to being published as an Internet Draft. It will then be posted to the NETFAX mailing list for comment after which the process will be started to get it into an Internet Draft so that the wider community can comment on it.

Clifford Lynch agreed, after the draft was revised to to make sure it was distributed to organizations involved in library projects transfering images over the internet and to try to get as many of these groups as possible to implement it for interoperability testing.

A discussion was then held on defining mechanisms for transporting fax in the internet and making use of the work of the Internet Message Extensions Working Group. Ned Freed gave a short presentation on the paper that group had produced. Alan Katz agreed to come up with a name for a tiff netfax body content type header. He will do this within two weeks so that it can be incorporated into the SMTP Extensions paper. Otherwise it will need to become a separate RFC as an add-on to that document.

Carl Malamud lead a discussion on addressing fax in the internet and what fax addresses should look like. Carl agreed to put together for the mailing list a proposal that could then be discussed and become the basis for a proposed RFC on the subject. Ned Freed also agreed to post the attributes that Innosoft uses to the list.

Alan Katz agreed to look at the possibility of defining body-type parts for the cover page and what elements would be required. This could then become another body-type as defined by the SMTP Extensions Working Group.

Attendees

Philip Budne	phil@shiva.com
John Cook	cook@chipcom.com
Tom Easterday	tom@cic.net
Ned Freed	ned@innosoft.com

Russ Hobby	rdhobby@ucdavis.edu
P. Allen Jensen	allen@audfax.audiofax.com
Alan Katz	katz@isi.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Jack Liu	liu@koala.enet.dec.com
Clifford Lynch	calur@uccmvsa
Carl Malamud	carl@malamud.com
Louis Mamakos	louie@ni.umd.edu
Keith Moore	moore@cs.utk.edu
Mark Needleman	mhn@stubbs.ucop.edu
William Nowicki	nowicki@legato.com
Geir Pedersen	geir.pedersen@use.uio.no
Jon Postel	postel@isi.edu

3.1.7 Network News Transport Protocol (nntp)

Charter

Chair(s):
Eliot Lear, lear@turbo.bio.net

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

Description of Working Group:

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

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

o User level and site designated authentication methods;

o Binary transfer capability;

o Minimization of line turnaround; and

o Stronger article selection capability.

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Eliot Lear/Intelligenetics

NNTP Minutes

There were three meetings of the Network News Transport Protocol Working Group.

There was also an area meeting which included the NNTP, LIST, 822-EXT, and SMTP Chairs, along with the Area Director. This is to say, Eliot Lear, David Lippke, Greg Vaudreuil, and Russ Hobby.

The following items were explored at the various meetings:

1. Differences Between Mail and News.

   We should consider moving towards a common user interface between mail and news. Similar message formats have made this possible in the past. With the advent of a new message format for mail, news will need to adopt some similar standard pretty quickly (ARE YOU READING THIS, NEWS READER PEOPLE?????). There was discussion of moving to unite the news and mail formats. While conceptually it sounds like a good thing, the details need to be kinked out, and the question needs to be discussed to death.

   Russ asked what a document be put out that describes the current news architecture. Erik Fair has volunteered to write that document. That document should almost certainly include a safe way to gateway news and mail. Conceivably another document will issue from the area recommending a course of action.

2. News Reader Capabilities.

   First, it was the consensus of the group that this topic is really part of the Charter of the NNTP group; we're just considering splitting new functionality into a separate protocol.

   The current version of the draft contains very little enhancements in the area of news reader capabilities. This is because the Chair does not have any real concrete language in front of him from what this group wants. The consensus, therefore, was to push on with the transport document,and explore further the reader issues, and in particular how this relates to Item 1.

   If we do produce an NNRP document, we must be careful that by its nature it would steer development away from useful areas (Ittai Hershman's paraphrased comments). In this vein, if we do produce a document, we should consider it an experimental effort rather than a standards track effort.

Along the lines of a news reader protocol, Stan Barber brought along a one page shopping list of items he would like to see in a reader protocol. We discussed how to define a search command so that it would be generally useful. Arguments for and against a specific syntax and mechanism were heard.

3. Authentication

   Theodore Tso is now the official "stuckee" for the SAAG in the NNTP Working Group. Issues of Common Authentication Technology (CAT) were discussed, particularly at the Thursday meeting. Text needs to be written into the document to take advantage of CAT. We are facing a problem with CAT because NNTP is one of the first protocols to use it. Currently CAT can only be used to access Kerberos and DEC SPX. Jeff Schiller suggested that a simple challenge/response method would be acceptable if someone did the footwork. Clear text, however, seemed to be right out, to the point where it was thought that the SAAG might hold things up. Jeff also discussed the evils of negotiating security methods.

   It turns out that some of the logic that was applied to mail standards can be applied to news. If we do, in fact, move the transport document to proposed standard, the impetus for authentication in the transport is greatly diminished.

4. Transporting Binary and Mixed Message Format.

   It turns out that simply adopting the mail standards as news standards may be a bit painful. With the introduction of binary, there needs to be a new canonical form. This in itself would be a minor irritation; however, the new mail format allows for mixed binary and text. This means that it could be necessary to switch between binary and text canonical forms in a single message. This makes transport a nightmare, and is a good argument for encoding. On the other hand, possibly the new binary canonical form might be able to handle the problems. Interested parties are URGED to read the draft mail documents and the archive of messages leading to their production.

5. CCITT

   Harri Salminen circulated a draft document that is CCITT's version of netnews. The document may be retrieved from nic.nordu.net, via anonymous FTP. Your comments are, of course, solicited.

6. Problems with the Current Document

   Several people have sent notes pointing out formatting problems, grammatical errors, and certain inconsistencies (like SIMPLE authentication descriptions). Please mail all such complaints directly to the Chair, and not to the list.

   It was the consensus of the group that the IMAGE and BINARY options be combined into a FORMAT option. Eliot Lear will write some text in to this effect. It was also agreed that the COMPRESSION and DATE commands would be removed, and that the NEWNEWS command be extended to deal with DATE's purpose (which is to say that NEWNEWS will both accept and deliver a cookie instead of a date). Text to be written and argued.

State diagrams need to be completed.

Default behavior needs to be defined and mandated.

We discussed eliminating the OPTION command. The problem with eliminating the OPTION command is that it gets hard to batch verbs, and we concluded that batching such things was a good idea.

7. Making the IETF lists available to the IETF via netnews.

   This issue was brought up both in the Working Group and in the area meeting. Some action is expected in this area Real Soon Now (tm). Social issues were discussed in the Wednesday meeting regarding the perceived stigma from which news suffers.

8. News MIB

   Russ Hobby stated that he would not require a News MIB from us. However, several people have indicated some interest in managing news objects, particularly Jim Thompson (not present at Atlanta). Jim should proceed to take comments and write up a document. One should be careful to study which functions are ubiquitous throughout the Internet, and which are implementation specific.

9. Timetable

   August 31, 1991 - We would like to see the NNTP document become an Internet Draft. All this does is expose the document to the Internet community. It can be changed from within the Working Group after that point.

   November, 1991 - Get architecture document out as an informational RFC.

   December, 1991 - After considering comments and making appropriate changes, let the NNTP document proceed to proposed standard.

   It is hoped that some code will be ready by December.

10. Multicasting Netnews

    Brief mention was made on a research effort under way to explore the possible use of multicast packets as a way for distribution of news. Interested parties should contact the Chair directly.

11. Next Meeting

    No next meeting date has been set as of yet. Depending on how we proceed with a news message format, we may meet at Interop (October).

12. General Information

    If you wish to be added to the ietf-nntp mailing list, you should send mail to ietf-nntp-request@turbo.bio.net.

    Drafts and message archives can be gotten from turbo.bio.net via anonymous FTP in the ietf-nntp directory. The format of draft document filenames is documentname.format-type.

Attendees

Stan Barber	
Robert Enger	enger@seka.scc.com
Erik Fair	fair@apple.com
Ned Freed	ned@innosoft.com
Olafur Gudmundsson	ogud@cs.umd.edu
Ittai Hershman	ittai@nis.ans.net
Russ Hobby	rdhobby@ucdavis.edu
Neil Katin	katin@eng.sun.com
Eliot Lear	lear@turbo.bio.net
David Lippke	lippke@utdallas.edu
Joseph Malcom	jmalcom@sura.net
Chris Myers	chris@wugate.wustl.edu
Michael Patton	map@lcs.mit.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
Jan Michael Rynning	jmr@nada.kth.se
Harri Salminen	hks@funet.fi
Theodore Tso	
Gregory Vaudreuil	gvaudre@nri.reston.va.us

3.1.8 Network Printing Protocol (npp)

Charter

Chair(s):
Glenn Trewitt, trewitt@nsl.dec.com

Mailing Lists:
General Discussion: print-wg@pa.dec.com
To Subscribe: print-wg-request@pa.dec.com
Archive:

Description of Working Group:

The Network Printing Working Group has the goal of pursuing those issues which will facilitate the use of printers in an internetworking environment. In pursuit of this goal it is expected that we will present one or more printing protocols to be considered as standards in the Internet community.

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Glenn Trewitt/DEC

NPP Minutes

Agenda

- Walk through LPR/LPD Protocol draft.
- Walk through Printer Access Protocol draft.

There was a lot to go through on both drafts. Either could have occupied the entire meeting. Because there were very few attendees interested in PAP, compared to LPR/LPD, we dealt only with that in the meeting. PAP was covered outside the meeting (actually, only Jim Jones, the author, and I showed up). This meeting was well-attended, and most of the people had read the draft before the meeting. Early preparation does pay off!

LPR/LPD Protocol Draft

We finally decided to make a clean break between the existing protocol and this specification, so it would be absolutely clear what was "in" the old and new. We paid full attention to old vs. new interoperability issues. Specifically:

1. The "new" features were removed from the "old" commands:
   - Zero-length for data files with EOF indicated by connection close.
   - Data file vs. control file ordering.
2. To provide these (very needed) features, we added two new opcodes:
   - Receive Data File With Unknown Length - receive a data file, terminated by connection close.
   - Receive Control File First - receive a control file before its associated data files.

For the new commands, we defined explicit rules for how a client should employ them, to achieve interoperability with both old and new LPD daemons. The new draft will have a very solid line between "old" and "new" LPD features. A new draft will be available in September.

Printer Access Protocol

Most of the work here was aimed at getting everything in the document that was appropriate for a protocol RFC – full information so that the protocol is implementable from the document, and removing implementation-specific verbiage. There is still significant work to do on this document.

Attendees

Charles Bazaar	bazaar@emulex.com
David Bridgham	dab@asylum.sf.ca.us
Philip Budne	phil@shiva.com
Cerafin Castillo	cec@emulex.com
James Jones	jones@regent.enet.dec.com
Charles Kimber	ckimber@dcp.com
Stev Knowles	stev@ftp.com
Louis Mamakos	louie@ni.umd.edu
Leo McLaughlin	ljm@wco.ftp.com
Keith Moore	moore@cs.utk.edu
Michael O'Dell	mo@bellcore.com
Richard Smith	smiddy@pluto.dss.com
Glenn Trewitt	trewitt@nsl.dec.com
Peter de Vries	peter@wco.ftp.com

3.1.9 TELNET (telnet)

Charter

Chair(s):

Dave Borman, dab@cray.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 6 years of technical advancements, and will determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the TELNET options, and decide which are still germane to current day implementations of the TELNET protocol.

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

Goals and Milestones:

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

Internet Drafts:

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

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

"Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-02.txt>

"Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-02.txt>

Request For Comments:

RFC 1116 "Telnet Linemode option"

RFC 1184 "Telnet Linemode Option"

CURRENT MEETING REPORT

Reported by David A. Borman/Cray Research, Inc.

TELNET Minutes

The Telnet Working Group met the morning of Tuesday, July 30, 1991.

An initial Agenda of possible topics included:

- Authentication Option
- Encryption Option
- Environment Option
- CAT (Common Authentication Technology)
- Possible Future Telnet Work
  - 8 bit NVT
  - International Characters
  - Option Negotiation Loop Avoidance
  - Compression
  - Terminal Types
  - MIB

The actual discussion that followed focused only on the Authentication option.

The first thing that was started was a quick walk-through of the draft Authentication option to do some final editing, before approving the draft for recommendation for being published as an RFC.

```
Page 1:
 The AUTH_WHO_CLIENT and AUTH_WHO_SERVER names will be changed to
 AUTH_CLIENT_TO_SERVER and AUTH_SERVER_TO_CLIENT to more accurately
 describe who is authenticating who.

 The words "(TCP LISTEN state)" will be added after "...that did the
 passive TCP open," to clarify things.

Page 2:
 Mid page, "authentication-type-list" will be changed to
 "authentication-type-pair-list" for consistentcy.

 The descriptions of IS and REPLY will be rewritten.

Page 3:
 In the first sentence, "has two values" will be changed to
 "is two octets".
```

This is about as much as was accomplished on the walk-through, as side discussions took us off on other more general issues.

A major point of discussion was:

1. Should we continue with the authentication option, and
2. Should the authentication option be providing "negotiation" for the type of authentication to be used?

The argument for the negotiation was that it provides a mechanism for the client and server to agree upon what type of authentication will be used. The argument against it was that the client should just pick one, and it probably will only have one type, and the server would either accept or refuse it. There was also fear that having a negotiation scheme would allow a weaker form authentication to be agreed upon that the user is willing to use.

After much discussion, it was decided that we would:

1. Leave the draft the way it is.
2. Add a description of the security concerns of doing negotiation of the authentication mechanism.

A shorter discussion was held on whether the IS and REPLY commands could be replaced with a single DATA command. It was decided that there was no benefit in changing it, so it was left as is.

Another discussion was whether we should be doing the Authentication option at all, in light of the work starting up in the Common Authentication Technology (CAT) Working Group. It was decided that since the CAT group is just starting up, and the Authentication option is already being implemented and used to get real work done, and that the Authentication option has the ability to evolve into CAT, that we would continue.

The draft will be modified as stated above, and circulated for one more round of review before being sent to the IESG.

Attendees

Steve Alexander	stevea@i88.isc.com
David Bolen	db31@nis.ans.net
David Borman	dab@cray.com
Johnny Eriksson	bygg@sunet.se
Barbara Fraser	byf@cert.sei.cmu.edu
Kenneth Goodwin	goodwin@psc.edu
Alton Hoover	hoover@nis.ans.net

Charles Kaufman	kaufman@dsmail.enet.dec.com
Mark Lewis	mlewis@telebit.com
John Linn	linn@zendia.enet.dec.com
Louis Mamakos	louie@ni.umd.edu
Matt Mathis	mathis@psc.edu
Greg Minshall	minshall@wc.novell.com
Clifford Neuman	bcn@isi.edu
Jason Perreault	perreaul@interlan.interlan.com
John Pickens	jrp@3com.com
P. Rajaram	rajaram@sun.com
Jan Michael Rynning	jmr@nada.kth.se
Mark Saake	saake@llnl.gov
Emil Sturniolo	emil@doe.dss.com
Theodore Tso	
Jesse Walker	walker@eider.enet@decpa.dec.com
David Ward	dward@chipcom.com

3.2 Internet Area

Director(s):

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

Area Summary reported by Noel Chiappa

Three Working Groups in the Internet Area met at the Atlanta IETF; the Router Requirements WG, the PPP Extensions WG and the IP over Appletalk WG. The results of these WG meetings will be briefly discussed below.

In addition, four BoF's in the Internet Area also met; these were the Variable Width Subnet Masks Bof (led by Philip Almquist), the IP Logical Networks BoF (also led by Philip Almquist), the IP Address Hacks BoF (led by Noel Chiappa), and the Internet Mobile Hosts BoD (led by Steve Deering). (The latter two groups met together.)

Finally, the Router Discovery protocol was moved forward to Proposed Standard status at this IETF, and that WG has now concluded it's work.

Router Requirements WG

The Router Requirements WG held what it hopes will be its final meeting at this IETF. A new (almost complete) draft of the Router Requirements document was discussed in detail, and new versions of the TOS document, the Next Hop document, and Route Leaking document were reviewed.

In an attempt to make definitive forward progress, the chair requested that only those who had prepared by reading the document speak up, and attempted to limit discussion on topics which had been thoroughly gone over in previous meetings.

PPP Extensions WG

The first order of business at the PPP Extensions WG meeting was to pick a new chair, as the previous chair (Stev Knowles of FTP Software) had been unable to give this group the attention he felt it needed. Brian Lloyd from Telebit agreed to serve as chair in replacement.

The group has 3 documents out: a basic LCP (including the basic IP support, IPCP), RFC-1171; initial options, both for the basic LCP and IPCP, RFC-1172; and support for Bridging, RFC-1220. In addition, a number of Internet-Drafts have been created, both for additional basic LCP options, and for support of further protocols. These include 32 bit checksums, Authentication, and a MIB, in the first class, and Appletalk, OSI CLNP, and DecNet IV in the second.

Implementation experience with the RFC-1172/1172 pair has shown a number of places where the documents are unclear or insufficiently precise; there were also minor difficulties with the state machine. As a result, an updated pair of documents was prepared, which also reformats the material along an LCP/IPCP split. While substantial experience has been gained with asynchronous implementations which match the new specification, synchronous experience is still lacking. It is hoped that this will be rectified shortly, so that the new versions can be advanced into the standards track.

It was also agreed that the various Internet Drafts (listed above), which had been languishing for lack of enthusiasm, be pushed forward. Two others which were written have been put on hold as no immediate need was seen for them, these being the SNAP and LLC documents. Finally, some technical discussion of the Authentication proposal was held.

IP over Appletalk WG

The MacIP document (IP over Appletalk) document is almost complete. The AURP protocol (Appletalk over IP) was discussed, and progress was reported. Operation of Appletalk over the PPP protocol was also covered, and some points to be forwarded to the PPP WG were brought out.

A proposal to run SNMP in the Appletalk suite was reviewed, and a document is in progress. Finally, two versions of the Appletalk MIB were debated; the question briefly addressed was whether or not to switch directly to the new one.

Variable Width Subnet Masks BoF

The subnets BoF reviewed a number of problematical cases brought up by the use of variable width subnet masks (i.e. use of more than one subnet mask in any given IP network). Consensus was reached on which configurations to allow and disallow.

The first question addressed was whether or not to allow two subnets in the same part of a network's address space to be topologically separate; the consensus was that this was necesary to get maximum use out of variable width masks.

The consensus was that the other three open questions should not be allowed. First, all masks must be contiguous and on the high end of the 'rest' field. Second, no 'subset' subnets would be allowed; no small subnet could have the same leading bits as a larger subnet. Finally, ambiguous subnets would not be allowed.

IP Logical Networks BoF

No information to report.

IP Address Enhancements BoF

This BoF discussed potential interim solutions to the near-term problem of the shortage of class B IP network numbers. Basically, at the current rate of usage these will run out in several years.

The two key questions were what kind of extra network numbers to create, and what portion of the existing IP address space to devote to them. After a commendably quick discussion, consensus answers did appear to both of these.

The answer to the first was to create a new class, intermediate between class B and C, with a 12 bit rest field, called class 'E' addresses. No definite answer was found for the second; the two possibilites for this field are to use either i) the entire current 'Reserved' space (i.e. 1111 prefix), or ii) the back half of the class C space (i.e. 1101 prefix). In addition, since the former would use the last reserved space, a new reserved space would be created, consisting of the back half of the existing class A space (i.e. 01 prefix).

Discussion in the hallways after the BoF concluded that the optimal choice was in fact the first one; it has also been suggested that the back half of the class C space be made reserved as well.

Mobile Hosts BoF

No information to report.

CURRENT MEETING REPORT

Reported by Andy Nicholson/Cray Research, Inc.

CBNR BOF Minutes

Description:

We held a BOF on this subject at the 20th IETF in St. Louis, and this is a continuation of that interchange. However, the only attendees from the St. Louis meeting present at this meeting were the Cray Research, Inc. representatives.

While working with circuit-switched T3 networks, developers at Cray Research, Inc., determined that there would be advantages to defining a standard way to control certain classes of network resources through the internet. In the case of a circuit-switched T3 line, the line should be switched on only when there are active transport connections which can fully utilize the service. Due to the high cost of the resource, under-utilization would be particularly undesirable. The developers believe that this capability might have other applications in the internet and that an effort should be made to define a standard protocol.

Minutes:

Due to the small size and informality of the meeting, no formal Minutes were taken. This record is believed to be reasonably accurate and proper credit given to the originators of the ideas and concepts presented. Andy Nicholson, who is preparing this report, apologizes for any errors or omissions.

A variety of new issues were brought up at this meeting, and it was encouraging to note that there were as many non-Cray people as Cray employees. Most of the discussion centered around the concrete example of the Circuit-Switched T3 services being prototyped by Cray Research, Inc.

The first issue raised centered on local routing to the T3 adapter. This would include routing to any controlled device. The prototypes assume that a particular network link will be used for transfer of data packets, thus static routing is implied. There is concern that this perspective may lead to the use of static routing between the requesting host and the controlled resource. There was general agreement that this should not happen.

Another issue concerned recursive conditioning of resources. A host in control of a link might need to pass a request to another host through the controlled link so that further downstream links may be conditioned. This should be possible.

Fred noted that some comparisons could be made with regard to this capability. For example, this is similar to the switching which takes place through the telco fabric as calls are routed. There is also a similarity to X.25. For example, TP0 will create a link over x.25 when a connection is established.

Matt did not think that this could be a widely deployed function; however, Fred noted that this might be useful in any kind of fundamentally switched service, i.e. ISDN or mobile hosts. This seems like something that is in the future.

In the Cray Research, Inc. prototype, most of the support code is in the kernel. Matt and Fred were concerned that perhaps this should all reside in user space. This leads to two fundamentally different approaches. For everything to be in user space, either special commands must be executed to condition the network before running applications, or the applications must make special library calls. If everything is in the kernel, then these services can be transparent to users and applications.

These discussions led us to a very different conclusion from the last meeting. All agreed that I would finish an informational RFC relating our experiences with the switched T3 services in time for review by the community before the next IETF. If possible, I will also document the protocol we are using.

At the 22nd IETF we will once again hold a BOF to gauge interest in these facilities. At that meeting we will determine whether to continue any work through the IETF.

Attendees

Fred Baker	fbaker@emerald.acc.com
David Borman	dab@cray.com
Matt Mathis	mathis@psc.edu
Andy Nicholson	droid@cray.com
John Seligson	johns@ultra.com
Jeff Young	jsy@cray.com

CURRENT MEETING REPORT

Logical Subnetwork BOF Minutes

Report not submitted. See Area Report for summary.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Tom Benkart	teb@saturn.acc.com
Randy Butler	rbutler@ncsa.uiuc.edu
Kurt Dobbins	dobbins@ctron.com
William Fink	bill@wizard.gsfc.nasa.gov
David Forster	forster@marvin.dec.com
Olafur Gudmundsson	ogud@cs.umd.edu
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Anders Klemets	klemets@cs.cmu.edu
Holly Knight	holly@apple.com
John Lekashman	lekash@nas.nasa.gov
Louis Leon	osll@emuvm1.cc.emory.edu
Shane MacPhillamy	slm@netrix.com
Andy Nicholson	droid@cray.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Lars Poulsen	lars@cmc.com
Harri Salminen	hks@funet.fi
Gershon Schatzberg	439-3582@mcimail.com
John Seligson	johns@ultra.com
Osamu Takada	takada@sdl.hitachi.co.jp
Glenn Trewitt	trewitt@nsl.dec.com
Richard Woundy	rwoundy@ibm.com

CURRENT MEETING REPORT

Reported by Noel Chiappa

ADRBOF Minutes

IP Address Hacks BOF

This BOF discussed potential interim solutions to the near-term problem of the shortage of class B IP network numbers. Of the three classes of IP network numbers, A, B and C, the class B numbers are being used up at the highest rate. There are 16K (2^14) of these, and over four thousand are already allocated (although not all are being advertised in the Internet). These are the most useful numbers, since most sites are not large enough to need a class A (24 bit rest), where most are too large to make good use of a class C (8 bit rest), particularly if subnetting is being used.

Depending on which exact model is used to predict future usage of these numbers, at the current rate of usage these will run out in several years. A straight exponential fits the curve so far quite well; it has been argued that this is not a useful model, but rather an S-curve model should be used. The problem is that no inflection point has yet appeared, so it is difficult to fit an S curve to the growth. In any case, there is general agreement that the problem is severe.

The two key questions were what kind of extra network numbers to create, and what portion of the existing IP address space to devote to them. After a commendably quick discussion, consensus answers did appear to both of these.

There were several potential answers to the first question. One option is to simply create more class B (i.e., 16 bit 'rest' field) numbers. The other was to create a new class of network numbers with a different size rest field, intermediate between class B and C. It was pointed out that most sites which are getting class B numbers do not need a whole class B space, but could easily use something a little smaller; this would reduce the usage of the class B space, thereby extending the lifetime of that space. Suggestions were made for a number of different sizes, including 14, 13 and 12 bits.

One thing going against more class B numbers was that to create a reasonable number of them would use a large chunk of the 32 bit IP address space. The current block of 16K used one quarter of the address space; all addresses with a '10' prefix. If another quarter were (somehow) devoted to class B, that would still only double the number. On the other hand, use of a smaller rest field would allow more network numbers to be packed into the portion of the address space allocated; since the available free (or reclaimable) spaces were mostly quite small, this weighed heavily in favor of the smaller fields.

A new class, with a 12 bit rest field (to called class 'E' addresses) was finally decided on, since it maximizes the number of network numbers that can be created. While a 12 bit rest field only allows for 4K hosts, this is still significantly more than a class C, and should be more than enough for most companies. Also, it is exactly equidistant between the class B and class C sizes. However, this decision (for 12 instead of 13 or more) needs to be reviewed carefully to make sure that a 12 bit rest field will actually be useful to a significant number of network number applicants.

This does point out the necessity of having hosts not pry into address formats. It is plausible to deploy a new network number format if only the routers have to be changed; doing so in a world where most types of host software have to be changed as well is clearly problematic.

There are two broad classes of solution to the question of where to allocate any new network numbers. The first is to use some or all of the currently reserved space; i.e. prefix 1111. The second is to recycle some of the currently "unlikely to be used" space; for instance the back half of the class A space (prefix 01), or the back half of class C (prefix 1101).

Considering the first, the question was whether or not to use the entire space, or to continue the practice of allocating a space whose prefix started with all 1's and ended with a 0; (i.e., allocate 11110 and reserve 11111). It was decided not to keep a part reserved if this space were used, but to use the entire space. The problem is that this practise is resulting in diminishing returns as far as the size of the portion of address space available to hold network numbers and the rest field; in other words, the overhead of the field dedicated to format identification was getting quite large (5 of 32 bits).

Use of the entire block would allow creation of 2^16 of these new network numbers. (4 bits of prefix + 16 of network number + 12 of rest = 32) This number, sixteen times larger than the number of class B's, could reasonably be expected to last quite a long time. Were this done, since it would use the last reserved space, a new reserved space should be created, consisting of the back half of the existing class A and/or C space.

Alternatively, if the back half of class C space (1101 prefix) were used to hold these new numbers, 2^16 of them could also be created here. It was pointed out that use of class C numbers would allow routers which did not understand class E addresses to treat them as a group (2^4, or 16) of class C numbers.

No definite answer was arrived at in the BOF for the question of where to place the new numbers, although there was general consensus that using all the reserved space or the back half of class C space were the two viable options. It was agreed that in either case the back half of the class A space should be reserved; it was felt that rather than move directly from one use to another, it would be best if a portion of the address space cycled through 'reserved', to allow use of the old meaning to disappear from the net before the new use was taken up.

Discussion in the hallways after the BOF concluded that the optimal choice was in fact to use the reserved space. It was felt that the ability to have older routers handle class E numbers as a group of class C numbers was not actually good, given the problems in the network with large numbers of network numbers. Also, it was felt that the argument above about cycling through reserve should apply to the back half of class C space as well.

Finally, and most important, it was pointed out that unless the new numbers were allocated from the reserved space, there would be less impetus on people's part to change their software. The ability to model a class E net as a group of C's would, from this viewpoint, be a problem, not an advantage. This is a weighty point, given the necessity of the change in the network; presumably people making the change to recognize E's would also put in the change to reserve the back half of A and C space, which might well be critical in the future.

CURRENT MEETING REPORT

Reported by Noel Chiappa

SUBNETBOF Minutes

Variable Width Subnet Masks BOF

The Subnets BoF reviewed a number of problematical cases brought up by the use of variable width subnet masks (i.e., use of more than one subnet mask in any given IP network). These cases all relate to the allocation of various subnetted addresses to various physical networks which are part of an IP network. Consensus was reached on which configurations to allow and disallow.

Before reviewing the specific points, it will be useful to include some terminology. Use of the subnet numbers "A, B.1 and B.2" means that A and B are differing values of a fixed part of the 'rest' field, and that 1 and 2 are differing values of a different, lower, fixed part of the 'rest' field.

For instance (using an 8 bit rest field), with the two masks 11100000 and 11111100, 'A' might be 001xxxxx, 'B' might be 010xxxxx, 'B.1' would be 010001xx and 'B.2' would be 010010xx. With this terminology in hand, the specific cases can now be reviewed in detail.

The first question addressed was whether or not to allow two subnets in the same part of a network's address space to be topologically separate. In other words, could subnets B.1 and B.2 be separated by subnet A? Looked at another way, if B.1 and B.2 are thought of as parts of a 'subnet' B, can that subnet be partitioned? If allowed, this would represent a divergence with the basic Internet philosophy, in which an IP network is not allowed to be partitioned. The argument for allowing this is to get maximum use out of variable width masks.

Variable width masks were added to the architecture to allow efficient use of address space. For example, if an enterprise, with a single IP network number, contains a single large LAN (with several thousand hosts), and a number of small LAN's (with tens of hosts), there is no single subnet mask that will efficiently use the address space of that network number. A wide mask, necessary to handle the single large LAN as a whole, will 'waste' space when used on the small LAN's. A small mask will force the single large LAN to be trated as a collection of small LAN's, with consequent forwarding overhead. An alternative approach would be to use a separate network number for the large LAN, but this will increase the number of network numbers in the system as a whole, with consequent global costs. If the enterprise is only singly connected to the rest of the Internet, there is no benefit to the rest of the system of having more than one network number for the enterprise. Thus, only with use of varying width masks can efficient use be made of address space, both in the network and the Internet as a whole.

The disadvantage to allowing this is that all the routers in a network must know where every subnet is (and what its mask is). For example, suppose B.1 and B.2 are on different sides of A (connected by routers R1 and R2 respectively), and a router R is attached only to subnet A and some outside network. In the current state of affairs, R will only know the subnet mask for A, on which it has an interface. Now, when a incoming packet for B.1 arrives at router R, knowledge of the mask for A (and thus B) is not sufficient; router R needs to be able to distinguish B.1 and B.2 as separate destinations if it is to forward the packet to the correct next hop router, R1 or R2. It is thus seen that, to function in the general case, all routers in a subnetted network now need to know the mask for every subnet in the system.

This is a substantial cost; however, it was felt that to make the restriction that all the small subnets in one piece of the network address space (i.e., B.1 .... B.N) must be contiguous worked against maximum utilization. Moreover, even this restriction does not necessarily remove the necessity for a router to know all the subnet masks in use in a given network. For example, if the router R above were connected to B.1, rather than A, it would still need the mask for A, unless it were for routing purposes to consider A as a large number of subnets of the same size as B.1.

Finally, the routing protocols which support variable length subnet masks do provide the necessary information to routers to do the forwarding correctly. The consensus thus was that allowing this configuration was necessary.

The next question to be addressed was whether all subnet masks must be contiguous and on the high end of the 'rest' field (i.e., have the form 11...1100...00). One argument that was put forward was that non-contiguous masks allowed more flexibility in extending the subnet mask when it ran out. It was pointed out that easy extension could be provided for by allocating subnet number bits from the high end of the rest field, and host bits from the low end, with unused space in the middle. Whenever either field became too small, it could be extended, as long as unused bits remained. Additionally, some versions of the Patricia tree algorithm do not work with non-contiguous masks.

While it was agreed that no good reason could be provided for not allowing other formats, no strong use could be seen for allowing them either, and in the interest of future flexibility the consensus was to not allow them.

The third question to be address was whether 'subset' subnets would be allowed; i.e., could a small subnet have the same leading bits as a larger subnet. For example, if one subnet is numbered B, could another subnet have the number B.1? Clearly, at a minimum, no hosts on subnet B could have a address which had B.1 as a prefix (i.e., addresses on subnets B.1 ... B.N which were in use could not appear on subnet B); this would leave routers unable to discover which subnet these hosts were on, unless each host was tracked separately.

It was initially thought that this was the only problem, which could be handled by correct configuration, and the feeling was that this should be allowed to optimize use of the address space. An additional cost would be that routers would have to do a 'best match' routing lookup. I.e., even after finding a mask and address that match, the router would still have to look for further potential matches that are more 'complete'. This cost exists now for routers that support host routes, however.

However, it was pointed out that a host H attached to subnet B would think that hosts attached to subnet B.1 (which host H would need to reach through a router) were in fact directly reachable by host H. No general fix (i.e., one that worked on all network technologies, not just those which used ARP) could be discovered for this problem. In addition, the chances for misconfiguration (e.g., a host on subnet B that appears to be on subnet B.1) are manifold. Given these points, the consensus was that this configuration should not be allowed.

Finally, ambiguous subnets were discussed briefly. This name refers to subnets masks (and numbers) which overlap in ways such that host addresses are not unambiguously on one network or another. For instance, consider two different subnets 5 and 6, with different subnet masks 5 and 6 (temporarily ignoring the fact that these are all 1's subnet numbers). Next, think of an address starting with 7; it matches the 5 address and mask, but also matches the 6 address and mask. Which one is better?

Since this case was ruled out by the fact that non-contiguous masks will not allowed, it was not discussed in detail. However, if that restriction is relaxed in the future, this question will need to be revisited.

Attendees

Steve Alexander	stevea@i88.isc.com
Philip Almquist	almquist@jessica.stanford.edu
Nagaraj Arunkumar	nak@3com.com
Karl Auerbach	karl@eng.sun.com
Tom Benkart	teb@saturn.acc.com
Arthur Berggreen	art@acc.com
David Borman	dab@cray.com
Scott Brim	swb@nr-tech.cit.cornell.edu
Rob Coltun	rcoltun@ni.umd.edu
Ralph Droms	droms@bucknell.edu
Robert Elz	kre@munnari.oz.au
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@canet.ca
Karen Frisa	karen.frisa@andrew.cmu.edu
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Phani Jujjavarapu	phani@cisco.com
Douglas Kerr	dougk@mtxinu.com
Nik Langrind	nik@shiva.com
John Lekashman	lekash@nas.nasa.gov
Tony Li	tli@cisco.com
Bill Manning	bmanning@rice.edu
Matt Mathis	mathis@psc.edu

Lars Poulsen	lars@cmc.com
Gershon Schatzberg	439-3582@mcimail.com
Osamu Takada	takada@sdl.hitachi.co.jp
Walter Wimer	walter.wimer@andrew.cmu.edu
Robert Woodburn	woody@cseic.saic.com
Richard Woundy	rwoundy@ibm.com

3.2.1 Connection IP (cip)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done	Produce a new specification of ST.
Done	Define common hardware and software platform.
Done	Implement hardware and software platform.
May 1991	Implement experimental modules and perform experiments.
May 1992	Produce a specification of a next generation connection oriented protocol.

Request For Comments:

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

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chair(s):

Ralph Droms, droms@bucknell.edu

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

Done	We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to: Exchange Internet packets with other hosts, Obtain packet routing information, Access the Domain Name System, and Access other local and remote services.
Done	We will summarize those mechanisms already in place for managing the information identified by Objective 1.
Jan 1991	We will suggest new mechanisms to manage the information identified by Objective 1.
Jan 1991	Having established what information and mechanisms are required for host operation, we will examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.
TBD	Write a bootp extensions document

Internet Drafts:

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

"Dynamic Host Configuration Protocol", 07/09/1991, R. Droms <draft-ietf-dhc-protocol-00.txt, .ps>

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell

DHC Minutes

Modifications and Extensions to DHCP

The Working Group agreed to the following changes and additions to DHCP. The DHCP Internet Draft will be edited to reflect these changes.

Changes to Protocol

The client–server protocol has been changed slightly, so that the client first broadcasts a message to locate available DHCP servers, and then selects one of the responding servers from which the client obtains its configuration parameters.

Protocol Messages

Corresponding to the changes to the protocol summarized in section 1.1, the DHCP messages have been redefined as shown in table 1.

Client–Server Interaction

1. Client broadcasts DHCPDISCOVER on local physical subnet. DHCP/BOOTP relay agents pass the broadcast on to DHCP servers not on the same physical subnet.

2. Servers respond with DHCPOFFER message with all configuration parameters including network address. Servers need not reserve the offered network address, although the protocol will work more efficiently if the server avoids allocating the offered network address to another client. The DHCPOFFER message is "unicast" to the client (using the BOOTP relay agent if necessary).

3. Client receives DHCPOFFER message from server. Client may choose to wait for multiple responses. Client chooses one server from which to request configuration parameters, based on offered configuration parameters in the DHCPOFFER messages. Client broadcasts a DHCPREQUEST message, specifying the desired server and desired network address in vendor extension fields. This DHCPREQUEST message is broadcast and relayed through BOOTP relay agents. Any DHCP/BOOTP relay agents must ensure that any messages from this client are forwarded to the same set of DHCP servers to ensure that the DHCPREQUEST message reaches the selected DHCP server.

   The client times out and retransmits the DHCPDISCOVER message if the client receives no DHCPOFFER messages.

Message	Use
DHCPDISCOVER	Client broadcast to locate available servers.
DHCPOFFER	Server to client in response to DHCPDISCOVER with offer of configuration parameters.
DHCPREQUEST	Client broadcast to servers requesting offered parameters from one server and implicitly declining offers from all others.
DHCPACK	Server to client with configuration parameters, including committed network address.
DHCPNAK	Server to client declining request for configuration parameters (e.g., requested network address already allocated).
DHCPDECLINE	Client to server indicating configuration parameters (e.g., network address) invalid.
DHCPRELEASE	Client to server relinquishing network address and cancelling remaining lease.

Table 1: DHCP Messages

4. Servers receive the DHCPREQUEST broadcast from the client. The servers not selected in the DHCPREQUEST message use the message as notification that the client has declined that server's offer. The server selected in the DHCPREQUEST commits the binding for the client to persistent storage and responds with a DHCPACK message containing the configuration parameters for the requesting client. The server inserts a unique lease identification cookie as a vendor extension.

   If the selected server is unable to satisfy the DHCPREQUEST (e.g., the requested network address has been allocated), the server responds with a DHCPNAK.

5. Client receives the DHCPACK with configuration parameters. The client performs a last minute check on the parameters (e.g., ARP for allocated network address), and notes the duration of the lease and the lease identification cookie specified in the DHCPACK message. At this point, the client is configured.

   If the client detects a problem with the parameters in the DHCPACK message, the client sends a DHCPDECLINE message to the server and restarts the configuration process.

Extension	Tag	Values	Length
DHCP message type	51	1=DHCPDISCOVER 2=DHCPOFFER 3=DHCPREQUEST 4=DHCPDECLINE 5=DHCPACK 6=DHCPNAK 7=DHCPRELEASE	[3]
Lease identifier cookie	52	address	[6]
Server identifier	53	address	[6]
Parameter request vector	54	256 bit (32 octet) vector	[34]
Parameter request list	55	n parameter tags	[n+2]

Table 2: New Vendor Extensions

If the client receives a DHCPNAK message, the client restarts the configuration process.

The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK or a DHCPNAK message.

6. A client may choose to relinquish its lease on a network address by sending a DHCPRELEASE to the server. The client identifies the lease to be released with the lease identification cookie.

New Vendor Extensions

The modifications to DHCP require some new vendor extensions, as listed in table 2.

Use of Vendor Extensions

A client can fill in vendor extensions in a DHCPDISCOVER and DHCPREQUEST to supply hints or request specific values from a server. For example, the client can fill in the 'IP address' vendor extensions to suggest a remembered network address The server fills in vendor extensions in DHCPDISCOVER and DHCPACK messages to supply specific configuration values to the client.

A client can also request specific configuration parameters without supplying hints through the "parameter request vector" and "parameter request list" vendor extensions. In the parameter request vector, a one bit in position n in the vector represents an explicit request for the vendor extensions parameter with tag n. The parameter request list is a list of vendor extension tags explicitly requested by the client.

Lease Durations and Clock Drift

The algorithm for lease duration interpretation given in subsection 6.1 of the DHCP Internet Draft is correct, assuming the client and server clocks are stable relative to each other. If there is drift between the two clocks, the server may still consider the lease expired before the client does. To compensate, the server may return a different lease duration to the client than the server commits to its local database of client information.

Lease Renewal Times

The client attempts to renew its lease from the allocating server beginning at time T1and from any available server at time T2. Times T1 and T2 are configurable by the server through vendor extensions. T1defaults to (0.5 * duration_of_lease). T2 defaults to (0.875 * duration_of_lease).

XID Field

The XID field must be interpreted by the server relative to individual clients, not as a globally unique value.

Retransmission

The client drives all retransmissions of the protocol. The protocol document still needs explicit descriptions of retransmission and exponential backoff algorithms.

"ciaddr" (Clarification)

The "ciaddr" field is to be filled in by the client only if the client has explicit knowledge of its network address. The client can supply a hint or a preferred network address through the IP address vendor extension.

If a server receives a DHCPDISCOVER or DHCPREQUEST message with an invalid "ciaddr", the server silently discards the message.

Use of DHCP in Hosts with Multiple Interfaces

A host with multiple network interfaces must use DHCP through each interface independently to obtain configuration information parameters for those separate interfaces.

DHCP and BOOTP Clients

Use of the vendor extensions defined in DHCP is not restricted to DHCP clients and servers. Existing BOOTP clients and servers may choose to use the newly defined vendor extensions. The one restriction is that BOOTP clients MAY NOT use the "DHCP client" vendor extensions. Only clients using DHCP may use the "DHCP client" vendor extension.

Implementations

Several members of the DHC Working Group indicated that they intend to work on independent implementations of DHCP. Completion of at least one of these implementations is expected before the Spring, 1992 IETF meeting.

Future Work

Greg Minshall agreed to develop a definition of the DHCP server-server protocol. Jesse Walker and Walt Wimer agreed to collaborate on the definition of a MIB for DHCP servers.

Attendees

Steve Alexander	stevea@i88.isc.com
Richard Basch	basch@mit.edu
David Bridgham	dab@asylum.sf.ca.us
Gregory Bruell	gob@shiva.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Steve Deering	deering@xerox.com
Ralph Droms	droms@bucknell.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Robert Gilligan	gilligan@eng.sun.com
Ron Jacoby	rj@sgi.com
Douglas Kerr	dougk@mtxinu.com
Michael Khalandovsky	mlk@ftp.com
Holly Knight	holly@apple.com
Nik Langrind	nik@shiva.com
Joshua Littlefield	josh@cayman.com
Greg Minshall	minshall@wc.novell.com
William Nowicki	nowicki@legato.com
Michael O'Dell	mo@bellcore.com
Bradford Parker	brad@cayman.com
Mark Saake	saake@llnl.gov
Jesse Walker	walker@eider.enet@decpa.dec.com
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	walter.wimer@andrew.cmu.edu
John Wobus	jmwobus@suvm.acs.syr.edu

3.2.3 IP over Appletalk (appleip)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

Done	Describe, in an RFC, the current set of protocols used to connect Macintoshes to IP internets.
Done	Define a MIB for the management of DDP/IP gateways.

Internet Drafts:

"The Transmission of IP Datagrams Over AppleTalk Networks", 03/08/1991, John Veizades <draft-ietf-appleip-ipoverappletalk-00.txt>

Request For Comments:

RFC 1243 "AppleTalk Management Information Base"

CURRENT MEETING REPORT

Reported by John Veizades/Apple

APPLEIP Minutes

AppleTalk over IP Tunneling

The Apple-IP Working Group met at the Atlanta IETF. These are the Minutes of that meeting.

The Agenda was as follows:

AURP

- Update
- Work in progress
- What is left to do

SNMP

- SNMP over DDP

AppleTalk MIB

- Old mib
- New mib
- Other mibs

PPP and AppleTalk

- Specification and description
- Issues
- Security

MacIP

- Protocol change
- Doc clean up
- Next step

Other Issues

- AA protocol
- Configuration management

The work with PPP over AppleTalk is being handled by Brad Parker. He has documented the NCP specifics for AppleTalk and is talking to the PPP Working Group about adding LCP support for dial back (one of the issues that the group felt must be added to the LCP). Brad is also working with Apple to get support of the protocol from Apple.

The SNMP and MIB work is progressing. The old MIB is now an RFC and has been implemented by several vendors. The new MIB is available from lancaster.andrew.cmu.edu the file name is appletlak-mib2.txt. The specification for SNMP over DDP is available from apple.com in the directory /pub/apple-ip. Mike Ritter of Apple and Greg Minshall from Novell are working on this. This specification should be finished by the next meeting.

The AURP work was presented by Alan Oppenheimer. The folks from Shiva and Apple have been working on prototypes of this specification and have been interoperating over the Internet. The Shiva folks brought up several issues as to the transport layer AURP provides. Alan will be changing the document to better differentiate the transport layer from the rest of the AURP protocol.

One change was made to the MacIP document and was presented to the group. Several issues as to the final format of the document were made and comments were presented to the author. A revised version of the document will be published shortly.

The meeting finished with the continuing discussion of the AA protocol. Phil Bunde still has the action to produce the actual document.

As Working Group Chair of this group I still have the concern that much of the work that must be done for AppleTalk to continue to grow as a protocol family is still not being accomplished by either Apple or the AppleTalk developer community at large. This work must be done and some forum should be created for this progress to occur.

Attendees

Gregory Bruell	gob@shiva.com
Philip Budne	phil@shiva.com
Richard Cherry	rcherry@novell.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Robert Elz	kre@munnari.oz.au
Karen Frisa	karen.frisa@andrew.cmu.edu
Douglas Kerr	dougk@mtxinu.com
Nik Langrind	nik@shiva.com
Brian Lloyd	brian@telebit.com
John Lynn	lynn@ttcllcat.cit.cornell.edu
Phil Mathis	phil@gatech.edu
Leo McLaughlin	ljm@wco.ftp.com
Greg Minshall	minshall@wc.novell.com
Chandy Nilakantan	csn@3com.com
Alan Oppenheimer	oppenheimer1@applelink.apple.com

Michael Ritter	mwritter@applelink.apple.com
John Scudder	jgs@merit.edu
Richard Telljohann	telljohann1@applelink.apple.com
John Veizades	veizades@apple.com
A. Lee Wade	wade@discovery.arc.nasa.gov
Jonathan Wenocur	jhw@shiva.com

3.2.4 IP over FDDI (fddi)

Charter

Chair(s):

Dave Katz, dmk@merit.edu

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

Request For Comments:

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

3.2.5 Multi-Media Bridging (mmb)

Charter

Chair(s):

Jeffrey Fitzgerald, `jjf@fibercom.com`

Mailing Lists:

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

Description of Working Group:

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

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

Goals and Milestones:

Done	Finalize Charter of Group
Aug 1991	Document mulit-media bridging technology and its affect on TCP/IP Internets.
Aug 1991	Document issues to be addressed by Working Group.

3.2.6 Point-to-Point Protocol Extensions (pppext)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

None specified

Internet Drafts:

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

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

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

"Proposed Point-to-Point Procotol for AppleTalk", 07/08/1991, S. Senum, J. Muchow, F. Slaughter, B. Parker <draft-ietf-pppext-appletalk-00.txt>

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

"The PPP Authentication Protocols", 07/25/1991, B. Lloyd, W.A. Simpson <draft-ietf-pppext-authentication-00.txt>

Request For Comments:

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

CURRENT MEETING REPORT

Reported Brian Lloyd/Telebit

PPPEXT Minutes

Noel Chiappa opened the Point-to-Point Protocol Extensions meeting and then handed it over to Brian Lloyd, the new Working Group Chair.

There was an early observation that no router vendors (other than Telebit) appeared to be present. This curtailed discussion of synchronous PPP. Most of the discussion that followed tended to address the needs/desires of the asynchronous PPP community.

The current list of documents was discussed. Here is the list of current documents:

RFC 1171 RFC 1172

- LCP draft Simpson
- IPCP draft McGregor
- Appletalk Parken - hold waiting for Appletalk wg
- ISO/CLNP Katz/Simpson - no interest yet?
- DECnet 4 Senum - not much said
- SNAP B? - no interest
- LLC Harvey - no interest
- Bridging Baker - already an RFC
- 32 bit FCS Harvey - general approval - no known implementations
- Authentication Lloyd/Simpson - most discussion here
- MIB Kastenholz - no implementations

Consensus indicated that the lcp and ipcp draft documents (these will supercede RFC 1171 and 1172 respectively) should proceed to the next level since there are numerous interoperable implementations. Bill Simpson will make very minor changes to his document and republish it.

None of the other documents were deemed ready to progress to the next level because of either a) lack of further interest, or b) no implementations.

Consensus also indicated a need for a catalog document to keep track of all the PPP related documents. Brian Lloyd is working on that.

Discussion was heavy on the new Lloyd/Simpson authentication protocol document that describes the Password Authentication Protocol (PAP) and the Challenge Handshake Authentication Protocol (CHAP). There was a discussion of where authentication should go. It was decided that link-level authentication in the Link Control Protocol (LCP) is acceptable so long as additional authentication may be used within the Upper Layer Protocols (ULPs – NCP or higher).

Strong discussion indicated a need for a mechanism to allow a called PPP system to indicate to its peer that it wishes to close the link and dial-back for purposes of authentication. This information was referred back to Simpson and Lloyd to research and add to the authentication document.

It was decided that the proper digest algorithm for CHAP should be MD5. As a result it was decided that references to MD2 and MD4 should be removed from the document.

James Galvin representing the Security Area Advisory Group (SAAG) Working Group strongly recommended adding a section on distribution of the "secret" used in CHAP.

More information is needed in the authentication document about bit and octet ordering and character sets used (in the case of legible passwords and secrets).

More detail is needed about the PAP message reply.

CHAP needs a mechanism (besides dropping the link) to indicate that the authentication has succeeded or failed. This is because some system will require the user to enter the secret value in real time so there may be errors and hence retries.

The size of the secret value was increased from 64 to 128 bits.

The challenge needs to be non-repeating. The document needs to discuss methods of generating good challenges.

The document should also remove all references to encryption.

The last item of the day was to generate a list of recommended PPP options to go into the router requirements document. The final list of suggested options for sync implementations:

- Support for the Link Quality Monitoring (LQM) option.
- Support for the magic number (loopback detection) option.
- No address/control field compression.
- No protocol field compression.

For Async Implementations:

- Do address/control field compression.
- Do protocol field compression.

Attendees

James Barnes	barnes@xylogics.com
Gregory Bruell	gob@shiva.com
Philip Budne	phil@shiva.com
Randy Butler	rbutler@ncsa.uiuc.edu
Cerafin Castillo	cec@emulex.com
Richard Cherry	rcherry@novell.com
Gary Ellis	garye@hpspd.spd.hp.com
Arlan Finestead	arlanf@ncsa.uiuc.edu
James Galvin	galvin@tis.com
Robert Griffioen	
Jack Hahn	hahn@umd5.umd.edu
Russ Hobby	rdhobby@ucdavis.edu
Ron Jacoby	rj@sgi.com
Frank Kastenholz	kasten@europa.clearpoint.com
Tom Kessler	kessler@sun.com
Kenneth Key	key@cs.utk.edu
Mark Knopper	mak@merit.edu
Anthony Lauck	lauck@tl.enet.dec.com
Joshua Littlefield	josh@cayman.com
Brian Lloyd	brian@telebit.com
Glenn McGregor	ghm@merit.edu
Bradford Parker	brad@cayman.com
Mike Petry	petry@ni.umd.edu
Richard Smith	smiddy@pluto.dss.com
Bruce Taber	taber@interlan.com
Theodore Tso	
John Veizades	veizades@apple.com
Preston Wilson	preston@i88.isc.com
L. Michele Wright	uncng!michele@uunet.uu.net
Joseph Zur	fibrontics!zur@uunet.uu.net

3.2.7 Router Discovery (rdisc)

Charter

Chair(s):
Steve Deering, deering@xerox.com

Mailing Lists:
General Discussion: gw-discovery@gregorio.stanford.edu
To Subscribe: gw-discovery-request@gregorio.stanford.edu
Archive:

Description of Working Group:

The Router Discovery Working Group is chartered to adopt or develop a protocol that Internet hosts may use to dynamically discover the addresses of operational neighboring gateways. The group is expected to propose its chosen protocol as a standard for gateway discovery in the Internet.

The work of this group is distinguished from that of the Host Configuration Working Group in that this group is concerned with the dynamic tracking of router availability by hosts rather than the initialization of various pieces of host state (which might include router addresses) at host-startup time.

Goals and Milestones:

Done	Created Working Group; established and advertised mailing list. Initiated email discussion to identify existing and proposed protocols, for router discovery.
Done	Held first meeting in Palo Alto. Reviewed 9 candidate protocols, and agreed on a hybrid of cisco's GDP and an ICMP extension proposed by Deering.
Done	Held second meeting in Tallahassee. Reviewed the proposed protocol and discussed a number of open issues.
Done	Held third meeting in Pittsburgh. Discussed and resolved several issues that had been raised by email since the last meeting. Draft specification of router discovery protocol to be ready by next meeting. Experimental implementations to be started.
Done	Meet in Vancouver. Review draft specification, and determine any needed revisions. Evaluate results of experimental implementations and assign responsibility for additional experiments, as required. Submit the specification for publication as a Proposed Standard shortly after the meeting.
Done	Revise specification as necessary, based on field experience. Ask the IESG to elevate the protocol to Draft Standard status. Disband.

Ongoing Gather implementation and operational experience, revise the specification to reflect lessons learned, and submit the protocol for Draft Standard.

Request For Comments:

RFC 1256 "ICMP Router Discovery Messages",

3.2.8 Router Requirements (rreq)

Charter

Chair(s):

Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:

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

Description of Working Group:

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

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

The purposes of this project include:

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

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

Goals and Milestones:

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

Internet Drafts:

"Requirements for Internet IP Routers", 09/17/1990, Philip Almquist <draft-ietf-rreq-iprouters-02.txt>

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

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

"Type of Service in the Internet Protocol", 07/25/1991, Philip Almquist <draft-almquist-tos-00.txt>

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

"IP Forwarding Table MIB", 08/14/1991, Fred Baker <draft-ietf-rreq-forwarding-00.txt>

CURRENT MEETING REPORT

Reported by Philip Almquist/Consultant

RREQ Minutes

The Router Requirements Working Group held extensive meetings in Atlanta in an attempt to resolve the remaining outstanding issues. We were mostly very successful: although a considerable number of loose ends remain, we expect that they can be handled on the mailing list (and, if necessary, in a videoconference). Our goal is to have a final version of the document available in October to be formally submitted to the standards process at the November IETF meeting.

The meetings focused on three activities:

1. Review of the Router Requirements draft (Monday, Wednesday, and Thursday).

2. Discussion of the theoretical basis of routing in border routers (Tuesday and Wednesday).

3. Review of a draft document describing IP type of service (Tuesday).

Some visitors from the IPLPDN Working Group asked us to consider the changes to the IP routing architecture that they are considering, but we did not have time to hold the extensive discussion which would have been necessary to reach consensus on that issue.

Each is described in more detail below. The Chair would like to thank Frank Solensky for recording the decisions reached during the meeting.

Review of Router Requirements Draft

CURRENT MEETING REPORT

Reported by Philip Almquist/Consultant

RREQ Minutes

The Router Requirements Working Group held extensive meetings in Atlanta in an attempt to resolve the remaining outstanding issues. We were mostly very successful: although a considerable number of loose ends remain, we expect that they can be handled on the mailing list (and, if necessary, in a videoconference). Our goal is to have a final version of the document available in October to be formally submitted to the standards process at the November IETF meeting.

The meetings focused on three activities:

1. Review of the Router Requirements draft (Monday, Wednesday, and Thursday).

2. Discussion of the theoretical basis of routing in border routers (Tuesday and Wednesday).

3. Review of a draft document describing IP type of service (Tuesday).

Some visitors from the IPLPDN Working Group asked us to consider the changes to the IP routing architecture that they are considering, but we did not have time to hold the extensive discussion which would have been necessary to reach consensus on that issue.

Each is described in more detail below. The Chair would like to thank Frank Solensky for recording the decisions reached during the meeting.

Review of Router Requirements Draft

The entire document was reviewed in detail. This process was considerably less contentious than at many previous meetings, since the most divisive issues had previously been resolved. Some of the issues included:

- The relationship between our document and the Host Requirements, and the extent to which our document ought to replicate material found in the Host Requirements.
- When (if ever) a router should believe ICMP Redirects.
- Metrics for static routes.
- Whether SNMP may be implemented via a proxy agent.
- The security of in-band configuration mechanisms.
- The allowability of IP multicast addresses in source routes.
- What still needs to be done to complete the document.

In regard to the first issue, we concluded that overlap is generally not desirable, and that we should work to eliminate it. The resolution of the other issues should be obvious from the next version of the draft.

Theory of Border Routers

The Working Group continued its discussion from previous meetings of how a router which is in multiple routing domains can choose from among routes to the same destination learned from different routing protocols ("route believability") and how can it pass routing information between multiple routing domains ("route leaking"). Our discussion of these difficult topics was guided by several papers by Working Group members (all available as Internet Drafts): "Ruminations on the Next Hop" and "Ruminations on Route Leaking", both by Philip Almquist, and "Some Thoughts on Multi-Domain Routing" by Ross Callon.

The Working Group concluded that the discussion in the papers was generally correct. However, two issues were raised concerning the first paper:

1. The meta-lookup algorithm always picks the most specific route to any destination; administrative policy controls are used only to choose among routes that are equally specific. Some felt that some network managers may wish to have policy influence route choice among routes that are not equally specific (in particular, some may wish to emulate the result of the "Rank Ordering of Routing Protocols" approach).

2. The preference of a route is not influenced by route leaking. Some felt that this is incorrect: the preference of a route ought to improve if the route is leaked into a routing protocol whose default preference value is better than the original preference value of the route.

Neither of these issues were completely resolved at the meeting, so the author was tasked to consider them in the next version of the paper.

IP Type of Service

Philip Almquist is attempting to write an RFC on the use of the TOS bits in the IP header. The group briefly reviewed the then-current draft ("Type of Service in the Internet Protocol", available as an Internet Draft). Although the group raised some editorial concerns, consensus on the technical content was reached with almost no debate.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Arthur Berggreen	art@acc.com
David Bridgham	dab@asylum.sf.ca.us
Gregory Bruell	gob@shiva.com
Ross Callon	callon@bigfut.enet.dec.com
Vinton Cerf	vcerf@nri.reston.va.us
Peter Chang	tpc@mtunm.att.com

Rob Coltun	rcoltun@ni.umd.edu
Kurt Dobbins	dobbins@ctron.com
Robert Elz	kre@munnari.oz.au
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@canet.ca
William Fink	bill@wizard.gsfc.nasa.gov
David Forster	forster@marvin.dec.com
Vince Fuller	vaf@stanford.edu
Deborah Futcher	dfutche@eco.twg.com
Chris Gunner	gunner@osicwg.enet.dec.com
Steven Hotz	hotz@isi.edu
B.V. Jagadeesh	bvj@3com.com
Frank Kastenholz	kasten@europa.clearpoint.com
Michael Khalandovsky	mlk@ftp.com
Paulina Knibbe	knibbe@cisco.com
Stev Knowles	stev@ftp.com
Richard Larkin	rblarkin@sprintf.merit.edu
John Lekashman	lekash@nas.nasa.gov
Mark Lewis	mlewis@telebit.com
Tony Li	tli@cisco.com
John Lynn	lynn@ttcllcat.cit.cornell.edu
Shane MacPhillamy	slm@netrix.com
Gary Malkin	gmalkin@ftp.com
Bill Manning	bmanning@rice.edu
Bernie May	bfm@houxa.att.com
April Merrill	abmerri@tycho.ncsc.mil
John Moy	jmoy@proteon.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
David O'Leary	oleary@sura.net
Radia Perlman	perlman@radia.enet.dec.com
Jason Perreault	perreaul@interlan.interlan.com
Lars Poulsen	lars@cmc.com
James Rees	rees@ifs.umich.edu
Michael Reilly	reilly@nsl.dec.com
Gershon Schatzberg	439-3582@mcimail.com
John Scudder	jgs@merit.edu
John Seligson	johns@ultra.com
Richard Smith	smiddy@pluto.dss.com
Frank Solensky	solensky@clearpoint.com
Emil Sturniolo	emil@dss.com
Osamu Takada	takada@sdl.hitachi.co.jp
Ross Veach	rrv@uiuc.edu
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov

John Wobus	jmwobus@suvm.acs.syr.edu
Robert Woodburn	woody@cseic.saic.com
Joseph Zur	fibrontics!zur@uunet.uu.net
Peter de Vries	peter@wco.ftp.com

3.2.9 Special Host Requirements (shr)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

General Discussion: ietf-hosts@nnsc.nsf.net
To Subscribe: ietf-hosts-request@nnsc.nsf.net
Archive:

Description of Working Group:

The Special-purpose Host Requirements Working Group is chartered to clarify application of the Host Requirements RFCs (1122 and 1123) to systems that are technically hosts but are not intended to support general network applications. These special-purpose hosts include, for example, terminal servers (a "Telnet host"), or file servers (an "FTP host" or an "NFS host").

The Host Requirements RFCs address the typical, general-purpose system with a variety of applications and an open development environment, and give only passing consideration to special-purpose hosts. As a result, suppliers of special-purpose hosts must bend the truth or make excuses when users evaluate their products against the Requirements RFCs. Users must then decide whether such a product is in fact deficient or the requirements truly do not apply. This process creates work and confusion, and undermines the value of the RFCs. The commercial success of the Internet protocols and their use in increasingly unsophisticated environments exacerbates the problem.

The Working Group must define principles and examples for proper functional subsets of the general-purpose host and specifically state how such subsets affect the requirements. The Working Group must determine the balance between an exhaustive list of specific special-purpose hosts and philosphy that remains subject to debate. For the most part, it should be possible to base decisions on existing experience and implementations. The special-purpose requirements will be stated as differences from the existing RFCs, not replacements, and will refer rather than stand alone.

Since they define strict subsets of the Host Requirements RFCs, the Special-purpose Host Requirements appear to be an easier job and can be developed and stabilized within 8-12 months. Most of the group's business can be conducted over the Internet through email.

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.

Done	First IETF Meeting: discussion and final approval of Charter; discussion and agreement on approach, including models, format, level and type of detail. Make writing assignments.
Oct 1990	First draft document.
Nov 1990	Second IETF Meeting: review first draft document, determine necessary revisions. Follow up discussion on mailing list.
Jan 1990	Revised document.
Feb 1990	Third IETF Meeting: make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail.
Apr 1991	Final draft document.
May 1991	Fourth IETF meeting: review final draft and if OK, give to IESG for publication as RFC.

3.3 Network Management Area

Director(s):

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

Area Summary reported by James Davin/MIT

At the Atlanta meeting of the IETF, seven Working Groups of the Network Management Area held one or more sessions throughout the week.

The SNMP Network Management Directorate reviewed four weighty MIB specifications that had been recently reported out of Working Groups: the DECNet Phase IV MIB, the FDDI MIB, the Bridge MIB, and the RMON MIB. The IESG announced its intention to consider these MIBs as candidates for Proposed Standard status after final text is available. In addition, the IESG announced its intention to consider the elevation of the Concise MIB specification (RFC 1212) to Draft Standard status and the elevation of MIB II (RFC 1213) to Standard status.

In addition, the Directorate conducted a special session for which the goal was identification of functional problems with the SNMP network management framework as represented in its core specifications. The meeting focused on identification of problems and did not attempt to formulate solutions. Many opinions were offered, and an account of the concerns expressed will be published.

Remote LAN Monitoring Working Group

This Working Group reached closure on its MIB for devices that provide for remotely monitoring ethernet LANs. Among other changes, the group decided to rename its effort to reflect its character more accurately. Thus, what was formerly called the "RLAN MIB" will be hereinafter called the "RMON MIB." A new draft of the RMON MIB, incorporating the final consensus of the Working Group and correcting deficiencies identified in Directorate review, will be issued soon. The IESG announced its intention to consider the RMON MIB document for Proposed Standard status.

With the production of this MIB, the work of the Working Group will be concluded. At the Atlanta meeting, members of the Working Group expressed interest in future projects to develop analogous MIB support for remote monitoring of media other than Ethernet. The Area Director and Mike Erlinger are discussing the possibility of chartering a new Working Group to focus on remote monitoring of 802.5 networks.

IEEE 802.3 Hub MIB Working Group

This Working Group met for the first time in Atlanta to consider the initial draft of an 802.3 Hub MIB for use with SNMP. This effort is based on work originally produced in the IEEE. A number of issues were resolved, but some contentious ones remain, and the

Working Group will meet again at the next IETF.

Internet Accounting Working Group

This Working Group met throughout the week at the Atlanta meeting. Of the three documents being considered in this group, the background document is completed. Discussion on the architecture document produced a list of additions and changes that will be made by the document editor. Work began on a document that defines a SNMP MIB to support accounting for IP packets.

X.25 MIB Working Group

This Working Group met for the first time in Atlanta to consider the initial draft of three MIB documents. One instruments LAPB functions; another instruments X.25 packet layer functions; the third instruments functions related to the encapsulation of IP over X.25. A number of issues were resolved, but a number remain. This Working Group will meet again at the next IETF.

FDDI MIB Working Group

This Working Group reached closure on its MIB for managing FDDI networks. The final draft of the FDDI MIB, representing the consensus of the Working Group and correcting deficiencies identified in Directorate review, will be issued soon. The Working Group also concluded that none of the previously contemplated trap definitions warranted standardization and agreed not to consider these further. The IESG announced its intention to consider the FDDI MIB document for Proposed Standard status.

SNMP Working Group

This Working Group discussed the most recent draft of the SMDS Interface Protocol (SIP) MIB and recommended that, with certain changes, it be considered for Proposed Standard status. An appropriately revised draft will be submitted for Directorate review.

This Working Group also discussed concerns recently raised by the IAB with respect to certain objects in the Ethernet MIB. The Working Group decided that the proposed course of simply deleting the identified objects was not acceptable, but did not, in the time available, come to any conclusion about how to treat those objects so that the Ethernet MIB could be progressed.

A special session of the Working Group was hastily convened later in the week in hopes of resolving the Ethernet MIB question. Although that session did formulate a tentative disposition for the relevant MIB objects, it is not clear that its conclusions enjoy broad community consensus. Accordingly, an additional, final meeting on the Ethernet MIB issue is planned for the next IETF. With this meeting, the business of the SNMP Working Group will be concluded.

Management Services Interface Working Group

This Working Group met at Atlanta to discuss its future. The Working Group elected to disband, because the daunting technical problems facing the Working Group were unlikely to be solved anytime soon, and because any broad community support for this project was unclear.

3.3.1 Bridge MIB (bridge)

Charter

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

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

Description of Working Group:

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

Goals and Milestones:

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

Internet Drafts:

"Definitions of Managed Objects for Bridges", 05/24/1991, E. Decker, P. Langille,, A. Rijsinghani, K. McCloghrie <draft-ietf-bridge-definitions-02.txt>

3.3.2 Character MIB (charmib)

__Charter__

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

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

Description of Working Group:

The Character MIB Working Group is chartered to define an MIB for character stream ports that attach to such devices as terminals and printers.

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

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

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

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

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discuss 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.

Internet Drafts:

"Definitions of Managed Objects for RS-232-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-rs232like-02.txt>

"Definitions of Managed Objects for Parallel-printer-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-parallelprinter-01.txt>

"Definitions of Managed Objects for Character Stream Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-charmib-01.txt>

3.3.3 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):
Jonathan Saperia, saperia@tcpjon.enet.dec.com

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

Description of Working Group:

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

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

Goals and Milestones:

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

Internet Drafts:

"DECnet Phase IV MIB Extensions", 06/06/1991, Jon Saperia <draft-ietf-decnetiv-mibextensions-01.txt, .ps>

3.3.4 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

Description of Working Group:

The primary goal of the FDDI MIB Working Group is to define a MIB for FDDI devices with objects which are based on those defined in the ANSI FDDI specifications and are compliant with the Internet standard SMI, MIB, and SNMP.

Goals and Milestones:

Sep 1990	"Final" initial draft of required get/set variables.
Oct 1990	Initial implementations of required get/set variables.
Feb 1991	Revised "final" draft of required get/set variables.
Mar 1991	Adoption of draft of required get/set variables.
Mar 1991	Initial draft of traps (events) and actions.

CURRENT MEETING REPORT

Reported by Jeff Case/UTenn

FDDIMIB Minutes

The FDDI MIB Working Group met on Tuesday afternoon, July 30, 1991 in Atlanta, Georgia. The meeting was held in conjunction with the Twenty First meeting of the Internet Engineering Task Force.

The Chair reported on the status of the MIB document. The document has been submitted for review by the SNMP Directorate. The next step is consideration by the IESG, followed by the IAB.

There was a brief discussion of the experiences to date implementing the current document. Several vendors have implemented or are implementing the current draft and will be demonstrating interoperability at a trade show the first full week in October.

Most of the meeting was spent in detailed discussion of Version 0.1 of the Trap document. After considerable discussion, the net result was that there is no need for a trap document. It should be noted that this resolution would not have been possible in such a timely way without the statesman-like efforts of Ron Jacoby and others who yielded to the majority view in the interest of consensus.

The work of the committee is finished until the documents are ready for promotion, the IAB has questions or comments, or work in ANSI renders our document obsolete. The Chair's understanding of the Area Director's wishes is that the Working Group will continue to exist to monitor these activities and related implementation experience. However, the Working Group may not meet as actively as it has in the recent past.

The MIB document may be obtained via anonymous ftp from any of the Internet Drafts directories.

Attendees

Steve Bostock	steveb@novell.com
Jeffrey Case	case@cs.utk.edu
Cho Chang	chang_c@apollo.hp.com
Jeff Fried	jmf@relay.proteon.com
Phillip Hasse	phasse@honchuca-emh8.army.mil
Mark Hoerth	mark_hoerth@hp0400.desk.hp.com
Ron Jacoby	rj@sgi.com
Mike Janson	mjanson@mot.com
Kenneth Key	key@cs.utk.edu
Tim Lee-Thorp	ngc!tim@uunet.uu.net
Ron Mackey	rem@dsiinc.com

Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Norman Patty	
David Perkins	dperkins@synoptics.com
John Pickens	jrp@3com.com
James Reeves	jreeves@synoptics.com
Anil Rijsinghani	anil@levers.enet.dec.com
Marshall Rose	mrose@dbc.mtview.ca.us
Mark Schaefer	schaefer@davidsys.com
Anil Singhal	
William Versteeg	bvs@nrc.com
David Waitzman	djw@bbn.com
David Ward	dward@chipcom.com
Mark Wood	markl@dsiinc.com
Joseph Zur	fibrontics!zur@uunet.uu.net

3.3.5 IEEE 802.3 Hub MIB (hubmib)

Charter

Chair(s):

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

Mailing Lists:

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

Description of Working Group:

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

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

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

Although the definitions produced by the Working Group should be architecturally consistent with MIB-II and related MIBs wherever possible, the Charter of the Working Group does not extend to perturbing the conceptual models implicit in MIB-II or related MIBs in order to accommodate 802.3 Hubs. In particular, to the extent that the notion of a "port" in an 802.3 Hub is not consistent with the notion of a network "interface" as articulated in MIB-II, it shall modelled independently by objects defined in the Working Group.

Because the structure of 802.3 Hub implementations varies widely, the Working Group shall take special care that its definitions reflect a generic and consistent architectural model of Hub management rather than the structure of particular Hub implementations.

The IEEE Hub Mgmt draft allows an implementor to separate the ports in a hub into groups, if desired. (For example, a vendor might choose to represent field-replaceable units as groups of ports so that the port numbering would match a modular hardware implementation.) Because the Working Group Charter does not extend to consideration of fault- tolerant, highly-available systems in general, its treatment of these groups of ports in an 802.3 Hub (if any) shall be specific to Hub management and without impact upon other portions of the MIB.

Goals and Milestones:

Done	Distribute first draft of documents and discuss via E-mail.
Done	Working group meeting as part of IETF to review documents.
Sep 1991	Distribute updated documents for more E-mail discussion.
Nov 1991	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Jan 1992	Documents available with specified changes incorporated.

Internet Drafts:

"Definitions of Managed Objects for IEEE 802.3 Repeater Devices", 07/23/1991, Donna McMaster, Keith McCloghrie <draft-ietf-hubmib-mib-00.txt>

CURRENT MEETING REPORT

Reported by Keith McCloghrie/Hughes LAN Systems

IEEE 802.3 Hub MIB Minutes

Since this was the first meeting of the Hub MIB Working Group, the first item of business was discussion of the Working Group's charter. The discussion emphasised that the focus would be on development of a Repeater MIB. The possibility of writing a MIB to represent a modular chassis containing multiple repeaters/bridges/terminal-servers/etc. was seen as a potential future work item, depending on the wishes of the NM Area Director and the Working-Group members, but any such effort would take second place to the development of the Repeater MIB.

Donna McMaster then gave an overview of how an initial draft of a Repeater MIB had been produced as a literal translation from the work of the IEEE Hub Management Task Force's MIB. This draft and a set of proposed changes had been distributed on the mailing-list. Some of these had been accepted, and some had been the subject of questions and/or differing views. Thus, a new draft had been produced having one section (section 7) documenting its differences from the IEEE MIB, and another (section 8) documenting the outstanding issues. This draft had also been distributed to the mailing-list, and submitted as an Internet Draft.

After a review of the changes made so far, the Working-Group approved the consensus reached on the mailing-list in agreeing to these changes. A suggestion was also made to change the names of SelfTest1 and SelfTest2 to be more descriptive of their function, i.e., non-disruptive and disruptive, respectively.

The next item on the agenda was discussion of the outstanding issues. The discussion began with the most fundamental issue: whether or not the MIB should explicitly allow for an agent managing multiple repeaters. Various arguments were made for each side of the argument. Much of the discussion was intertwined with discussion of how to manage a chassis containing multiple cards (repeaters, bridges, etc.) and having multiple (separate) LAN segments on its backplane.

Arguments for defining a MIB to manage just one repeater:

- The simple agent is simpler,
- Multiple repeaters can be managed through multiple MIB views, through proxy, or through multiple agents.
- A single repeater doesn't necessarily (usually?) correspond to an individual card.
- For consistency with the Bridge-MIB, which manages only one Bridge.

Arguments for defining a MIB with the explicit ability to manage multiple repeaters:

- An agent which manages multiple repeaters is simpler.
- A single repeater doesn't necessarily correspond to an individual card.
- Consistent with MIB-II's interfaces group.

After much discussion, a straw poll has called and resulted in a 15-15 tie. Regretfully, the meeting decided that the Working-Group would have to try to resolve the issue on the mailing-list, and if that failed then to schedule a meeting in the September/October timeframe, possibly at Interop.

The issue of optional tables was also discussed, and the consensus was that the distinction between the basic and monitoring groups should be maintained.

At this point, the meeting ran out of time, and it was agreed to take up the discussion of the remaining issues on the mailing-list.

Attendees

Steve Bostock	steveb@novell.com
Howard Brown	brown@ctron.com
Jeffrey Case	case@cs.utk.edu
John Cook	cook@chipcom.com
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Gary Ellis	garye@hpspd.spd.hp.com
Mike Erlinger	mike@mti.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Phillip Hasse	phasse@honchuca-emh8.army.mil
Mark Hoerth	mark_hoerth@hp0400.desk.hp.com
Ron Jacoby	rj@sgi.com
Mike Janson	mjanson@mot.com
Ken Jones	konkord!ksj@uunet.uu.net
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Mark Kerestes	att!alux2!hawk@uunet.uu.net
Kenneth Key	key@cs.utk.edu
Cheryl Krupczak	cheryl@cc.gatech.edu
Chao-Yu Liang	cliang@synoptics.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Lynn Monsanto	monsanto@eng.sun.com
David Perkins	dperkins@synoptics.com

Jason Perreault	perreaul@interlan.interlan.com
John Pickens	jrp@3com.com
Anil Rijsinghani	anil@levers.enet.dec.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Anil Singhal	nsinghal@hawk.ulowell.edu
Bob Stewart	rlstewart@eng.xyplex.com
Emil Sturniolo	emil@dss.com
Bruce Taber	taber@interlan.com
Geoffrey Thompson	thompson@synoptics.com
Dean Throop	throop@dg-rtp.dg.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Philip C. Wang	watadn!phil@uunet.uu.net
Drew Wansley	dwansley@secola.columbia.ncr.com
David Ward	dward@chipcom.com
Joseph Zur	fibrontics!zur@uunet.uu.net

3.3.6 Internet Accounting (acct)

Charter

Chair(s):

Cyndi Mills, `cmills@bbn.com`
Gregory Ruth, `gruth@bbn.com`

Mailing Lists:

General Discussion: `accounting-wg@wugate.wustl.edu`
To Subscribe: `accounting-wg-request@wugate.wustl.edu`
Archive:

Description of Working Group:

The Internet Accounting Working Group has the goal of producing standards for the generation of accounting data within the Internet that can be used to support a wide range of management and cost allocation policies. The introduction of a common set of tools and interpretations should ease the implementation of organizational policies for Internet components and make them more equitable in a multi-vendor environment.

In the following accounting model, this Working Group is primarily concerned with defining standards for the Meter function and recommending protocols for the Collector function. Individual accounting applications (billing applications) and organizational policies will not be addressed, although examples should be provided.

Meter <--> Collector <--> Application <--> Policy

First, examine a wide range of existing and hypothetical policies to understand what set of information is required to satisfy usage reporting requirements. Next, evaluate existing mechanisms to generate this information and define the specifications of each accounting parameter to be generated. Determine the requirements for local storage and how parameters may be aggregated. Recommend a data collection protocol and internal formats for processing by accounting applications.

This will result in an Internet Draft suitable for experimental verification and implementation.

In parallel with the definition of the draft standard, develop a suite of test scenarios to verify the model. Identify candidates for prototyping and implementation.

Goals and Milestones:

Done	Policy Models Examined.
Done	Internet Accounting Background Working Draft Written.
Done	Collection Protocols Working Papers Written.
Done	Internet Accounting Background final draft submitted as an informational document
Done	Collection protocol working papers reviewed.
Done	Collection protocol recommendation.
Mar 1992	Architecture submission as Internet Draft.
Jul 1992	Architecture submission as RFC
Done	Architecture working papers written.

Internet Drafts:

"Internet Accounting: Background", 05/13/1991, C. Mills, D. Hirsh, G. Ruth
<draft-ietf-acct-background-00.txt>

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

ACCT Minutes

- Status Review
  - Internet Accounting Background Document
  - Internet Accounting Architecture Document
    - Reporting Format
    - Rule Table
    - Meter Control
  - Liason with other Activities
    - SMDS
    - RLANMON MIB/OPSTATS
    - Interop BOF
    - SNMP/Security
- Working Agenda
  - Consensus on Internet Accounting Background Document
  - Consensus on Internet Accounting Architecture (Reporting Format)
    - Work on Rule Table
  - SNMP Concerns
    - Discuss Sample MIB
    - Find a home for the MIB
    - Security Concerns

Internet Accounting Background Document

The Working Group agreed to recommend the advance of the document to informational RFC when the following corrections are made:

The language about security requirements needs to be made more precise. The security area will provide a review of the amended text. Mention that the collection protocol is responsible for providing the necessary security, and therefore details of the security mechanisms are outside the scope of the document.

An explanation of the trade-offs between accounting on entry and exit (or both) in a router should be included. This discussion should include examples of performance impact, billing for offered load vs. billing for delivered load, reconciliation of counts with neighbors, and fragmentation. Note that SMDS has chosen to count on exit only.

Internet Accounting Architecture

The Accounting Architecture Document will be edited and then placed in the Internet-Draft area for access before the next meeting. The Architecture Document currently covers much of the material which was originally intended for the Meter Services Document. Mark Seger contributed many of the ideas and deserves special mention.

In the reporting format, the current link-level address is an insufficient description. It is amended to be the "adjacent address", as in the address of the next lower protocol layer carrying the internet packet. This concept should be explained in the architecture. OSPF describes a similar adjacency and may provide a useful explanation.

The motivation for the polling vs. interval reporting discussion in the architecture document should be made clearer.

The rule table was reviewed and will be re-issued in another format. It was agreed that some form of rule pre-processor will probably be needed for network manager sanity.

The manner in which the rule table for internet accounting forms a tree structure should be related to routing trees.

The notation for the rule table should be amended. The binary pattern matching scheme should be made up of the character set "0 1 *", where "0" requires that the corresponding bit in the value to be searched be 0 for a match to occur, "1" that the bit be 1, and "*" where then corresponding bit may be either 1 or 0. The overall notation should be restructured to make the formation of a MIB easier.

Because of the limitations of SNMP, it is difficult to access accounting information in table format. The appropriate set of protocol contortions needs to be investigated to return accounting information accurately and efficiently.

The maximum lifetime of a flow should be determined by the managing entity rather than by the meter. The manager must be able to maintain a clean state - e.g., insure that a record is fetched and a new record is started for the same flow as an atomic operation. Also, the fetched record should be stored at least temporarily (for a "short" time, "short" to be determined by the manager) in order to allow for a retransmission (i.e., repeated request).

The mechanisms for controlling data loss should be simplified to one or two parameters. If a meter is in danger of buffer overflow, probably the buffers will overflow and data will be dropped before the manager can take action anyway. This need further consideration.

The flags grouping data need further definition. The notion that some group of flows may be categorized as "expendable" (discard these flows first) or "essential" (avoid discarding these flows if at all possible) should be further explained.

SNMP Concerns

Jesse Walker wrote a sample MIB illustrating techniques that will be needed for expressing the accounting reporting format in terms that are compatible with SNMP.

The possibility of including of the MIB as a subtree of the RLANMON MIB was discussed and rejected by the Working Group. The two chief reasons for this were:

1. The RLANMON MIB is currently progressing towards draft standard and the late addition of accounting might hinder their progress.

2. Devices other that remote LAN monitoring equipment also perform accounting functions, so it is inappropriate to demand that these other devices implement the remote LAN monitoring MIB. It was pointed out that the RLANMON MIB could be structured so that only the accounting option is present in the MIB, but this was not considered a sufficient reason for homing the accounting MIB to RLANMON.

Internet Accounting requires that SNMP be able to dump tables in a more efficient and atomic fashion. FTP might be a preferred method of reporting data, but not all meters can be expected to support FTP due to memory limitations and abbreviated protocol stacks.

There are concerns that SNMP security may not be sufficient for accounting purposes.

Note change in location of list.

To join the accounting Working Group list, please send mail to accounting-wg-request@wugate.wustl.ed with "SUBSCRIBE" in the subject line. To leave the accounting Working Group mailing list, send mail to the same address with the subject line "UNSUBSCRIBE".

Attendees

Gigi Chu	gigic@hpspd.spd.hp.com
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Alan Emtage	bajan@cc.mcgill.ca
Shawn Gallagher	gallagher@quiver.enet.dec.com
Phillip Gross	pgross@nis.ans.net
Ittai Hershman	ittai@nis.ans.net
Mark Hoerth	mark_hoerth@hp0400.desk.hp.com
Mike Janson	mjanson@mot.com
Deidre Kostick	dck2@sabre.bellcore.com
Tim Lee-Thorp	ngc!tim@uunet.uu.net
Peter Liebscher	plieb@sura.net
Joshua Littlefield	josh@cayman.com
Cyndi Mills	cmills@bbn.com
Dennis Morris	morrisd@imo-uvax.dca.mil
Bradford Parker	brad@cayman.com
Robert Reschly	reschly@brl.mil

Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com.kr
Gregory Ruth	gruth@bbn.com
John Scudder	jgs@merit.edu
Robert Shirey	shirey@mitre.org
Anil Singhal	
Roxanne Streeter	streeter@nsipo.nasa.gov
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov
Claudio Topolcic	topolcic@nri.reston.va.us
William Versteeg	bvs@nrc.com
Jesse Walker	walker@eider.enet@decpa.dec.com
Gerard White	ger@concord.com
Osmund de Souza	desouza@osdpc.ho.att.com

3.3.7 Management Services Interface (msi)

Charter

Chair(s):

Oscar Newkerk, newkerk@decwet.enet.dec.com
Sudhanshu Verma, verma@hpindbu.cup.hp.com

Mailing Lists:

General Discussion: msiwg@decwrl.dec.com
To Subscribe: msiwg-request@decwrl.dec.com

Description of Working Group:

The objective of the Management Services Interface Working Group is to define a management services interface by which management applications may obtain access to a heterogeneous, multi-vendor, multi-protocol set of manageable objects.

The service interface is intended to support management protocols and models defined by industry and international standards bodies. As this is an Internet Engineering Task Force Working Group, the natural focus is on current and future network management protocols and models used in the Internet. However, the interface being defined is expected to be sufficiently flexible and extensible to allow support for other protocols and other classes of manageable objects. The anticipated list of protocols includes Simple Network Management Protocol (SNMP), OSI Common Management Information Protocol (CMIP), CMIP Over TCP (CMOT), Manufacturing Automation Protocol and Technical Office Protocol CMIP (MAP/TOP CMIP) and Remote Procedure Call (RPC).

Goals and Milestones:

Done	Initial version of the Internet Draft placed in the Internet-Drafts directory
Done	Revised version of the draft from editing meetings placed in the Internet-Drafts directory
Aug 1990	Initial implementation of the prototype available for test.
Done	Revised draft based on the implementation experience submitted to the RFC editor.

CURRENT MEETING REPORT

Reported by Oscar Newkerk/DEC

MSI Minutes

Minutes of the Management Services API Working Group at the Atlanta meeting of the IETF.

The first item discussed at the meeting was the motion to dissolve the Working Group. The reason given for this motion was that the solution that we had been developing would require changes in areas outside of the Group's original Charter. These areas would include a general database/method of mapping any IETF defined MIB to both the SNMP SMI and the ISO SMI, and a general object naming scheme that would allow for the mapping of ISO Distinguished Names to and from SNMP variables.

The group voted unanimously to disband the Working Group and to allow the MSI API draft document to be purged from the internet-drafts directories when they expire.

The next item discussed was a proposal from the mailing list that the Working Group reform or reconstitute itself to produce a standard API for SNMP applications. It was suggested by the Chair that this should be addressed by the normal process of forming a Working Group. If there is interest in this activity, then the Area Director, James Davin, should be approached with a proposed Charter.

Attendees

Steve Bostock	steveb@novell.com
Jeffrey Case	case@cs.utk.edu
Gary Ellis	garye@hpspd.spd.hp.com
Shari Galitzer	shari@gateway.mitre.org
Mark Kepke	mak@hpcndk.cnd.hp.com
Oscar Newkerk	newkerk@decwet.enet.dec.com
Marshall Rose	mrose@dbc.mtview.ca.us
Mark Sleeper	mws@sparta.com
David Waitzman	djw@bbn.com

3.3.8 OSI Internet Management (oim)

Charter

Chair(s):

Lee LaBarre, cel@mbunix.mitre.org
Brian Handspicker, bd@vines.enet.dec.com

Mailing Lists:

General Discussion: oim@mbunix.mitre.org
To Subscribe: oim-request@mbunix.mitre.org
Archive:

Description of Working Group:

This Working Group will specify management information and protocols necessary to manage IP-based and OSI-based LANs and WANs in the Internet based on OSI Management standards and drafts, NIST Implementors Agreements and NMF Recommendations. It will also provide input to ANSI, ISO, NIST and NMF based on experience in the Internet, and thereby influence the final form of OSI International Standards on management.

Goals and Milestones:

TBD Develop implementors agreements for implementation of CMIP over TCP and CMIP over OSI.

TBD Develop extensions to common IETF SMI to satisfy requirements for management of the Internet using OSI management models and protocols.

TBD Develop extensions to common IETF MIB-II to satisfy requirements for management of the Internet using OSI management models and protocols.

TBD Develop prototype implementations based on protocol implementors agreements, IETF OIM Extended SMI and Extended MIB.

TBD Promote development of products based on OIM agreements.

TBD Provide input to the ANSI, ISO, NIST and NMF to influence development of OSI standards and implementors agreements.

TBD Completion of the following drafts: Implementors Agreements, Event Management, SMI Extensions, MIB Extensions, OSI Management Overview, Guidelines for the Definition of Internet Managed Objects.

Request For Comments:

RFC 1095 "Common Management Information Services and Protocol over TCP/IP CMOT"

RFC 1189 "The Common Management Information Services and Protocols for the Internet"

RFC 1214 "OSI Internet Management: Management Information Base"

3.3.9 Remote LAN Monitoring (rmonmib)

Charter

Chair(s):
Mike Erlinger, mike@mti.com

Mailing Lists:
General Discussion: rmonmib@mti.com
To Subscribe: rmonmib-request@mti.com
Archive:

Description of Working Group:

The LAN Monitoring MIB Working Group is chartered to define an experimental MIB for monitoring LANs.

The Working Group must first decide what it covers and what terminology to use. The initial thought was to investigate the characteristics of some of the currently available products (Novell's LANtern, HP's LanProbe, and Network General's Watch Dog). From this investigation MIB variables will be defined. In accomplishing our goals several areas will be addressed. These include: identification of the objects to place in the MIB, identification of the tree structure and corresponding Object ID's for the MIB elements, generation of the ASN.1 for these new elements, and a test implementation.

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discussion and final approval of Charter; discussion and agreement on models and terminology. Make writing assignments.
Done	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1991	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Michael Erlinger/Micro Technology

Remote LAN Monitoring Minutes

The RLANMIB Working Group met on Monday, July 29th at the Atlanta IETF.

Actions:

By consensus the group decided that the Working Group should change its name to RMON-MIB. This better reflects the structure of the MIB being developed by the Working Group. The Chair will change the mail list and inform the IETF administration. This change in no way is meant to change the Charter. Such activities may take place, but that is another issue.

By consensus the Working Group decided that the RMON MIB should begin the standards process. This action would take place after several changes are incorporated into a new draft.

What to Do Next

It is recognized that the RMON MIB will become a standard only through implementation. Users and vendors were encouraged to review the MIB by implementation. Recommendations for change and modification will be driven by implementation.

The Working Group discussed possible Token Ring and FDDI activities. There was definite interest (14 for Token Ring and 12 for FDDI), but the overlap of the two groups was high (8). It was decided that if given a choice between the two MIBs, Token Ring was of more immediate concern. The Chair took on the task of coordinating the development of a Charter for a Token Ring Extensions Working Group for the RMON MIB.

Attendees

Karl Auerbach	karl@eng.sun.com
Steve Bostock	steveb@novell.com
Theodore Brunner	tob@thumper.bellcore.com
John Burruss	jburruss@wellfleet.com
Jeffrey Case	case@cs.utk.edu
Cho Chang	chang_c@apollo.hp.com
Gigi Chu	gigic@hpspd.spd.hp.com
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com
Nabil Damouny	nabil@tdd.sj.nec.com
James Davin	jrd@ptt.lcs.mit.edu
Gary Ellis	garye@hpspd.spd.hp.com

Mike Erlinger	mike@mti.com
Bill Fardy	fardy@ctron.com
Jeff Fried	jmf@relay.proteon.com
Phillip Hasse	phasse@honchuca-emh8.army.mil
Mark Hoerth	mark_hoerth@hp0400.desk.hp.com
Ron Jacoby	rj@sgi.com
Mike Janson	mjanson@mot.com
Mark Kepke	mak@hpcndk.cnd.hp.com
Charles Kimber	ckimber@dcp.com
Bobby Krupczak	rdk@cc.gatech.edu
Cheryl Krupczak	cheryl@cc.gatech.edu
Nik Langrind	nik@shiva.com
Tim Lee-Thorp	ngc!tim@uunet.uu.net
Chao-Yu Liang	cliang@synoptics.com
Ron Mackey	rem@dsiinc.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Donna McMaster	mcmaster@synoptics.com
Lynn Monsanto	monsanto@eng.sun.com
Norman Patty	
David Perkins	dperkins@synoptics.com
John Pickens	jrp@3com.com
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Anil Singhal	nsinghal@hawk.ulowell.edu
Geoffrey Thompson	thompson@synoptics.com
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@rm1.uucp
William Versteeg	bvs@nrc.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Philip C. Wang	watadn!phil@uunet.uu.net
Drew Wansley	dwansley@secola.columbia.ncr.com
Gerard White	ger@concord.com
Jeff Young	jsy@cray.com

3.3.10 Simple Network Management Protocol (snmp)

Charter

Chair(s):
Marshall Rose, mrose@dbc.mtview.ca.us

Mailing Lists:
General Discussion: snmp-wg@nisc.nyser.net
To Subscribe: snmp-wg-request@nisc.nyser.net
Archive:

Description of Working Group:

Oversee development of SNMP-related activity, especially the Internet-standard SMI and MIB. This Working Group is ultimately responsible for providing workable solutions to the problems of network management for the Internet community.

Goals and Milestones:

Aug 1990 Finish SNMP Authorization draft.

Ongoing Coordinate the development of various experimental MIBs.

Internet Drafts:

"SNMP Over IPX", 08/27/1990, Raymond Wormley <draft-ietf-snmp-snmpoveripx-00.txt>

"Definitions of Managed Objects for the Ethernet-like Interface Types", 09/26/1990, John Cook <draft-ietf-snmp-ethernetmib-05.txt>

"Use of the Community String for SNMP Proxys", 10/05/1990, Richard Fox <draft-ietf-snmp-proxys-01.txt>

"Definitions of Managed Objects for the SIP Interface Type", 11/07/1990, Kaj Tesink <draft-ietf-snmp-smdssipmib-03.txt>

"SNMP Communications Services", 04/23/1991, Frank Kastenholz <draft-ietf-snmp-commservices-00.txt>

"Comments on SNMP Proxy via Use of the @ sign in an SNMP Community", 04/29/1991, Jeff Case, et. al. <draft-ietf-snmp-proxycomments-00.txt>

Request For Comments:

RFC 1155 "Structure and Identification of Management Information for TCP/IP-based Internets"

RFC 1156 "Management Information Base for Network Management of TCP/IP-based internets"

RFC 1157 "A Simple Network Management Protocol (SNMP)"

RFC 1158 "Management Information Base for Network Management of TCP/IP-based internets: MIB-II"

RFC 1161 "SNMP over OSI"

RFC 1162 "Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542) Management Information Base"

RFC 1212 "Concise MIB Definitions"

RFC 1213 "Management Information Base for Network Management of TCP/IP-based internets: MIB-II"

RFC 1215 "A Convention for Defining Traps for use with the SNMP"

RFC 1229 "Extensions to the Generic-Interface MIB"

RFC 1230 "IEEE 802.4 Token Bus MIB"

RFC 1231 "IEEE 802.5 Token Ring MIB"

RFC 1232 "Definitions of Managed Objects for the DS1 Interface Type"

RFC 1233 "Definitions of Managed Objects for the DS3 Interface Type"

RFC 1238 "CLNS MIB - for use with Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542)"

CURRENT MEETING REPORT

Reported by James Davin/MIT

SNMP Minutes

SIP MIB

Tracy Cox of Bellcore presented the current MIB for the SMDS Interface Protocol. It was noted that the document largely exports those objects found in an SMDS switch that conforms to the Bellcore TA for SMDS. Numerous minor editorial comments were made and accepted.

Once this editing is complete, the Working Group recommends this document to the IESG for consideration as a Proposed Standard.

Ether-Like MIB

Chuck Davin, the Area Director, reported that the IAB had concerns regarding the ether-like MIB which had been produced by the Working Group and recommended by the IESG for entry onto the standards track. These concerns were in the area of 14 mandatory objects which might not be available on a chipset that is minimally 802.3-compliant.

The Working Group met in an ad hoc fashion the next day to edit the document to include the objects as optional. A liaison statement for the IAB, expressing the concern of the Working Group was unanimously adopted. (The minutes of this ad hoc group are below).

At the open plenary, the liaison statement was read. Although the ensuing discussion lasted nearly 1-1/2 hours, some points were crystalized: the IAB Chair, Vint Cerf, noted that the "IAB had consulted it's 802.3 expert to consider the matter". Vint continued by saying that the IAB's expert, IAB member Tony Lauck of DEC, had been on vacation, and this delayed the disposition of the MIB. Jeffrey Case, a member of the SNMP Working Group, disagreed with Cerf, noting that the IAB's 802.3 expert was not on vacation–as "the IAB's expert in this matter was not a member of the IAB nor any of his employees, but rather the IESG and the Working Group which produced the MIB." Another SNMP Working Group member, Karl Auerbach, observed that to a mean-spirited observer it might appear that one vendor, i.e., DEC, had purposely interfered with the MIB out of self-interest. Auerbach prefaced his remarks by indicating that he wasn't advocating that perspective. Nonetheless this raised several issues concerning the propriety of the IAB's actions and their reporting structure, which was subsequently discussed ad nauseum.

List of MIBs

Dave Perkins of SynOptics presented a comprehensive list of all known MIBs. This document will be revised on a regular basis and likely made available via ftp/mail.

Implementation Reports

The Chair queried the group of their implementation of various management technologies on the standards track:

```
Manager Agent
Concise Definitions high high
MIB-II high high
Token Bus ~15 1/2
Token Ring ~15 4
Interface Extensions ~10 5
DS1 ~5 1-1/2
DS3 ~5 0
```

Because of the near-universal implementation of Concise/MIB-II, these documents are being recommended by the Working Group to the IESG for advancement. On the remaining documents, Manager vendors complained of a lack of agents to test with. This issue will be raised again after the Interop SNMP Solutions Showcase. Hopefully this will provide better implementation experience.

Minutes of the Ad Hoc Meeting
Chair: Frank Kastenholz
Date: Thursday, August 1, 1991

An ad hoc meeting of the SNMP Working Group convened the morning of August 1, 1991 to revise the Ethernet MIB according to feedback from the IAB.

The IAB, after reviewing the Ethernet MIB, expressed concern about several of the variables of the MIB and their mandatory status. Specifically, the IAB felt that implementation of the following variables not be required in order for an implementation to claim compliance with the MIB:

```
dot3MacSublayerStatus
dot3MulticastReceiveStatus
dot3TxEnabled
dot3StatsSQETestErrors
dot3StatsDeferredTransmissions
dot3StatsLateCollisions
dot3StatsExcessiveCollisions
dot3StatsInternalMacTransmitErrors
```

```
dot3StatsCarrierSenseErrors
dot3StatsExcessiveDeferrals
dot3StatsFrameTooLongs
dot3StatsInRangeLengthErrors
dot3StatsOutOfRangeLengthFields
dot3StatsInternalMacReceiveErrors
```

The Working Group discussed several different strategies for revising the MIB per the IAB's comments. These included:

- Splitting the dot3Table and dot3StatsTable each into two tables, giving a total of four tables. Each of the original tables would be split into a table containing the variables acceptable to the IAB and a table containing the variables with which the IAB had concerns.

  This option was rejected when it was realized that the dot3Table would be split into two tables, one containing two variables and the other containing three. Such small tables were considered to be very inefficient by the Working Group.

- Assigning optional STATUS to the variables that the IAB had an issue with.

  This option was rejected when it was explained that the unit of conformance for MIBs is the mib group.

The Working Group settled on the following strategy:

- The Ethernet MIB will be divided into two tables. Each table will comprise a separate MIB Group. One table will contain the objects which the IAB found acceptable and implementation of that group will be mandatory. The second table will contain all variables with which the IAB had an issue and implementation of that table will be optional.

  The text on the MIB document that allows returning 0 for variables for which there is no underlying hardware support will be removed.

The Working Group believes that this MIB structure is less useful than the structure as submitted to the IAB because the requirements for conformance have been reduced. Vendors can implement only the mandatory group and be able to claim compliance with the MIB. However, the Working Group believes that this group does not contain all variables that are necessary for effective management of Ethernet-like interfaces.

The Working Group also believes that several extremely useful variables, as determined by implementation experience, will not be widely implemented because those variables have been moved to the optional portion of the MIB. Specifically, implementation and operational experience have shown the following five variables to be extremely valuable in detecting and diagnosing network problems:

1. dot3StatsDeferredTransmissions
2. dot3StatsLateCollisions
3. dot3StatsExcessiveCollisions
4. dot3StatsExcessiveDeferrals
5. dot3StatsFrameTooLongs

The Working Group then developed the following statement for including in the the Minutes of the meeting. The Working Group also directed the Chair to read this statement in the open IESG meeting on August 1, 1991.

LIAISON STATEMENT

The IETF meetings, unlike others, are open. All parties, even IAB members, are encouraged to attend and make their positions known and have them argued in the open. It is important to appreciate that development of technology does not occur in a vacuum for a community as large as the internet. It is also important to appreciate that the needs of the community must be met in a timely fashion.

IETF has worked over two years on the Ethernet MIB. The MIB was approved by the Working Group, the IETF plenary, and the IESG. Only at that late date, the IAB unilaterally raised issues that had been previously resolved. Because of the pressing needs of the community we feel coerced into revising the ethernet MIB as we have done today. Unfortunately the technical merit of this document has suffered owing to the lack of 802.3 expertise available to us on such short notice. As a result of this decline in technical quality, this document is less useful for management, and the manageability of the internet will suffer.

It must be emphasized that during the two years of preparation, numerous 802.3 experts contributed to the process and all the issues later raised by the IAB were thoroughly discussed, repeatedly argued, and resolved to the satisfaction of the experts. We feel that the editing session, hastily convened owing to the misplaced paternalism of the IAB, has undone some of the work of those experts.

Attendees

Steve Alexander	stevea@i88.isc.com
Karl Auerbach	karl@eng.sun.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
David Bridgham	dab@asylum.sf.ca.us
John Burruss	jburruss@wellfleet.com
Jeffrey Case	case@cs.utk.edu
Chi Chu	chi@sparta.com
Gigi Chu	gigic@hpspd.spd.hp.com
Henry Clark	henryc@oar.net
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com
Dave Cullerot	cullerot@ctron.com
James Davin	jrd@ptt.lcs.mit.edu
Gary Ellis	garye@hpspd.spd.hp.com
Mike Erlinger	mike@mti.com
Jeff Fried	jmf@relay.proteon.com
Shari Galitzer	shari@gateway.mitre.org
Shawn Gallagher	gallagher@quiver.enet.dec.com
Kenneth Goodwin	goodwin@psc.edu
Phillip Hasse	phasse@honchuca-emh8.army.mil
Ron Jacoby	rj@sgi.com
Mike Janson	mjanson@mot.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Mark Kepke	mak@.hpcndk.cnd.hp.com
Kenneth Key	key@cs.utk.edu
Christopher Kolb	kolb@psi.com
Deidre Kostick	dck2@sabre.bellcore.com
Bobby Krupczak	rak@cc.gatech.edu
Cheryl Krupczak	cheryl@cc.gatech.edu
Chao-Yu Liang	
Shane MacPhillamy	slm@netrix.com
Ron Mackey	rem@dsiinc.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Robert Meierhofer	
Lynn Monsanto	monsanto@eng.sun.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Michael Patton	map@lcs.mit.edu
David Perkins	dperkins@synoptics.com
John Pickens	jrp@3com.com
Brian Price	brian@bss.com

Marshall Rose	mrose@dbc.mtview.ca.us
Jonathan Saperia	saperia@decwrl.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Timon Sloane	peernet!timon@uunet.uu.net
Bob Stewart	rlstewart@eng.xyplex.com
Emil Sturniolo	emil@doe.dss.com
Bruce Taber	taber@interlan.com
Maurice Turcotte	dnmrt@rm1.uucp
David Waitzman	djw@bbn.com
Drew Wansley	dwansley@secola.columbia.ncr.com
David Ward	dward@chipcom.com
Gerard White	ger@concord.com
Linda Winkler	lwinkler@anl.gov
Mark Wood	markl@dsiinc.com
Joseph Zur	fibrontics!zur@uunet.uu.net

3.3.11 X.25 Management Information Base (x25mib)

Charter

Chair(s):
Dean Throop, throop@dg-rtp.dg.com

Mailing Lists:
General Discussion: x25mib@dg-rtp.dg.com
To Subscribe: x25mib-request@dg-rtp.dg.com
Archive: dg-rtp.dg.com:x25mib/Current.Mail

Description of Working Group:

This Working Group will produce a set of three documents that describe the Management Information Base for X.25. The first document will specify the objects for the X.25 Link Layer. The second document will specify the objects for the X.25 Packet Layer. The third document will specify the objects for managing IP over X.25. The Working Group need not consider the Physical Layer because the "Definition of Managed Objects for RS-232-like Hardware Devices" already defines sufficient objects for the Physical Layer of a traditional X.25 stack. Any changes needed at the Physical Layer will be addressed as part of that activity.

The X.25 object definitions will be based on ISO documents 7776 and 8208 however nothing should preclude their use on other similar or interoperable protocols (i.e., implementations based on CCITT specifications).

The objects in the Link and Packet Layer documents, along with the RS-232-like document, should work together to define the objects necessary to manage a traditional X.25 stack. These objects will be independent of any client using the X.25 service. Both of these documents assume the interface table as defined in MIB-II contains entries for the Link and Packet Layer interfaces. Thus these documents will define tables of media specific objects which will have a one to one mapping with interfaces of ifType ddn-x25, rfc877-x25, or lapb. The objects for the IP to X.25 convergence functions will be defined analogously with the ipNetToMedia objects in MIB II.

The Working Group will endeavor to make each layer independent from other layers. The Link Layer will be independent of any Packet Layer protocol above it and should be capable of managing an ISO 7776 (or similar) Link Layer provider serving any client. Likewise the X.25 Packet Layer objects should be independent of the Link Layer below it and should be capable of managing an ISO 8208 (or similar) Packet Layer serving any client.

The Working Group will also produce a third document specifying the objects for managing IP traffic over X.25. These objects will reside in their own table but will be associated with the X.25 interfaces used by IP. These objects will not address policy decisions or other implementation specific operations associated with X.25 connection management decisions except as explicitly described in existing standards. These objects will manage the packet flow between

IP and the X.25 Packet Layer specifically including observation of packet routing and diagnosis of error conditions. Progress on the Link and Packet Layer documents will not depend on progress of the IP over X.25 document. The IP over X.25 document will proceed on a time available basis after work on the Link and Packet Layer documents and as such the Link and Packet Layers may be completed before the IP over X.25 work.

All documents produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions (such as Concise MIB format). To the extent feasible, the object definitions will be consistent with other network management definitions. In particular ISO/IEC CD 10733 will be considered when defining the objects for the X.25 Packet Layer.

Goals and Milestones:

Done	Distribute first draft of documents and discuss via E-mail.
Done	Working Group meeting as part of IETF to review documents.
Sep 1991	Distribute updated documents for more E-mail discussion.
Nov 1991	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Jan 1992	Documents available with specified changes incorporated.

CURRENT MEETING REPORT

Reported by Dean Throop/Data General

X25MIB Minutes

The X25mib Working group met on Tuesday, July 30, 1991 at the Atlanta IETF meeting. The Working Group considered several aspects of the documents.

The Working Group first discussed adding an object identifier to identify the network to which a LAPB interface was connected. While it was agreed such an identifier could be provided, there wasn't enough justification presented by any participant to warrant adding the object. The Working Group was satisfied with the text stating that the ifdesc field for the interface should name the network. While this isn't very useful for network management software, it does make the information available in some form. An explicit object identifier can be added in the future if needed.

The Working Group then discussed the name of the x25PktStatInProviderInitiatedClears object. This is the only counter of received clear packets and as such the name could be simplified to x25PktStatInClears. Alternatively other clear counters could be defined. The attendees felt that remotely initiated clears were enough a part of normal operation that they need not be counted. Provider initiated clears however indicated a problem with the service from the provider and did justify being counted. The consensus of the attendees was not to add other objects. As for simplifying the name, the counter name does reflect what the object counts. The current name also sets precedence for naming other clear counters should future experience justify their existence. The Working Group decided to leave the name of the object as current defined.

The Working Group also discussed the differences indicated by the different types defined for the ifType field of an X.25 packet layer interface. A type of ddn-x25(4) indicates a simple interface using an algorithm for translating between X.25 address and IP addresses. An interface type of rfc877-x25(5) indicates a table is used for the address translation.

The Working Group then discussed expanding the IP over X25 MIB to include objects for X.25 call parameters. This would allow a manager to examine and change the X.25 parameters the IP over X.25 software would use to initiate an X.25 call. It was observed by the group that all users of X.25 would need a similar table. As an example, the IPX over X.25 interface will have X.25 call parameters that may be part of a future MIB. Rather than have each user of X.25 define their own objects, a table will be added to the X.25 MIB for X.25 call parameters. This table will include call user data, packet size, window size, charging information, and other parameters. The table will not include the destination X.25 address. This will allow one set of parameters to apply to several different destinations.

Other discussions in Atlanta concerned broadening the use of the LAPB MIB. Since LAPB and other Data Link protocols are very similar, the name of the LAPB MIB will be changed to HDLC and it will be presented to other Working Groups. It maybe possible to use the

HDLC MIB for more than just LAPB.

The HDLC, X.25 Packet Layer, and IP over X.25 draft MIBs will be updated and distributed to the x25mib@dg-rtp.dg.com mailing list for further discussion.

Attendees

Steve Alexander	stevea@i88.isc.com
Steve Bostock	steveb@novell.com
James Davin	jrd@ptt.lcs.mit.edu
Mark Kepke	mak@.hpcndk.cnd.hp.com
Evan McGinnis	bem@3com.com
John Pickens	jrp@3com.com
Dean Throop	throop@dg-rtp.dg.com

3.4 OSI Integration Area

Director(s):

- Ross Callon: callon@bigfut.enet.dec.com

Area Summary reported by Ross Callon/DEC

Change in Area Director

Rob Hagens has resigned as OSI Area Co-Director, effective at the end of the IETF meeting. I have greatly enjoyed working with Rob over the last three years, and am very sorry to see him resign as Co-Director (although Rob will fortunately continue to participate in IETF meetings, and Co-Chair the OSI General and X.400 Working Groups). The great growth in activity in the OSI Integration area, along with the progress that we are making in deployment of OSI in the Internet, could not have happened without Rob's help.

Working Groups Complete Their Mission

With the publication of the NSAP Guidelines (RFC 1237), the NSAP Working Group has now completed its task, and has been disbanded. I would like to thank Richard Collela of NIST for a job well done and completed. The further task of deploying NSAP addressing will proceed along with the related task of deploying CLNP, and will therefore be coordinated by the Network OSI Operations (NOOP) Working Group.

Also, the X.400 and X.400 Operations Working Groups have been merged. The new merged X.400 Working Group will be Co-Chaired by Rob Hagens and Alf Hansen.

Working Group Summaries

- Network OSI Operations (NOOP)

  The NOOP Working Group meeting was split into two parts: (1) A presentation by Walt Lazear of Mitre about Mitre's efforts on deploying OSI in their in-house network; and (2) A general discussion of routing and addressing plans that have been prepared for CLNP and NSAP addressing.

  Walt Lazear's presentation was very interesting. The bottom line appears to be that substantial progress is being made, but that we still have a way to go before OSI will offer a complete multi-vendor networking solution.

  "Routing plans" (plans that outline the relationship between NSAP addressing and routing for a particular environment) were distributed for Mitre, for CICNet, for DCA (now called DISA), and for the DoD Internet. There was also a general discussion of routing and addressing issues.

- OSI-General

  The OSI-General Working Group reviewed two sets of specifications: A specification for an FTAM/FTP gateway, and two documents specifying operation of Connection-Oriented/ConnectionLess interworking (CO/CL).

  Robert Slaski (author of the FTAM/FTP gateway specification) was present and led an in-depth review of the specification. The primary focus of the review was to take into account a comprehensive list of issues that were submitted by Larry Friedman of Digital Equipment Corporation. Robert had already incorporated a large number of the comments received into an updated draft brought to this review. A number of additional updates were identified, and Robert agreed to produce an updated version of the specification.

  The CO/CL specifications comprise two documents which talk about how to interwork between OSI applications operating over three different lower layer protocol stacks:

  1. TP4 over CLNP;
  2. TP0 over X.25;
  3. TP0 over RFC 1006 over TCP over IP.

  Again, a set of changes were identified, and will result in updated documents.

- Office Document Architecture (ODA)

  The ODA Working Group is developing guidelines for the use of the Office Document Architecture for the exchange of Compound documents (for example, including formatable text, bit-map graphics and geometric graphics, possibly Spreadsheets). The Working Group is defining how to use both SMTP and X.400 for interchange of ODA documents, and maintains close liaison with the SMTPEXT and osix400 Working Groups. A major part of the task of this Working Group is coordination of an ODA pilot project, which is currently underway.

  As a background for the discussions on Pilots, the current status of implementations was reviewed. Five different implementations from six different vendors were identified, and discussed in some detail (see Working Group report for specifics). Requirements for use of each implementation in the pilot project were also discussed in detail.

- X.400 Operations

  Alf Hansen and Rob Hagens are now Co-Chairs of the Working Group. The most significant work item completed by the old X.400 Working Group was an RFC describing how to use DNS to store RFC1148 mapping information. The status of this RFC is that it is awaiting proof of concept through implementation.

Because the two X.400 Working Groups have merged, the Working Group Charter will be updated to add a new goal: The Working Group will attempt to drive X.400 deployment in the Internet.

A roundtable discussion of the status of current X.400 services was provided by those Working Group members who are currently operating X.400 services.

Work continued on the X.400 Operational Requirements RFC, on "An X.400 Internet Strategy" and on "Requirements for X.400 PRMD's Operating in the Internet". Steve Hardcastle-Kille presented his draft RFC on '88->'84 downgrading. He accepted comments from the Working Group and will make some minor changes to the document. The working group recommended that this document, after minor edits, be progressed as an RFC on the standards track. There was a detailed discussion of the use of X.500 directory services for managing X.400 routing/mapping information.

- RFC1148bis Editing Meeting

  Steve Hardcastle-Kille led an ad-hoc editing meeting to discuss the Internet Draft on mapping between X.400 (1988)/ISO 10021 and RFC 822 (aka RFC-1148bis).

  The attendees unanimously agreed that the recommendation from this meeting was to proceed the Internet Draft on Mapping between X.400(1988)/ISO 10021 and RFC 822 to the RFC status.

  This document has also been reviewed by RARE WG1 in Europe, and the Chair of WG1 was present at this meeting.

- OSI Directory Services

  The OSI Directory Services Working Group is now jointly in the OSI and Applications areas.

  The OSI Directory Services Working Group is active in monitoring extensive X.500-based directory services pilot projects, and in producing a series of documents outlining operation of X.500 in the Internet.

  Liaison reports were provided from a number of organizations, including RARE WG3, ISO/CCITT, OSI Implementors Workshop, North American Directory Forum, the FOX project, the PSI White Pages project, the PARADISE project, the Australian Academic Research Network, NORDUNET, and the Coalition for Networked Information.

  Vint Cerf and Ross Callon reported on the status of progression of seven documents (currently internet drafts) which are basically complete and being progressed though the IESG/IAB review process. Vint remarked that he expected the documents to be published as RFCs quickly ("within a few weeks").

Also, the general overview and strategy documents were discussed. These should be merged. Steve agreed to combine and re-write these documents, in association with appropriate IESG and IAB members. There was also detailed discussion of a number of new documents. One document (Naming Guidelines) is nearing completion, and will be submitted as an RFC following an electronic mail discussion of the final editing. See the Working Group report for additional details.

- DISI

  The DISI Working Group is jointly in the OSI and User Services areas. It is covered in the User Services area report.

CURRENT MEETING REPORT

Reported by Erik Huizer/SURFnet

RFC1148 Editing BOF Minutes

Opening Remarks

Introduction:

Steve Hardcastle-Kille presented the most significant changes with respect to the one-before-last version of the Internet Draft on mapping between X.400 (1988)/ISO 10021 and RFC 822 (aka RFC-1148bis). These are summarized in Appendix J.

Discussion:

There was a short discussion on the effects of X.400 addresses encoded (on purpose or by accident) in DDA.RFC-822. The conclusion was that gateways should be able to cope with this by unpacking the address from the DDA, and then resubmitting the address to their tables. In a really good gateway, the body of the message will in such a case not be converted from X.400 to RFC 822 and back again. Steve explained the proper use of the gateway table.

Recommendation:

The attendees unanimously agreed that the recommendation from this meeting was to proceed the Internet Draft on Mapping between X.400(1988)/ISO 10021 and RFC 822 to the RFC status.

This document has also been reviewed by RARE Working Group 1 in Europe, and the Chair of Working Group 1 was present at this meeting. This editing meeting was also given authority by the X.400 Operations Working Group, which met earlier on the same day.

Attendees

Claudio Allocchio	claudio.allocchio@cosine-gw.infn.it
David Brent	brent@CDNnet.ca
Urs Eppenberger	eppen@v
Ned Freed	ned@innosoft.com
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Erik Huizer	huizer@surfnet.nl
Jim Knowles	jknowles@trident.arc.nasa.gov
Jack Liu	liu@koala.enet.dec.com
Peter Yee	yee@ames.arc.nasa.gov

3.4.1 Assignment of OSI NSAP Addresses (osinsap)

Charter

Chair(s):

Richard Colella, colella@osi3.ncsl.nist.gov

Mailing Lists:

General Discussion: ietf-osi-nsap@osi3.ncsl.nist.gov
To Subscribe: ietf-osi-nsap-request@osi3.ncsl.nist.gov

Description of Working Group:

The OSI NSAP Guidelines Working Group will develop guidelines for NSAP assignment and administration (AKA, the care and feeding of your NSAPs).

Assuming use of existing NSAP address standards, there are two questions facing an administration:

- Do I want to be an administrative authority for allocating NSAPs?
  - how do I become an administrative authority?
    - what organizations should expect to be an "administrative authority" in the GOSIP version 2.0 address structure?
    - where do I go to become an administrative authority?
  - what are the administrative responsibilities involved?
    - defining and implementing assignment procedures?
    - maintaining the register of NSAP assignments.
    - what are the advantages/disadvantages of being an administrative authority?
- Whether NSAPS are allocated from my own or some other administrative authority, what are the technical implications of allocating the substructure of NSAPs?
  - what should be routing domains?
    - implications of being a separate routing domain (how it will affect routes, optimality of routes, firewalls and information hiding).
    - organizing routing domains by geography versus by organization versus by network topology....
  - within any routing domain, how should areas be configured?
    - (same implications as above).

Goals and Milestones:

Done	Produce a paper describing guidelines for the acquisition and administration of NSAP addresses in the Internet.
Dec 1990	Have the paper published as an RFC.

Dec 1990 Have the paper incorporated, in whole or in part, into the "GOSIP User Guide" and the FNC OSI Planning Group document.

CURRENT MEETING REPORT

Chair: Richard Colella/NIST

Assignment of OSI NSAP Addresses

Upon publication of RFC 1237:"Guidelines for OSI NSAP Allocation in the Internet" (July 23, 1991), the Assignment of OSI NSAP Addresses Working Group concluded it's business.

3.4.2 Network OSI Operations (noop)

Charter

Chair(s):

Susan Hares, skh@merit.edu

Mailing Lists:

General Discussion: noop@merit.edu
To Subscribe: noop-request@merit.edu
Archive:

Description of Working Group:

The working group is chartered to work on issues related to the deployment of CLNP in the Internet. Initial activities include both deployment planning and education of regional and other conencted networks.

Initial planning efforts include the development of routing and management plans.

Goals and Milestones:

Aug 1991 Create tutorials for CLNP OSI routing protocols, including ES-IS, CLNP, IS-IS, and IDRP.

Aug 1991 Collect OSI Routing and Addressing plans into a Repository. Make the plans available at Merit.edu:/pub/iso/noop/plan

Ongoing Provide a forum to discuss these OSI Routing plans by email or in group discussions

Nov 1991 Collect 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.

Nov 1991 Collect list of OSI Network Layer NOC tools and publish a list.

Nov 1991 Collect Methods of OSI Network Layer Debugging and write a document describing these methods.

CURRENT MEETING REPORT

Chair(s): Sue Hares/Merit

NOOP Minutes

Report not submitted. See Area Report for summary.

Attendees

William Barns	barns@gateway.mitre.org
Randy Butler	rbutler@ncsa.uiuc.edu
Ross Callon	callon@bigfut.enet.dec.com
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Tom Easterday	tom@cic.net
Arlan Finestead	arlanf@ncsa.uiuc.edu
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Robert Hagens	hagens@cs.wisc.edu
Tony Hain	alh@es.net
Susan Hares	skh@merit.edu
Eugene Hastings	hastings@psc.edu
Paulina Knibbe	knibbe@cisco.com
Bobby Krupczak	rdk@cc.gatech.edu
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@nic.near.net
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Joe Ragland	jrr@concert.net
Yakov Rekhter	yakov@watson.ibm.com
Ron Roberts	roberts@jessica.stanford.edu
John Scudder	jgs@merit.edu
Mark Sleeper	mws@sparta.com
Michael St. Johns	stjohns@umd5.umd.edu
Ross Veach	rrv@uiuc.edu
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov

Agenda for NOOP working Group

1:30-1:35PM Discussion of extending NOOP for one more meeting

1:35-2:05PM Mitre's work with OS Walt Lazear

2:05-2:25PM Interop 91 OSI demo Sue Hares

2:25-2:40PM Break

2:40-3:15PM Routing Plans
ESNET
CICNET Questions
Mitre Questions
MERIT
Open Floor

3:15-3:30PM Network Debuggin Tools and Future

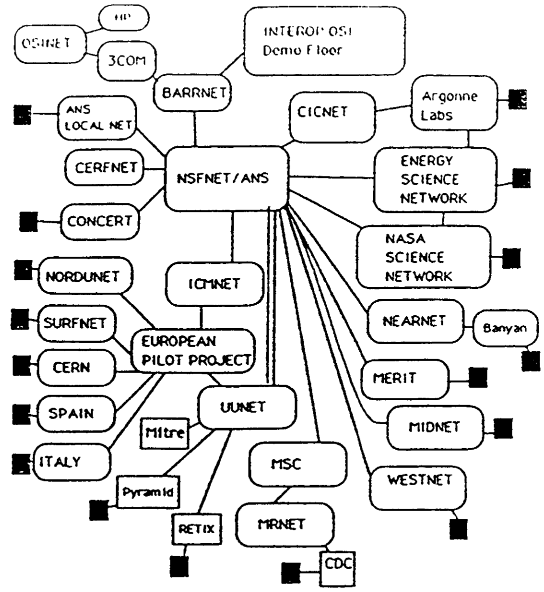

OSI InfraStructure in the Internet

Sue Hares NSFNET

Potential Networks in Wide Area INTEROP '91 OSI Demo - Part 1

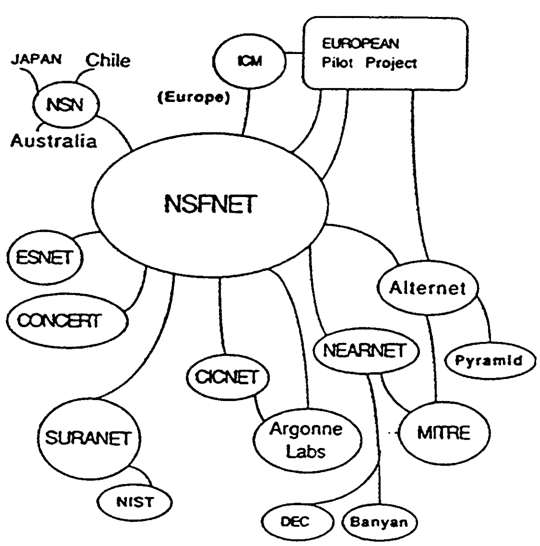

Concert OSI

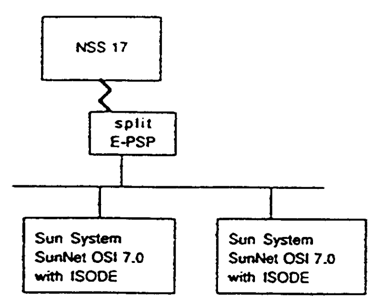

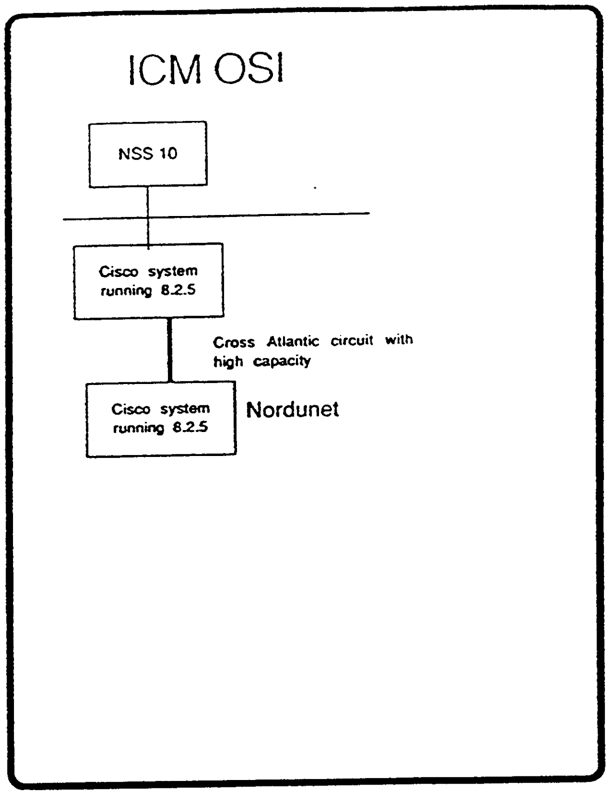
ICM OSI
NSS 10
Cisco system running 8.2.5
Cross Atlantic circuit with high capacity
Cisco system running 8.2.5
Nordunet

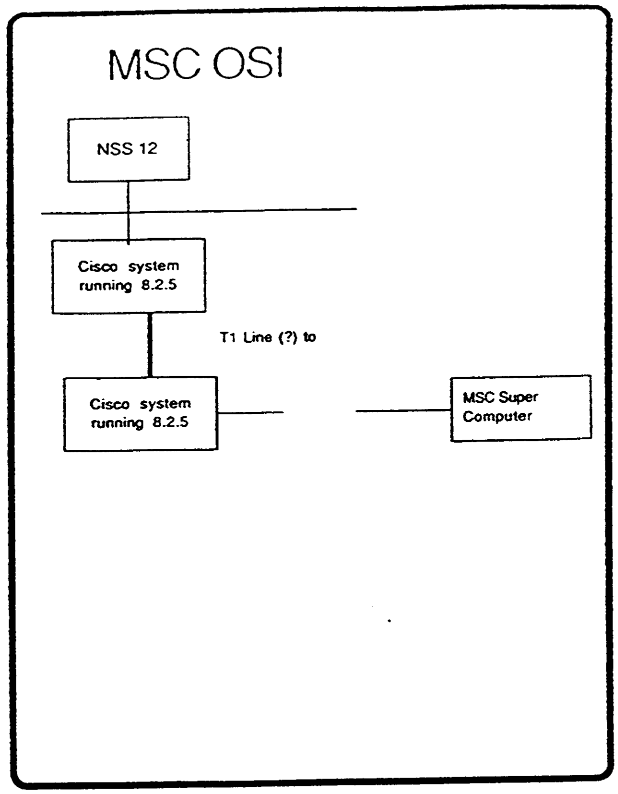
MSC OSI
NSS 12
Cisco system running 8.2.5
T1 Line (?) to
Cisco system running 8.2.5
MSC Super Computer

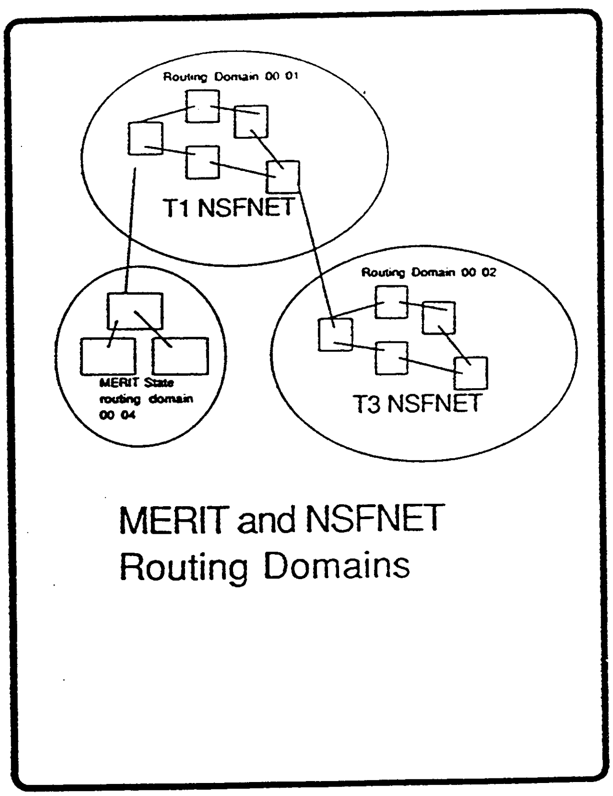
Routing Domain 00 01
T1 NSFNET
MERIT State routing domain 00 04
Routing Domain 00 02
T3 NSFNET
MERIT and NSFNET
Routing Domains

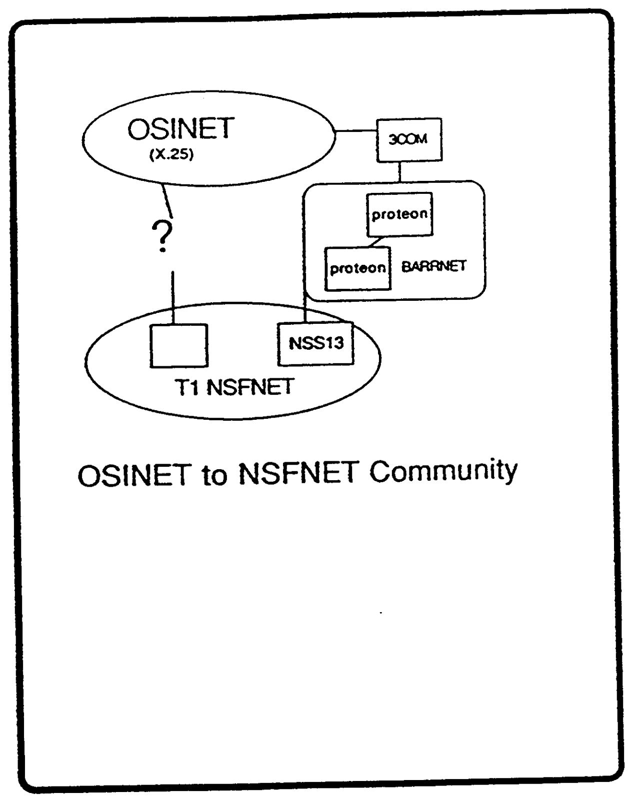
OSINET
(X.25)
?
3COM
proteon
proteon
BARRNET
NSS13
T1 NSFNET
OSINET to NSFNET Community

IDRP Demonstration
- IDRP routers

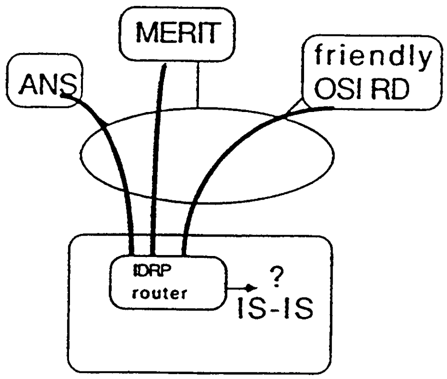

Interop Show Floor
Network

OSI Applications at Interop 91

X.500 with ISODE additions
X.400 - autoreply and delivery
report
FTAM
VT
X windows - Sun -> Cray

Mitre did Sun Code
Cray did UNICOS code
UNICOS 6.0

Other ideas: X.400-> FAX interface

3.4.3 OSI Directory Services (osids)

Charter

Chair(s):

Steve Kille, `S.Kille@cs.ucl.ac.uk`

Mailing Lists:

General Discussion: `ietf-osi-ds@cs.ucl.ac.uk`
To Subscribe: `ietf-osi-ds-request@cs.ucl.ac.uk`
Archive:

Description of Working Group:

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

Goals and Milestones:

Done — Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.

Done — Study the relationship of the OSI Directory to the Domain Name Service.

Ongoing — Maintain a Schema for the OSI Directory on the Internet

Ongoing — Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

Internet Drafts:

"X.500 and Domains", 01/31/1990, S.E. Kille <draft-ucl-kille-x500domains-03.txt, or .ps>

"A String Encoding of Presentation Address", 01/31/1990, S.E. Kille <draft-ucl-kille-presentationaddress-03.txt, or .ps>

"An Interim Approach to use of Network Addresses", 01/31/1990, S. Kille <draft-ucl-kille-networkaddresses-04.txt, or .ps>

"The COSINE and Internet X.500 Schema", 11/26/1990, P. Barker, S. Kille <draft-ietf-osids-cosinex500-05.txt>

"Replication and Distributed Operations Extensions to Provide an Internet Directory using X.500", 11/26/1990, S. Kille <draft-ietf-osids-replsoln-03.txt, or .ps>

"Using the OSI Directory to Achieve User Friendly Naming", 11/26/1990, S. Kille <draft-ietf-osids-friendlynaming-02.txt, or .ps>

"Replication Requirement to Provide an Internet Directory Using X.500", 11/26/1990, S. Kille <draft-ietf-osids-replication-03.txt, or .ps>

"Handling QOS (Quality of service) in the Directory", 03/20/1991, S.E. Kille <draft-ietf-osids-qos-01.txt, or .ps>

"Naming Guidelines for Directory Pilots", 03/21/1991, P. Barker, S.E. Kille <draft-ietf-osids-dirpilots-01.txt, .ps>

"DSA Naming", 03/21/1991, S.E. Kille <draft-ietf-osids-dsanaming-00.txt, or .ps>

"Schema for Information Resource Description in X.500", 06/14/1991, Chris Weider <draft-ietf-osids-resdescripx500-00.txt>

"Schema for NIC Profile Information in X.500", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-osids-nicprofilex500-00.txt>

"Interim Directory Tree Structure for Network Infrastructure Information", 06/14/1991, Chris Weider, Mark Knopper, Ruth Lang <draft-ietf-osids-treestructure-00.txt>

"Directory Requirements for COSINE and Internet Pilots (OSI-DS 18)", 07/09/1991, S.E. Hardcastle-Kille <draft-ietf-osids-requirements-00.txt, .ps>

"Generic Security Service Application Program Interface Overview and C bindings", 07/10/1991, John Wray <draft-ietf-cat-secservice-00.txt>

INTERIM MEETING REPORT

Reported by Steve Kille/UCL

OSIDS Minutes

Minutes of the Third IETF Directory Services (OSI-DS) Working Group Videoconference (April 11, 1991).

Agenda

This meeting was held as a videoconference at four sites: BBN; SRI (RIACS facility); ISI; UCL. Minutes were taken at each site, and this note is a compilation of those Minutes. In addition, there was a phone-in from Merit. The meeting was an interesting "first" in use of the videoconference technology in that:

- It was not a videoconference about videoconferencing
- Four sites were involved, one not in the US
- There were more than one or two participants at each site

This is a joint meeting with members of RARE WG3.

Pilot Activity

- PARADISE (David Goodman)

  PARADISE is a sub-project of the broader COSINE project sponsored under the umbrella of EUREKA by eighteen participating countries and aimed at promoting OSI to the academic, industrial and governmental research and development organizations in Europe. The countries involved are those of the EC, EFTA plus Yugoslavia.

  The partners funded by PARADISE besides UCL are:

  - The Networks Group at the University of London Computer Centre (ULCC), which is a service-oriented organization providing a range of facilities to the academic community in London and the entry point into the UK for IXI, the COSINE international X.25 backbone;

  - X-Tel Services Ltd, a software company based in Nottingham which currently provides service support to the UK Academic X.500 pilot; and

  - PTT Telematic Systems from the Netherlands, which in turn has subcontracted the Swiss and Finnish PTTs, and whose involvement is to create a forum for discussion on X.500 among the European carrier administrations.

The project also aims to have representation from all the participating countries, which in the majority of cases are the existing X.500 national pilots. Of the 18 countries involved, 12 are registered in the tree, including Ireland and Italy whose nodes were taken up this month. Most countries are using the QUIPU implementation developed at UCL. However, a French group have developed PIZARRO, which will form the basis of the emerging French pilot and, in Italy, a Torino-based company Systems Wizards are using DirWiz, which is currently the sole representative from Italy in the tree.

PARADISE recently announced an operational service providing a central configuration DSA with connectivity via IPSS, IXI, JANET (UK Joint Academic Network) and the Internet. This DSA contains the "root of the world node" and provides the glue at the top of the international DIT. By this summer a central DUA will be installed with public access via ULCC. Multilingual versions of this interface will be made available later in the project. Both these central services will be provided by ULCC, which will be offering a help desk with telephone and e-mail support.

- FOX (Bob Braden, Steve Hotz)

  Bob Braden remarked that the Internet funding agencies, as well as the IAB, were anxious to see an X.500 directory service infrastructure in the Internet, and that the FOX project was working toward this goal. He further noted that the FOX project wants to make every effort to make certain that it's effort are aligned with X.500 activities in other communities.

  Steve Hotz commented on the recently released directory activities report (for Internet and other North American efforts) that appeared in the March Internet Monthly Report. He asked for comments regarding contents of the report, additional efforts that should be contacted, and ideas on where else it should be distributed, in addition to the IMR.

  Steve announced that the FOX project has scheduled a phone conference for Wed. April 17th.

  The FOX project is a DARPA and NSF funded effort to provide a basis for operational X.500 deployment in the NREN/Internet. This work is being carried out at Merit, NSYERNet/PSI, SRI and ISI. ISI is the main contractor and responsible for project oversight.

  There are two primary thrusts of the FOX project:

  1. X.500 Infrastructure

     It is important that multiple interoperable platforms be available for deployment. FOX plans to examine and test the interoperability of the Quipu and

NIST-X.500 (Custos) implementations, and DNANS-X.500 if possible. In addition, FOX will explore X.500 interfaces to conventional database systems (one target is Sybase), an alternate OS platform (VM) for X.500 servers, and X-window based user interfaces.

2. X.500 Applications

   A long-range goal is to facilitate the use of X.500 for real Internet applications. FOX will first focus on making network infrastructure information available through X.500. This includes network and AS site contacts, topology information, and the NIC WHOIS service.

   A centrally managed X.500 version will be the first phase of a WHOIS service. Providing an X.500 version of a well-known widely-used service should promote the use of X.500 by Internet users. In addition, this effort will provide experience in designing X.500 applications. However, the manageability of this scheme will be short-lived, so the next step will be a design for a distributed version of WHOIS.

- RARE WG3 (Erik Huizer)

  Erik pointed out that WG3 was not a pilot activity, but rather an engineering group whose activities parallel those of IETF OSI-DS. WG3 is the directory services subgroup of the COSINE project, whose purpose was to handle technical aspects of directory service deployment. In the future, issues such as privacy, data management, and data update will receive more focus.

  He mentioned the efforts of the P2.2 project in user information services to build a meta-information server, which would contain data about network services worldwide. A commercial company (Level-7) has been contracted to provide this service.

- NADF (Einar Stefferud)

  Einar announced the release of NADF-123, a document on the organization of the North American DIT, and that they are currently soliciting comments.

  NADF-123 specifies that the current civilian infrastructure be used to organize the DIT, and pointed out some of the difficulties with other structures. In particular, U.S. organizations are registered at a state level, so difficulties arise if one were to normally place entities under the country level. NADF-123 proposes multiple-attribute RDNs to allow organizations that, in addition, want to be listed at the country level. This scheme deals with possible name conflicts that may result from multiple entities registered in different states.

S.K. Asked about timetable to build conforming directory services? E.S. Replied that different service providers vary widely in the stage of development of their services. What matters is the time when someone mounts the first shared DS.

NADF has also gotten directory providers to agree that they will share information about the DIT. Einar commented that this was a significant milestone.

Monthly Reporting

Hotz is working to coordinate the US submission; he offered that he had not had a chance to coordinate the International report with Goodman.

Goodman suggested model that complete status be given every six months and that incremental reports be given bimonthly.

Discussion followed regarding whether reports should be given by country. The Internet is international, whereas the DIT is structured by country. Goodman suggested that each country's efforts be summarized and an Internet summary be included as well. Hotz is working with OIW-DS to include their report as well.

Management of experimental object identifiers

Problem identified – experimental ids admitted to schema are changed; this forces a fast update cycle of document

Points:

1. No fundamental need to change oid when put into schema, but is a management problem.
2. Changing oid gives it an identity with schema.
3. Mixing concept of registry vs library of oids. Suggestion that library id numbers be created and given out with each.

Kille moved that Barker reflect the idea of 1 plus 3 in the schema document.

Discussion continued regarding:

1. The transition of oids from informal to formal. No conclusions.
2. IANA model. IANA process is mechanical, Kille feels that a purely administrative approach to the schema is not advisable – technical and aesthetic concerns must be

incorporated as well. No conclusions.

Activities Documents

Hotz discussed further status of the North American and Internet activities activities report. He indicated that he was talking with Einar about including entries for ANSI USA RAC and SG-D MHS-MD, and Youbong Weon-Yoon about OIW DSSIG reports.

David Goodman discussed the tentative plans for an international report, which is to be produced either every two or three months. Goodman asked Hotz if he would provide a summary for the Internet and other US activities. Hotz agreed, and asked for guidance in what was needed. Hotz and Goodman will continue off-line.

Management of Experimental Object IDs

There is some question whether (and how) provisions should be made for very fast allocation of OIDs for experimental efforts, in light of the consequent revocation problems. This is not facilitated by the current mechanisms for including new OIDs into the standard schema.

Stefferud commented that a plan which included reassignment of OIDs is a bad idea, as has been seen before with other assigned numbers.

Braden suggested that IANA be assigned an OID space and that mechanisms, already in place to assign Internet numbers, be used to allocate OIDs.

A comment from UCL was that this approach would lead to many name spaces, and this could make to various problems in managing the globally standard OIDs. How would one know where to find all of those currently supported?

Paul Barker noted that different OID requests and their intended applications had different characteristics, and that it might be possible to decide on a case-by-case basis which mechanism should be used. ACTION ITEM: Paul Barker will write this idea up.

Kille pointed out that the OID aliasing mechanism in Quipu could be used to facilitate transition when OIDs are reassigned. He added that maybe this mechanism should be required in directory pilots.

Einar has a document concerning number assignment. He will distribute it via email, where this topic can be further pursued.

Braden commented that a directory services requirements document, in a similar vein as the host and gateway requirements documents, would be useful. Among other things, this could document the OIDs required to interoperate, and solve the question of where to look for the officially required OIDs.

Kille expressed concern that this would only document Internet requirements and not be sufficient for international needs.

Braden pointed out that one needed to start somewhere and that this was an IETF working group meeting. He went on to note that the Internet is an international activity, and is growing more so. Einar Stefferud commented that the Internet will only remain an American effort so long as the European community insists that it is.

Kille asked about who should be responsible for producing a requirements document. Braden replied that this decision should be taken up with the IESG coordinator.

Document Status

Steve Kille organized this topic into three areas: strategy document, IETF OSI-DS documents, and others.

- Strategy Document

  Steve Kille noted that this has been submitted as an RFC. Bob Braden, who is serving as interim RFC editor will help see this along. Braden and Kille will follow this up off-line.

- IETF OSIDS Documents

  Steve Kille enumerated these seven documents. Braden inquired about the plans for progressing these documents.

  - scheme document - standards track
  - interim network names - standards track
  - representing presentation addresses - info only, maybe standards??
  - replication requirements - statement, info only
  - replication solutions - standards
  - user friendly naming - standards
  - X.500 and domain names - experimental, maybe standards track later??

  Braden indicated that he believed some of these should be offered as experimental RFCs now. Kille ask for a clarification of experimental versus standards track RFCs. Braden pointed out that there was not a strong relationship between experimental and standards track RFCs. It is not the case that standards track RFCs always (or never) start out as experimental.

Other documents

- Naming Guidelines

  Paul Barker discussed the addition of support for multilingual names, adding that it requires considerable effort. As an example, one can consider names of organizational units and departments. One would want people worldwide to be able to understand these attributes. This suggests multi-lingual tagging of commonly used names. The various structuring of human names is another issue to be resolved.

Einar Stefferud remarked that it would be an unacceptable burden to have every directory understandable in every other language.

It was suggested that a language attribute could be included to indicate what languages are supported. This raises the need for OIDs for each language; national OIDs would not be appropriate since there are many more languages and dialects than countries.

The question was raised about how one would name multi-national organizations. Einar commented that NADF-123 document dealt with multi-state organizations in the U.S., and that an analogous scheme could be used for international organizations.

Kille commented that any structure could work, but was concerned with how well they would work, and the technical impact that they might have.

- Representing Network Information

  Mark Knopper asked if there were any questions or comments on the Network Infrastructure Schema document that was distributed some time ago.

  Kille commented that the flat space was not scalable, and that it should match hierarchical network number structure.

  Braden pointed out that there was no hierarchical structure in Internet network numbers; it is a flat space.

  Ruth Lang commented that it is recognized as an interim scheme to serve current needs.

  The question of how to name networks was raised.

  Einar suggested that network names were user friendly, and the NICs names would be bad choices. Mark pointed out that most networks do not have official names, and using an ad hoc name for the RDN was not suitable.

  Kille questioned whether numbers were more friendly than network names. He pointed out that network numbers were not technologically independent, and expressed concern that this could lead to inconsistent naming of networks.

  Hotz commented that the lack of network names was perhaps a more general problem that the Internet needed to address. A mechanism for mapping network numbers to names exists within the DNS, but is not frequently used.

Einar suggested that the network number be used as the RDN, and the name be included for searching. Kille suggested that the opposite would work just as well, and would make for more user-friendly names.

This is to be discussed further off-line.

- NADF-123

  This was discussed somewhat during the NADF report.

  Kille remarked that using the old structure (civilian infrastructure) could put entities in very unnatural places, making it difficult for those outside the structure to find things.

  Einar emphasized that everything in the U.S. has a registered name already in the current infrastructure, and that renaming/registration expressly for purposes of directory services would be unlikely.

  Einar pointed out that the underlying notion is that the right to register and obtain a name is different the the right to be listed in parts of the DIT. Organizations will naturally want to be listed in the places where others will look for them.

  Kille commented that he would like to see experience with this architecture before incorporating it into the naming guidelines.

DISI and OSI-DS

Knopper raised question regarding the roles of both groups. Kille responded that he sees DISI tackling operational issues, technical administration and issuing related technical specifications. OSI-DS deals with technical issues related to DS.

Meeting Administrivia

Steve Kille asked for comments about the usefulness of the videoconference meeting.

Bob Braden said that this videoconference was unusually bad. Usually a videoconference rates a 7 or 8 on a scale from 1 (email) to 10 (in person), this one only rated a 4.

Most other comments ranked the videoconference somewhere between email and in person. Opinions varied on its usefulness compared to a phone conference.

One of the UCL folks commented that traveling to a teleconference site was unsatisfying, particularly with the quality of this one. If one had to make the effort to travel, one might as well meet in person.

This raised the subject of U.S./European collaboration. Someone noted that IETF meetings are rather well attended by Europeans, but conferences and working group meetings in Europe do not receive a similar level of U.S. participation.

Braden pointed out that many U.S. participants traveled on government funds, and that the cost of European trips is, unfortunately, not viewed in a particularly favourable light.

Steve Kille will take comments about the videoconference into consideration when deciding if and when it would be appropriate again.

To wind up there was a discussion to see if people thought the meeting useful.

BBN: Not as good as a face to face meeting, but better than email.

RIACS: Might be more effective to choose a few items and discuss to focus on the issues.

ISI (Bob): Technical quality apalling - too much delay. Echo annoying. Sound poor. Scale: email – 1, in person – 10, then generally video – 7, but this time – 4 due to the delay and quality. on line terminal may help.

UCL (SEK): "interesting", some useful discussion. Presentations did not work. If material becomes too technical, the interchange did not work.

AOB

DUA on VMS – one will be publically available soon. It was developed in Spain.

Next Meeting

This will be at the IETF Meeting in Atlanta, in the week of 29th July.

I believe that the meeting was useful, although it did not fulfill all expectations. The long delays were a serious problem.

I found the meeting very stressful to chair, despite a very high level of cooperation from each site.

Getting Minutes taken at each site was a disaster. The major reason for the delay in producing these Minutes was the problem of merging four similar but different pieces of text. There should only be one minute-taker, perhaps supplemented by notes from each site. Comments on the videoconference were provided by a number of people, and this was useful.

Attendees

Paul Barker

Robert Braden	braden@isi.edu
Jim Cragie	
Nick Emery	
David Goodman	
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Steven Hotz	hotz@isi.edu
Erik Huizer	huizer@surfnet.nl
Ruth Lang	rlang@nisc.sri.com
Peter Mierswa	mierswa@smaug.enet.dec.com
Daniel Molinelli	moline@trw.com
Julian Onions	J.Onions@cs.nott.ac.uk
Colin Robbins	
Einar Stefferud	stef@ics.uci.edu
Steve Titcombe	
Peter Williams	
Charles Wolverton	ctw@aero.org
Russ Wright	wright@lbl.gov

CURRENT MEETING REPORT

Reported by Steve Hardcastle-Kille/UCL and Tim Howes/UMICH

OSIDS Minutes

Agenda

Introduction and Welcome

Chair Steve Hardcastle-Kille called the meeting to order after some furniture moving at 9:42am.

Previous Minutes

The Minutes from the February meeting at SRI and the May video conference were accepted without change.

Liaison Reports

- RARE WG3 (Erik Huizer - RARE)

  RARE WG3 has had one meeting since the last OSI-DS meeting. Erik reported the following:

  - PARADISE has now taken over operation of the COSINE X.500 pilot project.
  - The next RARE WG3 meeting will be in Zurich from 9/30 to 10/2. The meeting will include demos of lots of X.500 DUAs (for Macs, Unix, VMS, etc). Others are encouraged to attend and to contact Erik about bringing a DUA to demo. Erik also mentioned that there is the possibility of making funds available to someone from the US for the trip.

- ISO/CCITT (Steve Hardcastle-Kille in lieu of a representative)

  There will be one more meeting before the 1992 white book comes out. This meeting will be in Berlin.

- OIW (Russ Wright - LBL)

  Russ reported that at the OSI Implementors Workshop the following agreements had been reached:

  - Change the maximum APDU size to 256K (up from 32K),

  - Up the 6 line by 30 character postaladdress limit to 6 lines by 60 characters (see mtr's report below).

- NADF (Marshall Rose - unaffiliated)

  Marshall passed out a revised copy of the US Naming Scheme document now known as NADF 175 and reported the following:

  - NADF 175 will be submitted as an RFC soon.
  - NADF naming documents can be obtained from tymer@mcimail.com

  - NADF naming documents can be obtained from tymer@mcimail.com
  - NADF naming documents can be obtained from tymer@mcimail.com

  - The NADF expects to sponsor a pilot to test its agreements sometime in first quarter 1992.
  - The NADF was not pleased to hear about the OIW agreement about a postaladdress. Apparently, postaladdress limits are determined by a separate international agreement and cannot simply be changed to suit X.500. After some discussion, it was decided that OSI-DS should support the standard (i.e., 6 by 30) definition.

- FOX (Steve Hotz - ISI)

  Steve reported the following FOX activities:

  - Approximately 85K NIC WHOIS entries are now online in SRI's DSA.
  - SRI is also working on a lightweight application to access this information in X.500.
  - SRI and Merit will work together to provide replication of the WHOIS information.
  - There are now RFC and FYI document subtrees under o=Internet.
  - Merit has put a Site Contacts database online under o=Internet, which the NSFNet network operations center is using.
  - Merit has also begun to put NIC profile information online.

- PSI White Pages (Wengyik Yeong - PSI)

  Weng reported that the PSI WPP would be focusing on four areas in the near future:

  - Increasing reliability (through probing, e.g.)
  - Transition to NADF naming scheme
  - Transition to the OSI-DS DSA naming scheme
  - Upgrading all pilot members to ISODE/QUIPU 7.0

- PARADISE (Steve Hardcastle-Kille - UCL)

  Steve reported that the PARADISE project was doing the following:

  - Running the world-wide root DSA at ULCC (Giant Tortoise)
  - Running a DUA service at ULCC for European organizations
  - Producing a glossy report ¡how can people get a copy of this???¿
  - Helping other implementations, in particular Pizarro, join the pilot

- AARN (Steve Hardcastle-Kille in lieu of a representative)

  Steve reported that there is a funded directory service pilot started in the Australian Academic Research Network and that they will be sending a representative to future OSI-DS meetings. Graham Rees at the University of Queensland is the AARN contact person.

- NORDUNET (Geir Pedersen - NORDUNET)

  Geir reported that a directory services group is operating within NORDUNET. The group focuses on promotion of the directory within the Nordic countries.

- CNI (George Brett - CNI)

  Coalition for Networked Information:

  - Members include Association for Research Libraries, CAUSE (administrative computing) and EDUCOM (academic computing).

- 7 Working Groups: commercial publishing, non-commercial publishing, standards and architecture, management and professional education, K-20 education, directories and networked research centers (user services people).

- George Brett and Peggy Seiden are involved in the directory Working Group. They are interested in a top node or directory of directories, and bibliography of bibliographies. They currently have a flat file, and are now working on an X.500 implementation with Merit. Also looking at a WAIS database. The Merit work involves a schema to represent the library of congress enhanced MARC record.

- Documentation on CNI's activities includes an article in Cause & Effect, v14no2, summer 1991. The Minutes of their last meeting can be made available to the list by George. There is also a listserv list called cnidir-l.

- Art St. George of CNI is interested in K-12 use of networking.

Internet Resource Descriptions in X.500 Document (OSI-DS 17)

After much discussion, the general consensus on this document was that it needs to be rewritten with a more object-oriented approach. The feeling was that the object class for an Internet Resource should be broken up more in line with the standard's description of auxiliary and structural object classes.

Document Progression

Vint Cerf reported that all the technical RFCs were close to publication and would be out in "a matter of weeks". Because of the general importance to the Internet community of deploying a Directory Service, the overview and strategy documents should be merged and their scope widened. It was agreed that Steve H-K should revise the document, in association with appropriate IESG and IAB members. Review should be done via email, and that a separate subgroup not be formed.

The scope of the Working Group was also considered, and the group was favourable to moving the activity out of the OSI area and into the Applications area.

OID Assignment (OSI-DS 10, RFC 1239)

It was decided that the existing OIDs will stay. A small debate followed about whether the current naming authority (under UCL) was ok, or whether it should be changed to the IANA. It was agreed that OIDs are just numbers, so the assignment authority is less important than stability. Currently assigned OIDs should not be adjusted. This had already been agreed between Steve H-K and Jon Postel. New OIDs would probably be assigned under an IANA subtree, subject to consensus with non-Internet users of this schema. It was agreed that this document could and should progress rapidly to RFC.

Requirements Specification Document (OSI-DS 18)

This document had been drafted after an agreement at the videoconference. After some discussion, it was agreed that a document of this nature would be important in the future, but was premature. It should be put on ice for now.

More Resources Description Document Discussion

Mark Knopper of Merit said he was interested in putting NSAPs in the DIT and doing reverse lookups. Along with Chris Weider, he volunteered to write something up by mid-September on how to do this.

NIC Profile Information Document (OSI-DS 16)

After some discussion, it was agreed that Chris Weider should rewrite this document along the same object-oriented lines as discussed previously. There was also some discussion about how this information should be organized in the DIT, since some NICs are real organizations, others aren't, some are listed in other parts of the tree, others aren't. Steve H-K drew a diagram describing a structure that could accommodate both situations, and it was generally approved.

K-12 Schema Document

It was agreed that this document suffered from the same object-oriented concerns as the others and should be rewritten. Also, it was decided that a companion document addressing DIT structure for these objects should be produced. Chris Weider and Mark Knopper were elected for both tasks.

Network Infrastructure Information in X.500

This document is already being rewritten and was not discussed.

Pictures in the Directory

Russ Wright presented a brief overview of the problem (summary: the g3fax format is bad), and a potential new format that is better (JPEG). It was agreed that JPEG is a step forward, but more study is needed on the transition path, potential size limits, etc. Russ Wright, Peter Yee, Tim Howes, and Mark Smith (in absentia) volunteered to look into these issues.

Quality of Service (OSI-DS 15)

The QOS definitions were generally accepted, and the next step is to start making use of these attributes now that the syntax handlers are available in QUIPU 7.0. Russ Wright, Erik Huizer, and Tim Howes volunteered to try incorporating QOS into some DUAs to gain some experience with its use. All of the represented pilots agreed to install the appropriate attributes into their DITs. Both efforts were needed to make an effective test of the Internet Draft. The Internet Draft should not be submitted as an RFC until results of this piloting, and probable modifications to the Internet Draft, had been done.

NADF Naming Document

NADF 175 is now considered Stable and a Good Thing by NADF, and will be released as an RFC Real Soon Now.

Naming Guidelines Document (OSI-DS 12)

After some discussion about multi-national organization naming, it was agreed that this document should be progressed to RFC status pending ironing out of some minor issues which will be done via email.

AOB

Erik Huizer made a request for someone from the US to look into issues involving security/privacy laws in the US that might relate to X.500. This is something that has come up several times in Europe.

Next Meeting

Steve would like to have the next OSI-DS meeting around Interop.

Summary of Action Items

- George Brett: Find out how to get CNI documents and send this information to the osi-ds list.
- Mark Knopper, Chris Weider: Write a paper describing how to store NSAP information in the DIT (mid-September).
- Chris Weider: Rewrite Resource Description paper.
- Chris Weider: Rewrite NIC Profile paper.
- Mark Knopper, Chris Weider: Rewrite K-12 Schema paper.
- Mark Knopper, Chris Weider: Write a companion paper to the K-12 Schema paper describing the suggested DIT structure.
- Russ Wright, Peter Yee, Tim Howes: Experiment with JPEG photos in the Directory.
- Russ Wright, Erik Huizer, Tim Howes, others: Incorporate QOS into DUAs and pilot exercises.
- Steve Hardcastle-Kille/Paul Barker: Initiate an email discussion on the Naming

Guidelines document and progress it to an RFC when all concerns have been addressed.

- Steve Hardcastle-Kille: Schedule the next OSI-DS meeting at or around Interop.
- Steve Hardcastle-Kille: Produce new strategy/overview document.

Attendees

Claudio Allocchio	claudio.allocchio@cosine-gw.infn.it
William Biagi	bbiagi@cos.com
Peter Boos	peterb@bnr.ca
David Brent	brent@CDNnet.ca
George Brett	pls@psulias
Ross Callon	callon@bigfut.enet.dec.com
Vinton Cerf	vcerf@nri.reston.va.us
Richard Cherry	rcherry@novell.com
Chi Chu	chi@sparta.com
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Ralph Droms	droms@bucknell.edu
Urs Eppenberger	eppenberger@switch.ch
Arlene Getchell	getchell@nersc.gov
Jack Hahn	hahn@sura.net
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Susan Hares	skh@merit.edu
Ittai Hershman	ittai@nis.ans.net
Russ Hobby	rdhobby@ucdavis.edu
Steven Hotz	hotz@isi.edu
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Holly Knight	holly@apple.com
Mark Knopper	mak@merit.edu
Christopher Kolb	kolb@psi.com
Ruth Lang	rlang@nisc.sri.com
Anthony Lauck	lauck@tl.enet.dec.com
Louis Leon	osll@emuvm1.cc.emory.edu
Peter Liebscher	plieb@sura.net
Carl Malamud	carl@malamud.com
Bill Manning	bmanning@rice.edu
John McGuthry	mcguthry@gateway.mitre.org
Joe Peck	peck@ms1.pa.dec.com
Geir Pedersen	geir.pedersen@use.uio.no
Mel Pleasant	pleasant@hardees.rutgers.edu
Marshall Rose	mrose@dbc.mtview.ca.us

A. Minick Rushton	rushton@stsci.edu
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Keld Simonsen	keld.simonsen@dkuug.dk
Subu Subramanian	subu@qsun.att.com
Jesse Walker	walker@eider.enet@decpa.dec.com
Chris Weider	clw@merit.edu
Brian Wheeler	wheeler@mbunix.mitre.org
Linda Winkler	lwinkler@anl.gov
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov
Wengyik Yeong	yeongw@psi.com

3.4.4 OSI General (osigen)

Charter

Chair(s):

Robert Hagens, hagens@cs.wisc.edu
Ross Callon, callon@bigfut.enet.dec.com

Mailing Lists:

General Discussion: ietf-osi@cs.wisc.edu
To Subscribe: ietf-osi-request@cs.wisc.edu
Archive: janeb.cs.wisc.edu:/pub/archives/ietf-osi

Description of Working Group:

Help facilitate the incorporation of the OSI protocol suite into the Internet, to operate in parallel with the TCP/IP protocol suite. Facilitate the co-existence and interoperability of the TCP/IP and OSI protocol suites.

Goals and Milestones:

Done	Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet. Coordinate addressing format with GOSIP version 2 and possibly other groups.
TBD	Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley.
TBD	Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.
TBD	Determine what should be used short-term for (i) intra-domain routing; and (ii) inter-domain routing.
TBD	For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Determine if there are any outstanding issues here.
TBD	Review short-term issues involved in adding OSI gateways to the Internet. Preferably, this should allow OSI and/or dual gateways to be present by the time that Berkeley release 4.4 comes out.

Request For Comments:

RFC 1139	"Echo function for ISO 8473"

CURRENT MEETING REPORT

Chair(s): Rob Hagens/UWisc and Ross Callon/DEC

OSIGEN CO/CL Interworking Review Minutes

Report not submitted. See Area Report for summary.

Attendees

William Biagi	bbiagi@cos.com
Randy Butler	rbutler@ncsa.uiuc.edu
Ross Callon	callon@bigfut.enet.dec.com
Chi Chu	chi@sparta.com
Henry Clark	henryc@oar.net
Tom Easterday	tom@cic.net
Shari Galitzer	shari@gateway.mitre.org
Tony Genovese	genovese@es.net
Robert Hagens	hagens@cs.wisc.edu
Christian Huitema	Christian.Huitema@MIRSA.INRIA.FR
Kenneth Key	key@cs.utk.edu
Dalè Land	land@lanl.gov
Mike Okulski	ssds.com
Mark Sleeper	mws@sparta.com
Osamu Takada	takada@sdl.hitachi.co.jp
Preston Wilson	preston@i88.isc.com
Cathy Wittbrodt	cjw@nersc.gov

OSIGEN FTP-FTAM Gateway Specification Minutes

Report not submitted. See Area Report for summary.

Attendees

Robert Griffioen	rgriff@bnr.ca
Robert Hagens	hagens@cs.wisc.edu
Susan Hares	skh@merit.edu

3.4.5 OSI X.400 (osix400)

Charter

Chair(s):

Rob Hagens, `hagens@cs.wisc.edu`

Mailing Lists:

General Discussion: `ietf-osi-x400@cs.wisc.edu`
To Subscribe: `ietf-osi-x400-request@cs.wisc.edu`

Description of Working Group:

The IETF OSI X.400 Working Group is chartered to identify and provide solutions for problems encountered when operating X.400 in a dual protocol internet. This Charter includes pure X.400 operational issues as well as X.400 <-> RFC 822 gateway (ala RFC 987) issues.

Goals and Milestones:

Jul 1990 Develop a scheme to alleviate the need for static RFC 987 mapping tables.

CURRENT MEETING REPORT

Chair(s): Rob Hagens

OSI X.400

The OSI X.400 Working Group has been merged with the X.400 Operations Working Group. For additional information, refer to the Minutes of the X.400 Operations Working Group.

3.4.6 Office Document Architecture (oda)

Charter

Chair(s):

Peter Kirstein, kirstein@cs.ucl.ac.uk

Mailing Lists:

General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk
Archive:

Description of Working Group:

The ODA Working Group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g., Spreadsheets. The Working Group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liason with the SMTP and X.400 Working Groups.

This Working Group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Done	Inaugural meeting.
Done	Produce a paper stating what ODA standards or profiles still need completing.
Jul 1991	Produce paper on how both SMTP and X.400 message systems should be supported.
Done	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Dec 1991	Provide first feedback on the ODA Pilot.
Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.

CURRENT MEETING REPORT

Reported by Peter Kirstein/UCL

ODA Minutes

Current Status of Implementations

As a background for the discussions on Pilots, the current status of implementations was reviewed. The following were known to be available, potentially, to the IETF Office Document Architecture Working Group:

PROVIDER	PACKAGE
BBN/UCL	SLATE/ODA
Bull	WORD for WINDOWS/ODA
DEC	DECWRITE/ODA
UPC/ICL	WORDPERFECT/ODA
XEROX	VIEWPOINT/ODA

The status of each is discussed below:

BBN/UCL-SLATE/ODA

There has been a Release of v1.1 of the BBN SLATE/UCL ODA software; it converts between SLATE v1.2 and ODA/ODIF Q112. This is freely available to anyone who has a license for BBN SLATE 1.2. The software is made available currently on SPARCstations, but is believed to be easily portable to IBM RISC 6000 machines and DEC Ultrix workstations. There is documentation for the system on the normal ietf-osi-oda infoserver. At present the system operates with the UCL PP message (v5.0 or later versions), and thus can operate over SMTP (with UUencode) or X.400; later versions will work with the extensions to SMTP proposed in the SMTP IETF working party. It is possible to interoperate with any other SMTP mail systems which does UUencoding.

It has been agreed with BBN, that they will provide for the IETF Pilot 250 copies of SLATE v1.2, and will maintain it with later releases. It is restricted to "academic and research institutes only"; others must purchase the SLATE from BBN. The software will be updated as later releases of SLATE become available. The whole documentation will be provided by UCL - who will include the BBN SLATE documentation. The BBN portion of the software will be provided to US participants by a "Shrink-wrapped Licence"; non-US sites will have to sign a BBN license supplied by UCL. In both cases, UCL will keep a register of copies supplied, and must furnish that to BBN. UCL will make a small handling charge for the distribution.

Bull-Word for Windows/ODA

There will be a version of this software made available to the Working Group; it will run on a DOS PC, and must be integrated by the using site with a mail system. The Bull software is designed for conversion between RTF and ODA Q112, but they use it only with WORD for WINDOWS. To date there are still some slight problems with the software, but an improved version is expected by the end of August 1. If this is up to expectations, UCL expect that it should be available to the Working Group during the 4th quarter of 1991.

The software requires a PC/AT with PC-DOS or MS-DOS v 3.10 or above - with at least 1 MB of EMS, hard disc and floppy. It requires also the editor, i.e., WORD for Windows v1 or other editor supporting RTF. The programs include Q112 <--> RTF convertors, a formatter, filter, and a browser. It also includes filters and test documents. There will be appropriate documentation from Bull.

The license agreement is for use on a single DOS system for R and D. One should report on the usage. It for Universities and Public Research laboratories for evaluation, research and demonstration. It is initially until June 1992. UCL will distribute the software and documentation, but users will have to sign a license agreement with Bull.

DEC This package is regarded as a Gateway product between their CDA products and OSI. The VMS release was made in April, the ULTRIX release is on Extended Field Trial (EFT). They run on all current DEC machines. Mitre has tested the VMS release, and found some problems with one of the directions of conversion. UCL received the EFT of the ULTRIX version on the day of the IETF meeting. While there have been no discussions with DEC on how it will be made available to the IETF, this is not usually a problem for educational or research organisations.

UPC/ICL The University Polytechnic of Catalonia (UPC) has offered a version of their convertor between WORDPERFECT 5.0 and ODA Q112 for evaluation by UCL. It will run under DOS and UNIX. The DOS version will leave it to the users to have the ODA files included under the users' favourite mailsystem; the UNIX version will be available both for SUN-3 and SUN-4 systems. UCL has tested one version, and expect a further version early in August. The version tested by UCL would not be suitable for release; it is expected that a suitable one will become available during the fourth quarter of 1991.

The RTF-ODIF convertor comes from ICL, so that availability conditions must be agreed both with ICL and UPC. The exact terms and conditions for availability are not yet clear, but are expected to be similar to those of Bull. It is expected that UCL will distribute the software for UPC.

XEROX They have confirmed the availability of their VIEWPOINT/ODA software; it runs under UNIX on SPARCstations. UCL has not been given a copy yet, and so must still evaluate it. While XEROX does have an X.400 which is integrated with their system, this requires the availability of XEROX hardware; none of the Working Group plans to test this integrated system. UCL will give further information on this package when it has evaluated it.

The XEROX software is a standard product. There has been no discussion yet on the terms under it will be made available to the Working Group. UCL plans to initiate such a discussion after the evaluation.

UCL has made available an Autonomous Active Mailbox for testing purposes. This will allow people to send test documents by X.400 or SMTP, and to Store, Retrieve or Modify them. A brief description is given in documentation available on the infoserver at UCL.

Interest was keen also in MAC software, and the Chair agreed to contact Apple - since it was believed they also had software in some relevant state. It was agreed that in view of the imminence of so much of the software, it was important to update this list regularly. The Chair would provide an updated status both at the end of September, and for the next meeting at the next IETF.

Discussion on Interests of Working Group Members in the Available Software

From the discussion, it became clear that for any serious use by participants, it would be necessary to have several copies of any software in each user group. This was consistent with the current BBN policy in the way they licensed three copies of SLATE, but was possibly at variance with the individual licenses assumed by Bull. Moreover, some of the Working Group participants said that even three was quite inadequate for their purposes. After some discussion, it was agreed that we should go back to those licensing the software, and get them to agree to provide three licenses to a group as part of the IETF Working Group activity. If individual groups felt they needed more copies, then they should approach the software provider, and discuss the terms under which additional copies were provided.

There was a discussion for what large-scale Pilots the facilities would be particularly useful. Various proposals were made such as specific Working Groups of the IETF, access to certain large databases, and alternate representation for RFCs. It was agreed that prior to any serious Pilots, it was necessary for Working Group members to get wider experience of the current software available - or soon to become available. This phase of activity should be called "Experiments" rather than "Pilots". It was hoped that this stage could be completed by the end of the year. The initial experimentation would be inside the Working Group, and the following volunteered their participation:

First Surname Organisation to be Investigated

Ned Freed	Innosoft	DEC VMS/PMDF/DECWRITE
Peter Kirstein	UCL	SLATE/ODA, WORD, WORDPERFECT
Peter Kirstein	UCL	DEC-Ultrix/DECWRITE, XEROX
Jim Knowles	NASA-AMES	DEC-Ultrix/DECWRITE, Sun 3,4, DOS, MAC
David Lippke	U of Texas	SLATE and WORDPERFECT
Carl Malamud	CONS	DOS-based systems
Brien Wheeler	MITRE	DEC-Ultrix/DECWRITE, VMS, Sun 3,4, DOS, MAC
Greg Vaudreuil	CNRI	SLATE and WORDPERFECT (later)

The Chair will try to organise the availability of the Bull and UPC/ICL software for Knowles, Malamud and Wheeler; Kirstein will, of course, be getting it in any case. Kirstein will also discuss with DEC how US participation in the IETF should be handled as regards the DEC software. These original users would try to get experience prior to the next meeting. Freed would plan to organise some integration of the DEC software with his PMDF Mail product. Wheeler agreed to provide some documentation on their experience with the DEC product.

Compatability with Mail Systems

The ODIF is not immediately compatible with SMTP, because of the need of handling binary contents. The v1.1 of the SLATE/UCL software has provision for UUencoding the data for use with SMTP - and even with X.400 if the implementations required it. It was agreed that the first experiments would be made with this form of encoding. It was noted that the Internet Mail Extensions Working Group was adopting a different encoding of binary. It was agreed that in the next version of the software to be used by this Working Group, the coding agreed in the Internet Mail Extensions Working Group would be used. Vaudreuil agreed to put some appropriate parameters for ODA into the Internet Mail Extensions next version of the document, including Profile (currently only Q112, but eventually others also), Originating Site, Version Number, and possibly receiving site. Details would be discussed by e-mail. It was agreed that the Chair should discuss with the X.400 Working Groups both compatibility with their formats, and possible interest in usage of ODA in their Pilots.

Documents Available on the Infoserver

A number of documents are currently available on the Infoserver at infoserver@cs.ucl.ac.uk. The documents are accessed by standard message systems, giving a message body of the form:request: ietf-osi-oda topic: xxxx where xxxx is the name of the document required. The list of documents currently in the collection is listed in a document called INDEX. The documents are currently all in text form. Some will be made available in ODA/ODIF format shortly.

Next Meeting

It was proposed to hold the next meeting prior at the Santa Fe IETF meeting, during the week of November 18-22.

Attendees

Jill Foster	jill.foster@newcastle.ac.uk
Ned Freed	ned@innosoft.com
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
P. Allen Jensen	allen@audfax.audiofax.com
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Mark Leon	leon@nsipo.arc.nasa.gov
David Lippke	lippke@utdallas.edu
Carl Malamud	carl@malamud.com
John McGuthry	mcguthry@gateway.mitre.org
Geir Pedersen	geir.pedersen@use.uio.no
John Scudder	jgs@merit.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Brian Wheeler	wheeler@mbunix.mitre.org

3.4.7 X.400 Operations (x400ops)

Charter

Chair(s):

Alf Hansen, Alf.Hansen@delab.sintef.no

Mailing Lists:

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

Description of Working Group:

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

Goals and Milestones:

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

CURRENT MEETING REPORT

Reported by Kevin Jordan/CDC

X400OPS Minutes

Welcome.

There were no additional comments against the St. Louis meeting Minutes.

The IETF X.400 Operations Working Group.

Alf Hansen and Rob Hagens are now Co-Chairs of the Working Group. Alf is returning home to Norway. The old X.400 Working Group has been merged with the X.400 Operations Working Group. The most significant work item completed by the old X.400 Working Group was an RFC describing how to use DNS to store RFC1148 mapping information. The status of this RFC is that it is awaiting proof of concept through implementation.

Because the two X.400 Working Groups have merged, the Working Group Charter will be updated to add a new goal: the Working Group will attempt to drive X.400 deployment in the Internet.

The X.400 Operational Requirements RFC was originally scheduled to be available for comment in July. This schedule needs to be revised because a lot of work is left to be done (especially considering the comments and resolutions discussed in Atlanta).

The following questions were asked: "Is XNREN a U.S. or an international PRMD? How would an organization outside of the U.S. join?"

Alf attempted to provide an answer by indicating that the IETF should find a way to register XNREN as a PRMD in each country. It is not clear exactly how this would be accomplished, but extensive cooperation from the international Internet community is required.

Status of the document, "An X.400 Internet Strategy".

Work on the document continues. It is slightly behind schedule.

Roundtable presentation of current X.400 service status.

At this point, the Working Group members who are currently operating X.400 services described the status of those services:

SURFNet — (Netherlands) The SURFNet operations team is currently working to improve the robustness of the service by providing live backups for key service elements, i.e., redundant WEP's and RFC987 gateways.

An international agreement is needed on how to define backup WEP's with associated priorities (like the MX concept in SMTP and DNS) so that MTA's can try alternate backup connections. Note: RARE WG1 has begun work on this concept.

SURFnet currently serves about 800 active X.400 users, and the number of users is growing rapidly.

X.400 implementations for Mac's and PC's are being evaluated, as are X.400 gateway products for Mac/PC LANS (e.g. cc:Mail, Banyan).

SURFNet Observation: Most currently available X.400 user interfaces are still quite primitive.

COSINE

Cooperation for OSI Networking in Europe. COSINE is a program funded by a number of European Governments (basically the European Community plus European Free Trade Association countries) plus the Commission of European Communities. Broadly, the mission is to provide OSI based services for the European research community. The prime contractor entrusted to fulfill this mission is RARE.

COSINE includes:

RARE Reseaux Associe pour la Recherche Europeenne

EEMA European Electronic Mail Association EEMA is an association whose membership is comprised of a number of European organizations, some very large (almost exclusively non R&D based). They come together to discuss issues related to electronic messaging in Europe. RARE/COSINE decided to become a member of EEMA, with a view to feed back the experiences learned by the RARE/COSINE MHS services into industry, (i.e., act as an experience pool), To make the views of the COSINE user community felt in this forum.

Y-Net OSI Services for ESPRIT researchers Y-NET is a project with its primary aim being to provide OSI based services to European Community SMEs (Small and Medium sized Enterprises) involved in the ESPRIT program. COSINE MHS is mandated to coordinate with Y-NET. An aim of COSINE MHS is to provide a seamless service between the Y-NET and COSINE MHS user communities.

EurOpen

COSINE-MHS

is a project which was chartered to drive deployment of X.400 in the European research community. Transport service stacks include:

TP0/CONS/X.25/LAPB, TP0/CONS/X.25/LLC2, TP0/RFC1006/TCP, TP4/CLNS.

X.400 '84 is used universally within the COSINE-MHS community. Some organizations are experimenting with X.400 '88, but there is no wide spread use of '88 yet. The European public service providers (the PTT's) offer '84 service only.

The COSINE-MHS project is currently comprised of between 20 and 25 WEP's. Connectivity between WEP's is not universal. Even with this relatively small number of WEP's, the amount of static configuration information which must be maintained and coordinated is approaching an unmanageable level. There is a very urgent need for dynamic configuration management via X.500 directory services and/or DNS.

Some countries consider COSINE-MHS to be an operational service, and some countries consider it to be a pilot service. Consequently, varying degrees of support and administration are provided.

A universal gateway service, COSINE-GW, is being implemented in Trieste, Italy. This gateway will provide connectivity between practically all commonly used electronic mail networks including: X.400, RFC822, BITNet/EARN, HEPNET (Mail-11), and SPAN (Mail-11). Connectivity with XNREN is also being implemented.

SWITCH

(Switzerland) SWITCH has one main WEP which provides access to the Swiss research community. This main WEP has ADMD connectivity. SWITCH serves about 8000 real end users. About 50 academic and research organizations are connected. Five commercial organizations are connected. Commercial organizations must connect as independent PRMD's.

UK

Two main X.400 services operate within the UK academic/research community: EAN and MHS-Relay (PP-based). Connectivity with 3 ADMD's is provided. Most UK sites are operating X.400 '84 services, but 3 sites are experimenting with '88 internally.

GARR

(Italy) GARR is registered as an official Italian ADMD, but it primarily services the academic/research community and is not a public service provider. GARR is connected with 2 public service ADMD's in Italy. GARR's potential user community numbers between 10,000 and 100,000 people.

GARR provides one principal access point (WEP) to COSINE. Backup WEP's are planned, pending international agreements on how to define and configure prioritized alternative MTA's for X.400 destinations.

X.400 '88 deployment is being considered, but GARR currently has no time table in place for deployment.

ARC

(NASA-Ames Research Center) The primary WEP for ARC was recently transferred from a microVAX to a SUN. While the transfer was in progress, connectivity to ARC was lost. ARC is working on connectivity to SPRINT. A fax gateway is planned. ARC is considered an experimental rather than an operational X.400 mail service.

CDC

Control Data operates its on PRMD named CDC. Control Data has a connection to XNREN via Internet and is also a subscriber to ADMD ATTMail. CDC is connected to ATTMail via AT&T's public X.25 network named Accunet.

Internally, CDC operates an X.400 network which currently interconnects 3 principal corporate locations: Arden Hills, Minnesota; Bloomington, Minnesota; and Santa Clara, California. It is estimated that well over 1000 X.400 messages per day are exchanged between and within these three locations. The number of users served is in the hundreds.

CDC has produced two main X.400 implementations. These are marketed as Control Data products, and they are also used very heavily within the company. One of the implementations, named MHS/4000, runs on the Control Data 4000 series of computer systems (based on the MIPS RISC chipset and running CDC's variant of UNIX named EP/IX). The other implementation, named Mail/VE, runs on the Control Data CYBER 180 series of mainframe computer systems under the NOS/VE operating system.

Several of CDC's customers in Europe (particularly Germany) are taking advantage of CDC's connection with XNREN. They are able to exchange true X.400 mail between their sites and Customer Support analysts at CDC in Minnesota. One of the customers even sends periodic X.400 "pings" from his X.400 mailbox in Germany to an autoforwarding mailbox at CDC in Minnesota. The autoforwarding mailbox forwards mail back to his mailbox in Germany. This allows him to verify that the international X.400 network is fully operational (between Germany and Minnesota at least).

CDC has implemented an X.400-based fax gateway and is currently using it internally. This gateway will be released as a CDC product in the Fall. The gateway can accept messages containing IA5-Text, Unidentified (aka Bilateral), and G3-Fax body parts. IA5-Text body parts can contain plain text, PostScript, uuencoded digital imagery, and digital imagery encoded using Macintosh BinHex format. TIFF, GIF, PICT, MacPaint, XBM, XWD, PBM, PPM, PGM, and Sun Raster digital image formats are recognized. Unidentified type body parts may also contain any of these formats (without having to be uuencoded or BinHex encoded). The gateway provides access controls and accounting, it honors deferred delivery requests, and it generates X.400 delivery reports. It also allows the network administrator to design customized cover sheets. It can receive as well as send faxes, and, of course, it can serve non-X.400 users across an RFC987 gateway.

ESNet

ESNet is implementing X.400 connectivity with XNREN. Internally, ESNet is providing pilot services to energy research labs. As ESNet must meet U.S. GOSIP requirements, the internal ESNet OSI backbone will be based on TP4/CLNS. The potential ESNet user community is about 4500 people.

CDNNet

(Canada) CDNNet is topologically organized as a star. A WEP and RFC987 gateway are located at the center of the star. About 40 organizations, Canada-wide, are connected to CDNNet. CDNNet has connections with XNREN and Envoy. CDNNet is considering becoming an ADMD. CDNNet participates in support of EAN. The number of X.400 end users served by CDNNet numbers in the thousands.

MITRE

MITRE's X.400 service is '84 based. MITRE's X.400 network has two main relays. One is located in Bedford, Massachusetts, and the other is located in Washington, D.C. Routing is hierarchical and concentrated at the two main relays. Departmental MTA's are configured with simple default routes to the central relays.

MITRE's X.400 service is confined within the corporation. MITRE does not yet exchange X.400 mail with other organizations because MITRE has not yet integrated the OSI protocol suite into its security perimeter mechanisms.

The MITRE X.400 service is currently operating as a mail backbone transport service only. X.400 mail is not yet exchanged with end users directly, i.e. X.400 user agents have not been deployed.

X.500 (Quipu) is being used for address lookup and distribution list expansion. The principal user agent in use is MH 6.7 with the enhancements that support X.500.

MITRE's current view of OSI: "It's tough to show the added value of OSI at this time." OSI products are immature. GOSIP is incomplete. It requires only IA5-Text support in X.400 body parts, it does not require X.500, and it does require any sort of network management support. The standards are incomplete. For example, the standards do not specify standard textual representations for host names or email addresses.

XNREN

X.400 traffic passing through the main XNREN relay steadily increases, but it is still relatively light. In June, between 7000 and 8000 X.400 messages were processed. Most traffic originates from X.400 and is destined for SMTP users.

PP (release status) Steve Hardcastle-Kille provided the following information about forthcoming releases of PP: A beta release of PP was distributed very recently. PP 6.0 is currently scheduled for general release in September or October. PP 6.0 will include a fax channel. At the present time, the fax channel works in the outbound direction only, but inbound should be working soon. In addition, the fax channel is currently implemented to interface with a fax modem which is not currently sold in the U.S. A Mail-11 channel will become available. 88-¿84 downgrading will be supported per Steve's draft RFC. The Domain table has been revised to look like the O/R table. An implementation of Message Store and an X Window User Agent will become available. The X Window UA will probably be released with PP 6.0, but its quality will be VERY beta in that release. It will be suitable for experimentation, but not for end user usage.

The PP API will be published. Note: this API will not be compatible with the X/Open API for X.400, and there are currently no plans to implement an X/Open conformant API for PP.

Review of draft RFC, "Requirements for X.400 PRMD's Operating in the Internet."

The Working Group went through the draft RFC section by section, discussing issues and resolving them. One major outcome of the dialogue is that the focus of the document has changed from being U.S.-centric to being international in scope.

- Status of this Memo: It was pointed out that the format of this section may not follow the approved wording format for Internet RFC's.

- Introduction: It was suggested that this section does not really introduce the reason for the existence of the document. It dives into technical details too quickly. This section should provide answers to the following questions:
  - What is the rationale for deploying X.400 on the Internet?
  - How does X.400 deployment relate to the forthcoming enhancements to SMTP?
  - Why is this document being written?

  One justification for deploying X.400 on the Internet is that there are a number of Internet-connected organizations which are beginning to operate internal X.400 services (in compliance with U.S. GOSIP), and it should be possible to use the Internet to interconnect these services.

  Among other things, the document should provide a boilerplate which describes how to connect an organization to the Internet X.400 network.

  After considerable discussion, the following conclusions were drawn:

- We probably need to produce a separate document which clearly lays out the rationale for deploying X.400 on the Internet.

- The document needs to be expanded such that it accommodates the international R&D community. In particular, the document must accommodate both of the XNREN and RARE/COSINE communities.

- Our basic goal is to foster an international X.400 service for the Internet.

Profiles

The intent of the profile section was to document the upper layer X.400 profiles which must be supported by participating organizations. It was agreed that the document should merely refer to other documents which define standardized inter/national profiles because, in practice, existing X.400 implementations are interoperable, and they conform to standardized inter/national profiles.

- Management Domains: Given that the document will be revised to accommodate international requirements, and given that a variety of management domain schemes are already in use, this section should describe existing practices. In particular, it should describe the existing variety of PRMD's and ADMD's and point out that management domain interconnection requirements will vary from one country to the next.

Lower Layer Stack Incompatibilities

Discussion of this section prompted a long and lively dialogue concerning what the definition of "Internet" truly is, whether it is still necessary to retain the I-WEP concept, and whether it should be a requirement that all I-WEP's and/or WEP's be capable of direct interconnection. In the process of resolving these issues, it was pointed out that the IAB has revised the definition of "Internet" as follows:

> Internet is a multiprotocol community which shares a common name service.

Given this definition, the Working Group produced the following proposals for the definition of "Internet X.400 Service" or "Internet X.400 Community". The proposals were produced in the order indicated below, each was discussed thoroughly, and then a revised proposal (based on the discussion) was generated.

- p1: The Internet X.400 Community consists of X.400 communities which are connected with X.400 to the international R&D X.400 community without the assistance of a third party intermediate ADMD.

- p2: The X.400 Internet includes all sites where you can send an X.400 e-mail and get a non/delivery report in response.

- p3: An organization is a member of the X.400 Internet Community if it meets the connectivity requirements defined in this RFC.

These proposals made it clear that our basic goal is to use the Internet as a vehicle for maximizing X.400 connectivity. Given that agreement was reached on this goal, it became obvious that we should allow organizations to connect to the X.400 Internet using any of the following three lower layer stacks: TP0/RFC1006/TCP, TP0/CONS, TP4/CLNS

Furthermore, it should not be a requirement that every MTA or PRMD directly support all three stacks, but if a particular stack is not directly supported by a PRMD, the PRMD will need to make bilateral agreements with other PRMD(s) in order to assure that connectivity from all stacks is available.

The final agreed definition of "Internet X.400 Sevice" became:

> The Internet X.400 Service includes all organizations meeting the international requirements described in RFCxxxx.

where RFCxxxx is an RFC which describes requirements for connecting to the international Internet X.400 network. As mentioned above, the lower layer protocol stacks supported by the international Internet X.400 network are:

> TP0/RFC1006/TCP, TP0/CONS, TP4/CLNS

Connection requirements include:

- An organization must support at least one of the above stacks.
- An organization must insure that it is reachable from all stacks. An organization can achieve universal reachability by:

  * Directly supporting all stacks

  * Negotiating bilateral agreements with other organizations which share a common stack and which either:

    · Support a stack not common to both organizations, or

    · Are willing to relay mail to organizations which do support other stack(s)

Editorial note: The TP0/CONS stack should probably be subdivided into the following two stacks: TP0/CONS/X.25/LLC2, TP0/CONS/X.25/LAPB. These two stacks qualify as TP0/CONS, but their link layer solutions prevent them from being interoperable, so they are effectively as different as TP0 and TP4.

Having made the resolutions described above, it was agreed that all references to the term I-WEP should be changed to WEP. It was agreed that the I-WEP concept is no longer necessary.

In conjunction with the decisions made above, new proposals were made for the structure of routing tables maintained for the X.400 network. Two tables, an MTA table and a Domain table, will be defined. The MTA table will define the names of well known MTA's (WEP's) and their associated connection data including selector values, NSAP addresses, supported protocol stacks, and supported X.400 protocol version(s) (i.e., 1984, 1988, 1992, etc.).

Each entry in the proposed Domain table will consist of an X.400 address, followed by a list of MTA's which are willing to accept mail for the address or provide a relay service for it. Each MTA name will be associated with a priority value. Collectively, the list of MTA names make the address reachable from all protocol stacks. In addition, the list may provide redundant paths to the address, so in this case, the priority value indicates the preferred path, or the preferred order in which alternative routes should be tried. The format of a Domain table entry might look like:

```
C=CH; ADMD=ARCOM; PRMD=SWITCH
  PRIO=1, MTANAME=switch.ch
  PRIO=2, MTANAME=relay.dbp.de
  PRIO=3, MTANAME=mhs-relay.cs.wisc.edu
```

Architectural Principles
This section will be removed as it is no longer required that all WEP's be directly interconnected.

- Description of PRMD policies: All references to PRMD will be changed to MD (Management Domain). This will allow ADMD's to operate within the Internet X.400 Service.

X.400 address registration

This section will be updated such that it supports the specification of numeric country codes, ADMD names, and PRMD identifiers. Support of numeric identifiers is required by the X.400 standards and implementors agreements.

The description of "unique address" will be softened. The basic requirement is that all originator addresses transmitted into the Internet X.400 Service must be universally "repliable". In support of this requirement, the document will recommend that users align their addresses with exactly one ADMD name in cases where they have a choice of ADMD names.

It was pointed out that the requirement that organization names be nationally unique is not justified. Organization names must be unique within the context of the subscribed PRMD or ADMD, but they need not be nationally unique. The document will be updated accordingly.

The document will include a strong recommendation about the syntax of PRMD, O, and OU names. Specifically, such names should consist of letters, digits, and hyphens only. Also, a hyphen should neither occur as the first nor the last character of a name, nor should a name begin with a digit.

The document needs to contain information about officially supported DDA's. In particular, the supported DDA's should be listed along with their required syntaxes and semantics. The document must indicate the DDA's for which support is mandatory.

The document should reference the forthcoming RFC which describes '88-¿'84 downgrading, and it should indicate that support of that RFC is mandatory for organizations connected to the Internet X.400 Service.

- An organization with no defined X.400 address space

  This section will be reworded such that it clarifies the fact that the address of an RFC987 gateway need not be precisely: C=US; ADMD= ; PRMD=Internet

  In particular, the country name C=US is not mandatory. Each country is free to choose its own well known RFC987 gateway address. For example: C=CH; ADMD= ; PRMD=Internet

General comments/issues:

- The document should mention that issues concerning X.400 '88 are, in general, left for further study. This leaves a hook for future work.

- The document should reference a separate RFC which will describe the details of routing. Section 4.3 of the current draft will be moved into the routing RFC.

- The "6A" concept described in the current draft needs to be revised such that it reflects the new international flavor of the document.

- X.400 network coordination and administration will need to be distributed between continents. The X.400 Working Group, in concert with RARE/COSINE, will need to document administrative responsibilities and how they are divided between countries.

- We must determine how the commercial ADMD service providers relate to the

Internet X.400 Service.

Use of distributed databases for routing/mapping purposes:

Claudio Allocchio presented his experiences in experimenting with DNS as a solution for managing RFC987 routing/mapping tables. First, Claudio experimented with using PTR records for storing and managing mapping tables. His conclusion is that this is a reasonable short-term solution (pending a better X.500-based solution).

Next, Claudio experimented with using MX records for managing X.400 routing information. Again, he concluded that this is a reasonable short-term solution.

Claudio is planning to implement and make generally available a portable tool (written in C) which will allow an administrator to create the standard RARE/COSINE routing/mapping tables from information stored in DNS.

Kevin Jordan reminded the Working Group about the description he distributed after the previous IETF meeting of CDC's use of X.500 directory services for managing X.400 routing/mapping information. Kevin agreed to update this information and redistribute it to the Working Group as a formal proposal.

X.400 84/88 downgrading:

Steve Hardcastle-Kille presented his draft RFC on '88-¿'84 downgrading. He accepted comments from the Working Group and will make some minor changes to the document.

Future issues:

No additional future issues were discussed.

Summary of conclusions and actions:

R. Hagens, A. Hansen. The RFC authors will revise the document in accordance with the comments and conclusions generated at this meeting. A new draft will be distributed prior to the next IETF meeting, no later than November 11.

K. Jordan: Kevin will update his previous white paper which described CDC's usage of X.500 directory services in support of X.400 routing/mapping. He will distribute the updated paper to the Working Group as a formal proposal.

Kevin will also distribute a proposal for mapping X.400 O/R addresses to X.500 distinguished names. This mapping will allow X.500-based routing/mapping information to be distributed easily across the Internet, in a fashion similar to the way in which DNS information is distributed.

C. Allocchio, E. Huizer, U. Eppenberger: This team will distribute a proposal for using DNS and/or FTP-based services for managing X.400 routing/mapping information.

S. Hardcastle-Kille: Steve will update the '88-¿'84 downgrading RFC and work with EWOS to make support of DD.COMMON well defined and mandatory.

P. Yee: Peter will do some research into North American groups such as EMA and NADF. He will discover what they are currently doing and recommend a level of involvement for XNREN and/or the X.400 Working Group.

Future meetings:

The next general IETF meeting is scheduled for November 18 - 22 in Santa Fe, New Mexico. The X.400 Operations Working Group will meet on Wednesday and Thursday (November 23 and 24). Also, if there is sufficient interest, a BOF meeting may be organized.

Attendees

Claudio Allocchio	claudio.allocchio@cosine-gw.infn.it
William Biagi	bbiagi@cos.com
Peter Boos	peterb@bnr.ca
David Brent	brent@CDNnet.ca
Vinton Cerf	vcerf@nri.reston.va.us
Henry Clark	henryc@oar.net
Curtis Cox	ccox@wnyose.nctsw.navy.mil
Urs Eppenberger	eppen@v
Tony Genovese	genovese@es.net
Arlene Getchell	getchell@nersc.gov
Robert Hagens	hagens@cs.wisc.edu
Alf Hansen	Alf.Hansen@delab.sintef.no
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Phillip Hasse	phasse@honchuca-emh8.army.mil
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
P. Allen Jensen	allen@audfax.audiofax.com
Kevin Jordan	kej@udev.cdc.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Walter Lazear	lazear@gateway.mitre.org
Jack Liu	liu@koala.enet.dec.com
Carl Malamud	carl@malamud.com
Joseph Malcom	jmalcom@sura.net
John McGuthry	mcguthry@gateway.mitre.org
Harvey Shapiro	shapiro@wnyose.nctsw.navy.mil
Keld Simonsen	keld.simonsen@dkuug.dk
Linda Winkler	lwinkler@anl.gov
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov

3.5 Operational Requirements Area

Director(s):

- Susan Estrada: Estradas@cerf.net
- Phill Gross: pgross@nri.reston.va.us
- Bernhard Stockman: boss@sunet.se

Area Summary reported by Susan Estrada/CERFnet

Bernhard Stockman, Susan Estrada, and Phill Gross share the responsibility for the Operational Requirements Area Directorship.

There are four active Working Groups and one Directorate in the Operational Requirements Area. The Operational Requirements Area Directorate (ORAD), met this week for the first time. ORAD is not an official Working Group, but the steering group for ORA. The Working Groups active are user connectivity, operational statistics, network joint management, and benchmarking.

The results of the ORAD meeting were really heartening. There is support for development of operational standards in the Internet. The ORAD defined and agreed to some goals for the Operational Requirements Area. An open mailing list was set up called orad@sdsc.edu. To add your name to the list, send a note to **orad-request@sdsc.edu.**

There will be discussions on the ORAD mailing list over the next couple of months to help define what the Operational Requirements Area's priorities should be for this year.

Network Joint Management:

Network Joint Management is Chaired by Gene Hastings. NJM discussed network maps and the difficulties in keeping a central repository up to date. It was decided to change the central repositories to have pointers to where maps actually reside (usually at NICs) instead of actually having the maps in the central repositories

Third party trace route which apparently has disappeared from some applications. One application is available from LBL and can be retrieved via FTP from NIC.NEAR.NET. Third party trace route should be standard in LTRIX 4.2

As the Internet is increasing in size, statistics are becoming too numerous to collect and analyze. NJM has identified a need for either developing a good sampling technique or providing a new platform to collect the data. This is a good action item for the ORAD.

There is a potential fallback problem in the T3 network. What happens when a T3 breaks and the network falls back to T1 Nearnet has a traffic selection mechanism which prioritizes certain traffic and kills the rest.

NJM also observed that there is an increasing number of low speed users on the net. Low speed users accessing the Internet via dial up services can generate a unique situation. There is a churn factor as these users appear and then disappear from the network. It is an area which should be studied.

NJM is developing a template for standardized up-down reports.

User Connectivity Problems:

UCP is Chaired by Dan Long. UCP did not meet at this session of the IETF. UCP has written a specification and is starting implementations of trouble ticket sharing.

Benchmarking:

Benchmarking is Chaired by Scott Bradner.

For results see Minutes.

Operational Statistics:

Opstats is jointly Chaired by Bernhard Stockman and Phill Gross.

During prior IETFs, the Opstats group developed a model to report operations statistics using metrics, exchange of information, and definition of presentation formats. At this IETF, Opstats discussed the gathering of metrics from remote sites.

The conclusion was made that those kind of metrics should be at least mentioned as a kind of general abstraction level of important metrics to be gathered sooner or later. For these reasons, we have put together four strawmen papers to be used as the base for the continuation of this work.

The first paper concentrates on the model as such, the general overview of this prospect, and metrics, gathering methods, and polling periods. A discussion was held on a test that was done at JVNC on variations in metrics based on variations of polling periods. More study is needed to fully understand the proposed statistics gathering model.

The second paper addresses information exchange. It should be possible for a client to connect to a server to ask for statistical information using SNMP language should that be transported back in some kind of transport mechanism. It also discusses the presentation form; especially, the possibility to display the results in a common way for comparison purposes.

There should be guidelines or rules for good behavior of networks with regards to the statistical performance. Once you have the statistics being presented in this way, you could also compare it to the general rule guidelines to have knowledge of network performance.

CURRENT MEETING REPORT

Reported by Bernhard Stockman/NORDUnet

ORAD Minutes

Minutes from Operations Area Directorate meeting at the Atlanta IETF.

Agenda:

- Introduction
- Presentations
- The issues for ORAD
- ORAD representatives
- Current and near future activities

Introduction:

Proposed goal for this session:

- To get a feeling for what the IETF Operational Area Directorate is, and is not, intended to be.

- To get a consensus on a set of prioritized activities to be initiated

Remarks for this meeting:

- We are all in it together, to make it work we need to cooperate on a set of basic things.

- With the common experience we should be able to identify what we can deploy of the existing technology

- We should also be able to identify and prioritize development needs

Presentations

- The IETF ORAD and the OPS Area, initial thoughts and views. (Phill Gross). The goal for an IETF Operations Area Directorate:

  - To improve the quality of:
    * The Internet
    * The tools
    * Objectives
    * Procedures

    * Methodologies
    * Interactions between operators and users
    * International cooperation

  - Advise the developers to create meaningful statistics and protocols.

  - International coordination. The ORAD as a forum for international interconnections.

- Ongoing coordination activities in the R &D networking field. (Bernhard Stockman NORDUNET).

  Relevant activities within the IEPG. The IEPG (a technical subgroup within CCIRN) has had two meetings so far. At these meetings, a list of prioritized items was defined:

  - Interncontinental link coordination
  - Global routing
  - Global DNS connectivity
  - Global address registration
  - NOC/NIC coordination aiming at common methods and practices and minimal basic services definitions.

  The IEPG is viewed as an Agenda setting group, i.e., IEPG will normally not undertake these items directly within itself but try to find relevant bodies for such actions.

  Example of ongoing activities:

  - IETF OPSTAT WG
  - RIPE mapping WG

- Commercial service providers' view on coordination. (Susan Estrada, CERFnet).

  It is important to realize that new network providers are entering the marketplace often. The Operational Requirements Area of the IETF and, in particular, the ORAD, have a large role to play to ensure the integrity of the network as new players enter. Additionally, the operations folks have a responsibility to provide input and advice to the protocol developers to insure that future implementations meet defined operational requirements.

  The ORAD should help set the agenda for the Operational Requirements Area. It should define the most pressing operational problems and seek common solutions and recommendations to solve those problems. The ORAD should also undertake the education of new network operators through the publication of guidelines for sensible operations of IP networks.

- A Nordic angle on the coordination issues, NETF Operations Coordinations WG - A Nordic "NFIX/NCIX"? - Mats Brunell NORDUNET

  The current (yesterday) situation:

  - R&D newtork operators only like NORDUnet, UNINETT, FUNET, SUNET and SURIS.

  - NETF (NORDUNET Technical and Engineering Forum) a possible way like the IETF for inclusion of all relevant parties.

  The new situation: Commercial service providers introduced on the Nordic scene like DataNET/Finland, SWIPnet/Sweden, TIPSnet/Sweden.

  The today provision have a growth of 100% or more/year.

  - There exist no long term planning.
  - There is limited resources.

  The today provision have a growth of 100% or more/year. This creates a need for coordination methods and procedures.

  Two possible approaches, "Like the situation" or "Do something about it".

  1. Like the situation: The existing situation with a multitude of uncoordinated network giving huge problems with regards to routing, name services, etc. This means we have to develop routing protocols etc. that can still work in the messy situation.

  2. A coordinated approach with planned routing and name services as well as new tools to aid in the collaborative process. The obvious answer is that we need to do both.

     A generally accessible Nordic interconnection point (NFIX/NCIX) has been proposed for implementation. The NFIX/NCIX will be available for both R &D and commercial network service providers.

     Sensitive to try to coordinate commercial service providers.

     The European scene:

     PTT's have a different view on have to provide Internet IP services than "we" do, the are used to the X.25/X.75 situation, and setting up a multitude of links to everywhere.

There is a need for knowledge transfer between different providers and a requirement for development to achieve scalability and operational stability.

The Issues for ORAD

To set off the discussion:

- Routing, what needs, what protocols, which topology?
- DNS connectivity to all world wide, how?
- IP address and name registration issues

ORAD Representatives

The structure of ORAD.

The question was expressed if ORAD should be a small group of people or if it should be formed from a large bunch of people.

There should be an election of ORAD members, for example 3 persons that together with the two co-chairs for 1 year is to form a working Executive with the responsibility for:

- Follow up on actions
- Promote membership and active work in between meetings
- Prepare meetings

Current and Near Future Activities

Below Working Groups exist today within IETF Operations Area:

- OPSTAT WG
- Benchmarking WG
- User Connectivity WG
- DDN WG
- Network Joint Management WG
- Topology and Engineering WG

Possible other operations WGs

- Routing Coordination
- DNS Coordination
- OSI Operations
- X.400 Operations
- X.500 Operations

A mailing list for the ORAD will be created by Susan Estrada named:

- orad@sdsc.edu
- orad-request@sdsc.edu.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Jordan Becker	becker@nis.ans.net
Eric Carroll	eric@utcs.utoronto.ca
Henry Clark	henryc@oar.net
James Conklin	conklin@bitnic.educom.edu
John Curran	jcurran@bbn.com
Tom Easterday	tom@cic.net
Robert Elz	kre@munnari.oz.au
Susan Estrada	Estradas@cerf.net
Peter Ford	peter@lanl.gov
Vince Fuller	vaf@stanford.edu
Shari Galitzer	shari@gateway.mitre.org
Phillip Gross	pgross@nis.ans.net
Tony Hain	alh@es.net
Martyne Hallgren	martyne@theory.tn.cornell.edu
Eugene Hastings	hastings@psc.edu
Ittai Hershman	ittai@nis.ans.net
Ken Jones	konkord!ksj@uunet.uu.net
Christopher Kolb	kolb@psi.com
Ruth Lang	rlang@nisc.sri.com
Louis Leon	osll@emuvm1.cc.emory.edu
Daniel Long	long@nic.near.net
April Marine	april@nisc.sri.com
Matt Mathis	mathis@psc.edu
David O'Leary	oleary@sura.net
Philippe Park	ppark@bbn.com
Marsha Perrott	mlp@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Erik Sherk	sherk@nmc.cit.cornell.edu
Bernhard Stockman	boss@sunet.se
Roxanne Streeter	streeter@nsipo.nasa.gov
Subu Subramanian	subu@qsun.att.com
Claudio Topolcic	topolcic@nri.reston.va.us
Ross Veach	rrv@uiuc.edu
Rudiger Volk	rv@informatik.uni-dortmund.de
Chris Waters-Pierandozzi	waters@jvnc.net

Wengyik Yeong	yeongw@psi.com
Osmund de Souza	desouza@osdpc.ho.att.com

3.5.1 Benchmarking Methodology (bmwg)

Charter

Chair(s):

Scott Bradner, sob@harvard.edu

Mailing Lists:

General Discussion: bmwg@harvisr.harvard.edu
To Subscribe: bmwg-request@harvisr.harvard.edu
Archive:

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

Done — Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Done — The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.

TBD — Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

Request For Comments:

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices"

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

BMWG Minutes

The BMWG met on Tuesday, July 30th in Atlanta during the IETF meeting.

The single topic of the discussions was to explore ways to more closely relate the design of tests for routers and bridges to the conditions found in the real world.

We explored the issues of bi-directional traffic, mixed protocols and random address and came to the conclusion that it would be difficult, at the least, to simulate a real-world network but that most of the above issues should be included in the test design.

It was concluded, in the absence of actual tests, that the choice of routing protocol probably did not make any performance difference to the routed protocol after the next-hop address had been learned and added to the routing cache. Tests should be performed to see if this is true.

We agreed to hold a video conference in mid September to continue refining the actual procedures that should be used to do throughput tests.

Attendees

James Barnes	barnes@xylogics.com
Jim Beers	beers@nr-tech.cit.cornell.edu
Scott Bradner	sob@harvard.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Arlan Finestead	arlanf@ncsa.uiuc.edu
Kenneth Goodwin	goodwin@psc.edu
Olafur Gudmundsson	ogud@cs.umd.edu
Jack Hahn	hahn@umd5.umd.edu
Phani Jujjavarapu	phani@cisco.com
Frank Kastenholz	kasten@europa.clearpoint.com
Michael Khalandovsky	mlk@ftp.com
Mark Lewis	mlewis@telebit.com
Bill Manning	bmanning@rice.edu
Bradford Parker	brad@cayman.com
Jason Perreault	perreaul@interlan.interlan.com
K.K. Ramakrishnan	rama@kalvi.enet.dec.com
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com.kr
A. Minick Rushton	rushton@stsci.edu
Gershon Schatzberg	439-3582@mcimail.com
John Scudder	jgs@merit.edu

Martha Steenstrup	msteenst@bbn.com
Osamu Takada	takada@sdl.hitachi.co.jp
Jonathan Wenocur	jhw@shiva.com
Cathy Wittbrodt	cjw@nersc.gov

3.5.2 Network Joint Management (njm)

Charter

Chair(s):

Gene Hastings, hastings@psc.edu

Mailing Lists:

General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu
Archive:

Description of Working Group:

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the Agenda of this group will include the IAB and its task forces, and groups within FARnet. In particular FARnet has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of'the group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing. The group will meet at the next IETF and determine the other schedules. Sub-groups may meet between IETF meetings.

Goals and Milestones:

None specified

CURRENT MEETING REPORT

Reported by Kenneth Goodwin/PSC

NJM Minutes

Instead of putting actual copies of maps into central repositories, pointers as to where to get the maps should be placed there. This would make it easier and quicker to update maps, since the network manager does not have to distribute maps, just update the local copy.

It seems as though the very bade traceroute has disappeared from many peoples' list of tools. Matt Mathis said that it appears that traceroute is included in Ultrix 4.2, although it is unknown whether it does third party or not. Sources for traceroute are available via anonymous ftp from ftp.lbl.gov (may not include the third party mods) and from nic.near.net (in the pub or src directory).

Also, Matt M. provided some examples of traceroutes into the backbone that failed. The common denominator being that if the endpoint or the corner of the third party traceroute is in the backbone, it will fail. Merit says that this may be by design.

Merit/ANS discussed their architecture plans for the CNSS [see also slides from the plenary technical presentation by Elise Gerich and Jordan Becker - EFH]. Currently, they decrement the TTL through the routers, but would like to make the CNSS's in a POP appear to be one router by only decrementing the TTL on entry or exit. This would make the separate CNSS's invisible. One point against this is that in a tightly coupled FDDI, this could create a giant loop. (??) Merit also said that future cards will receive routing updates from the CPU.

Fallback paths were also discussed. (In particular fallback paths of lower capacity) Some important points were that if a backup path exists, then it should be carrying some traffic all the time, so that the backup path's status is known at all times. Backup paths of lesser capacity can easily be flooded by production traffic, so some means of limiting traffic must be made. Things like mail can be MX'ed to the backup path, and all other traffic could be blackholed.

As nets become faster, monitoring is becoming harder. Merit is currently using periodic sampling of 1 in every 50 packets on the T3. Merit would like to use a more stochastic process. More work needs to be put into this to determine a more accurate sampling process.

With more and more nets appearing, some better method of reporting outages is needed. An outage of a campus with large nets could easily flood people. A template for reporting outages is needed, so that a database can parse these messages and store the information. Thus, one need not even read the mail, but query the database for the net in question.

Merit is working on something like this for trouble ticket tracking.

Since high schools are entering the internet, two problems are occurring. Under what name should the schools appear. (us or edu) Also, how do we get them started? Nearnet offers a full service option for new people, that completely orients the newcomer. User Services should also target the end user and not just the network operators.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Jordan Becker	becker@nis.ans.net
Eric Carroll	eric@utcs.utoronto.ca
Henry Clark	henryc@oar.net
Tom Easterday	tom@cic.net
Susan Estrada	Estradas@cerf.net
Vince Fuller	vaf@stanford.edu
Maria Gallagher	maria@nsipo.arc.nasa.gov
Kenneth Goodwin	goodwin@psc.edu
Jack Hahn	hahn@umd5.umd.edu
Eugene Hastings	hastings@psc.edu
Ittai Hershman	ittai@nis.ans.net
Daniel Long	long@nic.near.net
Matt Mathis	mathis@psc.edu
Philippe Park	ppark@bbn.com
Marsha Perrott	mlp+@andrew.cmu.edu
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Roxanne Streeter	streeter@nsipo.nasa.gov
Ross Veach	rrv@uiuc.edu
Chris Waters-Pierandozzi	waters@jvnc.net
Gerard White	ger@concord.com
Cathy Wittbrodt	cjw@nersc.gov

3.5.3 Operational Statistics (opstat)

__Charter__

Chair(s):

Bernhard Stockman, boss@sunet.se
Phillip Gross, pgross@nis.ans.net

Mailing Lists:

General Discussion: oswg-l@wugate.wustl.edu
To Subscribe: oswg-l-request@wugate.wustl.edu
Archive:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The Working Group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Done	Agreement on a model.
Done	Survey for most useful and popular metrics.
Done	Survey for most useful and popular presentation formats.
Dec 1990	Identify similar efforts being performed by other groups.
Done	Define a common minimal set of metrics.
Mar 1991	Propose a MIB for metrics not already there.
Done	Define a common storage format to facilitate data sharing.
Done	Define common presentation formats to make data comparable.
Mar 1991	Develop outline, and make writing assignments for paper (Opstat1) documenting March 91 milestones.
May 1991	Complete paper Opstat1.

May 1991	Possible mid-term meeting to review Opstat1.
May 1991	Submit Opstat1 as Internet Draft.
Jul 1991	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Jul 1991	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Jul 1991	Propose old tools to be retrofitted.
Jul 1991	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools
Sep 1991	Complete paper Opstat2
Sep 1991	Possible mid-term meeting to review Opstat2
Sep 1991	Submit Opstat2 as Internet-Draft
Dec 1991	Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Osmund de Souza/AT&T

OPSTAT Minutes

The proposed Agenda for the meeting was:

- Administration
- Metrics - old, new, and exotic
- Measurement polling periods
- Query language and exchange protocol
- Report formats
- Closing

The proposed Agenda was approved. Osmund de Souza volunteered to take the Minutes.

Metrics: The group reviewed the metrics that have been proposed so far. Bernhard had written a "metrics paper" before the meeting to help focus the discussion. We agreed that the basic set of measurements should include the following:

- Octets in/out, unicast packets in/out, non-unicast packets in/out for each interface.

- IP packets forwarded, IP packet discarded, (similar counts for other network layer protocols), for each router.

We then talked about whether we should expand this set to include new and exotic metrics and if so, what they should be. There were ideas about measuring availability, stability, delay performance, congestion, and line errors. We decided that rather than try to come to a conclusion at the meeting we would take it to the mailing list. We did agree that the metrics paper should not explicitly include or exclude these new measures, but rather should leave the door open for expanding the base set later. We also agreed that initially it would be wise to work with variables within the current MIB structure, though this should not limit us as the opstat architecture matures.

Measurement: As usual, we had a lively discussion about measurement polling intervals. Vikas Aggarwal presented the results of a study he did on JvNCNet. In the study he periodically polled each interface in the network to download traffic data to a central manager. Since it took three minutes to poll all the interfaces, the polling period was three minutes. He computed the utilization of each link in the network for each three minute interval, and also for integer multiples of three minutes by aggregating the measurements. The results showed that as the "integration" interval increased, the average of the average utilization in each interval remained unchanged (to within some small error) but the variation in average utilization for each interval decreased. Hence, for instance, the three minute utilizations had several peaks above 70

The discussion led to the understanding that measurement collection should ideally be done on a small enough time scale to capture short-term variations and peaks in traffic loads. Burstier traffic requires a smaller polling interval. However, small interval polling on a large network may be difficult for most network management systems, and there is the problem of storing the measurements for later processing. While no definite conclusion was reached on this one, the group is close to settling the issue by gravitating towards a polling interval of 5-15 minutes.

Query Language: We reviewed the paper that Bernhard had put together (and circulated on the mailing list) to define a query language syntax for the statistics client/server model. The language is geared around the following simple commands: login, exit, help, format, and select. There was some discussion about the meaning of the "format" command and what help the "help" command was supposed to give. We decided that "format" should let the user know the storage/presentation format of the data in the query, and that "help" should explain the meaning and syntax of commands as well as the nature of the data named in the query. Bernhard agreed to incorporate our comments in the document before resending it to the list.

David O'Leary said that he may be able to get a (more?) graduate student to work on a statistics client/server protocol. He did not have a clear idea of the direction in which the work would proceed, but would keep the group updated through the mailing list.

Report Formats: We did not have enough time to spend discussing the format of the reports. We reviewed the ideas that had come up at previous meetings and that Bernhard had summarized in a "reports paper". We did agree that we may have spent too much effort in the past defining the "reports for upper management", aka The McDonald's Report. We felt that we should concentrate on engineering reports for network operations and planning. These reports could then be stripped of an appropriate amount of useful information to present to upper management.

Closing: An item that came up during the discussion was the need for us to interface more with other groups, notably SNMP, and Internet Accounting. We were also pleased about the progress we had made since our previous meeting in St. Louis.

Finally the meeting agreed to use the papers produced by Bernhard as a basis for the continued work.

Attendees

Thomas Brisco	brisco@rutgers.edu
Gigi Chu	gigic@hpspd.spd.hp.com
Henry Clark	henryc@oar.net
Shari Galitzer	shari@gateway.mitre.org
Kenneth Goodwin	goodwin@psc.edu
Phillip Gross	pgross@nis.ans.net
Michael Khalandovsky	mlk@ftp.com

Tim Lee-Thorp	ngc!tim@uunet.uu.net
Chris Myers	chris@wugate.wustl.edu
David O'Leary	oleary@sura.net
Robert Reschly	reschly@brl.mil
Mark Saake	saake@llnl.gov
Timothy Salo	tjs@msc.edu
Erik Sherk	sherk@nmc.cit.cornell.edu
Frank Solensky	solensky@clearpoint.com
Bernhard Stockman	boss@sunet.se
Roxanne Streeter	streeter@nsipo.nasa.gov
David Waitzman	djw@bbn.com
Chris Waters-Pierandozzi	waters@jvnc.net
Gerard White	ger@concord.com
Osmund de Souza	desouza@osdpc.ho.att.com

3.5.4 Topology Engineering (tewg)

Charter

Chair(s):
TBD ,

Mailing Lists:
General Discussion: tewg@devvax.tn.cornell.edu
To Subscribe: tewg-request@devvax.tn.cornell.edu

Description of Working Group:

The Topology Engineering Working Group monitors and coordinates connections between networks, particularly routing relationships.

- Monitor interconnectivity among national and international backbones and mid-level networks.
- Monitor interconnection policies with a view of moving toward a common scheme for managing interconnectivity.
- Act as a forum where network engineers and representatives of groups of networks can come together to coordinate and tune their interconnections for better efficiency of the Internet as a whole.

Goals and Milestones:

Ongoing	Reports to the Internet community will be given reflecting what we learn each quarter. This periodic report will be of use to the IETF, to FARnet, and to the CCIRN members.
Dec 1990	An immediate project is to produce an RFC which will help mid-level networks when changing their interconnectivity.

CURRENT MEETING REPORT

TEWG

The Topology Engineering Working Group has been officially concluded as of August 2, 1991.

3.5.5 User Connectivity (ucp)

Charter

Chair(s):

Dan Long, long@nic.near.net

Mailing Lists:

General Discussion: ucp@nic.near.net
To Subscribe: ucp-request@nic.near.net
Archive:

Description of Working Group:

The User Connectivity Working Group will study the problem of how to solve network users' end-to-end connectivity problems.

Goals and Milestones:

Done	Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.
TBD	Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate.

Internet Drafts:

"FYI on an Internet Trouble Ticket Tracking System for addressing Internet User Connectivity Problems", 02/11/1991, M. Mathis, D. Long <draft-ietf-ucp-connectivity-00.txt>

"NOC Internal Integrated Trouble Ticket System Functional Specification Wish-list", 02/26/1991, Dale S. Johnson <draft-ietf-ucp-tt-01.txt>

3.6 Routing Area

Director(s):

- Bob Hinden: hinden@bbn.com

Area Summary reported by Bob Hinden/BBN

Border Gateway Protocol Working Group, Chair - Yakov Rekhter

The BGP Working Group met on Wednesday and discussed the following items:

- Supporting Default Routes
- Interactions between OSPF and BGP
- Multicast and BGP
- NEXT_HOP_SNPA Path Attribute

The first item discussed dealt with supporting default routes in BGP. The group decided that the default would be supported by announcing net 0.0.0.0. No changes to the protocol are required to support this. The selection of this value for default is compatible with what is currently being done on the MILNET with EGP. This is a good step towards replacing EGP with BGP in the Internet.

BGP met with the OSPF group to work on the revisions to a document, written by Kannan Varadhan, describing interactions between OSPF and BGP. John Moy will pass the comments of the BGP Working Group to the author. The group expects to have an Internet Draft ready for the next IETF meeting.

The group discussed a proposal by Scott Brim on how BGP can be used to support interautonomous system multicast. This approach is compatible with the work that the Multicast OSPF group is doing.

The last item discussed, was a a proposal for propagating MAC layer addresses with BGP in conjunction with the IP network addresses. The purpose of this path attribute is to eliminate the need to do additional address resolution on network like SMDS, where there are charges for traffic.

IS-IS for IP Internets Working Group, Chair - Ross Callon

The ISIS group met briefly and worked on updates to the Integrated IS-IS spec. Most of the changes were updating the references to the final OSI version of the IS-IS specification. The group also worked on the MIB for IS-IS. It is based on the OSI MIB, with extensions to support variables for IP.

Open Shortest Path First IGP Working Group, Chair - John Moy

The group reviewed the OSPF Trap MIB document. They expect it to be ready to submit to the network management directorate. They discussed a proposal for a new OSPF option called "Not So Stubby Areas" (NSSA). This option will provide for improved support for RIP clouds attached to OSPF domains and help with the transition of domains from RIP to OSPF. The group also meet jointly with the BGP Working Group and discussed revisions on the document describing how OSPF and BGP should interoperate.

Multicast Extensions to OSPF Working Group, Chair - Steve Deering

The group met for one session and resolved their last significant technical issue dealing with the routing, within an OSPF domain, of multicast packets originating in other domains. The solution selected is based on reverse path forwarding, which allows for several different multicast routing schemes at the inter-domain level. This approach is compatible with the proposed multicast extension to BGP. The group is working on an Internet Draft protocol document and expects to have, now that the technical issues are resolved, a draft at the next IETF meeting.

The group also worked on limiting the scope on multicast addresses. Currently multicast addresses are global. They worked on defining a property to allow the Multicast addresses to be limited to a defined area. This would allow for multicast groups to be restricted to a single administrative domain, such as a single campus or corporate site. This would provide for better scalability of multicast routing protocols and better security than the current scheme. An Internet Draft will be written to describe this proposal.

IP Over Large Public Data Networks Working Group, Chair - George Clapp

The group completed work on drafts for the operation of IP over Frame Relay and recommended that these be submitted for Proposed Standard.

The group next worked on two approaches to address resolution and routing on large public data networks. Work on this will continue. They also discussed IP over both circuit and packet (X.25) ISDN networks. They are considering a backwards compatible update to RFC 877. They plan to enhance the negotiation procedures for IP Over Frame Relay and investigate their applicability to IP over Circuit ISDN.

Inter Domain Policy Routing Working Group, Chair - Martha Steenstrup

The group met and worked on two Internet Draft documents relating to IDPR. These are a configuration guide and a MIB. The configuration issues related to the use of the Domain Name System (DNS) to resolve IP address to administrative domain identifiers. MIB issues discussed were:

- How much configuration information to include in the MIB, and
- How to represent policy information.

3.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):
Yakov Rekhter, yakov@watson.ibm.com

Mailing Lists:
General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu
Archive:

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done	Complete development of version 2 of the Border Gateway Protocol (BGP).
Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
Done	Develop a MIB for BGP.
Done	Work with the Security Area to enhance the provision for security in BGP.
Done	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

Internet Drafts:

"Definitions of Managed Objects for the Border Gateway Protocol (Version 3)", 07/17/1990, Steven Willis, John Burruss <draft-ietf-iwg-bgp-mib-02.txt>

"A Border Gateway Protocol 3 (BGP-3)", 01/25/1991, Yakov Rekhter, Kirk Lougheed <draft-ietf-bgp-bgp3-00.txt>

"Border Gateway Protocol NEXT-HOP-SNPA Attribute", 04/15/1991, Paul Tsuchiya <draft-ietf-bgp-nexthop-00.txt>

"Experience with the BGP Protocol", 05/08/1991, Yakov Rekhter <draft-ietf-bgp-experience-00.txt>

"BGP Protocol Analysis", 05/08/1991, Yakov Rekhter <draft-ietf-bgp-analysis-00.txt>

"Default Route Advertisement In The Border Gateway Protocol", 08/09/1991, Dimitry Haskin <draft-ietf-bgp-defaultroute-00.txt>

"Multicast Communications Using BGP", 08/26/1991, Scott Brim <draft-ietf-bgp-multicast-01.txt>

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"

CURRENT MEETING REPORT

Reported by Yakov Rekhter/IBM

BGP Minutes

The Border Gateway Protocol Working Group met on Wednesday. We discussed the following Agenda items:

1. Supporting default route in BGP.
2. Interaction between OSPF and BGP.
3. Inter-autonomous system multicast with BGP.
4. NEXT_HOP_SNPA path attribute.

On (1) it was decided that the default route will be supported in BGP by announcing net 0.0.0.0. Support for the default route in BGP is needed because some of the routers in MILNET have very tight memory constraints. We expect to generate a document that will describe details of how default is supported in BGP, and what are the implications of generating and accepting detault route propagated via BGP.

On (2) we discussed the document written by Kannan Varadhan. John Moy agreed to pass all the comments produced during the meeting to the author of the document. We expect to have an Internet Draft version of the document before the next IETF.

On (3) Scott Brim agreed to generate an Internet Draft that will describe a specific approach for inter-AS multicast, and what are the implications of that approach on BGP.

On (4) we pointed out the need to clarify certain aspects of the document. Once clarifications will be in place, the Working Group recommends to advance the document to a Proposed Internet Standard.

Attendees

Sudhanshu Aggarwal	wlsonu@watson.ibm.com
Nagaraj Arunkumar	nak@3com.com
Helen Bowns	hbowns@bbn.com
Scott Brim	swb@devvax.tn.cornell.edu
Dennis Ferguson	dennis@canet.ca
Susan Hares	skh@merit.edu
Dimitry Haskin	dhaskin@bbn.com
Jeffrey Hoenig	jch@risci.cit.cornell.edu
Kathleen Huber	khuber@bbn.com
Phani Jujjavarapu	phani@cisco.com
Walter Lazear	lazear@gateway.mitre.org
Louis Leon	osll@emuvm1.cc.emory.edu
Mark Lewis	mlewis@telebit.com

Tony Li	tli@cisco.com
Matt Mathis	mathis@psc.edu
April Merrill	Merrill@dockmaster.ncsc.mil
Dennis Morris	morrisd@imo-uvax.dca.mil
Charles Perkins	perk@watson.ibm.com
Yakov Rekhter	yakov@ibm.com
Martha Steenstrup	msteenst@bbn.com
Ross Veach	rrv@uiuc.edu
Robert Woodburn	woody@cseic.saic.com

3.6.2 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

George Clapp, meritec!clapp@uunet.uu.net

Mailing Lists:

General Discussion: iplpdn@nri.reston.va.us
To Subscribe: iplpdn-request@nri.reston.va.us
Archive: /ietf.mail.archives/iplpdn.mail.archive

Description of Working Group:

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

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

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

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

Goals and Milestones:

Done	Establish priorities and dates of completion for documents.
TBD	Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.
TBD	Routing of IP datagrams across very large internets implemented SMDS and on other PDNs.
TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.

Internet Drafts:

"Discovery and Routing over the SMDS Service", 06/17/1991, Paul Tsuchiya <draft-tsuchiya-routingsmds-01.txt>

"Multiprotocol Interconnect over Frame Relay Networks", 06/17/1991, T. Bradley, C. Brown, A. Malis <draft-ietf-iplpdn-ipoverframerelay-03.txt>

"Inverse Address Resolution Protocol", 06/17/1991, T. Bradley, C. Brown <draft-ietf-iplpdn-inarp-02.txt>

"Management Information Base for Frame Relay DTEs", 06/17/1991, Caralyn Brown, Fred Baker, Charles Carvalho <draft-ietf-iplpdn-frmib-01.txt>

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

IPLPDN Minutes

Opening Remarks

This was the third meeting of the IP over Large Public Data Networks Working Group. The following was the Agenda of the meeting:

```
Monday, July 29, 1991
     PM      IP over Frame Relay

Tuesday, July 30, 1991
     AM      IP over Frame Relay
     PM      IP over Frame Relay

Wednesday, July 31, 1991
     AM      Address Resolution and routing
     PM      Address Resolution and routing

Thursday, August 1, 1991
     AM      IP over ISDN
```

Monday, July 29, Through Tuesday, July 30, 1991

After brief introductory remarks by the Chair, Andy Malis opened with a presentation giving an overview of the current status of encapsulation over Frame Relay (copies of the slides are included with these minutes). The Working Group then turned to a close review of the following three documents:

- Multiprotocol Interconnect over Frame Relay Networks
- Inverse Address Resolution Protocol
- Management Information Base for Frame Relay DTEs

(The revised documents are available as Internet Drafts entitled "draft-ietf-iplpdn-ipoverframerelay-03.txt," "draft-ietf-iplpdn-inarp-02.txt", and "draft-ietf-iplpdn-frmib-01.txt," respectively.) Prior to the meeting, email discussion of these documents had progressed to the point where there was general satisfaction with the contents, and this was the final review prior to submitting the drafts to the IAB for approval.

There were some changes in content to the first document, "Multiprotocol Interconnect over Frame Relay Networks." The group agreed that the preferred method of indicating protocols other than IP was via a Network Layer Protocol ID (NLPID) value of 0x80 and SNAP. The option of indicating the Ethertype via an NLPID value of 0xCE was de-emphasized, and text describing this approach was moved to an appendix of the document.

A second change to the document was the explicit depiction of the encoding of bridged MAC frames. This text was added to minimize the possibility of incompatible implementations.

The two other documents were reviewed and adopted without significant modifications.

Keith Mader asked the group to consider developing a protocol by which end points can negotiate configuration and service parameters over Frame Relay permanent virtual circuits (PVCs). Keith also asked whether the OUI/PID value of 0x00-80-C2 0x00-0E could be used to indicate this protocol. The group agreed to undertake this work with the caveat that the work should not be extended into the realm of signaling. Fred Baker suggested that a useful distinction between negotiation and signaling is that the former is between end stations and the latter is between an end station and the network. Keith offered to develop a baseline document which could accompany a request to the IEEE 802.1 Working Group for a PID value.

By the end of Tuesday, July 30, the group was satisfied with the modified documents and agreed to submit them to the IAB for approval. Caralyn Brown was scheduled to give a presentation to the plenary on Wednesday evening describing the approach taken by the group.

Wednesday, July 31, 1991

All of Wednesday was spent discussing address resolution across large public data networks, with an emphasis on SMDS. John Hagan of the University of Pennsylvania and John Garrett of AT&T Bell Laboratories presented a discussion of "directed ARP," a technique in which ARP requests are sent to the source of routing information rather than broadcast to all participants in a Logical IP Subnetwork (LIS).

"Directed ARP" brought into question a fundamental assumption of the Internet, which is that if the network portion of the IP addresses of two devices differ, then these devices are attached to different networks and can only communicate via an intermediary router. The group debated the benefits and costs of violating this assumption and ended by recognizing that the IPLPDN Working Group could not resolve this issue and that other Working Groups of the IETF should become involved.

Paul Tsuchiya then presented a discussion of an alternative approach which makes use of BGP, and he expressed approval of the directed ARP technique. The group felt that a combination of directed ARP with an enhanced BGP may represent a solution to the address resolution and routing issues. Paul offered to work with John Hagan and John Garrett to investigate a synthesis of their approaches before the next IETF meeting.

Thursday Morning, August 1, 1991

Dory Leifer led a discussion of IP over circuit ISDN which focused on the relative merits of the Point-to-Point Protocol (PPP) and Frame Relay for use over the B channel. No resolution of reached, though members of the group volunteered to augment the Frame Relay approach with negotiation procedures, and the group agreed to discuss the topic again during the next meeting.

As a final topic, the draft written by Robert Ullman for IP over X.25 was discussed. Andy Malis and George Clapp volunteered to contact Robert and to investigate the possibility of updating RFC 877 during the interim.

In closing, the Chair congratulated the Working Group for a very productive meeting and thanked members for their hard work in drafting and revising the documents. The group then adjourned.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Fred Baker	fbaker@emerald.acc.com
Tom Benkart	teb@saturn.acc.com
Arthur Berggreen	art@acc.com
Helen Bowns	hbowns@bbn.com
Caralyn Brown	cbrown@wellfleet.com
Carter Bullard	carter@gatech.edu
Eric Carroll	eric@utcs.utoronto.ca
Charles Carvalho	charles@sage.acc.com
Cho Chang	chang_c@apollo.hp.com
John Chang	jrc@uswest.com
Peter Chang	tpc@mtunm.att.com
Richard Cherry	rcherry@novell.com
George Clapp	meritec!clapp@uunet.uu.net
Richard Cogger	rhx@cornellc.cit.cornell.edu
Nabil Damouny	nabil@tdd.sj.nec.com
Jill Foster	jill.foster@newcastle.ac.uk
Stephen Gabe	spgabe@bnr.ca
John Garrett	jwg@edsel.att.com
Robert Griffioen	rgriff@bnr.ca
John Dotts Hagan	hagan@dccs.upenn.edu
Tony Hain	alh@es.net
Joel Halpern	jmh@network.com
Frank Heath	heath@cmc.com
B.V. Jagadeesh	bvj@3com.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Tom Kessler	kessler@sun.com
Paulina Knibbe	knibbe@cisco.com
Joseph Lawrence	jcl@sabre.bellcore.com

John Leddy	jleddy@ibm.com
Dory Leifer	Dory_Leifer@um.cc.umich.edu
Chao-Yu Liang	cliang@synoptics.com
Daniel Long	long@nic.near.net
John Lynn	lynn@ttcllcat.cit.cornell.edu
Shane MacPhillamy	slm@netrix.com
Keith Mader	keith@python.eng.microcom.com
Andrew Malis	malis@bbn.com
Allison Mankin	mankin@gateway.mitre.org
Bill Manning	bmanning@rice.edu
Bernie May	bfm@houxa.att.com
Chandy Nilakantan	csn@3com.com
Michael O'Dell	mo@bellcore.com
David O'Leary	oleary@sura.net
Philippe Park	ppark@bbn.com
Bradford Parker	brad@cayman.com
Radia Perlman	perlman@radia.enet.dec.com
James Philippou	japhilippou@eng.xyplex.com
David Piscitello	dave@sabre.bellcore.com
Lars Poulsen	lars@cmc.com
K.K. Ramakrishnan	rama@kalvi.enet.dec.com
James Rees	rees@ifs.umich.edu
Ron Roberts	roberts@jessica.stanford.edu
Frank Solensky	solensky@clearpoint.com
Wayne Staats	staats@adm.csc.ncsu.edu
Martha Steenstrup	msteenst@bbn.com
Osamu Takada	takada@sdl.hitachi.co.jp
Mike Truskowski	truskowski@cisco.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Chris Waters-Pierandozzi	waters@jvnc.net
L. Michele Wright	uncng!michele@uunet.uu.net
Wengyik Yeong	yeongw@psi.com
Chin Yuan	cxyuan@pacbell.com

Frame Relay Multiprotocol Encapsulation Issues

IETF
7/29/91

Andrew G. Malis
BBN Communications

Issues to Discuss

- Summary of recent changes
- Use of 0xCE NLPID
- XID - mandatory or optional
- Proposal by Microcom (Keith Mader)

Meeting with ANSI X3S3.3

- In early June, C. Brown, J. Halpern, A. Malis met with ANSI X3S3.3 to discuss encapsulation
- X3S3.3 suggested use of ISO/IEC TR 9577, "Protocol Identification in the Network Layer"
- TR 9577 defines Network Layer Protocol IDs (NLPIDs) for CLNP (ISO 8473), IP, X.25, IEEE SNAP, Q.931, plus other ISO and CCITT protocols
- Use of TR 9577 brings encapsulation into ISO/CCITT conformance
- Decided to use undefined NLPID (0xCE) for Ethertype/Format ID escape

IP Encapsulation using NLPID

DLCI	
Control = 03 (UI)	NLPID = CC (IP)
IP Packet . .	

Using IEEE 802.1 Bridging PIDs

- In June, IEEE 802.6 & SMDS Interest Group proposed remote bridging over SMDS via LLC and SNAP, using PIDs defined for IEEE 802.1's Org. ID of 00-80-C2.
- Last week, in Hawaii, IEEE 802.1 approved SNAP PIDs for bridging via SNAP using 802.1's Org ID.

802.1 Remote Bridging PIDs

	With FCS		Without FCS
00-01	802.3/Ethernet	00-07	802.3/Ethernet
00-02	802.4	00-08	802.4
00-03	802.5	00-09	802.5
00-04	FDDI	00-0A	FDDI
00-05	802.6	00-0B	802.6
00-06	802.9	00-0C	802.9

Bridging Encapsulation Using SNAP

DLCI	
Control = 03 (UI)	Optional pad = 00
NLPID = 80 (SNAP)	(three octet) SNAP
Protocol ID or Org Code = 00-80-C2	
IEEE 802.1 PID	
Bridged Frame . .	

Fragmentation Format Using SNAP

DLCI		
Control = 03 (UI)	Optional pad = 00	
NLPID = 80 (SNAP)	(three octet) SNAP	
Protocol ID or Org Code = 00-80-C2		
IEEE 802.1 PID = 00-0D		
Fragment Sequence Number		
F	Reserved	Offset
Fragment Data . .		

Use of NLPID 0xCE

- Use of 0xCE NLPID controversial because it is not yet included in TR 9577
- It is not really needed - SNAP encapsulation can be used
- It may make it difficult to submit spec to ANSI and ISO
- Frame Relay Forum technical committee prefers not using 0xCE

Ethertype Encapsulation using NLPID

DLCI	
Control = 03 (UI)	NLPID = CE (new)
Ethertype	
Protocol Packet . .	

Ethertype Encapsulation using SNAP

DLCI	
Control = 03 (UI)	Optional pad = 00
NLPID = 80 (SNAP)	(three octet) SNAP
Protocol ID or Org Code = 00-00-00	
Ethertype	
Protocol Packet . .	

XID Exchange (Section 8)

- Section 8 specifies optional use of Exchange Identification (XID) exchange at PVC initialization.
- This allows dynamic negotiation of the maximum framesize supported by the PVC endpoints, and is also useful for future use of acknowledged I-frames (rather than UI-frames).
- Q.921 has an inconsistency regarding XID. Section 3.6.12 requires XID responses, but Appendix IV (Q.922 App. III) specifies that you use configured defaults if response does not arrive in 7.5 seconds (two retransmissions, each after 2.5 seconds)
- Optional use seems OK, due to the Appendices

XID Format

DLCI	
Control = AF (XID)	Format ID (FI) = 82
Group ID (GI) = 80	Group len high octet
Group len low octet	Xmit frame size = 5
Param len = 2	Size high octet
Size low octet	Rec frame size = 6
Param len = 2	Size high octet
Size low octet	Xmit window size = 7
Param len =1	Window size (1-127)
Retrans. timer = 9	Param len =1
Timer (0.1-25.5 sec)	

Proposal by Microcom

- Keith Mader of Microcom would like to present suggested language to replace Section 14 of the specification.

3.6.3 ISIS for IP Internets (isis)

Charter

Chair(s):
Ross Callon, callon@bigfut.enet.dec.com

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

Description of Working Group:

The IETF IS-IS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

Done	Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.
TBD	Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
TBD	Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

Request For Comments:

RFC 1195 "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"

CURRENT MEETING REPORT

Reported by Ross Callon/DEC

IS-IS Minutes

The IS-IS Working Group met at the IETF meeting in Atlanta, Georgia. There were two topics of discussion: A brief overview of the status of the IS-IS spec (led by Ross Callon), and a presentation and longer discussion of the SNMP MIB for IS-IS (led by Chris Gunner).

1. Status of IS-IS

   Ross reported that the OSI IS-IS Intra-Domain routing protocol (ISO DIS 10589) has completed the Draft International Standard (DIS) ballot, and all ballot comments were successfully resolved at a recent ISO meeting. This implies that the ISO IS-IS will be progressing to final International Standard state relatively quickly. This, in combination with the completion of a couple of Integrated IS-IS implementations means that it is a good time to start think about issuing an update to RFC 1195.

   Ross then gave a quck overview of some minor changes that would be involved:

   - Reference to ISO standard

     RFC 1195 reference the DP version of ISO IS-IS. This clearly needs to be updated to reference the final International Standard version, when available. This would also imply that Annex B (Encoding of Sequence Number Packets) can be removed. It turns out that we were either lucky or good, and the sequence number format in the current ISO document is compatible with Annex B.

   - RIP (or other external routes) at level 1

     Currently the spec says that this is not allowed. There are good technical reasons why we don't want fully general external connections at both level 1 and level 2. However, there may be many cases where we have a small RIP "island" which is only reachable via a level 1 area. For example, this is very likely to occur during transition from a RIP routing domain to an Integrated IS-IS routing domain. No technical change is needed, but the document should be upgraded editorially to specify that this is permissible.

   - Default IP route at level 1.

     There will be some cases where level 1 routing is IP-capable (using Integrated IS-IS) but level 2 routing is not (such as using OSI-only IS-IS at level 2, or possibly during phase 4 to phase 5 DECnet(TM) transition). In this case, there needs to be a way for level 1 routers to know where to send traffic destined to outside of the area (for example, one single level 2 router might be running RIP with external routers). The solution to this is to allow IP Default Route (subnet mask of all 0's) at level 1, and to specify that for level 1 only routers which see the default route advertised in level 1 LSPs, this takes precedence over forwarding traffic to level 2 routers.

- Compatibility with earlier versions of IS-IS

  There should be a "for information only" annex which specifies the differences between RFC 1195, and the updated RFC. This will also specify how to ensure interoperability between old and new routers.

- IS-IS / BGP interaction

  Yakov Rekhter brought up the issue of interaction between IS-IS and BGP. Ross and Yakov will work on this issue off-line, and report results back to the Working Group.

- Encoding of Authentication Field

  Someone brought up the issue that RFC 1195 and DIS 10589 both have an authentication field, in which the encoding and use is identical but the code value is different. The Working Group agreed that this was an unnecessary redundancy, and that we should use the value from 10589.

- Ships in the Night Operation

  RFC 1195 currently has sufficient functionality to allow operating two instances of IS-IS in "Ships in the Night" mode – one instance would be for IP-only routing, and one for OSI-only routing. However, just how to do this is not written down anywhere. It was agreed that this should be writted down, with the approach "you don't have to be capable to run two instances of IS-IS, but if you do run two instances then this is how you do it". Generally, you demultiplex on the "Protocols Supported" field, and optionally may use authentication to protect against accidental merging of the two logical routing domains by a misconfigured router.

2. MIB for IS-IS

Chris Gunner then gave a detailed presentation of the proposed MIB for IS-IS. This MIB allows management of Integrated IS-IS (including full management of both ISO 10589 and RFC 1195) using SNMP. This is based on the GDMO (i.e., ISO format network management information) contained in DIS 10589, with additional objects added for management of RFC 1195.

The recent progression of 10589 in ISO will result in some changes to the GDMO in 10589. Chris will need to produce an update of the MIB in order to maintain alignment with the ISO document.

There was a discussion of the size of the MIB. In particular, there are situations where several similar things are in different tables For example, different sorts of circuits currently are managed using different tables. There is substantial overlap between these different tables. The alternative is to have one type of table for all circuits, with some fields not always used. This implies slightly more bits will be transmitted on the wire, but allows a smaller MIB and less software code (e.g., data structures are simpler). The Working Group agreed that the latter approach was preferable, at least in those cases where the overlap is relatively large.

The group agreed that the MIB should permit multiple instances of Integrated IS-IS and/or IS-IS to be managed in a system. This means turning single instance objects in groups into table objects. The group also agreed that all such table entries should be capable of creation and deletion to mirror the creation and deletion capabilities of the DIS 10589 managed objects to which they are equivalent.

3. Other Issues

   Yakov Rekhter pointed out that the ISO GDMO of IS-IS does not allow measurement of routes coming from external protocols to IS-IS. Chris and Ross agreed to bring up this issue with the folks working on the ISO specification.

   Outside of the Working Group, a couple of folks brought up the issue of how to handle the "3rd party router" case (a single routing domain having several routers on a broadcase or general-topology network with only one router running BGP). Ross will write up a proposal on how to deal with this and discuss it within the Working Group.

Attendees

Nagaraj Arunkumar	nak@3com.com
William Barns	barns@gateway.mitre.org
Scott Barvick	sbarvick@wellfleet.com
William Biagi	bbiagi@cos.com
John Burruss	jburruss@wellfleet.com
Ross Callon	callon@bigfut.enet.dec.com
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@canet.ca
Robert Griffioen	
Chris Gunner	gunner@osicwg.enet.dec.com
Robert Hagens	hagens@cs.wisc.edu
Susan Hares	skh@merit.edu

Manu Kaycee	kaycee@trlian.enet.dec.com
Paulina Knibbe	knibbe@cisco.com
Dale Land	land@lanl.gov
Chao-Yu Liang	
Shane MacPhillamy	slm@netrix.com
Bill Manning	bmanning@rice.edu
Dennis Morris	morrisd@imo-uvax.dca.mil
Yakov Rekhter	yakov@watson.ibm.com
Mike Truskowski	truskowski@cisco.com
Rick Wilder	rick@gateway.mitre.org
L. Michele Wright	uncng!michele@uunet.uu.net

3.6.4 Inter-Domain Policy Routing (idpr)

Charter

Chair(s):

Martha Steenstrup, msteenst@bbn.com

Mailing Lists:

General Discussion: idpr-wg@bbn.com
To Subscribe: idpr-wg-request@bbn.com
Archive:

Description of Working Group:

The Inter Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Ongoing	Gain experience with the prototype in "real networks".
TBD	Develop gated version.
TBD	Add a small set of additional features and submit protocol into IETF standards process.

Internet Drafts:

"An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-idpr-architecture-03.txt>

"Inter-Domain Policy Routing Protocol Specification and Usage: Version 1", 03/05/1991, M. Steenstrup <draft-ietf-idpr-specv1-00.txt, or .ps>

"Definitions of Managed Objects for the Inter-Domain Policy Routing Protocol (Version 1)", 07/22/1991, R.A. Woodburn <draft-ietf-idpr-mib-00.txt, .ps>

"Inter-Domain Policy Routing Configuration and Usage", 07/25/1991, H. Brown, M. Steenstrup <draft-ietf-idpr-configuration-00.txt>

Request For Comments:

RFC 1126	"Goals and functional requirements for inter-autonomous system routing"

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

IDPR Minutes

The Inter-Domain Policy Routing Working Group met for two sessions in Atlanta. As usual, we conducted a tutorial for the first half of the first session. We spent the rest of the time discussing two new IDPR documents recently released as Internet Drafts. The documents are a configuration guide and a MIB. As the documents became available shortly before the IETF meeting, we did not expect that people would have had time to read them. Instead, we gave an overview of the content of each document, and we raised the issues that are not yet closed.

The main open configuration issue is the following. We advocate using the existing domain name servers to resolve addresses to administrative domain identifiers. This feature is not yet available in the DNS but would not be difficult to add.

The open questions are:

- Is the DNS in fact the correct place for this functionality?
- Should the address input be at the level of network, subnet, or host?
- How do we handle the traversal of the DNS hierarchy?

There are two open MIB issues.

1. What information to put in the MIB.

   IDPR has more configuration parameters than most routing protocols and more potential quantities to measure, because of policy support. The MIB contains a first cut at what information we expect that people will want to measure. However, some of the information is of general use, while other pieces of information are of most use to experts who know the details of the protocols. We want to make sure that we provide enough information in the MIB to allow problem detection, but we don't want to overload the MIB with information only useful for detailed fault detection and correction.

2. How to represent policy information.

   IDPR allows a domain administrator to configure transit policies that apply to traffic traversing its domain and source policies that apply to traffic generated by hosts in its domain bound for hosts in other domains. The current draft of the MIB does not contain policy information, because we had not at that time decided on how best to represent this information. At the Working Group meeting, we tentatively decided on an approach, and the next version of the MIB will contain this information.

We are seeking constructive criticism on both the MIB and configuration guide documents. Please send all comments to idpr-wg@bbn.com.

Attendees

Nagaraj Arunkumar	nak@3com.com
James Beers	
Helen Bowns	hbowns@bbn.com
John Chang	jrc@uswest.com
George Clapp	meritec!clapp@uunet.uu.net
Richard Cogger	
Dimitry Haskin	dhaskin@bbn.com
Tony Li	tli@cisco.com
Keith McCloghrie	kzm@hls.com
Paul Mockapetris	pvm@darpa.mil
Andy Nicholson	droid@cray.com
Chandy Nilakantan	csn@3com.com
Yakov Rekhter	yakov@ibm.com
Mark Sleeper	mws@sparta.com
Martha Steenstrup	msteenst@bbn.com
Robert Woodburn	woody@cseic.saic.com
L. Michele Wright	uncng!michele@uunet.uu.net
Osmund de Souza	desouza@osdpc.ho.att.com

3.6.5 Multicast Extensions to OSPF (mospf)

Charter

Chair(s):

Steve Deering, deering@xerox.com

Mailing Lists:

General Discussion: mospf@devvax.tn.cornell.edu
To Subscribe: mospf-request@devvax.tn.cornell.edu
Archive:

Description of Working Group:

This Working Group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done — Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.

Done — Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.

Done — We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using the existing BSD OSPF code, written by Rob Coltun of the University of Maryland, as an example.

Done — Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations. The specification should now be ready to submit as an RFC.

Internet Drafts:

"Multicast Extensions to OSPF", 07/25/1991, J. Moy <draft-ietf-mospf-multicast-00.ps>

CURRENT MEETING REPORT

Reported by Dino Farinacci/cisco Systems

MOSPF Minutes

The Multicast Extensions to OSPF Working Group met Thursday, August 1, at the Atlanta IETF. The meeting Agenda was:

- Introductions, roster, assign note taker.
- Progress of draft.
- Review of July draft.
- Progress of implementation.
- OSIfication of specification.
- Multicast scoping.

Progress of Draft.

John Moy (author of spec), felt the specification was ready for protocol review but not quite ready to implement from. The following are missing sections:

- System requirements. (i.e., manipulating multicast filters in media controllers).
- Data structures.
- Detailed description for forwarding multicast datagrams.
- Detailed description for forwarding cache construction.
- Additions to base OSPF specification.
  - How to flood group-membership LSAs.
  - Database Description process with new option bits.
  - Generation of Summary LSAs.

Review of July Draft.

Changes from previous draft (distributed at St. Louis IETF) included:

- Added concept of a "Inter-area multicast forwarder". The reason for this is that all ABRs do not need to forward multicast traffic.
- Inter-AS multicast section is new.
- Requirements on Inter-AS protocol are identified. Must determine if a multicast datagram came from outside of AS or from its own AS going outside.

- Reverse costs will be used for inter-area and inter-AS multicasting. John Moy described a scenario where problems will occur when reverse costs are used (in the case of the above bullet) in combination with forward costs used inside an area. It was determined that the problem can be corrected if reverse costs are used everywhere.

Progress of Implementations.

Initially there will be two implementations from Proteon and public domain source gated (Cornell).

Steve Deering, from Xerox, is working on modifications to the BSD Unix kernel for forwarding multicast datagrams. He briefly described the data structure for the forwarding cache. The key to access the cache is based on the tuple (source network, destination group, TOS).

OSIfication of Specification.

Steve Deering presented IP multicasting to the ANSI X3S3.3 committee. The group is interested to use this research for OSI multicasting but no work has begun.

Multicast Scoping.

- TTL is used in multicasting to limit scope. The problem with TTL usage is that it does not take into account administrative boundaries.
- Scope should be based on the group address.
- It was noted that the first 256 assigned multicast addresses are meant for local wire only.
- Scott Brim, from Cornell, has proposed to provide scoping in IGMP.
- Scoping boundaries should have multiple levels. For example, (Site, AS, Country, Continent).

Attendees

Jim Beers	beers@nr-tech.cit.cornell.edu
Scott Brim	swb@nr-tech.cit.cornell.edu
Chi Chu	chi@sparta.com
Richard Colella	colella@osi3.ncsl.nist.gov
Steve Deering	deering@xerox.com
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@canet.ca
Robert Griffioen	rgriff@bnr.ca
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Phani Jujjavarapu	phani@cisco.com

April Merrill	abmerri@tycho.ncsc.mil
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Bob Stewart	rlstewart@eng.xyplex.com
David Waitzman	djw@bbn.com
Robert Woodburn	woody@cseic.saic.com
L. Michele Wright	uncng!michele@uunet.uu.net
Osmund de Souza	desouza@osdpc.ho.att.com

3.6.6 Open Shortest Path First IGP (ospf)

Charter

Chair(s):

Mike Petry, petry@ni.umd.edu
John Moy, jmoy@proteon.com

Mailing Lists:

General Discussion: ospfigp@trantor.umd.edu
To Subscribe: ospfigp-request@trantor.umd.edu
Archive:

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

Done	Design the routing protocol, and write its specification.
Done	Develop multiple implementations, and test against each other.
Done	Obtain performance data for the protocol.
Done	Make changes to the specification (if necessary) and publish the protocol as a Draft Standard RFC.
TBD	Gather operational experience with the OSPF protocol and submit the

Internet Drafts:

"OSPF Version 2 Traps", 07/23/1991, Rob Coltun <draft-ietf-ospf-trapmib-00.txt>

Request For Comments:

RFC 1131	"OSPF specification"
RFC 1245	"OSPF Protocol Analysis"
RFC 1246	"Experience with the OSPF Protocol"
RFC 1247	"OSPF Version 2"

RFC 1248 "OSPF Version 2 Management Information Base"

RFC 1252 "OSPF Version 2 Management Information Base"

RFC 1253 "OSPF Version 2 Management Information Base"

CURRENT MEETING REPORT

Reported by John Moy/Proteon

OSPF Minutes

The OSPF Working Group met on Monday, July 29th at the Atlanta IETF. The following topics were discussed:

OSPF trap MIB

Rob Coltun led a discussion of his OSPF Trap MIB document. Briefly, this document began as a list of traps to help implementors debug their code. But Rob has now changed it to include only those traps that would be of use to a network manager. This has decreased the size of the Trap MIB significantly. Also, the Trap MIB was separated from the rest of the OSPF MIB in order to smooth the OSPF MIB's standardization path (traps are still somewhat controversial).

It was decided that the following changes should be made to the Trap MIB document. After it is updated, the document will then be submitted to the Network Management Directorate.

- The ospfIfStateChange and ospfVirtIfStateChange traps will only occur when either a) the new interface state is one of the terminal states ("DR", "Backup" or "Other") or b) the interface state regresses (e.g., goes from "DR" to "Down").
- The ospfNbrStateChange and ospfVirtNbrStateChange traps will only occur when either a) the new neighbor state is one of the terminal states ("2-Way" or "Full") or b) the neighbor state regresses (e.g., goes from "2-Way" to "1-Way").
- A new authentication failure trap will be created, splitting off from the existing ospfConfigError trap. This is because for net- works that are on the boundary of two ASes, authentication failures may be configured intentionally in order to separate two OSPF domains.
- The reasons for the ospfRxBadPacket trap will be enumerated, just as the is currently done for the ospfConfigError trap.
- The ospfOriginateLSA trap will not be invoked for simple refreshes of LSAs (which happen every 30 minutes), but instead will only be invoked when an LSA is (re)originated due to topology change.
- The ospfMaxAgeLSA trap will only be invoked for those LSAs that the router itself ages to MaxAge (either normally or prematurely). It will not be invoked when the

router receives a MaxAge advertisement from a neighbor.

- The ospfFreeLSA trap will be removed, since its functionality is (pretty much) identical to that of the ospfMaxAgeLSA trap.

Not so stubby areas

Next, Vince Fuller led a discussion of the proposed new "not so stubby area" option for OSPF. Briefly, the intent of this option is to create a new type of area (the "not so stubby area" or NSSA), which would not receive type 5 external LSAs from the backbone (and so would have a small database size), but would be allowed to itself originate a small number of external advertisements for distribution to the backbone. This would allow small RIP clouds to be hung off of the NSSAs.

Vince Fuller and Rob Coltun are writing a document defining NSSAs. The intent is to add them in a backward compatible way to OSPF Version 2.

As far as mechanisms for implementing NSSAs, there were two competing proposals, each differing on how the NSSA would represent the externals it exports to the backbone.

- Option 1.

  The externals would be originated as regular type 5 LSAs. Flooding of type 5s between an NSSA and the backbone is unidirectional. Type 5s can be flooded from NSSAs to the backbone, but not vice versa.

  Advantages: 1) No conversion of the LSA would be necessary at the area border routers connecting the NSSA to the backbone.

  Disadvantages: 1) There would no longer be a global type 5 database. In fact, the type 5 database in the area border routers connecting the NSSAs to the backbone would be split into several pieces: one for each NSSA, and one for those type 5s originated by the backbone. Maintaining this split may prove difficult.

- Option 2.

  The externals would be originated as a new LSA type (call it type 6 LSAs). The flooding of type 6 LSAs would be restricted to a single NSSA. The area border routers connecting the NSSA to the backbone would, in essence, convert the type 6 to a type 5 for distribution to the backbone.

  Advantages: 1) the type 5 database remains intact. In addition, the flooding of type

6s is similar to the flooding of all other LSAs that are specific to a single area (i.e., type 1-4s).

Disadvantages: 1) The "conversion" of type 6s to type 5s in the border routers may be fairly complicated (editor's note: at lunch Rob Coltun pointed out that the conversion could be made trivial by requiring that all type 6s be originated with forwarding addresses). 2) When there are multiple area border routers connect- ing the NSSA to the backbone, multiple type 5 LSAs may be produced for a single type 6 LSA. It was thought that this could be overcome by an election algorithm, if desired.

Vince and Rob are going to further weigh the two approaches and then document their conclusions.

Non-broadcast Networks

Several people have noticed that OSPF's non-broadcast support could be made more robust in the face of misconfiguration, and that the amount of configuration (especially address translations) could be reduced by using some of the mechanisms in Paul Tsuchiya's SMDS routing and addressing internet draft (like ARP servers). We attempted to find some- one to write a document discussing these issues, but have as yet been unsuccessful.

Joint Session with BGP Working Group

We also met in a joint session with the BGP Working Group, where we reviewed Kannan Varadhan's document on BGP and OSPF interaction. Kannan's document give rules for exporting OSPF routes to BGP, importing BGP routes into OSPF, and defines how to set the OSPF external route tag in a vendor-independent manner. It also mandates that the OSPF router ID and the BGP router ID be set identically, and explains the circumstances where OSPF and BGP forwarding addresses should be used.

Attendees

Nagaraj Arunkumar	nak@3com.com
Jim Beers	beers@nr-tech.cit.cornell.edu
Tom Benkart	teb@saturn.acc.com
Nelson Bolyard	nelson@sqi.com
Scott Brim	swb@nr-tech.cit.cornell.edu
Henry Clark	henryc@oar.net
Rob Coltun	rcoltun@ni.umd.edu
Dave Cullerot	cullerot@ctron.com
Dino Farinacci	dino@cisco.com
Dennis Ferguson	dennis@canet.ca
Arlan Finestead	arlanf@ncsa.uiuc.edu
Vince Fuller	vaf@stanford.edu

Robert Griffioen	
Jeffrey Honig	jch@nr-tech.cit.cornell.edu
Phani Jujjavarapu	phani@cisco.com
Mark Lewis	mlewis@telebit.com
Joshua Littlefield	josh@cayman.com
Brian Lloyd	brian@telebit.com
April Merrill	Merrill@dockmaster.ncsc.mil
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Andy Nicholson	droid@cray.com
Karen O'Donoghue	kodonog@relay.nswc.navy.mil
Norman Patty	
Joe Peck	peck@ms1.pa.dec.com
Jason Perreault	perreaul@interlan.interlan.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Mike Truskowski	truskowski@cisco.com
Brian Wheeler	wheeler@mbunix.mitre.org

3.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker

One of the themes in the Security Area is the interaction with other areas. There is a high level of activity, which is not reflected in the number of security related Working Groups. These activities include liasons between the SNMP Security, Trusted NFS, Telnet, and Point to Point Working Groups among others.

Security Guidelines

The Security Guidelines and the Site Security Handbook went to completion. A few little details to wrap up on the guidelines, but these are major accomplishments. Thanks to Joyce Reynolds and Paul Holbrook, as well as Rich Pethia, Barbara Fraser and Paul Holbrook.

Commercial IP Security Option

The CIPSO Working Group comes to us by the courtesy of TSIG. TSIG reports that it was extremely useful to meet in concert with the IETF, and that this provided access to people and expertise that they have not had access to before, particularly router vendors. This interaction was stimulating but in some respects difficult to deal with the influx of new ideas and points of view. It looks like the joint meeting was an overall success. Nobody is throwing tomatoes just yet.

Common Authentication Technology

The Common Authentication Technology Working Group led by John Linn had a successful meeting. The Working Group examined the protocols itself in the morning, and in an afternoon session, met with a number of Working Groups to see if that technology is applicable. Things are rolling right along, however, this is work that will take considerable time. There are several parts to the effort. There are several handbook Internet drafts that are out, including Kerberos and the GSSAPI specifications and we're expecting FTX to follow along, probably in the November timeframe. The Privacy Enhanced Mail is a quite mature, fully fleshed out protocol developed on the research side of our family and it moved into the IETF here for the first time in a formal way.

A number of larger issues about the fundamental basis of what we're trying to do which involved not only technical issues, but social infrastructure issues and how to establish high levels of trust throughout the community and so forth. In the words of international diplomacy, frank and useful discussions were held. It is encouraging that we have a lot of the right issues out on the table and in relatively short order should have the basis of consensus though. Incidentally, security is converging rapidly. There is a very strong orientation toward the SNMP community.

So my understanding is that they would like to have SNMP authentication working and tested even before progressing to proposed standard stage or having an interim draft in some cases. There are a small number of technical things related to access control via MIB views with respect to tables.

The CAT Working Group is providing active assistance to a number of other Working Groups. The Telnet folks have an authentication option and the consensus there appears to be that they'd like to gain some experience with it. The idea is instead of trying to go directly from specification to proposed standards, status, the specification will be used as the basis for experimentation.

The PPP Extension Working Group is considering a number of things including authentication and they're rapidly converging and should have an Internet draft out shortly. While drafting the requirements has been a lengthy process with a large, large document, they seem to be converging on a next Internet draft and hope to push it forward.

The CAT Working Group agreed to provide support to the Call Accounting Working Group. One need that is critical is the notion of distinguished names or a common naming system across the Internet. This issue is larger than any particular area, particularly larger than the Security Area. The IESG has committed to work toward a system to fulfill this need.

3.7.1 Commercial Internet Protocol Security Option (cipso)

Charter

Chair(s):
Ron Sharp, rls@neptune.att.com

Mailing Lists:
General Discussion: cipso@wdl1.wdl.loral.com
To Subscribe: cipso-request@wdl1.wdl.loral.com
Archive: archive-server@wdl1.wdl.loral.com

Description of Working Group:

The Commercial Internet Protocol Security Option (CIPSO) Working Group is chartered to define an IP security option that can be used to pass security information within and between security domains. This new security option will be modular in design to provide developers with a single software environment which can support multiple security domains.

The CIPSO protocol will support a large number of security domains. New security domains will be registered with the Internet Assigned Numbers Authority (IANA) and will be available with minimal difficulty to all parties.

There is currently in progress another IP security option referred to as IPSO (RFC 1108). IPSO is designed to support the security labels used by the U.S. Dept of Defense. CIPSO will be designed to provide labeling for the commercial, U.S. civilian and non-U.S. communities.

The Trusted Systems Interoperability Group (TSIG) has developed a document which defines a structure for the proposed CIPSO option. The Working Group will use this document as a foundation for developing an IETF CIPSO specification.

Goals and Milestones:

Done	Review and approve the Charter for the IETF CIPSO Working Group. Review revised TSIG CIPSO Specification.
Done	Review outstanding comments/issues from mailing list. Continue work on specification and prepare it for submission as an Internet Draft by the end of May.
Jul 1991	Review outstanding comments/issues from mailing list. The specification will be submitted to the IESG for consideration as a Proposed Standard.
Mar 1992	Submit specification to the IESG for consideration as a Draft Standard. There must be at least two interoperable implementations by this time.
Ongoing	Review outstanding comments/issues from mailing list. Continue the process to advance the Draft Standard to a Standard.

CURRENT MEETING REPORT

Reported by Ron Sharp/AT&T

CIPSO Minutes

Here are the Minutes for the Commercial Internet Protocol Security Option Working Group meeting held in Atlanta, Georgia. Most of these Minutes were provided by Noel Nazario who recorded them and sent me an electronic version. Thanks again Noel. The Working Group met for 1.5 days with the first half day being spent on new issues described below. The second day we addressed and closed nearly all of the old issues.

Quick CIPSO Summary

Option type 134. One option per packet. Only sensitivity tags are currently defined in the document. The document provides a common format and minimum configuration parameters required for interoperability. Interpretation of values within the option are DOI-dependent (Domain of Interpretation).

Description of Open Issues

1. Clarifications to the Spec
2. Multiple DOI's
3. Sort Categories
4. Policy on unrecognized tags
5. Exclusionary tag types
6. Multiple sensitivity tags
7. Minimum RFC Compliance
8. Tag alignment
9. Change exclusionary tag type number
10. Error condition definition
11. Configuration parameters
12. New tag types
13. Tags 128-255, Vendor/DOI defined
14. Move security level out of tags
15. Vendor tag types - router problem
16. 8-bit security level and 2-bit errors
17. Header space
18. Should DOI #3 be included in spec for testing
19. Canonicalization of encodings
20. Whether to allow category 0
21. Routing based on nationality caveats

New issues

Mike St. Johns proposed the following change to the CIPSO format:

```
        8          8                16 bit
|----------------------------------------------|
|  Type     | Length  |          Level         |
|----------------------------------------------|
|                  DOI                         |
|----------------------------------------------|
|  Cat ID   |             Categories           |
|----------------------------------------------|
| (8 bits)  |             (24 bits)            |
|----------------------------------------------|
|   255     |                                  |  <-- Sys Hi
|----------------------------------------------|
```

This format will effectively eliminate all tags. Its basic merit is that it is simpler for routers to handle. The 16-bit security level is to be encoded for certain Hamming distance and not open to definition of 2 to the 16th levels. It would be easy for a router to implement. No flexibility in specifying different security policies.

It was voted 9:4 to continue the discussion of the other issues and table this proposal and discuss it more electronically before the next meeting.

Discuss and Close Issues

- Issue 1: Spec Clarifications

  Section 4, eliminate the non-security IP options from the document. This section will be replaced with a reference to the Internet Assigned Numbers Administration (IANA).

  Note that the length fields and tag numbers in the document should be reviewed and corrected.

  The requirement to transmit between DOI should be done by IP gateways and not by routers. Pro: 11, Con:0

  Clarify the concept of classes of tags (i.e., sensitivity tags). Pro: 12, Con: 0.

- Issue 3: Sort Categories

  Categories enumerated in type-2 tags should appear in ascending order. Pro: 12, Con: 0

- Issue 5: Exclusionary tag type change

  Eliminate the exclusionary flag from the enumerated tag in favor of adding a new range tag type with lower and upper bounds. Pro: 14, Con: 0

  Redesigning Tag Type 2

```
      8           8           8
|----------------------------------------------------------|
|   Type    |  Length  |    Level    |  Enumerated Categories |
|----------------------------------------------------------|

It was voted to eliminate the flags field from the Tag Type
2 format.  Pro: 13, Con: 1

Design Tag Type 5  (Range type)

      8       8       8      16    16
|---------------------------------------------------|
| Type= 5 | Len | Level | lh | ll |    |...|  | nh | nl |
|---------------------------------------------------|

                                          Optional, 0
                                           assumed if
                                              missing
```

  nh is the range high value and nl is the range low value. ranges are sorted high to low non-overlapping.

  The use of this format for the new Tag Type 5 was voted Pro:11, Con: 2.

- Issue 8: Tag alignment

  Compliant implementations will support unaligned tags. Pro: 12, Con: 0

- Issue 13: Tag types 128-255: Vendor/DOI assigned

  It was agreed that these tag types would be defined by the DOI and not the vendor. For example, a vendor should not implement a new tag format with a hard coded type number >127 unless the implementation is solely for a particular DOI.

  The need of these tag types was questioned during the meeting. There was concern that allowing users to define their own label formats could lead to non-interoperability issues later. It was decided to table the discussion until we had more time to look at the problem.

- Issue 2: Multiple DOI's

  2a. Implementations must support one DOI and may/should support more than one. Pro: 13, Con: 0

  2b. Recommend to administrators that a common DOI should be understood by all hosts on a subnetwork. Pro: 12, Con 0

  2c. For outgoing communications the DOI is selected based on either the network interface that will be used or by the address of the destination. This insures that the DOI selected on outgoing packets is not just a host level default but can be configured based on either the network interface (i.e., network default DOI) or by the destination host address. If it is on the destination host address then a DOI could be configured for the destination IP subnetwork or the host itself. Pro: 8, Con: 0

- Issue 18: Reserve DOI number 3 for use in testing

  Not necessary to include in the RFC. DOI's are configurable. Pro: 12, Con:0

- Issue 6: Multiple sensitivity tags

  Only one sensitivity tag included in a CIPSO option. Pro: 8, Con: 2.

- Issue 19: Canonicalization

  Require a common deterministic algorithm in CIPSO implementations where each label can be represented by only 1 sensitivity tag type.

  - Increases speed of equality check
  - Possible algorithms
    - Ascending order of tag numbers (1 first then 2 . . .)
    - Minimize space (use tag that requires least bytes)
  - Possible loss in speed due to time required for new algorithm
  - Goes against concept of maximize what you will accept

  Tabled

- Issue 4: What to do on unrecognized tag types

  The behavior on unrecognized tag types should be configurable with the default being not to accept the packet. Pro: 8, Con: 4.

  Make configuration optional, the default being to generate an error on unrecognized tags. Pro: 10, Con: 1

Issues not discussed due to time constraints

1. Minimum RFC compliance
2. Error conditions and responses. Brian Yasaki and Debbie Futcher wrote up a section for error conditions. A new copy will be sent out the mailing list and comments should be placed back on the mailing list.
3. Configuration parameters. Minimal configuration parameters required for each CIPSO host. Some of these changed with the issues discussed and a new version will be put on the mailing list before the next meeting.

Conclusion

As you can see we made it through a lot of issues and closed most of them. This is great. Part of the reason many were closed so quickly is that we discussed them at the previous CIPSO IETF meeting but we agreed to hold them open until this meeting.

A request was made for a new CIPSO IETF editor. Mark Powers has been doing a great job but due to job requirements he has not been able to make many of the meetings. Mark Christenson from Cray volunteered to do the work. I will make the appropriate changes to the document and submit it to Internet-Drafts and then give the document to Mark. Thanks Mark.

The next meeting of the CIPSO IETF will be at the HP complex in Cupertino, CA September 24th - 26th of this year. It will be held in conjunction with the TSIG meeting. It will start around 9am on the 24th and will have a morning wrap-up on the 26th ending around noon. I will send out more details as I get them.

If you have comments on any of the issues discussed please put these comments out on this mailing list now. Do not wait until the next meeting. If you have any new issues then again please bring them out now. New issues brought up at the next meeting will take a lower priority to existing issues that have had a chance to be digested and discussed on the mailing list.

Attendees

Richard Balcon	rbalcon@oracle.com
Tom Benkart	teb@saturn.acc.com
Nelson Bolyard	nelson@sqi.com
Doug Brown	cdbrown@sandia.gov
Mark Christenson	mgc@cray.com
Daylan Darby	daylan@ssc_bee.boeing.com
Deborah Futcher	dfutche@eco.twg.com
Amir Hudda	amir@sware.com

Barry Miracle	miracle@sctc.com
Noel Nazario	nazario@csdserv.ncsl.nist.gov
Oscar Newkerk	newkerk@decwet.enet.dec.com
Jackie Proveaux	jackie@csd.harris.com
John Seligson	johns@ultra.com
Ron Sharp	rls@neptune.att.com
Rick Slade	ricks@csd.harris.com
Richard Smith	smiddy@pluto.dss.com
Frank Solensky	solensky@clearpoint.com
Michael St. Johns	stjohns@umd5.umd.edu
Emil Sturniolo	emil@doe.dss.com
Bruce Taber	taber@interlan.com
Dean Throop	throop@dg-rtp.dg.com
Gerard White	ger@concord.com
Brian Yasaki	bky@eco.twg.com

3.7.2 Common Authentication Technology (cat)

Charter

Chair(s):

John Linn, linn@zendia.enet.dec.com

Mailing Lists:

General Discussion: cat-ietf@mit.edu
To Subscribe: cat-ietf-request@mit.edu
Archive: /cat-ietf/archive@bitsy.mit.edu

Description of Working Group:

The goal of the Common Authentication Technology working group is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security-oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accomodating evolution of underlying technologies.

In support of these goals, the Working Group will pursue several interrelated tasks. We will work towards agreement on a common service interface allowing callers to invoke security services, and towards agreement on a common authentication token format, incorporating means to identify the mechanism type in conjunction with which authentication data elements should be interpreted. The CAT Working Group will also work towards agreements on suitable underlying mechanisms to implement security functions; two candidate architectures (Kerberos V5, based on secret-key technology and contributed by MIT, and X.509-based public-key Distributed Authentication Services being prepared for contribution by DEC) are under current consideration. The CAT Working Group will consult with other IETF Working Groups responsible for candidate caller protocols, pursuing and supporting design refinements as appropriate.

Goals and Milestones:

Done — Preliminary BOF session at IETF meeting, discussions with Telnet and Network Printing Working Groups.

Done — Distribute Generic Security Service Application Program Interface (GSS-API) documentation through Internet-Draft process.

Done First IETF meeting as full Working Group: review Charter, distribute documents, and status of related implementation, integration, and consulting liaison activities. Schedule follow-on tasks, including documentation plan for specific CAT-supporting security mechanisms.

Oct 1991 Update mechanism-independent Internet-Drafts in response to issues raised, distribute additional mechanism-specific documentation including Distributed Authentication Services architectural description and terms/conditions for use of the technology documented therein.

Nov 1991 Second IETF meeting: Review distributed documents and status of related activities, continue consulting liaisons. Discuss features and characteristics of underlying mechanisms. Define scope and schedule for follow-on work.

Dec 1991 Submit service interface specification to RFC standards track.

Ongoing Progress Internet Draft and RFC publication of mechanism-level documents to support independent, interoperable implementations of CAT-supporting mechanisms.

Internet Drafts:

"Generic Security Service Application Program Interface", 06/12/1991, John Linn <draft-ietf-cat-genericsec-00.txt, .ps>

"The Kerberos Network Authentication Service", 07/01/1991, John Kohl, B. Clifford Neuman <draft-ietf-cat-kerberos-00.txt, .ps>

CURRENT MEETING REPORT

Reported by John Linn/DEC

CAT Minutes

The Common Authentication Technology Working Group met for two sessions in Atlanta, and held discussions building on the three Internet Drafts issued on behalf of the group in advance of the meeting. John Linn led a discussion on CAT and GSS-API concepts, and Jeff Schiller and Charlie Kaufman gave presentations on implementations of CAT atop (respectively) Kerberos V5 and SPX mechanisms; slides from these presentations will be submitted along with these minutes for inclusion in the IETF Proceedings. Representatives from some protocol Working Groups were available to comment on issues related to integration and use of CAT within their protocols.

CAT concepts were generally well received. Some areas of potential refinement and discussion were raised, and discussions are expected to continue on the CAT mailing list. One key area of technical discussion was the interrelationship among CAT, underlying mechanisms, and alternative naming architectures; a related area was alternative types of authenticated principals (users, hosts, processes) and means for their distinction. It was noted that the fact of implementation of a particular mechanism in support of CAT should not be taken as IETF endorsement of the strength of that mechanism. It was also noted that multiple mechanisms may in principle be incorporated beneath a single GSS-layer implementation, though no such implementations have yet been developed.

Identification of Shared Mechanism

One major discussion topic was the question of how to identify a CAT mechanism which is shared with a peer CAT system. Options include combinations of negotiation, directory entries, configuration data, and user/caller input; it was agreed that CAT should seek to make suitable determinations internally where possible so as to ease burdens on its callers and to avoid replicating common security-oriented features separately within a variety of caller protocols. This implies, for example, that CAT callers' requests for the "default" mechanism type could result in exchange of tokens in order to resolve a common mechanism; the feasibility of such a scheme warrants investigation. Whenever negotiation is used to establish a mechanism, it should be carried out against an acceptable set defined by configuration data and/or caller input, to prevent blind acceptance of authentication schemes weaker than those intended by a CAT peer.

Naming Issues

As the Internet evolves to a multi-protocol environment, it also evolves to an environment where multiple naming architectures must coexist. Prominent examples include DNS names for hosts, mailbox identifiers for users, and X.500 Distinguished Names. This variation causes problems in many areas of technology (and is engendering discussion in several parts of the IETF and the TSIG, as well as other groups), and security is among those bitten.

Since authentication mechanisms typically authenticate principals in conjunction with name forms native to those mechanisms, mismatches are likely to emerge when CAT callers oriented to operation in particular naming environments are served by CAT mechanisms employing different native forms. It was agreed that CAT would benefit from broader IETF-defined approaches to handle such mismatches; in the interim, mechanism designers will have to anticipate, observe, and provide case-by-case resolutions to specific problems. In the interests of portability between alternative mechanisms both capable of authenticating a common name format, it was observed to be preferable for identification of the mechanism used to authenticate a name to be carried in a separate parameter rather than being encoded within the name itself.

Mechanism Discussions

(See also presentation slides.)

Jeff Schiller led a discussion on Kerberos GSS-API implementation. MIT believes that it is appropriate for all services which run as root on a given host to use a common set of verifier credentials in /etc/srvtab; the Athena DISCUSS service has a different identity with credentials in a different file. Distinction between client and server principals is made based on examination of names.

Jeff also observed that MIT intends to relinquish control of the Kerberos V5 specification (distributed to Internet-Drafts before the meeting) to the CAT Working Group for evolution and standards-track progression, and cited Ted Tso and Cliff Neuman as additional relevant contacts. A Kerberos V4 specification will also be submitted as an informational RFC.

Charlie Kaufman led a discussion on SPX GSS-API implementation, emphasizing implementors' agreements made in order to enable application portability (though not the broader issue of interoperability) between Kerberos and SPX. Internal names were accepted to be opaque (preserving flexibility for mechanism implementors), although use of a standardized format at this level could offer value if callers were positioned to use the same format across other interfaces besides the GSS-API. The target applications chosen to validate the portability concept were Telnet and rlogin; since DNS-style textual names are native to these applications, conflicts with SPX's use and certification of X.500 DNs needed to be resolved.

Protocol Integration Issues

It was observed that error cases resulting from inability to process a transferred and received token cannot always be reflected to a CAT peer before that peer believes that the context establishment sequence is complete; for CAT callers to be assured that their tokens have been successfully processed on receipt, mutual authentication must be performed. Error-indicating tokens received after context establishment is complete can still be processed, by being passed to a different primitive (process_`context_token`). It was observed that it might be preferable to incorporate more messages in mechanisms' context establishment sequences so that COMPLETE status is never returned before positive acknowledgment by the peer. No conclusive decision was made on this issue.

The Telnet Working Group plans to issue the Telnet authentication option as an experimental RFC; it was anticipated that migration to CAT as an additional Telnet-visible type (which would likely supplant other Telnet-visible type indicators over time) would be appropriate. Terminal servers cannot be assumed to maintain configuration data corresponding to arbitrary "walk-up" users, so raise special issues with regards to integration with user interfaces and CAT infrastructure.

The Network Printing Working Group is seeking to employ CAT. Discussion indicated that different types of authentication semantics (users, hosts, daemon processes) would be most appropriate in different circumstances; unfortunately, prioritized needs for the different alternatives were not available.

Possible CAT applications arise in the Network News Transport Protocol (NNTP). Primary requirement areas raised at the CAT meeting include host-granularity authentication for sessions between NNTP peers and user-granularity authentication for individuals associated with NNTP newsreaders. Ted Tso is engaging in additional discussion with the NNTP group regarding potential CAT usage.

The LIST group may wish to employ CAT-based authentication for those cases where list maintenance commands are transferred across on-line connections rather than within messages.

Possible Extension Areas

Various candidate CAT extension areas were discussed, and are likely to be discussed further on the CAT mailing list.

Means for provision of long-term signature capabilities were considered only briefly, in part because of unclear requirements for non-repudiation services outside the messaging paradigm. The following observations were noted:

1. Since such signatures are intended to be validatable over an extended period and by other than the single peer associated with a context, such extensions are not well suited to modeling via the Quality-of-Protection (QOP) parameters to existing GSS-API per-message protection primitives,

2. That alternative primitives might utilize common credentials, and

3. That long-term signature capabilities would not likely be portable to other than public-key mechanisms.

Interest was expressed in making the set of intermediary entities which had been involved in a CAT authentication visible to a caller, presumably by providing means to extract such a name list from a context's data structures. It was unclear whether callers would be likely to make use of such a list in a mechanism-independent manner.

We also discussed the idea of an overlay veneer ("init_sec_context_stream()") to provide CAT with a communications path over which to pass tokens rather than returning the tokens for caller manipulation and transfer, an extension facility which could simplify integration of CAT-based authentication into certain caller protocols. Such an overlay would be analogous to Kerberos's send_auth interface; follow-up mailing list discussion is anticipated.

David Bolen	
David Borman	dab@cray.com
Stephen Crocker	crocker@tis.com
Peter Deutsch	
James Ellis	jte@cert.sei.cmu.edu
Arlan Finestead	arlanf@ncsa.uiuc.edu
James Galvin	galvin@tis.com
Joe Godsil	jgodsil@ncsa.uiuc.edu
Russ Hobby	rdhobby@ucdavis.edu
Alton Hoover	
Ken Jones	konkord!ksj@uunet.uu.net
Charles Kaufman	kaufman@dsmail.enet.dec.com
Peter Kirstein	kirstein@cs.ucl.ac.uk
Dale Land	land@lanl.gov
Eliot Lear	lear@turbo.bio.net
John Linn	linn@zendia.enet.dec.com
Louis Mamakos	louie@ni.umd.edu
Ellen McDermott	emcd@osf.org
Glenn McGregor	ghm@merit.edu
Clifford Neuman	bcn@isi.edu
Oscar Newkerk	newkerk@decwet.enet.dec.com
Richard Parker	rp@mbunix.mitre.org
Geir Pedersen	geir.pedersen@use.uio.no
Mel Pleasant	pleasant@hardees.rutgers.edu
P. Rajaram	rajaram@sun.com
Michael Reilly	reilly@nsl.dec.com
Jan Michael Rynning	jmr@nada.kth.se
Jeffrey Schiller	jis@mit.edu
Robert Shirey	shirey@mitre.org
Mark Sleeper	mws@sparta.com
Mark Stein	marks@eng.sun.com
Brad Strand	bstrand@cray.com
Glenn Trewitt	trewitt@nsl.dec.com
Theodore Tso	
Preston Wilson	preston@i88.isc.com

Common Authentication Technology (CAT) Working Group Meeting

John Linn

Linn@zendia.enet.dec.com

IETF, Atlanta, July 1991

CAT Motivations: How Did We Get Here?

- Many protocol architects identifying security requirements
- Want portability among security mechanisms
- Want to satisfy corresponding requirements of different protocols in a uniform manner
- Want to separate security implementation from protocol integration tasks

CAT WG Charter Areas

- Cryptographic authentication technologies to support real-time, peer-peer ("on-line") protocols (Telnet, Network Printing, ...)
- Token representations and other protocol integration issues
- Interfaces (though novel IETF topic) and mechanisms to implement CAT services
- Consulting liaisons with integrators

CAT-Related Documentation

- Mechanism-independent interface, service definition, and top-level token wrapper format (I-D proposals exist)
- Mechanism-specific definitions to enable interoperable CAT implementations over individual mechanisms (I-Ds emerging)
- Protocol-specific procedures for integration of CAT tokens into caller protocols

July 1991 Agenda Topics

- Discussion of Internet-Draft interface definitions
- Discussion of implementation experience over alternate mechanisms
- Discussion of protocol integration topics
  - -- token representations
  - -- candidate protocols' concerns and issues

GSS-API Internet-Drafts

- Two Internet-Drafts propose interface and service definition
  - -- Base spec (draft-ietf-cat-genericsec-00) defines concepts and constructs in language-independent manner
  - -- C bindings spec (draft-ietf-cat-secservice-00) realizes GSS-API for C language environment

CAT Framework

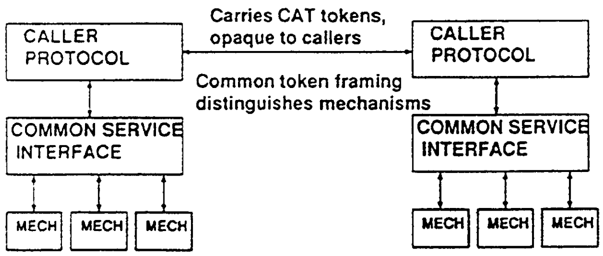

Mechanisms define interior token formats

GSS-API Security Contexts

- Security context establishment (Init_sec_context(), Accept_sec_context()) transfers context token(s), establishes short-term key(s)
  - -- achieves peer entity authentication (mutual, if requested)
  - -- optional delegation to target
  - -- contexts and service primitives logically independent of protocol associations

GSS-API Data Message Protection

- Context establishment creates shared keys
- Protection services (Sign(), Verify(), Seal(), Unseal()) apply keys to messages on contexts
  - -- security services offered: data origin authentication, integrity (per-message and stream-oriented), confidentiality
- Extensions to support stream orientation?

GSS-API Credentials

- Credentials (held locally, based on user login, server configuration, or incoming delegation)
  - -- provide prerequisites for principals to establish contexts under identities
  - -- can enable use of more than one mechanism
  - -- management characteristics (linkage with login, sharing semantics, selection of default) are system-dependent local matters

GSS-API Tokens

- Tokens (transferred between peers, in-band or out-of-band)
  - -- emitted and processed by GSS-API implementation, opaque from caller viewpoint
  - -- enclose mechanism-specific data elements
  - -- identify associated mechanism (in initial token)

GSS-API and Mechanisms

- Multiple mechanisms supported
  - -- Interoperation requires common type shared with peer
  - -- Overridable default type selection
- Can insulate callers from different mechanisms' needing different numbers of transactions to set up contexts, IF callers' protocols can transfer variable numbers of messages and tokens

Multi-Step Context Setup Example

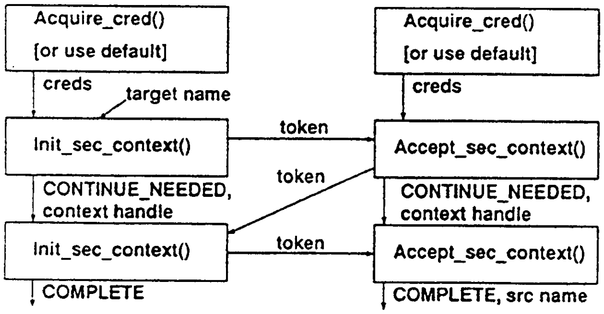

Naming Issues

- Need to select (set of) naming formats
  - -- What to name (and hence authenticate)? Users, hosts, services, ...
  - -- How to name? X.500 Distinguished Names vs. domain-qualified internet host specifiers?
- Need means to overcome mismatches between native forms of caller environments and those authenticated by mechanisms

GSS-API and Naming

- Support routines import to, display from, and compare internal names; goal is portability for "non-parsing" applications
- Internal forms opaque, mechanism-defined
- Default external printable syntax is local matter
- Mechanism implementors' agreements on supported (and OID-tagged) external string forms
  - -- e.g., `service:rlogin@bar.foo.com`

Channel Bindings

- Caller-invoked optional feature, to protect against substitution of intercepted tokens onto different associations
- Peers input intended channel description for validation against values (or signatures thereon) carried in tokens
- Content dependent on protocol environment: initiator and acceptor address types (or NULL indicators), addresses, application-specific data

Status Codes

- Major_status drives caller control flow, has mechanism-independent interpretation
  - -- generic function errors (8 bit numeric code)
  - -- binding-specific errors (8 bit numeric code)
  - -- supplemental status (16 bit mask)
- Minor_status is mechanism-defined

Other Issues

- Trust models have mechanism-specific aspects which can't be abstracted as generic
- Portability is good, but interoperability would be better
  - -- desirable to avoid proliferation of incompatible mechanisms
  - -- need means (negotiation, directory, ...) to identify common mechanism shared with peer

Token Representation

- To disambiguate tokens of multiple underlying mechanisms, need to wrap first context establishment token in self-descriptive wrapper:

  -- thisMech OBJECT IDENTIFIER

  -- innerContextToken ANY DEFINED BY thisMech

- Interior of this token, and contents of subsequent tokens, are mechanism-defined

IETF/CAT

GSSAPI Use with Kerberos

Jeffrey L Schiller

Manager System and Operations
MIT Network Services and Distributed Computing

IETF/CAT

GSSAPI Use with Kerberos

- Mapping of GSSAPI Operations and Kerberos
- Contents of GSSAPI Tokens
- Show which transactions are part of the GSSAPI mediated transaction and which are not

IETF/CAT

GSSAPI Authentication Steps

- gss_acquire_cred
  - Get a handle on your credentials
- gss_init_sec_context
  - Initiate Context with a Peer
- gss_accept_sec_context
  - Accept Context from a Peer

IETF/CAT

GSSAPI Authentication Steps

gss_acquire_cred

gss_init_sec_context → gss_accept_sec_context

IETF/CAT

GSSAPI Steps and Kerberos/Client

- gss_acquire_cred
  - Client Side: Noop
- gss_init_sec_context
  - Get a ticket for the service (either from cache or from Kerberos via TGT Service)
  - Build and Authenticator
  - Encapsulate the Ticket and Authenticator into the GSSAPI Token
  - User program gets token and sends it over wire
  - If MUTUAL authentication is requested, gss_init_sec_context will return gss_s_continue_needed

IETF/CAT

GSSAPI Step and Kerberos

- gss_acquire_cred
  - Server Side: Fetch the necessary keys from /etc/srvtab
- gss_accept_sec_context
  - Parse token received from client side
  - Return success or failure
  - Generate a new token if MUTUAL authentication is desired

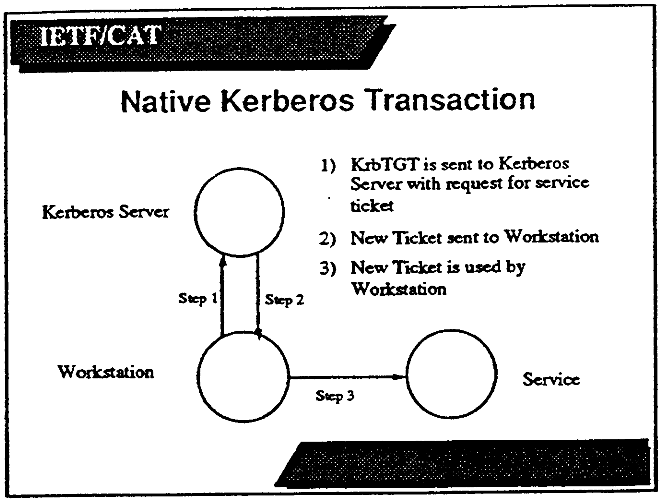
IETF/CAT
Native Kerberos Transaction
1) KrbTGT is sent to Kerberos Server with request for service ticket
2) New Ticket sent to Workstation
3) New Ticket is used by Workstation
Kerberos Server
Step 1
Step 2
Workstation
Step 3
Service

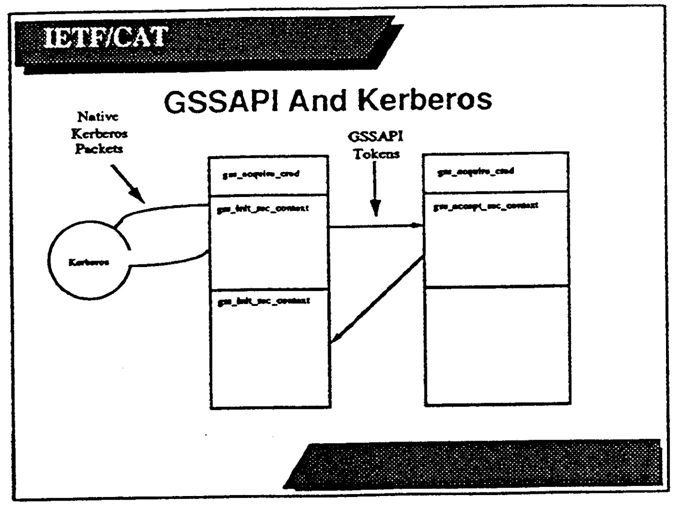
IETF/CAT
GSSAPI And Kerberos
Native Kerberos Packets
GSSAPI Tokens

GSSAPI Use with SPX

Charles W. Kaufman

Digital Equipment Corporation

IETF/CAT Atlanta July 1991

The Challenges:

- Function Mapping
- Missing Functions
- GSSAPI "wrapper"
- Implementers agreements to assure application portability with Kerberos V5.

Function Mapping

- SPX create_token created a token to pass and a mutual authentication token to compare; SPX accept_token accepted a token and produced a mutual authentication token. Each produced a shared key.
- The GSSAPI layer had to create a "security context" structure in which to store the mutual authentication token and the shared key.
- SPX had the concept of credential handles, but the types had to be mapped.

Missing Functions

- SPX has no mechanisms for signing and encrypting user data. It was envisioned for use over node to node encrypted network connections.
- By saving the shared key produced by create_ and accept_token in security context structures, we left a hook for adding this functionality later.

GSSAPI Wrapper

- GSSAPI requires that tokens contain information not present in the initial SPX design (e.g. name of principal, whether mutual authentication is desired); it also recommends a header to distinguish tokens from different mechanisms.
- We could have put the GSSAPI "wrapper" around the SPX token and the additional data, but chose to change the format of the SPX token to include it.

Implementers' Agreements

- The GSSAPI is sufficiently rich that putting a mechanism under it does not assure that applications will be portable.
  - -- GSSAPI supports an arbitrary number of messages in an authentication exchange; some protocols may not.
  - -- GSSAPI does not specify a syntax for names; mechanisms will have definite requirements.

- We worked with Project Athena on a syntax for server names that could be mapped to either Kerberos V5 or SPX:
  "**server:rlogind@nodename.whereever.edu**"
- User principal names do not match between Kerberos V5 and SPX, but applications can be portable by reading such names from configuration files (e.g. .rhosts) and using the **gss_compare_name** function.
- More work will be needed in the area of implementers' agreements to fully realize the portability potential of the GSSAPI.

3.7.3 Internet Security Policy (spwg)

Charter

Chair(s):

Richard Pethia, rdp@cert.sei.cmu.edu

Mailing Lists:

General Discussion: spwg@nri.reston.va.us
To Subscribe: spwg-request@nri.reston.va.us
Archive:

Description of Working Group:

The Security Policy Working Group (SPWG) is chartered to create a proposed Internet Security Policy for review, possible modification, and possible adoption by the Internet Activities Board. The SPWG will focus on both technical and administrative issues related to security, including integrity, authentication and confidentiality controls, and the administration of hosts and networks.

Among the issues to be considered in this Working Group are:

- Responsibilities and obligations of users, database administrators, host operators, and network managers.
- Technical controls which provide protection from disruption of service, unauthorized modification of data, unauthorized disclosure of information and unauthorized use of facilities.
- Organizational requirements for host, local network, regional network and backbone network operators.
- Incident handling procedures for various Internet components.

Goals and Milestones:

Done — Review and approve the Charter making any necessary changes. Begin work on a policy framework. Assign work on detailing issues for each level of the hierarchy with first draft outline.

Done — Revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.

Done — Prepare first draft policy recommendation for Working Group review andmodification.

Sep 1990 Finalize draft policy and initiate review following standard RFC procedure.

Internet Drafts:

"Guidelines for the Secure Operation of the Internet", 03/27/1991, Richard Pethia, Steve Crocker, Barbara Fraser <draft-ietf-spwg-secureop-01.txt>

3.7.4 Privacy-Enhanced Electronic Mail (pem)

Charter

Chair(s):
Stephen Kent, kent@bbn.com

Mailing Lists:
General Discussion: pem-dev@tis.com
To Subscribe: pem-dev-request@tis.com
Archive: pem-dev-request@tis.com

Description of Working Group:

PEM is the outgrowth of work by the Privacy and Security Research Group (PSRG) of the IRTF. At the heart of PEM is a set of procedures for transforming RFC 822 messages in such a fashion as to provide integrity, data origin authenticity, and optionally, confidentiality. PEM may be employed with either symmetric or asymmetric cryptographic key distribution mechanisms. Because the asymmetric (public-key) mechanisms are better suited to the large scale, heterogeneously administered environment characteristic of the Internet, to date only those mechanisms have been standardized. The standard form adopted by PEM is largely a profile of the CCITT X.509 (Directory Authentication Framework) recommendation.

PEM is defined by a series of documents. The first in the series defines the message processing procedures. The second defines the public-key certification system adopted for use with PEM. The third provides definitions and identifiers for various algorithms used by PEM. The fourth defines message formats and conventions for user registration, Certificate Revocation List (CRL) distribution, etc. (The first three of these were previously issued as RFCs 1113, 1114 and 1115. All documents have been revised and are being issed first as Internet Drafts.)

Goals and Milestones:

Done	Submit first, third, and fourth documents as internet drafts
Done	Submit second document as internet draft.
Done	First IETF working group meeting to review internet drafts.
Sep 1991	Submit revised internet drafts based on comments received during working group meeting, from pem-dev mailing list, etc.
Nov 1991	Submit internet drafts to IESG for consideration as Proposed Standards.

Ongoing Revise Proposed Standards and submit to IESG forconsideration as Draft Standard, and repeat for consideration as Internet Standard.

Internet Drafts:

"Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", 03/26/1991, John Linn <draft-ietf-pem-msgproc-01.txt>

"The MD5 Message-Digest Algorithm", 07/08/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md5-01.txt>

"The MD2 Message-Digest Algorithm", 07/10/1991, B. Kaliski <draft-rsadsi-kaliski-md2-00.txt>

"The MD4 Message-Digest Algorithm", 07/10/1991, R. Rivest, S. Dusse <draft-rsadsi-rivest-md4-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV: Notary, Co-Issuer, CRL-Storing and CRL-Retrieving Services", 07/10/1991, B. Kaliski <draft-ietf-pem-notary-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", 07/17/1991, Steve Kent <draft-ietf-pem-keymgmt-00.txt>

"Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", 08/22/1991, David Balenson <draft-ietf-pem-algorithms-00.txt>

CURRENT MEETING REPORT

Reported by Steve Kent/BBN

PEM Minutes

The Privacy Enhanced Mail (PEM) Working Group, Chaired by Steve Kent (BBN) met for the first time at the Atlanta IETF meeting. The Working Group met on Monday and Tuesday afternoon and both sessions were well attended. The purpose of these sessions was to discuss unresolved issues related to the Internet drafts which describe PEM processing and the proposed PEM key management system. The goal of this meeting was to reach consensus on these issues so that revised Internet Drafts could be issued and the Internet Drafts could progress to RFCs. A number of the issues which populated the meeting agenda were articulated in postings to the PEM mailing list in previous weeks.

During the first session progress was made in resolving an issue related to PEM message processing details and John Linn (the editor of the relevant Internet Draft) agreed to make the necessary edits to this Internet Draft.

There also appeared to be consensus on how to accommodate, in an orderly fashion, a growing list of distinguished name attributes in PEM certificates. The proposal was made to register such attributes, for use in all Internet applications which make use of distinguished name, in the periodic IANA RFCs.

Other issues regarding public key management for PEM were discussed including:

- The role of RSADSI in the certification process,
- Constraints on cross-certification, and
- Representation of object class identification in certificates.

However, consensus was not achieved in the discussion of these issues and other contentious topics associated with the key management Internet Draft were not addressed due to insufficient time over both days of the PEM Working Group sessions.

During the Tuesday afternoon session RSADSI presented proposed organization and user registration procedures and corresponding legal documents. Copies of the organizational agreement were not available for distribution at the time of the meeting. RSADSI also described the proposed agreement for distributing source code for the RSA cryptosystem as part of a freely available, Internet implementation of PEM. Because some aspects of this agreement are involved with unresolved issues as cited above, finalization of the agreement details, etc., will require resolution of these issues.

As a result of the lack of progress at these Working Group meetings, an interim meeting was proposed (later in the week) for September 11th at BBN in Cambridge, MA. This Working Group also expects to convene again in November at the next IETF meeting.

Attendees

Doug Barlow	barlow@decwet.dec.com
Jim Bidzos	jim@rsa.com
David Bolen	db31@nis.ans.net
Nelson Bolyard	nelson@sqi.com
Mats Brunell	mats.brunell@sics.se
Vinton Cerf	vcerf@nri.reston.va.us
Stephen Crocker	crocker@tis.com
Steve Dusse	spock@rsa.com
Jeffrey Edelheit	edelheit@mitre.org
James Ellis	jte@cert.sei.cmu.edu
Barbara Fraser	byf@cert.sei.cmu.edu
Deborah Futcher	dfutche@eco.twg.com
Maria Gallagher	maria@nsipo.arc.nasa.gov
James Galvin	galvin@tis.com
Joe Godsil	jgodsil@ncsa.uiuc.edu
William Griffeth	billg@usl.com
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Alton Hoover	hoover@nis.ans.net
Neil Katin	katin@eng.sun.com
Charles Kaufman	kaufman@dsmail.enet.dec.com
Stephen Kent	kent@bbn.com
Vincent Lau	vincent.lau@eng.sun.com
John Linn	linn@zendia.enet.dec.com
David Lippke	lippke@utdallas.edu
Jack Liu	liu@koala.enet.dec.com
Louis Mamakos	louie@ni.umd.edu
Ellen McDermott	emcd@osf.org
April Merrill	abmerri@tycho.ncsc.mil
Paul Mockapetris	pvm@darpa.mil
Clifford Neuman	bcn@isi.edu
Richard Parker	rp@mbunix.mitre.org
P. Rajaram	rajaram@sun.com
Michael Reilly	reilly@nsl.dec.com
Jeffrey Schiller	jis@mit.edu
Aaron Schuman	schuman@sgi.com
Ron Sharp	rls@neptune.att.com
Robert Shirey	shirey@mitre.org
Mark Stein	marks@eng.sun.com
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov
Theodore Tso	tytso@mit.edu
Brian Yasaki	bky@eco.twg.com
Peter Yee	yee@ames.arc.nasa.gov

3.7.5 SNMP Security (snmpsec)

Charter

Chair(s):

James Galvin, `galvin@tis.com`
Keith McCloghrie, `kzm@hls.com`

Mailing Lists:

General Discussion: `snmp-sec-dev@tis.com`
To Subscribe: `snmp-sec-dev-request@tis.com`
Archive: `snmp-sec-dev-request@tis.com`

Description of Working Group:

The SNMP Security Working Group is chartered to determine the set of security services needed by the SNMP. The specification of those services, the supporting mechanisms, and the adjunct infrastructure will become an enhancement to the SNMP and eventually an Internet standard.

The specification must not alter the fundamental SNMP network management philosophy and must not entail changes to existing SNMP standards or framework.

Goals and Milestones:

Done	Publish internet-draft specifications.
Jul 1991	Submit specification to IESG for consideration as a Proposed Standard.
Dec 1991	Submit specification to IESG for consideration as a Draft Standard.
Ongoing	Submit specification to IESG for consideration as a Standard.

Internet Drafts:

"SNMP Administrative Model", 04/09/1991, James Galvin <draft-ietf-snmpsec-admin-00.txt, .ps>

"Definitions of Managed Objects for Administration of SNMP Parties", 04/09/1991, Keith McCloghrie, James R. Davin, James M. Galvin <draft-ietf-snmpsec-mib-00.txt>

"SNMP Security Protocols", 04/09/1991, James M. Galvin, Keith McCloghrie, James R. Davin <draft-ietf-snmpsec-protocols-00.txt, .ps>

CURRENT MEETING REPORT

Reported by James Galvin/TIS and Keith McCloghrie/Hughes

SNMPSEC Minutes

Status of the Documents Reviewed:

- All three: the SNMP Administrative Framework, SNMP Security Protocols, and SNMP Party MIB, were published as Internet Drafts immediately after the previous IETF (in St. Louis).

- An update to the SNMP Party MIB was distributed to the snmp-sec-dev mailing-list at the beginning of July.

The Outstanding Issues were Discussed:

- Mike St.Johns suggested consideration of the use of "threshold keying", in the distribution of initial secrets. Threshold keying is a standard security technique (see Denning's book on Computer Security), in which the keys are split into multiple "shadow" parts. The parts could be distributed separately and then recombined to obtain the initial secret. Use of this technique would allow an administration to, for example, have a single shadow key which would be manually entered into each agent at install time, and another shadow key calculated by the nms so as to be agent-specific and distributed to the agent; these two parts could then be combined to get the initial secret. The advantages would be the ability to have the manually distributed secret information be a) the same for all agents, and b) different from the secret used as the initial key. The disadvantage being the special first-time-only processing the agent would need to recombine the keys. The meeting agreed to consider the suggestion in parallel with other activities.

- The differences between MD4 and MD5 were discussed, and the pros and cons of using each. A suggestion was made to update the text of the SNMP Security Protocols document to replace occurrences of "SNMP MD4 Authentication Protocol" by "SNMP Digest Authentication Protocol" in discussions of all parts of the protocol except the particular digest algorithm used, where the use of "MD4" would be retained. This suggestion was accepted since it would minimize the text (e.g. to one page) which would be needed in a future memo specifying alternative digest algorithms.

- A question on "wildcard" parties (analogous to the "public" community) was answered by discussing the "initial" noAuth,noPriv parties defined by convention in the Party MIB. A lively discussion ensued on the access rights to be afforded to this out-of-the-box noAuth,noPriv party. Some argued for allowing read-access to everything in the MIB (except SNMP security's secret information); others for allowing read-access to nothing, or just to MIB-II's system group. The consensus of the discussion seemed to be for this working group to stay silent on the issue, and let the various Requirements working groups make device-type specific recommendations. The Router Requirements WG. is making such a recommendation for use of "public" communities, and knows it will have to update that recommendation as and when the SNMP Security documents are further along.

- A discussion was held on the protocol's use of ASN.1 tags instead of a version number field. The same conclusion was reached as in previous discussions of the same topic.

- The term "random values" in the section of the SNMP Security Protocols document discussing what to do when an agent loses its knowledge of a secret, was clarified as being the need to set the values to non-valid or non-guessable values.

There was discussion of the implementation experience gained so far:

- Three separate implementations were in various stages of incompletion, and one other person had spent some preparing for an implementation. Two of these implementations interoperated with each other using noAuth,noPriv. Two had implemented MD4. One was using DES but was unsure that the encrypted data was correct. To date, there is no experience with multiple MIB views, proxy, clock synchronization, nor SNMP access to the Party MIB.

- A couple of ASN.1 definitions were discussed for possible optimizations:
  - The replacement of ANY by a CHOICE in types of AuthInformation,
  - The specification of a fixed length for the OCTET STRING containing the digest value, and
  - The rearrangement of the authentication information and the source/destination party fields leading to the removal of one of the levels of serialization.

There was also discussion of the present access-control granularity, and its ability to scale. The definition of MIB subviews does allow access control on individual instances, but at the cost of entering each object instance in the View Table. There is a legitimate requirement to support several Views each containing all the variables in, for example, the ifTable for just one interface. This requires a large number of entries in the View Table even with only a moderate numbers of interfaces.

The document editors agreed to update the documents to reflect the (minor) changes resulting from the above discussions. These updates are expected to be available by the end of August.

Finally, there was discussion of where to go next. The general consensus of the meeting was that SNMP Security was too important and central to the technology for us to recommend progression in the standards track with the present incomplete levels of implementation experience. When asked how many other implementation efforts were planned for the near future, a half a dozen attendees raised their hands. These and others were strongly encouraged to proceed with these implementations in order to gain the required experience. Interoperability testing of such implementations across the Internet, and at the Interop '91 SNMP-demo "staging" event were discussed and encouraged.

Attendees

Steve Alexander	stevea@i88.isc.com
Karl Auerbach	karl@eng.sun.com
Doug Barlow	barlow@decwet.dec.com
James Barnes	barnes@xylogics.com
Steve Bostock	steveb@novell.com
Howard Brown	brown@ctron.com
Theodore Brunner	tob@thumper.bellcore.com
John Burruss	jburruss@wellfleet.com
Jeffrey Case	case@cs.utk.edu
Gigi Chu	gigic@hpspd.spd.hp.com
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com
Emil Datability	
James Davin	jrd@ptt.lcs.mit.edu
Jeffrey Edelheit	edelheit@mitre.org
Gary Ellis	garye@hpspd.spd.hp.com
Bill Fardy	fardy@ctron.com
Barbara Fraser	byf@cert.sei.cmu.edu
Jeff Fried	jmf@relay.proteon.com
Deborah Futcher	dfutche@eco.twg.com
Maria Gallagher	maria@nsipo.arc.nasa.gov
Shawn Gallagher	gallagher@quiver.enet.dec.com
James Galvin	galvin@tis.com
Ron Jacoby	rj@sgi.com
Mike Janson	mjanson@mot.com
Frank Kastenholz	kasten@europa.clearpoint.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Mark Kepke	mak@hpcndk.cnd.hp.com
Kenneth Key	key@cs.utk.edu
Christopher Kolb	kolb@psi.com
Deidre Kostick	dck2@sabre.bellcore.com
Bobby Krupczak	rdk@cc.gatech.edu
Cheryl Krupczak	cheryl@cc.gatech.edu
Nik Langrind	nik@shiva.com

Anthony Lauck	lauck@tl.enet.dec.com
Tim Lee-Thorp	ngc!tim@uunet.uu.net
Ron Mackey	rem@dsiinc.com
Keith McCloghrie	kzm@hls.com
Evan McGinnis	bem@3com.com
Lynn Monsanto	monsanto@eng.sun.com
Bradford Parker	brad@cayman.com
David Perkins	dperkins@synoptics.com
John Pickens	jrp@3com.com
Brian Price	brian@bss.com
Anil Rijsinghani	anil@levers.enet.dec.com
Kary Robertson	kr@concord.com.kr
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
John Seligson	johns@ultra.com
Ron Sharp	rls@neptune.att.com
Anil Singhal	nsinghal@hawk.ulowell.edu
Mark Sleeper	mws@sparta.com
Michael St. Johns	stjohns@umd5.umd.edu
Bob Stewart	rlstewart@eng.xyplex.com
Bruce Taber	taber@interlan.com
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov
Glenn Trewitt	trewitt@nsl.dec.com
Theodore Tso	tytso@mit.edu
William Versteeg	bvs@nrc.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Drew Wansley	dwansley@secola.columbia.ncr.com
David Ward	dward@chipcom.com
Mark Wood	markl@dsiinc.com
Brian Yasaki	bky@eco.twg.com
Jeff Young	jsy@cray.com
Joseph Zur	fibrontics!zur@uunet.uu.net

3.7.6 Site Security Policy Handbook (ssphwg)

Charter

Chair(s):

J. Paul Holbrook, holbrook@cic.net
Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:

General Discussion: ssphwg@cert.sei.cmu.edu
To Subscribe: ssphwg-request@cert.sei.cmu.edu

Description of Working Group:

The Site Security Policy Handbook Working Group is chartered to create a handbook that will help sites develop their own site-specific policies and procedures to deal with computer security problems and their prevention.

Among the issues to be considered in this group are:

1. Establishing official site policy on computer security:
   - Define authorized access to computing resources.
   - Define what to do when local users violate the access policy.
   - Define what to do when local users violate the access policy of a remote site.
   - Define what to do when outsiders violate the access policy.
   - Define actions to take when unauthorized activity is suspected.
2. Establishing procedures to prevent security problems:
   - System security audits.
   - Account management procedures.
   - Password management procedures.
   - Configuration management procedures.
3. Establishing procedures to use when unauthorized activity occurs:
   - Developing lists of responsibilities and authorities: site management, system administrators, site security personnel, response teams.
   - Establishing contacts with investigative agencies.
   - Notification of site legal counsel.
   - Pre-defined actions on specific types of incidents (e.g., monitor activity, shut-down system).
   - Developing notification lists (who is notified of what).
4. Establishing post-incident procedures
   - Removing vulnerabilities.
   - Capturing lessons learned.
   - Upgrading policies and procedures.

Goals and Milestones:

Done	Review, amend, and approve the Charter as necessary. Examine the particular customer needs for a handbook and define the scope. Continue work on an outline for the handbook. Set up an SSPHWG “editorial board for future writing assignments for the first draft of document.
Done	Finalize outline and organization of handbook. Partition out pieces to interested parties and SSPHWG editorial board members.
Done	Pull together a first draft handbook for Working Group review and modification.
Oct 1990	Finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor for publication.
Oct 1990	Finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor forpublication.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI and J. Paul Holbrook/CERT

SSHWG Minutes

The "Site Security Handbook", RFC 1244/FYI 8, was released just prior to the Atlanta IETF meeting.

The SSH was the topic of a plenary presentation by Holbrook at the Atlanta IETF meeting.

The SSPHWG met in Atlanta with the goal of deciding what to do next. The meeting was attended by about five people, so the discussion was limited.

The group did agree to generate an executive summary of the document. Al Hoover of ANS volunteered to work with Paul Holbrook on this document. The intent is to turn the summary into an informational RFC as well. The summary would be something you could give to local decision makers to explain what the Site Security Handbook is all about and what the general approach is.

There was some agreement that once the executive summary is available, we can start pushing getting the word out about this document. We will almost certainly contact trade press, for example.

The group was vaguer about whether a workshop would be useful. Unless there are people to push on that issue, it won't happen.

At the conclusion of the meeting, Joyce K. Reynolds announced that she would be 'retiring' as Co-Chair of the SSHWG. Joyce was truly the glue that held the SSPHWG together; without her efforts, the Site Security Handbook would never have been published. Thanks, Joyce!

Attendees

John Cook	cook@chipcom.com
Jill Foster	jill.foster@newcastle.ac.uk
Arlene Getchell	getchell@nersc.gov
J. Paul Holbrook	holbrook@cic.net
Alton Hoover	hoover@nis.ans.net
Darren Kinley	kinley@crim.ca
Joyce K. Reynolds	jkrey@isi.edu
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov

3.8 Transport and Services Area

Director(s):

- Dave Borman: dab@cray.com

Area Summary reported by Greg Vaudreuil/CNRI

The Transport and Services Area of the IETF is new this meeting. It was formed to respond to the growing need for attention to the transport level services increasingly demanded by applications developers, and to address the evolution of new and existing transport level protocols. This area closely resembles the Host and User Services area originally Chaired by Craig Partridge/ BBN.

Working Groups in this area include: Distributed File Systems, Domain Name System, Service Location Protocol, and Trusted Network File Systems.

Domain Name System

The Domain Name System is encountering growth pains, both in terms of the protocol and common implementations and in operational coordination. The DNS Working Group met at this meeting to discuss how to re-commit itself to accomplishing the work, both to operations coordination and to evolving the protocols to meet new needs.

Trusted Network File Systems

This Working Group is a joint working group between the IETF and the Trusted Security Interoperability Group. So far this effort have been a useful collaboration between the two groups.

Service Location Protocol

The Service Location Protocol Working Group met and worked on defining parameters and goals for either choosing an existing approach or formulating a new approach to resource discovery.

3.8.1 Distributed File Systems (dfs)

Charter

Chair(s):
Peter Honeyman, honey@citi.umich.edu

Mailing Lists:
General Discussion: dfs-wg@citi.umich.edu
To Subscribe: dfs-wg-request@citi.umich.edu
Archive:

Description of Working Group:

Trans- and inter-continental distributed file systems are upon us. The consequences to the Internet of distributed file system protocol design and implementation decisions are sufficiently dire that we need to investigate whether the protocols being deployed are really suitable for use on the Internet. There's some evidence that the opposite is true, e.g., some DFS protocols don't checksum their data, don't use reasonable MTUs, don't offer credible authentication or authorization services, don't attempt to avoid congestion, etc. Accordingly, a Working Group on DFS has been formed by the IETF. The Working Group will attempt to define guidelines for ways that distributed file systems should make use of the network, and to consider whether any existing distributed file systems are appropriate candidates for Internet standardization. The Working Group will also take a look at the various file system protocols to see whether they make data more vulnerable. This is a problem that is especially severe for Internet users, and a place where the IETF may wish to exert some influence, both on vendor offerings and user expectations.

Goals and Milestones:

May 1990 Generate an RFC with guidelines that define appropriate behavior of distributed file systems in an internet environment.

3.8.2 Domain Name System (dns)

Charter

Chair(s):

Michael Reilly, reilly@nsl.dec.com

Mailing Lists:

General Discussion: dns-wg@nsl.dec.com
To Subscribe: dns-wg-request@nsl.dec.com
Archive:

Description of Working Group:

The DNS Working Group is concerned with the operation of name servers on the Internet. We do not operate name servers but serve as a focal point for the people who do operate them. We are also concerned with the Domain Name System itself. Changes to the existing RFC's, for example, are discussed by the Working Group. If changes to the RFC's or additional DNS related RFC's are deemed necessary the Working Group will propose them and will prepare the associated documents.

Because we intend to serve as the focal point for people operating name servers, one of our projects will be to assist anyone bringing up a name server by publishing a collection of useful hints, tips and operational experience learned by the people already running name servers.

The DNS Working Group will also take an active role in the dissemination of solutions to problems and bugs encountered while running various name server implementations. We will also provide guidance to anyone writing a new name server implementation, whenever possible.

Goals and Milestones:

TBD Adding DNS variables to the MIB.

TBD Hints, tips, and operations guide for DNS software

TBD Implementation catalog for DNS software.

TBD Discussion of adding load balancing capability to the DNS.

TBD Discussion of adding a Responsible Person Record.

TBD Discussion of adding network naming capability to the DNS.

none specified

CURRENT MEETING REPORT

Reported by Michael Reilly/DEC DNS Minutes

Administration

We began this meeting of the DNS Working Group by discussing administrative details. The group decided to merge the current DNS Working Group mailing list with the namedroppers mailing list. (This was completed during the week following the Atlanta meeting.) The proposed DNS Working Group charter was discussed and resulted in a decision to state the operational and protocol goals of the Working Group in separate paragraphs. This will serve to emphasize both aspects of the charter.

The working group has begun an examination of some of the security aspects of the DNS. In working with this sensitive information, many questions were raised. It was noted by some WG members that this type of information has not been made widely available in the past. It was asked whether there were any existing Internet or IETF policies either limiting or encouraging wide dissemination of this type of security information. The WG Chair will attempt to answer this question and will edit the proposed charter as appropriate. An edited version of the charter will be posted to the namedroppers mailing list before the end of August.

Splitting the DNS Working Group

A proposal to split the current DNS Working Group into two separate groups was discussed. The proposal calls for a WG responsible for the operational aspects of the DNS in the Internet and a second WG responsible for additions and changes to the existing DNS RFC's (aka a protocol group). The members of the Working Group present felt that the current group should not be split. It was noted that should the Working Group be split the membership of each new group would consist mostly of the members of the current DNS WG. It was also noted that commitments to other Working Groups may prevent many from attending two DNS related Working Groups.

Current Operational Problems

The Working Group then began discussions of problems which have been observed in the operation of the DNS in recent months. As a part of this discussion Philip Almquist spent a few minutes presenting an overview of the changes he has made to the current version of BIND (4.8.3) to solve some of the more critical problems. Philip is engaged in a project to produce a version of BIND which closely follows the DNS RFC's while being more robust and easier to configure than the current version. All of his work is based on the freely available BIND source code released by the University of California at Berkeley. The results of Philip's work will also be freely available. Philip expects his version of BIND (current known as BIND version 4.8.4) to begin beta testing in September or October.

A DNS MIB

Joseph Peck presented his "strawman" proposal of the DNS information which should be made available via SNMP. Several useful comments and suggestions were made by members of the Working Group concerning the contents of a DNS MIB as well as a proposal for organizing the MIB into a generic and an implementation specific portion. A revised proposal will be posted to the mailing list within a few weeks. Several WG members expressed an interest in working on the MIB document itself. This effort will be organized via the mailing list as soon as the revised proposal has been discussed.

Several people suggested that the DNS MIB may need to contain variables which should be "settable" i.e., "flush the cache" or "reload the database from permanent storage". The revised proposal mentioned above will include a separate section of these settable variables.

A proposal for the DNS MIB is expected to be presented at the next IETF meeting.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Stevan Belair	swbelair@sprintf.merit.edu
Thomas Brisco	brisco@rutgers.edu
Mats Brunell	mats.brunell@sics.se
Stephen Crocker	crocker@tis.com
John Curran	jcurran@bbn.com
James Ellis	jte@cert.sei.cmu.edu
Johnny Eriksson	bygg@sunet.se
Erik Fair	fair@apple.com
Russ Hobby	rdhobby@ucdavis.edu
Steven Holt	
Alton Hoover	hoover@nis.ans.net
Tim Howes	Tim.Howes@umich.edu.
Tom Kessler	kessler@sun.com
Darren Kinley	kinley@crim.ca
Richard Larkin	rblarkin@sprintf.merit.edu
Walter Lazear	lazear@gateway.mitre.org
John Lekashman	lekash@nas.nasa.gov
Louis Leon	osll@emuvm1.cc.emory.edu
Louis Mamakos	louie@ni.umd.edu
Bill Manning	bmanning@rice.edu
April Marine	april@nisc.sri.com
Paul Mockapetris	pvm@darpa.mil
Keith Moore	moore@cs.utk.edu
Michael Patton	map@lcs.mit.edu
Joe Peck	peck@ms1.pa.dec.com
Jon Postel	postel@isi.edu
Michael Reilly	reilly@nsl.dec.com

Jan Michael Rynning	jmr@nada.kth.se
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Erik Sherk	sherk@nmc.cit.cornell.edu
Keld Simonsen	keld.simonsen@dkuug.dk
Bernhard Stockman	boss@sunet.se
Glenn Trewitt	trewitt@nsl.dec.com
Rudiger Volk	rv@informatik.uni-dortmund.de

3.8.3 Service Location Protocol (svrloc)

Charter

Chair(s):

John Veizades, `veizades@apple.com`

Mailing Lists:

General Discussion: `srv-location@apple.com`
To Subscribe: `srv-location-request@apple.com`
Archive: `pub/srv-location/svr-loc-archive`

Description of Working Group:

The Service Location Working Group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based services location protocols. Protocols that would be designed by this group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the services location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this group to understand the operation of services location and then determine the correct action in their view whether it be to use current protocols to suggest a services location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

Done	Open discussion and determine if a Working Group should be formed.
Done	Continue discussion trying to refine the problem statement and possible resolutions.
Jul 1991	Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

SVRLOC Minutes

Minutes of the Atlanta Services Location Working Group

This meeting started with the discussion that has started many of these meetings, with a discussion of the solutions that exist in the area of services location protocols. The examples, as previously given, are the XNS network binding protocol, AppleTalk's Name Binding Protocol and of course the TCP/IP Resource Location Protocol.

Discussion lead to some of the problems in the propagation of request through IP based internets. IP based internets do allow for the propagation of datagrams with the use of multicast groups but the topography of the IP internet does not necessarily match the propagation pattern of the requests which may be organizational or geographical in nature.

The group decided to first concentrate on the local network case and then concentrate on the problem of extending the architecture to an arbitrary internet.

In dealing with the local network situation the group came up with the following list of parameters that maybe of interest in the predicate portion of look ups:

- Name
- Service type
- Organization
- Geographic location
- Authentication
- Next protocol (the next protocol used in finding the service)
- Query originator
- Where is the service in relation to where am I (delta distance) the predicate protocol to be used.

John Veizades and Leo McLaughlin will be looking into all of this and posting something to the list by the next meeting.

This group will also continue to work with Mike Schwartz and the IRTF group on resource location.

The next meeting will be at the November IETF.

Attendees

Steve Alexander	stevea@i88.isc.com
Karl Auerbach	karl@eng.sun.com
David Borman	dab@cray.com

David Bridgham	dab@asylum.sf.ca.us
Gregory Bruell	gob@shiva.com
Richard Cherry	rcherry@novell.com
Richard Cogger	rhx@cornellc.cit.cornell.edu
Peter Deutsch	peterd@cc.mcgillica
Ralph Droms	droms@bucknell.edu
Alan Emtage	bajan@cc.mcgill.ca
Holly Knight	holly@apple.com
Stev Knowles	stev@ftp.com
John Lekashman	lekash@nas.nasa.gov
Joshua Littlefield	josh@cayman.com
Carl Malamud	carl@malamud.com
Leo McLaughlin	ljm@wco.ftp.com
Bill Melohn	melohn@sun.com
William Nowicki	nowicki@legato.com
Geir Pedersen	geir.pedersen@use.uio.no
Mike Petry	petry@ni.umd.edu
John Veizades	veizades@apple.com
A. Lee Wade	wade@discovery.arc.nasa.gov
Peter de Vries	peter@wco.ftp.com

3.9 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Five Working Groups met at the IETF in Atlanta:

- Directory Information Services (pilot) Infrastructure Working Group (DISI), Chaired by Christopher Weider.

  DISI is a Working Group that provides a forum to define user requirements in X.500. It is an offshoot of the OSI Directory Services group and is a combined effort of the User Services Area and the OSI Integration Area of the IETF.

  Paper 1, "Exeutive Summary" (Weider, Reynolds, Heker). Intends to define issues DISI should be working on. This current draft was withdrawn for revisions, as the paper tended to blur the distinction between the X.500 standard and the Quipu implementation of that standard.

  Paper 2, "Survey" (Lang, Wright). Ruth Lang and Russ Wright were thanked for all of their work in this document. Additional DUAs will be added that were inadvertently left out.

  Paper 3, "Advanced Usages". General consensus is this third paper is still not ready to be a writing assignment of this group. There are other documents that need to be written before this one.

  There was a brief discussion regarding any areas which needed attention which were not being covered by either the DISI or OSI-DS forums. The general consensus was that there was as yet no justification for spinning off an operations group.

- Network Information Services Infrastructure Working Group (NISI), Chaired by Dana Sitzler and Patricia Smith.

  Review of this group's Internet Draft "Building a Network Information Services Infrastructure" included a discussion of including a security "verification" statement in the document and in the NIC Template.

  Also, discussion on where this group should go from here - there was justification of additional action items/tasks that fall in NISI's realm.

- NOC-Tool Catalogue Revisions Working Group (noctool2) Chaired by Robert Enger and Gary Malkin

  This group is on target. It is continuing to accept additional "vendor gathering". The official deadline for catalog submissions will be one month before the next IETF, with the document going to Internet Draft status after the next IETF.

  The issue of funding for a project in which the "living documents" (such as the NOCTOOLS catalogue), problem would be addressed was raised. This idea remains under consideration.

- Site Security Policy Handbook Working Group (SSPHWG), Chaired by J. Paul Holbrook and Joyce K. Reynolds.

  The SSPHWG published its document: FYI8, RFC1244, "Site Security Handbook", July 1991.

  Please consult the Security Area report for further information on this group's future progress.

- User Services Working Group (USWG), Chaired by Joyce K. Reynolds.

  Agenda items included:

  - Special Presentation:

    A special presentation on RARE WG3 (Association of European Research Networks) was provided by Jill Foster/Newcastle, U.K.

  - Discussions regarding:

    User-Doc Bibliography Revised Working Group Charter (To update FYI 3, RFC 1175).

    After a one year hiatus, the USER-DOC Working Group has been reborn. The two new Chairs who graciously volunteered to lead this endeavor are:

    * Lenore Jackson - NASA/Goddard
    * Ellen Hoffman - Merit, Inc.

    A revised Charter is currently underway.

    Internet-Draft:
    Aggarwal, V.,"Mid-Level Networks; A Minimum Services Provider",
    I-D: draft-aggarwal-services-00.txt, .ps

The RFC Editor tasked the review of this document to the USWG. In this forum, the USWG provided various input about the wording of the document, its relation to the efforts of the NISI Working Group, and the focus and goal of the document's recommendations. In general, it is a good document, but requires additional work before the USWG will send its "approval" to the RFC Editor.

- Internet User Glossary Working Group (userglos), Chaired by Karen Roubicek and Tracy LaQuey Parker. Did not meet.

3.9.1 Directory Information Services Infrastructure (disi)

Charter

Chair(s):

Chris Weider, clw@merit.edu

Mailing Lists:

General Discussion: disi@merit.edu
To Subscribe: disi-request@merit.edu
Archive: pub/disi-archive@merit.edu

Description of Working Group:

The Directory Information Services (pilot) Infrastructure Working Group (DISI) is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

The Directory Information Services (pilot) Infrastructure Working Group is an offshoot of the OSI Directory Services group, and, accordingly, is a combined effort of the OSI Integration Area and User Services Area of the IETF. The current OSIDS Working Group was chartered to smooth out technical differences in information storage schema and difficulties in the interoperability and coherence of various X.500 implementations. The DISI group is concerned solely with expanding the Directory Services infrastructure. As DISI will be providing infrastructure with an eye towards truly operational status, DISI will need to form liasons with COSINE, Paradyse, and perhaps the RARE WG3.

As a final document, the DISI Working Group shall write a Charter for a new Working Group concerned with user services, integration, maintenance, and operations of Directory Services, the Internet Directory User Services Group.

Goals and Milestones:

Done — First IETF Meeting: review and approve the Charter making any changes necessary. Examine needs and resources for the documentation to be produced, using as a first draft a document produced by Chris Weider, MERIT, which will be brought to the IETF. Assign writing assignments. Further work will be done electronically.

Jul 1991 Second IETF Meeting: review and approve documentation; review and approve Charter for the IDUS group.

Aug 1991 Electronically review final draft of documentation, and, if acceptable, submit to IESG for publication.

Dec 1991 Third IETF Meeting: Declare success and reform DISI group as IDUS group.

Internet Drafts:

"Interim Schema for Network Infrastructure Information in X.500", 06/14/1991, Chris Weider, Mark Knopper <draft-ietf-disi-netinfrax500-00.txt>

"A Catalog of Available X.500 Implementations", 07/25/1991, R. Lang, R. Wright <draft-ietf-disi-catalog-00.txt>

CURRENT MEETING REPORT

Reported by Chris Weider/ISI

DISI Minutes

A list of attendees follows the minutes.

The Agenda, as stated at the beginning of the meeting, was as follows:

1. Charter Issues
2. Scope of DISI
3. Paper 1 (Weider, Reynolds, Heker)
4. Paper 2 (Lang, Wright)
5. Third paper assignment
6. Transition to Ops group

The way it went:

1. Chris W. mentioned that the charter had been revised as recommended in St. Louis.

2. A freewheeling discussion about the scope of DISI then ensued. Chris W. asked Steve Kille to speak a bit about the chartered purpose of OSI-DS, and then asked the crowd if there were any areas which needed attention which were not being covered by either DISI or OSI-DS. The general consensus was that there was as yet no justification for spinning off an operations group. Thus, this covered both items 2 and 6.

3. Paper 1 was then mentioned. Chris W. stated that the paper would be withdrawn for revisions, as the paper tended to blur the distinction between the X.500 standard and the Quipu implementation of that standard. Chris promised that a revised version of the paper would be out in several weeks.

4. Paper 2 was then mentioned. The group publicly thanked Russ Wright and Ruth Lang for all the work they'd done on the paper. Several people then mentioned that there were some DUAs which did not appear in the paper, and Russ and Ruth stated that they would include them.

5. Discussion then turned to the third paper. Many people felt that there were some other papers which needed to be assigned before the 'Advanced Usages' paper; one which was agreed upon was a paper 2.5, a 'How to get Connected' paper. Steve Kille agreed to post to the list several papers written by European X.500 groups which could be assimilated and used as a springboard for the new paper; a 'How to join a pilot' paper, and a 'How to set up a DSA' paper.

6. Point 6 was covered in the discussion of point 2.

Attendees

William Biagi	bbiagi@cos.com
David Brent	brent@CDNnet.ca
Randy Butler	rbutler@ncsa.uiuc.edu
Chi Chu	chi@sparta.com
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgillica
Alan Emtage	bajan@cc.mcgill.ca
Robert Enger	enger@seka.scc.com
Urs Eppenberger	eppenberger@switch.ch
Jill Foster	jill.foster@newcastle.ac.uk
Arlene Getchell	getchell@nersc.gov
Jack Hahn	hahn@umd5.umd.edu
Steve Hardcastle-Kille	S.Kille@cs.ucl.ac.uk
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
Steven Hotz	hotz@isi.edu
Tim Howes	Tim.Howes@umich.edu.
Erik Huizer	huizer@surfnet.nl
Lenore Jackson	jackson@dftnic.gsfc.nasa.gov
Mark Knopper	mak@merit.edu
Ruth Lang	rlang@nisc.sri.com
Louis Leon	osll@emuvm1.cc.emory.edu
Brian Lev	lev@dftnic.gsfc.nasa.gov
Peter Liebscher	plieb@sura.net
Bill Manning	bmanning@rice.edu
April Marine	april@nisc.sri.com
Clifford Neuman	bcn@isi.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Mel Pleasant	pleasant@hardees.rutgers.edu
A. Minick Rushton	rushton@stsci.edu
Tom Sandoski	tom@concert.net
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Subu Subramanian	subu@qsun.att.com
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov
Wengyik Yeong	yeongw@psi.com

3.9.2 Internet User Glossary (userglos)

Charter

Chair(s):

Tracy LaQuey Parker, tracy@utexas.edu

Mailing Lists:

General Discussion: usergloss@ftp.com
To Subscribe: usergloss-request@ftp.com
Archive:

Description of Working Group:

The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the Charter as necessary. Discussion of Userglos Working Group Chair nominations submitted by USWGers.
TBD	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
TBD	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.

3.9.3 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@seka.scc.com
Gary Malkin, gmalkin@ftp.com

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu
Archive:

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.
- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.
- Identify additional new or improved tools that have become apparent since the last the compilation of the reference document.
- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Done Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.

Aug 1991 Draft of catalog will be prepared, draft to be reviewed and modified.Initiate IETF Internet Draft review process by submission of a "Son of NOCTools" catalog draft to IESG Secretary.

Dec 1991 Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI

RFC for publication

CURRENT MEETING REPORT

Reported by Darren Kinley/CRIM

NOCtools2 Minutes

Agenda

The NOCtool2 Working Group held a relatively short meeting at the 21st IETF in Atlanta. The meeting served mostly to report on the revision of RFC 1147/FYI 2, as well as to discuss the maintenance of "living documents". On the agenda,

- Summary of Working Group charter by Gary Malkin
- Status of the document revision process by Gary Malkin
- Discussion of how to best complete the revision process
- Discussion of on-going maintenance for "living documents"

Discussions

Deadlines and Schedules	The official deadline for catalog submissions will be one month before the next IETF, leaving almost a month to prepare a draft document. Following an editing session in Santa Fe the document will be put up as ID, and after a period of two months will be either reviewed and remain ID for another two month period, or be submitted to the RFC editor.
Living Documents	The group still hasn't made much headway on how to deal with this complex problem. Realizing that we are probably not best suited to solve this problem ourselves, we are beginning to look to the IETF and beyond for some input.
Funding	The issue of funding for a project in which the "living documents" problem would be addressed was raised. Much discussion about requirements for funding, IETF policy, and obligations to funding agencies ensued. Nothing has been decided, the idea remains under consideration.
Santa Fe Meeting	Draft document editing session and revisit the "living documents" problem.

Other Items

- Contact the Help Desk mailing list.
- Find alternative ways for submitting entries (i.e., kermit).
- Will we apply "diff"s?

Action Items

Robert Enger, Darren Kinley and Gary Malkin Contact remaining people with entries, continue to solicit and accept new entries, prepare draft document for Santa Fe meeting.

Joyce Reynolds Make available guidelines for solicitation of funds to interested persons.

Gary Malkin Consider funding for maintenance of "living documents" problem vis-a-vis the Internet Society.

Joyce Reynolds Make initial contact with the applications area director concerning the maintenance of "living documents" and consider a joint meeting with appropriate Working Groups.

Attendees

Henry Clark	henryc@oar.net
Robert Enger	enger@seka.scc.com
Shari Galitzer	shari@gateway.mitre.org
Kenneth Goodwin	goodwin@psc.edu
Jack Hahn	hahn@sura.net
Darren Kinley	kinley@crim.ca
Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Marsha Perrott	mlp@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com
Tom Sandoski	tom@concert.net

3.9.4 Network Information Services Infrastructure (nisi)

Charter

Chair(s):

Dana Sitzler, dds@merit.edu
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 WG will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. The work will begin with existing nic functions and services and should build upon work already being done within the Internet community. It should address areas such as common information formats, methods of access, user interface, and issues relating to security and privacy of Internet databases.

Goals and Milestones:

Done	Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of 'cooperative agreements' which will allow NICs to work together more effectively to serve users.
Done	Complete draft for phase 2 suggesting cooperative agreements for NICs.
Done	Revised draft document ready for wg 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	WG discussed current Internet draft and suggested minor revisions. Decision made to continue wg activity beyond this document.
Nov 1991	First document released as informational RFC. Outline and discuss new nisi tasks at IETF meeting.
Jul 1992	Write a document explaining the security issues of privacy and accuracy in Internet databases. Publish as an informational rfc.

Internet Drafts:

"Building a Network Information Services Infrastructure", 07/15/1991, D. Sitzler, P. Smith, A. Marine

<draft-ietf-nisi-infrastructure-00.txt>

CURRENT MEETING REPORT

Reported by Dana Sitzler/Merit

NISI Minutes

Agenda

- Review Activities
  - Draft document available
  - NSF nic solicitation
- Review Draft
  - Security issues
  - Information obligations
  - Other issues
- Implementing Doc Suggestions
  - Standard address (nic@domain)
  - Info validity suggestions
  - Nic-forum
  - Nic profiles
  - Discussion list
- What's Next
  - Publish and disband
  - Continue?
    - X.500
    - Archives
    - User interface recommendations

Discussion:

1. Review Activities

A document from this Working Group has been submitted as an Internet Draft. The draft document was used by NSF as one of the 'inputs' in preparing the NSFNET NIC solicitation.

2. Review Draft

The Working Group reviewed the draft document. We focused on the areas of where internet community members expressed concern either via email (as a result of the draft being available as a i-d) or in person during the Working Group. The following list outlines the group consensus which will be incorporated into the document.

2.1 Security Issues

There were some concerns expressed by the security area about the proper role for NICs in this area. The Working Group came up with a list of functions which it felt a NIC should deal with. This list will be shared with the Security Area Director and some agreement reached before these changes will be made in the document.

To deal with security issues, a NIC:

- Should be aware of security-related information/Educate users about security issues.
- Should be aware of security advisories.
- May serve as the first point of contact for an end-user and should know how to refer/escalate, etc.
- Should provide 'new' users with information about security such as referring to the Site Security HB.
- Should establish procedures for dealing with security 'emergencies' through coordination with NOCs.
- Can provide pointers to 'security' software such as PC virus disinfectant sw.
- Should be aware of and refer users (if appropriate) to security organizations such as CERT.

2.2 Personal/Organizational Information

The Working Group discussed the responsibilities and obligations of a NIC in providing personal or organizational information to the general public. We had a rather long and interesting discussion of this topic. We talked about the need to differentiate between different types of information, privacy issues, the expense involved with information collection and verification, and the trade-offs of having the info vs. having 'correct' info.

In terms of dealing with personal or organizational information, we decided to provide a mechanism to inform the information provider about what info is needed, what it will be used for, and if it will be made widely available. Here's what we came up with:

- When collecting personal/organizational information, NICs should provide a form which includes a 'disclosure statement'.
- The 'disclosure statement' should include:
  - What information is needed?
  - What it will be used for?
  - The consequences of supplying the information?
  - How widely available (and which info; some pieces may be made more widely available than other pieces?
    - Procedure for updating/correcting/disputing.

    * Frequency of update.
    * How to return the form (receipt of form w/requested information would be considered an acknowledgement that the info supplier agrees to the terms stated in the 'disclosure statement'.

- NICs should have a defined mechanism in place to update information collected @ some time frame.

- The date of last update/verification should be included with the information made available to others in the network community.

- NICs should understand and respect 'levels of security' for information – if info should not be widely available to the public, steps should be taken to make sure that the info is not accessible by anyone.

2.3 Other Issues:

- NICs have different audiences – emphasize in document the idea of working with other NICs (in terms of referring to another NIC if appropriate) and strengthen the idea of a 'primary' audience for a NIC which may have been funded by a specific group for a specific purpose

- Increase examples in section outlining current NIC services. For example, examples of archives, a specific online service, etc.

- Some discussion of recommending in document a common address for NIC ftp servers – lots of discussion here; no real consensus – folks with strong opinions about this may want to lobby to continue this discussion until some agreement can be made.

- Need to add section numbers.

- Shorten history section - does not add to document.

3 Implementing Draft Suggestions

3.1 Standard Email Address

The group had no problem with this recommendation. It was stressed that NIC people involved with this Working Group have to start the process of implementing this – and informing users about it.

3.2 Info Validity Check Info

Much of this discussion was covered in the previous section dealing with personal/organizational information. The basic suggestion to have all information include a contact (which may be the NIC) and some indication of the last verification

3.3 nic-forum

The group discussed the two components of the nic forum; the nic profiles and the discussion list. The nic profile information sheet was discussed and it was recommended that this sheet be made more 'user friendly!'. At present the profile sheet reflects the naming conventions necessary for X.500 but not the ones common to all of us human creatures. The profile sheet will be changed.

There was quite a bit of discussion about the discussion list aspect of the nic forum. Who is the audience? Is it an open list? Should it be moderated? Etc. A consensus was not reached on these issues. This meeting was the first time the actual implementation of this suggestion was discussed. The group agreed to continue discussion on the NISI mailing list.

4 What's Next?

The group discussed the possibilities for the next step for NISI. The following ideas were generated:

- Explore privacy issues.
- Develop an international profile database
- Develop an appropriate use document which addresses issues like privacy, how to use services, starting 'unsolicited stuff', etc.
- Define mechanisms for the exchange of information between groups such as nics and nocs.
- Access mechanisms; X.500, Z.39.50.
- Define requirements for user interface.
- Archive
- New user nethelp system - start nethelp pilot.
- Expand ideas presented in existing document including how nics and nocs interact; maintaining referral information; defining core information at nic.

The general consensus was that the last item on this list was probably an appropriate next step.

5 ACTIONS

- Update document.
- Review RARE Working Group profile.

- Discuss and agree to nic profile info and form (Sept).
- Discuss with USWG Chair – NISI next step (given list above).

Attendees

James Conklin	conklin@bitnic.educom.edu
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgillica
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.arc.nasa.gov
Arlene Getchell	getchell@nersc.gov
Joe Godsil	jgodsil@ncsa.uiuc.edu
Jack Hahn	hahn@sura.net
Martyne Hallgren	martyne@theory.tn.cornell.edu
Ittai Hershman	ittai@nis.ans.net
Ellen Hoffman	esh@merit.edu
J. Paul Holbrook	holbrook@cic.net
Darren Kinley	kinley@crim.ca
Christopher Kolb	kolb@psi.com
Dale Land	land@lanl.gov
Ruth Lang	rlang@nisc.sri.com
Brian Lev	lev@dftnic.gsfc.nasa.gov
Peter Liebscher	plieb@sura.net
Daniel Long	long@nic.near.net
April Marine	april@nisc.sri.com
Karen McKelvey	karen@cerf.net
Clifford Neuman	bcn@isi.edu
Marsha Perrott	mlp@andrew.cmu.edu
Joyce K. Reynolds	jkrey@isi.edu
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov
Theodore Tso	tytso@mit.edu
Rudiger Volk	rv@informatik.uni-dortmund.de
Chris Weider	clw@merit.edu
Wengyik Yeong	yeongw@psi.com

3.9.5 User Services (uswg)

Charter

Chair(s):
Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:
General Discussion: us-wg@nnsc.nsf.net
To Subscribe: us-wg-request@nnsc.nsf.net
Archive:

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF Working Groups created to perform certain projects, or outside organizations such as SIGUCCS.

- Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should:
  - Clearly address user assistance needs;
  - Produce an end-result (e.g., a document, a program plan, etc.);
  - Have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
  - Not duplicate existing or previous efforts.
- Create Working Groups or other focus groups to carry out projects deemed worthy of pursuing.
- Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

Ongoing This is an oversight group with continuing responsibilities.

Request For Comments:

RFC 1150 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. notes"

RFC 1177 "FYI on Questions and Answers - Answers to Commonly Asked "New Internet User" Questions"

RFC 1206 "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions"

RFC 1207 "Answers to Commonly asked "Experienced Internet User" Questions"

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

USWG Minutes

Announcements

- DISI - met Monday, 7/29, 4:00-6:00pm
- NISI - met Tuesday 7/30, 9:00-12Noon
- NOCTools Revised - met Monday evening, 7/29, 7:00-10:00pm
- Misc. Agenda Items
- Site Security Handbook published: FYI8, RFC1244

Special Presentation:

A special presentation on RARE WG3 (Association of European Research Networks) was provided by Jill Foster/Newcastle, U.K.

(See attached slides.) Please note: these slides were pulled from various other presentations, and were not specifically made for this session.

Discussions regarding:

User-Doc Bibliography Revised WG Charter (To update FYI 3, RFC 1175)

After a one year hiatus, the USER-DOC WG has been reborn. The two new chairpersons who graciously volunteered to lead this endeavor are:

- Lenore Jackson - NASA/Goddard
- Ellen Hoffman - Merit, Inc.

A revised charter is currently underway.

Internet-Draft: Aggarwal, V., "Mid-Level Networks; A Minimum Services Provider", I-D: draft-aggarwal-services-00.txt, .ps

The RFC Editor tasked the review of this document to the USWG. In this forum, the USWG provided various input about the wording of the document, its relation to the efforts of the NISI Working Group, and the focus and goal of the document's recommendations. In general, it is a good document, but requires additional work before the USWG will send its "approval" to the RFC Editor.

The author, Vikas Aggarwal, agreed to create a mailing list of the USWG session attendees in order to continue a discussion of this document due to time constraints. He will revise the I-D, and resubmit a new version to the I-D process for additional comments.

Final Comments from your semi-fearless USWG Chair:

The USWG had a "most excellent" session in Atlanta! Thanks to all who attended. We had super participation and contributions!

Special thanks to Jill Foster for enlightening us about RARE and its Working Groups.

KUDOS!!!!!!!!!!!!! You all deserve it!

Attendees

Vikas Aggarwal	vikas@JVNC.net
William Biagi	bbiagi@cos.com
George Brett	pls@psulias
Eric Carroll	eric@utcs.utoronto.ca
John Curran	jcurran@bbn.com
Peter Deutsch	peterd@cc.mcgillica
Alan Emtage	bajan@cc.mcgill.ca
Jill Foster	jill.foster@newcastle.ac.uk
Maria Gallagher	maria@nsipo.arc.nasa.gov
Jack Hahn	hahn@umd5.umd.edu
Martyne Hallgren	martyne@theory.tn.cornell.edu
Ellen Hoffman	esh@merit.edu
J. Paul Holbrook	holbrook@cic.net
Lenore Jackson	jackson@dftnic.gsfc.nasa.gov
Darren Kinley	kinley@crim.ca
Louis Leon	osll@emuvm1.cc.emory.edu
Brian Lev	lev@dftnic.gsfc.nasa.gov
Peter Liebscher	plieb@sura.net
Gary Malkin	gmalkin@ftp.com
April Marine	april@nisc.sri.com
Karen McKelvey	karen@cerf.net
Clifford Neuman	bcn@isi.edu
Marsha Perrott	mlp+@andrew.cmu.edu
Jack Pope	pope@uscdsv.acusd.edu
Joyce K. Reynolds	jkrey@isi.edu
A. Minick Rushton	rushton@stsci.edu
Tom Sandoski	tom@concert.net
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Subu Subramanian	subu@qsun.att.com
Ronald Tencati	tencati@nssdca.gsfc.nasa.gov
John Wobus	jmwobus@suvm.acs.syr.edu

RARE Working Group 3 User Support and Information Services Subgroup

Jill Foster

JNC

JillFoster@Newcastle.ac.uk

RARE :

Association of European Research Networks

RARE Working Groups:

WG1:	Message Handling Systems
WG2:	File Transfer Access and Management
WG3:	User Support, Information Services and Directories
WG4:	Network Operations and X.25
WG5:	Full Screen Services
WG6:	High Speed Communications and ISDN
WG8:	Security and Management of Network Applications

JNC

JillFoster@Newcastle.ac.uk

Level-7

COSINE

- Co-operation for Open Systems Interconnection Networking in Europe
- Establishing communications infrastructure for academic and research users
- 225000 target users

Level-7

COSINE Projects

Services	S1	IXI – X.25 infrastructure
	S2.1	Pan-European X.400
	S2.2	X.400 gateway to North America
Pilots	P1.1	FTAM gateway
	P2.1	Pan-European X.500
	P2.2	European Information Service
	P3	Support for International User groups

The Role of RARE WG3

- RARE:
  - Association of European Research Networks
- COSINE:
  - Cooperation for OSI Networking in Europe
- RARE / COSINE
  - Standards *and* Services
- RARE Working Group 3
  - Directory Services Subgroup
  - User Support & Information Services Subgroup
- WG3 USIS Subgroup
  - Liaison activities
  - International user group support
  - European Information Service

JNC

JillFoster@Newcastle.ac.uk

Liaison Activities

- Representatives from most RARE member countries:

  Austria, Belgium, Finland, France, Germany, Great Britain, Greece, Ireland, Netherlands, Norway, Spain, Switzerland, Sweden, Yugoslavia

  - Others welcome to join
  - EARN represented too
- "Network" of user support people
- Exchanging information and ideas
- Encouraging national user support

JNC

JillFoster@Newcastle.ac.uk

Real Networks

Don't Need

User Support!

JANET Networkshop '91

JillFoster@UKACNewcastle

Why Do We Need User Support?

- The success of COSINE depends on it

The users are the customers.

- Network users need good support:
  - Human Advisers / Help Desk
  - Networked Services that are easy to use
  - On-line information
  - Paper Documentation
  - Enhanced facilities for:
    - group communication
    - information dissemination

JNC

JillFoster@Newcastle.ac.uk

Why Do We Need User Support?

- The Network is not just for networkers but for academics of *ALL* disciplines.
  - Is easy to lose sight of this goal when:
    deciding on protocols
    writing applications
  - User Support is often forgotten or given low priority
- Good User Support is vital to the success of our networks
  - The Users are the Customers!

JANET Networkshop '91

JillFoster@UKACNewcastle

Who are the REAL Users?

- No longer just the networkers and a few scientists.
- Increasing number of users from *ALL* academic disciplines.
  (ISI Database used by staff & students)
- Also research support staff:
  - Librarians
  - University Admin. staff
- Users coming from PC environment

JANET Networkshop '91

JillFoster@UKACNewcastle

European Information Service

- WG3 Survey & Report on User Support
- WG3 Proposal for a Pilot EIS
- COSINE Sub-project P2.2
- Central Information Services
  - On-line Information
  - Help Desk
  - Enhanced facilities for group communication

JNC

JillFoster@Newcastle.ac.uk

Level-7

COSINE P2.2:

Support

and

Information Services

Level-7

Information Service

- Central Focal Point
- Based in Bracknell, UK
- Up to 20 interactive users
- Help desk
- Academic and research users

Level-7

Database Information

- COSINE
- Information servers
- Distribution lists
- Networks
- Special interest groups
- Conferences
- Help desks
- Contacts

... etc

Level-7

Access

- X.400
- FTAM
- VTP
- X.29
- Dial-up

Level-7

Other Activity

- Network co-ordination meetings
- Publicity
- Report on evolution of service
- Report on self sustaining service

Level-7

Plans

X.400 access	3rd q 1991
FTAM access	3rd q 1991
Interactive access	1st q 1992
Self sustaining service possible	1st q 1993

IUG Support Project Proposal

- WG3 Information Services Subgroup:

  Proposal for a COSINE sub-project:

  Pilot activities to support international user groups.
  - General mechanisms usable by all IUGs

- First define "User Group"

A group of (potential) users of the network facilities who are linked either formally or informally by their common area of research. (E.g. WHO Diabetes group)

Already use various methods to communicate:

- Postal mail
- Telex
- Fax
- Telephone
- Meetings
- Conferences / Workshops
- Email
- Bulletin Boards
- Email Discussion Lists

COSINE USER Forum

Jill.Foster@Newcastle.ac.uk

International User Group Support

- Main Objectives of Project
  - To provide a general framework for support of international user groups :
    "Network" of international, national and local user support
  - Co-ordination of activities at international level
  - To use and provide feedback on Pilot Projects:
    - P2.1 - Directory Services
    - P2.2 - Information Services (EIS and Tools)
  - To build on experience of existing user groups
  - To encourage the use of COSINE services by the European research community for enhanced group communication

COSINE USER Forum

JillFoster@Newcastle.ac.uk

Overview of project

- Phase I: Requirements Study
  - Communications requirements of different communities
  - Experience of groups already using network
  - Reasons why other groups do not use the network
  - "Ground rules" for Phase II
- Phase II: Demonstrator Projects
  - 8 Groups chosen according to certain "criteria":
    - spread of research interests across groups
    - existing reason to communicate
    - well motivated to move to electronic means
  - Activities to overcome main problems

COSINE USER Forum

JillFoster@Newcastle.ac.uk

Main Problems & Possible Solutions

- No natural pan-European organisational framework
  - Need coordination of the group's activities at an international level
  - Need "network" of international, national and local experts and advisers
- IUG Project Support Framework
  - IUG Project Team
  - Project Officers (one for each user group)
  - EIS Staff
  - National Network User Support Staff

COSINE USER Forum

JillFoster@Newcastle.ac.uk

- IUG Project Team is responsible for:
  - Coordination of the project
  - Ensuring effective liaison between:
    - Project Officers
    - EIS Staff
    - National network user support staff
  - Training and support of Project Officers
- IUG Project Officer is responsible for:
  - Knowledge dissemination within their group
  - Demonstrations and seminars
  - Documentation
  - Training: educate a "network" of national and local experts within their user group
  - Encouraging electronic discussions
  - Other promotional activities
- IUG Project Officer - Important Role

COSINE USER Forum

JillFoster@Newcastle.ac.uk

Main problems & possible solutions

- Users don't know that networks exist!
  - Need for publicity and promotional activities:
    - On paper: *"What networks can do for you"*
    - Demonstrations at Conferences
    - Articles in leading journals
    - Target and train key people in the group
    - Main function of IUG Project Officer
- No central point of contact
  - EIS to play a major role in the IUG Project
  - EIS+ could provide central IUG support:
    - Manned Help Desk
      - User referred to national user support
    - On-line Information
    - Directory of services

COSINE USER Forum

JillFoster@Newcastle.ac.uk

Main problems & possible solutions

- No generally available tools for group communication
  - EIS is to provide tools for IUGs:
    - Conferencing / Discussion Servers
    - File Servers
    - Database Servers
- Lack of tools for IUGs to manage their own information
  - EIS to provide tools for:
    - Data maintenance
    - Management of conferences, discussions etc.
    - Accounting / Statistics
    - Security: authorisation and access controls
  - Feedback from real user groups invaluable

COSINE USER Forum

JillFoster@Newcastle.ac.uk

IUG Project: Summary

- Develop a framework for IUG Support
- Develop a set of guidelines and services usable by other IUGs
- Investigate mechanisms for continued funding or cost recovery

COSINE USER Forum

Jill.Foster@Newcastle.ac.uk

Current NISP (Mailbase) Project

- Coming to end of two year period
- Developed prototype server: Mailbase for group communication
  - Enhanced Mailing List support
  - Group's discussions take place on a set of Lists
  - Files associated with the Lists
  - Access: Email and Interactive Session
- Test Service in use by:
  - 10 different groups (Libraries, CHEST,IUIC)
  - 59 Lists
  - 2200+ distinct users

JANET Networkshop '91

Jill.Foster@UK.AC.Newcastle

NISP II

- Provide a production version of Mailbase and run a Central Production Service
- To enhance the functionality of Mailbase
  - In line with user requirements
  - Improve the interactive interface
  - Add free text retrieval capability
- To help, encourage and train special user groups to use JANET & JANET Services
- Liaison with others working in the same area as NISP II

JANET Networkshop '91

Jill.Foster@UK.AC.Newcastle

NISP II: User Group Support

- Aim: to help user groups help themselves
  - Target 4 or 5 major user groups within the UK academic community
    - Wide range of research areas, teaching and support groups
  - Encourage and help them to use the tools provided by the NISP Server and other JANET Services
  - Role of Group Leader vital
  - User groups to fund Leader for NISP II

JANET Networkshop '91

Jill.Foster@UK.AC.Newcastle

User Group Support Activities

- Liaison
  - To develop a "virtual network" for user group support
- Training of group leader
  - To train members of the user group
    - in a language they understand
    - using examples that are relevant
- Publicity
  - Poster/brochure *"What networks can do for you"*
  - Presentations at users' conferences

JANET Networkshop '91

Jill.Foster@UK.AC.Newcastle

User Group Support Activities

- Training/Workshops
  - Active support and encouragement
  - Full day of presentations, demos and hands on workshop for each group
  - Material adapted to the specific group by the group "leader"
- Documentation Pack
- On-going encouragement
  - Role of group leader very important
    - Encouraging discussions
    - Providing "seed" information
    - Info on relevant Info Services
    - Admin functions

JANET Networkshop '91

Jill.Foster@UK.AC.Newcastle

Chapter 4

Network Status Briefings

4.1 CERFnet Report

Presented by Susan Estrada/CERFnet

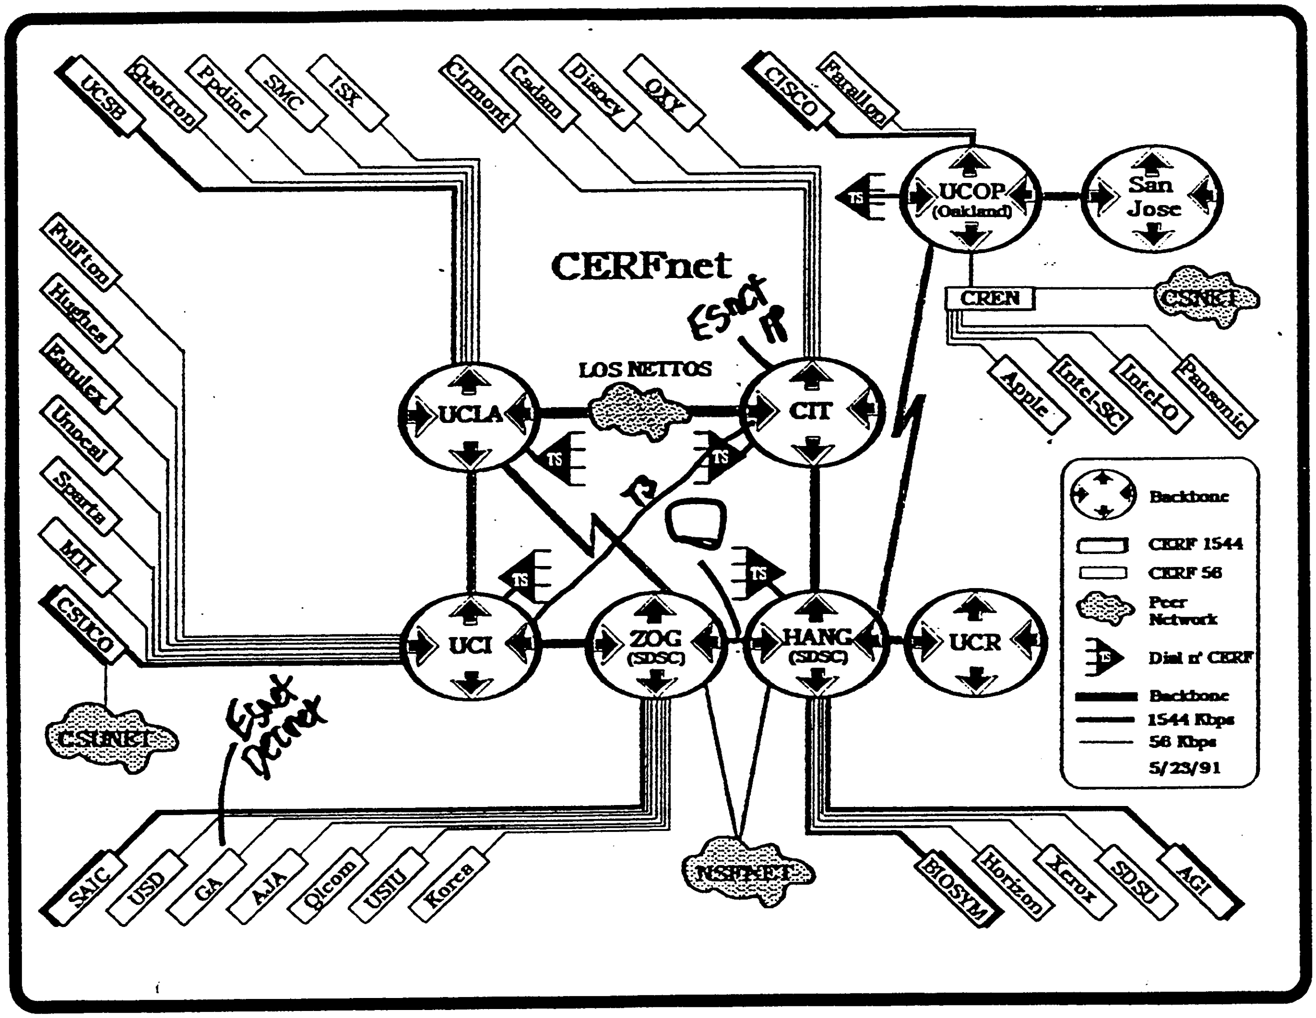
CERFnet
UCSB
Quotron
Ppdine
SMC
ISX
Clrmont
Cadam
Disney
OXY
CISCO
Farallon
UCOP
(Oakland)
San Jose
CREN
CSUNET
Apple
Intel-SC
Intel-O
Pansonic
Fullerton
Hughes
Emulex
Unocal
Sparta
MTI
CSUCO
LOS NETTOS
UCLA
CIT
UCI
ZOG
(SDSC)
HANG
(SDSC)
UCR
TS
CSUNET
NSFNET
SAIC
USD
GA
AJA
Qlcom
USIU
Korea
BIOSYM
Horizon
Xerox
SDSU
AGI
Backbone
CERF 1544
CERF 56
Peer Network
Dial n' CERF
Backbone
1544 Kbps
56 Kbps
5/23/91
Esnet IP
Esnet Decnet

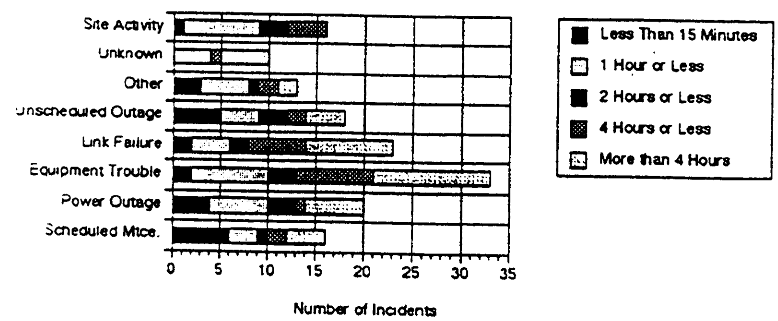

CERFnet Outages by Type
Data from 6-4-90 to 12-24-90

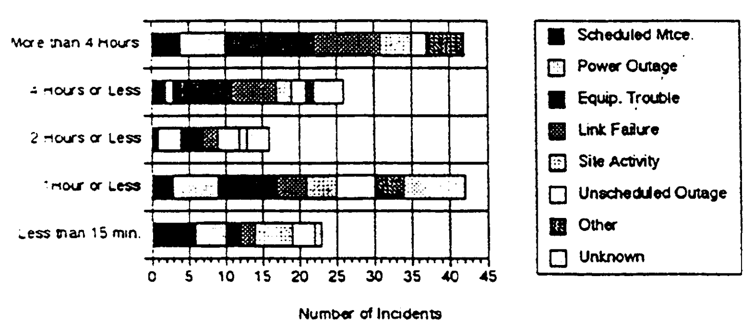

CERFnet Outages by Type Data
from 6-4-90 to 12-24-90

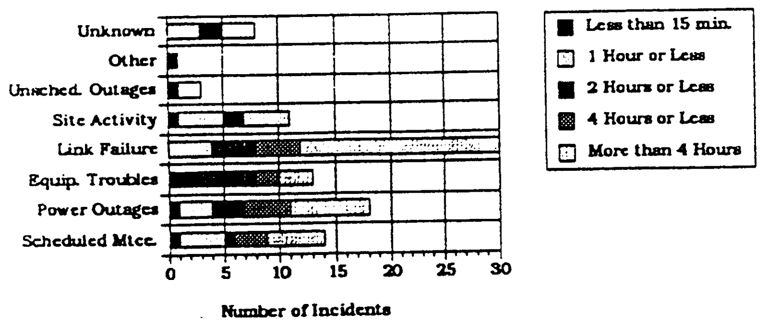

CERFnet Outages by Type
Data from 2-4-91 to 7-14-91

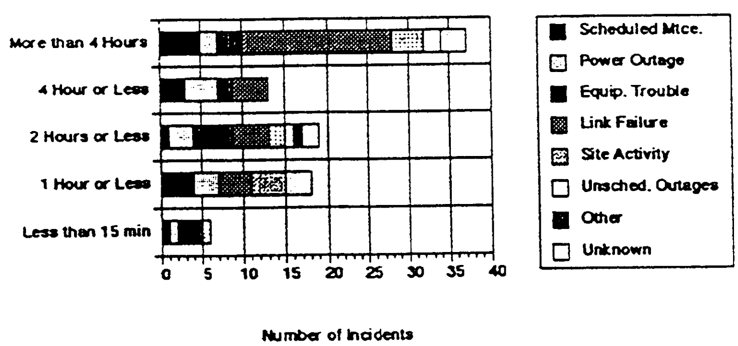

CERFnet Outages by Type
Data from 2-4-91 to 7-14-91

4.2 CICNet Report

Presented by J. Paul Holbrook/CICNet

CICNet Overview

J. Paul Holbrook
Technical Services Manager
CICNet, Inc.

Tom Easterday
CICNet Network Engineer
The Ohio State University
Academic Computing Services
CICNet Network Operations Center

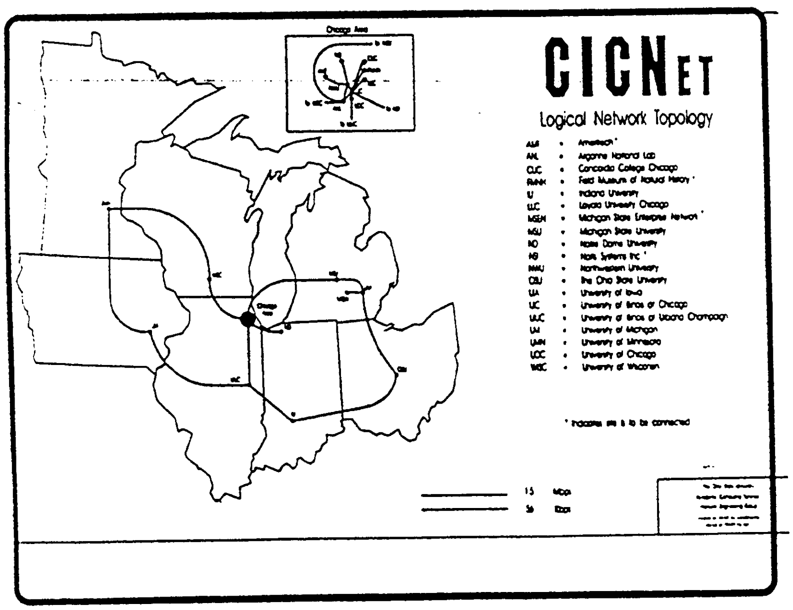

History of CICNet

- Began as a Project of the Committee on Institutional Cooperation (CIC) Panel on Computing and Telecommunications Policy
- Became Operational in March, 1989

CICNet, Inc.

Organizational Structure of CICNet

- 12 Member Board of Directors - Charter Members and the Director of the CIC
- Central Staff Located in Ann Arbor, MI
- Technical Board - Representatives From Each Member Organization

CICNet, Inc.

NOC Services Contract

- AT&T / The Ohio State University Team
- AT&T and Bell Labs
  Overall project management
  Monthly traffic reports
  Network Engineering
- The Ohio State University
  NOC
  Maintenance and upgrades
  New member connection support
  General Technical Analysis

CICNet, Inc.

CICNet Charter Members

Univ. of Chicago	Univ. of Minnesota
Univ. of Illinois-Urbana/Champaign	Univ. of Iowa
Univ. of Illinois-Chicago	Northwestern Univ.
Univ. of Michigan	Michigan State University
Univ. of Wisconsin - Madison	Indiana University

Ohio State University

Other CICNet Members

Loyola University of Chicago

Concordia University

Notre Dame University

Argonne National Labs

Field Museum

MCI

NOTIS Systems

Ameritech

Service Policies

- 7/24 Monitoring and Management Up To Ethernet Interface At Member Site
- Router On Member Site Is Owned and Managed By CICNet
- 800 Number for Trouble Reports
- Maintenance and Repair Procedures Require Designated Member Site Contact

Technical Features

- T-1 Backbone (1.5 mbps)
- cisco routers
- Majority of Circuits Supplied By MCI
- ESF Monitoring on all MCI Circuits
- Ring Architecture
- High Degree of Redundancy
- 2 T-1s To Each Backbone Node
- 3 NSFNET T-3 Nodes

CICNet, Inc.

Protocol Support

- TCP/IP
- DECNET
- OSI (Trial Stage)
- Peer With ESNET

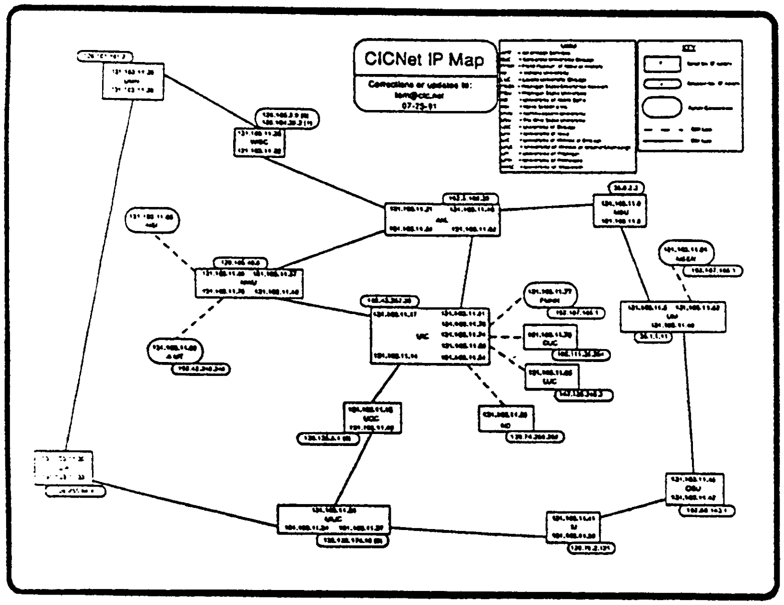

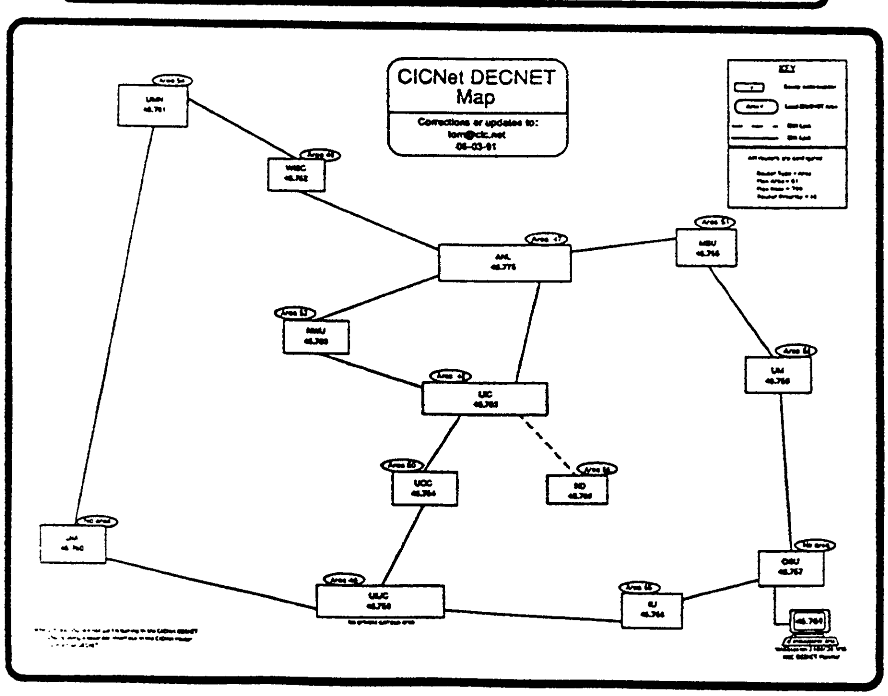

OVERALL NETWORK USAGE REPORT

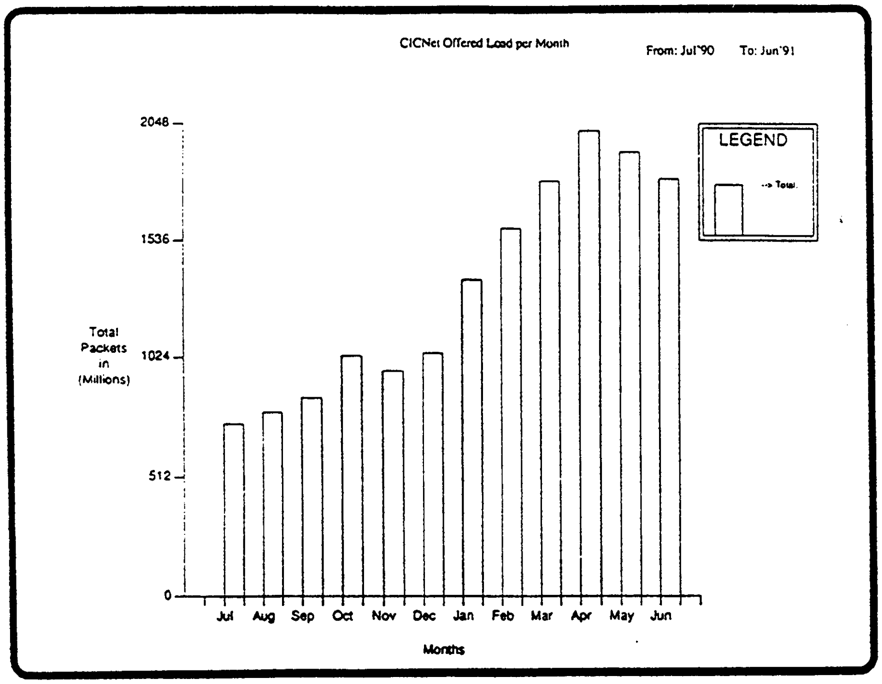

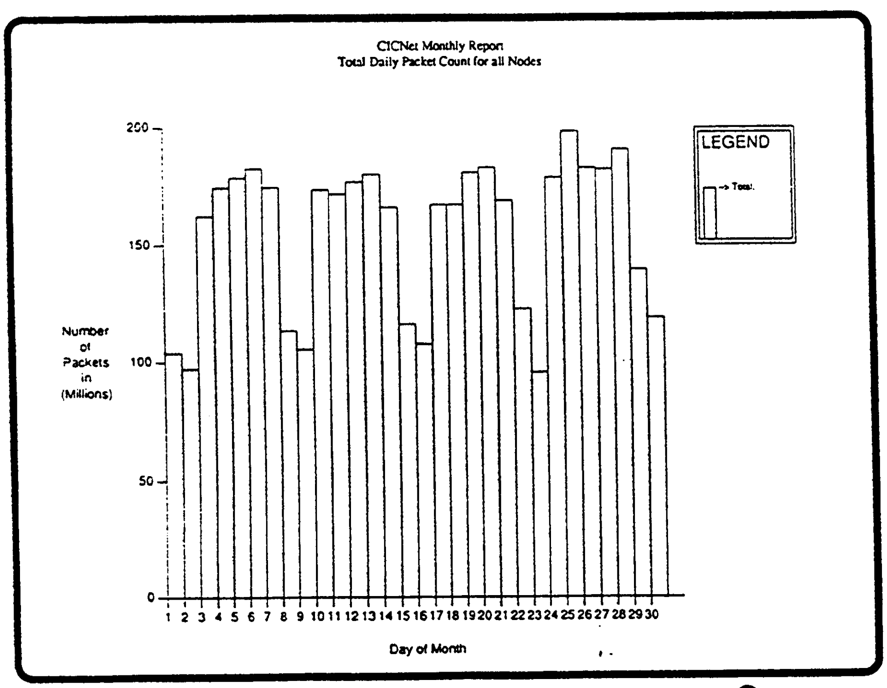

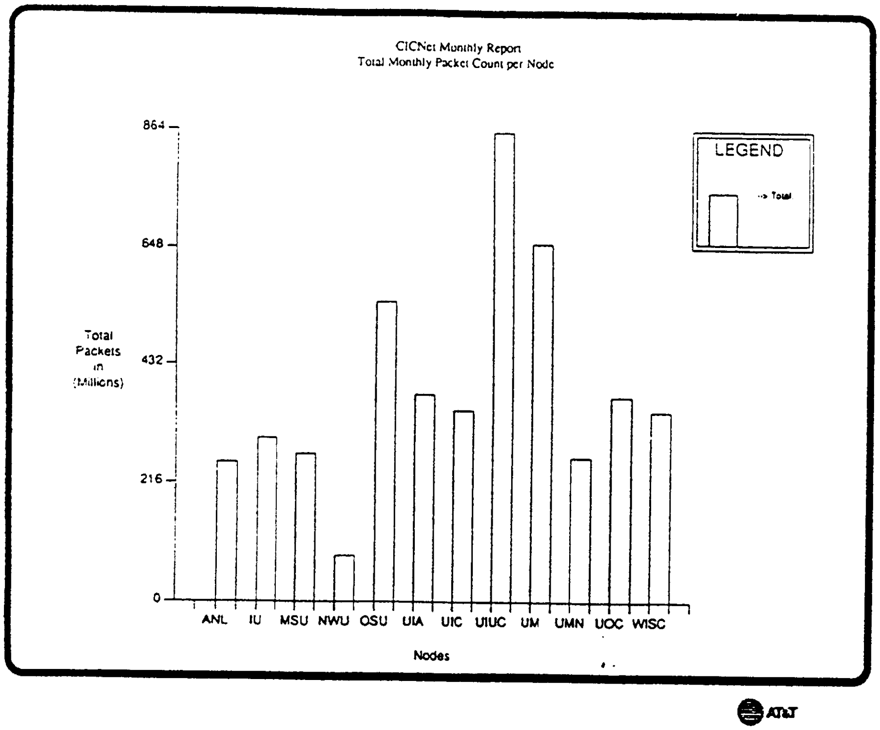
CICNet Monthly Report
Total Monthly Packet Count per Node
864
648
432
216
0
Total Packets in (Millions)
LEGEND
Total
ANL
IU
MSU
NWU
OSU
UIA
UIC
UIUC
UM
UMN
UOC
WISC
Nodes
AT&T

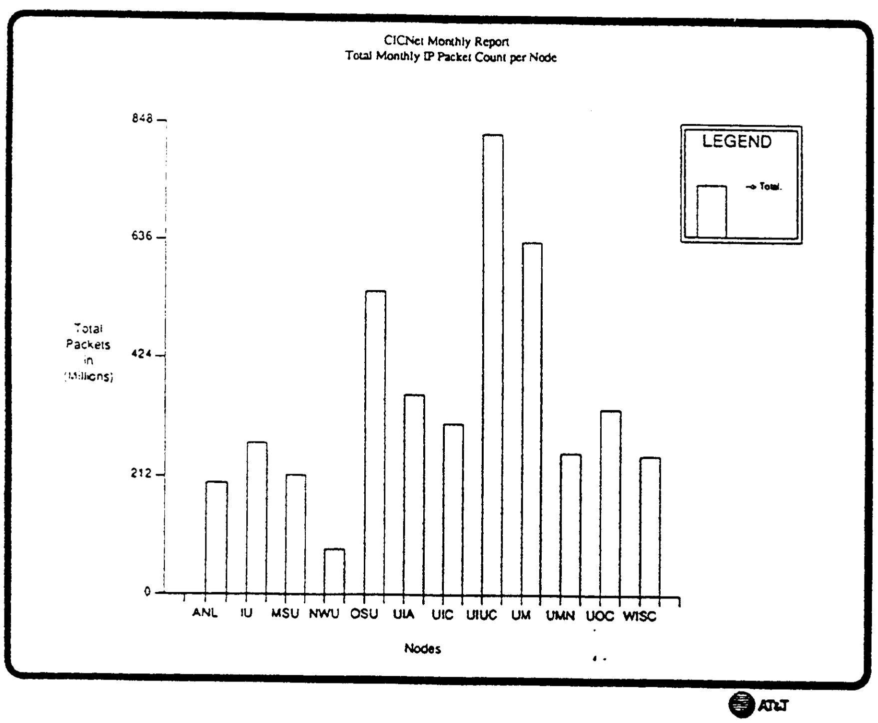
CICNet Monthly Report
Total Monthly IP Packet Count per Node
848
636
424
212
0
Total Packets in (Millions)
LEGEND
Total
ANL
IU
MSU
NWU
OSU
UIA
UIC
UIUC
UM
UMN
UOC
WISC
Nodes
AT&T

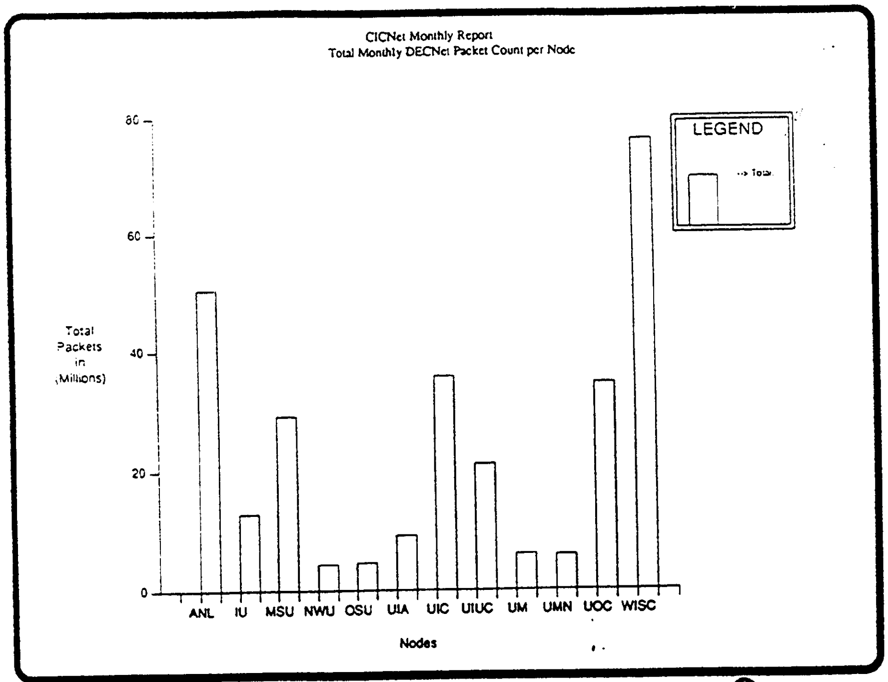
CICNet Monthly Report
Total Monthly DECNet Packet Count per Node
LEGEND
--> Total
Total Packets in (Millions)
80
60
40
20
0
ANL
IU
MSU
NWU
OSU
UIA
UIC
UIUC
UM
UMN
UOC
WISC
Nodes

AT&T

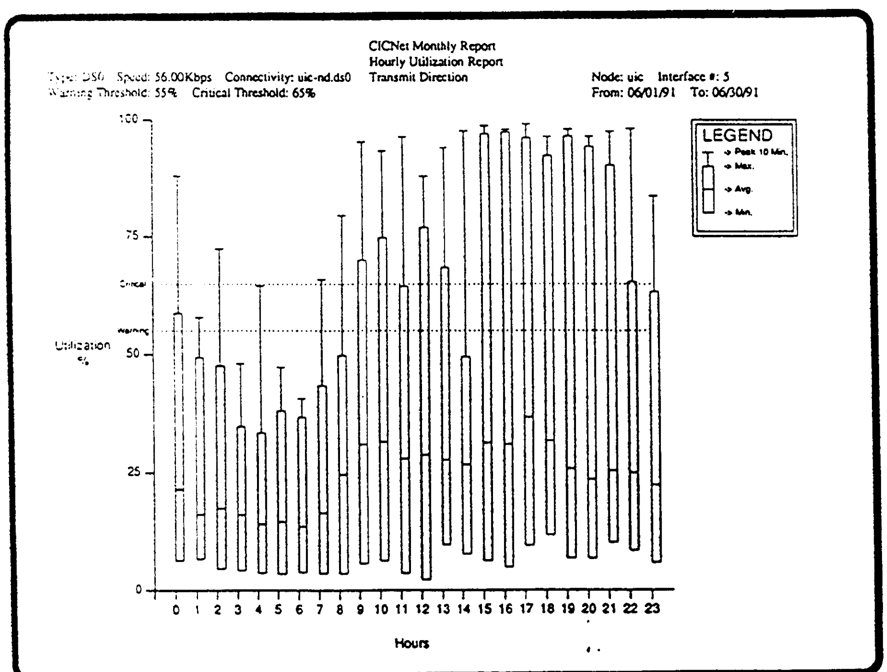
CICNet Monthly Report
Hourly Utilization Report
Transmit Direction
Type: DS0 Speed: 56.00Kbps Connectivity: uic-nd.ds0
Warning Threshold: 55% Critical Threshold: 65%
Node: uic Interface #: 5
From: 06/01/91 To: 06/30/91
LEGEND
-> Peak 10 Min.
-> Max.
-> Avg.
-> Min.
Utilization %
100
75
Critical
Warning
50
25
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Hours

AT&T

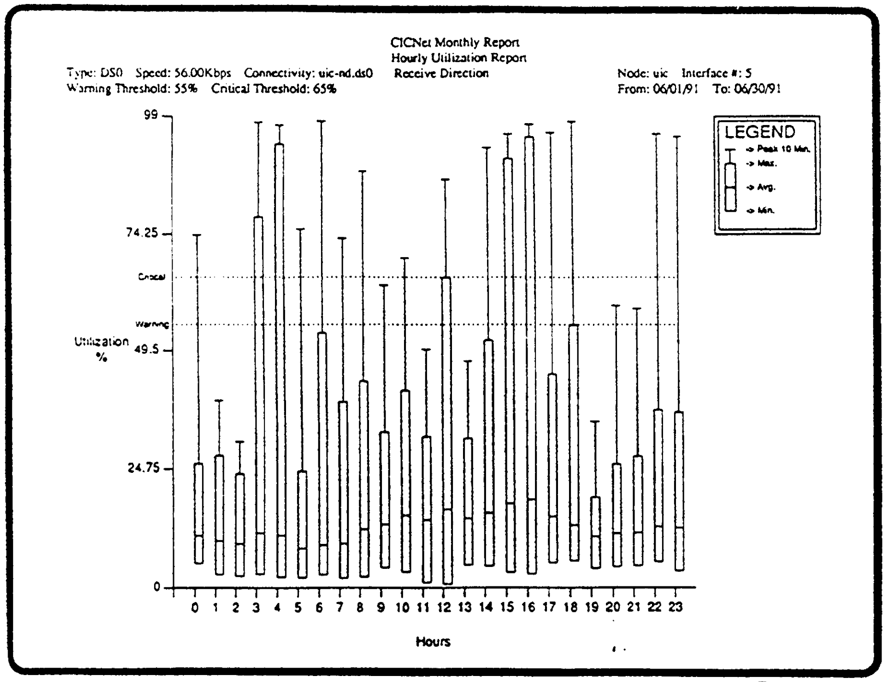
CICNet Monthly Report
Hourly Utilization Report
Type: DS0 Speed: 56.00Kbps Connectivity: uic-nd.ds0 Receive Direction
Node: uic Interface #: 5
Warning Threshold: 55% Critical Threshold: 65%
From: 06/01/91 To: 06/30/91
LEGEND
-> Peak 10 Min.
-> Max.
-> Avg.
-> Min.
Utilization %
99
74.25
Critical
Warning
49.5
24.75
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Hours

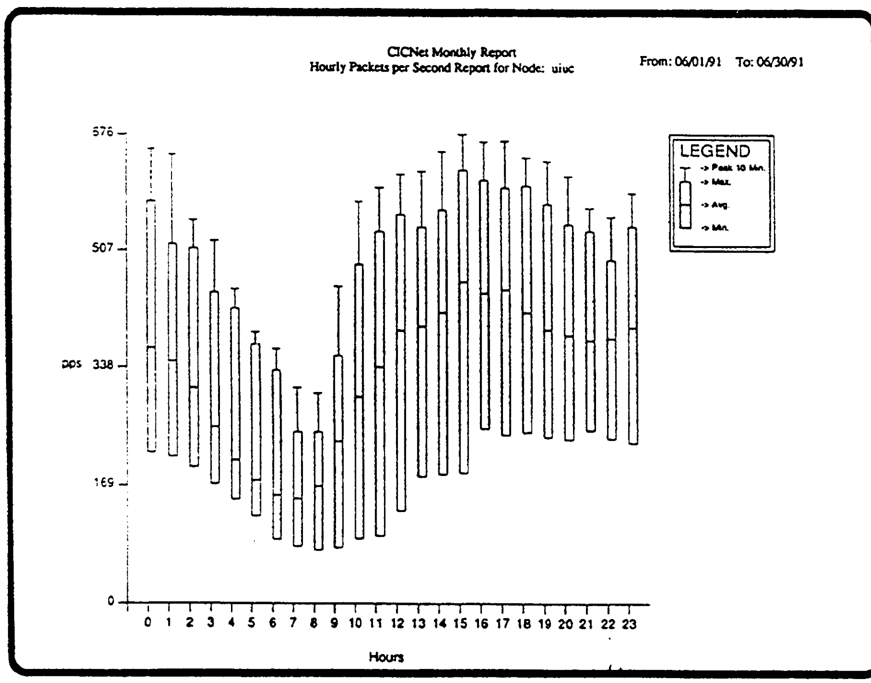
CICNet Monthly Report
Hourly Packets per Second Report for Node: uiuc
From: 06/01/91 To: 06/30/91
LEGEND
-> Peak 10 Min.
-> Max.
-> Avg.
-> Min.
pps
576
507
338
169
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Hours

4.3 DDN MILNET Report

Presented by Kathleen Huber/BBN

DDN MILNET

Kathleen Huber

August, 1991

BBN Communications
A Division of Bolt, Beranek and Newman, Inc.

AGENDA

- MILNET Statistics
- Internet Growth
- Mailbridges

2

MILNET Statistics

3

Milnet Statistics / Devices, Trunks

- 227 PSN's
- 1729 Hosts
  - 335 IP Routers
- 476 Trunks -- 75 at 9.6 kbps
  26 at 19.2 kbps
  32 at 50 kbps
  322 at 56 kbps
  15 at 64 kbps
  5 at 100 kbps

4

Current MILNET Throughput

- Packet Statistics
  - Week-long average kbit /second (snd/rcv) = 705.2
  - Peak-hour average kbit /second (snd/rcv) = 1068.6

5

INTERNET GROWTH

6

INTERNET GROWTH SUMMARY

- 2619 Networks Advertised By BMILBBN on 7/24/91 - 12:00 Noon
- 9605 Networks Registered
- 30,146 Network Numbers Assigned

7

NUMBER OF NETWORKS

LINEAR

DECEMBER 1983-JULY 1991

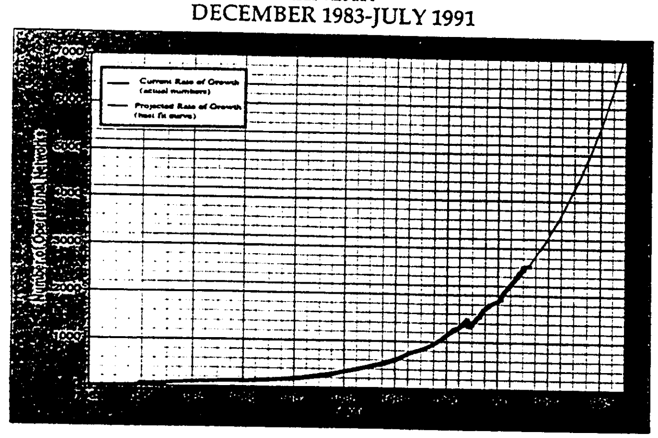

NUMBER OF NETWORKS

LOGARITHMIC

DECEMBER 1983-July 1991

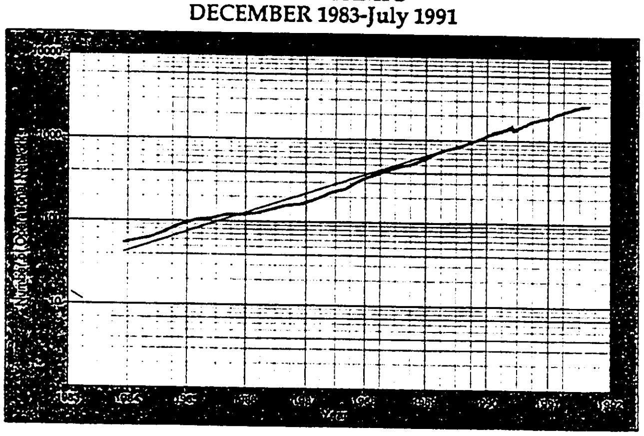

MAILBRIDGES

10

EGP NEIGHBOR COMPARISON

DIRECT NEIGHBORS

	Aug 90'	Oct	Dec	Feb 91'	Apr	Jun	Jul
BMILAMES	55	54	58	61	52	54	51
BMILBBN	26	38	40	62	78	89	89
BMILDCEC	105	85	89	88	95	155	155
BMILISI	52	61	62	60	58	63	69
BMILOAK (BMILLBL)	72	70	59	78	114	177	180
BMILMTR	61	58	63	56	57	57	54
BMILRAN	33	42	38	50	70	76	76

11

TRAFFIC SUMMARY COMPARISON

	Avg. Pkts/Day Forwarded		Avg. Pkts Dropped	
	Dec/Feb	March/July	Dec/Feb	March/July
BMILAMES	4,971,389	4,417,258	0.7%	0.06%
BMILBBN	173,642	302,469	0.8%	0.0%
BMILDCEC	350,733	432,644	8.2%	0.0%
BMILISI	291,263	238,492	2.9%	0.0%
BMILOAK (BMILLBL)	176,996	300,352	6.1%	0.05%
BMILMTR	2,091,064	3,661,255	1.2%	0.03%
BMILRAN	151,498	232,519	2.4%	0.02%

12

CURRENT STATUS

- Implementing strategy for long-term growth effects
- April deployment Patch 10, "Token Queue"
  - Token loss problem causing EGP message failure corrected
  - EGP Peer poll queue management streamlined
  - Support for 2700 networks
  - Support for 510 external gateways
  - EGP update construction time cut in half

13

CURRENT STATUS (cont.)

- July deployment Patch 11, "3800 Nets"
  - Support for 3800 networks
  - Limited route distribution - non-MILNET networks excluded from EGP update; default gateway advertised
  - Excessive polling of peers corrected by adjustment to Mailbridge clock

14

SUMMARY

- Current Efforts
  - Implementing strategy for long-term growth effects
  - Advertisement of network 0 to limit size of EGP updates

15

4.4 ESnet Report

Presented by Tony Hain/LLNL

ESnet Status Report

IETF - Atlanta

July 1991

Anthony L. Hain

Associate Network Manager

ESnet / NERSC

ESnet

PAST ACTIVITIES:

Replaced NNT circuits with FTS-2000

Added Connections To OSTI, NIFS via PACCOM / Hawaii

Deployed X.25 switching at FNAL, MIT, LLNL, BNL, SLAC

German 128kbps circuit...Up 6/3

Deployment of Initial CLNP Routing Across Backbone

ESnet

STATS:

34 Routers Managed

85 Directly Connected Networks

798 Regional Connected Networks

809 Networks via Other Backbones

.95G Packets Received

78% IP / 22 % DECnet

ESnet

PLANNED ACTIVITIES:

Investigating Participation in InterOP OSI Demonstration

Frame Relay Trial

ESnet

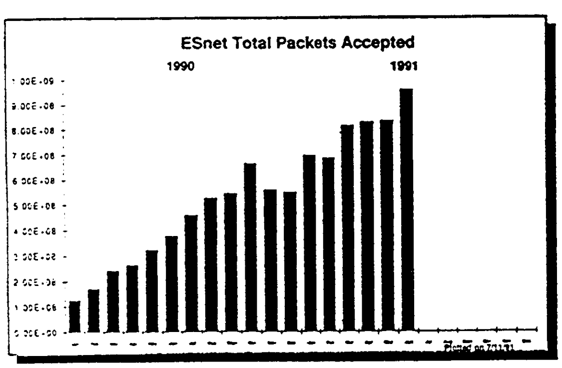

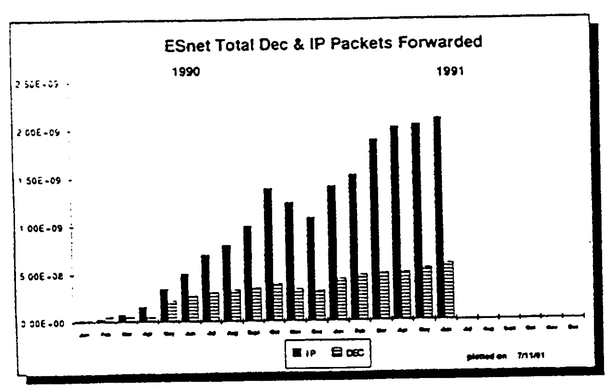

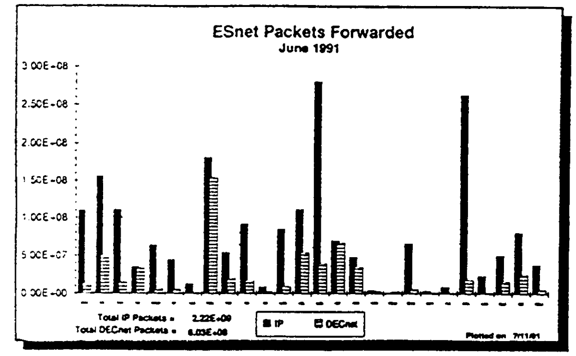
ESnet Packets Forwarded
June 1991
3.00E+08
2.50E+08
2.00E+08
1.50E+08
1.00E+08
5.00E+07
0.00E+00
Total IP Packets = 2.22E+09
Total DECnet Packets = 6.03E+08
IP
DECnet
Plotted on 7/11/91

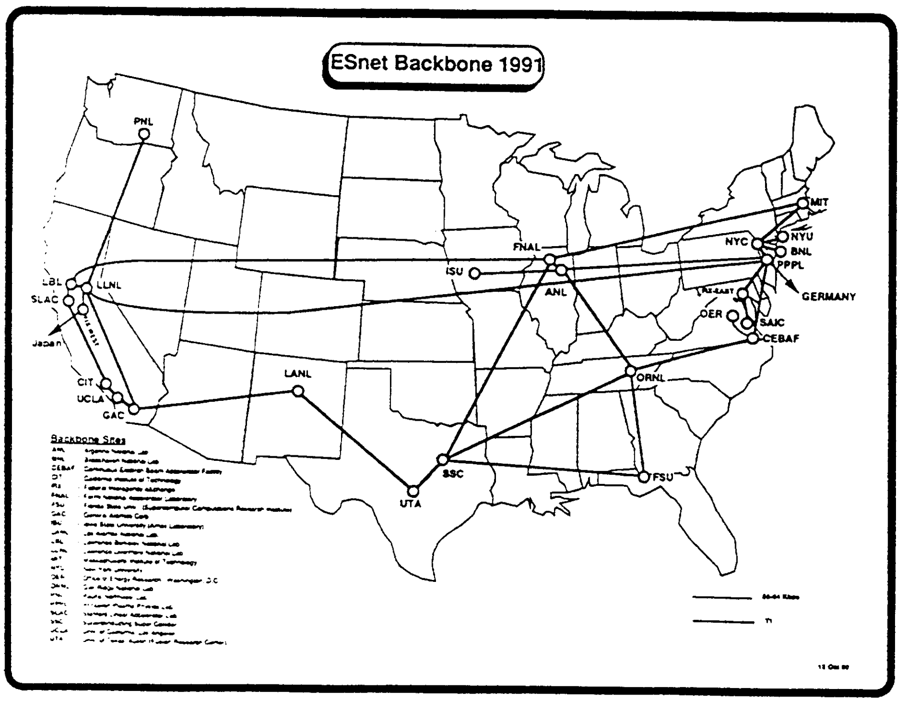
ESnet Backbone 1991
PNL
LBL
SLAC
LLNL
Japan
CIT
UCLA
GAC
LANL
UTA
SSC
ISU
FNAL
ANL
ORNL
FSU
NYC
MIT
NYU
BNL
PPPL
GERMANY
OER
SAIC
CEBAF
Backbone Sites

4.5 RIPE Report

Presented by Rudiger Volk/University of Dortmund

RIPE Report

Rüdiger Volk

University of Dortmund, Germany

Network Topology
Statistics - database
- DNS

Root Name Server @ Stockholm
Eastern Europe (and others)
RIPE NCC

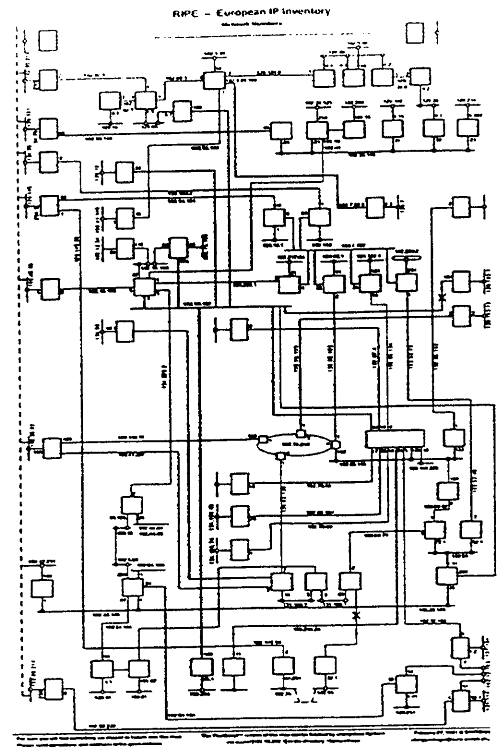

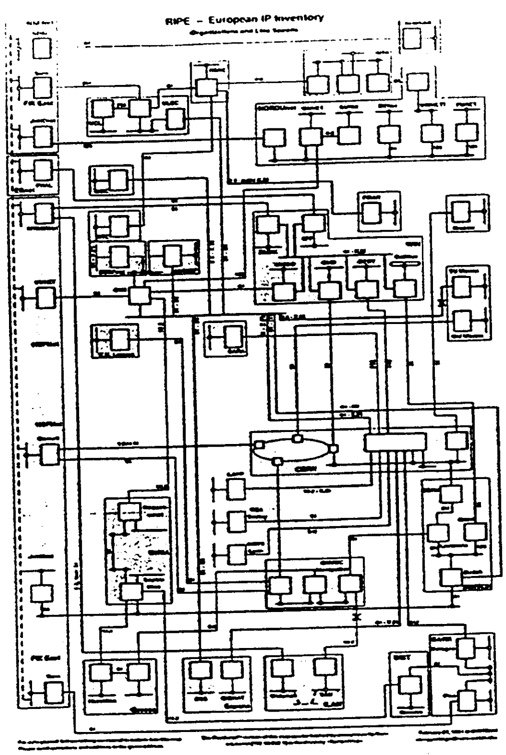

Legend to the RIPE European IP Inventory – Abbreviations

Abbreviation	Country	Organisation, Town
BelWü	DE	Baden Württembergs Extended Lan
CEA	FR	Commissariat à l'Energie Atomique
CERN		Organisation Européenne pour la Recherche Nucléaire, Geneva
CIEMAT	ES	
CNUSC	FR	Centre National Universitaire Sud de Calcul, Montpellier
CWI	NL	Centrum voor Wiskunde en Informatica, Amsterdam
DENet	DK	Danish Ethernet Network
DESY	DE	Deutsches Elektronen Synchrotron, Hamburg
DFN	DE	Deutsches Forschungsnetz
DIST	IT	University of Genova
EASInet		European Academic Supercomputer Initiative
ESnet	US	Energy Science Network
EUnet		European UNIX Network
EurOpen		European Forum or Open Systems (formerly known as EUUG)
FGAN	DE	Forschungsgesellschaft f. Angew. Naturwissenschaften, Wachtberg
FNAL	US	Fermi National Accelerator Laboratories
FORTH	GR	Foundation of Research and Technology, Hellas
FUNET	FI	Finnish University Network
GARR	IT	Gruppo Armonizzazione Reti della Ricerca
GMD	DE	Gesellschaft für Mathematik und Datenverarbeitung
ILAN	IL	Israeli Academic Network
INRIA	FR	Institut National de Recherche en Informatique et Automatique
IN2P3	FR	Institut National de Physique Nucléaire et Physique des Particules, Lyon
IPP	DE	Max Planck Institut für Plasmaphysik, Garching
IRIS	ES	Spanish Academic & Research Network
IXI		International X.25 Infrastructure
JvNCnet	US	John von Neumann Supercomputer Center Consortium Network, Princeton
K. U. Leuven	BE	Katholieke Universiteit Leuven
LAPP	FR	
NIKHEF	NL	National Institute for Nuclear and High-Energy Physics
NORDUnet		Nordic Research and Development network
NORSAR	NO	
NSFNET	US	National Science Foundation Network
NTA	NO	Norwegian Telecommunications Administration
NYSERNet	US	New York State Educational Research Network
RAL	GB	Rutherford Appleton Laboratory
RIPE		Réseau IP Européenne
RSRE	GB	Royal Signals Research Establishment, Malvern
SARA	NL	Academic Computer Services Amsterdam
SDAC	US	
STC	NL	SHAPE Technical Center
SUNET	SE	Swedish University Network
Sura	US	
SURFnet	NL	Dutch higher Educational and Research Network
SURIS	IS	Association of Research Networks in Iceland
SWITCH	CH	Swiss Telecommunication System for higher Education and Research
TAU	IL	Tel Aviv University
UCL	GB	University College London, London
UKC	GB	University of Kent at Canterbury, Canterbury
ULCC	GB	University of London Computer Center, London
UNIDO	DE	Universität Dortmund, Dortmund
UNINETT	NO	University Network
UUNET	US	
WIN	DE	Wissenschafts Netz
XLINK	DE	eXtended Lokales Informatik Netz Karlsruhe

Please send additions and modifications to: <denggenhager@verw.switch.ch> Feb 27, 1991 © SWITCH

Legend to the RIPE European IP Inventory – Pictures

Legend to the "Speed-Map"

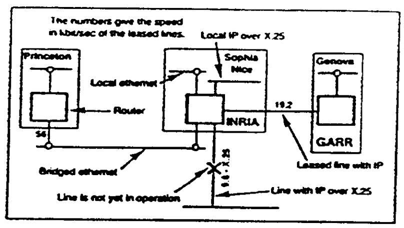

Legend to the "Netnum-Map"

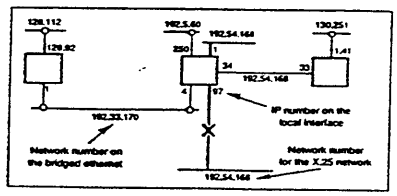

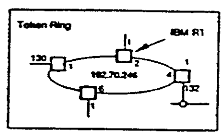

RIPE Statistics - hosts counted by DNS

	11/90	6/91	7/91
at	783	1728	1921
be	2	147	*
ch	6284	9047	9592
de	3936*	14692	16028
dk	858	805	1080
es	3	512	597
fi	3781	7330	7740
fr	3436	7628	9536
gr	105	161	96*
is	16	160	172
it	649	1298	*
il	661	1393	1576
nl	5082	6404	7057
no	3317	6340	7983
se	5190*	4191?	3800?
uk	206	1433	2764

MX only TLDs: CS, HU, IE, PL, SU, YU, PT

others: AL, BG, RO, TN, tiny ones

TLDs by Size of Host Population (6/91)

DE	16028
CH	9592
FR	9536
NO	7983
FI	7744
NL	7057
SE	
UK	2764
AT	1921
IL	1576
IT	1298
DU	1080
ES	597
GR	161
BE	147

RIPE Database

	11/90	6/91
Persons	643	1270
Networks with connectivity	670	1053
LOCAL	393	470
RIPE	265	539
NSF	183	373

average 10 updates/day

Chapter 5

IETF Protocol Presentations

5.1 IP over Frame Relay Report

Presented by Caralyn Brown/Wellfleet

Frame Relay is a wide area technology which connects many destinations with virtual circuits, creating mesh of point-to-point connection between end stations. These virtual circuits are identified by a Data Link Connection Identifier (DLCI) which has only local significance. That is, the DLCI identifying a particular virtual circuit may be different at each end.

The goals for developing the Multiprotocol Over Frame Relay document were simple. We need a method for transferring data from multiple protocols across the same virtual circuit as the same time. It makes no sense to require a customer to set up (and pay for) multiple connections to the same destination just to be able to send several different types of data.

The second goal was to include bridged data as just another "protocol" being transported. Therefore, the encapsulation scheme must also allow the ability to distinguish between protocol (routed) and bridged data.

The encapsulation format decided upon is shown in the Frame Format slide. The DLCI and flagging sequence are part of the definition of frame relay. We added the control field for compatibility with other LAPD networks, but also for later expansion and possible use of LLC2 acknowledged transfer mode.

The Optional Pad field is just that OPTIONAL. It is used to force alignment of the data packet that follows. The choice of alignment is arbitrary and is used only for convenience. There may be multiple pads included.

The Network Layer Protocol Identifier (NLPID) is used to indicate what protocol follows in the data portion. NLPIDs are administered by ISO and CCITT. IP has NLPID 0xCC. Though many protocols have NLPIDs, many widely used protocols, such as Decnet and AppleTalk, do not. There is, however, a NLPID defined for the SNAP header, and we can use this to encapsulate those protocols that do not have a NLPID defined.

To send IP traffic over frame relay, one uses the encapsulation format shown in the format slide and used a NLPID of 0xCC. Notice, however, that there are two ways to encapsulate IP. The second is to use a SNAP encapsulation. In order to avoid confusion, we must select one method. It is always preferable to use the NLPID value when one is defined. If there is no NLPID defined for a protocol, use the SNAP encapsulation.

Notice in the drawing of IP encapsulated within the SNAP format, that the IP packet is not word aligned. We could have used a pad field just before the NLPID field to align the packet.

Bridging: There are provisions for indicating bridged packets by using a SNAP NLPID. 802.1 has assigned values for bridged packets.

Address Resolution within frame relay is an interesting problem because the stations do not necessarily have a source address. The drawing shows how a simple three node frame relay network may be configured. Notice how the stations have different addresses (DLCIs) depending on which station they are talking to. Because of the way that a frame relay network modifies addresses as a packet traverses the network, we can make ARP work even in this network. For example, if station A wishes to resolve address 192.32.1.3 (node C), it will send an ARP message with its protocol address and the protocol address it wishes to resolve and sends it to all relevant stations.

As the packet traverses the network from A to C the Q.922 address (and therefore the DLCI) will be modified so that upon arrival at C, the address will reflect where the packet came from (DLCI 60). Station C may use this address as a mapping to destination IP address 192.32.1.1. The ARP response works exactly the same. When the packet arrives at station A, it takes the address from the frame relay header and inserts it as the source hardware address of station C.

Yes we know this violates layering a little. But it does work! It works for both locally and globally addressed networks and it works equally well for Reverse ARP. If and when we get multicasting worked out, we may be able to use this instead of simulating it by sending the packets to each destination.

We tested the Multiprotocol Interconnect Over Frame Relay (including ARP) both in house and at ComNet West. In house we had a Williams Telecom switch with two connections. One went to a Proteon router and the other to a Wellfleet router. Both routers were using LMI with the switch. We were successfully able to pass IP data between the two routers. The same configuration was tried at ComNet using a Netrix switch. This test was also successful.

Multiprotocol Interconnect Over Frame Relay Networks

Caralyn Brown
Wellfleet Communications, Inc.
15 Crosby Drive
Bedford, MA 01730
cbrown@wellfleet.com

C8-07-25 WELLFLEET

Interconnect Issues - Goals

- Interconnection of many protocols over one DLC
- Ability to transmit both bridged and routed packets over the same DLC

C8-07-25 WELLFLEET

Frame Format

LAPD flag (7E hexadecimal)
DLCI
Control (UI = 0x03)
Optional Pad (0x00)
NLPID
. . . Data . . .
LAPD Frame Check Sequence . (two octets)
LAPD flag (7E hexadecimal)

DLCI
10-bits encoded in two octets. In some networks DLCIs may optionally be increased to three or four octets.

Optional Pad
Used to align data portion of the message. This is NOT required.

NLPID
Values are defined in ISO/IEC TR 9577.

IP	0xCC
Voice and fax	0x44

C8-07-25 WELLFLEET

IP Over Frame Relay

- Using the encapsulation described, there are ~~three~~ two distinct ways to encapsulate IP
  - NLPID of 0xCC
  - ~~NLPID of 0xOE and ethertype~~
  - NLPID of 0x80 and SNAP header
- Must select one method as the preferred method.

C8-07-25 WELLFLEET

IP Over Frame Relay (Continued)

1. NLPID value indicating IP

Q.922 Address	
Control 0x03	NLPID 0xCC
IP Packet	

2. NLPID value indicating SNAP

Q.922 Address	
Control 0x03	NLPID 0x80
SNAP header	
IP Packet	

Bridging Over Frame Relay

- Format packet using 0x80 indicating SNAP
  - set OUI to 0x0080C2
  - PID fields reserved by 802.1 Indicating Origin Media

PID Values

with preserved FCS	w/o preserved FCS	Media
0x0001	0x0007	802.3/Ethernet
0x0002	0x0008	802.4
0x0003	0x0009	802.5
0x0004	0x000A	FDDI
0x0005	0x000B	802.6
0x0006	0x000C	802.9

- Encapsulated MAC headers using the formats specified in RFC 1220.
- Origin media type is media from which packet was forwarded to Frame Relay NOT where the packet originated.

C8-07-25 WELLFLEET

Address Resolution

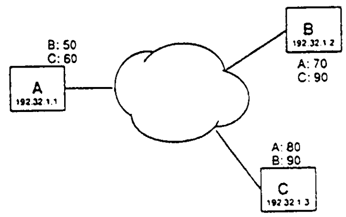

A → B	DST DLCI = 50
A → C	DST DLCI = 60
B → A	DST DLCI = 70
B → C	DST DLCI = 90
C → A	DST DLCI = 80
C → B	DST DLCI = 90

WELLFLEET communications

CB-07-25

Address Resolution

- Station A ARPs for station C
- No source hardware address

request

ar$sha - 0
ar$spa - 192.32.1.1
ar$tha - ?
ar$tpa - 192.32.1.3

- Station C swaps target and source source addresses to form response

response

ar$sha - 0
ar$spa - 192.32.1.3
ar$tha - 80
ar$tpa - 192.32.1.1

- Station C receives request
- Use DLCI from header as source hardware address

request after modification

ar$sha - 80
ars$spa - 19232.1.1
ar$tha - ?
ar$tpa - 192.32.1.3

- Station A receives response
- Use DLCI from header as source hardware address

response after modification

ar$sha - 60
ar$spa - 192.32.1.3
ar$tha - 80
ar$tpa - 192.32.1.1

WELLFLEET communications

CB-07-25

Other Address Resolution Items

- Allows dynamic resolution of protocol addresses using methods already defined in the ARP RFC.
- Works for both locally and globally addressed networks with minimal layer violations.
- Reverse ARP (RARP) works the same way.
- Broadcasting requests
  - Must simulate if no multicast
  - Use multicast if DLCI is modified through network.
- Possible extension for Inverse ARP

WELLFLEET communications

CB-07-25

Interoperability Test

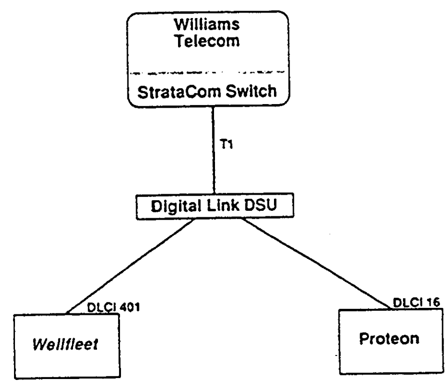

WELLFLEET communications

CB-07-25

Inverse Address Resolution Protocol

- Describes additions to ARP (and AARP) to allow a station to request a protocol address given a hardware address.
- Specifically applies to Frame Relay where a Data Link Connection Identifier (DLCI) corresponds to a hardware address.
- May apply to any network that provides destination hardware addresses without indicating corresponding protocol addresses.

WELLFLEET communications

CB-07-25

Inverse ARP Packet Format

ar$hrd	16 bits	Hardware type
ar$pro	16 bits	Protocol type
ar$hln	8 bits	Byte length of each hardware address (n)
ar$pln	8 bits	Byte length of each protocol address (m)
ar$op	16 bits	Operation code
ar$sha	nbytes	Source hardware address
ar$spa	mbytes	Source protocol address
ar$tha	nbytes	Target hardware address
ar$tpa	mbytes	Target protocol address

-- Possible values for hardware and protocol types are the same as those for ARP and may be found in the Assigned Numbers RFC [3].

-- Length of the hardware and protocol address are dependent on the environment in which InARP is running. Frame Relay networks will use 2, 3, or 4.

-- Operation code indicates the type of message, request or reply.

InARP request = 8 ~~InAARP request = 8~~
InARP reply = 9 ~~InAARP reply = 9~~

WELLFLEET communications

CB-07-25

InARP Operation

- Operates essentially the same as ARP.
- No broadcasting necessary because destination is already known.
- InARP Request formatting
  - fill in source hardware address
    source protocol address
    target hardware address
  - Zero fill the target protocol address
- Encapsulate for specific network and send the request directly.
- Upon receipt of a request, a station should
  - store the protocol address/hardware address mapping of requester
  - Format a reply by filling in the protocol address requested.
  - Send the reply.

InARP Operation

Properly Formatted InARP Request

- ar$hrd - 0x00F the value assigned to Frame Relay
- ar$pro - protocol type for which you are searching (i.e., IP = 0x0800)
- ar$hln - up to 4 for extended DLCI addressing
- ar$pln - byte length of protocol address for which you are searching (for IP = 4)
- ar$op - 8; InARP request
- ar$sha - DLCI of requesting station
- ar$spa - protocol address of requesting station
- ar$tha - DLCI corresponding to newly announced PVC
- ar$tpa - 0; This is what we are looking for

The InARP response will be completed similarly.

- ar$hrd - 0x00F the value assigned to Frame Relay
- ar$pro - protocol type for which you are searching (i.e., IP = 0x0800)
- ar$hln - up to 4 for extended DLCI addressing
- ar$pln - byte length of protocol address for which you are searching (for IP = 4)
- ar$op - 9; InARP response
- ar$sha - DLCI of responding station
- ar$spa - protocol address of requested
- ar$tha - DLCI if requesting station
- ar$tpa - protocol address of requesting station

5.2 Site Security Handbook Report

Presented by J. Paul Holbrook/CICNet

The Site Security Handbook

RFC1244 / FYI #8

A product of the Site Security Policy Handbook Working Group (SSPHWG)

J. Paul Holbrook
Co-chair, SSPHWG

Technical Services Manager
CICNet, Inc.

Site Security Handbook - 1

The Problem - I

Security is a People problem

Security is up to the users, but they can't and won't do it all themselves

Some sites are prepared technically but lack appropriate policies to deal with attacks

Lack of Policy can lead to kneejerk reactions and trouble

Site Security Handbook - 2

The Problem - II

Internet-specific issues:

- Threats: over-the-shoulder vs. across-the-ocean
- Rapid spread of information and threats
- Trust and sharing are central to Internet
- Cooperative approach to dealing with security problems
- Sites can be unwitting hosts to attacks on other sites

Commercial firms consider security a key issue

Site Security Handbook - 3

Site Security Policy Handbook Working Group

Co-chairs:
- J. Paul Holbrook, CICNet
- Joyce K. Reynolds, USC/ISI

Started March 90

Joint effort by User Services Area and Security Area

Site Security Handbook - 4

Site Security Handbook

RFC 1244 / FYI 8

Informational, not standard

Site Security Handbook - 5

Handbook structure

Define security policy

Define procedures and mechanisms to support policy

Handle incidents

Follow-up

Annotated bibliography

Site Security Handbook - 6

Security Policy

Sites are different: e.g., commercial vs. university

Must have buy-in from decision makers

Risk analysis key to cost-effective security

Security is no longer just an internal matter

Site Security Handbook - 7

Security mechanisms and procedures

Policy should guide mechanisms and procedures

Procedures & Policies for:

- accounts and passwords
- what happens when policy is violated

Use common sense: why use C2 system if passwords are weak?

Site Security Handbook - 8

Incident Handling

Who gets called? Who calls the shots?

Protect and proceed vs. pursue and prosecute

Legal / investigative issues

Understanding and preparation is the antidote to fear and loathing

Site Security Handbook - 9

All midlevels should:

- Have policies regarding assisting sites during security incidents
- Encourage sites to have policies
- Maintain up-to-date 24 hour security contacts

Site Security Handbook - 10

All sites should:

- Define security policies with buy-in from decision makers
- Have appropriate mechanisms for dealing with security
- Have an agreed upon plan for dealing with security incidents
- Have some central person or group who watches security
  - ☞ Contact must be available and have backup

Site Security Handbook - 11

Future efforts

Create an executive summary

Encouraging use and feedback

Revision in one year (hopefully)

Site Security Handbook - 12

Acknowledgments

The Wonderful Joyce K. Reynolds

Richard Pethia and the CERT

Dave Curry (Purdue)
Sean Kirkpatrick (Unisys)
Tom Longstaff (LLNL)
Greg Hollingsworth (Johns Hopkins University)
Jeffrey Carpenter (University of Pittsburgh)
Barbara Fraser (CERT/CMU/SEI)
Fred Ostapik (SRI NISC)
Allen Sturtevant (LLNL)
Dan Long (BBN)
Jim Duncan (Pennsylvania State University)
Frank Byrum (DEC)

The Problem - I

The Problem - II

Site Security Policy Handbook Working Group

Site Security Handbook

Handbook structure

Security Policy

Security mechanisms and procedures

Incident Handling

All midlevels should

All sites should

Future efforts

Chapter 6

Technical Presentations

6.1 Security Issues and Directions

Presented by Steve Crocker/TIS

Security Issues and Directions

Steve Crocker

Area Director for Security

Security is Vital

- o "Even" the academic community needs security.
- o Commercial success is impossible without security.
  - Companies either avoid the Internet or they build firewalls.
- o Advanced services and applications are impossible without adequate security.

Is Security Important in the Internet?

- o Morris worm occurred in November 1989.
- o Much activity to improve security.
- o Are we safer now?

Internet Security

- o What would it mean for the Internet to be "secure"?

Public Confidence

- No widespread penetration.
- No widespread denial of service.
- Known risks.
- Understandable and believable local controls.

How Bad Is It?

o Internet now consists of:

- Several thousand networks.
- Several hundred thousand hosts.
- A couple million users.

"This is a large system!"

o The SEI/CERT estimates one (?) new break-in per day.

o Products are getting better - but only slowly.

o Protocols are getting better - but only slowly.

o Users are getting better - but only slowly.

o Hackers are getting better - not very slowly.

Internet Activities

o Policies, Guidelines, etc.

- Security Guidelines
- Site Security Handbook
- Etc.

o Protocols

- Privacy Enhanced Mail
- SNMP Security
- Telnet Authentication
- PPP Authentication
- Common Authentication Technology
- IP security options

A Broader View

o Policies, protocols, CERTs all help.

o Are they enough? No.

o What else is needed?

- Broader attack.

- Cooperation from all segments of the community.

- Mixture: common sense, focus of attention, some technical developments.

Computer Emergency Response Teams

o SEI/CERT in Pittsburgh established as result of Morris worm.

o Similar groups in several communities.

o "CERT" is now generic. Family of CERTs have regular meetings and close cooperation.

Two Big Problems

Two "Easy" Fixes

o Poor password management.

o Inadequate configuration control.

"These account for the overwhelming majority of security problems today."

Password Management

o All accounts need passwords (or better).

- Routers

- Network operation centers

- Dial-up access terminal servers

* Terminal access machines without dial-up access are mildly debatable.

Insist On Good Password Hygiene

o Passwords must be long enough and not easily "guessable."

o Passwords should be changed regularly.

o Passwords should not be shared among people or machines.

o Password generation or checking programs should be used.

Configuration Control

o The software that comprises all of the security relevant operation of the machine.

o The setting of all of the protection bits.

o Access control on all of the network services.

- NFS

- finger

- rlogin

The Buyers Prerogative

o Insist on "network safe" products:

- Configured safe out of the box.
- Simple, coherent security model.
- Simple, understandable security controls.
- Verifiable controls.

Simple, Coherent Security Model

o An explicit security perimeter.

o A small number of subjects and objects.

o A limited set of capabilities.

Simple, Understandable Security Controls

o Clear connection between the user's policies and the available security mechanisms.

o Small number of easily "settable" controls.

"Bad"

o Unix umask, ugo + rwx

- Very flexible.
- Very awkward.
- Incomprehensible consequences.

"Good"

- o Each user is the only one who can see or change files in his directory.
- o Groups share project directories.
- o Everyone sees company directories.
- o No outsiders see anything.

Verifiable Controls

- o It should be easy to check that the machine, LAN, campus, etc., is buttoned up.
- o It should be easy to review audit trails.
- o Active checks are needed.

Education

- o Education of users is essential.
  - Symposia.
  - Local training.
  - Each one teach one.

"Set Goals! Test for them!"

- o Training for network and host administrators is also essential.

"Visit a CERT!"

Education...

- o Community awareness is also important.
- o An annual security assessment will focus attention.
  - What are the real problems?
  - Is it getting better or worse?

Longer Term

- o New criteria for network safe computers.
  - Relevant security models.
  - Attainable levels of assurance.
  - Completely integrated into main product line.
- o Widespread use of commercial cryptography.
  - PEM.
  - Authentication Servers.
  - Hardware improvements including good random number generators.

Next Steps?

6.2 Guidelines for the Secure Operation of the Internet

Presented by Steve Crocker/TIS

Security Guidelines

from the

Security Policy Working Group

Steve Crocker

Security Guidelines

o We needed a statement of principles.

- Tie together and motivate the technical work.
- Provide leadership and guidance to:
  - . Users
  - . Sites,campuses, companies
  - . Regional & backbone operators
  - . Vendors

o IAB has no formal responsibility and no formal operational authority.

Security Guidelines

History

o SPWG formed/met in Pittsburgh.

o Met in Vancouver, St. Louis, Boulder & Reston.

o Considered spectrum of policies.

- Lists of security services.
- Separate statements for vendors, users, network operators, etc.
- Focused on short list of basic principles.

o "Policy" replaced with "Guideline".

Basic Principles

(1) Everyone is obliged to behave.

- It's not O.K. to take advantage of another person, system, etc., even if he's careless.

(2) Everyone is obliged to protect himself, his users, his customers.

- Despite principle (1), it's a hostile world.

(3) Cooperation is encouraged.

- We want to build a sense of community - a civilized world.

(4) The technology needs improvement.

The Words

(1) Users are individually responsible for understanding and respecting the security policies of the systems (computers and networks) they are using. Users are individually accountable for their own behavior. System operators are obligated to provide their policies to their users.

(2) Users have a responsibility to employ available security mechanisms and procedures for protecting their own data. They also have a responsibility for assisting in the protection of the systems they use.

(3) Computer and network service providers are responsible for maintaining the security of the systems they operate.

(4) Vendors and system developers are responsible for providing systems which are sound and which embody adequate security controls.

(5) Users, service providers and hardware and software vendors are expected to cooperate in the provision of security.

- *Mutual assistance*
- *Notification of others*
- *CERT's*
- *Investigation & prosecution*

Big question (when penetration is discovered) close the door or try to trap the intruder?

(6) Technical improvements in Internet security protocols should be sought on a continuing basis. At the same time, personnel developing new protocols, hardware or software for the Internet are expected to include security considerations as part of the design and development process.

Next Steps

o Edit document to smooth out nits.

o Internet Draft.

o Forward to IAB to publish as Informational RFC.

o Publish widely.

- *Trade magazines*
- *Professional communications*

o Encourage other groups to develop policies consistent with these Guidelines.

6.3 Trusted Systems Interoperability Group

Presented by Paul Cummings/DEC

A little about the

Trusted
Systems
Interoperability
Group

Presentation to IETF
July 30, 1991
Paul T. Cummings
Digital Equipment Corporation

Digital Equipment Corporation

TSIG ORIGINS

- Vendor based initiative
- B1/CMW Unix-based Systems Developers Recognized
  - Unix-based systems popularity was based on interoperability
  - Independently developed secure implementations would not be interoperable
- !! Something had to be done !!

Digital Equipment Corporation

ATTENDANCE

- Companies that have recently been represented:

DEC	MITRE
IBM	CISCO
AT&T	PYRAMID
SUN	BULL
HP	CRAY
SGI	OSF
SANDIA	SYBASE
UNISYS	CONVEX
ADDAMAX	SECUREWARE
NCSC	SEQUENT
SCO	MIT
LORAL	

Digital Equipment Corporation

TSIG CHARTER

The TSIG is estabished as a forum where interested vendors can convene to:

- Identify and prioritize efforts to ensure multi-vendor interoperability among B1/CMW systems
- Develop implementation agreements
- Implement those agreements
- Demonstrate interoperability and compliance with TSIG implementation agreements

Digital Equipment Corporation

GOALS

- *Ensure near term multi-vendor interoperability for the following*
  - *Xwindow System*
  - *Network File System*
  - *Electronic Mail*
  - *Tar*
  - *r* utilities*
  - *uucp*
  - *FTP/Telnet*
  - *NCS*

Digital Equipment Corporation

WHO SHOULD ATTEND

- *Engineers who are working directly on the implementation of trusted software.*

Digital Equipment Corporation

STRUCTURE

- *Host*
  - *arranges logistics*
- *Work Group Leader*
  - *leads discussion,*
  - *provides minutes,*
  - *serves for preferably 2 or more meetings*
- *Facilities*
  - *mailings lists: maintained by LORAL*

Digital Equipment Corporation

6.4 NREN Architecture and Goals

Presented by Peter Ford/LANL

Outline

Purpose
Quick Overview of NREN
What is the NREN Engineering Group
What NEG is doing
Look at Current Architecture of U.S. Portion of Internet
Why the Architecture Will Change
How Might the Architecture Change
Solicit Input

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

Purpose

Update IETF on what is happening with NREN

Solicit input from IETF community on architectural evolution

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

What is NREN

NREN is a multi-agency program to improve and enhance U.S. portion of Internet

- NSF has role as coordinator for NREN
- Will start Interagency Interim NREN in FY92
- NREN is part of High Performance Computing Initiative (HPCI)
- Congress is still working on the Budget

NREN is not:
the Internet
NSFnet

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

NREN Engineering Group

Hired by NSF to do early engineering for NREN

Reports to Steve Wolff at NSF

NEG's activities coordinated with:
FEPG: Federal Engineering Planning Group
Working Group Chairs of the FNC

Currently Staffed from LANL/LBL:
Bob Fink, Bill Johnston, Stu Loken (LBL)
Peter Ford, John Morrison, Mitch Sukalski,
C. Phil Wood(LANL)

nren@cnt.lanl.gov

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

NEG's current activities

Learn requirements for IINREN

Evaluate Tech. for NREN: transmission, service, protocols, etc.

Where can we get the most leverage?
Govt. funds only a part of what is the NREN
Coordinate with other HPCI programs
Funding needs to fill the gaps.
Guide evolution of NREN, not dictate

Evaluate possible architectural changes
Coordination with FEPG and WGCs

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

Picture of Current U.S. Internet

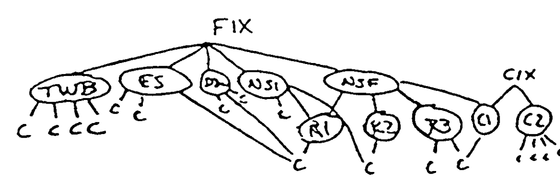

Note:
No international links shown

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

Why will Architecture Change?

Scaling:
- Large Increase in # of nets
- Changes in available bandwidth and services
- Connectivity
  - International
  - greater # of "affinity groups" on Internet
  - Build best possible infrastructure

Commercialization, more transit providers

Availability of greater # of services from/thru network

IETF has a lead role in making this possible!

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

More on current Architecture

Currently FIX centric

Tree structure for routing sanity

1 general transit network: NSFnet run by Merit

- Merit arbitrates routing conflicts
- Merit is the top of the routing default hierarchy
  - Very important since most want to use internal routes + default inside their domain

Routing arbiter role is critical

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

How might architecture evolve in next year or two?

Simply add bandwidth and connections using same technology, architecture remains the same

Change underlying technology (T-1, T-3 --> level 2 services), architecture remains pretty much the same at level 3.

Over time the "1 general transit net" will break down, we are already seeing this with commercial nets.

Need to plan for a robust, coherent, interconnection strategy with a routing arbiter until better tools come along.

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

NEX – NREN eXchange

Need fair access for all providers who want to carry traffic

NEX provides routing arbitration

Currently NEX is contained in NSFnet(arbiter) and the regionals(access). NSF provides access to the federal networks(ESnet, NSN, DDN) at the FIX.

This will change as more transit networks appear.

NEG recommends that NSF formally establish NEX authority for routing arbitration, to be coordinated with FEPG(FIX), CIX members and other transit nets.

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

Features of 2 market model

Familiar topology and routing: NEXes generate a single long haul network

2 markets established: long haul, regional
Both markets are competed

NEX could be run as a franchise

Long haul would be competed, put bandwidth where needed.

Not perfect, will evolve as source specific and policy routing emerges

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

ABSOLUTELY CRITICAL:

- STABILITY OF INTERNET
- HPCC BANDWIDTH
- RELIABLE (END TO END)
  - RAISE OPERATIONAL STD.S
- MORE INFRASTRUCTURE
  - CONNECTIVITY
  - NAME SERVICES (DNS, X.500)
  - SECURITY (AUTHENTICATION, AUTHORIZATION)
  - LOWER BARRIERS TO ACCESS
    - EASE OF USE, SERVICES

31 Jul 1991, Atlanta IETF   Los Alamos National Laboratory

Possible future model (call it 2 market model)?

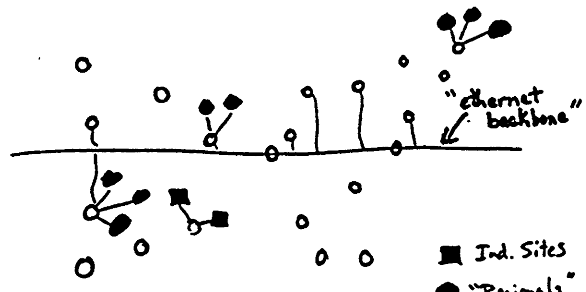

Interconnection of O?
currently level 3 a possibility
possibly level 2 at end of 93

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

NREN Report

Peter S. Ford

NREN Engineering Group

peter@lanl.gov
+1 505 665 0058
Distributed Systems Group
Los Alamos National Laboratory

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

Input?

Send mail to: nren@cnt.lanl.gov

Catch me during the IETF meeting.

BOF?

31 Jul 1991, Atlanta IETF *Los Alamos National Laboratory*

6.5 NREN Legislative Update

Presented by Mike Roberts/EDUCOM

NREN Legislative Update

Mike Roberts
EDUCOM

H.R. 656 - The High Performance Computing Act of 1991
(Passed House of Representatives unanimously July 11)

Section 6(a): Establishes "national multigigabit research and education computer network" (NREN), "which shall link research institutions and educational institutions, govt and industry in every state."

Section 6(a): Agency activities "shall not include purchasing switches, optical fiber, or any other networking hardware for purposes other than research and development."

Section 6(d): "the National Science Foundation shall be responsible for managing the Network according to policies established .....(by Director OSTP acting for the President).

"From sums otherwise authorized to be appropriated", a total of $2.9 billion over FY92-FY96 for all parts of bill.

Gephardt/Sabo amendment added - Buy American
Traficant amendment added - more Buy American
Walker amendment added - "make clear that the agencies participating in the ... program are responsible for implementation, not the Director of OSTP."

S.272 - The High Performance Computing and National Research and Education Network Act of 1991

(Staff work in process to compromise with S.343 and bring to Senate floor asap)

- no net management lead agency designation
- NSF has "primary responsibility" for connecting colleges, universities and libraries to extent they are are not covered by state & local govt, etc.
- references to FCCSET deleted
- DARPA lead agency for R&D changed to simple authorization to develop NREN technology
- Information services coordination responsibility moved from NSF to OSTP
- individual agency designations removed so that Director OSTP may decide who participates in NREN and other HPC program activities

Other NREN Related Developments

"Communications Competitiveness and Infrastructure Modernization Act of 1991" - H.R. 2546/S.1200

- develop and deploy a "nationwide, advanced, inter-active, interoperable, broadband telecommunications infrastructure on or before 2015"
- Amends Comm Act of 1934 - "each local exchange carrier shall prepare and carry out a broadband communications system implementation plan..."
- FCC to convene Joint Board with State PUC's to develop recommended implementation plan

"GPO Wide Information Network forData Online Act of 1991" - H.R.2772 (aka "GPO WINDO")

- Govt Printing Off to become single point of access to wide range of federal electronic databases
- access fees to approximate incremental cost of dissem-ination
- immediate GPO connection to Internet planned

6.6 BellSouth Telecommunication

Presented by Caroline Cranfill/BellSouth

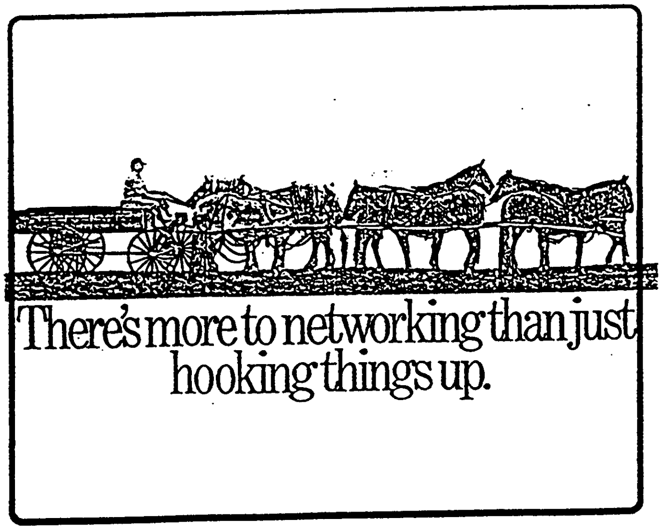
There's more to networking than just hooking things up.

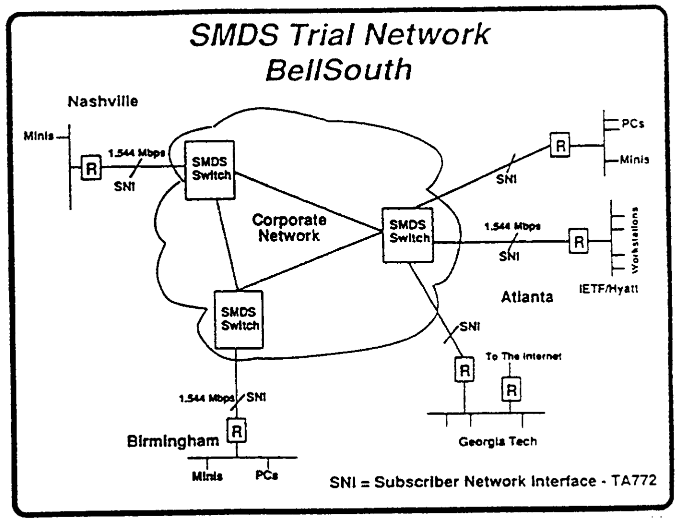
SMDS Trial Network
BellSouth
Nashville
Minis
1.544 Mbps
SNI
R
SMDS Switch
Corporate Network
SMDS Switch
PCs
Minis
1.544 Mbps
SNI
Workstations
IETF/Hyatt
Atlanta
SMDS Switch
SNI
To The Internet
1.544 Mbps
SNI
Birmingham
Minis
PCs
Georgia Tech
SNI = Subscriber Network Interface - TA772

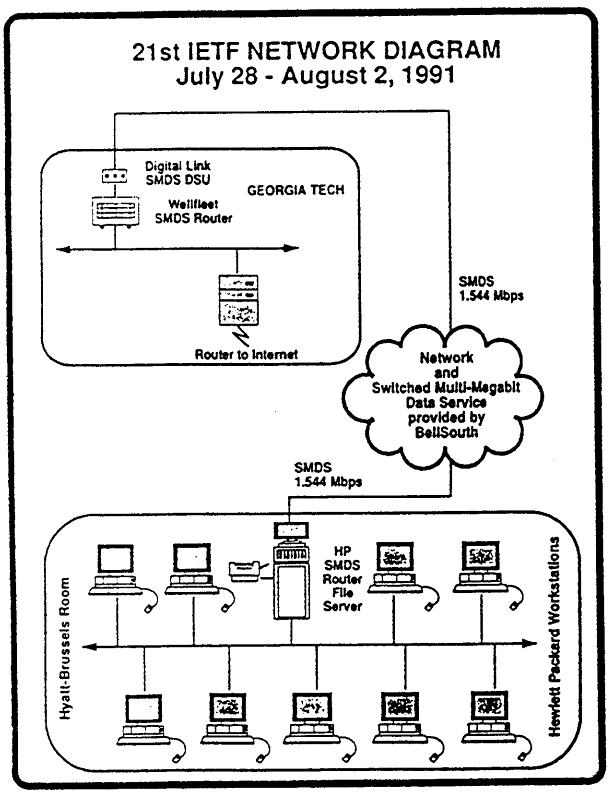
21st IETF NETWORK DIAGRAM
July 28 - August 2, 1991
Digital Link SMDS DSU
Wellfleet SMDS Router
GEORGIA TECH
Router to Internet
SMDS
1.544 Mbps
Network and Switched Multi-Megabit Data Service provided by BellSouth
SMDS
1.544 Mbps
HP SMDS Router File Server
Hyatt-Brussels Room
Hewlett Packard Workstations

6.7 Introduction to the Internet Society

Presented by Vinton Cerf/CNRI

The Internet Society will function as a professional society to facilitate, support and promote the evolution and growth of the Internet as a global research communications infrastructure. The suggestions and recommendations of all parties interested in the Internet are solicited to assist in making the Internet Society robust, productive and structured to meet the needs of its members.

The Internet Society

The Internet, is a collection of cooperating, interconnected, multiprotocol networks which supports international collaboration among thousands of organizations. Because of its current scope and rapid rate of growth, the Internet will benefit from a more organized framework to support its objectives. To this end, an Internet Society is being formed to foster the voluntary interconnection of computer networks into a global research and development communications and information infrastructure. The Internet Society will not operate the Internet. Internet operation will continue to be a collaborative activity which the Society will seek to facilitate. The Society will provide assistance and support to groups and organizations involved in the use, operation and evolution of the Internet. It will provide support for forums in which technical and operational questions can be discussed and provide mechanisms through which interested parties can be informed and educated about the Internet, its function, use, operation and the interests of its constituents.

Membership

The Internet Society will be a membership organization with voting individual members and non-voting institutional members. There will be several classes of institutional members. The society will produce a newsletter on a regular basis and hold an annual meeting to which all members and other interested parties will be invited. The topics of the annual meeting will vary, but are expected to focus on current research in networking, Internet functionality and growth, and other interests of the Society constituency. All members will receive the newsletter and an invitation to attend the annual meeting of the Internet Society.

Membership dues will vary according to class of membership. The amounts of these dues and the basis on which they are set will be determined by the Board of Trustees of the Society and may be revised from time to time as set forth in the By-Laws.

Charter

The Society will be a non-profit organization and will be operated for academic, educational, charitable and scientific purposes among which are:

- To facilitate and support the technical evolution of the Internet as a research and education infrastructure and to stimulate involvement of the academic, scientific and engineering communities, among others in the evolution of the Internet.

- To educate the academic and scientific communities and the public concerning the technology, use and application of the Internet.
- To promote scientific and educational applications of Internet technology for the benefit of educational institutions at all grade levels, industry and the public at large.
- To provide a forum for exploration of new Internet applications and to foster collaboration among organizations in their operation and use of the Internet.

Activities of the Society

- Support for Internet Technical Evolution

  The Internet Activities Board (IAB) has been concerned with the development and evolution of architectures supporting the use of multiple protocols in a networked environment. The Internet Society will incorporate the IAB and its functions into the operation of the Internet Society. The Internet Society will work with other interested organizations to support and assist efforts to evolve the multiprotocol Internet. The Internet Society will use the Internet Engineering and Research Task Forces to stimulate networking research and facilitate the evolution of the TCP/IP protocol suite and the integration of new protocol suites (e.g. OSI) into the Internet architecture. The Internet Society will work actively with parties and organizations interested in fostering improvement in the utility of the Internet for its constituent users.

- Meetings and Conferences

  Internet Society will convene an annual meeting and will organize and facilitate workshops and symposia, jointly with other organizations where appropriate, on specific topics of interest to the Society membership. The annual meeting will address issues of global and regional importance to the evolution and growth of the Internet. In particular, future INET conferences will be incorporated into the Society's annual meetings.

- Information and Infrastructure Services

  The Internet Society will publish an Internet Newsletter providing members with information about the international activities of Internet constituents. In addition, the Society will also provide assistance to and support for organizations responsible for maintaining the databases crucial to Internet function (e.g. the Domain Name System, X.500 Directory Services, etc.) and organizations concerned with the security of the Internet (e.g., the Software Engineering Institute Computer Emergency

Response Team (CERT) and its CERT-System). The Society will assist in the development of educational, advisory and informative materials of use to Society members. Where appropriate, the Society will organize or support activities which aid in the coordination among the organizations operating components of the Internet.

The Society will refer members to appropriate parties involved in operating the various parts of the Internet where they may be helpful with specific questions. Where possible, the Society would seek to provide access to its information on-line, but would also offer hard copy and, perhaps eventually, CD-ROM-based information resources.

Plans

The initial organizers of the Internet Society include the Corporation for National Research Initiatives (CNRI), EDUCOM and the Internet Activities Board. During the six month period from June - December 1991, the organizers will work with interested parties to prepare for beginning operation of the Society by the end of 1991. Computer networking has become a critical infrastructure for the research and development community and has the potential to become the basis for world-wide collaboration and cooperation in every field of human endeavor. The Internet Society will seek to solidify, enhance and encourage further international collaborative networking. Individuals joining the Society during its formation will receive special recognition as Society pioneers and will have the opportunity to shape the early agenda of Society activities. Opportunities for organizational and institutional participation are also available.

It is time. The technology is available. A global renaissance of scientific and technical cooperation is at hand. You are cordially invited to take part in an enterprise without precedent and an adventure without boundary. The Internet Society sets sail in January of 1992 on a voyage of internetwork discovery. Will you be aboard?

Internet Society

INTERNET SOCIETY

* Organizers: CNRI, EDUCOM, IAB
* Endorsements: RARE, CNI, ARL,
  CCIRN co-chairs, FNC members
* Seeking Other Interested Parties
* General Operations: end of 1991
* Pioneer Individual and...

  Founding Organizational Members are invited to join before 1992

INTERNET SOCIETY

GOALS

* Multiprotocol Evolution
  TCP/IP, OSI, ...
* Encourage Internet Growth
* Educate the Public
* Stimulate Provision of Service
* Recognize Individual Contrib.
* Promote and Explore Scientific, Educational, Business Use
* Facilitate Collaboration

INTERNET SOCIETY STRUCTURE

* International Professional Society
* Individual and Organization Members
* Elected Board of Trustees
  who appoint officers
* Staggered 3-year terms
  max: 2 contiguous terms
* Incorporates IAB/IETF/IRTF
  details: TBD

INTERNET SOCIETY

ACTIVITIES

* Evolution of Internet Technology
* Incorporate IAB, IETF, IRTF
* Newsletter and Prof. Journal
* Annual Conference: INET

  INET 92: June 15-19, 1992

  Kobe, Japan
* Possible Infrastructure Assistance
  - CERT-System
  - Crypto-Certificates
  - Internet Registry
  - Referral Services

INTERNET SOCIETY SUMMARY

* Nourish the Community now growing up around networks.
* Our packets cross borders freely and our sense of community should be equally open
* Continue the Grand Collaboration now linking 5,000+ nets and 3,000,000+ people

For More Information...

ISOC@NRI.RESTON.VA.US

FAX: +1 703 620 0913

TEL: +1 703 620 8990

6.8 Toward a New Routing Architecture

Presented by Noel Chiappa/Consultant

A New IP Routing And Addressing Architecture

"One ring to rule them all, one ring to find them, one ring to bring them all, and in the darkness bind them."
J.R.R.Tolkien

"Perfection has been attained, not when there is nothing left to add, but when there is nothing left to take away."
A. de St. Exupery

METHOD

- *Act as if a blank slate exists - plan for the best design & then worry about how/if it can be deployed.*
- *This is a general plan of attack - a broad brush architecture, based on a study of the fundamental problems of routing and addressing in a large network.*
- *This is not an engineering design - detailed mechanisms and optimizations have not been thought through.*
- *This is not complete - details are still being tweaked at this time.*

GOALS

- *Incrementally deployable.*
- *Large size - basically indefinite - solve all existing problems.*
  - *Exhaustion of IP network numbers.*
  - *Exhaustion of IP address space.*
  - *Routing meltdown.*
- *Policy controls.*
  - *Access control*
  - *Trust model*
  - *Information hiding*
- *Firewalls/Robustness.*
- *Minimize configuration.*

Routing Fundamentals

- *DV Algorithms*
  - *Distributed computation.*
  - *Intermediate results passed around.*
- *LS Algorithms*
  - *Distribute map.*
  - *Computation local, may be delayed.*
- *LS is best*
  - \+ *LS makes policy routing feasible.*
  - *Size is a problem.*
- *Terms*
  - *Compression*
  - *Thinning*

  Abstraction

o Compression is insufficient.

o Thinning must be used.

o Data being discarded causes non-optimal routes.

There are two costs to balance:

1) Cost of distributing routing information.
2) Cost of using non-optimal routes.

o Balancing these two is a trade-off we must leave to the future, since it is a cost/benefit tradeoff policy issue.

Addressing Fundamentals

o Object spaces

- Nodes
- Network attachment points (NAP)
- Networks

o Names

- One or more possible names or types of names for each object.
- Different names may or may not have structure, e.g.,

  - Ethernet 48 bit hardware address
  - IP address

o Problems appear when object classes are confused or mixed, or when the name or form of a name is not considered separately from the object it refers to.

o Structure is generally to help something else do its job.

o Name for nodes
- Short, flat, fixed length.
- Useful to efficiently identify sources/destinations in packets.

o Names for NAP's - "Addresses"
- Topologically related NAP's get related addresses.
- Allows the number of "destinations" tracked by the routing to be minimized.

o Topologically related addresses may also
- Allow quick location of NAP's on map.
- Provide representation for topology distribution.
- Provide framework for abstraction process.
- Be used in "no brainer" routes.

o Note that these are all logically separate tasks

o Handling the size issue is the most difficult of the goals.

o Addresses should take <u>whatever</u> form makes the routing (especially the size issue) easiest.

Architecture Outline

- o Routing scheme is multi-level (i.e., hierarchical) LS.
  - Multi-level for size.
  - LS for policy.
  - Discards IGP/EGP split.
- o Forwarding path is source-specified, not hop-by-hop.
  - Needed to allow complex source policies.
  - Needed to allow incremental deployment of new link attributes.
  - Allows "local abstraction control". (LAC)
- o Routing algorithm is not part of the specification.
  - Sample algorithm will be available as appendix to the eventual specification.
  - Allows experimentation and incremental deployment of new algorithms.

- o Abstraction of topology is not part of the specification.
  - Allows new algorithms.
  - Allows trade-offs of costs of varying abstractions.
  - Sample algorithm will be available as appendix.
- o Clients not mandated to accept an abstraction - LAC.
  - Allows users to tune cost/ benefit tradeoffs.
  - Allows non-hierarchical routing.
- o A new kind of address, for use in routing and forwarding.
  - Variable length, variable number of levels.
- o Old addresses become node ID.
  - Allows interoperation with existing code and incremental deployment.
  - Can be extended via incremental mechanism.

Reflections on the Architecture

- o Note what is <u>not</u> part of the architecture:
  - Routing algorithm
  - Abstraction algorithm
  - LAC algorithm
- o Key Advantages:
  - Limits amount of work to be done.
  - allows future improvement and easy deployment of new algorithms.
  - reduces scope for errors and bad design by minimizing global mechansims.
- o What <u>is</u> part of the architecture is:
  - Method of representing topology.
  - Method of distributing topology information.
  - Method of setting up flows.

6.9 NSFNET T3 Deployment

Presented by Elise Gerich/Merit and Jordan Becker/ANS

T3 Network Status Report - IETF (Atlanta)
August 1, 1991

Network Architecture and Evolution

- 7/88 T1 NSFNET Backbone
- 12/90 Phase-I T3 Network
- 10/91 Phase-II T3 Network
- 1Q92 Phase-III T3 Network

Phase-I T3 Network Architecture and Experiences

Phase-II T3 Network Architecture & Schedule

- Phase-I -> Phase-II Deployment
- Latest Performance & Usage Measures

T1 -> T3 Transition Experiences Summary

Planned Phase-III T3 Network Architecture

- New DS3/FDDI Adapter Technology
- New Switch Architecture

NSFNET T1 Network Monthly Packet Traffic

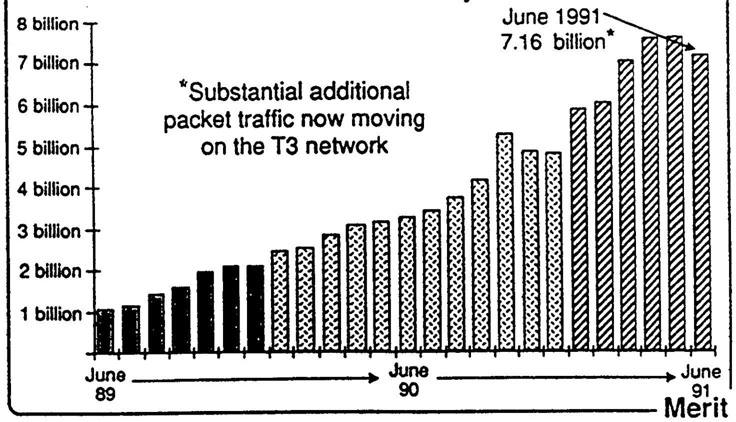

Phase-I T3 CNSS

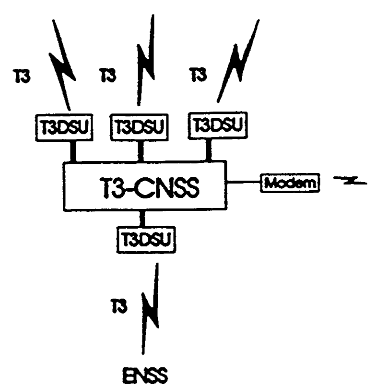

RS/6000 CPU Provides Packet Forwarding, Routing

Four DS3 Interface Per CNSS
- Microchannel Bus Master Adapter with I386
- 0.5MB Memory
- T3 Technologies DSU Interface

NSFNET T3 Network Monthly Packet Traffic

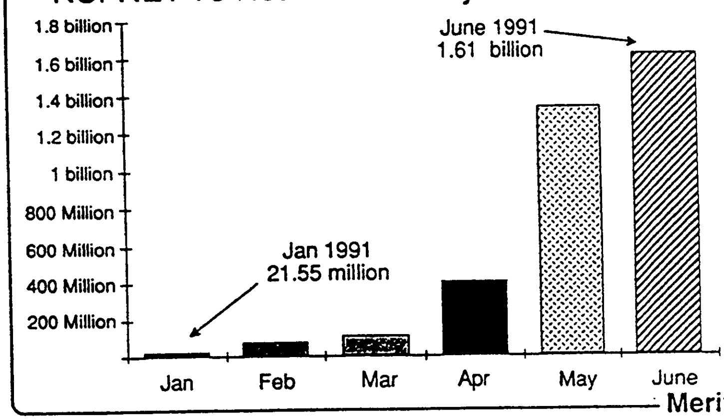

Phase I and II T3 ENSS

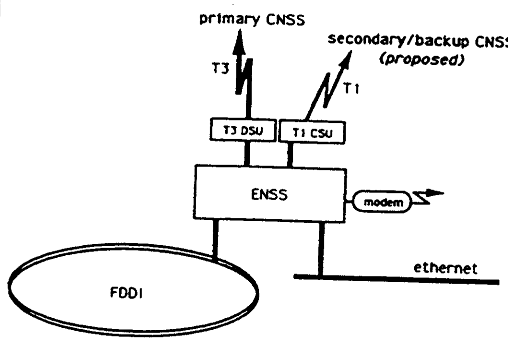

NSFNET T1/T3 Networks 6/91

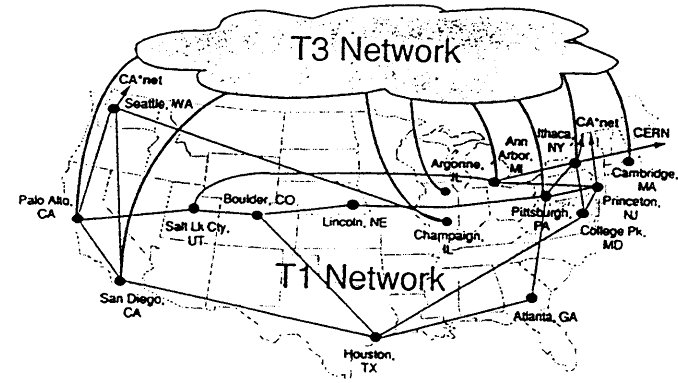

Merit

Phase I T3 Network - 5/91

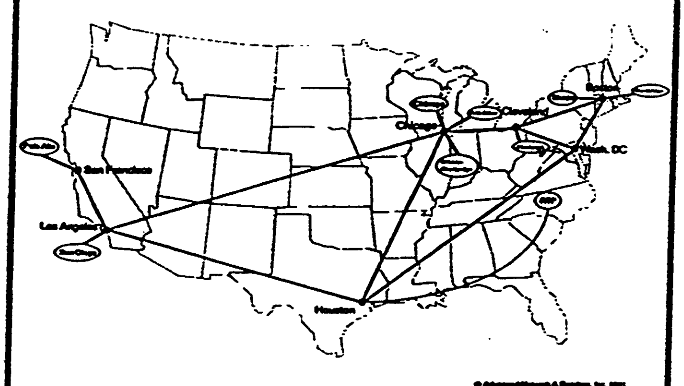

T3 Network Delay Matrix Report for 7/24/91

1st Quartile Sample

DelayTo		128	129	130	131	132	133	134	135	150
1q	128	--	27	27	30	32	43	41	8	43
1q	129	27	--	2	6	7	18	16	23	23
1q	130	27	2	--	5	7	17	15	22	23
1q	131	30	6	5	--	9	20	18	26	26
1q	132	32	7	7	9	--	14	13	28	19
1q	133	43	18	17	20	14	--	5	38	15
1q	134	41	16	16	18	13	5	--	37	13
1q	135	8	23	22	26	28	38	37	--	39
1q	150	51	26	25	24	26	29	27	46	--

3rd Quartile Sample

DelayTo		128	129	130	131	132	133	134	135	150
3q	128	--	28	27	31	32	51	50	8	47
3q	129	28	--	2	6	8	31	31	23	55
3q	130	27	2	--	5	7	30	30	22	55
3q	131	31	6	5	--	10	37	35	26	34
3q	132	32	8	7	10	--	26	25	28	20
3q	133	51	31	30	37	26	--	5	47	15
3q	134	50	30	29	35	25	6	--	46	14
3q	135	8	23	22	26	28	47	45	--	44
3q	150	53	26	26	25	27	29	29	48	--

Notes:

- Results are one way delays in milliseconds (1/1000 sec)
- Delays are measured using external addresses

Phase-I Experiences Summary

- Gained Experience with DS3 Non-Channelized Framed Format
- T3 Router/DSU Problems
  - Host Reboot, Packet Loss, Grey Link Problems
- T1 & T3 Routing Configurations
- Expect T3 Network to Stabilize with Field Upgrades & Burn-In
  - Phase-II Network Will Support All Field Upgrades

Phase-II T3 Network Design Requirements

- Switch Architecture Requirements for Phase-II Expansion
  - Require 6-8 DS3 Interfaces Per CNSS
    - Backbone (CNSS) Trunks & Customer Site (ENSS) Tails
    - Single Router Supports Up to 4 DS3 Interfaces
  - Minimize New Technology Beyond Phase-I
  - Multiple DS1 / 56Kbps ENSS Interface Concentrators
    - Support Switched T1 / 56Kbps Backup Services
    - Enable the Phase-Out of T1 Backbone
  - Use Commercially Available Technology Where Possible
- No Single Router Product Satifies All Requirements Today
- Phase-II CNSS Switch Design
  - Multi-Router Hierarchical Switch
    - RS/6000's & Cisco AGS+ Routers in All POPs
    - T3 Trunk Router, T3/T1/56 ENSS Concentrators

Phase II CNSS

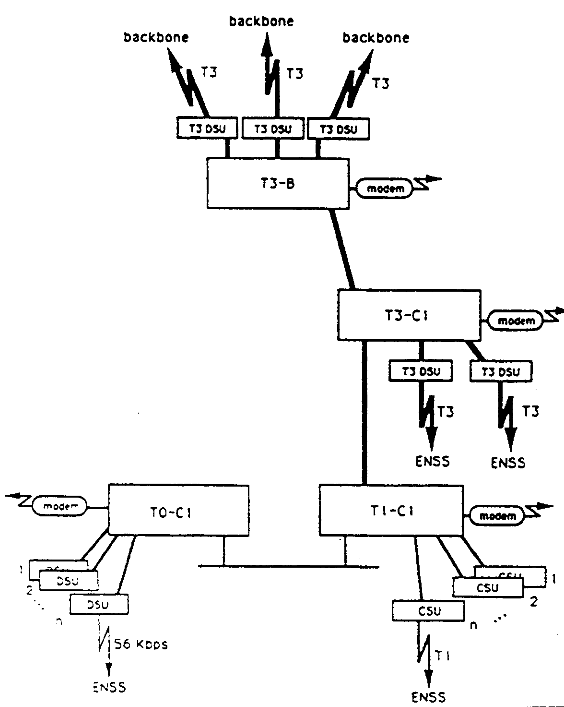

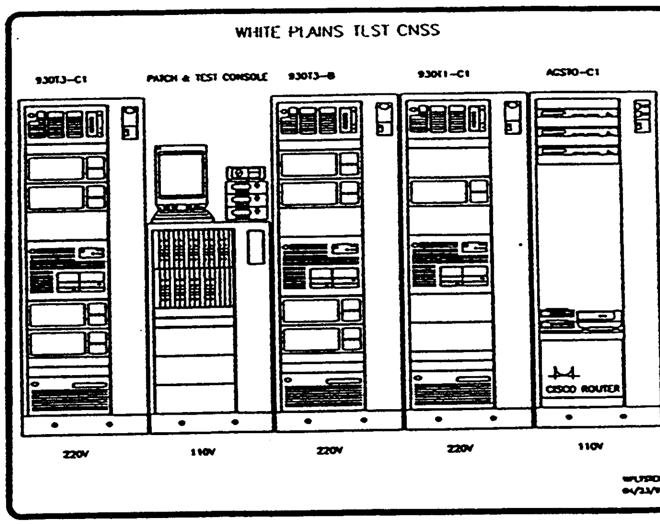

Phase-II T3 Topology Design Objectives

- o **Expand T3 Topology To Support 12 CNSS, 16 ENSS Sites**
  - Average CNSS Transit Hop Count < 3 (2.6 Actual)
  - Maximum CNSS Transit Network Diameter of 5 Hops
  - Minimize CNSS Transit Traffic
    - Average = 65%, Min = 25%, Max = 85%
  - Full CNSS Link Redundancy
  - Build Parallel East-West & Northeast-Southeast Routes

Phase II T3 Network - 10/91

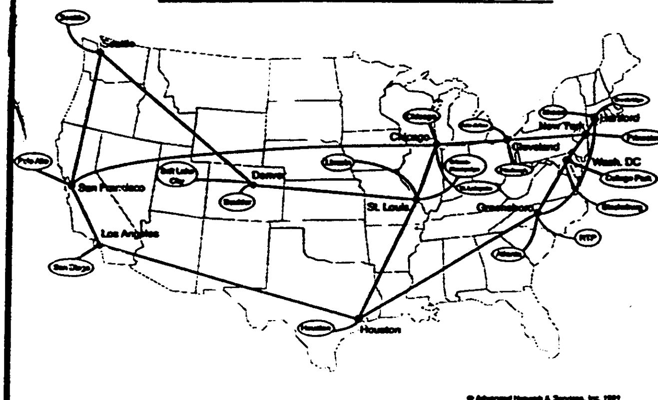

Phase-II T3 Network Deployment Plan

- o Stable Transition of Live Phase-I into Live Phase-II
  - Build New CNSS Facilities Before ENSS Cutovers
  - Field Upgrades for All Phase-I Problems
  - Avoid Flash CNSS Cut-Overs Where Possible
  - Reuse of Existing Transmission Facilities
  - 3 Parallel Installation Teams
    - 2 Day Installation (POP or Customer Site)
    - 2 Day Cut-in and Test
- o Staged Deployment
  - 5 Stages of Phase-II Deployment
  - East Coast & Southwest Stages Complete
  - Midwest & West & Northwest Stages Pending
- o Schedule for Remaining Stage Completion
  - Midwest (Cleveland & Chicago) 8/9
  - West (L.A., Hayward) 8/30
  - Northwest (Seattle, Denver) 9/15

Phase II T3 Network Deployment

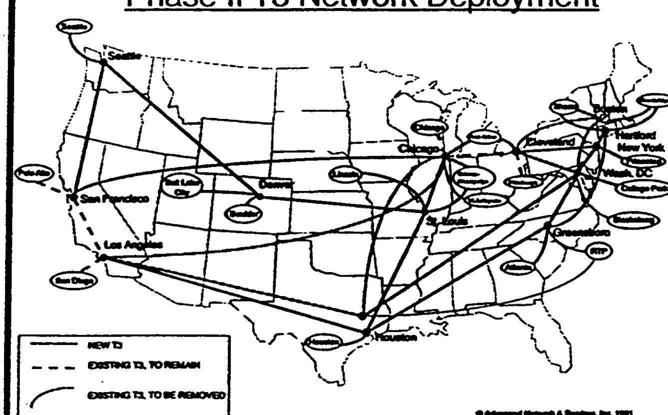

Switched Services on the Phase II T3 Test Network

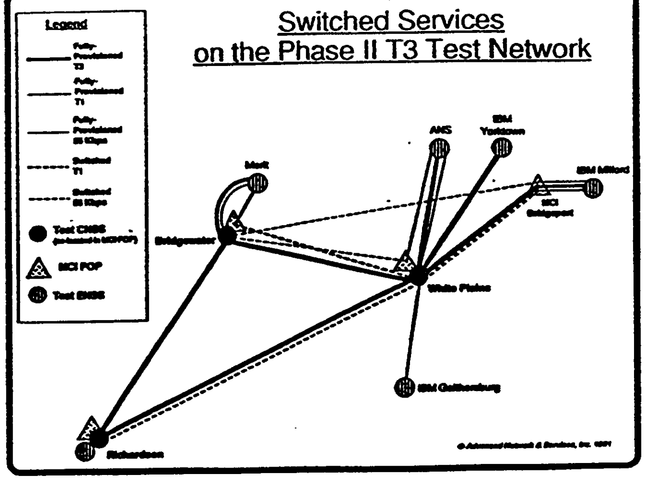

Phase-III Network Plan

- o On-Going Development of Performance Enhancements
- o Upgrade RS/6000 T3 Packet Switching Technology
  - New High Performance Router Interface Adapters
    - i960CA On-Board IP Packet Forwarding
    - DS3, FDDI, Ethernet Interface Support
    - Improved Performance (20KPPS Per Card)
    - Direct Card-To-Card Data Transfers
- o New CNSS Switch Architecture To Satisfy Requirements
  - Redundant Multi-Router Switch (NSS Style)
    - FDDI Internal Backbone
  - Five i960 Adapters Per RS/6000 T3 Router
    - 4 DS3, 1 FDDI Configuration Possible
  - Cisco AGS+ Participation Via FDDI Backbone
  - Proposed Software Enhancements
    - Migrate to Common IGP Within Phase-III Backbone
    - Support for Connectionless OSI Network Protocols (CLNP)
      - Intra-Domain IS 10589 Routing (IS-IS)
      - Inter-Domain Routing (IDRP)
    - Being Implemented in GATED
- o Implement Switched T1 & 56 Backup Capability

Summary

Rapid Evolution of Technology

Multi-RT NSS T1 Network

Phase-I Eight Node T3 Network 12/90

Phase-II Sixteen Node T3/T1/56K Network Target Completion 9/91

Phase-III Performance Enhancements To Be Scheduled

Several New Problems Addressed

DS3 Transmission Facilities for Un-Channelized Data Application

Router & DSU Technology Stabilized

Increased Routing Complexity

New Services and Technology To Come

- Enhanced T3/FDDI Performance
- Switched Backup Facilities

6.10 Statistics from the DDN NIC

Presented by April Marine/SRI

A quick picture providing an idea of the growth of the Internet is presented via the slides from SRI International.

The first slide shows IP Address Assignments, broken down by Class of network, and reflects the large number of TCP/IP networks being established. Not every number assigned reflects a network with access to the Internet.

The second slide shows a minimum number of hosts and domains known to the Domain Name System (DNS). This data is collected by the ZONE program. ZONE runs on a DEC2065 mainframe and starts with a list of top-level domains and their servers. For each domain, the program attempts to make a TCP connection to one of the servers. Once connected, it requests a zone transfer for that domain. If a name server record refers to a subdomain, the ZONE adds that subdomain to the list of domains it searches. The program thus descends through the entire domain tree trying to find all existing domains. ZONE cycles through its list of domains left to search until it has gone through the entire list without receiving any new information.

The third slide presents both linear and logarithmic graphs of the data showing the number of hosts and domains collected by the ZONE program from several times it was run over the past three years.

IP ADDRESS ASSIGNMENTS

Date	Class A	Class B	Class C	Total
6/91	43	5,654	24,449	30,146
1/91	37	4,354	21,872	26,153
7/90	34	2,533	16,214	18,781

DOMAIN SURVEY RESULTS

Date	Hosts	Domains
7/91	535,000	16,000
10/90	313,000	9,300
10/89	160,000	4,800

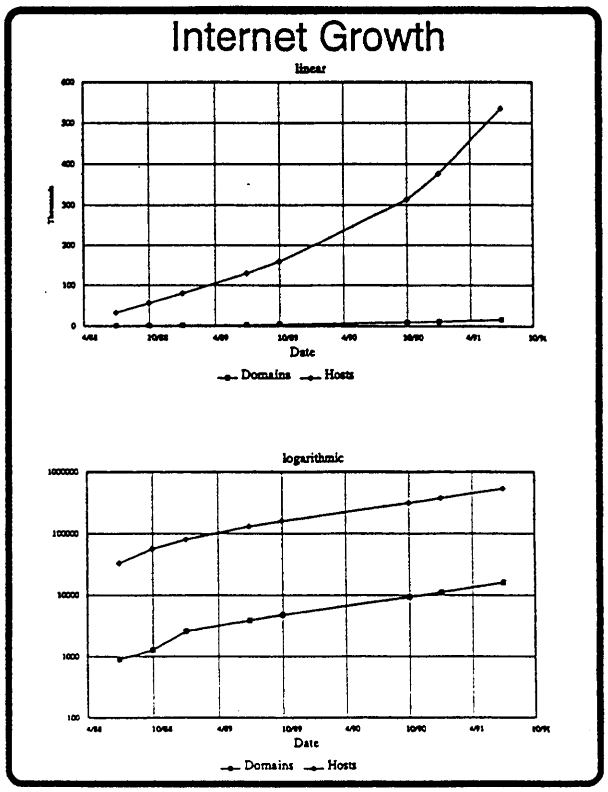

6.11 Introduction to Archie

Presented by Peter Deutsch/McGill University

Lessons Learned So Far:

- *The Internet needs infodiscovery tools*
- *You don't have to break the bank to get started.*
- *Solutions can (and should) be built incrementally, with as much feedback from users as possible.*
- *Archie is a great "loss lender" to demonstrate network potential. (Users Love It!)*

(Philosophical Rant)

- *We need "professionally run" Internet services.*
- *Who pays for running these is still an open question.*

Cooperative Efforts

- *Prospero*
- *WAIS (Z39.50)*
- *Perhaps WWW, X.500*
- *Others? (contact us!)*

o Volunteers have started providing:

- *Prospero clients*
- *WAIS server*
- *Documentation (e.g., man page)*

Current Availability:

- *McGill University (Montreal,Canada)*
- *AARnet (Australia)*
- *funet (Finland)*

Others With Source:

- *CERFnet*
- *JVNCnet*
- *MIDnet*
- *SURAnet*
- *EFF*

** Plus interest from others*

Our Policy on Source Code

- *Archie is NOT public domain.*
- *Archie is currently available at no charge to regionals and other institutional service providers.*
- *We are seeking long-term support to build a reliable service.*
- *Depending upon funding, we MAY charge for future maintenance and updates. (TBD)*

What Is Archie?

- Archie is a "card catalogue" for anonymous ftp archive sites.
- Archie can be seen as a "helpdesk" for the Internet.

Architecture

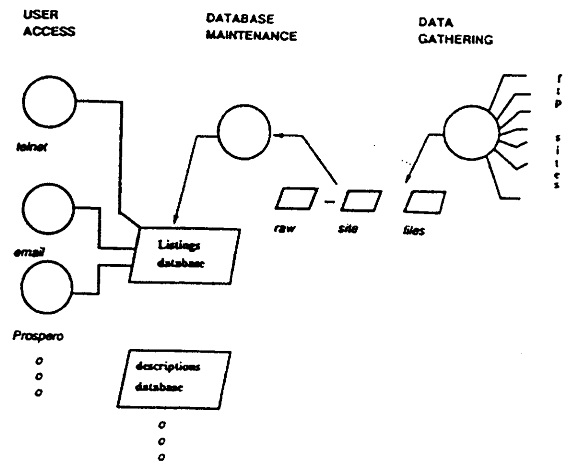

What Can You Do?

- Search for file names.
- Search secondary database for package descriptions.
- List contents & information on individual sites.

What's Missing?

o Archie is a "proof of concept"

o Archie is still under development

To come (soon):

- More work on client-server model (WAIS, better front-ends)
- A method to coordinate updates among multiple Archies
- Parsers for VMS, TOPS.20, VM
- Speed improvements
- Better documentation
- "Finish" email interface
- Better instrumentation
- A home <:sniff:>

Follow-on work ("Real Soon Now" TM)

- o Additional databases for:
  - Mailing lists & their descriptions.
  - Mailing list archives.
  - On-line library catalogues.
  - Usenet newsgroup archives.
  - ???

  (Basically, any list of information.)
- o An integrated archie-ftp tool
- o GUI front-ends
- o ???

"Son of Archie"

% setenv HANDWAVING

Architecture

- Multiple information servers for different types of information. (Let users choose the partitioning.)

--> Need for a "top level" infoserver for locating groups of servers.

- Servers for dispensing as well as locating information.
- A registry mechanism for service providers to enter information into the servers.

Basic Idea

The idea is to partition the search space and lower "buy-in" cost.

We want to avoid, not solve, the resource discovery problem (for now).

%unsetenv HANDWAVING

Operations Wishlist

- o A standardized format for storing information on anonymous ftp sites.
- o A "Recommended Operating Guide" for anonymous ftp sites.
- o Promote use of registries. (e.g., "whatis" database).
- o We'd like to see *lots* of archie-like services, run by lots of different people, and an architecture tying them all together.

Design Criteria

1) We wanted universal access.

2) We wanted a simple, easy to explain tool.

3) We wanted low "entry cost". (=>telnet, email)

4) We *NEEDED*

- Short development time.
- Low development cost.

5) Don't be afraid to spend cycles.

6) *MINIMIZE DEPENDENCE UPON OTHERS!* (Don't rely upon 700 volunteers in 24 time zones to make it work :-)

Some Numbers

Size

- *We currently track~one million files (totally~50 Gigabytes). on~700 sites.*

- *Index is ~70 Megabytes.*

Logins

- *~200,000 logins since December 1990.*

- *We now receive ~1,600 logins a day. (We're saturated!)*

From

- *At least 33 countries on at least 6 continents. (So, is the Antarctic on the Internet?)*

Request for Proposal

o *Archie is NOT the comic character!*

o *Archie is NOT from "All in the Family".*

We Need A New Name!

o *Archie is still growing and will add functionality in the coming months. This is probably the last chance we'll have to change the name...*

- *Send suggestions to "archie-l@cs.mcgill.ca"*

"Credit Where Credit Is Due" Dept.

Implementors

- *Alan Emtage*
- *Peter Deutsch*
- *Bill Heelan*
- *Mike Parker*

Collaborators

- *Cliff Neuman (Prospero)*
- *Brewster Kahle (WAIS)*
- *John Granrose (ftp site list)*
- *Ed Vielemetti (comp. archives)*
- *Khun Yee Fung (perl-based Prospero client)*
- *Jerry Peek (discussions)*

Getting There

telnet quiche.cs.mcgill.ca

Login as "archie" (no password)

email archie@quiche.cs.mcgill.ca"

try "help" in Subject or Body.

Prospero anon ftp: June.cs.washington.edu

email to: Cliff Neuman "bcn@isi.edu"

Documentation on Archie

anon ftp to quiche.cs.mcgill.ca

subdirectory "archie/doc"

- Man page
- Archie blurb
- Etc...

Chapter 7

IAB Workshop Report

Presented by the Members of the IAB

IAB Architecture Retreat

7.1 Introduction - Bob Braden/ISI

The Internet architecture, the grand plan behind the TCP/IP protocol suite, was developed and tested in the late 1970s by a small group of network researchers. Several important features were added to the architecture during the early 1980's – subnetting, autonomous systems, and the domain name system; more recently, IP multicasting has been added. Within this architectural framework, the Internet Engineering Task Force has been beavering away with great energy and effectiveness, engineering, defining, extending, testing, and standardizing protocols for the Internet. Three areas of particular importance have been routing protocols, TCP performance, and network management. Meanwhile, the Internet infrastructure has continued to grow at an astonishing rate. Since January 1983 when the ARPANET first switched from NCP to TCP/IP, we have all been laboring mightily to survive our success!

The Internet Activities Board (IAB) has evolved from a technical advisory group set up in 1981 by DARPA, with its membership drawn from the ranks of the network researchers who developed the original Internet architecture and protocols. IAB membership has changed somewhat over the years, to better represent the changing needs and issues in the Internet community, and more recently, to reflect the internationalization of the Internet.

The IAB created the Engineering Task Force to carry out protocol development and engineering for the Internet. To manage the burgeoning IETF activities, the IETF Chair set up a steering group, the IESG. The IAB and IESG work closely together in ratifying protocol standards developed within the IETF. The IAB itself retains an institutional concern for the protocol architecture. For a number of years, Dave Clark served as IAB Chair with the informal title of "Internet Architect".

Over the past few years, there have been increasing signs of strains on the fundamental architecture, mostly stemming from continued Internet growth. Discussions of these problems reverberate every day on many of the mailing lists we all love. Recognizing the growing cracks in the foundations, the IAB and the IESG scheduled a joint meeting for January 1991, to include a full day's discussion of Internet architectural issues. The framework of the January architecture meeting was set by Dave Clark. [His slides are included in the summary that is available for anonymous FTP from venera.isi.edu: pub/IABmins.jan91Arch.txt.] The discussion was spirited, provocative, and at times controversial, but little was agreed upon.

The group therefore decided to meet in June 1991 at SDSC to devote three full days to Internet architecture issues. This meeting, which was called somewhat perversely the "Architecture Retreat", was convened with a strong resolve to advance the architecture. Besides the IAB and IESG, the group of 32 people included the members of the Research Steering Group (IRSG), plus a few special guests.

In January, there had been a lot of soul-searching and blood-letting over questions of relevance and future direction. One view has been that we should just let the TCP/IP suite strangle in its success, and switch to ISO protocols. However, others who have worked hard and successfully on Internet protocols, products, and service were anxious to at least try to solve the new problems in the existing framework. In the long run, ISO is likely to suffer from many of the same problems. The January meeting reached a fairly solid consensus, which was led to the basic assumptions underlying the Architecture Retreat:

1. TCP/IP and OSI will coexist for a long time;

2. The Internet will not become homogeneous, but will continue to include diverse networks and services;

3. Commercial and private networks will be incorporated, but we cannot expect the common carriers to provide the entire service;

4. We need to be able to scale to 10**9 networks.

Guided by the results of the January meeting, the June Retreat was organized into 5 separate discussion areas:

1. Routing and Addressing: "The Mother of all IP architecture issues."

2. Multi-Protocol Architecture: "Making the problem harder for the good of mankind."

3. Security

4. Traffic Control and State

5. Advanced Applications.

On the second day, the Retreat broke into groups, one for each topic; groups reported their conclusions to the plenary on the third day. These are complex and difficult issues. We feel that considerable progress was made, although of course much remains to be done.

At the Atlanta IETF meeting, each of the group Chairs presented a brief summary of their conclusions (in reverse order). Summaries of these reports follow.

IAB RETREAT ON THE

FUTURE OF THE INTERNET ARCHITECTURE

SDSC June 11-13, 1991

Report to the Internet Engineering Task Force

Atlanta, Georgia August 1, 1991

AGENDA

- The Internet: We labor mightily to survive our success!
- Phenomenal growth ==>

  strains in the underlying protocol architecture.

  "Cracks in the foundations..."

  "Driving full steam into a swamp..."

- **WHO is going to get us out of this mess?**

  ==> Y * O * U <== **the energetic members of the IETF**

  with a little help from the IRTF... !

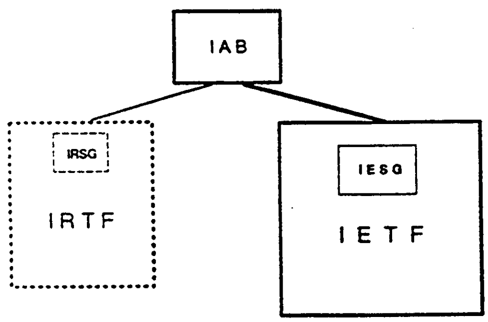

THE IAB . . .

They got us into this mess, they ought to try to help get us out of it...!

IAB-SPONSORED MEETINGS ON INTERNET ARCHITECTURE

1) IAB & IESG met for one day: January 1991.

   Lots of heat, limited light...

   [See: pub/IABmins.jan91Arch.txt on venera.isi.edu]

2) IAB RETREAT ON THE FUTURE OF THE INTERNET ARCHITECTURE.

   3 days in June 1991; IAB, IESG, IRSG...

   Results reported in this session.

IAB ARCHITECTURE RETREAT

- IAB:

Bob Braden	ISI	Hans-Werner Braun	SDSC
Vint Cerf	CNRI	Lyman Chapin	BBN
Dave Clark	MIT	Phill Gross	ANS
Christian Huitema	INRIA (France)	Steve Kent	BBN
Tony Lauck	Digital	Barry Leiner	ADS
Dan Lynch	Interop	Jon Postel	ISI

- IESG:

David Borman	Cray	Ross Callon	Digital
*Vint Cerf	CNRI	Noel Chiappa	MIT
Dave Crocker	Digital	Steve Crocker	TIS
Chuck Davin	MIT	*Phill Gross	ANS
[Rob Hagens	UWisconsin]	Bob Hinden	BBN
Russ Hobby	UCDavis	Joyce Reynolds	ISI
Greg Vaudreuil	CNRI		

- IRSG:

*Bob Braden	ISI	*Dave Clark	MIT
*Steve Kent	BBN	Dave Mills	UDelaware
Deborah Estrin	USC	Cliff Lynch	UC Pres. Office
Mike Schwartz	UColorado		

- Guests:

Van Jacobson	LBL	Paul Mockapetris	DARPA
Claudio Topolcic	CNRI	Peter Ford	LANL
[Craig Partridge	BBN/SICS]		

- Scribe:

  Elise Gerich Merit

* = listed earlier

Architecture Retreat:

ASSUMPTIONS

o TCP/IP and OSI will coexist "indefinitely"

o Must continue to support diverse networks and services.

o Incorporate commercial and private networks, but don't depend on common carriers to solve all the problems or provide all the services.

o Able to scale to 10**9 networks.

Architecture Retreat:

POSSIBLE OUTPUTS:

o Architectural principles —> document(s)

o Recommendations to community —> document(s)

o New IETF WGs

o New R&D areas

GROUPS

1. ADDRESSING AND ROUTING

   Clark[chair], Braun, Chiappa, Estrin, Gross, Hinden, Jacobson, Lauck, Gerich (scribe).

2. MULTI-PROTOCOL ARCHITECTURE

   Chapin[chair], Callon, DCrocker, Huitema, Leiner, Postel.

3. SECURITY

   Cerf[chair], SCrocker, Kent

4. TRAFFIC CONTROL & STATE

   Braden[chair], Davin, Mills, Topolcic.

5. ADVANCED APPLICATIONS

   Hobby[chair], Borman, CLynch, DLynch, Reynolds, Schwartz, Schatz, Vaudreuil.

Architecture Retreat:

GROUPS AND ISSUES:

1. Routing and Addressing

"The Mother of All IP architecture issues."

- Add hierarchy to IP addresses
- Transition strategy
- Policy-based routing — architecture?
- Route servers to supplement/replace distributed routing
- AS/AD model retained for IP?
- Host identifiers: mobile hosts
- IP Address mask generalization

2. Multi-Protocol Architecture

Architectural issues which transcend particular protocol suite.

- How do we architect an Internet with n>1 protocol suites, regardless of what the suites are?
- How exactly will we define "the Internet"?
  - o Architecture for Multi-suite Internet
  - o Definition of Internet
  - Strategy —
    - o Emphasis: TCP/IP vs. OSI
    - o Merge/share suites? [Callon]
- Should we architect for partial or filtered connectivity?
  - Administratively-restricted connectivity
  - Relay architecture [Clark]
  - "Closed User Groups"
  - End-to-End security implications

3. Security:

This is very important to many areas.

o General architecture ?

o Secure network/Internet control mechanisms

o Global authentication

o End-to-end private communication

4. Traffic Control & State

The mechanics of change impinge on Addressing and Routing (only want to change IP once). It may also be somewhat premature, relative to research results.

- o New service classes — "real time": protected bandwidth, delay
- o Congestion control
- o CO vs. CL
- o Link sharing
- o Phase out fragmentation?
- o Accounting & billing

5. Advanced Applications

It was generally agreed in January that we need to work on applications and architecture for applications.

- Tools for building distributed systems
  - o transactions, multicasting, ...
- Important applications: [Hinden]
  - o Make email good enough for commerce.
  - o Desktop conferencing with video.
  - o Bulletin board paradigm: It is powerful and should be exploited more.
  - o Information collection (Knowbots?)
  - o Video retrieval and libraries.
  - o Distributed simulation.

7.2 Routing and Addressing: Dave Clark, MIT

Members: Clark, Braun, Chiappa, Estrin, Gross, Hinden, Jacobson, Lauck.

This group considered what changes need to be contemplated in the addressing and routing structure of the Internet to deal with its anticipated growth and functional evolution.

We concluded (as have others) that we must plan now for the following:

- We will run out of certain classes of network addresses, e.g., B addresses.
- We will run out of the 32 bit address space all together, as it is currently subdivided and managed.
- The total number of network numbers will grow to the point where we cannot run reasonable routing algorithms that route using network numbers as a basis.
- There will be a need for more than one route from a source to a destination, to permit variation in TOS and policy conformance. Both new applications and diverse transit services will drive this need. The source, or an agent acting for the source, must control the selection of the route options.

There was total agreement on the general approach needed to deal with these facts.

- We must move to an addressing scheme in which network numbers are aggregated into larger units, as a basis for routing. An example of an aggregate is the Autonomous System, or the Administrative Domain (AD). This aggregation is for several reasons: to define regions where policy is applied, to control the number of routing elements, and to provide elements for network management. Most of us believe that it must be possible to further combine aggregates, so that there is a nesting of ADs.
- We must provide some efficient means to compute common routes, and some general means to compute "special" routes. The general approach to special routes will be some form of route setup specified by a "source route". (We liked the general idea proposed by Deborah Estrin, based on her work with Yakov Rekhtor, of a mixed scheme in which the common routes are pre-computed and the special routes are provided on demand. Since the better methods for special route computation are based on link-state algorithms, it is unclear if this should also be the means for pre-computed routes.)

  There was disagreement on how we expect ADs to be aggregated, and how routing protocols should be organized to deal with the aggregation boundaries. Noel Chiappa proposed a very general scheme, while some would prefer a scheme which more restricts and defines the expected network model.

To deal with the address space exhaustion, we must either expand the address space, or reuse the 32 bit field (the 32BF) in different parts of the net. We conclude that there are several possible address formats that might make sense (see below). Perhaps more interesting is the question of how we migrate. All schemes will require that some routers (or other components inside the Internet) be able to rewrite headers so that hosts that expect the old or new format header are properly handled. Unless the need for conversion can be driven algorithmically, migration by itself will require some sort of setup of state in the conversion element.

We thus conclude that, for several reasons (special routes and address conversion, as well as accounting and resource allocation), we are moving from a "stateless gateway model, where only precomputed routes are stored in the gateway, to a model where at least some of the gateways have per-connection state.

We conclude that we should not plan a series of "small" changes to the architecture. We should embark now on a plan that will take us past the exhaustion of the 32BF. This is a more long-range act of planning than the Internet community has undertaken recently, but the problems of migration will require a long lead time, and it is hard to see an effective way of dealing with some of the more immediate problems, such as class B exhaustion, in a way that does not by itself take a long time. So, once we embark on a plan of change, it should take us all the way to replacing the current 32 bit global address space. (This conclusion is subject to revision if, as is always possible, some very clever idea surfaces that is quick to deploy and gives us some breathing room. We do not mean to discourage creative thinking about short-term actions. We just want to point out that even small changes take a long time to deploy.)

Conversion of the address space by itself is not enough. We must at the same time provide a more scalable routing architecture, and tools to better manage the Internet.

We propose the following immediate directions:

- Construct a specific set of estimates for the time at which the various problems above will arise, and construct a corresponding time-line for development and deployment of a new addressing/routing architecture. Use this time line as a basis for evaluating specific proposals for changes. This is a matter for the IETF.

- Take steps to make network aggregates (ADs) the basis of routing. We already have partial means to do this. IDPR does this. The OSI version of BGP (IDRP) does this. BGP could evolve to do this. The additional facility needed is a global table that maps network numbers to ADs. There are several options for this, which should be explored. This direction is a matter for the IETF.

- Continue the current work on policy based routing. There are several specific objectives.

- Seek ways to control the complexity of setting policy (this is a human interface issue, not an algorithm complexity issue).
- Better understand the issues of maintaining connection state in gateways.
- Better understand the issues of connection state setup.

- Explore, as a research activity, how we should aggregate ADs into still larger routing elements.
  - Consider whether the architecture should define the "role" of an AD or an aggregate.
  - Consider whether one universal routing method or distinct methods should be used inside and outside ADs and aggregates.
- Explore the options for a next generation address format. Develop a plan for migration. Specifically, construct a prototype gateway that does address mapping. Understand the complexity of this task, to guide our thinking about migration options.

Existing projects planned for DARTnet will help resolve several of these issues: state in gateways, state setup, address mapping, accounting and so on. Other experiments in the R &D community also bear on this area.

Appendix: address formats.

We considered three possible address formats. Briefly, they are as follows:

1. Replace the 32 bit field (32BF) with a field of the same size but with different meaning. Instead of being globally unique, it would now be unique only within some smaller region (an AD or an aggregate of ADs). Gateways on the boundary would rewrite the address as the packet crossed the boundary. Issues: addresses in the body of packets must be found and rewritten; the host software need not be changed; some method (perhaps a hack to the DNS) must set up the address mappings. (This scheme due to Van Jacobson. See also the work by Paul Tsuchiya on NAT.)

2. Replace the 32BF with a 64BF(or some other new size). Issues: must change the IP header, so host software must change.
   - Use the 64 bits to hold a global host address and an AD for that host. This makes possible a trivial mapping from the host to the value (the AD) which is the basis of routing. Common routes (those selected on the basis of destination address, without taking into account the source address as well) can be selected directly from the packet address, as is done today, without any prior setup.

- Use the 64 bits as a "flat" host identifier. Use connection setup to provide routers with the mapping from host id to AD as needed. The 64 bits can now be used to simplify the problem of allocating host ids, as in Ethernet addresses.

Each of these has in common the need to provide an address re-writing module as a part of migration.

ADDRESSING AND ROUTING

GENERAL APPROACH

1) We must aggregate networks into larger units (call them Administrative Domains, or ADs).
   - Specification of policy
   - Drive routing algorithms
   - Support network management

2) Need efficient means to compute "common" routes.

3) Need general means to compute and specify "special" routes.

At least for special routes, the means will be route setup based on source routes.

SLIDE 3

ADDRESSING AND ROUTING

ADDRESS EXHAUSTION

1) Reuse the 32 bit address in different parts of the network (e.g. the Van Jacobson proposal or the Paul Tsuchiya "NAB" proposal.)
   - Must "rewrite" the 32 address at the boundary of region.
   - If 32 address is sent undetected in body of packet, method fails.
   - Seems like a hack. Security people are concerned.
   - Host software "might not change".

2) Change IP header to have 64 bit address.

Migration to the new scheme is the hard part. In both cases, a "packet rewriter" is required.

SLIDE 4

ADDRESSING AND ROUTING

STATE IN GATEWAYS

This proposal implies that gateways will change.

Today: "connectionless" -- know about routes, but not flows of packets.

Tomorrow: some state about flows will be required.
- Rewriting addresses
- Accounting
- Resource management
- Remembering the route

SLIDE 5

ADDRESSING AND ROUTING

WHAT'S IN AN ADDRESS?

We propose that routing be done on ADs, not network numbers.

Today, packets have network numbers in them.

The gateway must convert, to route the packet.

1) Construct a (rather static) table and propagate.
2) Put the AD into the address (in the 64 bits?)
3) Always do a route setup, even for common routes, and have the gateway remember.

SLIDE 6

ADDRESSING AND ROUTING

AGGREGATES OF NETWORKS

We agree that there must be aggregates of networks (ADs) to control network explosion.

But are ADs themselves aggregated? Disagreement here.

- If nesting of ADs is needed, is it a hierarchy (a partition) or a general graph?
- Must relate to (and simplify) how routes are computed.

SLIDE 7

ADDRESSING AND ROUTING

GROUPING NETWORKS INTO ADs

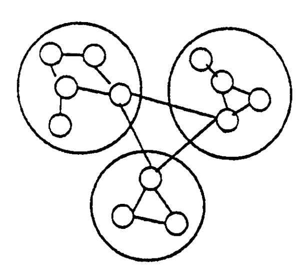

Except for interdomain links, a strict partition.

SLIDE 8

ADDRESSING AND ROUTING

GROUPING ADs INTO BIGGER AREAS

Option 1: **strict nesting.** **Option 2: overlapping regions.**

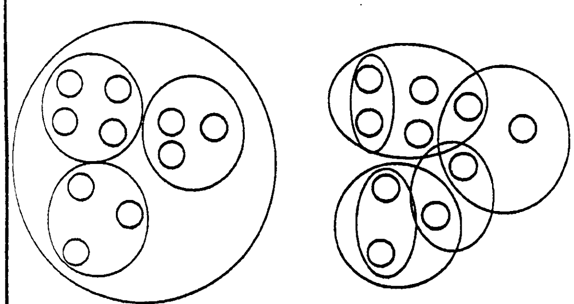

SLIDE 9

ADDRESSING AND ROUTING

LOOK AT THE BIG PICTURE

It is not enough to find a hack to extend the address space.

Unless we resolve the routing problem (size, source controlled routes, etc.) we will still die of excess growth.

Even simple hacks take a long time to deploy.

Our vote: take the big step. (Some disagreement here; we all hope for the easy way out.)

Migration is probably the hardest part of the problem.

SLIDE 10

ADDRESSING AND ROUTING

SPECIFIC ACTIONS

Work out a **specific time line. (For the problem and the** solution.) **Bound the time for random proposals.**

Make ADs the **basis of routing.**

Start to move **away from the current address classes** (address **masks in routing packets).**

Continue work **on policy routing.**

Study **aggregates of ADs. (research)**

Explore **options for next generation addresses. Relate to** routing.

Consider **problems of net management.**

SLIDE 11

7.3 Multi-Protocol Architecture: Lyman Chapin, BBN

Members: Chapin, Callon, D. Crocker, Huitema, Leiner, Postel

This group was asked to consider the general issue of "the multiprotocol Internet", and to specifically address three questions:

1. How do we architect an Internet with n¿1 protocol suites, regardless of what the suites are?

2. How exactly will we define "the Internet"?

3. Should we architect for partial or filtered connectivity?

It is very difficult to deal constructively with the issue of "the multiprotocol Internet" without first determining what we believe "the Internet" is (or should be). We distinguish the Internet (a set of communicating systems) from the Internet community (a set of people and organizations). Most people would accept a loose definition of the latter as "the set of people who believe themselves to be part of the Internet community"; however, no such "sociological" definition of the Internet itself is likely to be useful.

Not too long ago, the Internet was defined by IP connectivity (IP and ICMP were - and still are - the only "required" Internet protocols). If I could PING you, and you could PING me, then we were both on the Internet, and a satisfying working definition of the Internet could be constructed as a roughly transitive closure of IP-speaking systems. This model of the Internet was simple, uniform, and - perhaps most important - testable. The IP-connectivity model clearly distinguished systems that were "on the Internet" from those that were not.

As the Internet has grown, and the technology on which it is based has gained widespread commercial acceptance, the sense of what it means for a system to be "on the Internet" has changed, to include

- Any system that runs the TCP/IP protocol suite, whether or not it is actually accessible from other parts of the Internet.

- Any system that can exchange RFC-822 mail (without the intervention of mail gateways or mail object transformations).

- Any system with e-mail connectivity to the Internet, whether or not a mail gateway or mail object transformation is required.

These definitions of "the Internet", however, are still based on the original concept of connectivity, just "moving up the stack".

We propose instead a new definition of the Internet, based on a different unifying concept:

- "Old" Internet concept: IP-based; the organizing principle is the IP address (common network address space).

- "New" Internet concept: Application-based; the organizing principle is the domain name system and directories (common - albeit necessarily multiform - application name space).

This suggests that the idea of "connected status", which has traditionally been tied to the IP address (via network numbers), should instead be coupled to the names (and related identifying information) contained in the distributed Internet directory.

A naming-based definition of "the Internet" implies a much larger Internet community, and a much more dynamic (and unpredictable) operational Internet. This argues for an Internet architecture based on adaptability (to a broad spectrum of possible future developments) rather than anticipation. Rather than specify a particular "multi-protocol Internet", embracing a pre- determined number of specific protocol architectures, we propose instead a process-oriented model of the Internet, which accommodates different protocol architectures according to the traditional "things that work" principle. This model includes, as a basic postulate, the assertion that there is no *steady-state* "multiprotocol Internet". The most basic forces driving the evolution of the Internet are pushing it not toward multiprotocol diversity, but toward the original state of protocol-stack uniformity (although it is unlikely that it will ever actually get there). We have represented this tendency of the Internet, as a complex system, to evolve in favor of homogeneity as the most "thermodynamically stable" state by describing four components of a new process-based Internet architecture:

1. The core Internet architecture.

   This is the traditional TCP/IP-based architecture. It is the "magnetic center" of Internet evolution, recog- nizing that (a) homogeneity is still the best way to deal with diversity in an internetwork, and (b) IP connectivity is still the best basic model of the Internet (whether or not the actual state of IP ubiquity can be achieved in practice in a global operational Internet).

   > "In the beginning", the Internet architecture consisted only of this first part. The success of the Internet, however, has carried it beyond its uniform origins; ubiquity and uniformity have been sacrificed in order to greatly enrich the Internet "gene pool". Two additional parts of the new Internet architecture express the ways in which the scope and extent of the Internet have been expanded.

2. "Link sharing".

   Physical resources (transmission media, network inter- faces, perhaps some low-level (link) protocols) are shared by multiple, non-interacting protocol suites. This part of the architecture recognizes the necessity and convenience of coexistence, but is not concerned with interoperability; it has been called "ships in the night". Coexisting protocol suites are not, of course, genuinely isolated in practice; the ships passing in the night raise issues of management, non-interference, coordination, and fairness in real Internet systems.

3. Application interoperability.

   Absent ubiquity of interconnection (i.e., interoperability of the "underlying stacks"), it is still possible to achieve ubiquity of application functionality, by arranging for the essential semantics of applications to be conveyed among otherwise non- interconnected communities of Internet systems. This can be accomplished by application relays, or by combined user agents (which present a uniform virtual access method to different application services that expresses only the shared semantics). This part of the architecture emphasizes the ultimate role of the Internet as a basis for communication among applications (rather than as an end in itself); to the extent that it enables a population of applications (and their users) to move from one underlying protocol suite to another without unacceptable loss of functionality, it is also a "transition enabler".

   > Adding parts 2 and 3 to the original Internet architecture is at best a mixed blessing. Although they greatly increase the scope of the Internet and the size of the Internet community, they also introduce significant problems of complexity, cost, and management, and they usually represent a loss of functionality (particularly with respect to part 3). Parts 2 and 3 represent unavoidable, but essentially undesirable, departures from the homogeneity represented by part 1; some functionality is lost, or additional system complexity and costs are endured, in order to expand the scope of the Internet. In a perfect world, however, the Internet would evolve and expand without these penalties. There is a tendency, therefore, for the Internet to evolve in favor of the homogeneous architecture represented by part 1, and away from the compromised architectures of parts 2 and 3. Part 4 expresses this tendency.

4. Hybridization/integration.

   This part expresses the tendency of the Internet, as a system, to attempt to return to the original "state of grace" represented by the uniform architecture of part 1. It is a force acting on the evolution of the Internet, rather than a process whereby the Internet actually returns to a uniform state at some point in the future. Part 4 recognizes the desirability of integrating similar elements from different Internet protocol architectures to form hybrids that reduce the variability and complexity of the Internet system. It also recognizes the desirability of leveraging the existing Internet infrastructure to facilitate the absorption of "new stuff" into the Internet, applying to "new stuff" the established Internet practice of test, evaluate, adopt.

According to this dynamic model, running X.400 mail over RFC 1006 on a TCP/IP stack, integrated IS-IS routing, transport gateways, and the development of a single common internetwork protocol successor to IP and CLNP are all examples of "good things" - they represent movement away from the non-uniformity of parts 2 and 3 in the direction of greater homogeneity, under the influence of the "magnetic field" asserted by part 1, following the hybridization dynamic of part 4.

Multi-protocol Internet Architecture

IAB/IESG Internet Architecture Workshop

San Diego • June 11–13, 1991

Lyman Chapin
Dave Crocker
Barry Leiner
Jon Postel
Ross Callon
Christian Huitema

MPIA Charter

- ❑ How do we architect an Internet with n>1 protocol suites, regardless of what they are?
- ❑ How exactly will we define "the Internet"?
- ❑ Should we architect for partial or filtered connectivity?

"The Internet is the community of people who ought to be interested in what the IETF ought to be working on."

IP Connectivity

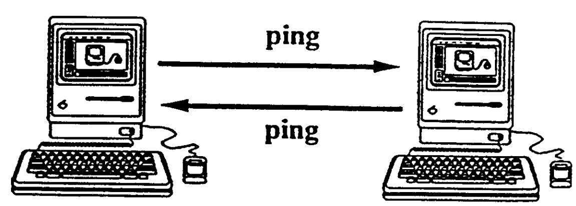

- ❑ simple
- ❑ uniform
- ❑ testable

Internet as a Set of Systems

set membership based on:

- ❑ IP connectivity
  - A is on the Internet (by definition);
  - B can exchange IP packets with A;
  - therefore B is on the Internet
- ❑ "It does IP"
- ❑ 822 mail (no mail gateways or transformations)
- ❑ e-mail connectivity (some mail object transformation)

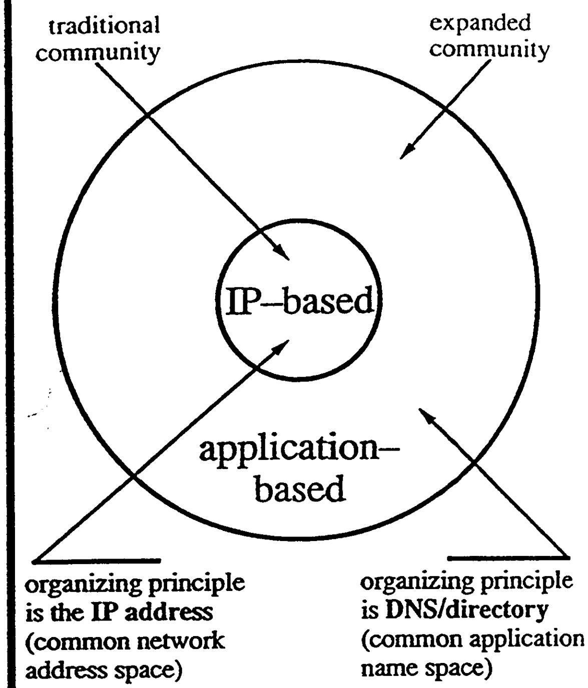

Process Model of the Internet

adaptability, not anticipation

the "**things that work**" principle

no **steady–state** "multiprotocol Internet"

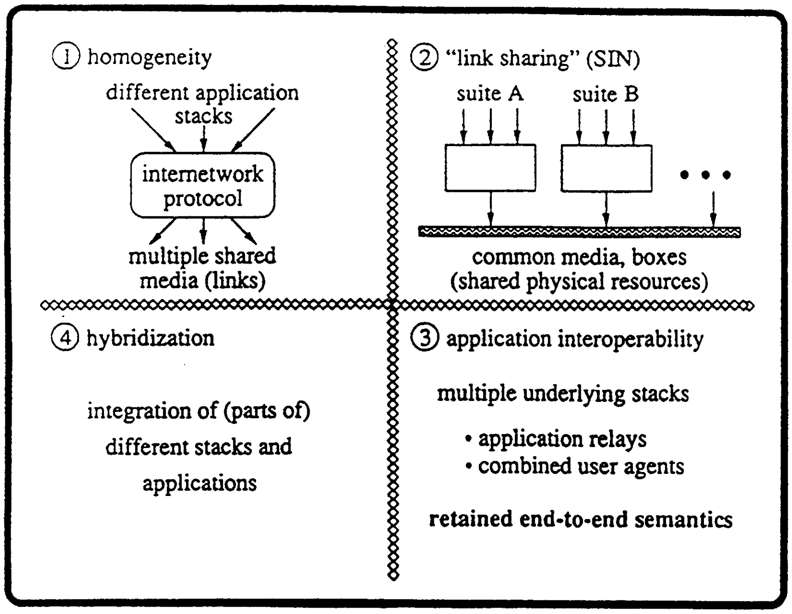

1 Core Internet Architecture

homogeneity is **still** the best way to deal with diversity

IP connectivity is **still** the best basic model of the Internet

the "magnetic center" of Internet evolution

2 Link Sharing

multiple, non-interacting protocol suites

"When one horse is no longer the only game in town, it's either 'this town ain't big enough for both of us', or let the ships pass in the night . . ."

co-existence, not interoperability

non–oblivious ships

(3) Application Interoperability

convey essential application semantics between two otherwise non-interconnected communities

~~ubiquity of interconnection (interoperability of underlying "stacks")~~

ubiquity of application functionality

- application relays
- combined user agents (access methods)

Internet population/extent

Transition enabler

(4) Hybridization / Integration

attempt to return to the original state of grace

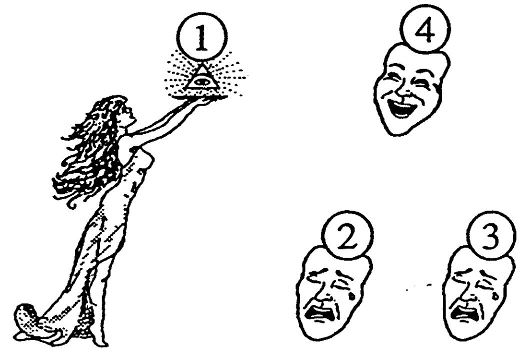

crippled functionality
complexity, cost, manageability

(4) Hybridization / Integration

— Examples —

Moving away from (3)

- running X.400 over an RFC 1006 stack
- adoption of NNTP as an Internet protocol

Moving away from (2)

- integrated IS-IS routing
- transport gateways
- Internet integration of Berkeley r-commands
- common successor to IP and CLNP

Architect for the integration / hybridization dynamic, not multi–protocol steady state

What Next?

- Document the core Internet architecture
- What does it really mean to say that the Internet is based on the concept of a common application name space?
- Apply hybridization/integration principle ("magnetic field") to guide Internet evolution

7.4 Security: Vint Cerf, CNRI

Members: Cerf, S. Crocker, Kent, Mockapetris

Philosophical Guidelines:

The principal themes were simplicity, testability, trust, technology and security perimeter identification. It was emphasized that there is more to security than protocols and cryptographic methods. The security architecture and policies should be simple enough to be readily understood. Complexity breeds misunderstanding and poor implementation. The implementations should be testable to determine if the policies are met. We are forced to trust hardware, software and people to make any security architecture function. We assume that the technical instruments of security policy enforcement are at least as powerful as modern personal computers and work stations; we do not require less capable components to be self-protecting (but might apply external remedies such as link level encryption devices). Finally, it is essential to identify security perimeters at which protection is to be effective. Several possible perimeters were identified.

Plausible Security Perimeters:

There were four possible perimeters: link level, net/subnet level, host level and process/application level. Each imposes different requirements, can admit different techniques, and makes different assumptions about what components of the system must be trusted to be effective. Privacy Enhanced Mail is an example of a process level security system; providing authentication and confidentiality for SNMP is another example. Host level security typically means applying an external security mechanism on the communication ports of a host computer. Network or subnetwork security means applying the external security capability at the gateway/router(s) leading from the subnetwork to the "outside." Link level security is the traditional point to point or medium (e.g., for Ethernet) level encryption mechanism.

Applying protection at the process level assumes that the underlying scheduling and operating system mechanisms can be trusted not to prevent the application from applying security when appropriate. As the security perimeter moves downward in the system architecture towards the link level, one must make many assumptions about the security threat to make an argument that enforcement at a particular perimeter is effective. For example, if only link level encryption is used, one must assume that attacks come only from the outside via communications lines and that hosts, switches and gateways are physically protected, the people and software in all these components are to be trusted.

Desired Security Services:

It was concluded that we need authenticable, distinguished names if we are to implement discretionary and non-discretionary access control at application and lower levels in the system. In addition, we need enforcement for integrity (anti-modification, anti-spoof and anti-replay defenses), confidentiality, and prevention of denial of service. It was speculated that we also need ways to prevent repudiation of message transmission and covert channel protection for some situations.

We have some building blocks with which to build the Internet security system. Cryptographic algorithms (e.g., Data Encryption Standard, RSA, El Gamal and possibly other public key and symmetric key algorithms) are available, as are hash functions such as MD2 and MD5. We need Distinguished Names (in the OSI sense) and are very much in need of an infrastructure for the assignment of such identifiers together with widespread directory services for making them known. Certificate concepts binding distinguished names with public keys and binding distinguished names to capabilities and permissions may be applied to good advantage.

At the router/gateway level, we can apply address and protocol filters and other configuration controls to help fashion a security system. It was agreed that the proposed OSI Security Protocol 3 (SP3) and Security Protocol 4 (SP4) be given serious consideration as possible elements of an Internet security architecture. Finally, it was observed that we have no good solutions to safely storing secret information (such as the secret component of a public key pair) on systems not designed to enforce secure storage (such as a random PC or laptop).

Proposed Next Steps:

For Privacy Enhanced Mail, the most critical steps seem to be getting the certificate generation and management infrastructure in place, together with X.500 directory services to provide access to public keys via distinguished names. Serious attention needs to be placed on any limitations imposed by patent and export restrictions on the deployment of this system.

An examination of methods for dealing with security in distributed systems applications is called for in both simple (client/server) and complex (distributed computing environment) cases. The use of certificates granting permissions/capabilities to objects bound to distinguished names should be examined for utility, for example.

For host-oriented security, it was proposed to evaluate SP4 in particular but also to consider SP3 as the protocol basis for this architectural level of security protection.

There were many open questions about network/subnetwork security protection, not the least of which was a potential mismatch between host level (end/end) security methods and network/subnetwork level. Moreover, it was observed that network level protection does not deal with threats arising within the security perimeter.

A Security Reference Model for the Internet is needed and should be developed expeditiously. The model should establish the target perimeters and document the objectives of the security architecture.

We should move ahead quickly with deployment of a distinguished name service (e.g., X.500) and should carry out implementations of application level security services both for their immediate utility (e.g., PEM, SNMP authentication) and to gain valuable practical experience which can inform the refinement of the Internet security architecture.

SECURITY ARCHITECTURE

S. Kent, S. Crocker
P. Mockapetris and V. Cerf

* Philosophical Guidelines

  Simplicity of Design and Policy

  Testability

  Trust

  Technology (PC/WS)

  Security Perimeter Identification

* There is more to Security

  than Protocols!

PLAUSIBLE SECURITY PERIMETERS

* Process level
  Application
  Object
  Note: SNMP/Routing -> appl.

* Host level

* Administrative Domain (Net?)

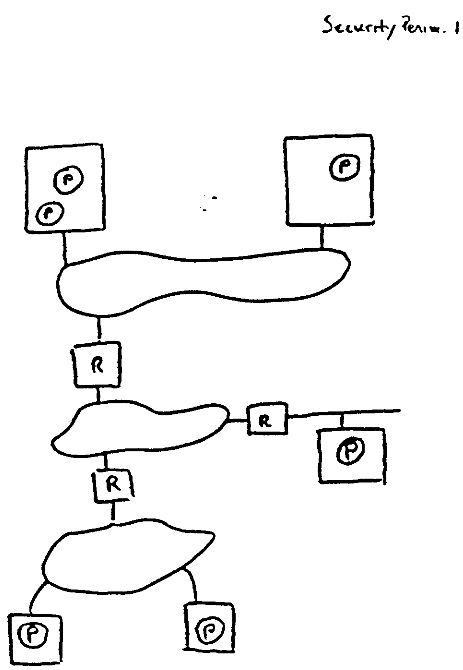

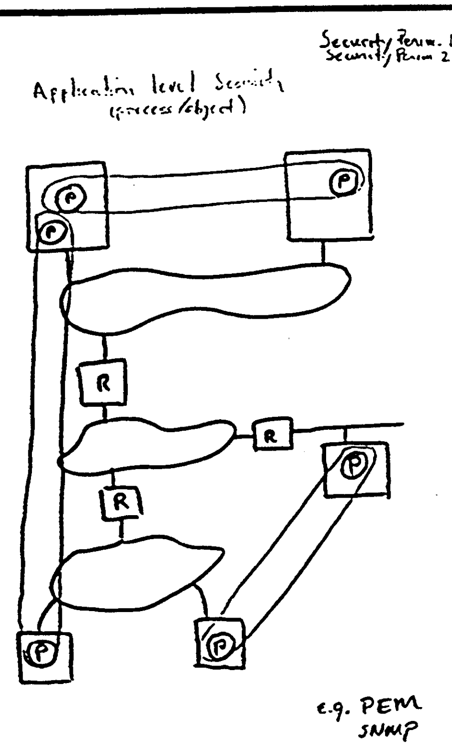

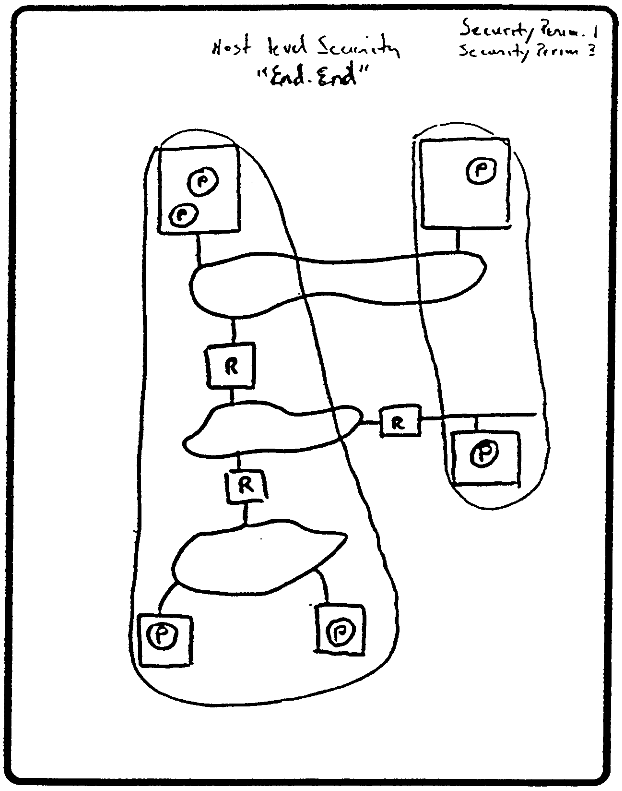

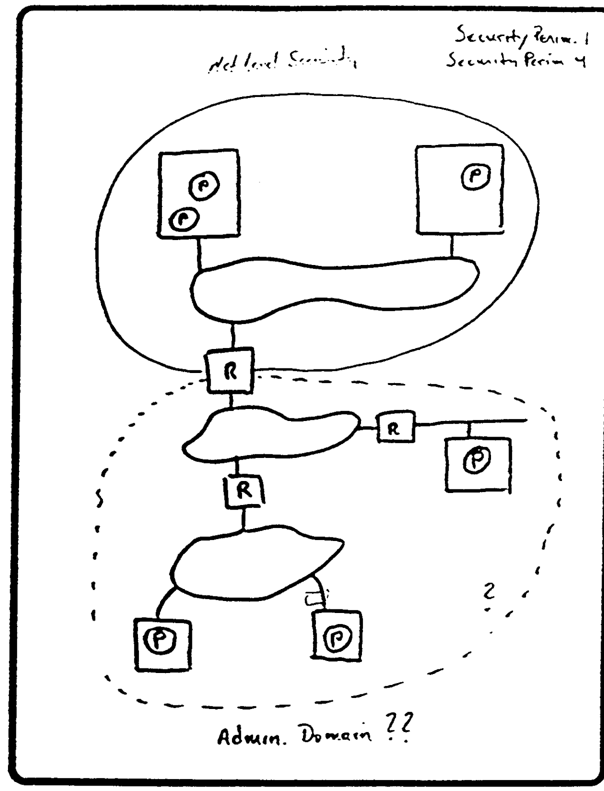

CURRENT ACCESS CONTROL IN THE INTERNET

* Login -> you have an account
  (Poor Password Protection)

* Non-discretionary Services

  - Anonymous FTP
  - Time Service
  - EMAIL

* Need/Want Discretionary, Distributed application-level access control

  e.g. Knowbot® programs w/authoriz. application level guards

Note: if we want Access Control Lists, need Authenticable Distinguished names!

DESIRED SECURITY SERVICES

* Authentication
* Integrity (Anti-mod, -spoof, -replay)
* Confidentiality
* Non-repudiation
* Continuation of Service*
* Covert Channel Protection(?)

"non-denial of Service"

BUILDING BLOCKS

* Cryptography (DES, RSA, ElGamal?)
* Hash Functions (MD2, MD5?)
* DISTINGUISHED NAMES
* Certificates:
  signed (DN, PubKey)
* Private Attribute Certificate:
  signed(DN,Attrib List)
* SP3 (OSI draft document)
* SP4 (OSI less drafty document)
* Addr/protocol filters (in routers)
* Configuration controls

A MISSING ELEMENT

Wanted:

Architecture for safe storing of secret information on systems not designed to enforce secure storage...

HOST-ORIENTED SECURITY

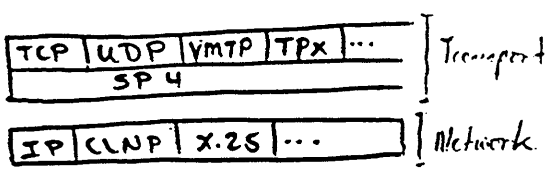

- SP4 does not define/specify key management
- SP4 encapsulates Transport
- SP4 provides integrity, confidentiality, Authenticity

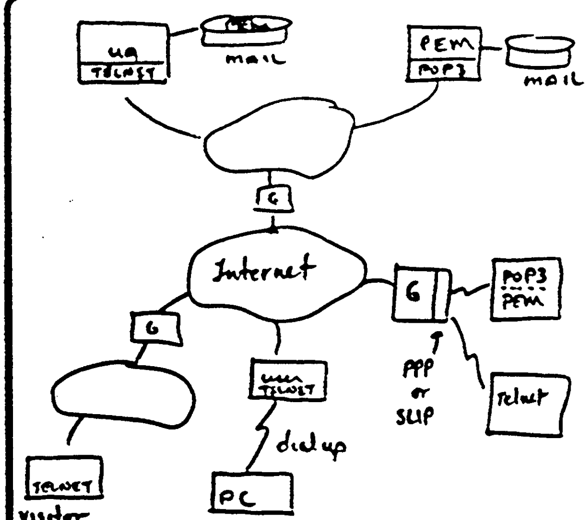

POP3 → ok
Telnet → insecure

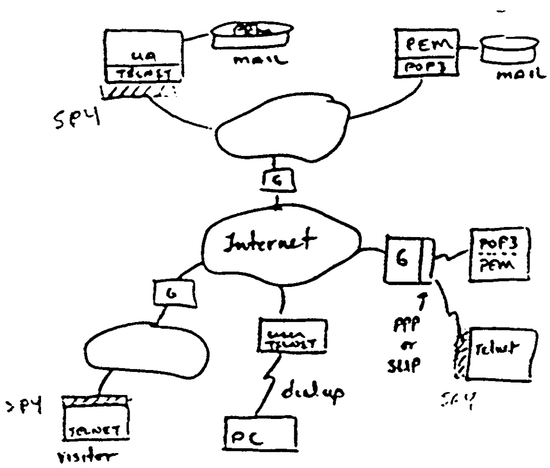

POP3 → ok
Telnet → insecure

issue: Key distribution!

NETWORK/-ORIENTED SECURITY

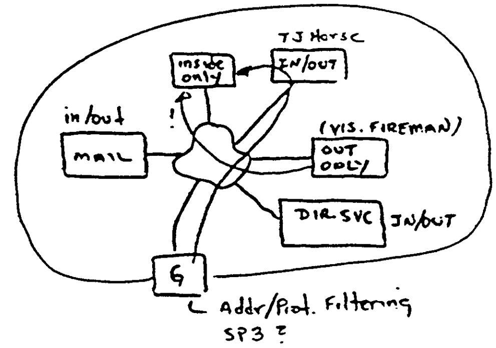

Issues: Performance (video!...)
Configuration mgmt

Network Perimeter doesn't deal with intranet access control and trojan horses. Need Host/Application level Security.

TOP-LEVEL ACTIONS NEEDED

1. Security Reference Model
   establish perimeters
   document objectives
2. Distinguished Name Service
3. Application Level Security Services
4. Proceed with SP4
   TP4/SP4 avail. from NIST
   Key Mgmt Protocol (KMP) for SP4 uses X.509. Could use PEM Certificates
5. Use Kerberos in addition to KMP
6. Need Common Authentication Tech. This API needs well-defined Distinguished Names.

MORE TOP-LEVEL ACTIONS

7. PEM needs certificate hierarchy and certificate generation mechanism. Could Internet Society or IETF hand them out??
8. Need proving ground for secure, distributed applications
9. Compare ISO Security Architecture with Internet requirements.

   * Multiprotocol requirements esp!
10. Begin protecting infrastructure:
    Design security for OSPF
11. Develop dependency model:
    who/what is trusted?
    To do what?
12. Re-think Operational Security

SECURITY FOR APPLICATIONS

* PEM: Authentication
  Integrity
  Confidentiality

  RSA, DES, Certificates, DN's

  -> Certificate Management Struct.
  -> Patent, Export limitations

* Directory Services (X.500)
  Distinguished Name Registr.
  Certificate Distribution

MORE ON APPLICATIONS

* Distributed Systems and Remote Client/Servers

  Privileged Attribute Certificates:

  Giver Signs: [DN, Attr. List]

  Use DN to find associated Public Key and perform 3-way handshake challenge/response.

  Issue: is secret key readily available? (to application)

7.5 Traffic Control and State

Members: Braden, Davin, Mills, and Topolcic

In the present Internet, all IP datagrams are treated equally. Each datagram is forwarded independently, regardless of any relationship it has to other packets for the same connection, the same application, the same class of applications, or the same user class. There are Type-of-Service and Precedence bits defined in the IP header, but these are not generally implemented and in fact it is not clear how to implement them. No promises are made; forwarding is strictly "best-effort".

It is now widely accepted that the future Internet will need to support important applications – e.g., packet video and voice for teleconferencing – for which best-effort is not sufficient. This will require some "traffic control" mechanism in routers, controlled by additional state, to handle "real-time" traffic.

Assumptions and Principles

- ASSUMPTION: The Internet will need to support performance guarantees for particular subsets of the traffic.

  However, we are far from being able to give precise meanings to the terms "performance", "guarantees", and "subsets" in this statement. Reasonable people differ, and much R &D is still necessary.

- The default service will continue to be the current "best- effort" datagram delivery.

- The mechanism of a router can be separated into the forwarding path, which is highly optimized, sometimes with hardware- assist, and therefore costly and difficult to change; and the routing and resource control computations that take place in background. We will have at most one shot at changing the forwarding paths of routers, so we had better get it right the first time.

- The new extensions must operate in a highly heterogeneous environment, in which some parts will never support guarantees. For some hops of a path (e.g., a high-speed LAN), over-provisioning must be usable instead of explicit resource reservation for real-time traffic.

- Multicast distribution is probably essential.

Technical Issues

- "Hard" state (as in ST-2) vs. "soft" (i.e., cached) state in routers?

- Resource binding vs. route binding?

The group discussed alternative designs that might allow the logical separation of allocating resources in nodes route determination.

- Alternative multicast models

  IP multicasting uses a model of logical addressing in which targets attach themselves to a group. In ST-2, each host in a multicast session includes an explicit list of target addresses in its setup packet. Each of these approaches has advantages and drawbacks; it is not currently clear which will prevail for n-way teleconferences.

- Resource Setup vs. Inter-AD routing

  Resource guarantees of whatever flavor must hold across an arbitrary end-to-end path, including multiple ADs. Hence, any resource setup mechanism needs to mesh smoothly with the path setup mechanism incorporated into IDPR.

- Accounting

  The resource guarantee subsets ("classes") may be natural units for accounting.

Suggested Actions

The actions proposed by the group are generally concerned with R &D in this area. DARTnet, the DARPA Research Testbed network, will play an important role.

IAB RETREAT ON THE FUTURE

OF THE INTERNET ARCHITECTURE

GROUP 4: TRAFFIC CONTROL & STATE

Braden, Davin, Mills, & Topolcic.

- Outline:
  - Assumptions
  - Principles
  - Technical Issues
  - Suggested Actions

Traffic Control & State: ASSUMPTIONS

- The Internet architecture will need to support *performance guarantees* for traffic *subsets.*
  - *"Performance":* bandwidth, delay, jitter ...
    [we really don't know]
  - *"Guarantees":* roughly: isolation or protection.
    ("best-effort guarantees"?)
  - *"subsets":* => "flows" or "classes", ranging from broad user classes [link-sharing] to individual packet video streams.

- The default service will continue to be "best-effort" datagram delivery, ubiquitously.
- Two parts of the problem in a router:
  - The forwarding path
    - Implements queueing discipline
    - hard and expensive to change; do it < 2 times!
  - Background code (like routing calculation)
    - will implement resource controller
    - easier to change.

PRINCIPLES

- Accomodate Internet heterogeneity
  - Some parts of Internet will never support "guarantees".
  - Every part must at least say "I can't" politely.
- Liberal "guarantees" must be possible
  - "I want X, but give me what you can..."
  - Can use over-provisioned links, even if no guarantees.
- Multicasting is probably an essential part of the service.

TECHNICAL ISSUES

- "Hard" state [ST-2] or "soft" (cached) state in routers?
- Resource binding vs. route binding? *Three approaches:*
  - (1) "crank-back mechanism [ST-2]
    => lots of state, coupling.
  - (2) Allocate at source
    => Link-state tree has current resource allocations on each link.
    Convinced ourselves: race and deadlock problems are solvable.
  - (3) Decouple
    => Pick route, then allocate resources. May fail unnecessarily.

TECHNICAL ISSUES (2)

- IP Multicast vs. source-list multicast.
- Resource setup vs. IDPR.
- Long-term service scheduling
  - Management function
  - Issue certificate for time slot
  - Preempt non-scheduled flows

TECHNICAL ISSUES (3)

- Accounting
  - Accounting unit should be resource guarantee unit.
  - Could be service given [counts] or simply service guaranteed.
  - Best-effort traffic: may not do detailed accounting on aggregated flows.

SUGGESTED ACTIONS

A. FORWARDING MECHANISM: *ongoing R&D*

1. DARTnet: try out alternative forwarding-path mechanisms — Van's, Flow Protocol, etc.

B. SETUP/RESOURCE CONTROL:

1. Use DARTnet for experiments.
2. Develop soft-state setup for Van's mechanism.
3. Experiment with ST-II hard-state setup.

4. Use IDPR to as experimental platform for resource setup experiments.
5. Fold data capture ("accounting") into these efforts.

7.6 Advanced Applications

Members: Hobby, Borman, CLynch, Reynolds, Schatz, Schwartz, Vaudreuil.

The first questions asked by the Advanced Applications Group was "What are the network based applications that we want and why don't we have them now?" The group came up with a large list of potential applications, most of which were client/server based. However, the more interesting part of the question was why haven't people done them already.

The answer that the group came up with was that the tools to make application writing easy just do not exist. The next step was to define the tools that would be necessary for application writers to do their job. One of the most basic tools needed is a Common Interchange Format for a number of data items that will be used across the network. The applications have to know the format of information that they are exchanging for the information to have any meaning. To this end, several types of data formats were defined.

Common Interchange Formats

- Text - Of the formats in this list, text is the most stable, but today's international Internet has to address the needs of other character sets as well.
- Image - As we enter the "Multi-media Age", images will become increasingly important, but we need to agree on how to represent them in "bits".
- Graphics - Like images, vector graphic information needs a common definition. With such a format we could exchange things like architectural blueprints.
- Video - Before we can have a video window running on our workstation we need to know the format of that video information coming over the network.
- Audio/Analog - Of course, we also need the audio to go with the video, but such a format would be used for representation of all types of analog signals.
- Display - Now that we are opening windows on our workstation, we want to open a window on another person's workstation to show her some data pertinent to the research project, so now we need a common window display format.
- Data Objects - For inter-process communications we need to agree on the formats of things like integers, reals, strings, etc.

Many of this formats are being defined by other, often several other, standards organizations. We need to agree on one format per category for the Internet.

Once we have defined the common formats of the data, we need tools that the applications can use to move the data easily. The group came up with several methods of data exchange that applications would require.

Data Exchange Methods

- Store and Forward - Not everyone is on the network all the time. We need a standard means of providing an information flow for these host's applications. Multicast capabilities would also be important to the store and forward mechanism.

- Global File Systems - Much of the data access over the network can be broken down to simple file access. If you had a real global file system where you access any file on the Internet (assuming you have permission) would you ever need FTP?

- Inter-process Communications - This includes RPC, API, etc. For true distributed computing environment, we need the means to allow processes to exchange data in a standard method over the network.

- Data Broadcast - Many of the applications need to send the same information to many other hosts. A standard, efficient method is needed to accomplish this.

- Database Access - For good information exchange, we need to have a standard means for accessing databases. The Global File System can get you to the data, but the database access methods will tell you about its structure and content.

Again, many of these items are being addresses by other organizations but for Internet interoperability, we need to agree on the methods for the Internet.

The Advanced Applications Group came up with requirements from two other groups at the Workshop. From the Traffic Control Group, applications need the ability to transmit real-time data. This means some sort of expectation level as far as data delivery within a certain time-frame. Applications also require from the Security Group global authentication and access control systems. Much of the usefulness of today's Internet applications is lost due to the lack of trust and security. This needs to be solved for tomorrow's applications.

San Diego IAB Workshop
Advanced Applications Group

Common Interchange Formats
- Text
- Image
- Graphics
- Video
- Audio/Analog
- Display
- Data Objects

Data Exchange Methods
- Store and Forward
- Global File Systems
- Inter-process Communications
- Data Broadcast
- Database Access

Needs From Other Groups
- Traffic Control
  - Real-time Data Transmission
- Security
  - Global Authentication
  - Global Access Control

Appendix A

Attendees

Vikas Aggarwal
Princeton University
(W)609-258-2403
vikas@JVNC.net

Steve Alexander
INTERACTIVE Systems Corporation
(W)708-505-9100
stevea@i88.isc.com

Claudio Allocchio
National Institute
for Nuclear Physics
(W)39 40 3758523
claudio.allocchio@cosine-gw.infn.it

Philip Almquist
Consultant
(W)415-752-2229
almquist@jessica.stanford.edu

Nagaraj Arunkumar
3Com Corporation
(W)408-264-5104
nak@3com.com

Karl Auerbach
Sun Microsystems, Inc.
(W)415-336-2090
karl@eng.sun.com

Fred Baker
Advanced Computer
Communications
(W)805-963-9431
fbaker@emerald.acc.com

Richard Balcon
Oracle Corporation
(W)415-506-2924
rbalcon@oracle.com

Stan Barber
Baylor College of Medicine
(W)713-798-6042
sob@bcm.tmc.edu

Doug Barlow
Digital Equipment Corporation
(W)206-562-8552
barlow@decwet.dec.com

James Barnes
XYLOGICS, Inc.
(W)617-272-8140
barnes@xylogics.com

William Barns
The MITRE Corporation
(W)703-883-6832
barns@gateway.mitre.org

Scott Barvick
Wellfleet Communications, Inc.
(W)617-275-2400
sbarvick@wellfleet.com

Richard Basch
Massachusetts Institute of
Technology
(W)617-253-0100
basch@mit.edu

Charles Bazaar
Emulex Corporation
(W)714-668-5416
bazaar@emulex.com

Jordan Becker
Advanced Network & Services, Inc.
(W)914-789-5315
becker@nis.ans.net

Jim Beers
Cornell University
(W)607-255-9917
beers@nr-tech.cit.cornell.edu

Stevan Belair
GTE
(W)508-880-4285
swbelair@sprintf.merit.edu

Tom Benkart
Advanced Computer
Communications
(W)301-290-8100
teb@saturn.acc.com

Arthur Berggreen
Advanced Computer
Communications
(W)805-963-9431
art@acc.com

William Biagi
Corporation for Open Systems International
(W)703-883-2806
bbiagi@cos.com

David Bolen
Advanced Network & Services, Inc.
(W)914-789-5327
db31@nis.ans.net

Nelson Bolyard
Silicon Graphics, Inc.
(W)415-335-1919
nelson@sqi.com

Peter Boos
Bell Northern Research
(W)613-763-5366
peterb@bnr.ca

David Borman
Cray Research, Inc.
(W)612-683-5571
(W)612-452-6650
dab@cray.com

Steve Bostock
Novell, Inc.
(W)408-473-8203
steveb@novell.com

Helen Bowns
Bolt, Beranek, and Newman
(W)617-873-2411
hbowns@bbn.com

Robert Braden
Information Sciences Institute
(W)213-822-1511
braden@isi.edu

Scott Bradner
Harvard University
(W)617-495-3864
sob@harvard.edu

David Brent
University of British Columbia
(W)604-822-6537
brent@CDNnet.ca

David Bridgham
Epilogue Technology
(W)415-594-9268
dab@asylum.sf.ca.us

Scott Brim
Cornell University
(W)607-255-9917
swb@nr-tech.cit.cornell.edu

Thomas Brisco
Rutgers University
(W)908-932-2351
brisco@rutgers.edu

Caralyn Brown
Wellfleet Communications, Inc.
(W)617-275-2400
cbrown@wellfleet.com

Howard Brown
Cabletron Systems, Inc.
(W)603-332-9400
brown@ctron.com

Gregory Bruell
Shiva Corporation
(W)617-252-6300
gob@shiva.com

Mats Brunell
Swedish Institute of
Computer Science
(W)46 87521500
(W)46 87521563
(W)8126254 SICS S
mats.brunell@sics.se

Theodore Brunner
Bell Communications Research
(W)201-829-4678
tob@thumper.bellcore.com

Philip Budne
Shiva Corporation
(W)617-252-6300
phil@shiva.com

Carter Bullard
Georgia Institute of Technology
(W)404-853-0202
carter@gatech.edu

John Burruss
Wellfleet Communications, Inc.
(W)617-275-2400
jburruss@wellfleet.com

Randy Butler
National Center for
Supercomputing Applications
(W)217-244-8285
rbutler@ncsa.uiuc.edu

Ross Callon
Digital Equipment Corporation
(W)508-486-5009
callon@bigfut.enet.dec.com

Eric Carroll
University of Toronto
(W)416-978-6134
(W)416-961-7187
eric@utcs.utoronto.ca

Charles Carvalho
Advanced Computer
Communications
(W)805-963-9431
charles@sage.acc.com

Jeffrey Case
University of Tennessee
(W)615-573-1434
case@cs.utk.edu

Cerafin Castillo
Emulex Corporation
(W)408-452-4777
cec@emulex.com

Vinton Cerf
Corporation for National
Research Initiatives
(W)703-620-8990
vcerf@nri.reston.va.us

Cho Chang
Hewlett-Packard
(W)508-256-6600
chang_c@apollo.hp.com

John Chang
US West Advanced Technologies
(W)303-541-6255
jrc@uswest.com

Peter Chang
AT&T Bell Laboratories
(W)908-946-7262
tpc@mtunm.att.com

A. Lyman Chapin
Bolt, Beranek, and Newman
(W)617-873-3133
Lyman@bbn.com

Richard Cherry
Novell Incorporated
(W)415-947-0998
rcherry@novell.com

J. Noel Chiappa
jnc@ptt.lcs.mit.edu

Chi Chu
SAIC
(W)703-448-0210
chi@sparta.com

Gigi Chu
Hewlett-Packard
(W)415-857-6240
gigic@hpspd.spd.hp.com

George Clapp
Ameritech Services
(W)708-806-8318
meritec!clapp@uunet.uu.net

David Clark
Massachusetts Institute of
Technology
(W)617-253-6003
ddc@lcs.mit.edu

Henry Clark
Ohio Supercomputer Center
(W)614-292-6483
henryc@oar.net

Richard Cogger
Cornell University
(W)607-255-9917
rhx@cornellc.cit.cornell.edu

Richard Colella
National Institute of
Standards and Technology
(W)301-975-3627
colella@osi3.ncsl.nist.gov

Rob Coltun
Consultant
(W)301-240-2168
(W)301-921-8600
rcoltun@ni.umd.edu

James Conklin
EDUCOM
(W)202-872-4200
conklin@bitnic.educom.edu

John Cook
Chipcom Corporation
(W)508-460-8900
cook@chipcom.com

Curtis Cox
NCTSWASH
(W)202-433-4026
ccox@wnyose.nctsw.navy.mil

Tracy Cox
Bell Communications Research
(W)908-758-2107
tacox@sabre.bellcore.com

Stephen Coya
Corporation for National
Research Initiatives
(W)703-620-8990
scoya@NRI.Reston.Va.Us

Caroline Cranfill
BellSouth
(W)404-420-8432
rcc@bss.com

Kenneth Crepea
cisco Systems
(W)407-857-6002
crepea@cisco.com

Stephen Crocker
Trusted Information Systems
(W)301-854-6889
crocker@tis.com

Dave Cullerot
Cabletron Systems, Inc.
(W)603-332-9400
cullerot@ctron.com

John Curran
Bolt, Beranek, and Newman
(W)617-873-4398
jcurran@bbn.com

Nabil Damouny
NEC America, Inc
(W)408-433-2956
nabil@tdd.sj.nec.com

Megan Davies
Corporation for National
Research Initiatives
(W)703-620-8990
mdavies@nri.reston.va.us

James Davin
Massachusetts Institute of
Technology
(W)617-253-6007
jrd@ptt.lcs.mit.edu

Steve Deering
Xerox Corporation
(W)415-494-4839
deering@xerox.com

Peter Deutsch
McGill University
(W)514-398-3709
peterd@cc.mcgillica

Kurt Dobbins
Cabletron Systems
(W)603-332-9400
dobbins@ctron.com

Hardy Doelfel
Siemens - Nixdouf
(W)+49 8963642268
uunet!unido!sinix!hardy

Ralph Droms
Bucknell University
(W)717-524-1145
droms@bucknell.edu

Steve Dusse
RSA Data Security, Inc.
(W)415-595-8782
spock@rsa.com

Tom Easterday
Ohio State University
(W)614-292-4027
tom@cic.net

Jeffrey Edelheit
The MITRE Corporation
(W)703-883-7586
edelheit@mitre.org

Gary Ellis
Hewlett-Packard
(W)415-857-8199
garye@hpspd.spd.hp.com

James Ellis
CERT/CC
(W)412-268-2475
jte@cert.sei.cmu.edu

Robert Elz
University of Melbourne
(W)+6133445225
kre@munnari.oz.au

Alan Emtage
McGill University
(W)514-398-8117
bajan@cc.mcgill.ca

Robert Enger
Contel Federal Systems
(W)703-818-5555
enger@seka.scc.com

Kent England
Bolt, Beranek, and Newman
(W)617-873-3617
kwe@bbn.com

Urs Eppenberger
Switch
(W)+41 12623141
eppen@v

Johnny Eriksson
NORDUnet
(W)+46 8 790 6517
bygg@sunet.se

Mike Erlinger
Micro Technology, Inc.
(W)714-870-2117
(W)800-999-9684
mike@mti.com

Susan Estrada
CERFnet
(W)(619) 534-5181
(W)(800) 876-CERF
Estradas@cerf.net

Fran Fadden
Digital Equipment Corporation
(W)603-881-2115
fran@decvax.dec.com

Erik Fair
Apple Computer, Inc.
(W)408-996-1010
fair@apple.com

Bill Fardy
Cabletron Systems
(W)603-332-9400
fardy@ctron.com

Dino Farinacci
cisco Systems
(W)415-688-4696
dino@cisco.com

Dennis Ferguson
University of Toronto
(W)416-978-3328
dennis@canet.ca

Arlan Finestead
National Center for
Supercomputing Applications
(W)217-244-5553
arlanf@ncsa.uiuc.edu

William Fink
NASA Goddard Space Flight Center
(W)301-286-9423
bill@wizard.gsfc.nasa.gov

Sean Fitzpatrick
BellSouth
(W)205-977-8243

Peter Ford
Los Alamos National Laboratory
(W)505-665-0558
peter@lanl.gov

David Forster
Digital Equipment Corporation
(W)0734 856818
forster@marvin.dec.com

James Forster
cisco Systems
(W)415-688-8245
forster@cisco.com

Jill Foster
University of
Newcastle-Upon-Tyne
(W)+44 912228250
jill.foster@newcastle.ac.uk

Barbara Fraser
CERT/CC
(W)412-268-5010
byf@cert.sei.cmu.edu

Ned Freed
Innosoft International, Inc.
(W)714-624-7907
ned@innosoft.com

Jeff Fried
Proteon, Inc
(W)508-898-2800
jmf@relay.proteon.com

Karen Frisa
Carnegie Mellon University
(W)412-268-8576
karen.frisa@andrew.cmu.edu

Vince Fuller
Stanford University
(W)415-723-6860
vaf@stanford.edu

Deborah Futcher
The Wollongong Group, Inc.
(W)703-847-4500
dfutche@eco.twg.com

Stephen Gabe
Bell Northern Research
(W)613-763-2911
spgabe@bnr.ca

Shari Galitzer
The MITRE Corporation
(W)703-883-6229
shari@gateway.mitre.org

Maria Gallagher
NASA
(W)(415) 694-6141
maria@nsipo.arc.nasa.gov

Shawn Gallagher
Digital Equipment Corporation
gallagher@quiver.enet.dec.com

James Galvin
Trusted Information Systems
(W)301-854-6889
galvin@tis.com

John Garrett
AT&T Bell Laboratories
(W)908-580-4719
jwg@edsel.att.com

Tony Genovese
Lawrence Livermore National
Laboratory
(W)415-423-2471
genovese@es.net

Arlene Getchell
Lawrence Livermore National
Laboratory
(W)415-423-6349
getchell@nersc.gov

Robert Gilligan
Sun Microsystems, Inc.
(W)415-336-1012
gilligan@eng.sun.com

Joe Godsil
National Center for
Supercomputing Applications
(W)217-244-7715
jgodsil@ncsa.uiuc.edu

Kenneth Goodwin
Pittsburgh Supercomputing Center
(W)412-268-4960
goodwin@psc.edu

William Griffeth
UNIX System Laboratories
(W)908-522-6531
billg@usl.com

Robert Griffioen
Bell Northern Research
(W)613-763-3969
rgriff@bnr.ca

Phillip Gross
Advanced Network & Services, Inc.
(W)914-789-5335
(W)914-789-5300
pgross@nis.ans.net

Olafur Gudmundsson
University of Maryland
(W)301-405-2767
ogud@cs.umd.edu

Chris Gunner
Digital Equipment Corporation
(W)+44-734-853516
gunner@osicwg.enet.dec.com

John Dotts Hagan
University of Pennsylvania
(W)215-898-9192
hagan@dccs.upenn.edu

Robert Hagens
University of Wisconsin-Madison
(W)608-262-5776
(W)608-262-1017
hagens@cs.wisc.edu

Jack Hahn
SURAnet
(W)301-982-4600
(W)301-982-3214
hahn@sura.net

Tony Hain
Lawrence Livermore National Laboratory
(W)415-422-4200
alh@es.net

Martyne Hallgren
Cornell University
(W)607-254-8767
martyne@theory.tn.cornell.edu

Joel Halpern
Network Systems Corporation
(W)612-424-1606
jmh@network.com

Alf Hansen
SINTEF DELAB
(W)+47 7 592982
Alf.Hansen@delab.sintef.no

Steve Hardcastle-Kille
University College London
(W)+44 71 380 729
S.Kille@cs.ucl.ac.uk

Susan Hares
Merit Network, Inc.
(W)313-936-2095
skh@merit.edu

Dimitry Haskin
Bolt, Beranek, and Newman
(W)617-873-8609
dhaskin@bbn.com

Phillip Hasse
U.S. Army
(W)602-533-2807
phasse@honchuca-emh8.army.mil

Eugene Hastings
Pittsburgh Supercomputing Center
(W)412-268-4960
hastings@psc.edu

Frank Heath
Rockwell CMC
(W)714-833-6864
heath@cmc.com

Ittai Hershman
Advanced Network & Services, Inc.
(W)914-789-5337
ittai@nis.ans.net

Robert Hinden
Bolt, Beranek, and Newman
(W)617-872-3757
hinden@bbn.com

Russ Hobby
University of California
(W)916-752-0236
rdhobby@ucdavis.edu

Mark Hoerth
Hewlett-Packard
(W)415-857-4867
mark_hoerth@hp0400.desk.hp.com

Ellen Hoffman
Merit Network, Inc.
(W)313-936-3335
esh@merit.edu

J. Paul Holbrook
CICNet, Inc.
(W)313-998-7680
holbrook@cic.net

Jeffrey Honig
Cornell University
(W)607-255-9917
jch@nr-tech.cit.cornell.edu

Alton Hoover
Advanced Network & Services, Inc.
(W)(914) 789-5300
hoover@nis.ans.net

Steven Hotz
Information Sciences Institute
(W)213-822-1511
hotz@isi.edu

Tim Howes
University of Michigan
(W)313-764-2278
Tim.Howes@umich.edu.

Kathleen Huber
Bolt, Beranek, and Newman
(W)617-873-2520
khuber@bbn.com

Erik Huizer
SURFnet
(W)+31 30 310290
huizer@surfnet.nl

Ron Hutchins
Georgia Institute of Technology
(W)404-894-3612
ron@cc.gatech.edu

Lenore Jackson
NASA Science Internet
(W)301-286-7251
jackson@dftnic.gsfc.nasa.gov

Ole Jacobsen
Interop, Inc.
(W)800-INTEROP
(W)415-962-2515
ole@csli.stanford.edu

Ron Jacoby
Silicon Graphics, Inc.
(W)415-335-1371
rj@sgi.com

B.V. Jagadeesh
3Com Corporation
(W)408-764-5169
bvj@3com.com

Mike Janson
Motorola, Inc.
(W)708-576-5403
mjanson@mot.com

P. Allen Jensen
Audiofax, Inc.
(W)404-618-4282
allen@audfax.audiofax.com

Bill Jerome
BellSouth
(W)404-529-3996

James Jones
Digital Equipment Corporation
(W)508-635-9265
jones@regent.enet.dec.com

Ken Jones
Concord Communications, Inc.
(W)508-460-4646
uunet!konkord!ksj

Kevin Jordan
Control Data Corporation
(W)612-482-6835
kej@udev.cdc.com

Phani Jujjavarapu
cisco Systems
(W)415-688-8186
phani@cisco.com

Frank Kastenholz
Clearpoint Research Corporation
(W)508-435-2000
kasten@europa.clearpoint.com

Neil Katin
Sun Microsystems, Inc.
(W)415-336-3718
katin@eng.sun.com

Alan Katz
Information Sciences Institute
(W)213-822-1511
katz@isi.edu

Charles Kaufman
Digital Equipment Corporation
(W)508-486-7329
kaufman@dsmail.enet.dec.com

Manu Kaycee
Digital Equipment Corporation
(W)603-884-5503
kaycee@ctron.com

Stephen Kent
Bolt, Beranek, and Newman
(W)617-873-3988
kent@bbn.com

Mark Kepke
Hewlett-Packard
(W)303-490-2866
mak@hpcndk.cnd.hp.com

Mark Kerestes
AT&T Bell Laboratories
(W)215-439-6046
att!alux2!hawk

Douglas Kerr
LightWork Technologies, Inc.
(W)415-549-8367
dougk@mtxinu.com

Tom Kessler
Sun Microsystems, Inc.
(W)415-336-3145
kessler@sun.com

Kenneth Key
University of Tennessee
(W)615-974-5067
key@cs.utk.edu

Michael Khalandovsky
FTP Software, Inc.
(W)617-246-0900
mlk@ftp.com

Charles Kimber
Dataproducts Corporation
(W)818-887-8009
ckimber@dcp.com

Darren Kinley
RISQnet / CRIM
(W)514-340-5700
kinley@crim.ca

Peter Kirstein
University College London
(W)44-71-380-7286
kirstein@cs.ucl.ac.uk

Anders Klemets
Swedish Institute of
Computer Science
(W)412-268-7694
klemets@cs.cmu.edu

Paulina Knibbe
cisco Systems
(W)415-688-8186
knibbe@cisco.com

Holly Knight
Apple Computer, Inc.
(W)408-974-5110
holly@apple.com

Mark Knopper
Merit Network, Inc.
(W)313-936-3335
mak@merit.edu

Jim Knowles
Sterling Software
Federal Systems
(W)415-604-3056
jknowles@trident.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
(W)617-246-0900
stev@ftp.com

Christopher Kolb
Performance Systems
International, Inc.
(W)703-620-6651
kolb@psi.com

Deidre Kostick
Bell Communications Research
(W)908-758-2642
dck2@sabre.bellcore.com

Bobby Krupczak
Georgia Institute of Technology
rdk@cc.gatech.edu

Cheryl Krupczak
Georgia Institute of Technology
cheryl@cc.gatech.edu

Dale Land
Los Alamos National Laboratory
(W)505-667-9923
(W)(505) 667-3310
land@lanl.gov

Ruth Lang
SRI International
(W)415-859-5608
rlang@nisc.sri.com

Nik Langrind
Shiva Corporation
(W)617-252-6300
nik@shiva.com

Richard Larkin
GTE
(W)508-880-4843
rblarkin@sprintf.merit.edu

Vincent Lau
Sun Microsystems, Inc.
(W)415-336-1069
vincent.lau@eng.sun.com

Anthony Lauck
Digital Equipment Corporation
(W)508-486-7644
(W)508-486-7417
lauck@tl.enet.dec.com

Joseph Lawrence
Bell Communications Research
(W)201-758-4146
jcl@sabre.bellcore.com

Walter Lazear
The MITRE Corporation
(W)703-883-6515
lazear@gateway.mitre.org

Eliot Lear
Intelligenetics, Inc.
(W)415-962-7300
(W)415-962-7302
lear@turbo.bio.net

John Leddy
IBM Corporation
(W)203-783-7212
jleddy@ibm.com

Tim Lee-Thorp
Network General
(W)415-688-2848
uunet!ngc!tim

Debra Legare
Corporation for National
Research Initiatives
dlegare@nri.reston.va.us

Dory Leifer
Merit Network, Inc.
(W)313-936-3335
Dory_Leifer@um.cc.umich.edu

John Lekashman
NASA Ames Research Center
(W)415-604-4359
lekash@nas.nasa.gov

Louis Leon
Emory University
(W)404-727-4782
osll@emuvm1.cc.emory.edu

Mark Leon
NASA
(W)415-604-6498
leon@nsipo.arc.nasa.gov

Brian Lev
NASA Science Internet
(W)301-286-9514
lev@dftnic.gsfc.nasa.gov

Mark Lewis
Telebit Corporation
(W)408-745-3232
mlewis@telebit.com

Tony Li
cisco Systems
(W)415-688-8186
tli@cisco.com

Chao-Yu Liang
Synoptics Communications, Inc.
(W)408-764-1443
cliang@synoptics.com

Peter Liebscher
SURAnet
(W)301-982-4600
plieb@sura.net

John Linn
Digital Equipment Corporation
(W)508-486-6588
linn@zendia.enet.dec.com

David Lippke
University of Texas at Dallas
(W)214-690-2632
lippke@utdallas.edu

Joshua Littlefield
Cayman Systems, Inc.
(W)617-494-1999
josh@cayman.com

Jack Liu
Digital Equipment Corporation
(W)603-881-2590
liu@koala.enet.dec.com

Brian Lloyd
Telebit Corporation
(W)408-745-3103
brian@telebit.com

Daniel Long
Bolt, Beranek, and Newman
(W)617-873-2766
long@nic.near.net

Clifford Lynch
University of California
(W)415-987-0522
calur@uccmvsa

John Lynn
Cornell University
(W)607-255-9917
lynn@ttcllcat.cit.cornell.edu

Shane MacPhillamy
Netrix Corporation
(W)703-742-6000
slm@netrix.com

Ron Mackey
Distributed Systems International, Inc.
(W)708-665-4693
rem@dsiinc.com

Keith Mader
Microcom, Inc.
(W)617-551-1211
keith@python.eng.microcom.com

Carl Malamud
Consultant
(W)303-440-7052
carl@malamud.com

Joseph Malcom
SURAnet
(W)301-982-4600
jmalcom@sura.net

Andrew Malis
Bolt, Beranek, and Newman
(W)617-873-3419
malis@bbn.com

Gary Malkin
FTP Software, Inc.
(W)617-246-0900
gmalkin@ftp.com

Louis Mamakos
University of Maryland
(W)301 405-7836
louie@ni.umd.edu

Allison Mankin
The MITRE Corporation
(W)703-883-7907
mankin@gateway.mitre.org

Bill Manning
Communications Strategy
(W)713-274-2441
bmanning@rice.edu

April Marine
SRI International
(W)415-859-5318
april@nisc.sri.com

Matt Mathis
Pittsburgh Supercomputing Center
(W)412-268-4960
mathis@psc.edu

Phil Mathis
Georgia Institute of Technology
(W)404-853-0254
phil@gatech.edu

Bernie May
AT&T Bell Laboratories
(W)908-949-7506
bfm@houxa.att.com

Keith McCloghrie
Hughes LAN Systems
(W)415-966-7934
kzm@hls.com

Ellen McDermott
Digital Equipment Corporation
(W)617-621-7266
emcd@osf.org

Evan McGinnis
3Com Corporation
(W)408-764-6064
bem@3com.com

Glenn McGregor
Merit Network, Inc.
(W)313-936-3000
ghm@merit.edu

John McGuthry
The MITRE Corporation
(W)703-883-6403
mcguthry@gateway.mitre.org

Karen McKelvey
CERFnet
(W)619-534-5077
karen@cerf.net

Leo McLaughlin
FTP Software, Inc.
(W)415-615-0500
ljm@wco.ftp.com

Donna McMaster
SynOptics Communications, Inc.
(W)408-764-1206
mcmaster@synoptics.com

Milo Medin
NASA
(W)415-604-6440
(W)415-964-9900
medin@nsipo.nasa.gov

Robert Meierhofer
Computer Network Technology
(W)612-550-8256
robert@cnt.mn.org

Bill Melohn
Sun Microsystems, Inc.
(W)415-336-2941
melohn@sun.com

April Merrill
National Security Agency
(W)301-684-7302
abmerri@tycho.ncsc.mil

Cyndi Mills
Bolt, Beranek, and Newman
(W)617-873-4143
cmills@bbn.com

Greg Minshall
Novell, Inc.
(W)415-975-4507
minshall@wc.novell.com

Paul Mockapetris
Defense Advanced Research Projects Agency
(W)(703) 696-2231
pvm@darpa.mil

Lynn Monsanto
Sun Microsystems, Inc.
(W)415-336-2946
monsanto@eng.sun.com

Keith Moore
University of Tennessee
(W)615-974-5067
moore@cs.utk.edu

Dennis Morris
Defense Communications Agency
(W)703-487-3035
morrisd@imo-uvax.dca.mil

John Moy
Proteon, Inc.
(W)508-898-2800
jmoy@proteon.com

Chris Myers
Washington University
(W)314-726-7390
chris@wugate.wustl.edu

Mark Needleman
University of California
(W)415-987-0530
mhn@stubbs.ucop.edu

Clifford Neuman
University of Southern California
(W)206-543-7798
bcn@isi.edu

Oscar Newkerk
Digital Equipment Corporation
(W)206-865-8913
newkerk@decwet.enet.dec.com

Gerard Newman
San Diego Supercomputer Center
(W)619-534-5076
gkn@sds.sdsc.edu

Andy Nicholson
Cray Research, Inc.
(W)612-683-5473
droid@cray.com

Chandy Nilakantan
3Com Corporation
(W)408-764-5213
csn@3com.com

William Nowicki
Legato Systems, Inc.
(W)415-329-7880
nowicki@legato.com

Michael O'Dell
Bell Communications Research
(W)201-829-5128
mo@bellcore.com

Karen O'Donoghue
Naval Surface Warfare Center
(W)703-663-1567
kodonog@relay.nswc.navy.mil

David O'Leary
SURAnet
(W)301-982-3214
oleary@sura.net

Mike Okulski
SSDS
(W)303-798-5520
ssds.com

Philippe Park
Bolt, Beranek, and Newman
(W)617-244-4045
ppark@bbn.com

Bradford Parker
Cayman Systems, Inc.
(W)617-494-1999
brad@cayman.com

Richard Parker
The MITRE Corporation
(W)617-271-3076
rp@mbunix.mitre.org

Michael Patton
Massachusetts Institute of Technology
(W)617-253-6061
map@lcs.mit.edu

Joe Peck
Digital Equipment Corporation
(W)415-688-1341
peck@ms1.pa.dec.com

Geir Pedersen
University of Oslo
(W)+47 2453470
geir.pedersen@use.uio.no

David Perkins
SynOpticsCommunications, Inc.
(W)408-764-1516
dperkins@synoptics.com

Radia Perlman
Digital Equipment Corporation
(W)508-486-7648
perlman@radia.enet.dec.com

Jason Perreault
Racal InterLan
(W)508-263-9929
perreaul@interlan.interlan.com

Marsha Perrott
PREPnet
(W)412-268-7871
mlp@andrew.cmu.edu

Mike Petry
University of Maryland
(W)301-405-2994
petry@ni.umd.edu

James Philippou
Xyplex, Inc.
(W)508-264-9900
japhilippou@eng.xyplex.com

John Pickens
3Com Corporation
(W)415-764-6981
jrp@3com.com

David Piscitello
Bell Communications Research
(W)908-758-2286
dave@sabre.bellcore.com

Mel Pleasant
Rutgers University
(W)201 932-2023
pleasant@hardees.rutgers.edu

Jack Pope
University of San Diego
(W)619-260-4810
pope@uscdsv.acusd.edu

Jon Postel
Information Sciences Institute
(W)213-822-1511
postel@isi.edu

Lars Poulsen
Rockwell/CMC
(W)805-968-4262
lars@cmc.com

Brian Price
BellSouth
(W)404-420-8351
brian@bss.com

Joe Ragland
Microelectronics Center
of North Carolina
(W)919-248-1404
jrr@concert.net

P. Rajaram
Sun Microsystems, Inc.
(W)415-336-5194
rajaram@sun.com

K.K. Ramakrishnan
Digital Equipment Corporation
(W)508-486-7267
rama@kalvi.enet.dec.com

James Rees
University of Michigan
(W)313-763-4403
rees@ifs.umich.edu

James Reeves
SynOptics Communications, Inc.
(W)408-764-1440
jreeves@synoptics.com

Michael Reilly
Digital Equipment Corporation
(W)415-688-1323
(W)415-688-1320
reilly@nsl.dec.com

Yakov Rekhter
IBM Corporation
(W)914-945-3896
yakov@watson.ibm.com

Robert Reschly
Ballistics Research Laboratory
(W)301-278-6808
reschly@brl.mil

Joyce K. Reynolds
Information Sciences Institute
(W)213-822-1511
jkrey@isi.edu

Anil Rijsinghani
Digital Equipment Corporation
(W)508-952-3520
anil@levers.enet.dec.com

Michael Ritter
Apple Computer, Inc.
(W)408-862-8088
mwritter@applelink.apple.com

Michael Roberts
EDUCOM
(W)202-872-4200
(W)(202) 872-4214
roberts@educom.edu

Ron Roberts
Stanford University
(W)415-723-3352
(W)415-723-3909
roberts@jessica.stanford.edu

Kary Robertson
Concord Communications, Inc.
(W)508-460-4646
kr@concord.com

Marshall Rose
Consultant
mrose@dbc.mtview.ca.us

A. Minick Rushton
Space Telescope Science Institute
(W)301-338-5537
rushton@stsci.edu

Gregory Ruth
Bolt, Beranek, and Newman
(W)617-873-3150
gruth@bbn.com

Jan Michael Rynning
Royal Institute of Technology
(W)+46 8 790 6288
jmr@nada.kth.se

Mark Saake
Lawrence Livermore National Laboratory
(W)415-422-3058
saake@llnl.gov

Harri Salminen
FUnet/NORDUnet
(W)+35804572288
hks@funet.fi

Timothy Salo
Minnesota Supercomputer Center
(W)612-626-0347
tjs@msc.edu

Tom Sandoski
Microelectronics Center of North Carolina
(W)919-248-1104
tom@concert.net

Jonathan Saperia
Digital Equipment Corporation
(W)508-496-8333
saperia@tcpjon.enet.dec.com

Mark Schaefer
Distributed Systems International, Inc.
(W)408-720-6804
schaefer@davidsys.com

Gershon Schatzberg
RAD Network Devices
(W)714-891-1964
439-3582@mcimail.com

Jeffrey Schiller
Massachusetts Institute of Technology
(W)617-253-8400
(W)617-253-0161
jis@mit.edu

Aaron Schuman
Silicon Graphics, Inc.
(W)415-335-1901
schuman@sgi.com

John Scudder
Merit Network, Inc.
(W)313-936-3335
jgs@merit.edu

John Seligson
Ultra Network Technologies
(W)408-922-0100
johns@ultra.com

Harvey Shapiro
NCTSW
(W)202-433-4048
shapiro@wnyose.nctsw.navy.mil

Ron Sharp
AT&T Bell Laboratories
(W)201-386-4479
rls@neptune.att.com

Daisy Shen
IBM Corporation
(W)914-945-3632
daisy@watson.ibm.com

Erik Sherk
Cornell University
(W)607-255-1679
sherk@nmc.cit.cornell.edu

Robert Shirey
The MITRE Corporation
(W)703-883-7210
shirey@mitre.org

Keld Simonsen
Danish Uxix Users' Group
(W)+45 33 130023
keld.simonsen@dkuug.dk

Anil Singhal
Frontier Software Development, Inc.
(W)508-851-8872
nsinghal@hawk.ulowell.edu

Dana Sitzler
Merit Network, Inc.
(W)313-936-3335
dds@merit.edu

Mark Sleeper
SPARTA, Inc.
(W)703-448-0210
mws@sparta.com

Sharon Sliz
Corporation for National
Research Initiatives
(W)703-620-8990
ssliz@nri.reston.va.us

Timon Sloane
PEER Networks
(W)408-727-4111
uunet!peernet!timon

Dale Smith
BellSouth
(W)404-420-8432

Patricia Smith
Merit Network, Inc.
(W)313-936-3335
psmith@merit.edu

Richard Smith
Datability Software Systems
(W)212-807-7800
smiddy@pluto.dss.com

Larry Snodgrass
EDUCOM
(W)202-872-4200
snodgras@bitnic.educom.edu

Frank Solensky
Clearpoint Research
Corporation
(W)508-435-2000
solensky@clearpoint.com

Lance Sprung
Standard Microsystems Corp.
(W)516-273-3100
sprung@sys.smc.com

Michael St. Johns
Department of Defense
(W)301-688-6741
stjohns@umd5.umd.edu

Wayne Staats
IBM Corporation
(W)919-859-4326
staats@adm.csc.ncsu.edu

Martha Steenstrup
Bolt, Beranek, and Newman
(W)617-873-3192
msteenst@bbn.com

Mark Stein
Sun Microsystems, Inc.
(W)415-336-2435
marks@eng.sun.com

Bob Stewart
Xyplex, Inc.
(W)508-264-9900
rlstewart@eng.xyplex.com

Bernhard Stockman
NORDUnet
(W)46 8 790 6519
boss@sunet.se

David Storey
BellSouth
(W)205-321-3413

Brad Strand
Cray Research, Inc.
(W)612-68e-5685
bstrand@cray.com

Roxanne Streeter
NASA Science Internet
(W)415-604-4845
streeter@nsipo.nasa.gov

Emil Sturniolo
Datability Software Systems
(W)212-807-7800
emil@dss.com

Subu Subramanian
AT&T Bell Laboratories
(W)908-949-2950
subu@qsun.att.com

Bruce Taber
Racal Interlan
(W)508-263-9929
taber@interlan.com

Osamu Takada
Hitachi, Limited
(W)81 44 966 9111
takada@sdl.hitachi.co.jp

Richard Telljohann
Apple Computer
(W)404-668-2011
telljohann1@applelink.apple.com

Ronald Tencati
STX/NASA
(W)301-286-5223
tencati@nssdca.gsfc.nasa.gov

Michael Tharenos
Interop, Inc.
(W)415-962-2521
tharenos@interop.com

Geoffrey Thompson
SynOptics Communications, Inc.
(W)408-764-1339
thompson@synoptics.com

Dean Throop
Data General Corporation
(W)919-549-8421
throop@dg-rtp.dg.com

Claudio Topolcic
Corporation for National
Research Initiatives
(W)703-620-8990
topolcic@nri.reston.va.us

Glenn Trewitt
Digital Equipment Corporation
(W)415-688-1324
trewitt@nsl.dec.com

Mike Truskowski
cisco Systems
(W)314-275-9174
truskowski@cisco.com

Theodore Tso
Massachusetts Institute of
Technology
(W)617-253-8091
tytso@mit.edu

Paul Tsuchiya
Bell Communications Research
(W)201-829-4484
tsuchiya@thumper.bellcore.com

Maurice Turcotte
Racal-Milgo Information Systems
(W)404-804-5190
dnmrt@rm1.uucp

Gregory Vaudreuil
Corporation for National
Research Initiatives
(W)703-620-8990
gvaudre@nri.reston.va.us

Ross Veach
University of Illinois
(W)217-244-4274
rrv@uiuc.edu

John Veizades
Apple Computer, Inc.
(W)408-974-2672
veizades@apple.com

William Versteeg
Versteeg Codeworks
(W)404-664-2858
bvs@nrc.com

Ruediger Volk
Universitaet Dortmund
rv@informatik.uni-dortmund.de

A. Lee Wade
Sterling Software Federal Systems
(W)415-964-9900
wade@discovery.arc.nasa.gov

David Waitzman
Bolt, Beranek, and Newman
(W)617-873-4323
djw@bbn.com

Steven Waldbusser
Carnegie Mellon University
(W)412-268-6628
waldbusser@andrew.cmu.edu

Jesse Walker
Digital Equipment Corporation
(W)508-486-7326
walker@eider.enet@decpa.dec.com

Philip C. Wang
Network Application Technology
(W)408-370-4300
(W)408-370-4220
uunet:watadn:phil

Drew Wansley
NCR
(W)803-739-7664
dwansley@secola.columbia.ncr.com

David Ward
Chipcom Corporation
(W)508-460-8922
dward@chipcom.com

Chris Waters-Pierandozzi
Princeton University
(W)609-258-1544
waters@jvnc.net

Chris Weider
Merit Network, Inc.
(W)313-936-3335
(W)313-936-2090
clw@merit.edu

Jonathan Wenocur
Shiva Corporation
(W)617-252-6300
jhw@shiva.com

Brian Wheeler
The MITRE Corporation
(W)617-271-2567
wheeler@mbunix.mitre.org

Gerard White
Concord Communications, Inc.
(W)508-460-4646
ger@concord.com

Michael White
BellSouth
(W)404-529-2822
mike@bss.com

Rick Wilder
The MITRE Corporation
(W)703-883-6174
rick@gateway.mitre.org

Preston Wilson
INTERACTIVE Systems Corporation
(W)708-505-9100
preston@i88.isc.com

Walter Wimer
Carnegie Mellon University
(W)412-268-6252
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
(W)708-910-6104
(W)708-972-7236
lwinkler@anl.gov

Cathy Wittbrodt
Lawrence Livermore National Laboratory
(W)415-466-4016
cjw@nersc.gov

John Wobus
Syracuse University
(W)315-443-3631
jmwobus@suvm.acs.syr.edu

Mark Wood
Distributed Systems International, Inc.
(W)708-665-4639
markl@dsiinc.com

Robert Woodburn
SAIC
(W)703-734-9000
(W)703-448-0210
woody@cseic.saic.com

Richard Woundy
IBM Corporation
(W)203-783-4308
rwoundy@ibm.com

L. Michele Wright
Timeplex
(W)213-443-5743
uncng!michele@uunet.uu.net

Russ Wright
Lawrence Berkeley Laboratory
(W)415-486-6965
wright@lbl.gov

Brian Yasaki
The Wollongong Group, Inc.
(W)703-847-4500
bky@eco.twg.com

Peter Yee
NASA Ames Research Center
(W)415-604-3812
yee@ames.arc.nasa.gov

Wengyik Yeong
Performance Systems International, Inc.
(W)703-620-6651
yeongw@psi.com

Jeff Young
Cray Research, Inc.
(W)612-683-5536
jsy@cray.com

Chin Yuan
Pacific Bell
(W)415-867-5016
cxyuan@pacbell.com

Joseph Zur
Fibronics
(W)972-4-313642
fibrontics!zur@uunet.uu.net

Osmund de Souza
AT&T Bell Laboratories
(W)908-949-1393
desouza@osdpc.ho.att.com

Peter de Vries
FTP Software, Inc.
(W)415-615-0500
peter@wco.ftp.com