Proceedings of the Sixteenth Internet Engineering Task Force Florida State University

February 6-9, 1990

Corporation
for
National
Research
Initiatives

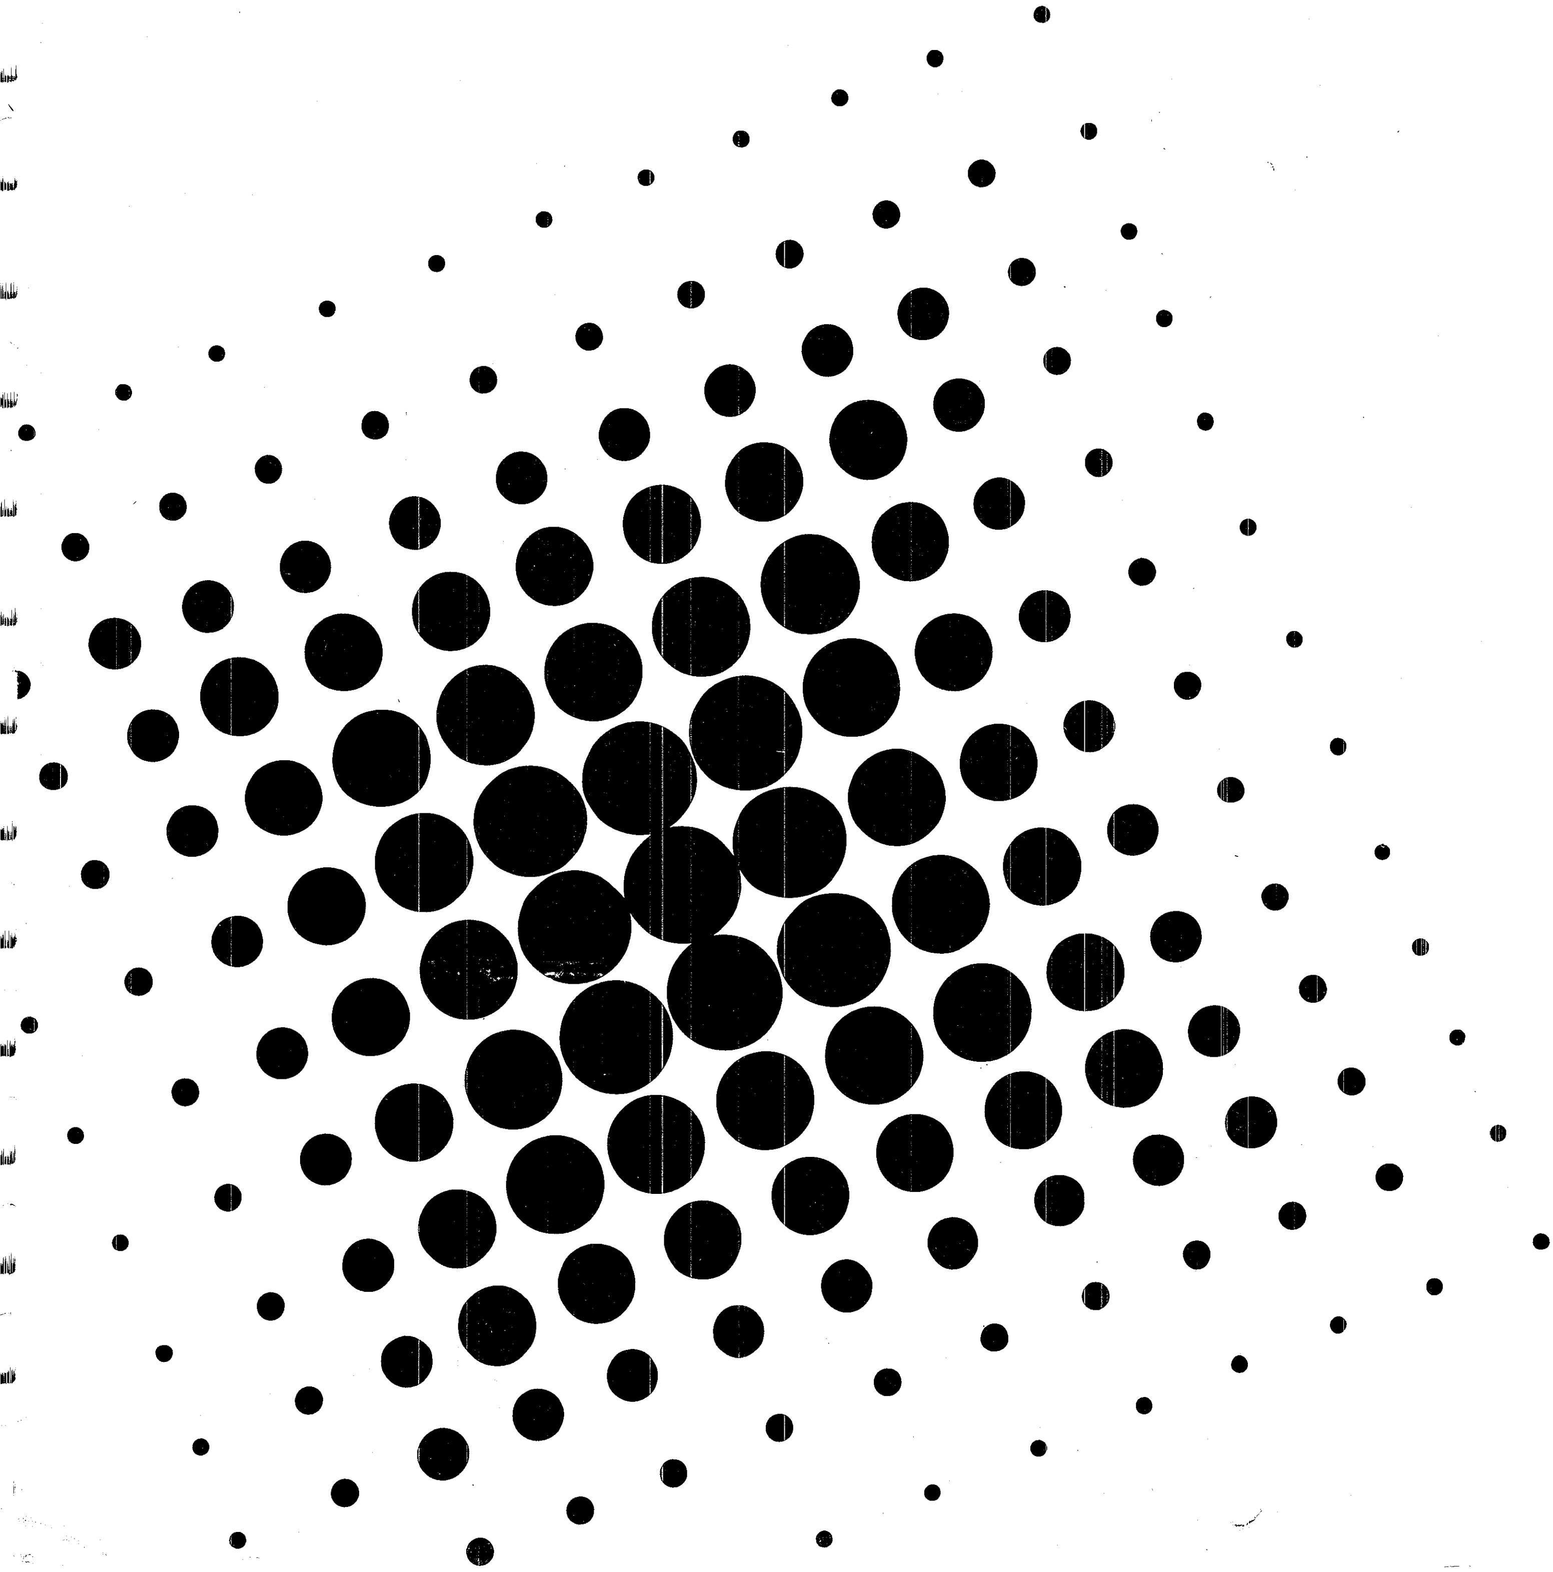

PROCEEDINGS OF THE SIXTEENTH INTERNET ENGINEERING TASK FORCE FLORIDA STATE UNIVERSITY

February 6-9, 1990

Compiled and Edited by
Phillip G. Gross
Gregory M. Vaudreuil

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

ACKNOWLEDGEMENTS

The Sixteenth IETF plenary at Florida State University was a great success. On behalf of the IETF, I would like to extend my thanks to Ken Hayes for the superb arrangements at Florida State University.

I would like acknowledge the technical speakers, John Moy, Radia Perlman, Philip Almquist, Martha Steenstrup, Mats Brunnell, and Brad Johnson for their time and excellent preparation. The presentations were well organized and very interesting. The informative network status reports were given by Chet Birger, Elise Gerich, and Tony Hain.

Thank you Susie Karlson and Valarie Collins for your help in keeping track of the increasingly complex arrangements, working groups, and presentations. I am especially grateful for your help distributing the proceedings of the Hawaii IETF.

Special thanks to Monica Hart, who worked tirelessly to improve the quality of these Proceedings while the content has drastically, increased. This document could not have been prepared without her.

Greg Vaudreuil / NRI

ACKNOWLEDGEMENTS

The Sixteenth IETF plenary at Florida State University was a great success. On behalf of the IETF, I would like to extend my thanks to Ken Hayes for the superb arrangements at Florida State University.

I would like acknowledge the technical speakers, John Moy, Radia Perlman, Philip Almquist, Martha Steenstrup, Mats Brunnell, and Brad Johnson for their time and excellent preparation. The presentations were well organized and very interesting. The informative network status reports were given by Chet Birger, Elise Gerich, and Tony Hain.

Thank you Susie Karlson and Valarie Collins for your help in keeping track of the increasingly complex arrangements, working groups, and presentations. I am especially grateful for your help distributing the proceedings of the Hawaii IETF.

Special thanks to Monica Hart, who worked tirelessly to improve the quality of these Proceedings while the content has drastically increased. This document could not have been prepared without her.

Greg Vaudreuil / NRI

Contents

Chairman's Message

The Sixteenth IETF meeting was held at Florida State University on February 6-9, 1990. The meeting was sponsored by the Super Computations Research Institute (SCRI) and the Department of Energy. The local host was Ken Hayes.

The agenda was very full. Approximately 35 working groups met for a total of 44 separate sessions during the five half-day working group periods. In addition to network status reports and technical presentations, there was a particular focus on Intra-Autonomous Domain routing.

IAB Participation

We were very pleased to have several members of the IAB in attendance at FSU. Not only did Vint Cerf/NRI (IAB Chair) and Dan Lynch/Interop, Inc. attend, but so did two new members of the IAB. Tony Lauck/DEC and Lyman Chapin/Data General (Chair of X3S3.3) were invited to join the IAB in January. As an IAB member, I am gratified to have these two new members on the IAB, and I was especially happy to see the participation of IAB members, new and old, at the IETF meeting.

New Working Groups

Several working groups had their first meetings at an IETF plenary at FSU:

Internet Services

- IP-over-SMDS
- MTU Discovery (also met prior to IETF)
- Router Discovery (also met prior to IETF)
- Router Requirements

Security

- Internet Security Policy

Routing

- IS-IS for Dual IP/OSI routing (also met prior to IETF)
- Multicast routing for OSPF

Host and User Services

- Distributed File System

OSI Integration

- OSI NSAP Guidelines

Applications

- Network Printing Protocol

Operations

- Topology Engineering
- User Connectivity

In other WG news:

The Interconnectivity WG (chaired by Guy Almes) will conquer by dividing. IWG has had two main activities in recent meetings - BGP, and operational routing and topology management. We have decided to create a new WG, Topology Engineering (tewg), to focus specifically on the second issue. Scott Brim (Cornell Theory Center) will chair the new TEWG. TEWG will have a specific goal of coordinating among the various relevant operational routing and topology management groups in the Internet. This includes regional networks, FARNET, national backbones, etc. Guy Almes will continue to chair IWG, which will now take BGP as its main focus. Please see the charters for IWG and TEWG, or contact the chairs for additional information.

There was also a "Birds-of-a-feather" (BOF) session on accounting in the Internet. The purpose of the BOF was to determine if there was enough interest and technical issues to form a WG in this area. Depending on the outcome and proposed focus, such a WG could be organized under the Network Management Area or the Operations Area.

The Joint Monitoring for Adjacent NSFnet Networks WG (JoMANN) has undergone a minor transformation. Sue Hares (Merit) organized JoMANN, at least partly, to assist Merit in interacting with the regional networks attached to NSFnet. JoMANN proved useful enough that we have decided to establish it as a mainstay of the new Operations Area. The WG will be renamed Network Joint Monitoring (NJM) to emphasize that the new focus will be monitoring issues beyond simply networks adjacent to NSFnet.

Several WGs have completed their charter. These WGs are essentially retired, although in some cases the WGs will simply be inactive until further activity develops under their charter. The WGs that have completed their charters include:

- NOC Tools (Internet-Draft complete, to be submitted as RFC)
- Performance and Congestion Control (Internet-Draft complete, to be submitted

as RFC)
- IP Authentication (Internet-Draft complete, to be submitted as RFC)

OSI Integration

The OSI area has changed its name but not its important focus. The original name "OSI Coexistence and Interoperability" was a cumbersome attempt by the IETF chair to capture the charter of the area in the title! The area's main focus was always intended to be the sound planning required for the integration of OSI protocols into the Internet. It was always intended for the OSI to "coexist" with the other protocol families now in the Internet. It was also intended for this area to consider methods for OSI protocols to interoperate with the current TCP/IP protocols. This charter is not new or unique. The DoD developed an OSI Implementation Plan several years ago, which had a similar focus. The Federal Research Internet Coordinating Committee (FRICC) also formed a planning group with a similar focus. The IETF OSI Integration Area hopes to act as a point of focus for technical OSI planning in the Internet. Please contact the co-Directors (Rob Hagens and Ross Callon) for a more information on activities in the OSI Integration Area.

IGP Policy

Perhaps the most important topic at the FSU IETF plenary was the discussions and presentations on Intra-AS routing protocols. As was advertised prior to the meeting, the IESG made the following recommendation to the IAB:

> "There is a pressing need for a high functionality *open* Intra-AS Interior Gateway Protocol (IGP) for the TCP/IP protocol family. Users and network operators have also expressed a strong need for routers from different vendors to interoperate.
>
> Based on these two requirements, the IESG hereby recommends that one high functionality routing protocol be designated as the "recommended" standard IGP for routers in the Internet. Other routing protocols may also be designated as "elective" standards.
>
> By this, it is the intent that all developers of Internet routers make the "recommended" standard IGP available in their products. However, it is not the intent to discourage the use of other routing protocols in situations where there may be sound technical reasons to do so. This recommendation is meant to *enable* multi-vendor router interoperation. It is not otherwise meant to dictate what routing protocol can be used in a private environment.
>
> Therefore, developers of Internet routers are free to implement, and network operators are free to use, other elective Internet standard routing

> protocols, or proprietary non-Internet-standard routing protocols, as they wish."

During the FSU IETF meeting (specifically at the IESG meetings of February 8th and 9th), the IESG discussed the question of choosing one routing protocol to become the "recommended" standard IGP for the TCP/IP protocol family. The two candidates under discussion were ISO's IS-IS, enhanced to support IP in tandem with CLNP, and OSPF. Both protocols use the SPF routing algorithms.

A preliminary recommendation is being forwarded to the IAB and will be announced to the IETF mailing list in early March.

IESG Meetings and Minutes

The IESG held an open meeting on Thursday at the FSU IETF. This has become a standard practice for exchanging information between the IESG and the IETF plenary, and will continue at future meetings. Meeting notes from the open IESG meeting are included in Chapter Two of these Proceedings.

Final Agenda of the Sixteenth IETF

(February 6-9, 1990)

TUESDAY, FEBRUARY 6

9:00 am - 9:15 am TECHNICAL PRESENTATIONS

- Clarification of GOSIP, Phill Gross/NRI

9:15 am - 12:00 pm MORNING WORKING GROUP SESSIONS

- Connection IP (Claudio Topolcic/BBN)
- Interconnectivity (Guy Almes/Rice)
- Distributed File Systems (Peter Honeyman/U-Michigan)
- Router Requirements (Philip Almquist/Stanford, Jim Forster/cisco)
- User Documents (Karen Roubicek/BBN, Tracy LaQuey/U-Texas)
- OSI Internet Management (Lee LaBarre/MITRE)
- Internet Security Policy (Richard Pethia/CERT)

1:00 pm - 4:00 pm AFTERNOON WORKING GROUP SESSIONS

- Connection IP (Claudio Topolcic/BBN)
- Interconnectivity (Guy Almes/Rice)
- User Services (Joyce K. Reynolds/ISI)
- Point to Point Protocol Extentions (Russ Hobby/UC-Davis)
- Multicast Routing for OSPF (John Moy/Proteon)
- OSI NSAP Guidelines (Richard Colella/NIST)

4:15 pm - 5:30 pm TECHNICAL PRESENTATIONS

- OSPF Routing, John Moy/Proteon (30 minutes)
- Open Routing Architecture, Martha Steenstrup/BBN (45 minutes)

5:30 pm - 7:00 pm EVENING WORKING GROUP SESSIONS

- Internet Accounting (Cyndi Mills/BBN)
- Network Joint Management (Gene Hastings/PSC)

WEDNESDAY, FEBRUARY 7

9:00 am - 9:15 am NETWORK STATUS REPORT

- ESnet Report, Tony Hain/DOE

9:15 am - 12:00 pm MORNING WORKING GROUP SESSIONS

- OSI General (Ross Callon/DEC, Rob Hagens/U-Wisc)
- Connection IP (Claudio Topolcic/ BBN)
- Topology Engineering (Scott Brim/Cornell)
- Router Requirements (Philip Almquist/ Stanford, Jim Forster/cisco)
- SNMP Authentication (Jeff Shiller/MIT)
- User Services (Joyce K. Reynolds/ISI)
- Point to Point Protocol Extentions (Russ Hobby/UC-Davis)
- Network Graphics (Craig Partridge/BBN)
- Open Routing (Marianne Lepp/BBN, Martha Steenstrup/BBN)

1:00 pm - 4:00 pm AFTERNOON WORKING GROUP SESSIONS

- Connection IP (Claudio Topolcic/BBN)
- Maximum Transmission Unit Discovery (Jeff Mogul/DEC)
- Router Discovery (Jeff Mogul/DEC)
- Benchmarking Methodology (Scott Bradner/Harvard)
- TELNET (Dave Borman/Cray)
- IP over FDDI (Dave Katz/Merit)
- OSI X.400 (Rob Hagens/U-Wisc)
- Management Services Interface (Oscar Newkerk/DEC)
- Open Routing (Marianne Lepp/BBN)
- End to End User Connectivity (Dan Long/BBN)
- Network Printing Protocol (Leo McLaughlin/Wollongong)

4:15 pm - 5:30 pm TECHNICAL PRESENTATIONS

- Use of OSI IS-IS in IP and Dual Environments Radia Perlman/DEC

6:00 pm - 8:00 pm EVENING WORKING GROUP SESSIONS

- ISIS for IP Internets (Steve Willis/Wellfleet)
- NOC Tools (Bob Enger/Contel, Bob Stine/Sparta)

THURSDAY, FEBRUARY 8

9:00 am - 9:15 am NETWORK STATUS REPORT

- Internet Report, Chet Birger/BBN

9:15 am - 12:00 pm MORNING WORKING GROUP SESSIONS

- Connection IP (Claudio Topolcic/BBN)
- Dynamic Host Configuration (Ralph Droms/Bucknell)
- Transmission MIB (John Cook/Chipcom)
- TCP Large Windows (Craig Partridge/BBN)
- Router Requirements (Philip Almquist/Stanford, Jim Forster/cisco)
- IP over Switched Megabyte Data Service (George Clapp/Ameritech, Mike Fidler/Ohio State U)

1:00 pm - 4:00 pm TECHNICAL PRESENTATIONS

- From Smart Drop to Congestion Control, Martha Steenstrup/BBN (45 minutes)
- NORDUNET, Mats Brunnell/NORDUNET (45 minutes)
- Report of the Open Software Foundation, Brad Johnson/OSF (45 minutes)
- The Interop 89 Network, Philip Almquist/Consultant (45 minutes)

4:30 pm - 7:00 pm OPEN IETF STEERING GROUP MEETING

7:00 pm EVENING WORKING GROUP SESSION

- ISIS Routing IP Implementation Issues (Ross Callon/DEC)
- Alert-MAN (Lou Steinberg/IBM)

FRIDAY, FEBRUARY 9

9:00 am - 9:15 am NETWORK STATUS REPORT

- NSFnet Report, Elise Gerich/Merit (15 minutes)

9:00 am - 11:30 am Working Group Area and Selected Working Group Presentations

- Applications Area (Russ Hobby/UC-Davis)
- Host and User Services Area (Craig Partridge/BBN)
- Internet Services Area (Noel Chiappa/Consultant-Proteon)
- Network Management Area (Dave Crocker/DEC)
- Operations Area (Interim - Phill Gross/NRI)
- OSI Integration Area
  (Ross Callon/DEC and Rob Hagens/U-Wisc)
- Routing Area (Bob Hinden/BBN)
- Security Area (Steve Crocker/TIS)

12:00 pm Adjourn

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) is a large open community of network designers, operators, vendors, and researchers concerned with the smooth operation of the Internet and evolution of the Internet protocol architecture. 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. It has grown into the primary focus for the evolution of the TCP/IP protocol suite and the management of the global Internet.

The IETF mission includes:

- Specifying the short and mid term Internet protocols and architecture for the Internet Activities Board,
- Making recommendations regarding Internet Protocol Standards for IAB approval,
- Identifying and solving 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, agency contractors, and network managers.

The IETF is organized into eight technical areas, each of which is led by a technical area director. Each director has primary responsibility for one area of IETF activity. These eight technical directors with the chair of the IETF compose the Internet Engineering Steering Group (IESG).

The current areas and directors are:

IETF and IESG Chair:	Phill Gross/ NRI
1 Applications:	Russ Hobby/ UC-Davis
2 Host and User Services:	Craig Partridge/ BBN
3 Internet Services:	Noel Chiappa/ Consultant to Proteon
4 Routing:	Robert Hinden/ BBN
5. Network Management:	Dave Crocker/ DEC
6 . OSI Integration:	Rob Hagens/ U-Wisc and
7	Ross Callon/ DEC
8 . Operations:	Phill Gross/ NRI (interim)
9. Security:	Steve Crocker/ TIS

The work of the IETF is conducted in Working Groups, each convened to solve a particular problem and work on an enhancement or exchange information vital to the operation of the Internet. The working groups conduct business via electronic mail on mailing lists established for each group and during plenary meetings of IETF and other meetings. A summary of all current WGs, containing detailed information like charter and mailing list address, is provided in section 1.1.

The IETF holds quarterly plenary sessions composed of working group sessions, technical presentations and network status briefings. Information and logistics about upcoming meetings of the IETF are distributed on the IETF mailing list. To join the list or for inquiries about the IETF, send a request to ietf-request@isi.edu.

1.1 IETF Working Group Summary (by Area)

Applications

Russ Hobby
rdhobby@ucdavis.edu

Network Printing Protocol (npp) WG mail: print-wg@pluto.dss.com Status: Continuing Work	Leo McLaughlin ljm@twg.com
TELNET (telnet) WG mail: telnet-ietf@cray.com Status: Continuing Work	Dave Borman dab@cray.com

Host and User Services

Craig Partridge
craig@bbn.com

Distributed File Systems (dfs)
WG mail: dfs-wg@citi.umich.edu
Status: New Group

Peter Honeyman
honey@citi.umich.edu

Dynamic Host Configuration (dhc)
WG mail: host-conf@sol.bucknell.edu
Status: Continuing, met Nov.

Ralph Droms
droms@sol.bucknell.edu

Internet Draft: "Dynamic Configuration of Internet Hosts",<draft-ietf-dhc-problem-stmt-00.txt and .ps>, Ralph Droms,

Internet User Population (iup)
WG mail: ietf@venera.isi.edu
Status: First meeting Nov.

Craig Partridge
craig@bbn.com

Network Graphics (netgraph)
WG mail: netgraph@nri.reston.va.us
Status: Concluded

Craig Partridge
craig@bbn.com

Network Information Services Infrastructure (nisi)
WG mail: unknown
Status: revived group

Dana Sitzler
dds@merit.edu

TCP Large Windows (tcplw)
WG mail: ietf@venera.isi.edu
Status: New Group

Craig Partridge
craig@bbn.com

User Connectivity (ucp)
WG mail: end2end@nic.near.net
Status: New Group

Dan Long
long@bbn.com

User Documents (userdoc)
WG mail: user-doc@nnsc.nsf.net
Status: Continuing

Karen Roubicek
roubicek@nnsc.nsf.net
Tracy LaQuey
tracy@emx.utexas.edu

User Services (uswg)
WG mail: us-wg@nnsc.nsf.net
Status: Continuing

Joyce Reynolds
jkrey@venera.isi.edu

Internet Services

Noel Chiappa
jnc@lcs.mit.edu

Connection IP (cip)
WG mail: cip@bbn.com
Status: Continuing

Claudio Topolcic
topolcic@bbn.com

IP MTU Discovery (mtudisc)
WG mail: mtudwg@decwrl.dec.com
Status: New Group

Jeff Mogul
mogul@decwrl.dec.com

IP over Appletalk (appleip)
WG mail: apple-ip@apple.com
Status: New Group

John Veizades
veizades@apple.com

IP over FDDI (fddi)
WG mail: FDDI@merit.edu
Status: Continuing

Dave Katz
dkatz@merit.edu

Internet Draft: "The Transmission of IP Datagrams over FDDI Networks",<draft-ietf-fddi-ipdatagrams-00.txt>, Dave Katz, 01/01/1990

IP over Switched Megabit Data Service (smds)
WG mail: smds@nri.reston.va.us

Status: Continuing

George Clapp
meritec!clapp@
bellcore.bellcore.com
Mike Fidler
ts0026@ohstvma.
ircc.ohio-state.edu

Point-to-Point Protocol Extentions (pppext)
WG mail: ietf-ppp@ucdavis.edu
Status: Continuing

Russ Hobby
rdhobby@ucdavis.edu
Stev Knowles
stev@ftp.com

Router Discovery (rdisc)
WG mail: gw-discovery@gregorio.stanford.edu
Status: New Group

Steve Deering
deering@pescadero.stanford.edu

Router Requirements (rreq)
WG mail:
Status: Continuing

Jim Forster
forster@cisco.com
Philip Almquist
almquist@jessica.stanford.edu

Network Management

Dave Crocker
dcrocker@nsl.dec.com

Alert Management (alertman)
WG mail: alert-man@merit.edu
Status: Continuing

Louis Steinberg
louiss@ibm.com

Internet Draft: "Managing Asynchronously Generated Alerts",<draft-ietf-alertman-asyncalertman-01.txt>, Louis Steinberg, 09/01/1989

Internet Accounting (acct)
WG mail:
Status: new group

Cyndi Mills
cmills@bbn.com

LAN Manager (lanman)
WG mail: lanmanwg@spam.istc.sri.com

Status: inactive 3/12/90

Jim Gruel
jimg%hpcndpc@
hplabs.hp.com

Management Services Interface (msi)
WG mail: MSI@nri.reston.va.us
Status: Continuing

Oscar Newkerk
newkerk@decwet.dec.com

Internet Draft: "Management Services Interface",<draft-ietf-msi-api-00.txt and .ps>, Oscar Newkerk, 03/01/1990

NOC-Tools (noctools)
WG mail: noctools@merit.edu
Status: Concluding

Bob Enger
enger@sccgate.scc.com

Internet Draft: "A Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices",<draft-ietf-noctools-debugging-01>, Robert Stine, 11/01/1989

OSI Internet Management (oim)
WG mail: oim@mbunix.mitre.org
Status: Continuing

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

SNMP (snmp)
WG mail: snmp-wg@nisc.nyser.net
Status: Continuing Work

Marshall Rose
mrose@psi.com

Internet Draft: "Management Information Base for Network Management of TCP/IP-based Internets",<draft-ietf-snmp-mib2-01.txt>, Marshall Rose, 09/01/1989

Transmission Mib (transmib)
WG mail:
Status: Continuing

John Cook
cook@chipcom.com

OSI Integration

Ross Callon
callon@erlang.dec.com
Rob Hagens hagens@cs.wisc.edu

Assignment of OSI NSAP addresses (osinsap)
WG mail: ietf-osi-nsap@osi3.ncsl.nist.gov
Status: New Group

Richard Colella
colella@osi3.ncsl.nist.gov

OSI General (osigen)
WG mail: ietf-osi@cs.wisc.edu
Status: Continuing

Rob Hagens
hagens@cs.wisc.edu
Ross Callon
callon@erlang.dec.com

Internet Draft: "An Echo Function for ISO 8473",<draft-ietf-osi-iso8473-00.txt>, Robert Hagens,

OSI-X.400 (osix400)
WG mail: ietf-osi@cs.wisc.edu
Status: Continuing

Rob Hagens
hagens@cs.wisc.edu

Operations

Phill Gross (Interim)
pgross@nri.reston.va.us

Benchmarking Methodology (bmwg)
WG mail: bmwg@harvisr.harvard.edu
Status: New group

Scott Bradner
sob@harvard.harvard.edu

Installation Checklist (check)
WG mail:
Status: proposed group

Martyne Hallgren
martyne@tcgould.tn.cornell.edu
Bob Enger
enger@sccgate.scc.com

Network Joint Management (njm)
WG mail: njm@merit.edu
Status: Continuing

Gene Hastings
hastings@psc.edu

Topology Engineering (tewg)
WG mail: tewg@devvax.tn.cornell.edu
Status: New Group

Scott Brim
swb@devvax.tn.cornell.edu

Routing

Bob Hinden
hinden@bbn.com

ISIS for IP Internets (isis)
WG mail: isis@merit.edu
Status: Continuing

Ross Callon
callon@erlang.dec.com

Internet Draft: "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments",<draft-ietf-isis-spec-00.ps>, Ross Callon, 01/01/1990

Interconnectivity (iwg)
WG mail: iwg@rice.edu
Status: Continuing

Guy Almes
almes@rice.edu

Internet Draft: "A Border Gateway Protocol (BGP)",<draft-ietf-iwg-bgp-00.txt>, K. Lougheed, Y. Rekhter 03/01/1990

Multicast Extentions to OSPF (mospf)
WG mail: ospfigp@trantor.umd.edu

Status: New Group

Steve Deering
deering@
pescadero.stanford.edu

Open Distance Vector IGP (odv)
WG mail: odvigp@rutgers.edu

Status: Inactive

Charles Hedrick
hedrick@
aramis.rutgers.edu

Open Shortest Path First IGP (ospf)
WG mail: ospfigp@trantor.umd.edu
Status: Concluded

Mike Petry
petry@trantor.umd.edu
John Moy
jmoy@proteon.com

Open Systems Routing (orwg)
WG mail: open-rout-interest@bbn.com
Status: Continuing

Martha Steenstrup
msteenst@bbn.com

Internet Draft: "An Architecture for Inter-Domain Policy Routing",<draft-ietf-orwg-architecture-00.txt>, Marianne Lepp, Martha Steenstrup 02/01/1990

Private Data Network Routing (pdnrout)
WG mail: pdn-wg@bbn.com
Status: Continuing Work

CH Rokitansky
roki@isi.edu

Internet Draft: "Assignment/Reservation of Internet Network Numbers for the PDN-Cluster",<draft-ietf-pdn-pdnclusternetassignm-00.txt>, Carl-Herbert Rokitansky, 06/01/1989

Internet Draft: "Application of the Cluster Addressing Scheme to X.25 Public Data Networks and Worldwide Internet Network Reachability Information Exchange",<draft-ietf-pdn-pdncluster-00>, Carl-Herbert Rokitansky, 08/01/1989

Security

Steve Crocker
crocker@tis.com

IP Authentication (ipauth)
WG mail: awg@bitsy.mit.edu
Status: Continuing Work

Jeff Schiller
jis@athena.mit.edu

Internet Draft: "The Authentication of Internet Datagrams",<draft-ietf-auth-ipauthoption-00>, Jeff Schiller, 08/01/1989

Internet Security Policy (isp)
WG mail: isp@nri.reston.va.us
Status: New Group

Richard Pethia
rdp@sei.cmu.edu

SNMP Authentication (snmpauth)
WG mail: awg@bitsy.mit.edu
Status: Reorganized group

Jeff Schiller
jis@athena.mit.edu

Internet Draft: "Authentication and Privacy in the SNMP",<draft-ietf-auth-snmpAuth-00.txt>, J. Galvin, K. McCloghrie J. Davin 01/01/1990

Site Security Policy Handbook (ssphwg)
WG mail: ssphwg@cert.sei.cmu.edu
Status: New Group

Paul Holbrook
ph@sei.cmu.edu
Joyce Reynolds
jkrey@venera.isi.edu

1.2 Future IETF Meeting Sites

Spring 1990

Pittsburgh Supercomputer Center
Host: Gene Hastings
May 1-4, 1990

Summer 1990

University of British Columbia
Host: John Demco
July 31- August 3, 1990

Fall/Winter 1990

National Center for Atmospheric Research (NCAR)
The University of Colorado
Host: Don Morris and Carol Ward
December 4-7 ,1990

Spring 1991

Washington University in St. Louis
Host: Guru Parulkar
March 1991 (tentative)

1.3 On Line IETF Information

The Interent Engineering Task Force maintains up-to-date on-line information on all its activities at NNSC.NSF.NET. On this host, there are two directories containing Internet-Draft documents and IETF working group information. All this information is available for public access.

The "IETF" directory has been created as an aid to both veteran IETF members and newcomers. It contains a general description of the IETF, summaries of ongoing working group activities and provides information on past and upcoming meetings. The directory generally reflects information contained in the most recent IETF Proceedings and Working Group Reports.

The "Internet-Drafts" directory has been installed to make available, for review and comment, draft documents that will be submitted ultimately to the RFC Editor to be considered for publishing as an RFC. 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.

In each directory there is a 00README file.

To access these directories, use FTP to NNSC.NSF.NET. After establishing a connection, Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

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

Individual files can then be retrieved using the GET command:

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

1.4 Guidelines to Authors of Internet Drafts

The Internet Drafts Directory is available to provide authors with the ability to distribute and solicit comments on documents they plan to submit as RFC's. Submissions to the Directory should be sent to **"internet-drafts@nri.reston.va.us"**. Unrevised documents placed in the Internet Drafts Directory 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 Directory with an announcement to that effect for an additional 6 months.

Internet Drafts (I-D's) are generally in the format of an RFC. This format is described in RFC 1111.

Following the practice of the RFCs, submissions are acceptable in postscript format, but we strongly encourage a submission of a matching ascii version (even if figures must be deleted) for readers without postscript printers and for online searches.

There are differences between the RFC and I-D 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 a RFC or a Draft RFC. The Internet Draft should not 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 document should have an abstract section, containing a two to three paragraph description suitable for referencing, archiving, and announcing the document. The abstract should follow the Status of this Memo section. If the 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 RFC editor as a standards document. Distribution of this memo is unlimited. Please send comments to ....
2. This draft document will be submitted to the RFC editor as an informational document. Distribution of this memo is unlimited. Please send comments to ....

If the draft is lengthly, please include on the second page a table of contents to make the document easier to reference.

1.5 Current Internet Drafts

This summary sheet provides a short synopsis of each Internet Draft available within the "Internet-Drafts" Directory at NNSC.NSF.NET.

"Managing Asynchronously Generated Alerts," edited by Louis Steinberg/IBM for the Alert Management Working Group, March 1990 <draft-ietf-alertman-asyncalertman-02.txt>

> This draft defines mechanisms to prevent a remotely managed entity from burdening a manager or network with an unexpected amount of network management information, and to ensure delivery of "important" information. The focus is on controlling the flow of asynchronously generated information, and not how the information is generated. Mechanisms for generating and controlling the generation of asynchronous information may involve protocol specific issues.
>
> There are two understood mechanisms for transferring network management information from a managed entity to a manager; request-response driven polling, and the unsolicited sending of "alerts". Alerts are defined as any management information delivered to a manager that is not the result of a specific query. Advantages and disadvantages exist within each method. This draft discusses these in detail.

"Authentication and Privacy in the SNMP", edited by J. Galvin/TIS, K. McCloghrie/Hughes LAN Systems and J. Davin/MIT for the Authentication Working Group, January 1990 <draft-ietf-auth-snmpAuth-00.txt>

> The Simple Network Management Protocol (SNMP) specification allows for the authentication of network management operations by a variety of authentication algorithms. This memo specifies alternatives to the trivial authentication algorithm defined in the SNMP specification. It also describes an abstract Authentication Service Interface (ASI) by which SNMP-based management applications or agents may in a convenient and uniform way benefit from the algorithms described here and a wide range of others. The terms of the ASI are used to describe three distinct algorithms, including one with support for privacy.

"The Authentication of Internet Datagrams", edited by Jeff Schiller/MIT for the Authentication Working Group, August 1989 <draft-ietf-auth-ipauthoption-00.txt>

> This draft RFC describes a protocol and IP option to allow two communicating Internet hosts to authenticate datagrams that travel from one to the other. This authentication is limited to source, destination IP address pair. It is up to host-based mechanisms to provide authentication between separate processes running on the same IP host. The protocol will provide for "authentication" of the datagram, not concealment from third party observers.

"Dynamic Configuration of Internet Hosts" edited by Ralph Droms/Bucknell, November 1989 <draft-ietf-dhc-problem-stmt-00.txt> and <draft-ietf-dhc-problem-stmt-00.ps>

> The purpose of this document is to lay out and define the dynamic host configuration problem. We expect that this document will provide a basis for the analysis of existing dynamic configuration mechanisms, and that this document might suggest new architectural features that are missing in current dynamic configuration mechanisms. The Dynamic Host Configuration Working Group of the IETF is currently investigating dynamic IP address assignment to hosts (excluding gateways and multi-homed hosts). An accompanying document to propose and define a new dynamic configuration protocol will be written based on this document.

"The Transmission of IP Datagrams over FDDI Networks" edited by Dave Katz for the IP over FDDI Working Group, January 1990 <draft-ietf-fddi-ipdatagrams-00.txt>

> The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams and ARP requests and replies over FDDI networks.

"Use of OSI IS-IS for Routing in TCP/IP and Dual Environments" edited by Ross W. Callon for the IS-IS for IP Internets Working Group, January 1990 <draft-ietf-isis-spec-00.ps>

This internet draft specifies an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments and dual environments. This specification was developed by the IS-IS working group of the Internet Engineering Task Force.

"A Border Gateway Protocol (BGP)" edited by K. Lougheed/cisco and Y. Rekhter/T.J. Watson Research Center, IBM Corp. for the Interconnectivity Working Group, March 1990 <draft-ietf-iwg-bgp-00.txt>

The BGP protocol supports Inter-Autonomous-System routing by providing a means for Autonomous Systems to exchange routing data. The network reachability information exchanged via BGP provides sufficient information to detect routing loops and enforce routing decisions based on performance preference and policy constraints.

"Application of the Border Gateway Protocol in the Internet" edited by J. Honig/Cornell, D. Katz/Merit, M. Mathis/PSC/ Y. Rekhter T.J. Watson Research Center and J. Yu/Merit for the Interconnectivity Working Group, March 1990 <draft-ietf-iwg-bgpapplication-00.txt>

The Border Gateway Protocol (BGP), described in a forthcoming RFC, is an interdomain routing protocol. The network reachability information exchanged via BGP provides sufficient information to detect routing loops and enforce routing decisions based on performance preference and policy constraints as outlined in RFC 1104.

This paper discusses the use of BGP in the Internet environment. Issues such as topology, the interaction between BGP and IGP's, and the enforcement of policy rules with BGP will be presented.

"Management Services Interface" edited by Oscar Newkerk/DEC for the Management Services Interface Working Group, March 1990 <draft-ietf-msi-api-00.txt> and <draft-ietf-mis-api-00.ps>

The Management Services API defines Application Programming Interfaces which provide a set of services for the management of the objects in a heterogeneous, multivendor distributed computing environment.

The Management Services API is designed to allow for the development of portable management applications. The Management Services API insulate management application developers from the details of the management protocol and from the transport services used to route the management directives to the managed objects. It provides facilities to manage both local and remote objects in a seamless fashion.

"Tutorial on OSI Event Management, Alarm Reporting, and Log Control for TCP/IP Networks" edited by Lee LaBarre/MITRE for the OIM Working Group, February 1990, <draft-ietf-oim-eventmanagement-00.txt> and <draft-ietf-oim-eventmanagement-00.ps>

This draft provides a tutorial on OSI mechanisms for event management, alarm reporting, and log control in TCP/IP networks. The mechanisms are based on ISO Draft Proposals and are expected to align with agreements developed by the National Institute of Standards and Technology (NIST) and the Network Management Forum (NMF). Also included is a mechanism for incorporating event flow control as defined in the Internet. It is proposed that systems implementing OSI management protocols for TCP/IP networks [1] should include the mechanisms described in this draft.

"OSI Internet Management: Management Information Base" edited by Lee LaBarre/MITRE for the OIM Working Group, January 1990, <draft-ietf-oim-mib2-00.txt>

This draft defines the management information base (MIB) for use with the OSI network management protocol in TCP/IP based internets. It formats the Management Information Base (MIB-II) in OSI templates and adds variables necessary for use with the OSI management protocol.

"An Architecture for Inter-Domain Policy Routing", edited by M. Lepp/BBN and M. Steenstrup/BBN for the Open Systems Routing Working Group, February 1990 <draft-ietf-orwg-architecture-01.ps>

We present an architecture for policy routing among administrative domains within the Internet. The objective of inter-domain policy routing is to synthesize and maintain routes between source and destination administrative domains, providing user traffic with the requested service within the constraints stipulated by the administrative domains transited. The architecture is designed to accommodate an Internet with tens of thousands of administrative domains.

"An Echo Function for ISO 8473" edited by Robert Hagens for the OSI Working Group, October 1989, <draft-ietf-osi-iso8473-00.txt>

This draft defines an echo function for the connectionless network layer protocol. Two mechanisms are introduced that may be used to implement the echo function. The first mechanism is recommended as an interim solution for the Internet community. The second mechanism will be progressed to the ANSI X353.3 working group for consideration as a work item.

"Internet Cluster Addressing Scheme", by Carl-Herbert Rokitansky/Fern Uni-Hagen, August 1989 <draft-ietf-pdn-clusterscheme-00.txt>

In this document, the new concept of an addressing scheme, similar, but inverse to the subnetting scheme, is proposed, in which a set of Internet networks is associated to an Internet cluster. This "Cluster Addressing Scheme" is of interest especially for wide-area networks, whose structure should be visible to the outside world for (global) routing decisions. In addition, the use of an address-mask (called "Cluster-Mask") for routing decisions within the cluster is discussed.

"Application of the Cluster Addressing Scheme to X.25 Public Data Networks and Worldwide Internet Network Reachability Information Exchange", by Carl-Herbert Rokitansky/Fern Uni-Hagen, August 1989 <draft-ietf-pdn-pdncluster-00.txt>

In this document, the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks is discussed and a new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is proposed.

"Assignment/Reservation of Internet Network Numbers for the PDN-Cluster", by Carl-Herbert Rokitansky/Fern Uni-Hagen, July 1989 <draft-ietf-pdn-pdnclusternetassignm-00.txt>

This document contains a proposal for the reservation of Internet network numbers for the PDN-cluster and the assignment of these PDN-cluster networks to all national X.25 public data networks (DNICs), which are worldwide already in operation.

"Gateway Congestion Control Policies", edited by A.J. Mankin/Mitre and K.K. Ramakrishnan/DEC, July 1989, <draft-ietf-perfcc-gwcc-00.txt>

The task remains for Internet implementors to determine effective mechanisms for controlling gateway congestion. This paper describes the characteristics of one experimental gateway congestion policy, Random Drop, and several that are better-known: Source Quench, Congestion Indication, Selective Feedback Congestion Indication, and Fair Queuing. Random Drop needs further study and does not offer solutions to the resource allocation problems that are the generalization of the congestion control problem. However, a motivation for documenting it now is that it has as primary goals low overhead and suitability for scaling up. Both of these are important goals for future gateway implementations that will have fast links, fast processors, and will have to serve large numbers of interconnected hosts.

"The Point-to-Point Protocol (PPP): A Proposed Standard for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links" edited by Drew Perkins/CMU for the PPP Working Group, March 1990 <draft-ietf-ppp-multidatagrams-00.txt>

The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three parts:

1. A method for encapsulating datagrams over serial links.
2. An extensible Link Control Protocol (LCP).
3. A family of Network Control Protocols (NCP) for establishing and configuring different network-layer protocols.

This document defines the encapsulation scheme, the basic LCP, and an NCP for establishing and configuring the Internet Protocol (IP) (called the IP Control Protocol, IPCP).

The options and facilities used by the LCP and the IPCP are defined in separate documents. Control protocols for configuring and utilizing other network-layer protocols besides IP (e.g., DECNET, OSI) are expected to be developed as needed.

"The Point-to-Point Protocol (PPP) Initial Configuration Options" edited by Drew Perkins/CMU for the PPP Working Group, March 1990 <draft-ietf-ppp-options-02.txt>

The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of

1. a method for encapsulating datagrams over serial links,
2. an extensible Link Control Protocol (LCP), and
3. a family of Network Control Protocols (NCP) for establishing and configuring different network-layer protocols.

The PPP encapsulating scheme, the basic LCP, and an NCP for controlling and establishing the Internet Protocol (IP) (called the IP Control Protocol, IPCP) are defined in The Point-to-Point Protocol (PPP) [1].

This document defines the initial options used by the LCP and IPCP. It also defines a method of line quality monitoring and a simple authentication scheme.

"Implementation Agreements for Transport Service Bridges" by M.T. Rose/Performance Systems International, Inc., February 1990 <draft-ietf-rose-TSBridge-00.txt>

This draft reports implementation experience when building transport service bridges for OSI applications. It does not specify a standard for the Internet community.

"Management Information Base for Network Management of TCP/IP-based internets", edited by M. T. Rose/ NYSERNET for the SNMP Working Group, December 1989 <draft-ietf-snmp-mib2-01.txt>

This memo defines the second version of the Management Information Base (MIB-II) for use with network management protocols in TCP/IP-based internets. In particular, together with its companion memos which describe the structure of management information (RFC 1065) along with the network management protocol (RFC 1098) for TCP/IP-based internets, these documents provide a simple, workable architecture and system for managing TCP/IP-based internets and in particular the Internet community.

"IP Routing Between U.S. Government Agency Backbones and Other Networks" by Scott Brim/Cornell University, January 1990 <draft-fricc-brim-BackboneRouting-01.txt>

This is an overview of how the agency backbones rout IP (Internet Protocol) packets at this time, with any generalizations that can be made and statements of their differences. Also included are recommendations from the agency backbones about how other networks that connect to them can best set up their inter-administration routing.

"OSI Connectionless Transport Services on top of the UDP: Version 1", edited by C. Shue/OSF, W. Haggerty/Wang and K. Dobbins/ Cabletron, November 1989 <draft-osf-shue-osiudp-00.txt>

This draft proposes a method for offering the OSI connectionless transport service (CLTS) in TCP/IP-based Internets by defining a mapping of the CLTS onto the User Datagram Protocol (UDP). If this draft becomes a standard, hosts on the Internet that choose to implement OSI connectionless transport services on top of the UDP would be expected to adopt and implement the methods specified in this draft. UDP port 102 is reserved for hosts which implement this draft. Distribution of this memo is unlimited.

This memo serves as a companion document to RFC 1006 "ISO Transport Service on top of the TCP, Version 3".

"The Knowbot Information Service" by Ralph Droms/Bucknell, December 1989 <draft-nri-droms-kis-00.txt> and <draft-nri-droms-kis-00.ps>

Within the metanetwork of networks that exchange electronic mail, there are many directory services that provide partial coverage of network users; that is, directories with information about some subset of a particular network's user population. Searching the collection of available directories is time-consuming and requires knowledge of each directory's user interface. Although X.500 is currently under study as a basis for an Internet-wide directory service, it is unlikely that a universal user registry will be in place in the near future. The Knowbot Information Service provides a uniform interface to heterogeneous directory services that simplifies the task of locating users in the combined network.

Chapter 2

Steering Group Report

2.1 Minutes of the February 8th Meeting

The Internet Engineering Steering Group makes recommendations to the IAB on Internet standards and policy. It is vitally important for the IESG to have input from the IETF as a whole in formulating these recommendations.

The IESG met in open session (i.e. with full participation of the plenary) on Thursday, February 8th. The main issues were discussion of the new IAB proposed standards process, a discussion of the grandfathered Internet Standards under the new process, and a specific discussion on the recommended policy and direction for the Inter-AS routing protocols (IGP's)

2.1.1 Proposed IAB Standards Process

The following are notes from Phill Gross's presentation of the IAB's new proposed standards process. (See attached slides) It must be emphasized that this is a proposed process that may be amended as a result of IETF input.

There were questions about the differences between required and recommended. It was felt that the distinction was further clouded by the perception that an RFC is a standard. While the participants in the IETF generally recognize RFC's as a document series, it is not so clear in the wider community. There was a desire to have informational RFC's clearly designated.

There is a difference between Experimental and Historical protocols. An experiment is aside from the standards process. There is a difference between a recent experiment and an old effort. Historical efforts are not recommended.

If there are conflicting protocol standards proposals, the IESG will serve the function of coordinating the protocols. Questions arose about the decision to give this responsibility to the IESG and not include the IRTF. It was felt that we need one body to make the standards, and the engineering arm is the more appropriate choice. One hallmark of the IAB is that something must be implemented to be designated an internet standard.

A question was raised whether all documents had to undergo the time-consuming standards process. Some felt the time constraints were a good thing because they give some element of protocol stability. Concern was expressed that if a standard originates in a more formal standards body and then has to go through the IETF process, that there mmay be even greater delay to get IETF approval for use on the Internet.

There was much discussion about the notion of linking a "status" with a "requirement" level. Opposition was expressed to putting a future recommendation into the Draft standards. Some felt that only Full standards should have a status.

After clarifying the point that Proposed and Draft standards have only an "anticipated requirement level", a concern was expressed that the status might be misleading especially if the protocol does not have significant experience and the requirement is subject to change. Others felt that draft and proposed standards should have a requirement, otherwise vendors could not be confident of the stability of the protocol. Vendors are feeling stretched out by the process, with multiple proposed standards, and have concerns about the amount of effort required to implement new protocols, the effect an anticipated requirement designation might have on a Request for Proposal.

IAB Proposed Internet Standards Procedures

Protocol specifications will have two attributes:

- a Status
  - EXPERIMENTAL
  - HISTORICAL
  - PROPOSED standard
  - DRAFT standard
  - FULL standard
- a Requirement Level
  - Not Recommended
  - Elective
  - Recommended
  - Required

A specification in the Proposed Standard or Draft Standard state is said to be in the "standards track".

Protocol Standards Track

```
       |                                 |
       V                                 V
+================+               +================+
|    Proposed    | <------------ |  Experimental  |
+================+               +================+
       |    |                            |
       |     ---------                   |
       V              |                  V
+================+    V          +================+
| Draft Standard | ---+--------> |   Historical   |
+================+    ^          +================+
       |              |
       |     ---------
       V    |
+================+
| Full Standard  |
+================+
```

The requirement level is to be assigned as follows:

Experimental
- Not Recommended

Historical
- Not Recommended

Proposed Standard
- Anticipated RL at full Standard status

Draft Standard
- Anticipated RL at full Standard status

Standard
- Elective, Recommended, or Required

44

Other. Issues:

- A protocol specification can enter Proposed Standard, Draft Standard, or Standard state only with the recommendation of the IESG and the approval of the IAB.
- Draft RFC's containing Internet protocol specifications may be introduced into the standards track by Working Groups of the IETF, by Research Groups of the IRTF, or by outside sources.
- A protocol specification that enters the Proposed Standard state must remain there at least 4 months, and in the Draft stage at least 6 months.
- Raising the Requirement Level on a Proposed, Draft or full Standard will require an additional waiting period to give vendors an opportunity to react and to adjust their planning.

Other Issues ...

- Protocol specifications may be published in Experimental or Historical state at the discretion of the RFC Editor, with appropriate review.

2.1.2 Protocol Standards Review

Old Protocols

Craig Partridge presented to the plenary the list of old protocols submitted to the IESG for review by the IAB for discussion. The following are the protocols discussed and the status to be recommended to the IAB.

The IAB's actions on these recommendations will be announced in future Internet Monthly reports, and in the next RFC on IAB official Standards (eg RFC 1130)

RFC 407 RJE Remote Job Entry

This protocol should be Historical.

RFC 569 NETED Network Standard Text Editor

This protocol should be Historical.

RFC 734 SUPDUP

This protocol should be a Draft Standard. When it becomes a standard, its implementation should be elective.

RFC 742 Finger

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective.

RFC 818 RTELNET Remote Telnet Service

This protocol should be Historical.

RFC 887 RLP Resource Location Protocol

This protocol should be Historical.

RFC 913 SFTP Simple File Transfer Protocol

This protocol should be Historical.

RFC 937 POP2 Post Office Protocol, Version 2

This protocol is Historical and has been obsoleted by POP3. Note we are still deciding on the exact status of POP3, but POP2 is clearly obsolete.

RFC 953 HOSTNAME Protocol

This protocol should be designated Historical. The DNS is now the standard way to get a hostname.

RFC 954 NICNAME WhoIs Protocol

This protocol should be a Draft Standard. When it becomes a standard, its implementation should be elective.

RFC 996 STATSVR Statistics Server

This protocol should be Historical. It was used on Fuzzballs in NSFNET Phase I.

RFC 987, 1026 Mapping between X.400 and RFC-822

This protocol should be Experimental. This advice from the author of the RFCs.

RFC 1057 Sun RPC

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective.

RFC 1058 RIP

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective.

RFC 1037 NFILE

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective. See note for SUN NFS.

RFC 1094 Sun NFS

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective. Note: NFile can not become a standard unless SUN NFS does. It does not make sense to standardize a minor file system when the most popular is not. A working group has been convened to examine distributed file systems.

RFC 1006 ISO Transport on TCP

This protocol should be a Draft Standard. When it becomes a standard, it should be recommended for all systems which run OSI connection oriented applications over TCP/IP.

RFC 1090 X.400 SMTP

This protocol should be Experimental.

Consensus was not reached on the following protocols.

RFC 1056 PCMAIL

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective. See note for POP3

RFC 1081, 1082 POP3 Post Office Protocol Version 3

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective. There was not consensus on either of the two mail protocols. These two protocols have merits but further recommendation should be studied by a working group.

RFC 977 NNTP Network News Transfer Protocol

This protocol should be a Proposed Standard. When it becomes a standard, its implementation should be elective. There are revisions underway, and the revised version is expected to become the standard.

RFC 1045 VMTP

This protocol should be experimental. There is the need for a transaction protocol and it was suggested to form a working group to examine that need.

New Protocols

Dave Crocker presented six documents to be assigned a status upon publication as an RFC. Additionally, the NOC tools Catalogue is expected to become an RFC in 1 month. The recommendations on these documents were:

RFC 1067 SNMP

> This document specifies a standard for the Internet community. It is recommended that TCP/IP implementations in the Internet which are network manageable adopt and implement this specification.

RFC 1065 SMI

> This document specifies a standard for the Internet community. It is recommended that TCP/IP implementations in the Internet which are network manageable adopt and implement this specification.

RFC 1066 MIB I

> This document is a definition of MIB objects and specifies a standard for the Internet community. It is recommended that TCP/IP implementations in the Internet which are network manageable adopt and implement this specification.

Internet Draft MIB II

> This document represents a Proposed Standard for the Internet community. When it becomes a full standard, its Requirement Level will be Recommended. It is expected that implementations supporting MIB I will also support MIB II.

Internet Draft SNMP OSI MIB

> This document is a definition of MIB objects under the experimental branch of the MIB. It does not specify a standard for the Internet community. It is not recommended for operational implementation.

Internet Draft SNMP over OSI

This document defines an experimental usage of SNMP. It does not specify a standard for the Internet community. It is not recommended for operational implementation.

2.1.3 IGP Discussion

Phill Gross presented the framework for the IGP debate.

There are two viable protocols. We are to gather information to assist in making a decision. OSPF is proposed elective. IS-IS is an Internet Draft. Both are on the standards track, the question is the level of requiredness. One standard will be recommended, and one will be Elective. The router requirements document is expected to require the recommended protocol.

Phill Gross presented the following important criteria for making the choice.

- State of protocol development and standardization effort
- Technical merit of protocol and component Algorithm
- Operational experience
- Status of Product Availability, or expected time to market
- OSI integration issues. Efficiencies may arise from converging multiple protocol functions
- Publicly available code (preferably BSD Unix 4.4)
- Provision for authentication

If we are to choose one routing protocol as recommended, it must be implemented and field tested. None of the protocols under consideration meets those requirements. This decision must be made carefully since we do not want to change the decision once made.

After the Phill Gross's introduction, discussion ensued. During the discussion several important points were raised.

There was a strong feeling among the attendees that the protocol chosen must have operational experience. The protocol must be chosen on technical merit, including careful consideration of the component algorithms and performance tradeoffs.

An emphasis on timely deployment was desired by many participants. A delay to wait for implementations and experience for IS-IS was felt by these participants to be excessive.

Many felt that using an OSI protocol for IP carried a lot of extra overhead. The IETF has a dual responsibility for the evolution of TCP-IP and the operational stability of the Internet. In the case of OSI transition, these responsibilities may conflict. The IS-IS protocol itself is optimized for OSI, but just as importantly, the IETF is not empowered to change the IS-IS protocol. The "Dual" portion of IS-IS, however is a

creation of the IETF and can be changed by implementers agreements administered by the IETF.

The integration of OSI protocols proved to be a difficult issue. Many participants felt that a delay in choosing a protocol may be painful in the short run, but would allow for more efficient OSI integration into the Internet. While the desire by the IAB and IESG for fully independent stacks at this stage in the Internet was made apparent, there may be efficiencies to be gained by sharing router resources with a single protocol. OSI IS-IS can be run as a fully independent stack for IS-IS, but the current dual IS-IS shares information between stacks.

After some debate, the IESG conducted a straw poll to gauge sentiments of the plenary on this issue. The options were presented as following:

- Choose the IGP now:
- Delay the decision until at least one IGP has significant field experience:
- Wait for field experience for both protocols:
- Do not decide, ie, let the market decide:
- Abstentions:

The clear preference of the plenary was to delay the decision but to choose as soon as at least one of the IGP's had significant field experience. Based on the experience, this might mean deciding to reject the first IGP and wait for experience with the second or it might lead to immediate adoption of the first.

2.2 The IGP recommendation to the IAB

Following the IETF meeting, the IESG formulated the following recommendation and forwarded it to the IAB:

General Recommendation on Standardizing Routing Protocols

There is a pressing need for a high functionality *open* Intra-AS Interior Gateway Protocol (IGP) for the TCP/IP protocol family. Users and network operators have also expressed a strong need for routers from different vendors to interoperate.

Based on these two requirements, the IESG hereby recommends that one high functionality routing protocol be designated as the "Recommended" Standard IGP for routers in the Internet. Other routing protocols may also be designated as "Elective" standards.

It is the intent that all developers of Internet routers make the "Recommended" standard IGP available in their products. To help ensure that this IGP is available to all users, the IETF Router Requirements Working Group will be directed to indicate in their document that conformant routers must implement the standard IGP.

However, it is not the intent to discourage the use of other routing protocols in situations where there may be sound technical reasons to do so. This recommendation is meant to *enable* multi-vendor router interoperation with a modern high functionality routing protocol. It is not otherwise meant to dictate what routing protocol can be used in a private environment.

Therefore, developers of Internet routers are free to implement, and network operators are free to use, other "Elective" Internet standard routing protocols, or proprietary non-Internet-standard routing protocols, as they wish.

Reference: Please see RFC 1140, "IAB Official Standards" (replaces RFC 1130) for a listing and status of current Internet standards. RFC 1140 also describes the Internet standards process established by the IAB.

Recommendation on Specific Intra-AS Routing Protocols

During the February 6-9 IETF meeting at Florida State University (specifically at the IESG meetings of February 8th and 9th), the IESG discussed the question of standardizing Intra-AS (i.e., IGP) routing protocols for the TCP/IP protocol family.

The two protocols under discussion were the Dual IS-IS and OSPF. Both protocols use the SPF routing algorithms.

OSPF was developed by the IETF OSPF Working Group. The OSPF specification was published as RFC 1131 in October 1989. There is a publicly available implementation for Berkeley Unix, and there is at least one vendor product which is now undergoing deployment in several regional networks.

IS-IS (ISO Draft Proposal 10589) is an OSI proposed protocol for Intra-AS routing. IS-IS products are not widely available, but variations of DP 10589 are being used operationally by at least two vendors.

Dual IS-IS is an enhancement of DP 10589 to support IP in tandem with CLNP. Dual IS-IS is being developed by the IETF IS-IS Working Group. The current specification of the Dual IS-IS is available in the Internet-Draft directories as file DRAFT-IETF-ISIS-SPEC-00.PS. There are plans in progress to develop a publicly available implementation for Berkeley Unix.

The IESG, reflecting the discussion in the IETF plenary at FSU, decided that both protocols need substantial operational experience before either could be made full Internet standards or recommended to the IAB as the "Recommended" IGP for the TCP/IP protocol family.

The practice within the IETF has been to allow a protocol to begin the standards process (i.e., be given the designation "Proposed Standard") prior to gaining field experience, but extensive field experience *is* required prior to advancing to either "Draft Standard" or full Internet Standard.

Therefore, the IESG recommends that OSPF be designated a Proposed Standard at this time. Further review and advancement as an Internet standard will await the outcome of current ongoing field trials.

The IETF and IESG have expressed interest in the integrated routing that is promised by the Dual IS-IS, but also expressed concern about potential complexity and side-effects. Other schemes for running ISO and IP side-by-side have been proposed and demonstrated in practice, and a comparison needs to be undertaken in a systematic manner.

Such issues can only be resolved through extensive field experience. The IESG will re-examine the issue of standardizing Dual IS-IS when the Dual IS-IS specification matures to the point of being published as an RFC and has had some field experience.

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director: Russ Hobby/UC Davis

The Internet has grown to the point where a vast number of people have access and they are now asking "What do we do with it?". Most TCP/IP implementations include three basic applications: remote login (Telnet), file transfer (FTP) and electronic mail (SMTP). These applications need to be looked at to see if they meet todays need, but people want more!

The main reason for TCP/IP's success has been its interoperability. Now that new applications are being looked at (and in some cases developed), we need to provide standards for these applications to insure continued interoperability. In the telephone world, the user does not care what is happening with the switching and circuits, he just wants to be able to talk to the person at the other end. This also needs to be true with network applications.

We already see proprietary network systems, particularly with microcomputer, that can not talk to each other. What we need are agreed upon standards at the network level for the applications, and the vendors can then sell their product because their's is the "best" implementation and user interface. Also, regardless of one's options on OSI, it will happen at some point and TCP/IP needs to work closely with the OSI groups to make sure that there will be interoperability at the application level.

Here are a few of the applications, old and new, that have produced some interest and questions.

Electronic Mail

There is no doubt that current email could use some improvements. How can we include image information in email. What about electronic signatures? Is what we really need an electronic document standard that will include these issues and more? How is X.400 going to fit into or work with the TCP/IP world? What about ANSI Z39.50

Network Printers

An IETF working group for this topic is forming now. We need to define a standard method of sending printer output to a printer connected to the network. Some items to consider are:

1. Authentication/security/accounting
2. Begin/end control of print job
3. Printing modes and options (postscript, plain text, page/line size, ...)
4. Scheduling priorities

Network Backups

Define a standard method of doing disk backups to a mass storage system on the network. This is becoming particularly important with the increase of PCs and workstations that do not have mass storage directly attached.

Distributed Network Bulletin Board System

Define a Bulletin Board System such that various parts of the information base can reside on different computers. This allows each provider of their information to provide the maintenance and computing resources for that part of the information base. Also as the information base grows, rather than having get a bigger computer to handle the growth, you add more computers.

One idea currently being looked at UC Davis is to use the USENET concept and NNTP, but use the Domain Name System to specify which computer provides NNTP service for a particular newsgroup.

Distributed Network Calendar/Scheduling System

Define a system such that one computer can maintain a calendar for a group of people/rooms/items, but can also communicated with calendars on other computers over the network for scheduling.

Network FAX

Define a standard method of sending FAX information over a network. If we can get email to include images, this need may decrease, but people what to do FAX now!

Network Interactive Conversations

Define a standard method for interactive conversations over the network. There are several programs that allow users to talk to each other, but no standards for it. UNIX "talk" or Internet Relay Chat (IRC) are probably the closest to defacto standards.

Network Database

Define a standard method of interacting with databases over a network. SQL seems to one option.

Directory Services

What is the best way to provide this service? Whois? DNS? X.500? We need an official way of doing it over TCP/IP.

3.1.1 Network Printing Working Group (npp)

CHARTER

Chairperson: Leo J. McLaughlin III, ljm@twg.com

Mailing List:

print-wg@pluto.dss.com
print-wg-request@pluto.dss.com

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.

Specific Objectives:

Provide a draft RFC which will describe the LPR protocol as it exists today. Describe printing specific issues on topics currently under discussion within other working groups (e.g. security and dynamic host configuration) and present our concerns to those working groups. Examine printing protocols which exist or are currently under development and assess their applicability to Internet wide use, suggesting changes if necessary.

Milestones:

- Write LPR specification as an RFC : By April 15
- Submit LPR RFC onto the standards track with an as yet to be decided suggested status. :After May Working Group meeting
- Compile draft list of printer issues for other working groups :By end of May IETF meeting
- Submit list of printer issues : By beginning of following IETF
- Decide if existing print protocols are sufficient or if suggested modifications are required. : By end of May IETF meeting
- Submit proposed printing protocol(s) as a draft RFC. :By end of following IETF

CURRENT MEETING REPORT

Reported by Leo McLauglin/The Wollongong Group

MINUTES

Three primary tasks were discussed as objectives for the working group:

1. A specification for LPR will be published as an RFC and submitted for inclusion in the standards track. There was a divergent set of opinions as to what status we should suggest for LPR.
2. A common, full featured network printing protocol was seen as a good, long term goal for the Internet. Paladium, the Project Athena printing protocol, was suggested. Copies of the specification will be made available electronically as soon as possible. We are looking for comments on the Paladium protocol as well as other suitable printing protocols to be presented in Pittsburg.
3. If this long term network printing protocol proves to be unsuitable for use in resource limited environments, an interim protocol (possibly based on LPR) may be necessary. For example, requiring an RPC mechanism was viewed as unacceptable for most terminal server or PC DOS environments. Further information on this topic is again expected to be presented in Pittsburg.

A number of other topics, including security issues and a single 'network' PostScript definition, were discussed and judged to be within the purview of other working groups or committees. We will follow these discussions to present our ideas and concerns.

Administrative Details:

The mailing list of this working group is print-wg@pluto.dss.com. We will be meeting in Pittsburg. A proposed charter will soon be presented for review.

ATTENDEES

Russ Hobby	rdhobby@ucdavis.edu	(Area director)
Leo J. McLaughlin III	ljm@twg.com	(Chair)
Dave Monachello	dave@pluto.dss.com	
John Wobus	jmwobus@suvm.acs.syr.edu	
Don Merritt	don@brl.mil	
Adrianne Glappa	aglappa@tdd.nec.com	
Brian D. Handspicker	bd@vines.dec.com	

3.1.2 TELNET Working Group (telnet)

CHARTER

Chairperson: Dave Borman/Cray, dab@cray.com

Mailing List:

telnet-ietf@cray.com
telnet-ietf-request@cray.com

Description of Working Group:

The TELNET working group is to look at RFC 854, "Telnet Protocol Specification", in light of the last 6 years of technical advancements, and determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the numerous TELNET options, and decide which of them are still germane to current day implementations of the TELNET protocol.

Specific Objectives:

- Reissue 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.
- Disband when the group has met the first two objectives, and re-convening as necessary to addres the third objective.

Estimated Timeframe for Completion:

- ENVIRONMENT Option: May of 1990
- AUTHENTICATION Option: December of 1990
- ENCRYPTION Option: December of 1990
- COMPRESSION Option: (tabled for now)
- Re-write of RFC 854: Early 1991

CURRENT MEETING REPORT

Reported by Dave Borman/Cray Research, Inc.

MINUTES

The TELNET working group meeting got off to a slow start, but gained momentum as the meeting went on. The following are the highlights from the meeting, in the order which they were discussed.

Assigned Numbers:

Joyce Reynolds brought up that a new assigned numbers document will be issued in six to eight weeks. This document contains a list of all the TELNET options, and their current status. This list needs to be updated.

ACTION: Dave Borman will send out the proposed list for comments, update the list as necessary, and forward it to Joyce.

Dan Bernsteins Q Method of Option loop avoidance:

Joyce also said that she and Jon Postel have decided to publish Dan Bernsteins Q method of option negotiation loop avoidance. Since the TELNET working group does not agree with all the technical points of this method, the working group needs to decide if it wants to issue a discussion RFC commenting on the Q method, or whether the group want to just ignore the issue for now, waiting for the revised TELNET spec to comment/clarify about option negotiation loop avoidance. No decision was made.

Re-Issuing the TELNET RFC:

It was decided that the TELNET RFC will need to be updated and re-issued. The reason for this decision was that there are several areas that need to be addressed, among them are: 8 bit NVT support, option negotiation loop avoidance, and DO/WONT vs DO/WONT/DONT option negotiation. The status section will also need to be redone to conform to the current standards for the status section.

Review of proposed options:

DONT-TELNET option:
Since Bill Westfield, the author, was not in attendance, the discussion was tabled until the next meeting.

ENVIRONMENT option:
The option, with revisions agreed upon at the last meeting, was discussed. It was decided that an INFO command, identical to the IS command, was needed. The IS is only sent in response to a SEND command, and an INOF can be sent spontaneously to indicate changes. The INFO is not to be use to indicate initial state; that is what the SEND/IS is for.

ACTION: Dave Borman will write up a new draft for review. It is hoped that by the next meeting it will be ready for RFC submission.

COMPRESSION option:
This option was reviewed in light of the comments from the mailing list. It was decided that: 1) this is a non-trivial option to define. 2) No one in attendance had a burning desire to have this option. Therefor, it was decided that this option will be put at the bottom of the list of things to do, unless someone else is willing to become a champion for this option.

AUTHENTICATION/ENCRYPTION options:
Midway through the meeting, Steve Crocker joined the group. Steve is the Security Area Director for the IETF. Since most of the people at the meeting were not security type people, and Steve is not a TELNET person, we spent some time telling Steve about what we were doing, and he spent some time telling us about security things.

Steve brought up some good points. Since we are not doing any key passing through TELNET, we could just as well do the decision about what type of encryption/authentication is being used out-of-band from TELNET. Then, these options just become a way to turn the stuff on/off, and not a negotiation about what form of encryption/authentication is to be used.

One fear that Steve brought up is that without having people who know about security designing/reviewing the options, there is a good chance that what is designed will not be useful. He also brought up that the privacy enhanced mail group has been thinking about ftp as its next step. Could their work be applied to TELNET also? Should our work be applied to FTP?

From the discussion, it was decided that to really be able to hammer out the solutions, we needed to get the security people and the TELNET people together. Several action items came out of this:

ACTION: Steve Crocker will be scheduling a joint security/TELNET meeting at the next IETF. This meeting will probably also be talking about FTP.

ACTION: Dave Borman will write up a short paper describing the motivation behind wanting the AUTHENTICATION and ENCRYPTION options. This would be something that the security people could look over before the next IETF meeting to help them understand why the TELNET working group is addressing these issues, and what the desired goal is. (Our goal is to avoid having clear-text passwords being sent over the Internet, and to obsolete rlogin.)

ATTENDEES

Dave Borman	dab@cray.com
Steve Crocker	crocker@tis.com
Louis A. Mamakos	louie@trantor.umd.edu
Greg Minshall	minshall@kinetics.com
Joyce Reynolds	jkrey@isi.edu
Keith Sklower	sklower@okeeffe.Berkeley.Edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov

3.2 Host and User Services Area

Director: Craig Partridge/BBN

User Services

The User Services WG has announced that it will begin to produce a new type of RFC, an F.Y.I. note. F.Y.I. notes are informational RFCs, designed to help users and managers better understand how to use the Internet.

The NISI working group has been re-instated. Dana Sitzler of MERIT is chair. The groups charter is to examine the on-line information services offered by the NIC, and consider what service protocols could be standardized, so all NICs could provide similar information using the same protocol.

Host Services

The TCP Big Windows WG has developed two possible ways to expand the TCP window size and sequence space to sizes appropriate for gigabit networks. Researchers at Cray Research and some of the national supercomputer centers have agreed to develop and test these options and report on which one seems more suitable.

The Distributed File Systems WG has started work to identify issues in operating distributed file systems over wide areas. There is reason to believe that existing DFS protocols are not well suited to this problem.

The Graphics WG decided to disband, in favor of trying to arrange a one-time workshop in which people interested in graphics, networking, and standards could discuss common concerns. The WG felt that it was just too difficult to persuade the graphics community to attend networking meetings, or the networking community to attend graphics meeting.

3.2.1 Distributed File Systems Working Group (dfs)

<u>**CHARTER**</u>

Chairperson: Peter Honeyman/honey@citi.umich.edu

Mailing Lists:

dfs-wg@citi.umich.edu
dfs-wg-request@citi.umich.edu

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 WG 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 WG 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.

Estimated Timeframe for Completion:

CURRENT MEETING REPORT

Reported by Peter Honeyman/University of Michigan

SYNOPSIS

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, dont offer credible authentication or authorization services, dont attempt to avoid congestion, etc.

Accordingly, a working group on DFS has been formed by the IETF. The WG 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 WG 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.

dfs-wg@citi.umich.edu is a mailing list for ongoing discussions of the WG; administrative matters, such as requests to be added or dropped from the list, should be addressed to dfs-wg-request@citi.umich.edu, not to the list as a whole.

MINUTES

The meeting was chaired by Peter Honeyman. At the meeting, plans were made to meet the following objectives.

OBJECTIVE: Produce a document for implementors and administrators, in the style of the Hosts Requirements RFCs.

Issues to be addressed include recommendations to be followed when UDP is used as the transport layer. Most of these recommendations come from experiences with TCP. The recommendations include:

- the use of the transport-layer checksum;
- techniques for congestion avoidance;
- techniques for fragmentation avoidance;
- retransmission strategy based on measured round-trip times.

The group intends to identify other recommendations and to flesh out the details in time for a review at the next IETF meeting.

OBJECTIVE: Standard for Kerberos authentication for NFS.

Several groups have deployed or are preparing to deploy NFS integrated with Kerberos. Among these are MIT, U Michigan, and Transarc, Inc.. These groups will describe the protocols they now use for establishing and maintaining Kerberos credentials in an NFS session. The intent is to agree on a common protocol, which will be described in an RFC.

Representatives from MIT, Michigan, and Transarc agreed to describe their protocols in the dfs-wg mailing list. At the next meeting of the IETF, substantive differences between the protocols will be discussed.

OBJECTIVE: Establish the requirements for Internet-friendly DFS protocols.

DFS protocols that were developed for a LAN environment can behave abysmally on a WAN. A well designed DFS will balance its performance needs with those of other users and uses of the network.

Many of the issues concerning the design of DFS protocols depend on one another, or on advances in other areas under study by the IETF. A partial list of the areas in which recommendations can be made includes:

- Naming
- Data representation
- Type management
- Locking
- Impact of design choices:
  - Statelessness
  - Cache management
  - Choice of transport
- Use of MTU discovery
- Authentication and authorization
- Trusted vs. untrusted client
- Time protocol

- User expectations

The first task is to establish concrete goals to guide the WG in this area.

GOALS FOR NEXT IETF MEETING

“Guidelines for DFS Administrators and Implementors” in draft form. Current status of Kerberized NFS implementations on paper. Further discussion on ”Guidelines for DFS Designers.”

ATTENDEES

Richard Basch	probe@mit.edu
Dave Borman	dab@cray.com
Peter Honeyman	honey@citi.umich.edu
Mike Karels	karels@berkeley.edu
Ole Jacobsen	ole@csli.stanford.edu
Dan Jordt	danj@washington.edu
Loius A. Mamakos	louie@trantor.umd.edu
Tony Mason	mason@transarc.com
Matt Mathis	mathis@pele.psc.edu
Leo J. McLaughlin	ljm@twg.com
Greg Minshall	minshall@kinetics.com
Don Morris	morris@ucar.edu
Drew Perkins	ddp@andrew.cmu.edu
Joel Replogle	replogle@ncsa.uiuc.edu
Dean Throop	throop@dg-rtp.dg.com
A. Lee Wade	wade@orion.arc.nasa.gov
Dan Wintringham	ydanw@osc.edu

3.2.2 Dynamic Host Configuration Working Group (dhc)

CHARTER

Chairperson: Ralph Droms/Bucknell, rdroms@nri.reston.va.us

Mailing List: host-conf@rutgers.edu

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 that which cannot (i.e., those that must be managed by network administrators).

Specific Objectives:

1. We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to:
   (a) Exchange Internet packets with other hosts (e.g., discover own Internet address).
   (b) Obtain packet routing information (e.g., discover local gateways).
   (c) Access the Domain Name System (e.g., discover a DNS server).
   (d) Access other local and remote services.
2. We will summarize those mechanisms already in place for managing the information identified by objective 1.
3. We will suggest new mechanisms to manage the information identified by objective 1.
4. 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.

Estimated Timeframe for Completion:

1. Problem statement will be submitted as an RFC.
2. New Protocol document in one year.

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell University

The meeting began with a report from Russ Hobby describing address allocation in PPP. The DHC WG concluded that there's not much overlap between PPP and DHC. Next, Mark Rosenstein reviewed the Gateway Discovery WG meeting. The DHC WG decided that a DHC mechanism could pass an initial gateway without interfering with gateway discovery, with the proviso that any information from a gateway discovery protocol should override any DHC gateway information.

The WG agreed to concentrate on the following configuration information:

- IP address
- Dynamic address allocation bounds
- Subnet mask
- Network broadcast address
- Initial gateway
- Subnet MTU

The WG further agreed to use bootp as a strawman protocol, with extensions for dynamic address allocation. The choice of bootp is intended to maximize compatibility with existing bootp clients and routers.

There are several specific issues that must be resolved if bootp is used as the DHC protocol:

- Is the bootp 576 byte packet limit a problem
- How to select and negotiate dynamically allocated IP addresses
- How should a host proceed when no bootp servers respond
- How can a host detect when its environment has changed and 'cached hints' (e.g., a locally recorded IP address) are no longer valid

The WG agreed *not* to consider:

- Dynamic DNS registration (we refer the problem to the DNS WG)
- Higher level configuration information; e.g., network file servers, boot image name
- Authentication (we refer the problem to the Authentication WG)

What's next:

- The WG would like to have a design complete by 9/1/90
- The WG should develop a detailed strawman protocol based on bootp by next meeting.

ATTENDEES

Basch, Richard	probe@mit.edu
Bowers, Karen L.	kbowers@nri.reston.va.us
Broersma, Ron	ron@nosc.mil
Hobby, Russ	rdhobby@ucdavis.edu
Jacoby, Ronald	rj@sgi.com
Johnson, Brad	bradcj@osf.org
Lauck, Tony	lauck@dsmail.dec.com
Mamakos, Louis A.	louie@trantor.umd.edu
McLaughlin, Leo J.	ljm@twg.com
Minshall, Greg	minshall@kinetics.kinetics.com
Petry, Michael	petry@trantor.umd.edu
Rosenstein, Mark	mar@athena.mit.edu
Roubicek, Karen	roubicek@nnsc.nsf.net
Seaver, Tim	tas@mcnc.org
Sklower, Keith L.	sklower@okeeffe.berkeley.edu
Stahl, Mary	stahl@nisc.sri.com
Throop, Dean	throop@dg-rtp.dg.com
VanBokkelen, James	jbvb@ftp.com
Veizades, John	veizades@apple.com
Yuan, Aileen	aileen@gateway.mitre.org

3.2.3 Internet User Population Working Group (iup)

CHARTER

Chairperson: Craig Partridge/BBN, craig@nnsc.nsf.net

Mailing List: ietf@venera.isi.edu (interim address)

Description of Working Group:

To devise and carry out an experiment to estimate the size of the Internet user population.

Specific Objectives:

We expect to produce two documents: (1) a description of the experimental procedure and (2) an RFC that gives the results of the experiment. We may also produce a short paper for publication in a networking magazine.

Estimated Timeframe for Completion:

The timeframe is being revised by the working group chair.

CURRENT MEETING REPORT

Did not meet.

3.2.4 TCP Large Windows Working Group (tcplw)

CHARTER

Chairperson: Craig Partridge/BBN, craig@bbn.com

Mailing List: ietf@venera.isi.edu

Description of Working Group:

This is a short term, ad hoc, single question working group chartered to make some progress on the various proposals for TCP in long fat pipes.

Specific Objectives:

Choose a proposed standard for the TCP extended window size option.

Estimated Timeframe for Completion:

The timeframe is being reconsidered by the working group chair.

CURRENT MEETING REPORT

Reported by Craig Partridge/BBN

MINUTES

The working group made progress on deciding what to do about the options to support large TCP windows.

Key concerns were making the window size big enough so that a full window could fill a gigabit pipe and expanding the sequence space large enough that it wouldn't be consumed too fast at a gigabit.

The WG made the following tentative decisions about how to deal with big windows:

1. The RFC 1072 window shift option will be used to expand the window size (everyone likes it, and implementation is a snap).
2. We'll create a new Urgent Pointer option so it can point to urgent data anywhere in the expanded window. (To use it, the URG bit is left off – when the receiver processes options, it will see that there is urgent data).
3. The WG was divided among two options for expanded the sequence and ack space. Thankfully some folks with supercomputers offerred to implement both options and report back this spring. The two options were:
   (a) since we need an URG option anyway, steal the urgent field in the TCP header to get 8 more bits for each of the sequence # and the ack #. These would be high order bits.
   (b) don't futz with the header (note that in most cases urgent data is likely to be within 2**16 of the current sequence number and thus the urgent field is still useful). Instead, put an option at the end of each segment header, which contains an additional high order 16-bits of the sequence number and the ack number.

Based on the implementation experience, the WG will decide which path to take.

People interested in participating in the experiment should contact Dave Borman (dab@cray.com).

ATTENDEES

Doug Bagnall	bagnall_d@apollo.hp.com
Dave Borman	dab@cray.com
Richard Fox	sytek!rfox@sun.com
Mike Karels	karels@berkeley.edu
Walt Lazear	lazear@gateway.mitre.org
Paul Love	loveep@sds.sdsc.edu
Gary Malkin	gmalkin@proteon.com
Paul E. McKenney	mckenney@sri.com
Don Merritt	don@brl.mil
Lee Oattes	oattes@utcs.utoronto.ca
Craig Partridge	craig@bbn.com
K.K. Ramakrishnan	rama@dsmial.dec.com
Joel Replogle	replogle@ncsa.uiuc.edu
Frank Solensky	solensky@interlan.com
Mike St. Johns	stjohns@umd5.umd.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Rick Wilder	rick@gateway.mitre.org
Brian Yasaki	bky@twg.com

3.2.5 User Connectivity Problems Working Group (ucp)

CHARTER

Chairperson: Dan Long, long@nic.near.net

Mailing Lists:

ucp@nic.near.net	information
ucp-request@nic.near.net	requests
nic.near.net:mail-archives	archive available

Description of Working Group:

The User Connectivity working group will study the problem of how to solve network users' end-to-end connectivity problems.

Specific Objectives:

1. Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.
2. 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.

Estimated Timeframe for Completion:

Completion by end of 1990

CURRENT MEETING REPORT

Reported by Dan Long/BBN

MINUTES

The User Connectivity Problems WG had it's first meeting at the Florida IETF. The meeting was well attended with representatives from all levels of the Internet hierarchy. A lively discussion ensued about how best to solve end-users' Internet connectivity problems. These are some notes collected from the meeting:

1. The scope of the WG should be broad. Don't restrict UCP focus to certain protocols or geographic areas. In particular, handle more than just IP and more than just USA connectivity.
2. Services offered by UCP system. There wasn't much discussion about the actual methods employed by whatever group/groups do the problem resolution. It was generally agreed, though, that results (both successes and failures) should be reported to provide balanced feedback on performance and to educate others facing the same problems) and that ticket formats and reporting should be standardized.
3. Most user connectivity problems are host problems. Fewer are site problems. Still fewer are regional problems. And fewer still are backbone problems. As a result of this "hierarchy", there should be filters to protect the UCP organization(s) from the users. In particular, users should contact their site rep(s). The problem of educating users to contact their site rep(s) is (hereby :-) delegated to SIGUCCS. [And Martyne Hallgren has agreed to provide a brief writeup on how SIGUCCS work ties in to UCP].
4. Other than the user-¿site-rep constraint, the UCP organzation(s) should be flexible in allowing entry points into the support system. UCP organizations should try to redirect people back to the appropriate level of the network hierarchy (especially when end-users call) but they should be flexible in the case where callers report that the redirect was a dead end.
5. The UCP service provider may not have the expertise to solve all problems but it should have enough knowledge to intelligently refer, track, and adjudicate.
6. The Internet community needs some UCP organization who will "own" user problems. Each person in the community needs to know of (at least) one UCP organization that will own his problem. That is, users should be shielded from the "not my problem" responses.
7. A single national or international 800 service could easily get swamped. If the UCP service is centralized, it needs filters.
8. It would be nice to hire 40 people to man the pohones but infinite funds are not available. This WG needs to consider the economics of proposed solutions.

However, if the only good solution involves hiring 40 people, this group should not shy away from proposing that. (We're not constrained by a budget at this stage.)

9. Does there need to be a "root" of the UCP organization hierarchy? If there is one, it would require filters to pervent being swamped. If there isn't one, need to have adjudication agreements among the top-tier organizations.
10. UCP organizations that can't fix a problem may hand the problem off to the next higher level but the higher level should be empowered to push the problem back down where appropriate.
11. It's very important to document everyone's responsibilities in a cooperative/distributed UCP structure. Perhaps someone (this group?) should write a "Network Troubleshooting Manual".
12. Several UCP org charts were presented and there was much discussion about their respective merits (no pun intended!). The proponents of the various plans were invited to write up their ideas. [Looking for brief writeups from Elise G., Craig P., Karen B., and Guy Almes]

Notes from Steve Goldstein, NSF:

13. Knowledge among UCP organizations needed. That is, each needs to know something about some [needs to be identified] of the others, though probably not *all* of the others. Specifying the needs for knowledge [distributed] among UCP's is not a trivial task.
14. Similarly, need to specify [menu of] communication modalities and "protocols" [procedures] among UCP's. Need for authentication here?
15. Specification of several levels of service/problem-handling abilities among NIC's and NOC's: 3-star, 4-star, 5-star NIC's and NOC's? This way, a prospective customer could choose among offered services– and in a commercial environment, pay accordingly.
16. NIC's and NOC's probably would not map identically among user communities. Some NIC's would be SIG-specific, such as planetary- or seismic-sciences NIC– speculative.
17. Anticipation that NIC services would extend to security and accounting interfaces on behalf of user, as well as operations (NOC) interface. Example: registration on authentication server and/or obtaining privacy keys; checking and verifying billing information. NIC is the standard user contact point, but NIC will have to hand off some of these inquiries to operations, security, accounting, etc.

ATTENDEES

Almes, Guy	almes@rice.edu
Armstrong, Karen	armstrong@sds.sdsc.edu
Aronson, Cathy	cja@merit.edu
Bowers, Karen L.	kbowers@nri.reston.va.us
Brunell, Mats	mats.brunell@sics.se
Easterday, Tom	tom@nisca.ircc.ohio-state.edu
England, Kent	kwe@bu.edu
Feridun, Metin	mferidun@bbn.com
Finkelson, Dale	dmf@westie.unl.edu
Fuller, Vince	fuller@jessica.stanford.edu
Gerich, Elise	epg@merit.edu
Goldstein, Steve	goldstein@note.nsf.gov
Hahn, Jack	hahn@umd5.umd.edu
Hallgren, Martyne M.	martyne@tcgould.tn.cornell.edu
Hastings, Gene	hastings@psc.edu
Jordt, Dan	danj@cac.washington.edu
Love, Paul	loveep@sds.sdsc.edu
Mathis, Matt	mathis@pele.psc.edu
Moore, Berlin	prepnet@andrew.cmu.edu
Morris, Dennis	morrisd@imo-uvax.dca.mil
Morris, Don	morris@ucar.edu
O'Leary, Dave	oleary@umd5.umd.edu
Pace, Donald	pace@fsu2.cc.fsu.edu
Partridge, Craig	craig@nnsc.nsf.net
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Roubicek, Karen	roubicek@nnsc.nsf.net
Sheridan, Jim	jsherida@ibm.com
Stahl, Mary	stahl@nisc.sri.com
Streeter, Roxanne	streeter@nsipo.arc.nasa.gov
Ward, Carol	cward@spot.colorado.edu
Wintringham, Dan	danw@igloo.osc.edu
Yuan, Aileen	aileen@gateway.mitre.org

3.2.6 User Documentation Working Group (userdoc)

CHARTER

Chairpersons:	Tracy LaQuey/University of Texas	tracy@emx.utexas.edu
	Karen Roubicek/BBN	roubicek@nnsc.nsf.net

Mailing List: user-doc@nnsc.nsf.net

Description of Working Group:

The USER-DOC Working Group will prepare a bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "how to use the Internet". (Target audience: those individuals who provide services to end users and end users themselves.)

Specific Objectives:

1. Identify and categorize useful documents/reference materials/training tools.
2. Publish both an on-line and hard copy of this bibliography.
3. Develop and implement procedures to maintain and update the bibliography. Identify an organization or individuals to accept responsibility for this effort.
4. As a part of the update process, identify new materials for inclusion into the active bibliography.
5. Set up procedures for periodic review of the biblio by USWG.

Estimated Timeframe for Completion:

1. Format for the bibliography will be decided upon by the July IETF session, as well as identification of "sources of information" (e.g., individuals, mailing lists, bulletins, etc.)
2. Draft bibliography will be prepared by mid-December 89.
3. Draft to be reviewed at February IETF, and installed in the Internet-Drafts Directory by March 21, 1990.
4. Bibliography submitted as a FYI RFC at the May IETF.

CURRENT MEETING REPORT

Reported by Karen Roubicek/BBN

AGENDA

1. Review Charter, Objectives, Goals.
2. Review current version of bibliography, assign remaining tasks.
3. Discuss remaining issues (appropriateness of categories, keywords, RFC format)
4. Determine publication schedule.
5. Discuss maintenance of bibliography, update procedures
6. Discuss distribution.
7. Determine action items for next meeting, future projects.

MINUTES

The meeting began with a review of the charter of the USER-DOC Working Group, and a check to see if the group was meeting its goals. A draft of the bibliography was successfully generated in December, thanks particularly to the efforts of the "Editorial Board". The Board met twice since the last IETF using the videoteleconferencing facilities at SRI, ISI, DARPA and BBN.

The current draft of the bibliography is significantly more complete than the previous version. The group went through the individual entries and made comments. There is still a small amount of additional research as well as some additions and corrections that are required before the bibliography is complete. These specific assignments will be done by members of the Editorial Board. The question of multiple repositories for RFCs was discussed. Although there are collections of RFCs available on several hosts in the Internet, the group decided that in the interest of centralizing access and ensuring that readers would be able to find all the RFCs mentioned in the "basic beige" list, only the main collection at SRI would be listed in the RFC section (Section 12). However, the Introduction to the bibliography will be revised to state that RFCs are available elsewhere (e.g., Merit, CSNET)

In discussing which items should be included or excluded from certain sections, the group decided that some form of disclaimer should be put into the Introduction. The concern is that an organization or author might complain if their entry were not included and the group wanted to state up front that the bibliography is not all-inclusive but rather a sampling of what resources are available. Karen Armstrong will write the disclaimer. On a similar topic, Jack Hahn expressed concern that even though the bibliography covers selected "basic materials", there is still a need for a short list, of perhaps 4 or 5 entries, for those users who don't have the time to go through the entire bibliography. Jack Hahn will provide Karen Bowers with a suggested list and she will change the Introduction accordingly.

All the attendees concurred on the value and need for an online, searchable version of the bibliography. However, the more pressing need is to make the document available as soon as possible, so the group agreed that a refer and formatted version will be available first, while a small group (Lee Oattes and Mary Stahl) will pursue the options for an online searchable form.

The current draft of the bibliography contains keywords, but there were strong feelings that those keywords are not effective and should not be included in this hardcopy version. Mary Stahl and Berlin Moore, as members of the Keyword Committee, will spend some time looking for the correct set of keywords.

For the correct formatting of the bibliography for submission as an RFC in the FYI series, Joyce Reynolds referred the Editorial Committee to RFC 1111, Request for Comments on Request for Comments.

Some organizational issues were raised, including the need to break up the Introduction for clarity's sake. Chapter 10, Online Files, was viewed as misleading, because several of the other documents listed throughout the bibliography are available online. This chapter was intended to refer to documents that were exclusively online. As an alternative, the items listed in Chapter 10 will be merged into the other chapters according to category and a separate appendix will list all online documents.

The publication schedule is as follows:

- March 7 - research/writing assignments due
- March 21 - Formatted version of bibliography to be put into Internet Drafts Directory
- March-May IETF - Comment period
- May IETF - Submission as formal FYI RFC

Mary Stahl, Tracy LaQuey and Karen Roubicek will be responsible for maintaining the bibliography. As was decided at the last IETF meeting, the bibliography will be updated annually.

In addition to the bibliography's availability online in the RFC repositories (at the NIC and in the shadow directories elsewhere on the Internet), the group brought up the option of additional hardcopy distribution. Mary Stahl will pursue the possibility of SRI providing and distributing paper copies and will let the group know what price the NIC would charge.

Final points in the discussion led into topics to be taken up in the User Services Working Group, including an RFC about distribution procedures for documents and information, standards for document formatting and retrieval, and the need for other networking documents (e.g. comprehensive glossary of networking terms) in the future.

ATTENDEES

Armstrong, Karen	armstrong@sds.sdsc.edu
Bowers, Karen L.	kbowers@nri.reston.va.us
Enger, Robert M.	enger@sccgate.scc.com
Hahn, Jack	hahn@umd5.umd.edu
Jacobsen, Ole	ole@csli.stanford.edu
Moore, Berlin	prepnet@andrew.cmu.edu
Oattes, Lee	oattes@utcs.utoronto.ca
Reynolds, Joyce K.	jkrey@venera.isi.edu
Roubicek, Karen	roubicek@nnsc.nsf.net
Smith, Pat	psmith@merit.edu
Stahl, Mary	stahl@nisc.sri.com
Sturtevant, Allen	sturtevant@ccc.nmfecc.gov
Wobus, John	jmwobus@suvm.acs.syr.edu
Yuan, Aileen	aileen@gateway.mitre.org

3.2.7 User Services Working Group (uswg)

CHARTER

Chairperson: Joyce K. Reynolds, jkrey@venera.isi.edu

Mailing Lists:

us-wg@nnsc.nsf.net
us-wg-request@nnsc.nsf.net

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 WGs created to perform certain projects, or outside organizations such as SIGUCCS.

Specific Objectives:

1. Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should
   - clearly address user assistance needs;
   - produce an end-result (e.g. a document, a program plan, etc);
   - have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
   - and not duplicate existing or previous efforts.
2. Create WGs or other focus groups to carry out projects deemed worthy of pursuing.
3. Provide a forum in which user services providers can discuss and identify common concerns.

Estimated Timeframe for Completion:

This is an operational and liason WG and, as such, has an indefinite lifetime.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI

MINUTES

Announcements

Steve Crocker (TIS), IETF Security Area Director and Rich Pethia (CERT) briefly presented to the USWG new plans for a Security WG, chaired by Rich. Steve and Rich discussed possible activities to be pursued jointly by the USWG and the new Security WG. The USWG will work with Steve and Rich in whatever endeavors they feel we could make contributions.

The NISI WG has been reinstated per Craig Partridge, our Area Director. The new Chair will be Dana Sitzler, MERIT.

RFC: F.Y.I. Series of Notes

The USWG progressed further on this topic from the last IETF meeting in Hawaii. There are five points that are now considered "well known" and define the initial plan for the information series:

- F.Y.I. is part of the RFC series of notes.
- Expansion of the "Status of this Memo" section will further depict the inclusion of F.Y.I. in the RFC scheme.
- Documents for F.Y.I. consideration will be submitted to USWG for review and approval via the RFC Editor.
- RFC: F.Y.I. is a NEW process, starting in 1990.
- Previously published "informational" RFCs will be gathered up in an F.Y.I. RFC that points to the old RFCs, "capturing" the history of informational RFCs prior to 1990.

Repository for F.Y.I. Memos will be in parallel with the current repository for RFCs. The F.Y.I. Memos will have their own separate index. This will be organized further with the RFC Editor.

The premier of the F.Y.I. Memos will be an RFC: F.Y.I. entitled, "F.Y.I. on F.Y.I." in the same spirit of RFC 1111, "Request for Comments on Request for Comments".

Discussion of F.Y.I. topics:

- F.Y.I. RFC
- NOCTools Catalog
- User-Doc Bibliography
- Q/A results (quarterly publication) (aka commonly asked questions and answers)
- Other concepts/ideas from USWGers:
  - How to Connect to the Internet (general procedures)
  - Connection Checklist (technical considerations)
  - Network Information and Information Services Currently Available

ACM Conference Paper Participation

Deadline for paper submission: 5-March-90. Contributors include: NNSC - Internet Resource Guide, ISI - Request for Comments Documents, ISI/NNSC - Internet Monthly Report and CNRI - Video Teleconferencing.

A tentative working group has been proposed to our Area Director - DAWG: Distribution and Announcement Working Group, to be chaired by Robert Enger and assisted by Karen Bowers. Its intent:

- recommend procedures on how to get information out.
- use existing methods and begin to address long term distribution methods.
- Note: this will be a follow-on effort of the Distribution and Announcement session facilitated by Martyne Hallgren at the Stanford IETF.

Current avenues to "get the word out" should include:

- The Internet Numbers and Assigned Numbers Process and RFCs
- "Hello" New User Text Servers
- Vendor Packages to Be Distributed with Hardware/Software

The concept of designing "Intro Packages" for new Internet users was discussed at great length. The purpose of publishing "Intro Packages" is to determine what information new users need immediately when they begin their connection to or begin to use the Internet.

Research will be done and reported at the next IETF on:

- what the information is going to be
- what already exists
- what needs to be written

Mary Stahl, Pat Smith, and Karen Roubicek volunteered to research and report.

Q/A List

The Q/A mailing lists/box will be set up at Proteon and maintained by Gary Malkin.

They will include:

quail: This is a discussion mailing list. Its primary use is for pre-release (to the USWG) review of the Q/A FYIs.

quail-request: This is how you join the quail mailing list. The USWG members will not be added by default (except those which have already volunteered to monitor mailing lists).

quail-box: This is where the questions (and hopefully answers) will be forwarded-and-stored. It is not necessary to be on the quail mailing list to forward to the quail-box.

Electronic mailing lists will be monitored for pertinent questions. James Van Bokkelen, Berlin Moore, Gary Malkin, and Joyce Reynolds voluteered to monitor particular lists. Any USWGer who routinely reads public mailing lists should keep an eye out for the types of questions relevant to the Q/A list, and is encouraged to contribute their findings.

We are specifically looking for commonly asked, or significant questions. Gary Malkin will send a message to the USWG mailing list with further details, once the mailing lists are set up.

It is projected that approximate every three months, an F.Y.I. RFC will be published with the Q/A results.

User-Doc Bibliography Update

Karen Roubicek briefly presented the User-Doc's current work and the status of their Bibliography.

ATTENDEES

Armstrong, Karen	armstrong@sds.sdsc.edu
Bennett, Derek	xndmis14@servax.bitnet
Bowers, Karen L.	kbowers@nri.reston.va.us
Crocker, Steve	crocker@tis.com
Enger, Robert M.	enger@sccgate.scc.com
Fidler, Mike	ts0026@ohstvma.ircc.ohio-state.edu
Finkelson, Dale	dmf@westie.unl.edu
Hallgren, Martyne	martyne@tcgould.tn.cornell.edu
Jacobsen, Ole	ole@csli.stanford.edu
Malkin, Gary	gmalkin@proteon.com
Moore, Berlin	prepnet@andrew.cmu.edu
Morris, Don	morris@ucar.edu
Oattes, Lee	oattes@utcs.utoronto.ca
Pace, Donald	pace@fsu1.cc.fsu.edu
Perkins, Dave	dave_perkins@3com.com
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Robertson, Jim	jar@bridge2.3com.com
Roubicek, Karen	roubicek@nnsc.nsf.net
Sheridan, Jim	jsherida@ibm.com
Shue, Chi	chi@osf.org
Smith, Pat	psmith@merit.edu
Stahl, Mary	stahl@nisc.sri.com
Stine, Bob	stine@sparta.com
Sturtevant, Allen	sturtevant@ccc.nmfecc.gov
Van Bokkelen, James	jbvb@ftp.com
Van Wyk, Ken	krvw@sei.cmu.edu
Ward, Carol	cward@spot.colorado.edu
Wobus, John	jmwobus@suvm.acs.syr.edu
Woodburn, Robert	woody@saic.com
Yuan, Aileen	aileen@gateway.mitre.org
Zimmerman, David Paul	dpz@convex.com

3.3 Internet Services Area

Director: Noel Chiappa/Consultant

Most of the activity to report in this area happened at the IETF meeting in Florida.

The Connection-Oriented IP WG heard a number of presentations about the new ATM (cell switching) networks, and coordinated their proposals in response to the BAA (roughly, an RFP) from DARPA in this area.

The Router Requirements WG held its organizational meeting, which reviewed the draft outline put together by the co-chairs, and solicited and assigned writing assignments to have the first drafts of the various sections completed. Some discussion of a few technical topics was also held.

The PPP Extensions WG dealt with the final issues in the Line Quality option area, and were ready to produce the draft of that RFC shortly.

The MTU Discovery WG, after hearing arguments from Van Jacobsen and others, decided to discard the initial approach that was agreed on at the January meeting (which used a mix of fragmentation detection and MTU probing) and revert to a mechanism based on fragmentation detection.

The IP over FDDI WG has a draft RFC, and discussed at some length some of the remaining technical issues, including handling of dual MAC stations. No agreement could be reached on this topic, and it was agreed to put the document out "as is", to get something out, leaving the issue of dual MAC's to be addressed later.

The IP over SMDS WG met for an organizational meeting, and after some introductions to SMDS managed to come up with a list of areas that needed to be looked at.

Finally, a new WG, IP over AppleTalk, has been organized. It will initially concentrate on formalizing the protocols pioneered by Kinetics, and now used by others, to interoperate IP and Appletalk networks.

3.3.1 Connection IP Working Group (cip)

CHARTER

Chairperson: Claudio Topolcic/BBN, topolcic@bbn.com

Mailing List: cip@bbn.com

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:

1. Produce a new specification of ST: April 1990
2. Define common hardware and software platform: May 1990
3. Implement hardware and software platform: October 1990
4. Implement experimental modules and perform experiments: May 1991
5. Produce a specification of a next generation connection oriented protocol: May 1992

CURRENT MEETING REPORT

Reported by Claudio Topolcic/BBN

MINUTES

The CO-IP Working Group met at the February 6-9 IETF Meeting at Florida State University. The Tuesday sessions were a presentation and discussion on ATM networks by Guru Parulkar of Washington University. The Wednesday morning session was a discussion of the issues, questions, and experiments raised by a guaranteed service network. The Wednesday afternoon session was canceled so that the working group members could attend the SMDS working group meeting. Work on the ST-2 protocols specification was dropped due to insufficient time.

The ATM presentation consisted of roughly three parts: a BISDN perspective, ATM networks, and SMDS. The Broadband Integrated Services Digital Network is intended to support voice and video, DS1/3, X.25, and Switched Multi-megabit Data Services, with the latter including connectionless 802.6. The video services would include broadcast/permanent connections (e.g., cable TV), point-to-point connections, and switched connections, including conferencing. Video rates include both NTSC (45 Mbit) and ATV (135 Mbit). It is predicted that after 1995 it will be cheaper to run fiber to a home than to run copper. The protocol stack consists of application supported by an adaptation layer (which would include segmentation and reassembly if required) over the ATM layer over a SONET physical layer. SONET STS-3c consists of a 155.520 Mbit channel divided into 125 microsecond frames, with each frame containing 90 bytes of overhead and nine "channels" of 260 bytes each (the channels are not all byte aligned).

An ATM frame consists of a 5-byte header followed by 48 bytes of data. The header format isn't yet standardized, but would most likely consist of 28 bits of combined Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI), an 8-bit checksum, a 2-bit priority field, and a 2-bit type field. A Virtual Path may inlude a number of Virtual Channels switched as a unit, so either the VPI or the VCI is used for cell forwarding on a hop-by-hop basis. The boundary between the VPI and VCI might vary at different interfaces. The VPI/VCI field might include other logical subfields, e.g., flow control information, etc.

The Adaptation layer consists of a convergence sublayer on top of a segmentation and reassembly sublayer. The convergence sublayer wraps the padded application data with a header and trailer; the segmentation and reassembly sublayer segments the wrapped application data and adds its own header and trailer before passing each segment to the ATM layer.

The services provided by ATM include point-to-point, multicast, and dynamic multicast callees, a QOS (which would probably be a fixed delay and loss specification within a homogeneous network), naked (aka dark) cells which will not be reordered by the network, and a bandwidth requirement specification. Bandwidth would be specified in terms of mean, peak, and burst characteristics, with the actual nature of the latter still unspecified. Bandwidth consumption would be limited to that requested.

With respect to CO-IP, there are two basic assumptions: the Internet will be heterogeneous for some time, and that LANs will not be ATM networks in the forseeable future. The conclusions were: since applications may generate packets larger than cell size, transparent fragmentation and reassembly should be supperted, CO-IP parameters should be consistent with ATM (at least in the voice context where the packet size is small) and CO-IP should try to be consistent with ATM naked cells (to minimize as much as possible the adaptation layer), the working group should make concrete plans for CO-IP experiments across ATM based high-speed networks, and to identify work that has/is being done in the ATM community for use in the CO-IP subnet dependent layer.

Wednesday morning's session consisted of a discussion of CO-IP issues, questions, services and parameters. Included were: adherence to a schedule, blocking and delay, chokepoints, effect of linear topology problems and multi-hop paths, enforcement to meet performance requirements, fairness, reuse of unclaimed reserved bandwidth, combining best-effort and resource-reservation algorithms, throughput, and traffic characterizations. The latter were described as duration relative to RTT (i.e., << 2 RTT, 2 RTT, and >> 2 RTT), flowrate (steady, compressed steady, bursty), and predictability (none, e.g., interactive, ASAP, e.g., mail, and scheduled, e.g., a conference).

A subset of the working group met Wednesday and Thursday evenings to discuss the practical details of future research collaboration. We agreed that such cooperation was possible, and would result in increased results with an overall decrease in effort. Since most participants felt most comfortable working with UNIX, we decided to adopt it as the experimental platform. We agreed to implement a basic protocol infrastructure in the UNIX-based DRI experimental gateway for experiments across the DRI testbed. MCHIP, ST-2, or other experimental protocols will be built on top of this infrastructure, and this would support experimentation and changes to the protocols. It will be possible to replace the modules that implement different functions, such as resource management or failure detection, relatively easily. By experimenting with them, we will gain practical experience in how different algorithms perform in various situations. These initial implementations will evolve to a single better protocol as we incorporate the better approaches. We are initially planning to implement a MCHIP gateway and host, and an ST-2 gateway and voice and video hosts.

ATTENDEES

Brim, Scott	swb@devvax.tn.cornell.edu
Casner, Steve	casner@isi.edu
Chatterjee, Samir	samir@nynexst.com
Clapp, George	meritec!clapp@bellcore.bellcore.com
Easterday, Tom	tom@nisca.ircc.ohio-state.edu
Fidler, Mike	ts0026@ohstvma.ircc.ohio-state.edu
Fox, Richard	sytek!rfox@sun.com
Gerich, Elise	epg@merit.edu
Goldstein, Steve	goldstein@note.nsf.gov
Lynn, Charles	clynn@bbn.com
McKenney, Paul E.	mckenney@sri.com
Medin, Milo	medin@nsipo.nasa.gov
Parulkar, Guru	guru@flora.wustl.edu
Piscitello, David	dave@sabre.bellcore.com
Ramakrishnan, K.K.	rama%erlang.dec.com@decwrl.dec.com
Su, Zaw-Sing	zsu@tsca.istc.sri.com
Topolcic, Claudio	topolcic@bbn.com
Wilder, Rick	rick@gateway.mitre.org
Yavatkar, Raj	raj@ms.uky.edu

3.3.2 MTU Discovery Working Group (mtudisc)

CHARTER

Chairperson: Jeffrey Mogul/DEC, mogul@decwrl.dec.com

Mailing List:

mtudwg@decwrl.dec.com
mtudwg-request@decwrl.dec.com

Description of Working Group:

The MTU Discovery Working Group is chartered to produce an RFC defining an official standard for an IP MTU Discovery Option. "MTU Discovery" is a process whereby an end host discovers the smallest MTU along a path over which it is sending datagrams, with the aim of avoiding fragmentation.

Specific Objectives:

1. Decide if the proposal in RFC 1063 is sufficient, or if there are flaws to be corrected, or possible improvements to be made. Or, decide that it is unwise to create an official standard.
2. Unless the proposal in RFC 1063 is acceptable, write a new RFC describing a different approach.
3. Encourage the participation of gateway implementors, since the MTU discovery process affects the design and performance of IP gateways.
4. Encourage sample implementations of end-host and gateway portions of MTU Discovery for popular software (BSD-derived kernels, primarily). (b) Encourage rapid implementation by major gateway vendors, since this option is relatively useless without widespread support.

Estimated Timeframe for Completion:

The first two objectives should be completed by April 1990. Objective 4a (sample implementations) should be attempted before the final RFC is released, to alert us to any pitfalls. Objective 4b (implementation by gateway vendors) may take longer.

CURRENT MEETING REPORT

Reported by Jeffrey Mogul/DEC

AGENDA

(a) Report on current draft (McCloghrie/Fox/Mogul)
(b) Review other alternatives
(c) Review goals and assumptions
(d) Obtain consensus on approach
(e) Focus on details
(f) What next?

MINUTES

This was the second meeting of the MTU Discovery Working Group.

We started with a quick presentation by Keith McCloghrie of the draft that he and Rich Fox wrote based on the apparent consensus of the December meeting. Some attendees had not read the draft, and we tried to ensure that everyone understood the basic outline. [Summary: senders occasionally attach an IP PTMU-Query Option to their datagrams. Routers update the PMTU value in the option; the last-hop router returns the PMTU to the sender using the ICMP Path-MTU message. If the destination host detects a change in the MTU (when a fragment is received), it sends an ICMP Unexpected Fragment Report message.]

We also reviewed the "Steve Deering" proposal from last year, as there was a realization that it might not be dead, after all. Among other things, we now know that there are not 1 but 4 spare bits in the IP header (there are 3 unused in the TOS field), and that the powers that be might therefore be likely to let us use one. [Summary of Deering proposal: senders often send datagrams with "RF" (Report Fragmentation) bit set in the IP header. A host receiving fragment-0 of a datagram with RF set sends an ICMP Fragmentation Occurred message.]

We then started a fairly unstructured discussion comparing the costs and benefits of the two approaches.

1. Lifetime of protocol: on the one hand, in principle MTU discovery should be obviated by the coming revolution in routing protocols. Within "a few" years, the routing protocols will provide path-MTU information, so MTU discovery will be unnecessary. Of course, we all know about things that are supposed to happen "real soon now"; we particularly all know about relatively new things

that "everyone" implements. Still, while avoiding the trap of assuming that the world will be perfect in just a couple of years, it may not be worth trying to solve the problem of MTU discovery for all time, since it may not be useful for that long.

2. Rapidity of deployment: Clearly, MTU discovery of any form only works for a sender if some subset of the other nodes (routers and/or destinations) suport it. Query-based schemes depend upon support from a large fraction of the routers; RF-style schemes only help if a large fraction of the end-hosts support it. There was some debate about which population is more likely to upgrade soon (routers or end-hosts). No consensus was reached.
3. Connection lifetimes: Van's data suggest that most non-local TCP connections are short (ca. 4 datagrams). This makes some sense (mostly SMTP) although this is only one sample point, and we agreed that more data would be useful. Van argued that this works against a query-based scheme, since by the time one has useful information, there's not much left to do with it. His argument in favor of the RF scheme was that the right way to use it is to assume that you can send large datagrams (sized by your first-hop MTU, or perhaps some estimate of the NSFNET PMTU, ca. 1500), and let the destination tell you if you are screwing up.

   In general, we realize that fragmentation is not inherently evil. Although it might create some extra overhead for the routers, what we really have to avoid is the "deterministic fragment loss" problem which causes connections to stall. Thus, (I hope I am correctly paraphrasing Van's argument) it is only worth doing for connections that last a while, either because they are carrying lots of data, or because they are stalled due to fragment loss. Query-based schemes waste router resources because processing IP options is expensive, and the payoff is unlikely.

   It was argued that, since the senders cache the MTU values learned by either scheme in the per-host routing entries, querying would not have to be done on every connection to be useful. Again, Van drew on his traffic studies to suggest that (even over a 12-hour period) there was generally little correlation between connections ... that is, just because one pair of hosts makes a connection does not mean that they will do so any time soon. Some of us did not believe that is necessarily true (for example, how much traffic comes from mail-hub machines like DECWRL and UUNET?) Again, we agreed that it would be nice to have more traffic data available.
4. Complexity: Now that the draft specification for the query-based scheme is done, we realized that it is a lot more complex than we thought. One problem is the number of tunable parameters. Since the RF scheme doesn't require the receiver to maintain any state about the sender [actually, this is not quite true, as noted later], doesn't require the sender to schedule when to send the option, doesn't cause the receiver to send notifications when intentional fragmentation

occurs [NFS would probably not set RF], and it requires no support at all from the routers, it appears to be simpler [but keep reading].

After this discussion, it was pretty clear that the consensus had shifted to trying to use the RF scheme. We made the assumption that we could get a header bit (Van argued that although the RF scheme could be done using an option, the cost/benefit analysis might be against it). The next step was to explore how well that would really work.

One problem that came up right away is that James VanBokkelen believes there to exist many PC-based systems that (1) do not reassemble fragments (2) do advertise MSS values of 1500 to non-local peers Currently, these hosts function because the 576-if-nonlocal rule observed by most non-PC hosts means that, given today's Internet, even when they advertise an MTU of 1500 to a non-local host, the host at the other end will not send datagrams big enough to be fragmented. [I suppose it is unlikely for two PCs to talk to each other over long distances.] However, if we use the simplest RF scheme, these hosts are going to get fragmented datagrams. Since we assume that any host which implements MTU discovery is also in conformance with the other rules (specifically, fragmentation reassembly), we therefore know that such sub-standard PCs won't send the ICMP Fragmentation Occurred message, and these connections would stall.

The obvious fix is to not invoke MTU discovery (i.e., not send segments > 576 bytes) unless you are sure that the other end supports it. This means that you have to have seen a datagram with RF set coming back to you from the destination before you can send large datagrams.

More subtly, since we don't want to mislead these stupid PCs (which apparently don't follow the 576-byte rule in either direction) you cannot even send an MSS > 576 to a non-local peer until you have seen an RF bit from it. Thus, since the TCP MSS option can only be sent on the SYN datagram, a host initiating a TCP connection may not be able to use MTU discovery (and large segments) unless it has talked with the other end recently. (The second host is in a better position; since it sees the RF bit before it has to sends its own MSS option, it can set a large MSS immediately. This is nice for FTP retrieves; it doesn't help for SMTP, alas).

The consensus was that this limitation was acceptable, since it erred on the conservative side. (Although it errs on the case of the most common connection-type [SMTP], since SMTP connections are normally short we wouldn't gain much anyway.) When two connections are made in quick succession, things work nicely (e.g., several mail messages, or the control connection of an FTP session followed by the data connection. The control connection will seldom carry large segments, but the exchange of

RF bits done then will allow the data connection to use large segments right away.)

Mike Karels proposed (off-the-cuff, not necessarily believing that it was right) that routers fragmenting a datagram with RF set could also send the fragmentation-occurred ICMP. This seemed to create problems given the requirement for handshaking imposed by the broken-PC crowd, so Mike agreed to go off and think about this one.

One question arose about the use of a previously unused bit in the IP header: what would current implementations do if they see it set? (We know that we can safely add options, since by definition these are ignored if not known.) While the IP spec says these bits must be zero, the "robustness principle" implies that routers and hosts should ignore them. Unfortunately, John Moy from Proteon admitted that Proteon routers drop such datagrams, and Noel Chiappa says that this is true of other implementations based on his old MIT "C-gateway" code. We have to find out just how bad this is going to be; perhaps Proteon will be able to upgrade all of its customers before MTU discovery is widely implemented.

[Side note: Clearly, implementations contrary to the basic IP spec are causing us serious grief. How much do we twist the protocol to accomodate them?]

An orthogonal issue is that in high-speed long-distance networks, there might be lots of packets in flight when the route changes to one with a lower MTU (e.g., on a satellite link with a half-second RTT, 4kb packets, and 100 Mbit/sec channel, this means 1500 packets per RTT!) Since the source cannot react to a Fragment Occurred message sooner than one RTT worth of packets after the one that triggered the message, we are concerned that setting the RF bit on every packet could lead to positive (i.e., anti-stability) feedback in a network that is loosing capacity.

This could be attacked in two ways: limit the rate at which the RF bit is sent, or limit the rate at which the ICMP is sent. The former could be done "once per RTT", once per some constant time period, or perhaps once per window. It's not clear if there is a convenient way of marking out the boundaries between windows

ACTION ITEMS

1. Noel Chiappa and Van Jacobson were assigned to try to get the IESG to free up an IP header bit.
2. Mike Karels was going to think more about having routers send ICMPs when they fragment.
3. We need to determine how many routers will drop packets with RF set, and how hard it will be to fix this. Is it any different if we use one of the bits in the

TOS area?

4. Ditto for end-hosts; are there any that drop such packets?
5. The Router Requirements WG was known to be considering changing the way that fragmentation was done (fragment into equal-size pieces; currently, routers are supposed to send N maximal-size fragments and one smaller one). This would make the RF scheme nearly useless. [Phil Almquist says that the RRWG will work with us on this, so it shouldn't be a problem].
6. Perhaps more traffic studies would be useful.
7. Someone has to write the next draft. Keith and Rich were thanked for their hard work, on their draft that is now tabled, and were not coerced into starting a different document. Since Van was the fiercest proponent of RF at the meeting, he was given responsibility to see to it that the draft is written. He agreed but said he was going to try to get Steve Deering to do the work (Steve was absent due to serious thesis time-pressure, so maybe Van is going to be stuck with it.) The chair requested a draft within one month (7 March 1990).
8. James VanBokkelen was going to see just how many hosts out there are unable to reassemble fragmented IPs, how hard it would be to fix this, how many vendors are involved, etc.

IESG ACTION

On Thursday, February 8, at the open IESG meeting, the IESG was asked to allow this bit to be used for MTU discovery. I was not there, but I understand that the IESG is willing to release this bit if we come to a consensus on a protocol that they think is reasonable.

SCHEDULE

We expect to meet again at the May IETF meeting.

At that point, we will probably either adopt one of the schemes, or give up.

ATTENDEES

Ballard Bare	bare%hprnd@hplabs.hp.com
Art Berggreen	art@sage.acc.com
Richard Bosch	probe@mit.edu
Ron Broersma	ron@nosc.mil
John Cavanaugh	John.Cavanaugh@StPaul.ncr.com
Noel Chiappa	jnc@LCS.MIT.EDU
James Davin	jrd@ptt.lcs.mit.edu
Farokh Deboo	sun!iruucp!ntrlink!fjd
Rich Fox	sytek!rfox@sun.com
Van Jacobson	van@lbl-csam.arpa
Mike Karels	karels@berkeley.edu
Mike Marcinkevicz	mdm@gumby.dsd.trw.com
Tony Mason	mason@transarc.com
Keith McCloghrie	sytek!kzm@hplabs.HP.COM
Bill Melohn	melohn@sun.com
Jeff Mogul	mogul@decwrl.dec.com
John Moy	jmoy@proteon.com
Drew Perkins	ddp@andrew.cmu.edu
Michael Petry	petry@trantor.umd.edu
Nuggehalli Pradeep	pradeep@orville.nas.nasa.gov
Mark Rosenstein	mar@athena.mit.edu
Tony Staw	staw@marvin.enet.dec.com
James VanBokkelen	jbvb@ftp.com
John Veizades	veizades@apple.com
Steve Willis	swillis@wellfleet.com
John Wobus	JMWobus@suvm.acs.syr.edu
David Zimmerman	dpz@convex.com

3.3.3 IP Over FDDI Working Group (fddi)

CHARTER

Chairperson: Dave Katz/Merit, dkatz@merit.edu

Mailing Lists:

fddi@merit.edu
fddi-request@merit.edu

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 group is specifically not chartered to provide solutions to mixed media bridging problems.

Specific Objectives:

To create Internet Standards for the use of IP, ARP, and related protocols on the FDDI medium.

To 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.

To maintain liason with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards.

This working group is not chartered to provide solutions to mixed- media bridging problems, although results produced by this working group should not preclude such solutions.

Estimated Timeframe for Completion:

An Internet Standard or Standards should be produced within six months, with an estimated completion date of May, 1990.

CURRENT MEETING REPORT

Reported by Dave Katz/Merit

MINUTES

The group met on the afternoon of Wednesday, February 7.

1. Document Overview: Dave Katz gave an overview of the current draft IP Over FDDI document, which had been distributed to the FDDI@MERIT.EDU mailing list, for the benefit of those new to the working group. Highlighted were differences between the current draft and RFC 1103.
2. Outstanding Technical Issues:
   (a) A and C Indicators: A discussion ensued on the issue of the A (Address Recognized) and C (Frame Copied) indicators. The current draft states that "the A and C indicators shall be ignored for IP and ARP packets." Objections were raised that this would appear to preclude ANY use of these indicators, such as the management of ARP cache entries. The document editor gave his view that a standard can only specify externally-visible behavior, and that implementation decisions such as ARP cache management could not be precluded.

   The intent of the language regarding A and C was to preclude the use of link-level retransmission in the face of apparent transient congestion in the receiver. The pros and cons of retransmission were debated. After some discussion, the group decided that the usage of the A and C bits would be specified as an implementation decision, with an explicit note that link level retransmission may in fact occur.

   (b) Dual-MAC Issues: Dave Katz provided an overview of the issues regarding the use of dual-MAC stations. Two basic approaches are possible:
      - Separate IP subnetworks on each ring
      - A single IP subnetwork spanning both rings, with both MACs using the same IP address (for load splitting)

   With separate IP subnetworks, the major technical requirement seems to be that all stations properly support subnetting (only sending ARPs for stations on the proper subnet, for example) so that the ring may wrap and unwrap without stations on the two rings learning each others' MAC addresses. A further issue is that if a dual-MAC station wraps the ring, the SMT Configuration Management state machine implies that one of the MACs may be disconnected for the duration of the wrap.

   When a single IP subnetwork is used, the current ARP protocol is insufficient to maintain knowledge of the binding between MACs and rings. In particular, if the ring is wrapped and an ARP is sent for an IP address, two responses may be received at each source MAC, and it becomes ambiguous when the ring unwraps as to which ring each MAC is connected. This problem is made more difficult in the face of the lack of a reliable event-driven indication of the wrap state of the ring (especially if two MAC-less concentrators are performing the wrap). Further complicating this problem are "translucent" bridges between Ethernets and FDDI rings.

It was generally agreed that both the single-subnetwork and dual- subnetwork configurations are desirable, and that they should both be defined, and configurable on a per-LAN basis.

Doug Hunt of Prime presented a straw-man proposal of how to deal with the single IP subnetwork case. It suggests the use of an extension to the ARP protocol that allows the unambiguous determination of the ring on which a MAC is present, even in the face of the ring wrapping and unwrapping. This proposal and other potential solutions were discussed by the group.

It was recognized that the development of the single-subnetwork solution, which is generally viewed as being desirable, is going to take a significant amount of work. No decision was made regarding the mechanism to be used.

3. Document Progression and Future Work: The question of the progression of one or more documents into the IETF standards track was discussed. The choices of action balance a need to produce a standard very quickly versus producing a complete standard.

   The choices are:

   (a) Progress the current document immediately as a single-MAC standard and begin work on a separate dual-MAC standard.

   (b) Quickly write a dual-subnetwork, dual-MAC solution, add it to the current document, progress it as a standard, and begin work on a separate single-subnet, dual-MAC standard.

   (c) Add single- and dual-subnetwork, dual-MAC solutions to the current document and progress it as a standard.

Choice a) has the advantage of starting the base document through the standards process most quickly, significantly moving up the date at which a standard could be published and conformant products could be produced by vendors. It has the disadvantage of being only a partial solution, and may give the impression of favoring single-MAC stations.

Choice b) includes support for dual-MAC stations, but delays the progression of the base document and gives the impression that the dual-subnetwork solution is the "right" solution for dual-MAC stations.

Choice c) provides the most even-handed document in terms of the various solutions, but seriously delays the publication of any sort of standard.

The group decided to pursue the following course:

- Make minor additions and corrections to the current draft, including a statement to the effect that a dual-MAC solution is to follow. Forward this draft into the February X3T9.5 meeting. Incorporate any additional comments from X3T9.5 into the draft and
- publish it immediately thereafter as an Internet Proposed Standard.
- Create a new working group to address "multi-rail" LANs, of which FDDI is a specific case, with the intent of producing an Internet Standard on the subject. Hope was expressed that generalizing this problem would not significantly delay the development of a solution for FDDI.

ATTENDEES

Doug Bagnall	bagnall_d@apollo.hp.com
Samir K. Chatterjee	samir@nynexst.com
Noel Chiappa	jnc@lcs.mit.edu
Dino Farinacci	dino@bridge2.3com.com
Ken Hays	hays@scri1.scri.fsu.edu
Binh Hua	no email
Doug Hunt	dhunt@enr.prime.com
Ronald Jacoby	rj@sgi.com
B.V. Jagadeesh	bvj@chamundi.esd.3com.com
Dave Katz	dkatz@merit.edu
Dave Piscitello	dave@sabre.bellcore.com
Michael Reilly	reilly@nsl.dec.com
Steve Senum	sjs@network.com
Steve Shibuyama	no email
Mary Jane Strohl	strohl@apollo.hp.com
Dean Throop	throop@dg-rtp.dg.com

3.3.4 IP over Switched Multi-Megabyte Data Service (smds)

CHARTER

Chairpersons:	George Clapp	meritec!clapp@bellcore.bellcore.com
	Mike Fidler	ts0026@ohstvma.ircc.ohio-state.edu

Mailing Lists:

smds@nri.reston.va.us
smds-request@nri.reston.va.us

Description of Working Group:

The SMDS Working Group is chartered to investigate and to specify the manner in which the Internet and the newly defined public network service, Switched Multi-megabit Data Service, will interact. The group will discuss topics such as addressing, address resolution, network management, and routing.

Objectives and Milestones:

- The goal of the group is to specify clearly an efficient interworking between the Internet and SMDS.

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech and
Mike Fidler/Ohio State University

MINUTES

The Switched Multi-megabit Data Service (SMDS) Working Group met for the first time for a single half-day session on Thursday morning, February 8. Co-chair Mike Fidler opened discussion by stating the purpose of the group, which is to clarify the manner in which IP may operate over SMDS, and then asked George Clapp to present a tutorial discussion of SMDS (a copy of the viewgraphs is enclosed with the minutes).

SMDS is a switched, connectionless, high-speed data service which will be offered on a nationwide basis by the public carriers. The service is intended to be the equivalent of an IEEE 802 LAN in functionality and performance and is designed to fit within the internet protocol stack as a transit network to IP. First trials may occur in late 1990; the first tariffed service may occur in 1991; and the service may be widely tariffed and deployed in 1992.

A number of questions arose as George progressed through the SMDS tutorial. The first was cost. Members of the group felt that they could not evaluate the service until they had an idea of the cost and of how that cost compared with the cost of a leased private line. George responded that the tariff structure was not yet determined but that the public carriers recognized that SMDS must be cost competitive with a leased private line. He then queried the group whether they would prefer flat or usage-based billing. Members answered that the essential feature to billing was predictability and that a flat fee was preferred since network administrators had little knowledge or control over the traffic generated by their network.

George ended the tutorial by presenting the following list of potential topics to be discussed by the working group.

- Addressing and Address Resolution
- Routing
- Network Management

With respect to addressing, SMDS uses a 60 bit address similar in format to a telephone number. It may be possible to extend ARP to handle the 60 bit SMDS address in response to a query for an internet address. The notion of ARP itself, however, may be extended to that of a "directory service," in which SMDS returns a 60 bit

SMDS address in response to a network address as well as to an internet address.

With respect to routing, this function may be done in a number of ways over SMDS. The routers of an organization may operate as before by exchanging link state packets via SMDS to build and maintain routing tables. The issue which arises in this approach is the cost of the multicast packets, which depends upon the number of routers and upon the frequency of the generation of the link state packets. If the cost grows too large, alternative approaches may be desirable. One alternative may be to use the previously mentioned directory service to build the routing table. Rather than exchange link state packets, a router may query the directory service for the SMDS address of an internet network. The approach, however, would require an extension to existing routing protocols.

The discussion of routing brought out two models in which SMDS may operate. One is a Private Virtual Network (PVN) in which SMDS interconnects a set of routers belonging to an existing organization. In this model, communication among devices is restricted to those devices which belong to the PVN and communication with devices external to the PVN would be carefully controlled. The issues which arise in this model are how it would be done and what SMDS features would enhance the service. The second model is that of a public network, analogous to the existing telephone network, in which an SMDS device may communicate with any other SMDS device. The issues which arise here are security concerns, such as restricted access and authentication, and the issue of scale, since existing algorithms may not operate in an environment with large numbers of devices. A third model was suggested in which existing leased lines of a private virtual network are kept and SMDS is accessed to provide additional capacity.

A number of other questions were raised.

- What will be the performance of SMDS and what are the kinds of services that SMDS may adequately support?
- What kind of network management features will SMDS support? Will SMDS "speak SNMP?"
- How would internet access the proposed directory service?
- To what extent will SMDS support multicast and how should multicast be used? For example, it would be necessary to limit the extent of an ARP multicast in the public network model for SMDS, in which there is universal connectivity among SMDS devices.

At the end of the meeting, the group tentatively scheduled a video conference for either March 27 or 28. (Note: we have subsequently learned that these dates are unavailable; new dates are not yet determined.)

Mike Fidler had asked those present to indicate on the sign up sheet whether they wished to participate in the SMDS mail list. This mail list will be built and communication established after the IETF meeting.

ATTENDEES

Chet Birger	cbirger@bbn.com
Scott Bradner	sob@harvard.harvard.edu
Mats Brunell	mats.brunell@sics.se
Ted Brunner	tob@thumper.bellcore.com
Steve Casner	casner@isi.edu
Samir Chatterjee	samir@nynexst.com
Steve Crocker	crocker@tis.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Kent England	kwe@bu.edu
Dino Farinacci	dino@bridge2.3com.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Dale Finkelson	dmf@westie.unl.edu
Der-Hwa Gan	dhg@bridge2.3com.com
Ella Gardner	epg@gateway.mitre.org
Herve Goguely	rvg@bridge2.3com.com
Steve Goldstein	goldstein@note.nsf.gov
Jack Hahn	hahn@umd5.umd.edu
Gene Hastings	hastings@psc.edu
Juha Heinanen	jh@funet.fi
Chris Hemrick	cfh@sabre.bellcore.com
Bob Hinden	hinden@bbn.com
Steven Hunter	hunter@ccc.nmfecc.gov
Tom Hytry	tlh@iwlcs.att.com
Dan Jordt	danj@cac.washington.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Walt Lazear	lazear@gateway.mitre.org
Dan Long	long@bbn.com
Charles Lynn	clynn@bbn.com
Milo Medin	medin@nsipo.nasa.gov
Berlin Moore	prepnet@andrew.cmu.edu
Dennis Morris	morrisd@imo-uvax.dca.mil
Don Morris	morris@ucar.edu
John Moy	jmoy@proteon.com
Dave O'Leary	oleary@umd5.umd.edu
Donald Pace	pace@fsu1.cc.fsu.edu

Guru Parulkar	guru@flora.wustl.edu
Dave Piscitello	dave@sabre.bellcore.com
Dave Pokorney	poke@nervm.nerdc.ufl.edu
Ira Richer	richer@vax.darpa.mil
Jim Showalter	gamma@mintaka.dca.mil
Martha Steenstrup	msteenst@bbn.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
Claudio Topolcic	topolcic@bbn.com
Greg Vaudreuil	gvaudre@nri.reston.va.us
Ross Veach	rrv@uiuc.edu
Steven Willis	swillis@wellfleet.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@igloo.osc.edu
Raj Yavatkar	raj@ms.uky.edu
David Zimmerman	dpz@convex.com

Switched Multi-megabit Data Service (SMDS)

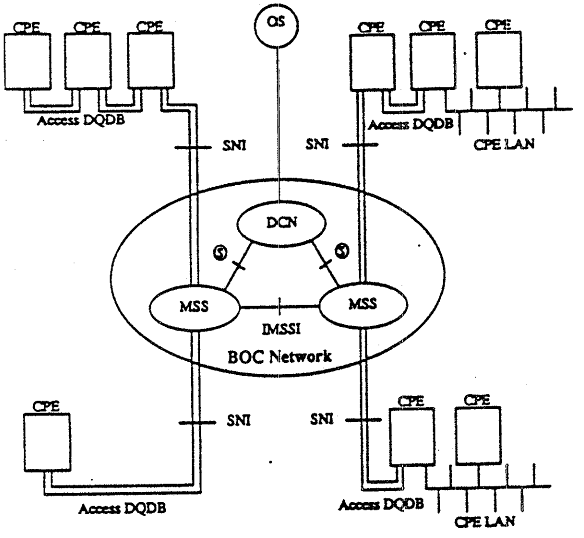

SNI: Subscriber Network Interface
MSS: MAN Switching System
IMSSI: Inter-MAN Switching System Interface
Ⓢ: Generic Interface for Operations
OS: Operations System

DCN: Data Communications Network
CPE: Customer Premises Equipment
CPE LAN: CPE Local Area Network
DQDB: Distributed Queue Dual Bus

George H. Clapp

AMERITECH SERVICES

ISSUES

ADDRESSING AND ADDRESS RESOLUTION.

- IP ROUTERS EXCHANGE NORMAL PACKETS TO FILL ROUTING TABLES.
- " Address Resolution ~~DIRECTORY~~ SERVICE " : STATION GIVES IP ADDRESS TO SMDS NETWORK AND RECEIVE SMDS ADDRESS.
  - INDIVIDUAL
  - NETWORK
  - SET OF NETWORKS
  - ALTERNATIVE MEANS USED BY ROUTERS TO FILL TABLES.

NETWOEK MANAGEMENT
SNMP-MIB issues.

1- Private Net
interconnecting router.

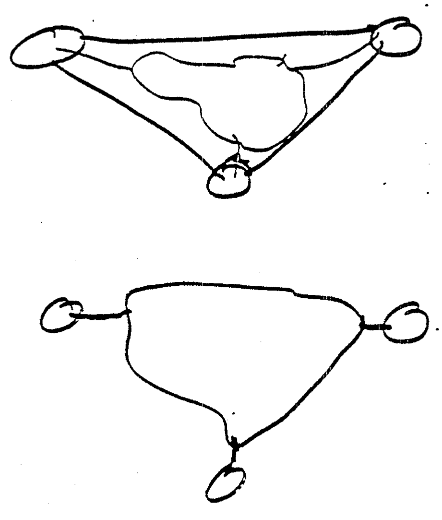

2- Open Connect
(public Net)

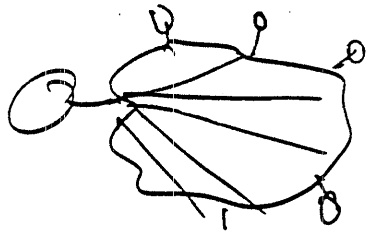

Target Broadband Data Architecture

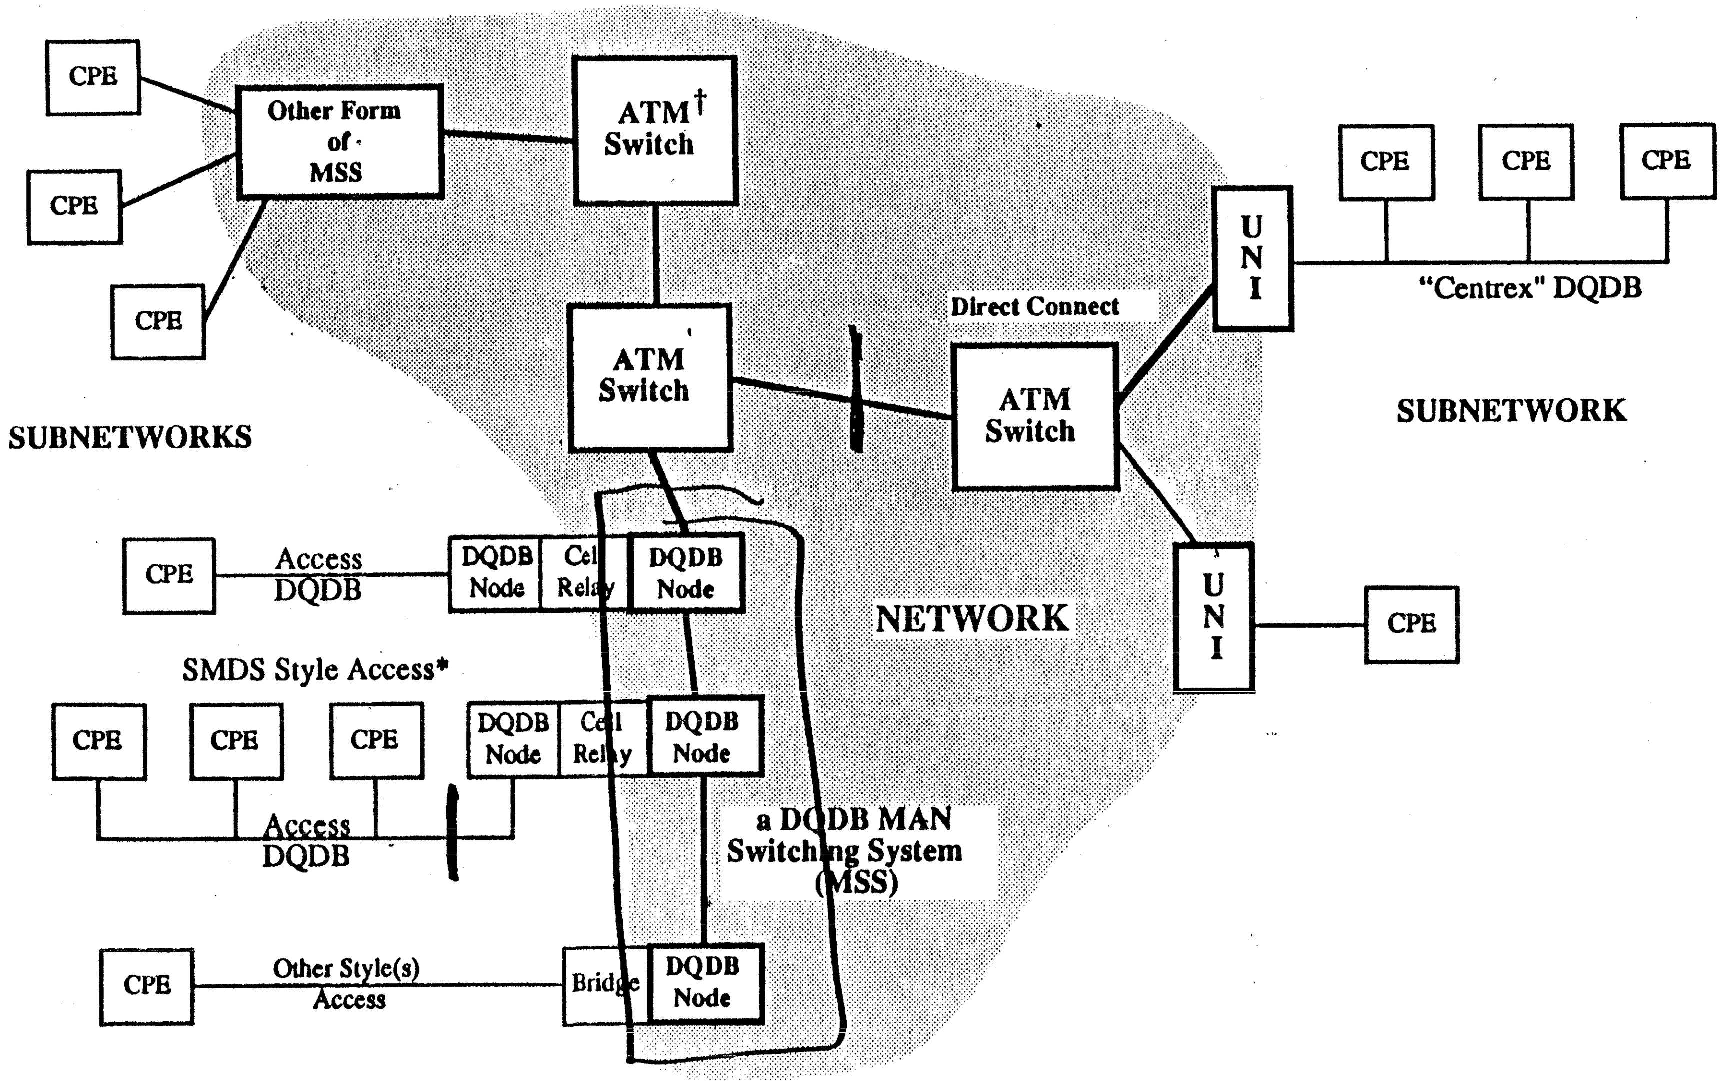

* Per Current Bellcore TA-772
† ATM Switches can initially be DQDB Bridges

George H. Clapp

SMDS, MAN Technology and BISDN

David M. Piscitello

Bellcore
Bell Communications Research

SMDS, MAN Technology, and BISDN

Outline

- **Introduction**
- Switched Multi-Megabit Data Service (SMDS)
- Early Availability SMDS via MAN Technology
- SMDS in an Integrated Services Broadband Network

Where are we coming from?

High-Speed Data Networking: A Typical Scenario

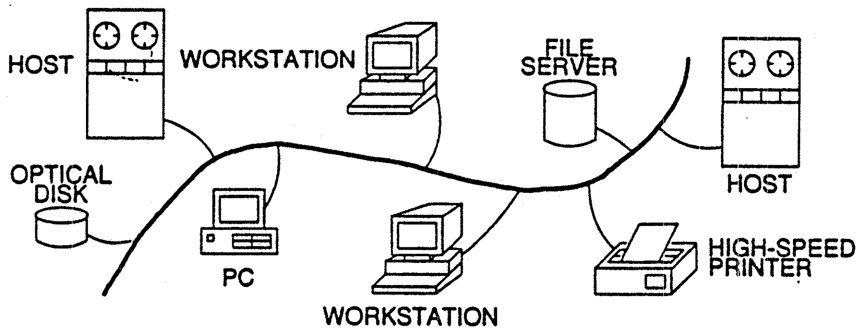

- Premises-based, high-speed, local area networking
- Powerful and increasing affordable CPU
- Shared, distributed resources and distributed applications
- Widespread use of "connectionless internetworking"

Where are we going?

Communications Technologies in the Next Millenium

- Control methods providing aggregate channels up to 2.4 Gb/s
- Packet switching fabrics providing 1.6 Gb/s *per channel*
- Fully optical switching
- Integrated voice, data, and video services
- Distributed applications we haven't even dreamed of yet...

What's wrong with this picture?

Inter-Premises Data Transport: Today

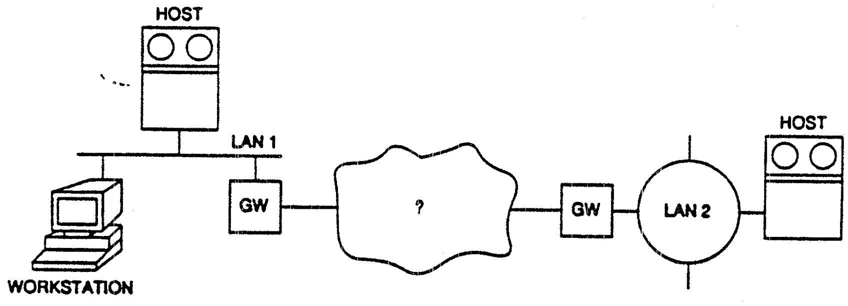

- Current options:
  - — circuit switched at ≤ 64 Kbps
  - — packet switched with ≤ 64 Kbps interfaces
  - — private lines at ≤ DS1 rate (1.5 Mbps)
  - — some DS3 (44 Mbps) private line services available
- Most options are either "slow" (relative to LAN speeds) or costly

What can be done today?

Capitalize on what's available

- Many key components available now:
  - — Optical fiber
  - — MAN technologies
  - — Innovative switching
- Take it out of the lab and put it into the telco plant now
- Why?
  - — Accelerate the development of new applications

SMDS

A Wide-Area, High-Speed Data Service Solution

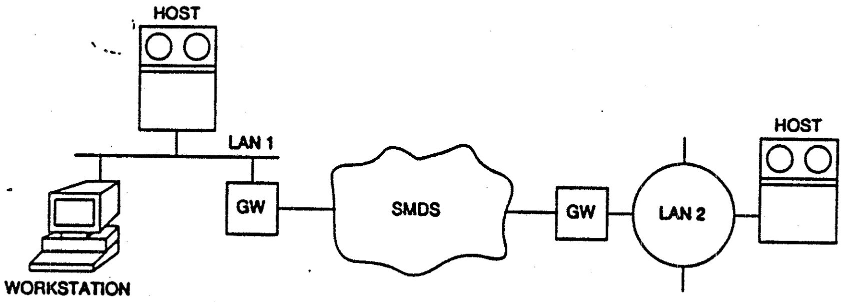

SMDS

A Wide-Area, High-Speed Data Service Solution

- A high-performance, public, packet-switched data service
  - Connectionless (datagram) transfer mode
  - LAN-like performance and features across a wide area
- Easily incorporated into customers' existing networks
  - Simple (MAC-level) interface to customer's data networking equipment
  - Synergy with many data networking architectures

Role Of SMDS In Target Customer's Internetwork

- SMDS viewed as a **subnetwork** in customer's internetwork
- Variety of data networking architectures can easily accommodate SMDS

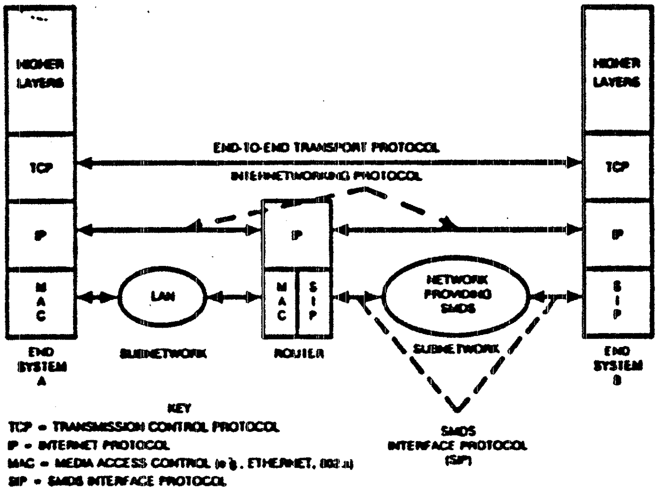

KEY
TCP = TRANSMISSION CONTROL PROTOCOL
IP = INTERNET PROTOCOL
MAC = MEDIA ACCESS CONTROL (e.g., ETHERNET, 802.x)
SIP = SMDS INTERFACE PROTOCOL

- Same treatment of SMDS in other architectures
  - OSI, DECnet, XNS . . .

Role of SMDS in Distributed Networked Systems

- LAN Interconnection
- High-performance Client/Server interconnect
- Window terminal/Window server interconnect
- Supercomputer access
- Channel extension

Advantages of SMDS

- $$$
  - — Economies of a shared network
  - — Pay per usage
- Customer invests in a **Service**, not a technology
  - — Fiber-based transmission facilities and switches deployed, maintained and operated by BOCs
  - — Considerable flexibility with respect to customer network expansion, changing traffic requirements, etc.
  - — Opportunity for future, integrated access to other BISDN services

SMDS, MAN Technology, and BISDN

Outline

- Introduction
- **Switched Multi-Megabit Data Service (SMDS)**
- Early Availability SMDS via MAN Technology
- SMDS in an Integrated Services Broadband Network

What is SMDS?

- A public, high-speed, packet-switched data service concept
  - An inter-premises (wide area) high speed data communications solution
  - Supportable by different technology platforms
  - First service to be supported by Broadband ISDN
- LAN analogous service features
  - High bandwidth, low delay
  - Large packet sizes
  - Multi-cast or group addressing

Subscriber-To-Network Interface (SNI)

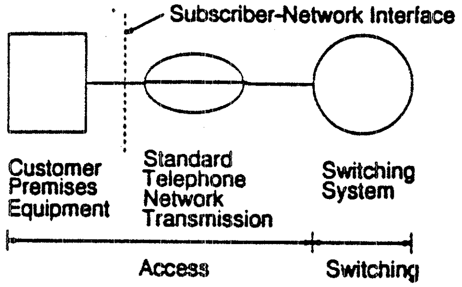
Subscriber-Network Interface
Customer Premises Equipment
Standard Telephone Network Transmission
Switching System
Access
Switching

- An access path dedicated to a single subscriber
- DS3 (44.21 Mbps) or DS1 (1.536 Mbp) based access
- For DS3 access path, single or multi-CPE access arrangement

135

Subscriber-Network Interface (SNI)

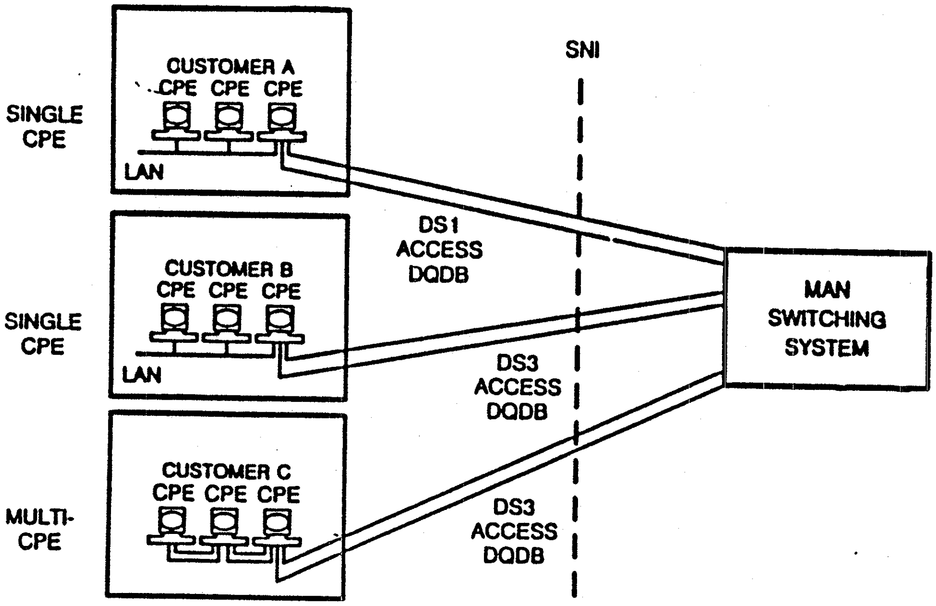

Data Transport

- Connectionless (Datagram) Service
  - — each packet addressed and transferred independently
  - — no connection set up or release
  - — no acknowledgements
  - — no explicit flow control
- Variable length packets
- Very large maximum packet size (9188 octets)
- Credit management based access classes

Addressing

- E.164 addresses (numbering plan for ISDN Era)
- Multiple addresses per SNI
- Address Validation
  - Network ensures that sender cannot indicate a fraudulent source address

Group Addressing

- Ability to use a single destination address to identify a set of destinations
  - Analogous to LAN Multi-cast addressing feature

Address Screening

- Source address screening
  - Ability to restrict delivery of SMDS data units coming from list of senders
- Destination address screening
  - Ability to restrict transmission of SMDS data units to list of senders
- Used by customers to establish higher degree of security/privacy over public data service
  - Forms a “logical private network”

Access Classes

- An access class defines average and peak rates of data flow into and out of network (chosen at subscription)
  - Access class is enforced by implicit flow control mechanism
  - No explicit protocol functions to enforce access classes
- Network always conforms to access class when sending to customer
- Performance objectives of service not guaranteed if CPE exceeds subscribed-to access class

SMDS Performance Objectives

Apply from SNI to SNI

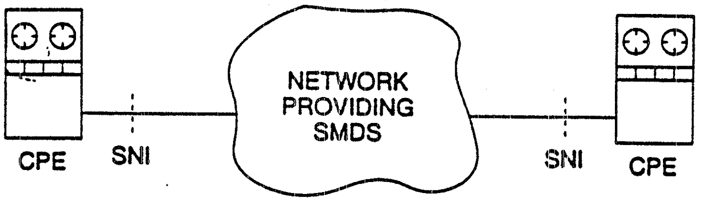

- Errored L3_PDU ratio < 5 x 10^{-13}
- Misdelivered L3_PDU ratio < 5 x 10^{-8}
- Not delivered L3_PDU ratio < 1 x 10^{-4}
- Average delay for two DS3 interfaces: 20 msec
- Average delay for a DS3 interface and a DS1 interface: 75 msec
- Average delay for two DS1 interfaces: 120 msec

SMDS, MAN Technology, and BISDN

Outline

- Introduction
- Switched Multi-Megabit Data Service (SMDS)
- **Early Availability SMDS via MAN Technology**
- SMDS in an Integrated Services Broadband Network

Basic Approach for Early Availability SMDS

- Respond to user needs by offering a service that could be deployed quickly
  - — Use existing MAN technology
  - — Use embedded plant of digital transmission facilities
  - — Minimize implementation impact on CPE

Application Of MAN Technology

- As the basis for the SMDS Interface Protocol (SIP) across the Subscriber-Network-Interface (SNI) between the CPE and switching system in BOC network
- As the basis for switching systems in the BOC network

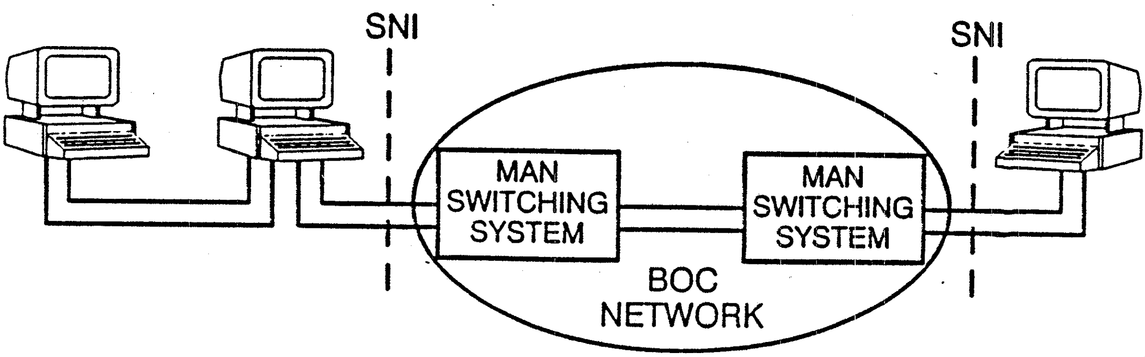

Two Applications of MAN Technology

As the basis for MAN Switching Systems (MSSs)

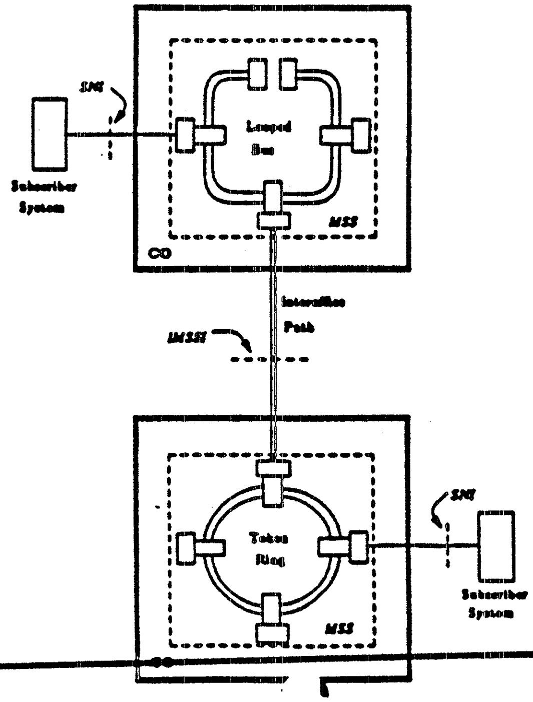

SMDS, MAN Technology, and BISDN

Outline

- Introduction
- Switched Multi-Megabit Data Service (SMDS)
- Early Availability SMDS via MAN Technology
- **SMDS in an Integrated Services Broadband Network**

Characteristics of Broadband ISDN

- Customer interface rates of 150 Mbps and 600 Mbps
- Traffic from all services (voice, data, video) carried in ATM cells
- ATM cells carried in Synchronous Optical NETwork (SONET) payload
- Single-mode optical fiber

Synchronous Optical Network (SONET)

A Standard Optical Interface

- SONET:
  - — establishes "building block" signal (OC-1 = 51.840 Mbps)
  - — uses synchronous byte-interleaved multiplexing methods
  - — provides a hierarchy of standard signals (integer multiples of OC-1; e.g., OC-3 = 3 x 51.840 Mbps)

SMDS service continuity

- SIP Level 3 protocol preserved
  - E.164 addressing
  - Maximum packet size
- Features preserved
  - Address screening
  - Group addressing
  - etc.

High Speed Connectionless Data Service Standard Broadband Access

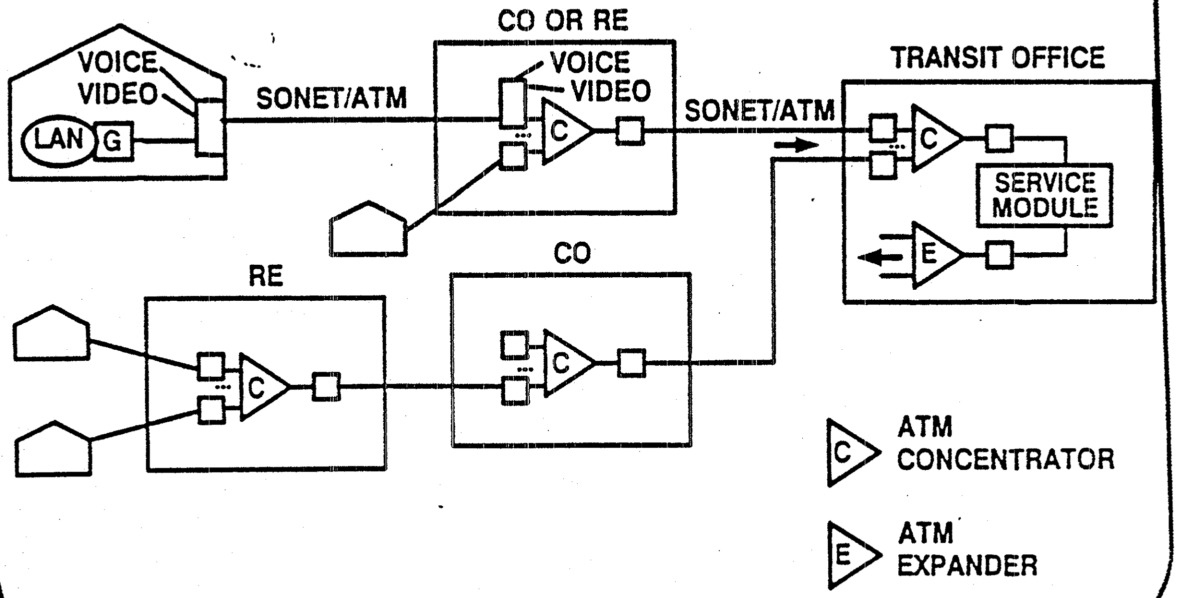

Summary

- SMDS is a **Service** that can be supported by several technology platforms
- LAN analagous attributes and performance are directed at
  - — distributed processing applications
  - — image processing
  - — supercomputer access
  - — channel extension
- 1991 is the target for early availability

BISDN

SMDS service continuity

SIP Level 3
SIP Level 2
DS3 or DS1

Early availability protocol stack

<table>
<tr><td>SIP Level 3</td><td>CO Data</td><td rowspan="2">Voice, Video, ...</td></tr>
<tr><td>AL-CL</td><td>AL-CO</td></tr>
<tr><td colspan="2">AL-Data</td><td>Other ALs</td></tr>
<tr><td colspan="3">Asynchronous Transfer Mode</td></tr>
<tr><td colspan="3">SONET</td></tr>
</table>

AL - Adaptation Layer
CL - Connectionless
CO - Connection Oriented

BISDN protocol stack

3.3.5 Point-to-Point Protocol Extensions Working Group (pppext)

CHARTER

Chairpersons:	Russ Hobby/UC Davis	rdhobby@ucdavis.edu
	Steve Knowles/ FTP	stev@ftp.com

Mailing lists:

ietf-ppp@ucdavis.edu
ietf-ppp-request@ucdavis.edu

Description of Working Group:

The Point-to-Point Protocol (PPP) was design 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. The group will also define new PPP options for the existing protocol definitions, such as stronger authentication and encryption methods.

Specific Objectives:

The main objective of the working group is to produce an RFC or series of RFCs to define the use of other protocols on PPP.

Estimated Timeframe for Completion:

The RFC(s) should be complete during the year.

CURRENT MEETING REPORT

Reported by Russ Hobby/UC Davis

MINUTES

The PPP Extensions WG met on February 6 and 7 at the IETF meeting at Florida State University. The primary subject of discussion was the method for Link Quality Monitoring in the PPP Options document. The main problem was how to clearly explain the process. Some changes were made in the way that error variable are exchanged so that the process can broken up into separate tasks, thus making it easier to understand and implement. These changes will be make and the document will be submitted to be an RFC.

Fred Baker presented his paper on how to do bridging over PPP. Fred will follow up on the comments make and work toward the goal of an RFC.

No progress was made on DECNET or XNS over PPP.

ATTENDEES

Baker, Fred	baker@vitalink.com
Berggreen, Art	art@sage.acc.com
Bryant, Stewart	bryant@janus.enet.dec.com
Jacobsen, Ole	ole@csli.stanford.edu
Jacobson, Van	van@helios.ee.lbl.gov
Jagadeesh, B.V.	bvj@chamundi.esd.3com.com
Kaufman, David	dek@proteon.com
Knowles, Stev	stev@ftp.com
Lauck, Tony	lauck@dsmail.dec.com
Merritt, Don	don@brl.mil
Monachello, Dave	dave@pluto.dss.com
Perkins, Drew	ddp@andrew.cmu.edu
Petry, Mike	petry@trantor.umd.edu
Reilly, Michael	reilly@nsl.dec.com
Rosenstein, Mark	mar@athena.mit.edu
Senum, Steve	sjs@network.com
Solensky, Frank	solensky@interlan.interlan.com
Throop, Dean	throop@dg-rtp.dg.com
Waldbusser, Steve	sw01@andrew.cmu.edu

3.3.6 Router Discovery Working Group (rdisc)

CHARTER

Chairperson: Steve Deering/Stanford, deering@pescadero.stanford.edu

Mailing Lists:

gw-discovery@gregorio.stanford.edu
gw-discovery-request@gregorio.stanford.edu

An archive of all mail to the list is available by anonymous FTP from host gregorio.stanford.edu, file gw-discovery/mail-log.

Description of Working Group:

The Gateway 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 gateway availability by hosts, as opposed to the initial configuration of hosts.

Specific Objectives:

1. Identify existing and proposed protocols, and if necessary develop a new protocol, for gateway discovery.
2. Evaluate the protocols identified in 1 for suitability as Internet standards, according to criteria to be agreed upon by members of the Working Group. For new protocols or extensions to existing protocols, the evaluation shall include prototype implementations before being proposed as a standard.
3. Produce an RFC recommending a standard protocol for gateway discovery.

Estimated Timeframe for Completion:

It is hoped that the Working Group can complete all of its objectives within 6 months of its initial meeting.

CURRENT MEETING REPORT

Reported by Steve Deering/Xerox Parc, from notes by Jeff Mogul and James VanBokkelen

MINUTES

This was the second meeting of the Router Discovery (nee Gateway Discovery) working group. Jeff Mogul served as acting chair, in Deering's absence.

The proposed protocol from the December meeting was reviewed. The significant features are:

- It is an ICMP extension.
- Routers periodically multicast router reports; hosts multicast router queries at boot time only.
- Use of broadcast instead of multicast is permitted but discouraged.
- A router report does not include a subnet field.
- A router report includes a holding-time field and a preference-level field.
- In cases where more than one subnet is assigned to the same physical network, a router may include multiple addresses (i.e., one for each subnet) in a single router report.

Jeff identified the following open issues:

1. Meaning of preference levels:
   Nobody seemed to know what to do with more than three levels (primary, backup and last chance?).
   Decision: use 8 or 16 bits, whatever fits in the packet format.
2. Choice of multicast period:
   Noted that ES-IS uses 10 seconds; we may want slower?
3. How should router respond to query, unicast or multicast?
   Mike Karels proposed that routers be allowed to attempt to avoid unnecessary replies, substituting a single broadcast or multicast for several unicast replies.
   Decision: "keep it simple", i.e., always send unicast replies.
4. Who writes the RFC?
   No volunteers, so it's still Deering's responsibility.
5. Do clients dally before sending queries?
   Yes, so that if a LANful of X terminals reboot from ROM at the same instant, they don't all hit the broadcast at the same instant.

Other issues raised:

- Use on non-broadcast media.
  Dismissed as too complicated.
- Do routers dally before replying?
  Someone suggested that the router dally randomly (mean time based on pref level) to avoid overwhelming client. We argued over whether the clients needed all the possible router responses right off. However, since we don't want to invent a protocol to stop the other N routers from responding, we realized that if we were already going to expend the resources for N+1 packets on the wire, we might as well try to make use of them at the client. So, dallying seems to be wanted.
- When to query?
  Drew Perkins proposed a minor shift of definitions to allow initial query to be sent when a gateway is first needed (i.e., when first sending to an off-subnet destination), rather than at boot time.

ATTENDEES

Ballard Bare	bare%hprnd@hplabs.hp.com
Art Berggreen	art@sage.acc.com
Richard Bosch	probe@mit.edu
Ron Broersma	ron@nosc.mil
John Cavanaugh	John.Cavanaugh@StPaul.ncr.com
James R. Davin	jrd@ptt.lcs.mit.edu
Farokh Deboo	fjd@interlink.com
Rich Fox	sytek!rfox@sun.com
Mike Karels	karels@berkeley.edu
Tony Mason	mason@transarc.com
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Bill Melohn	melohn@sun.com
Jeff Mogul	mogul@decwrl.dec.com
John Moy	jmoy@proteon.com
Drew Perkins	ddp@andrew.cmu.edu
Michael Petry	petry@trantor.umd.edu
Mark Rosenstein	mar@mit.edu
Tim Seaver	tas@mcnc.org
Tony Staw	staw@marvin.enet.dec.com
James VanBokkelen	jbvb@ftp.com
John Veizades	veizades@apple.com
Steve Willis	swillis@wellfleet.com
John M. Wobus	jmwobus@suvm.acs.syr.edu
David Paul Zimmerman	dpz@convex.com

3.3.7 Router Requirements Working Group (rreq)

CHARTER

Chairpersons:	Jim Forster/Cisco	forster@cisco.com
	Philip Almquist/Stanford	almquist@jessica.stanford.edu

WG Mailing List:

- ietf-rreq@Jessica.Stanford.EDU: general working group mailing list. To be added to or deleted from the list, please send a note to:
  ietf-rreq-requests@Jessica.Stanford.EDU.
- ietf-rreq-interest@Jessica.Stanford.EDU: mailing list primarily for announcements of new drafts. To be added to or deleted from the list, please send a note to:
  ietf-rreq-interest-requests@Jessica.Stanford.EDU.
- ietf-rreq-editor@Jessica.Stanford.EDU - mailing address for non- technical comments about the document (grammar, readability, spelling, etc). This address may also be used by authors who, for whatever reason, wish to submit comments or proposed text anonymously.

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 explictness levels of the Host Requirements RFC's, as well as b) including references to more recent work, such as the RIP RFC and others.

Specific Objectives:

- Produce a draft document for initial comment by the community by the summer of 1990.

Estimated TimeFrame for Completion:

The objective is to have a completed document ready to be made into an RFC by early in 1991.

CURRENT MEETING REPORT

Reported by F. Baker/Vitalink, S. Senum/Network Systems Corporation, P. Almquist/Consultant and J. Forster/cisco Systems

MINUTES

The IETF Router Requirements Working Group is a new working which met for the first time during the IETF meeting in Tallahassee. The IESG formed this working group to draft a standard for Internet IP routers which will be up to date and which will match the level of clarity and completeness achieved in the recent Host Requirements RFC's (RFC1122 and RFC1123). The group intends to submit its document to the Internet standards process in early 1991. If it is accepted, it will replace RFC1009, the current standard for Internet routers.

The group held three half-day meetings in Tallahassee. Topics discussed fall into four basic categories:

1. Group startup activities. This included such things as discussing the charter and the schedule and agreeing on exactly what is the task to be performed.
2. Creation of an outline. The co-chairs drew up a strawman outline for the document in advance of the meeting. The group adopted this outline with some modifications.
3. Distribution of writing assignments. Many of the group members agreed to draft sections of the document before next meeting, tentatively scheduled to be a videoconference in late March, 1990.
4. Initial discussion of technical issues. Close to half of the meeting was devoted to discussion of what the document should say about a variety of issues, such as ARP and IP option support.

UPCOMING SCHEDULE

- February 6-9, 1990: IETF meeting, FSU
- Late March, 1990: video conference
- May 1-4, 1990: IETF meeting, PSC
- mid-June, 1990: video conference
- July 31-August 3, 1990: IETF meeting, UBC
- August 15, 1990: first Internet draft version submitted
- mid-September, 1990: video conference

- early November 1990: IETF meeting, Washington University
- December 3, 1990: second Internet draft version submitted
- mid-December, 1990: video conference
- January 15, 1991: final Internet draft version submitted
- early February 1991: IETF meeting, NCAR
- February 15, 1991: final version submitted to IESG and RFC editor

MINUTES - 2/6/89

Introduction:

The charter of the working group calls for a descendant of the Router Requirements Document (RFC 1009), modelled on the general format and spirit of the Host Requirements Document (RFCs 1122 and 1123).

The initial submission is an outline, which may be reorganized, and will be fleshed out by the submissions of a number of authors. There is also a proposed schedule for the work to be done, culminating in a final RFC in early 1991.

In support of this, a commitment is requested and required of all participants, to expedite the delivery of this document.

Charter:

Several documents feed into this process:

- RFC 1009 Router Requirements
- RFC 1122 Requirements for Internet hosts - Communication Layers.
- RFC 1123 Requirements for Internet hosts - Application and Support.
- Many RFCs included by reference

There was general agreement on the charter as written, with the suggestion of some extra wording on the general subject of Interoperability.

Presentation Style:

The IESG proposes a document much like the Host Requirements Document, especially in the sense of flagging sections as 'required' and 'optional', flagging requirements stated in those sections as 'must [not]', 'should [not]', and 'may', and the concepts of 'full' and 'conditional' compliance. Compliant Routers (full or conditional) should interoperate by definition. This is accepted as the initial strategy, and the usage of those terms is intended to be as much like its predecessor as possible.

By way of example, ARP is generally an optional protocol, and should be stated as optional; However, if Ethernet interfaces are implemented, ARP implementation is required, and certain operational experience since the original RFC was written must be accounted for.

A postscript version of the document may come into existence; however, to facilitate world wide ease of access to the document, a text version WILL be available.

A number of points that came out in a general discussion (not all of which represent any kind of concensus of the group):

Vendors need the document in order to know how to implement a universally acceptable product, and Network Managers need to one to specify in RFQs. Net Managers ALSO need a document indicating how the features of a compliant Router are intended to be used. These are not necessarily the same document.

We need a Multi-protocol Router Document, although this is technically outside the working group's charter. This document should explicitly not preclude alternative architectures, and will attempt to state requirements that will allow multi-protocol routers to be compliant.

There needs to be a section regarding global traffic engineering. Congestion Management is a special case of traffic engineering. There will be a discussion of Congestion Management, especially Fair Queuing, but the details may be referred to another working group.

Verbiage needs to be clear and consistent, and consistent (where relevant) with the Host Requirements Document. A Glossary may be helpful.

IP Multicast needs to be dealt with on a MAC by MAC basis.

The 'All Subnets' Broadcast, although required by RFC1009, is probably not implemented by anyone, and its originally intended function is probably better served by IP multicast. Also, the all subnets broadcast would be dangerous if implemented, since hosts which are unaware that the network is subnetted may generate packets that look like all subnets broadcasts unintentionally.

There seem to be two options for dealing with the all subnets broadcast. The first is to require it, since it would not be useful unless widely enough implemented that applications could reasonably expect it to work. If we do that, we should also specify an/the agorithm for the flooding. An alternative is to declare the entire notion of the all subnets broadcast to be obsolete. Philip Almquist will write an RFC suggesting the latter.

Line protocols recommended by this document must have some sort of protocol discriminator field. Point-to-Point and ISO 7776/8880(3) both have this.

Apple Localtalk has an RFC in progress.

X.25 protocol:

- there are several alternatives for discriminating between protocols on X.25 - most notably using a discrimination octet on each packet and running separate Virtual Circuits by protocol, TOS, destination tuple.
- several de facto standards exist: DDN 'basic' and 'standard', Blacker, and European.

ARP needs to be fairly closely spelled out, especially regarding Proxy ARP and ARP Response to IP Multicast Address. (See minutes from 2-7-90)

Hyperchannel and ARCnet have inadequate current interest to justify a section in the document.

There needs to be a precedence statement indicating that this document takes precedence over the base RFCs for each feature, and over the Host Requirements Document in certain cases. Writers are expected to reference the Host Requirements Document regularly; in the draft versions, please include the relevant text to simplify reviewer's cross-referencing; this will be removed from the final draft.

The Objective (writers take note!) is to specify the external characteristics of an Internet standard Router, not the algorithms it uses to implement that behavior. For example, this document should state that the ARP cache should lose track of information about hosts that disappear from the network, and should do so reasonably expeditiously. Whether this depends on internal timers 'popping', or on entries being found to be invalid upon reference, or some other algorithm, is not for this document to specify.

Routing Protocols:

Thou shalt not digress into religion, politics, or the correct standard IGP! The IESG and IAB are expected to decide on the standard Internet IGP, and this document should reference their decision.

EGP2 and BGP should be documented; EGP3 should not be.

Re: Filters and Controls, the effects of these should be specified without specifying the specific syntax or semantics.

Network Management:

By default, routers should allow public SNMP read access to at least the subset of the MIB required for diagnosis of end-to-end connectivity problems. However, the manager of the router should be able to disable the public access. The security aspects of SNMP need to be examined in more depth.

There needs to be universal read access to a subset of the SNMP readable information, with significant control of the ability to write. However, attention should be given to the security aspects of SNMP readability.

Performance requirements should include or reference standard tests, network stability under load, the forwarding and timely processing of updates, and the priority (if any) those updates should enjoy.

There should perhaps be a vendor independent (potentially SNMP based) user interface to all Routers.

MINUTES - 2/7/90

Address Resolution Section

The day started out with a detailed discussion of ARP. Generally, people seemed to feel that MAC-specific details of ARP should be discussed in the relevant MAC layer sections, but a stand- alone section should cover common ARP issues. This section should be called Address Resolution, and should also cover X.213/X.25 addressing issues (referencing the PDN Working Group's work).

Several viewpoints were brought out on most of the following points, but these represent the endpoints of several (sometimes simultaneous) discussions:

- Routers are generally triggered to ARP by a message which needs to be forwarded to a currently unknown system. While the impact of not holding the triggering packet is not great, a Router 'should' hold and re-transmit a small number of such messages for a limited period of time.

- The ARP procedure calls for periodic refreshing of the ARP database, potentially using a broadcast in case the system has changed its MAC address. Generally, the first refresh attempt should be a unicast to the last known MAC address.
- Proxy ARP is useful in circumstances where the Network Administrator can't or doesn't choose to advise his hosts of the local subnet architecture, or where the architecture is ambiguous, as in a multi-rail FDDI with some single rail systems on each ring. It may be viewed as a simplistic Router Discovery Protocol or a subnet disambiguator. Use of Proxy ARP in normal networks, however, is discouraged. Routers 'may' provide it, it 'must' be configurable, it 'must' be disabled by default, and 'should' be configurable by interface.
- Systems 'must not' respond to an ARP for any recognizable Broadcast or Multicast address (Class D, 0.0.0.0, 255.255.255.255, Network or Subnet variants).
- Routers 'should' emit an ARP request for their own address upon startup, and log an error in the event that anyone responds. Similarly, during normal operation, any ARP Request or Response sourced from a Router's IP address and indicating a Hardware Address other than the Router's should be logged.
- In the event that a Router's Hardware Address is changed, it should broadcast a gratuitous ARP reply advising the world of the event. However, on startup in a multiprotocol Router, this should NOT occur when the Router changes from its native address to its protocol-specific MAC address; instead, the Router should wait for the completion of the configuration sequence to send its initial ARP reply.

Internet Protocol Section

Routers 'must' implement options:

- see RFC 1009 in case we forgot something
- Loose Source Route and Record
- Strict Source Route and Record
- Record Route
- Standard Security Option ('must' be first in option list)
- Timestamp
- NOP
- End of List

Routers 'may' implement

- Extended Security Option
- Detection of combined or multiple Strict/Loose/Record Route
- MTU Options

Routers 'may' ('should not'?) implement obsolete IP options

- SATNET Stream ID
- Revised Security Option

Routers 'must not'

- Combine or multiply Strict/Loose/Record Route Options

There are some computational order dependencies:

- most options only make sense after the forwarding decision has been made.
- Strict and Loose Source Route apply BEFORE the forwarding decision, but only in systems addressed by the Destination IP Address.
- Routers must fragment traffic. There was some feeling that the Router should make the first n-1 fragments the size of the MTU, and let the last be the modulus, and some feeling that the fragments should all be approximately the same size (we learned later that the currently proposed MTU discovery mechanism requires the first choice - ed).
- Time to Live must be decremented on each hop. No vendor present decrements in seconds, but there was some feeling that decrementing by two when presenting traffic to a congested queue was doable and not all bad.
- Routers should recognize and correctly deal with all recognized broadcast addresses (Class D, 0.0.0.0, 255.255.255.255, Network or Subnet variants) at the same time; configuration parameters deal with what the Router emits, not what it recognizes.
- Routers 'must' not route traffic directed to a MAC broadcast or multicast address back to the same MAC. Routers 'should' not forward traffic directed to a MAC broadcast or multicast address at all.

Initial Writing Assignments

Fred Baker:	Address Resolution, ARP section
Art Berggreen:	Address Resolution, X.25/X.213 section
Steven Senum:	Hyperchannel, IP Options
John Hamner:	Time to Live
Stev Knowles:	Fragmentation and Re-assembly
Stev Knowles:	Treatment of IP Broadcast Addresses

Stev Knowles:	Glossary
John Veizades:	LocalTalk
Mike Ride, Jeff Burgan, Roxanne Streeter:	Internet IGPs, IP Filtering, IGP Translation
Steve Willis:	IEEE 802 LANs, Ethernet
Martin Gross:	ICMP
Philip Almquist:	Introduction
Bill Melohn and Fred Baker:	Serial Line Protocols
Jeff Burgan:	External Gateway Protocols
Jim Forster:	Variable Length Subnet Masks
Steve Willis:	SNMP
Philip Almquist:	Forwarding
Jim Forster:	Congestion Management

MINUTES - 2/8/90

Additional requirements discussed:

- Ignore reserved bits in IP packets
- Router should not require network services (like DNS) to boot (Jeff Burgan will write a section on this)
- Specify that if a Routing Protocol is implemented, it must be fully implemented (must follow appropriate RFC)
- Discuss multiple subnets (addresses) on an interface.
- Require SNMP for a router
- Performance requirements (like IS-IS) (Steve willis will write a section on this)

Discussion on Broadcast issues: Proposal to disallow old zero-host type broadcasts. Routing vendors present indicated that they would always provide an option to control this. Discussion on how to restrict forwarding broadcast packets, and packets that are ill-formed (i.e. subet of zero, host of zero) Directed broadcast, should allow, should have option to disable Some discussion on default setting of option. All subnets broadcast. Point made that no router seems to be doing this. (Steve Willis will gather info). Suggestion made to disallow generic case, maybe allow only for specific protocols (like bootp). Proposal made by Jim Forster to do "broadcast mapping" for resource location.

Variable length subnet masks Current rfc (1009) requires ability to specify subnet mask per interface. Comment made that this feature might be eliminated if support for multiple (sub)nets per interface was required. Suggestion made that a document describing usage of subnets be written (new working group)? Discussion on non-contiguous subnet masks (disallow)? Discussion on non-contiguous subnet networks (disallow)?

Filtering (both packet forwarding filtering, and routing protocol filtering) Good idea. Another working group? Require routing protocol filters for inter-AS routing. Option to turn source routing off. Martian address filtering.

ATTENDEES

Aronson, Cathy	cja@merit.edu
Bare, Ballard	bare%hprnd@hplabs.hp.com
Berggreen, Art	art@sage.acc.com
Bryant, Stewart	bryant@janus.enet.dec.com
Burgan, Jeffrey	jeff@nsipo.nasa.gov
Cavanaugh, John	john.cavanaugh@stpaul.ncr.com
Chatterjee, Samir K.	samir@nynexst.com
Clapp, George	meritec!clapp@bellcore.bellcore.com
Colella, Richard	colella@osi3.ncsl.nist.gov
Coltun, Rob	rcoltun@trantor.umd.edu
Deboo, Farokh	sun!iruucp!ntrlink!fjd
Easterday, Tom	tom@nisca.ircc.ohio-state.edu
Fernandez, Louis	lfernandez@bbn.com
Froyd, Stan	sfroyd@salt.acc.com
Gan, Der-Hwa	no email
Goguely, Herve	rvg@bridge2.3com.com
Gross, Martin	martin@protolaba.dca.mil
Heinanen, Julia	jh@funet.fi
Hua, Binh K.	no email
Hytry, Tom	tlh@iwlcs.att.com
Jacobson, Van	van@helios.ee.lbl.gov
Jagadeesh, BV	bvj@chamundi.esd.ecom.com
Jensen, Phil	jensen@fsu1.cc.fsu.edu
Karels, Mike	karels@berkeley.edu
Kaufman, David	dek@proteon.com

Knowles, Stev	stev@ftp.com
Marcinkevicz, Mike	mdm@gumby.dsd.trw.com
Merritt, Don	don@brl.mil
Miller, Dave	dtm@mitre.org
Mills, Cyndi	cmills@bbn.com
Monachello, Dave	dave@pluto.dss.com
Petry, Michael	petry@trantor.umd.edu
Pokorney, Dave	poke@nervm.merdc.ufl.edu
Reilly, Michael	reilly@nsl.dec.com
Rekhter, Yakov	yakov@ibm.com
Senum, Steve	sjs@network.com
Shibuyama, Steven	no email
Staw, Tony	staw@marvin.enet.dec.com
Streeter, Roxanne	streeter@nsipo.arc.nasa.gov
Veizades, John	veizades@apple.com
Ward, Carol	cward@spot.colorado.edu
Willis, Steven	swillis@wellfleet.com
Wobus, John	jmwobus@suvm.acs.syr.edu
Youssef, Mary	mary@ibm.com

3.4 Network Management Area

Director: Dave Crocker/DEC

Creations

A new working group, Internet Accounting, has been formed by Cindi Mills/ BBN. It has been chartered and will be tackling the ambitious and tricky issues of acquiring information for monitoring network usage. While accounting typically provides input for usage charging, the group will also consider, statistical and possibly security (pattern analysis) applications.

The formation of the Architecture group is still in progress. I am having difficulty finding an OSI-knowledgeable/architecture-oriented person. Other groups under consideration are a DECNET Phase IV MIB working group led by John Saperia/ DEC, and a group to formalize guidelines on how to write a MIB led by Jeff Case/ U-Tenn, Brian Handspicker/ DEC, and Lee LaBarre/ Mitre

Progressions

The Alert Management Working group has a stable draft, ready for initial publication. The Management Services Interface Group has a first draft ready. There is work in progress to align the OIM and MIB II variables. Transmission MIB is continuing with work on defining MIBs for the many physical media.

Completions

SNMP continues to demonstrate it's usefulness in operational networks. The SNMP, SMI and MIB I documents have been submitted to the IAB for consideration as Full Standards. MIB II has been submitted as a Proposed Standard.

After a full and productive life, the NocTools Working group has delivered its final draft for publication. Further work on the catalogue will be done by the Distribution and Awareness group (DAWG) Bob Stine will continue on as Interim Editor of the catalogue.

Integrations

In the future all new protocols should make provisions for publishing companion documents describing the associated MIB variables. The SNMP authentication is beginning work on an authentication MIB.

3.4.1 Alert Management Working Group (alertman)

CHARTER

Chairperson: Louis Steinberg/IBM, louiss@ibm.com

Mailing Lists:

alert-man@merit.edu
alert-man-request@merit.edu

Description of Working Group:

The Alert Management Working Group is chartered with defining and developing techniques to manage the flow of asynchronously generated information between a manager (NOC) and its remote managed entities. The output of this group should be fully compatible with the letter and spirit of SNMP (RFC 1067) and CMOT (RFC 1095).

Specific Objectives:

1. Develop, implement, and test protocols and mechanisms to prevent a managed entity from burdening a manager with an unreasonable amount of unexpected network management information. This will focus on controlling mechanisms once the information has been generated by a remote device.
2. Write an RFC detailing the above, including examples of its conforment use with both SNMP traps and CMOT events.
3. Develop, implement, and test mechanisms to prevent a managed entity from generating locally an excess of alerts to be controlled. This system will focus on how a protocol or MIB object might internally prevent itself from generating an unreasonable amount of information; examples of such techniques might include limiting number of alerts per time period, delayed reporting of "good news" (as in the link up sgmp trap on NSFNET), or the use of thresholds.
4. Write an RFC detailing the above. Since the implementation of these mechanisms is protocol dependent, the goal of this RFC would be to offer guidance only. It would request a status of "optional".

Estimated Timeframe for Completion:

A draft of the first RFC (alert flow control) will be written and reviewed by the July 1989 IETF meeting, with final review expected at the October 1989 IETF meeting. The second RFC draft will be submitted for initial review at the October 1989 IETF meeting. A date for final review of this document has not yet been determined.

CURRENT MEETING REPORT

Reported by Lee Oattes/University of Toronto

AGENDA

- Administrivia
  - Attendance list, someone to write minutes
- Review of Charter
  - Objectives, goals of each document
- Document 1: FINAL REVIEW (again!)
  - Behind Timeframe objectives
  - WILL submit to RFC process on 22 Feb. (Thursday)
  - DRAFT with minor revisions will be posted to mailing list for final comments THIS MONDAY. Comments will be accepted through 20 Feb.
  - Revisions to date: primarily syntax, semantics from the draft.01. Major change was moving examples from the statement of problem to an appendix.
  - My questions: Agreement on logged alert is ASN.1 OCTET STRING? No implication as to contents, and the goal is to avoid the dreaded OPAQUE. String might or might not be a construct (mine aren't right now :-). Since SNMP says primitives, we would be treating the OCTET STRING like an OPAQUE; the guy asking for it has to know it's intern format. Failure to reach closure TODAY means I will mak it a CHOICE of the two.
- OPEN FOR QUESTIONS AND COMMENTS ON DRAFT 1, but ...
  - Protocol implementers are encouraged to voice concerns on specific deviations from SNMP/CMOT specs
  - Document syntax, layout is still fair game
  - Document goals have been agreed on for over 6 months, and are no longer a topic for debate (closure was reached).
  - Criticisms of mechanisms must be backed up with implementations (eg. I tried this, and it cost too much, didn't work, etc).

MINUTES

Meeting opened with Lou discussing the two documents that the WG is chartered with writing. The first deals with mechanisms that systems should use to control the flow of alert information. It is being authored by Lou. The second, written by John Cook, details specific techniques that an implementer of alerts may wish to consider. Its focus is on how alerts are generated.

Final Review of document 1

Logged alerts should be of the SMI type OPAQUE.

Lou was hesitant to use this, as OPAQUES use has been avoided. However, the SNMP constraint of not using construct ASN.1 types ruled out the option of encapsulating an SNMP trap in an OCTET STRING. The use of a simple OCTET STRING would be a violation of ASN.1, and would not result in further parsing of the trap.

A review of the DRAFT event and logging tutorial by the OIM Working Group followed. Lou agreed to contact Lee LaBarre about resolving any potential conflicts with the alert logging. The only problem centered on the Alert-Man requirement that a full log wrap. OIM shows this wrapping behavior as optional; halting behavior must be supported. The Alert-Man requirements might be presented as a subset of the OSI logging function. No problems were found with the OIM view of feedback/pin.

The latest set of changes to the DRAFT will be posted on the 11th of this month. Comments will be accepted through the 22nd. At that time, the DRAFT will be submitted to the RFC process, as a protocol standard. Its initial, requested status will be elective.

2nd Document

Lou will again post his sample technique/format to the mailing list. He also agreed to write a short overview of CMOT events that Brian will review for the second document. John discussed the format and status of the second document. Implementation reviews are being solicited for the following techniques:

- threshold hysteresis (time and value based)
- snapshots
- thresholds on exceptions
- thresholds built on counters, gauges, tidemarks
- sliding window "pins" on each threshold
- adaptive thresholds

We are also looking for any other techniques in use. Submissions should represent actual implementations (can be in-house code), and can be posted to the working group list in the format Lou is using.

John will post the DRAFT in its current state shortly.

ATTENDEES

Aronson, Cathy	cja@merit.edu
Cook, John	cook@chipcom.com
Feridun, Metin	mferidun@bbn.com
Handspicker, Brian	bd@vines.dec.com
Minshall, Greg	minshall@kinetics.kinetics.com
Newkerk, Oscar	newkerk@decwet.dec.com
Oattes, Lee	oattes@utcs.utoronto.ca
Perkins, David	dave_perkins@3com.com
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Sheridan, Jim	jsherida@ibm.com
Waldbusser, Steve	sw01@andrew.cmu.edu
Wittbrodt, Dave	dmw@cisco.com

3.4.2 Internet Accounting (acct)

CHARTER

Chairperson: Cyndi Mills/BBN Communications, cmills@bbn.com

Mailing List:

accounting-wg@bbn.com
accounting-wg-request@bbn.com: to join list

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

Specific Objectives:

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 / 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.

MAJOR MILESTONES:

- May 1990 Policy Models examined.
- Aug 1990 Meter Working Draft written.
- Nov 1990 Meter Revised Draft reviewed.
- Collection Protocol Working Papers written.
- Feb 1991 Meter Final Draft submitted. Collection Protocol Working Papers reviewed.
- May 1991 Collection Protocol Recommendation.

FIRST MEETING:

February 1990

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

AGENDA

- Bounding the Charter.
- Form a Working Group
- Requirements Discussion: Draft Minutes below.

Minutes

1. Summary
   Agreed to form an Internet Accounting working group. Cyndi Mills will chair it and write the charter. This working group is in the Network Management Area under Dave Crocker.
2. Bounding the Charter:
   We need to examine a wide range of policies to understand what set of information is required to satisfy the billing and reporting requirements, bearing in mind realistic requirements and restrictions regarding:
   - Availability of Information,
   - Performance, and
   - Accuracy.

   Policy Disclaimer: Neither issues surrounding how policies are set nor how they are formulated will be addressed by this group.
   2.1 OSI Accounting
   Brian Handspicker, ANSI X3T5.4 OSI Management Accounting Ad Hoc Group Leader, presented the OSI view of accounting. The OSI Accounting working group is defining the collection service and protocols. The OSI group is not addressing the content information to be measured and reported by the collection service. Suggest that the IETF working group coordinate with the OSI accounting group so as not to duplicate effort.
   Meter <--> Collector <--> Application
   Application: The application manipulates the billing data in accordance with policy, and determines which information will be requested from the metering devices.
   Collector: The collector is responsible for integrity and security of the data during transport from the meter to the application.
   Meters: Meters perform the measurement and aggregate the results. The characteristics of the meter may be implementation-specific.
   2.2 Data Generation vs. Data Collection vs. Billing Application

The generation of accounting data (the meter function) is the focus of this IETF group. First, we need to determine what information will satisfy the widest possible range of policies, and what the constraints are. Secondly, we should cover local storage and aggregation techniques.

Data collection protocols, i.e. methods for carrying accounting data, are under development in ANSI. Accounting data may be carried by a combination of protocols, including network management protocols such as OSI Accounting, SNMP, CMIP. The selection of collection protocol(s) should be deferred until the structure and constraints of the carried data are known.

The billing process, i.e. the processing of the accounting data, is beyond the scope of this group. Billing methods, tariffs, and exceptions tend to be unique to each organization.

2.3 Network-Level vs. Host-Level

The information available to the meter depends on its location in the network. One of the major issues here is attribution - with what granularity can we account for the source and destination of network traffic? Can we track the source/destination of a packet to the autonomous network, the network number, the host address, the user, or to a charge number on one of a user's many projects?

For network meters, a function attached to the routers, this information is limited to what can be extracted from the IP packet flow. Various counters may be implemented, but attribution of the packet to a source is limited to the information available in the IP address (and the protocol ID of the protocol carried). There is no unique identifier in the packet for a user.

Host meters are more flexible. They have direct knowledge of the user and his operation, and are in a position to implement user-level accounting in accordance with the behavior of a specific operating system.

This working group will concentrate on network-level meters. The discussion section covers a number of background arguments for this restriction.

3. Discussion

The Internet community is made up of:

- Network providers, e.g. backbone and regional networks, who usually own the transmission media, regulate or own the routers, but disown the hosts. Internet accounting is for the benefit of the network provider, an aid in the implementation of the network provider's policy. In networks with chargeback policies, accounting may be the sole source of funding for the network.
- Network users, e.g. hosts, individual users, and projects. These are the consumers of network services. From an accounting point of view, these are the end-users, the finest granularity of attribution.
- Stuck in the middle. These are the entities that are both providers and consumers of network services. Hosts and regional networks are frequently in

this category. They receive service from the network and provide network service to the user. In addition to compensating other network providers for network services rendered, they must assist in allocating responsibility for those services received and provided to end-users.

The phone company analogy was used frequently to illustrate several interrelated points.

- Regional/Local Operating Providers: The Bell Operating Companies (BOCs) serve as the network connection point for subscribers. They maintain directories and connectivity information, because they control the end-users' connections.
- Long-Distance Providers: AT&T and MCI are backbone services.
- PBX Installations: A subscriber may be a single telephone, or a private telephone network. The private telephone network is analogous to the LAN: it receives a bulk bill from the regional BOC and it is responsible for maintaining its own records to allocate costs back to its local users.

The potential billing models between a long-distance provider and a middleman (BOC) provider in the phone company model illustrated some of the issues.

Under the existing policy, the BOCs bill users for long-distance services as a courtesy to the long- distance companies, who set the rates. Two hypothetical models for implementing this service were discussed.

The long-distance company provides per-call detail to the BOC. The BOC maintains the accounting data and the association of usage data with its end-users. The BOC generates the bill.

The BOC provides per-call "tags" to identify its end users to the long-distance provider. The long- distance carrier maintains the accounting data and the association of usage data with those tags. The long- distance carrier generates the bills' contents. The BOC simply forwards the bill to the user associated with its "tags".

Under a hypothetical policy, BOCs receive an aggregate bill for long-distance services from the backbone provider. The BOC is treated as a single billable entity by the long-distance service. In this case, the BOC is solely responsible for maintaining the accounting data and policies which allocate those costs to users. The BOC provides no user-level information to the backbone provider, nor does the backbone provider give detailed per-call accounting to the BOC. (Not interactively, at least.)

DEFINE THE BILLABLE ENTITY FIRST. We are examining the nature of traffic, interesting but too much for simple accounting purposes. Start with the definition of the "billable entity" and build up to what you need.

DON'T INCLUDE NETWORK DESIGN AND ANALYSIS DATA. Accounting needs very precise data about certain kinds of traffic. Network design and analysis needs different data, and frequently works with sampling techniques inappropriate for accounting. Although much of the accounting information may be useful for network design and analysis, covering network design and analysis requirements will overburden the scope of this group.

NEED TO KEEP THE ENTITY MATRIX SIMPLE. There are inherent limits in the current situation. Routers can't handle keeping a matrix of counters for every possible user-user combination. Some kind of hierarchical billing is required. One division is for hosts to be billed in aggregate by the network, and leave the hosts responsible for allocating costs to users. However even host-host matrices can get very large. If each datagram entering a router is on a different source- destination pair, thrashing could be easily induced.

WATCH OUT FOR OVERHEAD. Accounting for every packet in a fine- grain way could result in 100may have more than 50it can be appropriately attributed to users. 50off point for feasibility.

DIFFERENT ALGORITHMS FOR LOCAL AND LONG-DISTANCE SERVICES: Note that the phone company uses different algorithms for local and long distance services. Long distance calls are handled with detailed call accounting or aggregate counts (message units). Local calls are handled with simple aggregate counts (message units) or flat fees regardless of usage.

The lesson here is that where the cost of accounting is huge in comparison with the cost of providing the basic service, many subscribers prefer a policy which allocates usage as a flat fee. Some subscribers, however, (message units), still want usage-based fees. Phone companies provide a wide variety of such combinations of service.

WHAT ABOUT SPECIAL END-USERS? Suppose I am a long-distance carrier and I want a particular research group to get a special rate. In the various models how can I ensure that their traffic and only their traffic is billed at the reduced rate? How do government clients get a bulk rate?

We need to consider the interaction between government and commercial entities, e.g., what does GEC Marconi do when it wants to communicate with NASA on commercial issues?

NEED A SET OF TEST QUESTIONS FOR PRELIMINARY VERIFICATION OF THE MODEL. What is an accountable unit? Examples of questions that should be answered are how to deal with rate periods (time-of-day), special end-users, etc. Need many more questions.

ON FORMING A WORKING GROUP: We will see commercial services in the Internet. This will require accounting. The IETF should get the process set up first. Good value for traffic and capacity planning, as well. Suggest we talk to people who are planning to offer commercial Internet service (PSI, UUNET, Finnish PTT, SMDS) to see what kinds of charging strategies they use. The RACE program, with Ira Richer, is also working on accounting issues.

ATTENDEES

Cerf, Vinton	vcerf@nri.reston.va.us
Crocker, Dave	dcrocker@nsl.dec.com
Crocker, Steve	crocker@tis.com
Fernandez, Louis	lfernandez@bbn.com
Handspicker, Brian D.	bd@vines.dec.com
Kirstein, Peter	kirstein@cs.ucl.ac.uk
Lazear, Walter	lazear@gateway.mitre.org
Little, Mike	little@saic.com
Morris, Dennis	morris@imo-uvax.dca.mil
Newkerk, Oscar	newkerk@decwet.dec.com
Pace, Donald	pace@fsu1.cc.fsu.edu
Saperia, Jon	saperia%tcpjon@decwrl.dec.com
Su, Zaw-Sing	zsu@tsca.istc.sri.com
Youssef, Mary	mary@ibm.com
Yuan, Aileen	aileen@gateway.mitre.org

3.4.3 LAN Manager Working Group (lanman)

CHARTER

Chairperson: Jim Gruel/HP, jimg@hpcndpc.cnd.hp.com

Mailing List: lanmanwg@spam.istc.sri.com

Description of Working Group:

To define and maintain the MIB and relevant related mechanisms needed to allow management overlap between the workgroup environment (LAN Manager based) and the enterprise environment (based on TCP/IP management).

Specific Objectives:

This translates into three basic objectives:

- Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
- Propose extensions to the TCP/SMI when appropriate.
- Develop requirements for additional network management information, as needed, and work to extend the LAN Manager interfaces to support such information.

Estimated Timeframe for Completion:

Objective 1: Version 1 of the LANMAN MIB has been completed and is awaiting consideration by the RFC editor (two RFCs have been proposed: LANMAN-MIB for "conventional" objects, and LANMAN-MIB-EXPER for objects related to LAN Manager alert handling). Subsequent versions will be worked on as necessary after further experience is gained with version 1. There is no definite timeframe set for work on version 2.

Objective 2: No extensions to the SMI have been proposed, and there are no immediate plans for making such a proposal.

Objective 3: No modifications to the LAN Manager interfaces were required for version 1 of the LANMAN MIB. This issue will be reconsidered after further experience is gained with version 1.

CURRENT MEETING REPORT

Did not meet.

3.4.4 Management Services Interface Working Group (msi)

CHARTER

Chairpersons: Oscar Newkerk/DEC and
Sudhanshu Verma/HP

Mailing List: msiwg@decwrl.dec.com
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).

Specific Objectives:

1. Determine the feasibility of a common interface across multiple management protocols.
2. Define the requirements for such an interface.
3. Define an architectural framework for such a service interface.
4. Define a specification that satisfies the architectural requirements.
5. Implement one or more prototypes of the interface.
6. Advance an RFC based on the specification and prototype experience.

Milestones:

Feb 1990 — Initial version of the Internet draft placed in the Internet Drafts directory.

May 1990 — Revised version of the draft from editing meetings placed in the Internet Drafts directory.

Aug 1990 — Initial implementation of the prototype available for test.

Dec 1990 — Revised draft based on the implementation experience submitted to the RFC editor.

CURRENT MEETING REPORT

Reported by Oscar Newkerk/DEC

MINUTES

A proposed draft API was presented by Oscar Newkerk followed by a question and answer session. Issues raised during the presentation were:

There seems to be a requirement for a set of services that are not addressed by the draft. This set was referred to as "MIB Services" and characterized as providing online access to the object and attribute definitions in the MIB documents. It was felt that this type of information would be required to allow the API to translate an operation against an object in a MIB into the appropriate parameters for the protocol that was being used to encode the operation. This work will be evaluated by the MSI working group either for inclusion in the API draft or as a separate document.

There is a section in the draft API document that deals with alert handling services. This needs to be evaluated in light of the output from the Alertman working group.

The draft API will be reformatted and submitted to the internet-drafts directory as an Internet draft, with the addition of a note that it is an interim draft and will be reworked by the MSI working group.

There were two other people at the meeting that expressed interest in working on the editing of the draft API. These were Sudhanshu Verma/HP and Dave Perkins/3COM. An editing meeting will be arranged for sometime early in March 1990.

In addition, Sudhanshu Verma volunteered to become the co-chair for the MSI working group.

ATTENDEES

Engineer, Hunaid	hunaid@opus.cray.com
Froyd, Stan	sfroyd@salt.acc.com
Glappa, Adrianne	no email
Handspicker, Brian D.	bd@vines.dec.com
Hunter, Steven	hunter@ccc.mfecc.arpa
Hytry, Tom	tlh@iwlcs.att.com
Mills, Cyndi	cmills@bbn.com
Minshall, Greg	minshall@kinetics.kinetics.com
Perkins, Dave	dave_perkins@3com.com
Robertson, Jim	jar@bridge2.3com.com
Saperia, Jon	saperia%tcpjon@decwrl.dec.com
Stine, Bob	stine@sparta.com
Verma, Sudhanshu	verma@hpindbu.hp.com
Waldbusser, Steve	sw01@andrew.cmu.edu
Woodburn, Robert	woody@saic.com
Yasaki, Brian	bky@twg.com

MSI Model

- Uses an *association object* to configure the 'connection' with the managed object's agent.
- Uses directive services to issue management operations requests to the agent and to receive replies to confirmed requests.
- Uses event subscription services to receive events.
- Uses the Management Operations Support Services (MOSS) to support both events and directives.

Slide No. 1

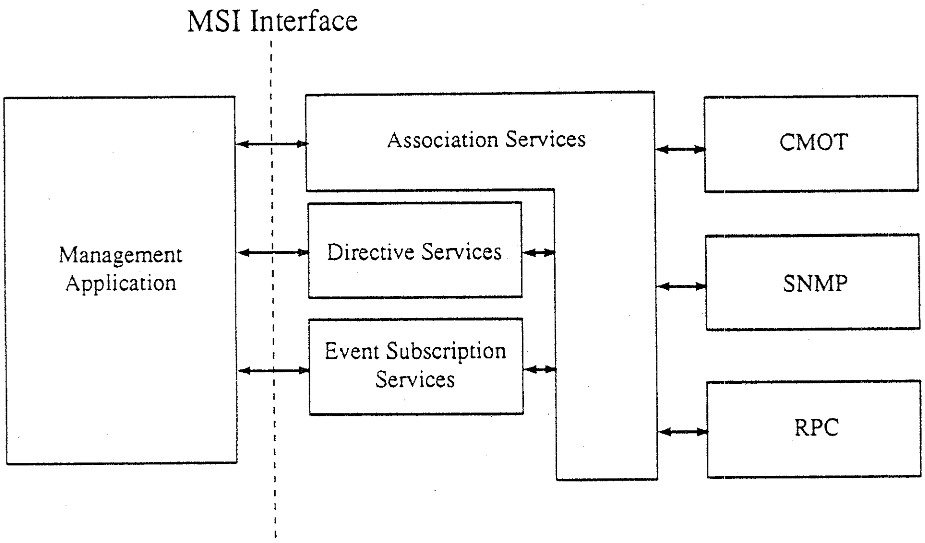

Slide No. 2

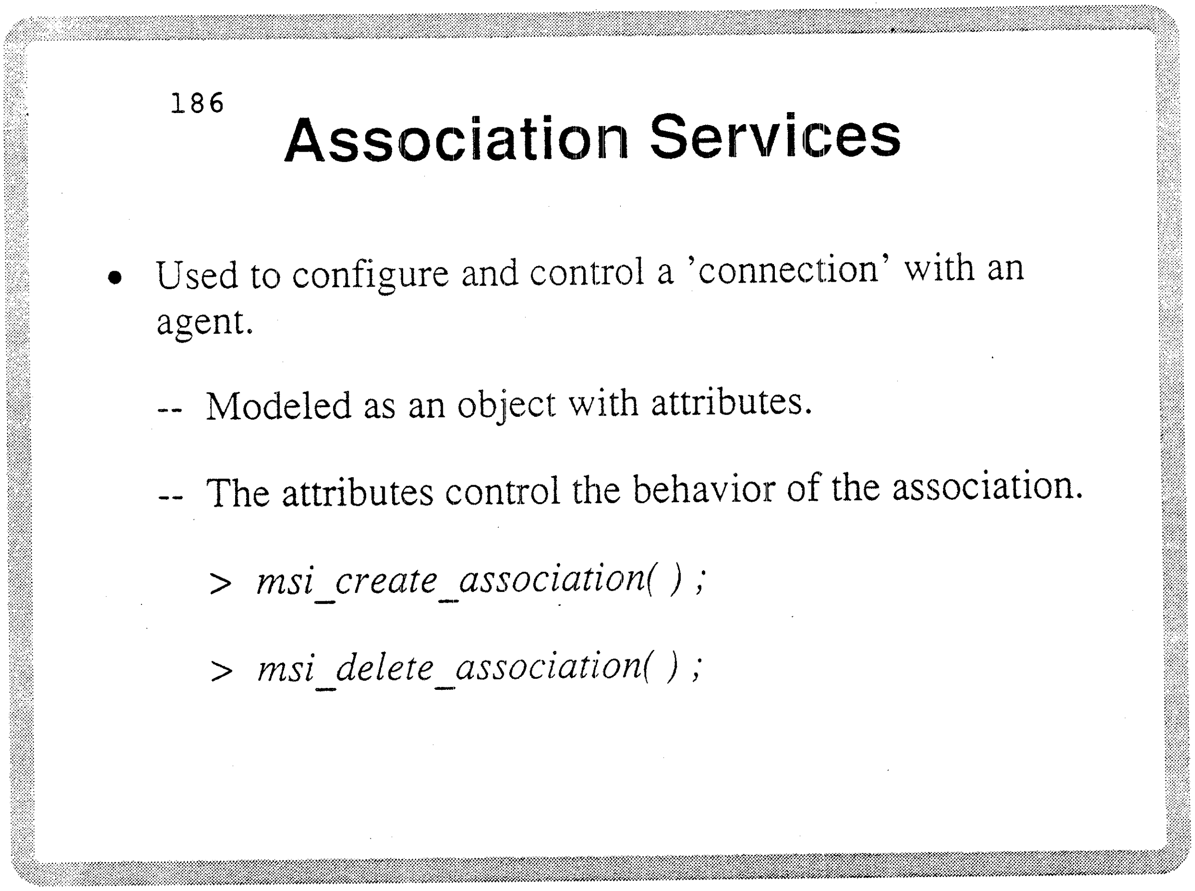

Association Services

- Used to configure and control a 'connection' with an agent.

  -- Modeled as an object with attributes.

  -- The attributes control the behavior of the association.

  > *msi_create_association( ) ;*

  > *msi_delete_association( ) ;*

Slide No. 3

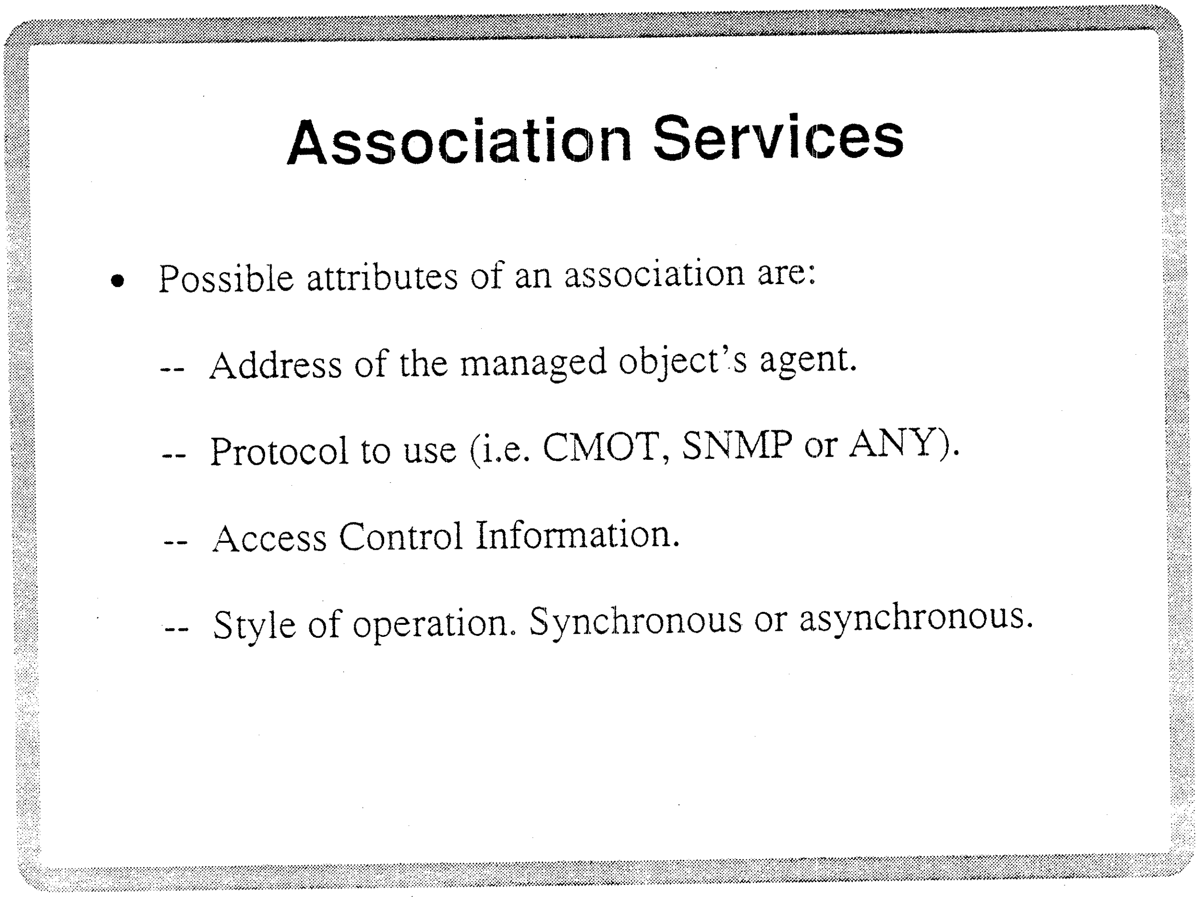

Association Services

- Possible attributes of an association are:

  -- Address of the managed object's agent.

  -- Protocol to use (i.e. CMOT, SNMP or ANY).

  -- Access Control Information.

  -- Style of operation. Synchronous or asynchronous.

Slide No. 4

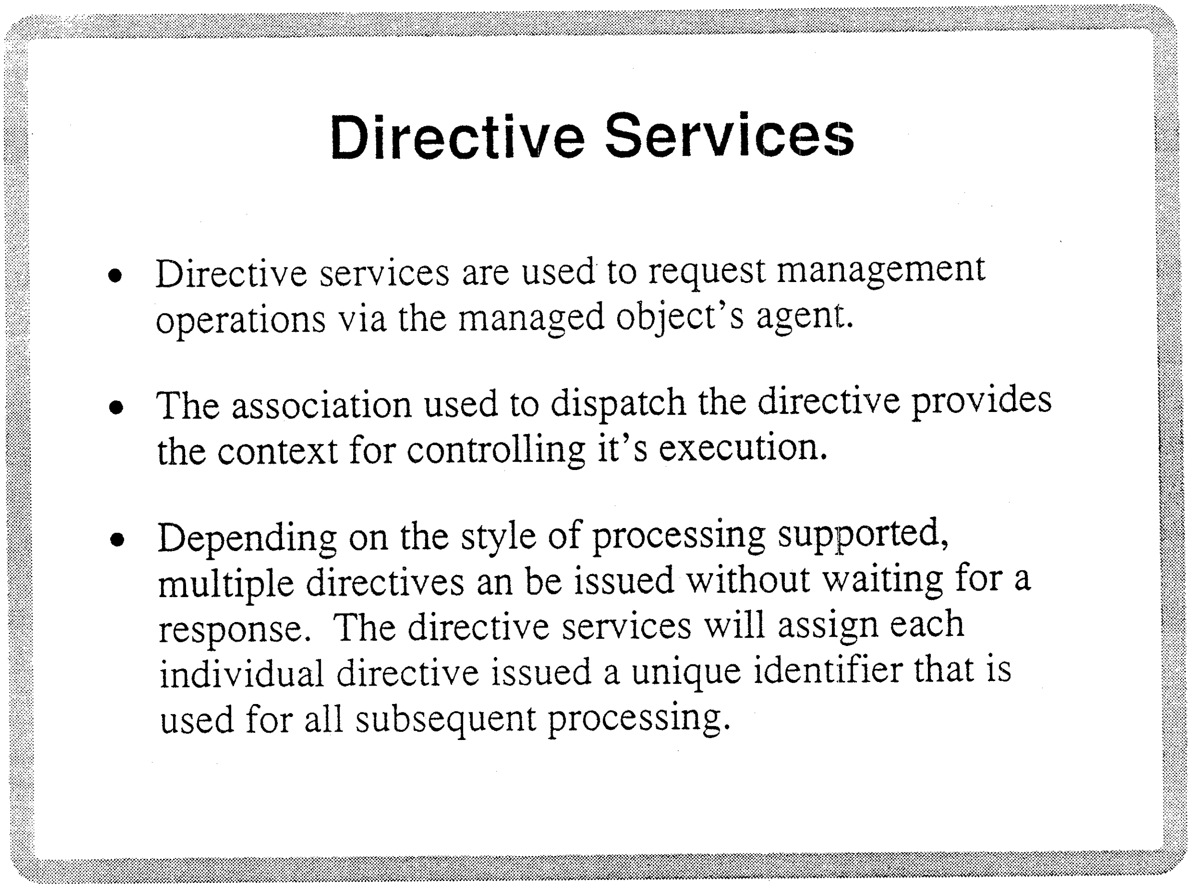

Directive Services

- Directive services are used to request management operations via the managed object's agent.
- The association used to dispatch the directive provides the context for controlling it's execution.
- Depending on the style of processing supported, multiple directives an be issued without waiting for a response. The directive services will assign each individual directive issued a unique identifier that is used for all subsequent processing.

Slide No. 5

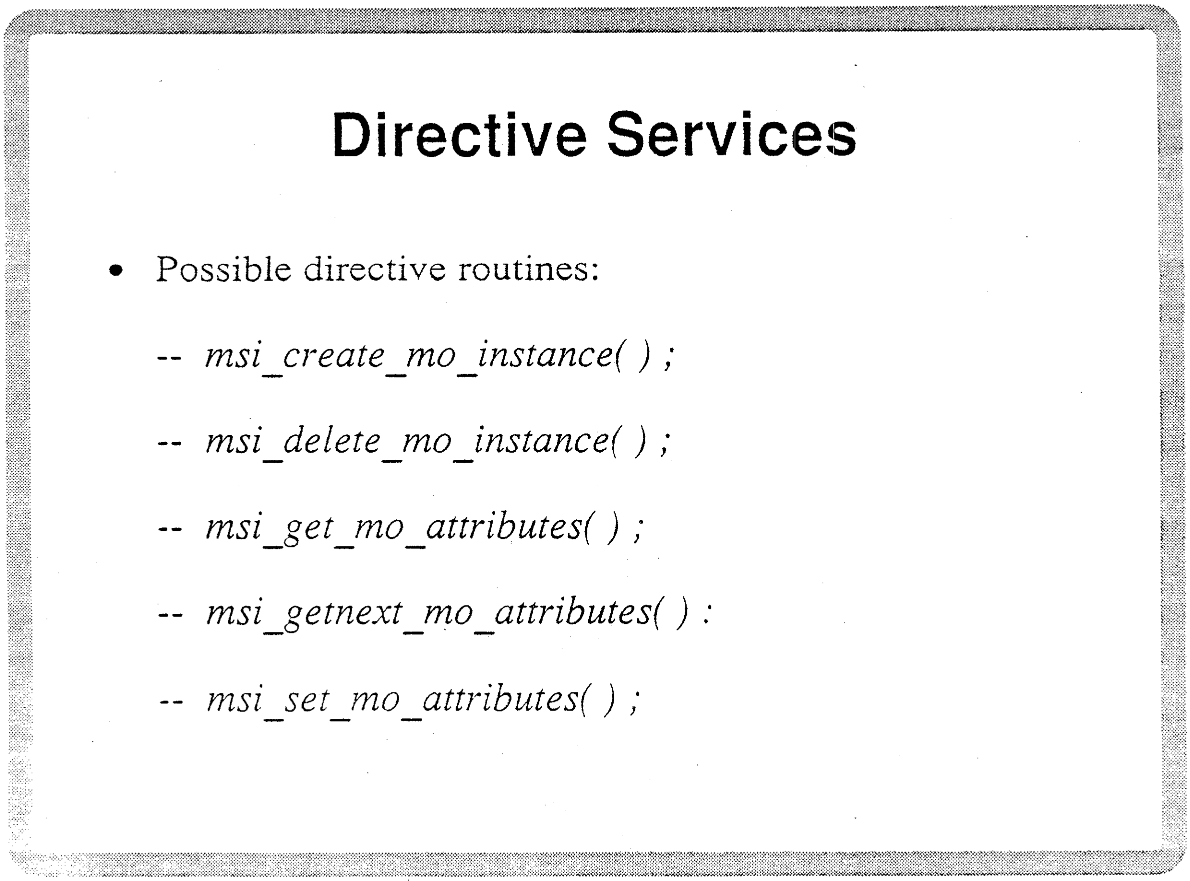

Directive Services

- Possible directive routines:

-- *msi_create_mo_instance( ) ;*

-- *msi_delete_mo_instance( ) ;*

-- *msi_get_mo_attributes( ) ;*

-- *msi_getnext_mo_attributes( ) :*

-- *msi_set_mo_attributes( ) ;*

Slide No. 6

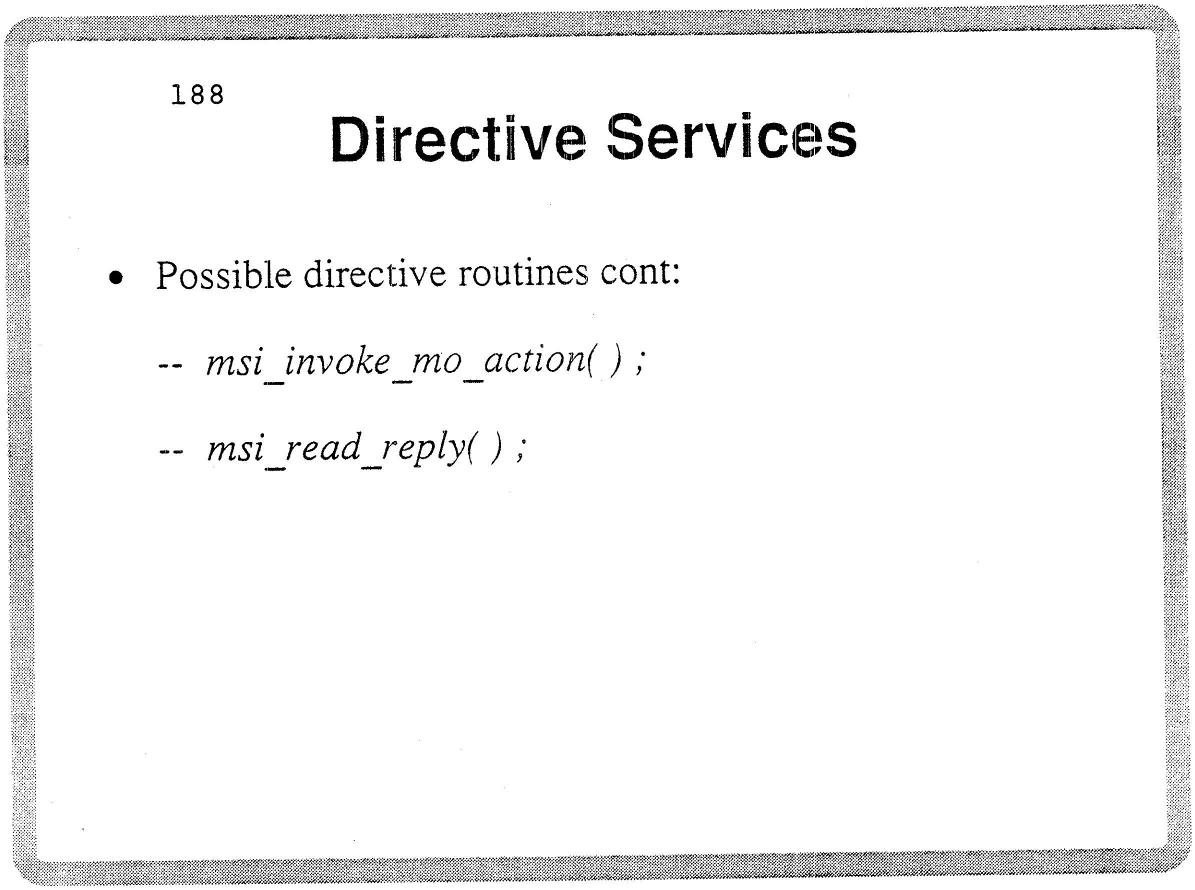

Directive Services

- Possible directive routines cont:
  - -- *msi_invoke_mo_action( ) ;*
  - -- *msi_read_reply( ) ;*

Slide No. 7

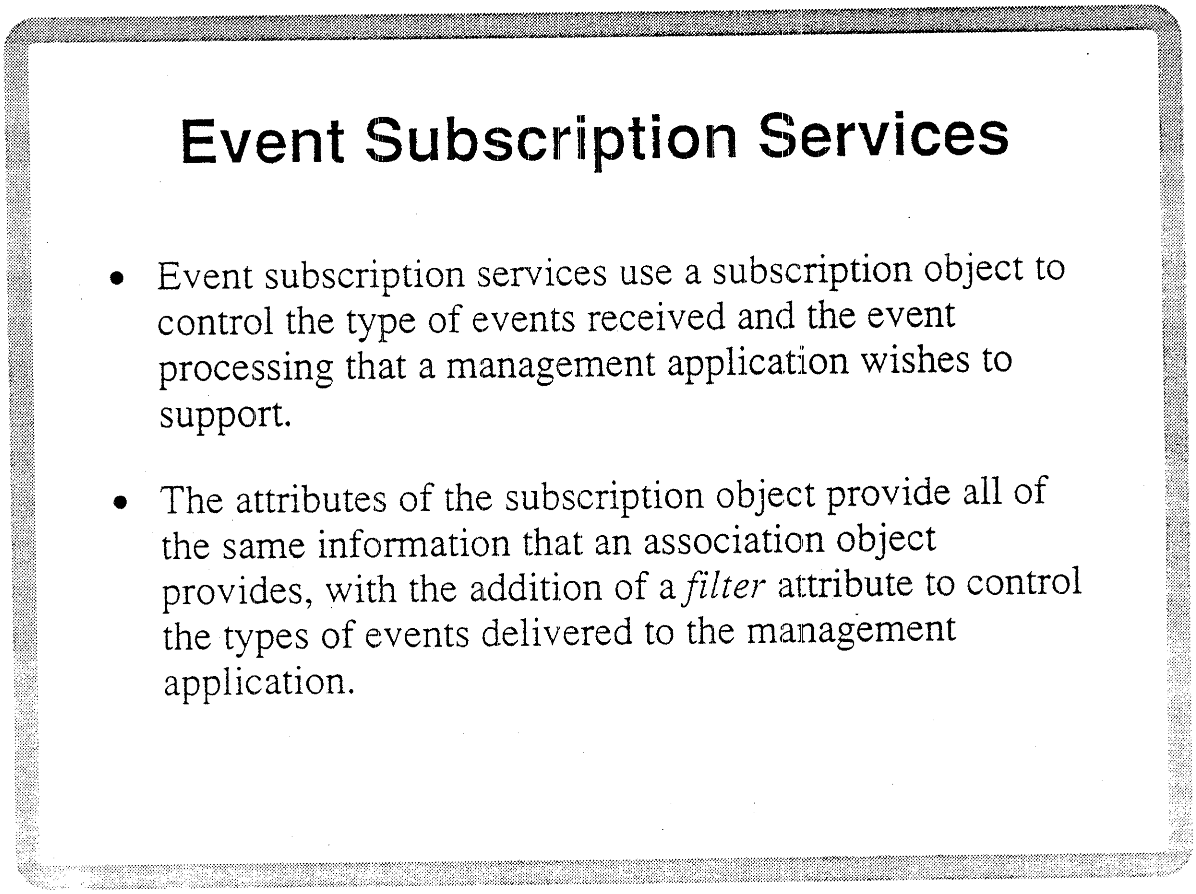

Event Subscription Services

- Event subscription services use a subscription object to control the type of events received and the event processing that a management application wishes to support.
- The attributes of the subscription object provide all of the same information that an association object provides, with the addition of a *filter* attribute to control the types of events delivered to the management application.

Slide No. 8

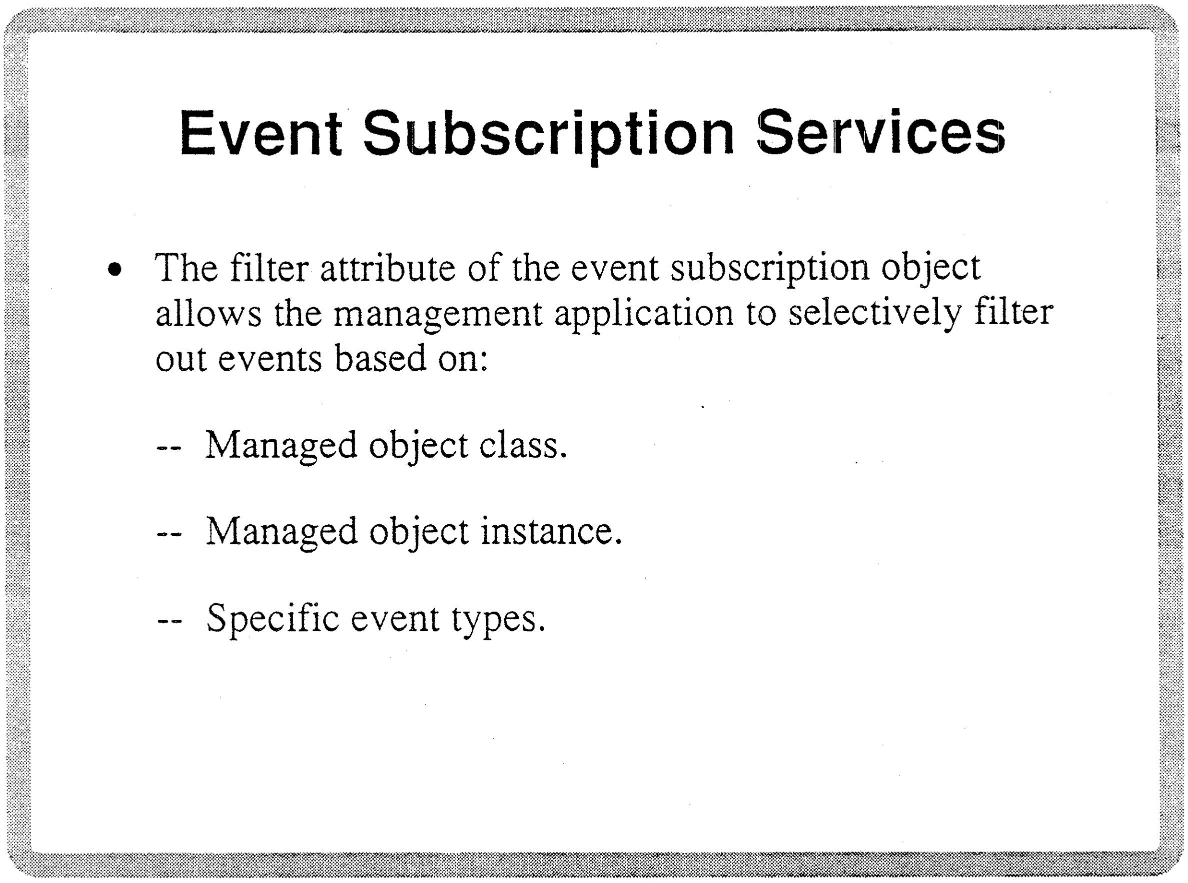

Event Subscription Services

- The filter attribute of the event subscription object allows the management application to selectively filter out events based on:

  -- Managed object class.

  -- Managed object instance.

  -- Specific event types.

Slide No. 9

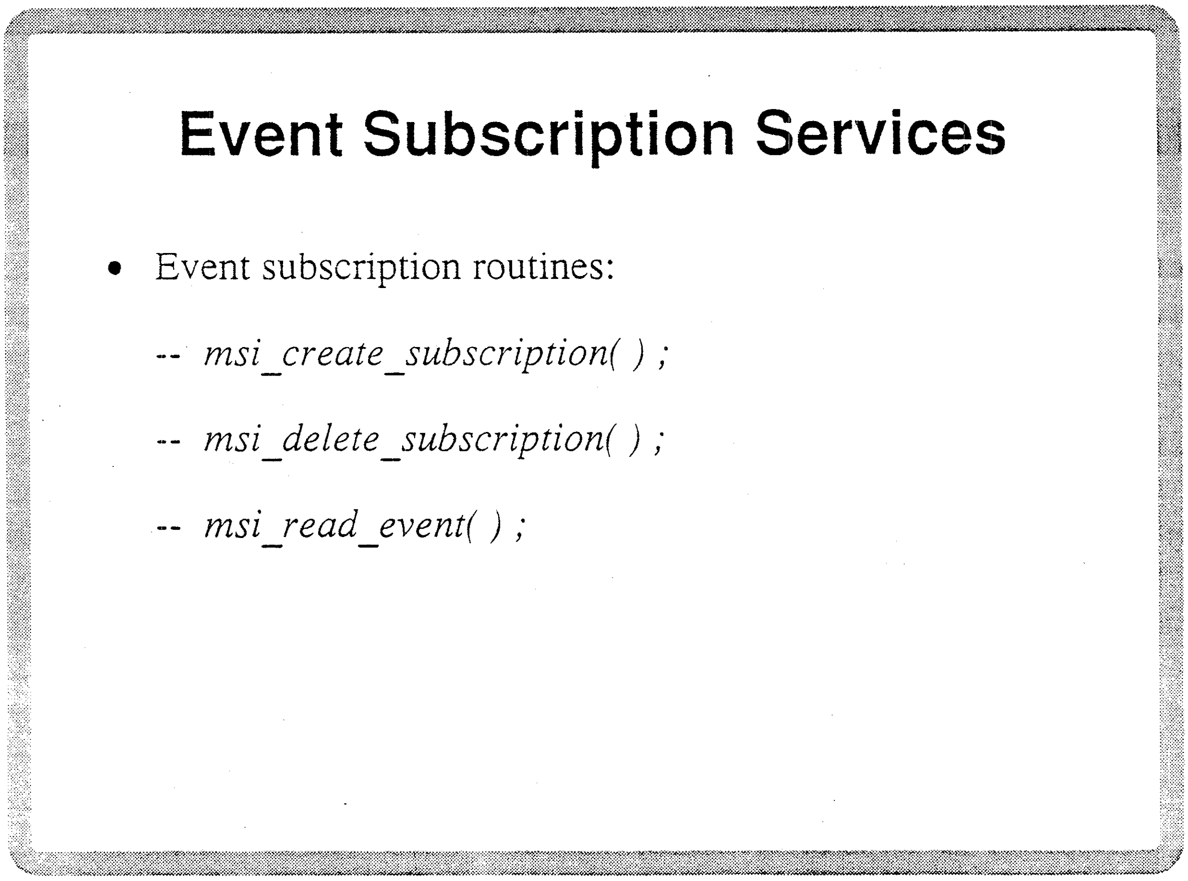

Event Subscription Services

- Event subscription routines:

  -- *msi_create_subscription( ) ;*

  -- *msi_delete_subscription( ) ;*

  -- *msi_read_event( ) ;*

Slide No. 10

Management Operation Support Services

- The management operation support services (MOSS) are a set of utility routines used to support the association, directive and event services. The routines are used for:

  -- Creating and manipulating attribute lists.

  -- Creating and manipulating object identifiers.

  -- Creating and manipulating directive filters.

  -- Creating and manipulating event filters.

Management Operation Support Services

- The attribute list data type (AVL) is used to support the creation and processing of variable length lists of variable length elements.
- The AVL routines allow the management application to process these lists without knowledge of their internal structure.
- All operations on the AVL type, including creation and deletion, are accomplished using MOSS procedures.

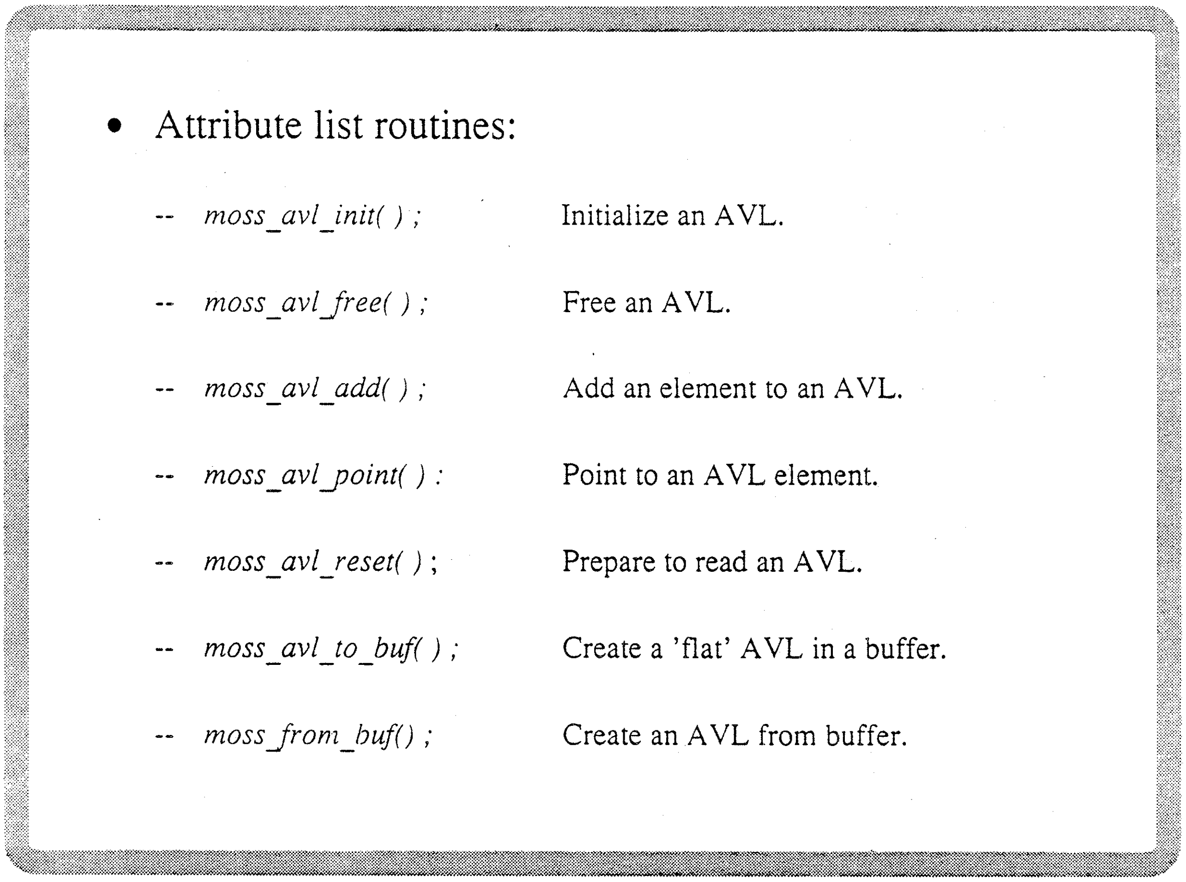

- Attribute list routines:

-- *moss_avl_init( ) ;*	Initialize an AVL.
-- *moss_avl_free( ) ;*	Free an AVL.
-- *moss_avl_add( ) ;*	Add an element to an AVL.
-- *moss_avl_point( ) :*	Point to an AVL element.
-- *moss_avl_reset( )* ;	Prepare to read an AVL.
-- *moss_avl_to_buf( ) ;*	Create a 'flat' AVL in a buffer.
-- *moss_from_buf() ;*	Create an AVL from buffer.

Slide No. 13

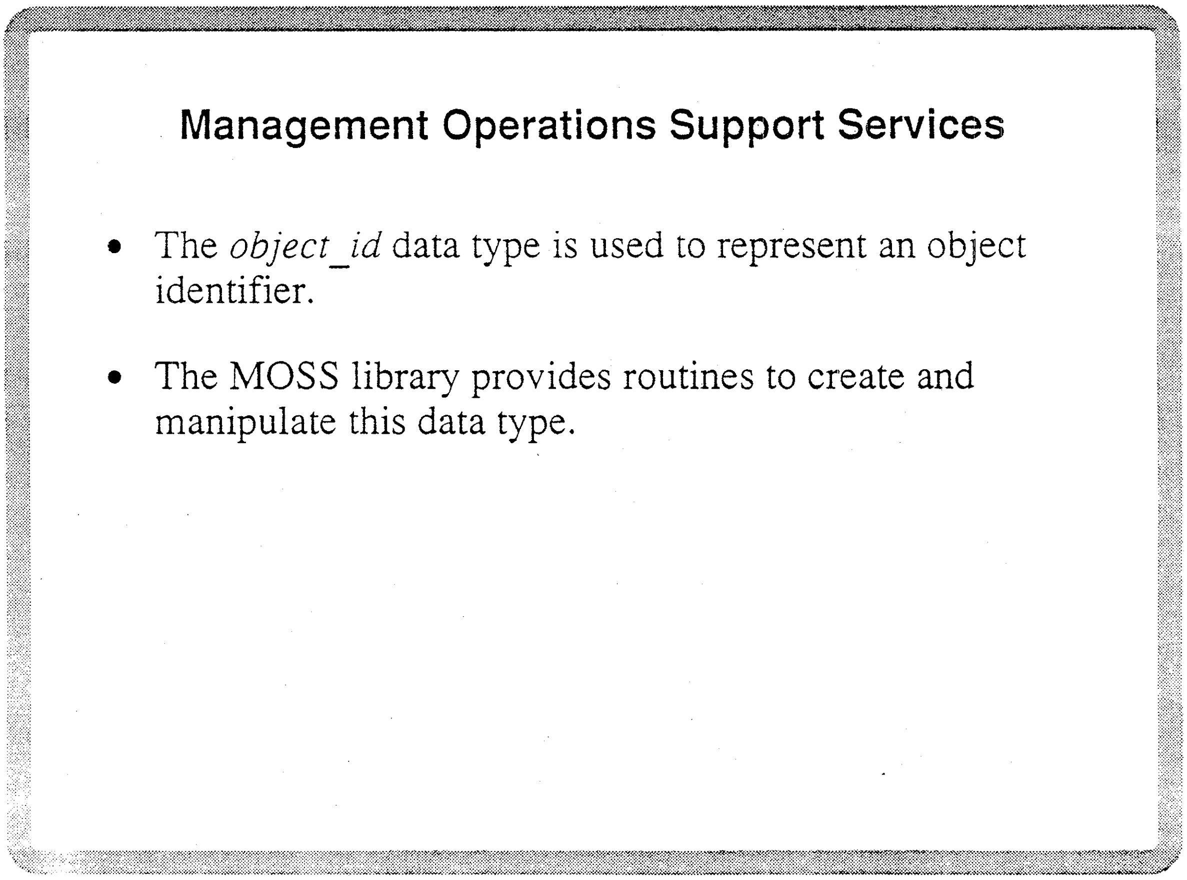

Management Operations Support Services

- The *object_id* data type is used to represent an object identifier.
- The MOSS library provides routines to create and manipulate this data type.

Slide No. 14

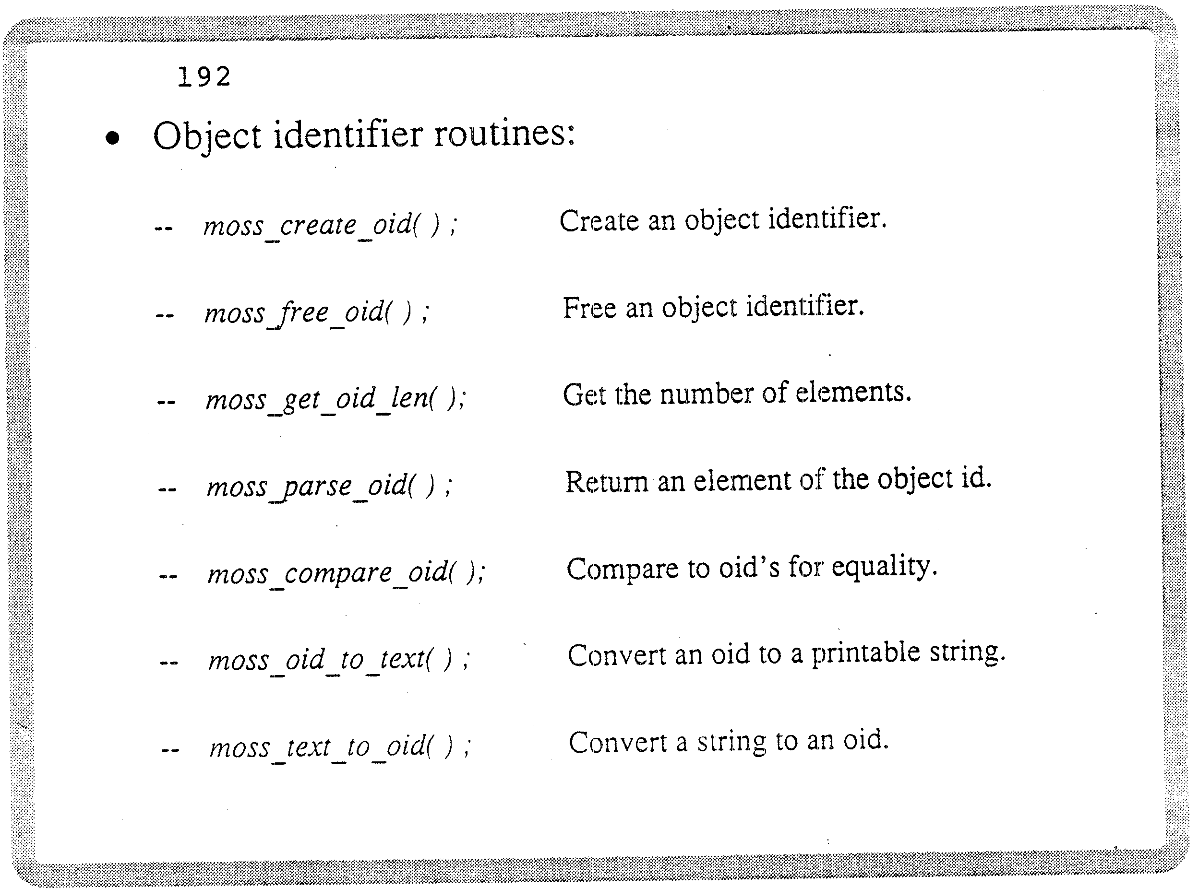

- Object identifier routines:

-- *moss_create_oid( ) ;*	Create an object identifier.
-- *moss_free_oid( ) ;*	Free an object identifier.
-- *moss_get_oid_len( );*	Get the number of elements.
-- *moss_parse_oid( ) ;*	Return an element of the object id.
-- *moss_compare_oid( );*	Compare to oid's for equality.
-- *moss_oid_to_text( ) ;*	Convert an oid to a printable string.
-- *moss_text_to_oid( ) ;*	Convert a string to an oid.

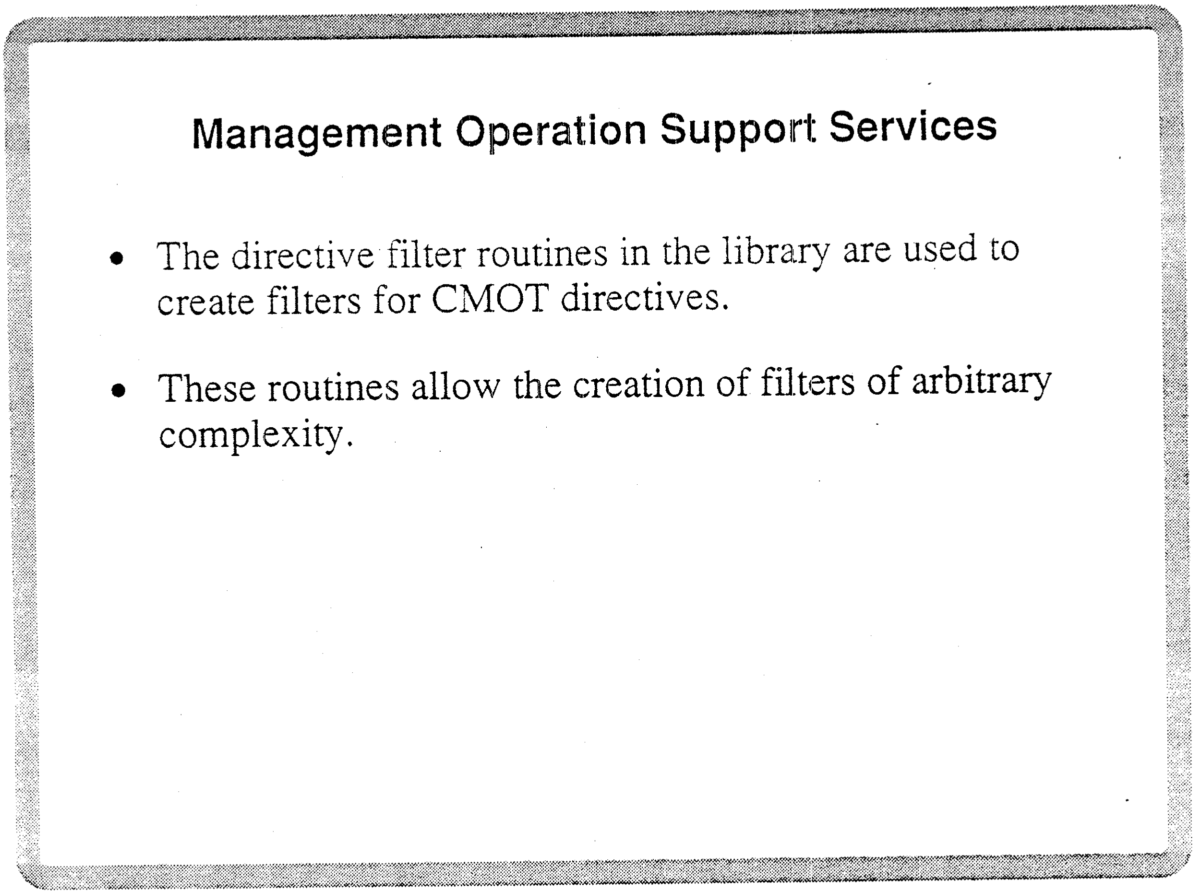

Management Operation Support Services

- The directive filter routines in the library are used to create filters for CMOT directives.
- These routines allow the creation of filters of arbitrary complexity.

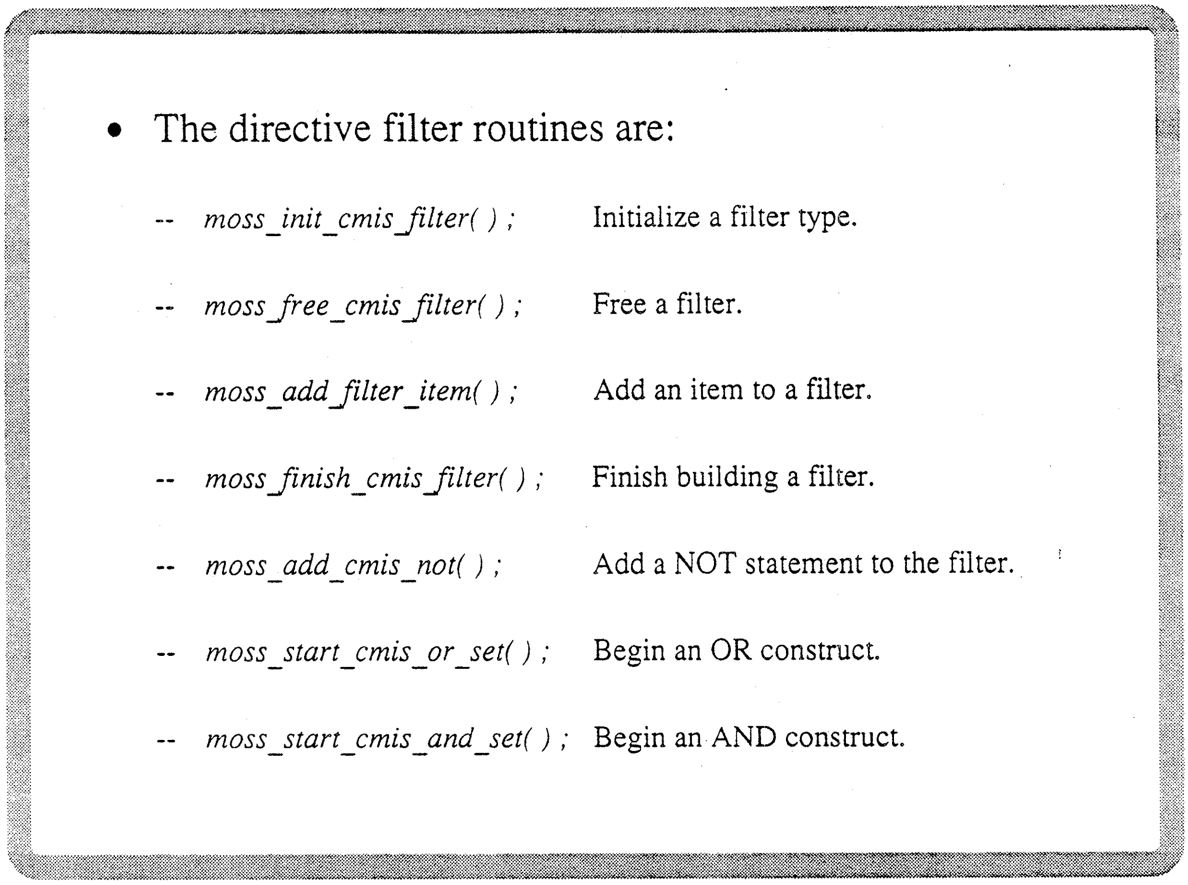

- The directive filter routines are:

-- *moss_init_cmis_filter( ) ;*	Initialize a filter type.
-- *moss_free_cmis_filter( ) ;*	Free a filter.
-- *moss_add_filter_item( ) ;*	Add an item to a filter.
-- *moss_finish_cmis_filter( ) ;*	Finish building a filter.
-- *moss_add_cmis_not( ) ;*	Add a NOT statement to the filter.
-- *moss_start_cmis_or_set( ) ;*	Begin an OR construct.
-- *moss_start_cmis_and_set( ) ;*	Begin an AND construct.

Slide No. 17

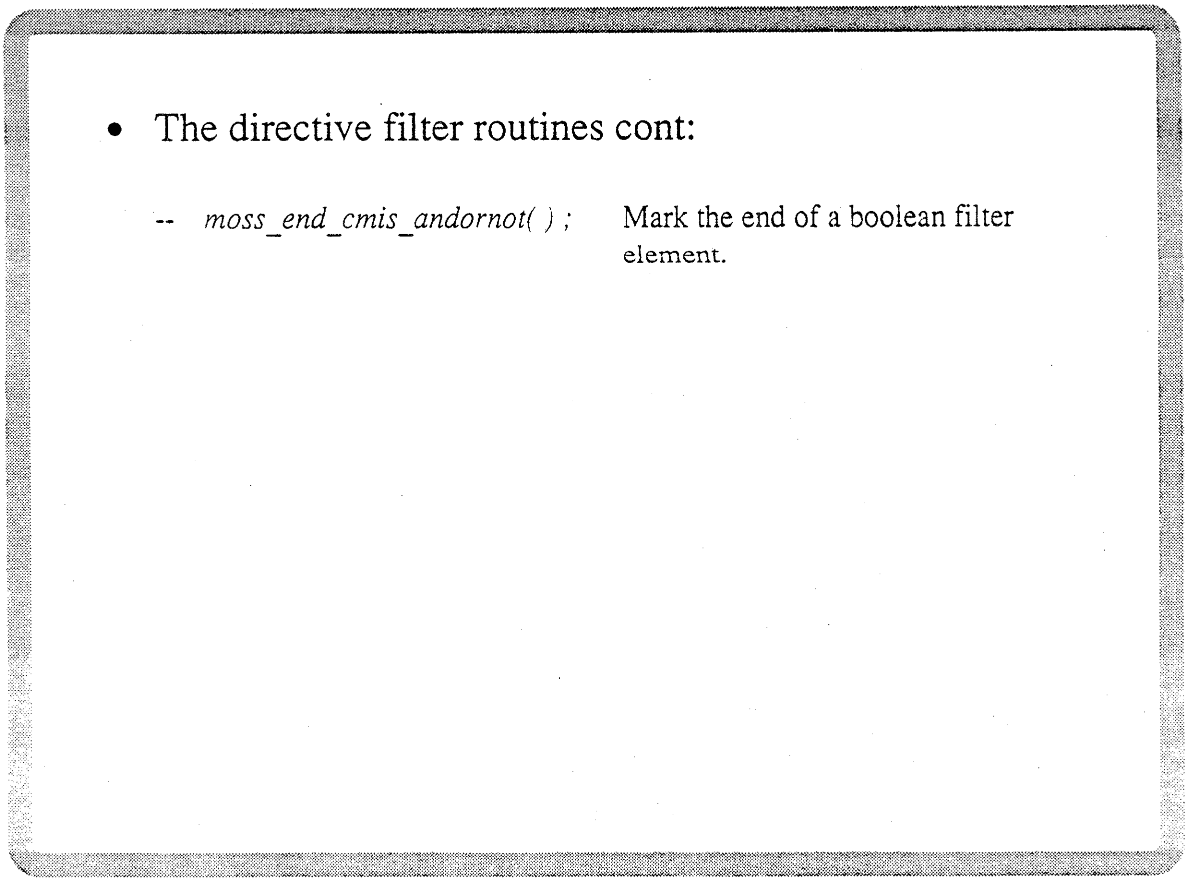

- The directive filter routines cont:

-- *moss_end_cmis_andornot( ) ;*	Mark the end of a boolean filter element.

Slide No. 18

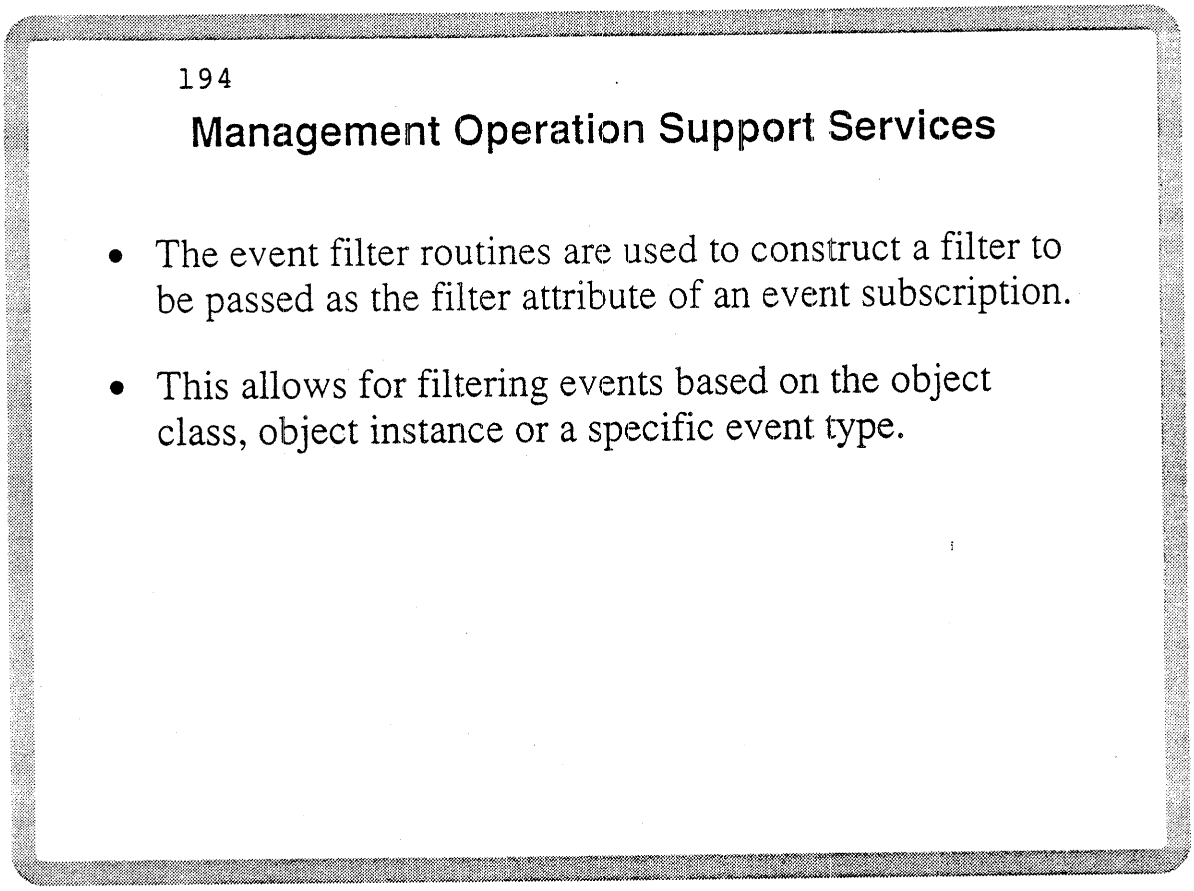

Management Operation Support Services

- The event filter routines are used to construct a filter to be passed as the filter attribute of an event subscription.
- This allows for filtering events based on the object class, object instance or a specific event type.

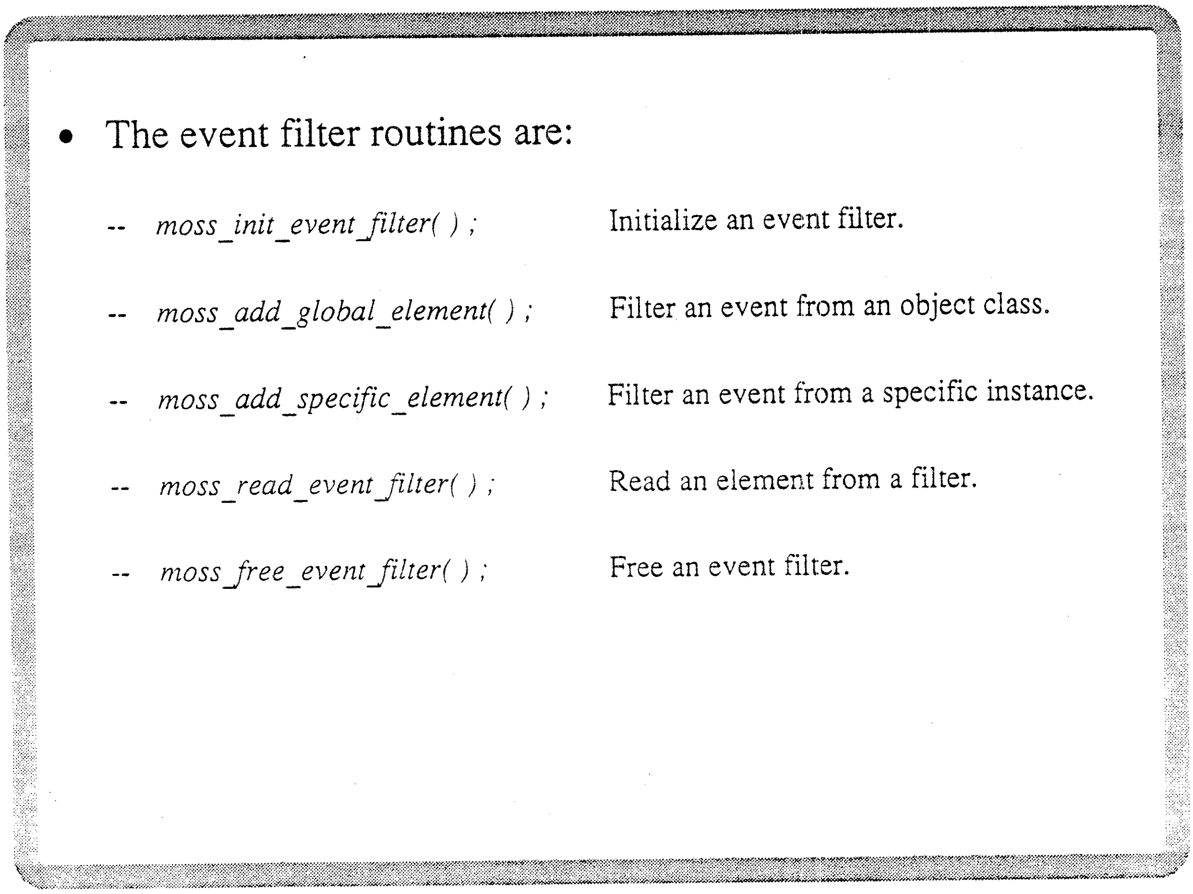

- The event filter routines are:

-- *moss_init_event_filter( ) ;*	Initialize an event filter.
-- *moss_add_global_element( ) ;*	Filter an event from an object class.
-- *moss_add_specific_element( ) ;*	Filter an event from a specific instance.
-- *moss_read_event_filter( ) ;*	Read an element from a filter.
-- *moss_free_event_filter( ) ;*	Free an event filter.

3.4.5 NOC-Tools Working Group (noctools)

CHARTER

Chairpersons: Robert Enger/Contel enger@sccgate.scc.com
Robert Stine/Sparta stine@sparta.com

Mailing List: noctools@merit.edu

Description of Working Group:

The NOC-Tools Working Group will develop a catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

Specific Objectives:

1. Identify tools available to assist network managers in debugging and maintaining their networks.
2. Publish a reference document listing what tools are available, what they do, and where they can be obtained.
3. Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
4. Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.
5. Identify the need for new or improved tools as may become apparent during the compilation of the reference document.

Estimated Timeframe for Completion:

The first edition of the catalog will be submitted for final review at the October 1989 IETF meeting. Preliminary versions will be made available earlier.

CURRENT MEETING REPORT

Reported by Robert Stine/SPARTA and Robert Enger/Contel

MINUTES

At this meeting, Bob Stine briefed the working group on the status of the NOCTools Catalog, and identified several issues concerning dissemination of the catalog, and production of the second edition.

Representatives of Sun Microsystems objected to the text of the YP/DNS discussion in the tutorial section. While they offered to assist in the re-write, time constraints resulted in Bob Stine doing the job alone.

The catalog will be published as an RFC in early March. The delay is to allow the draft to remain frozen, though subject to critique, for a respectable period of time. During this period, typos and other minor glitches will be quietly corrected.

Among the issues for the second edition are a few improvements to the tool description format, refinement of keywords, and improved record keeping on both entry sources and the solicitations that have been made for catalog submissions. Representatives from CMU have requested that some of their catalog entries be updated; new text is expected from them for the second edition.

Following the meeting, Steve Crocker, Dave Crocker, and co-chairs Bob Enger and Bob Stine decided that responsibility for disseminating the catalog could be assumed by DAWG, and that production of a new catalog could be handled by a new working group, when needed. At the closing plenary, Dave Crocker (head of the IETF Network Management Area), announced the disbanding of the NOCTools Working Group. In short, we declared victory and quit. Bob Stine will remain as the interim editor, until the formation of a working group to produce the second edition of the catalog.

ATTENDEES

Armstrong, Karen	armstrong@sds.sdsc.edu
Aronson, Cathy	cja@merit.edu
Bowers, Karen L.	kbowers@nri.reston.va.us
Brunner, Ted	tob@thumper.bellcore.com
Crocker, Dave	dcrocker@nsl.dec.com
Enger, Robert	enger@sccgate.scc.com
Gerich, Elise	epg@merit.edu
Hahn, Jack	hahn@umd5.umd.edu
Hays, Ken	hays@scri1.scri.fsu.edu
Jordt, Dan	danj@cac.washington.edu
Malkin, Gary	gmalkin@proteon.com
Moore, Berlin	prepnet@andrew.cmu.edu
Morris, Don	morris@ucar.edu
Pace, Donald	pace@fsu1.cc.fsu.edu
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Reynolds, Joyce	jkrey@venera.isi.edu
Roubicek, Karen	roubicek@nnsc.nsf.net
Saperia, Jon	saperia%tcpjon@decwrl.dec.com
Smith, Pat	psmith@merit.edu
St. Johns, Mike	stjohns@umd5.umd.edu
Stahl, Mary	stahl@nisc.sri.com
Throop, Dean	throop@dg-rtp.dg.com
Waldbusser, Steve	sw01@andrew.cmu.edu
Ward, Carol	cward@spot.colorado.edu
Wobus, John M.	jmwobus@suvm.acs.syr.edu
Yuan, Aileen	aileen@gateway.mitre.org
Zimmerman, David	dpz@convex.com

3.4.6 OSI Internet Management Working Group (oim)

CHARTER

Chairpersons: Lee LeBarre/Mitre cel@mbunix.mitre.org
Brian Handspicker/DEC bd@vines.dec.com

Mailing Lists: oim-request@mbunix.mitre.org
oim@mbunix.mitre.org

Description of Working Group:

- 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.
- 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.

Specific Objectives:

1. Develop implementors agreements for implementation of CMIP over TCP and CMIP over OSI.
2. Develop extensions to common IETF SMI to satisfy requirements for management of the Internet using OSI management models and protocols.
3. Develop extensions to common IETF MIB-II to satisfy requirements for management of the Internet using OSI management models and protocols.
4. Develop prototype implementations based on protocol implementors agreements, IETF OIM Extended SMI and Extended MIB.
5. Promote development of products based on OIM agreements.
6. Provide input to the ANSI, ISO, NIST and NMF to influence development of OSI standards and implementors agreements.
7. Completion of the following drafts:
    - Implementors Agreements
    - Event Management
    - SMI Extensions

- MIB Extensions
- OSI Management Overview
- Guidelines for the Definition of Internet Managed Objects

Estimated Timeframe for Completion:

Current specific objectives should be completed by December 1990.

CURRENT MEETING REPORT

Reported by Lee LaBarre/MITRE

MINUTES

The OIM WG met at the Florida IETF meeting for one morning. The meeting was primarily for dissemination of information and to seek feedback on the OIM activities. Topics discussed were:

- CMOT Agreements on CMIS/P IS
- Interoperability Lab at DEC's NSL
- Testing of CMOT Implementations
- The Internet MIB and MIB-II
- Management Functions for event reporting and logging

CMOT Agreements on CMIS/P IS

Brian Handspicker provided an overview of the draft agreements for amending the CMOT RFC in accordance with the International Standard for CMIS and CMIP. The OIM will reference the agreements coming out of the OSI Implementors Workshop NMSIG for CMIS and CMIP. We are requesting review of those NMSIG agreements so that we may provide input to the NMSIG meeting the week of 12 March. Comments should be in to Brian and Lee by 6 March. Of course you will need a copy of the agreements to comment on. We will send a copy to the distribution list ASAP. It is important that these agreements be in final form for the March meeting since a request will be made to put them into the OIW stable agreements. Once there, they cannot be changed without great difficulty. We need stable agreements on which to base stable implementations.

The agreements will cover CMIS/P, the underlying services required for use with TCP/IP (CMOT) and for use with OSI protocols, and policy on the use of associations.

Agreements will be developed in two phases: Phase I will include agreements on CMIP IS and its use with MIB-I. The current RFC 1095 rules for identifying MIB instances will be rertained. Phase II will include agreements on the use of the Phase I agreements on CMIP, agreements on the use of MIB-II and the OSI SMI, and agreements on the management functions. Again we will seek alignment with agreements developed by the OIW - where appropriate.

Interoperability Lab at NSL

Dave Crocker described the NSL Lab that has been set up for the purpose of interoperability testing - called OpenLab. The lab will provide space, power, air conditioning, etc. only. Access to the lab is convenient - no escort is needed. Room is currently available for 5 to 10 people to work in the lab. The area can be quickly expanded to accomodate 30 - 50 people. The lab will be open on 15 March, but Dave says that anyone wishing to use it before that date may do so by making arrangements with him. Currently no charge is required. But that may change.

Although the lab space is being provided by DEC, DEC will not control or oversee the operations of the lab. Security for the lab will be provided by a guard who controls access to the facilities by authorized individuals.

For further information contact Dave Crocker at dcrocker@nsl.dec.com or dcrocker@decwrl.dec.com or call (415) 688-6820.

Brian Handspicker will be developing a plan for use of the lab to test CMOT implementations. The plan will allow for testing of current CMOT implementations based on RFC1095 as well as CMIP IS based implementations. Contact Brian for more information especially if you want to do testing of RFC 1095 implementations. Brian is at bd@vines.enet.dec.com.

Testing

Tom Halcin of HP is developing a test plan for interoperability testing of CMOT implementations. Phase I testing will include the MIB I defined in RFC 1066, and RFC 1095 CMOT and IS based CMOT. Essential for this testing are agreements on marching rules for MIB-I objects, as discussed below.

MIB

Lee LaBarre provided an overview of the OIM MIB-II Internet-Draft which is available at the NIC. The draft recasts the proposed MIB-II into the OSI SMI in accordance with the ISO DP 10165-4 Guidelines for Managed Objects. Attributes are added to align with the ISO requirement to assign distinguished attributes to every managed object class. Inheritance from the object "top" is specified. SNMP traps are included in the MIB as events assigned to specific objects.

The OIM MIB-II draft specifies matching rules to each attribute in the MIB for the purpose of applying CMIS Filtering. Such specification is also required for MIB-I attributes (subset of MIB-II) for the purpose of phase I agreements and interoperability testing. Comments are requested ASAP on the Matching Rules specified for the MIB-I subset of attributes specified in the OIM MIB-II draft. Please send comments to

Tom Halcin (halcinLee LaBarre (cel@mbunix.mitre.org).

Management Functions

Lee Labarre has written a tutorial (specification) on event reporting, logging, and alarm management functions. It describes the objects required by ISO for these functions, and the models by which the functions were developed. The Internet Alertman WG effort on event flow control is incorporated into the paper. The paper will be submitted as an Internet-Draft shortly. It was offered as input ot the IETF Alertman WG.

ATTENDEES

Brunner, Theodore	tob@thumper.bellcore.com
Case, Jeff	case@utkux1.utk.edu
Crocker, Dave	dcrocker@nsl.dec.com
Feridun, Metin	mferidun@bbn.com
Gardner, Ella	epg@gateway.mitre.org
Glappa, Adrianne	no email
Halcin, Tom	halcin%hpinddm@hplabs.hp.com
Handspicker, Brian	bd@vines.dec.com
Hunter, Steven	hunter@ccc.mfecc.arpa
Hytry, Tom	tlh@iwlcs.att.com
Kellen, Daniel	kellen@eglin.af.mil
LaBarre, Lee	cel@mbunix.mitre.org
Malkin, Gary	gmalkin@proteon.com
McCloghrie, Keith	sytek!kzm@hplabs.hp.com
Newkerk, Oscar	newkerk@decwet.dec.com
Nitzan, Becca	nitzan@nsipo.nasa.gov
Perkins, Dave	dave_perkins@3com.com
Robertson, Jim	jar@bridge2.3com.com
Saperia, Jon	saperia%tcpjon@decwrl.dec.com
Sheridan, Jim	jsherida@ibm.com
Showalter, Jim	gamma@mintaka.dca.mil
Stine, Bob	stine@sparta.com
Stursa, Scott	xndmis14@servax.bitnet
Waldbusser, Steve	sw01@andrew.cmu.edu
Woodburn, Robert	woody@saic.com
Zimmerman, David	dpz@convex.com

3.4.7 SNMP Working Group (snmp)

CHARTER

Chairperson: Marshall T. Rose/NYSERNet, mrose@nisc.nyser.net

Mailing List: snmp-wg@nisc.nyser.net

Description of Working Group:

The SNMP Working Group has the goal of producing necessary SNMP centric RFCs especially in the area of the Management Information Base (MIB) and the Structure of Management Information (SMI) to provide for both critical operational management requirements and cooperative experimental work.

Specific Objectives:

Provide a draft RFC for an enhanced backwardly compatible MIB in 4Q89 which can be implemented and interoperability tested by 1Q90 to address critical operational requirements. After multivendor testing, draft will be submitted to the RFC Editor for standardization.

Milestones

		SCHEDULED	ACTUAL
GOAL	Prepare MIB-II draft		
o TASK	- Initial meeting to assign actions	89-08-18	89-08-18
o TASK	- Actions due	89-09-01	89-09-08
o TASK	- Edit draft	89-09-15	89-09-22
o TASK	- QC draft and release	89-09-22	89-10-29
GOAL	Examine and tentatively agree		
o TASK	- Discussion meeting to review draft	89-10-16	89-10-16
o TASK	- Edit drafts and release		
	- MIB-II draft	89-10-20	
	- Ethernet-like draft	89-10-20	
	- T1-carrier draft	89-10-20	
	- Token-ring draft	89-10-31	
	- other drafts	TBD	
GOAL	Implement and report back		
o TASK	- Incremental editing of drafts	throughout	
o TASK	- 90 percent implimentation of relevant portions	89-12-01	
	- along with interoperability testing		
GOAL	Evaluate and possibly iterate		
o TASK	- Determine if concensus is reached	89-12-01	
o TASK	- Final edit of drafts	89-12-08	
o TASK	- Submit drafts for standardization		
	- MIB-II draft	89-12-08	
	- Ethernet-like draft	89-12-08	
	- T1-carrier draft	89-12-08	
	- Token-ring and other drafts	N/A	

CURRENT MEETING REPORT

Did not meet.

3.5 OSI Integration Area

Directors: Ross Callon/DEC and Rob Hagens/University of Wisconsin

The OSI general WG has reviewed the following documents:

- RFC 1006
- Internet Draft: DRAFT-UCL-KILLE-NETWORKADDRESSES-00.PS.1
- Internet Draft: DRAFT-UCL-KILLE-PRESENTATIONADDRESS-00.PS.1
- Internet Draft: DRAFT-OSF-SHUE-OSIUDP-00.TXT.1

We have determined which should be progressed in the RFC Standards Track, and for each document to be progressed, the anticipated requirement level. These recommendations are listed in the OSI general meeting report.

The OSI-NSAP working group had their initial meeting. The group accepted their charter, to develop guidelines for NSAP assignment and administration, and identified eight issues/questions that need to be resolved. These issues are detailed in the meeting report.

The OSI-X.400 Working Group met to consider the transition to, and operation of an Internet X.400 Private Management Domain. Their work can be summarized into 2 points:

1. There is no need to specify a transition ORAddress structure for the Internet; domain defined attributes will suffice.
2. There is a real need to allocate funds to administer and operate a PRMD on behalf of the National Research and Education Network, NREN.

The group is preparing a detailed statement to this effect that also includes details on the administration and operation of an NREN PRMD.

3.5.1 OSI NSAP Guidelines Working Group (osinsap)

CHARTER

Chairpersons: Richard Colella, COLELLA@OSI3.NCSL.NIST.GOV

Mailing Lists: ietf-osi-nsap@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).

Specific Objectives:

This working group will produce a paper describing guidelines for the acquisition and administration of NSAP addresses in the Internet. The goal is to have the paper incorporated, in whole or in part, into the "GOSIP Users Guide". Assuming use of existing NSAP address standards, there are two questions facing an administration:

1. 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?
2. 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)

Estimated Timeframe for Completion:

Three or four IETF meetings.

CURRENT MEETING REPORT

Reported by Richard Colella/NIST

MINUTES

This was the initial meeting of the OSI NSAP Guidelines Working Group. The chairman, Richard Colella opened the meeting by distributing and reviewing the working group charter. A question was raised about the application of the groups work to the European part of the Internet. Although this remains an open issue the discussion in the group, at this time, will be limited to the US Internet. Another issue was raised about possible duplication of work with respect to the GSA guidelines on administration authorities. It was determined that the group planed to go beyond the GSA guidelines and that there should not be any major duplication of current work.

Since the groups orientation is towards the NSAP structure specified in GOSIP Version 2, Richard Colella presented a review of the NSAP structure. From the review and discussion the group agreed that all issues dealing with the fields from the Administrative Authority to the left are administrative and those from the Routing Domain to the right are technical.

The group determined eight issues or questions that need to be resolved or answered. The issues are:

1. The use of 47 0005 NSAPs by US Non-Government organizations.
2. The use of 47 0005 NSAPs by Non-US organizations.
3. The DoD NSAP structure vs the Internet structure and possible interoperability problems.
4. The relationship of GSA to the Internet.
5. Possible technical repercussions of NSAP deployment.
6. Who gets the Administrative Authority and the authority to define routing domains and area?
7. How can the Internet topology be mapped onto IS-IS routing?
8. The lack of an inter-domain protocol.

The remaining time in the meeting was spend discussing issues 1,2,3,4 and 7. The main points of the discussions are outlined below.

Issue 4 - Relationship of GSA to the Internet.

- GSA is the authority for Administrative Authority field assignment under 47

0005.

- Government organizations request an AA field assignment from GSA.
- GSA returns an AA value or a reason why a value was not assigned.

Issue 1 - The use of 47 0005 by US non-Government organizations.

- There is no problem with US non-Government organizations using NSAPs under 47 0005.
- A government agency may delegate authority for an AA to the Internet. It is not known who would administer that AA but the group should probably make a recommendation.

Issue 2 - The use of 47 0005 by non-US organizations.

- European countries require End Systems to have NSAPs derived from their country codes.
- The AA for the US non-Government organizations could be extended for use by non-US organizations as long as they prescribed to the rules and procedures. There are no technical problems associated with this solution but there may be administrative ones.
- Need to figure out how to deal with this issue.

Issue 3 - DoD NSAP structure vs Internet structure and interoperability

- 47 0005 will probably work for DoD fixes assets but might not work for mobile ones.
- DoD has not determined what to do with 47 0006.

Issue 4 - Mapping Internet Topology onto IS-IS Routing

- Need to take a piece of the Internet and map it into OSI. It is an ugly process but it needs to be done.
- One problem is that no one has the big picture.
- Discussed the transition from DECNET Phase 4 to Phase 5 and the problems associated with the transition.
- Tony Hain presented DoE's model.

The meeting adjourned with plans to hold an informal lunch meeting to discuss these issues further.

ATTENDEES

Almquist, Philip	almquist@jessica.stanford.edu
Ballard, Bare	bare%hprnd@hplabs.hp.com
Case, Jeff	case@utkux1.utk.edu
Cerf, Vint	vcerf@nri.reston.va.us
Colella, Richard	colella@osi3.ncsl.nist.gov
Deboo, Farokh	sun!iruucp!ntrlink!fjd
Farinacci, Dino	dino@bridge2.3com.com
Galvin, James	galvin@tis.com
Gardner, Ella	epg@gateway.mitre.org
Gan, Der-Hwa	no email
Glappa, Adrianne	no email
Goguely, Herve	rvg@bridge2.3com.com
Gross, Martin	martin@protolaba.dca.mil
Hain, Tony	hain@nmfecc.arpa
Heinanen, Juha	jh@funet.fi
Hua, Binh	no email
Hytry, Tom	tlh@iwlcs.att.com
Jensen, Phil	jensen@fsu1.cc.fsu.edu
Katz, Dave	dkatz@merit.edu
Lazear,Walt	lazear@gateway.mitre.org
Love, Paul	loveep@sds.sdsc.edu
Marcinkevicz, Mike	mdm@gumby.dsd.trw.com
McCloghrie, Keith	sytek!kzm@hplabs.hp.com
Mills, Cyndi	cmills@bbn.com
Minshall, Greg	minshall@kinetics.kinetics.com
Morris, Dennis	morrisd@imo-uvax.dca.mil
Nitzan, Becca	nitzan@nsipo.nasa.gov
Sheridan, Jim	jsherida@ibm.com
Showalter, Jim	gamma@mintaka.dca.mil
Shue, Chi	chi@osf.org
Stursa, Scott	xndmis14@servax.bitnet
Winkler, Linda	b32357@anlvm.ctd.anl.gov
Youssef, Mary	mary@ibm.com

3.5.2 OSI General Working Group (osigen)

CHARTER

Chairpersons:	Ross Callon/DEC	callon@erlang.dec.com
	Rob Hagens/UWisc.	hagens@cs.wisc.edu

Mailing Lists:

ietf-osi@cs.wisc.edu
ietf-osi-request@cs.wisc.edu

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.

Specific Objectives:

The following are specific short-term goals and objectives for the OSI WG. Other mid-term objectives have also been identified and are available from the chairs.

- 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.
- Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley folks.
- Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.
- What routing should be used short term for (i) intra-domain routing; and (ii) inter-domain routing?
- For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Are there outstanding issues remaining here?
- 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.

Estimated Timeframe for Completion:

This is an operational and liason WG and, as such, has an indefinite lifetime.

CURRENT MEETING REPORT

Reported by Rob Hagens/University of Wisconsin

AGENDA

Discuss status of RFC 1006 and related documents.

MINUTES

The meeting was convened by co-chairman Rob Hagens.

The following documents were reviewed at the meeting:

1. RFC 1006
2. Internet Draft: DRAFT-UCL-KILLE-NETWORKADDRESSES-00.PS.1
3. Internet Draft: DRAFT-UCL-KILLE-PRESENTATIONADDRESS-00.PS.1
4. Internet Draft: DRAFT-OSF-SHUE-OSIUDP-00.TXT.1

The group discussed which of these documents should be progressed in the RFC Standards Track, and for each document to be progressed, the anticipated requirement level.

The outcome of this discussion was:

- RFC1006: progress to Draft Standard with an anticipated requirement level of "Recommended for all systems which run OSI connection- oriented applications over TCP/IP"
- Internet Draft: DRAFT-UCL-KILLE-NETWORKADDRESSES-00.PS.1: progress to Proposed Standard with an anticipated requirement level of "Required of all systems which implement RFC 1006"
- Internet Draft: DRAFT-UCL-KILLE-PRESENTATIONADDRESS-00.PS.1: progress to Proposed Standard with an anticipated requirement level of "Required of all systems which accept or display OSI addresses in textual form"
- Internet Draft: DRAFT-OSF-SHUE-OSIUDP-00.TXT.1: progress to Experimental Status

ATTENDEES

Bennett, Derek	xndmis14@servax.bitnet
Cerf, Vinton	vcerf@nri.reston.va.us
Colella, Richard	colella@osi3.ncsl.nist.gov
Curci, Raymond	curci@gw.scri.fsu.edu
Easterday, Tom	tom@nisca.ircc.ohio-state.edu
Engineer, Huniad	hunaid@opus.cray.com
Gardner, Ella	epg@gateway.mitre.org
Gan, Der-Hwa	no email
Glappa, Adrianne	no email
Goguely, Herve	rvg@bridge2.3com.com
Goldstein, Steven	goldstein@note.nsf.gov
Hagens, Robert	hagens@cs.wisc.edu
Handspicker, Brian	bd@vines.dec.com
Hytry, Tom	tlh@iwlcs.att.com
Jacobsen, Ole	ole@csli.stanford.edu
Kellen, Daniel	kellen@eglin.af.mil
Kirstein, Peter	kirstein@cs.ucl.ac.uk
LaBarre, Lee	cel@mbunix.mitre.org
Love, Paul	loveep@sds.sdsc.edu
Mills, Cyndi	cmills@bbn.com
Nitzan, Becca	nitzan@nsipo.nasa.gov
Perlman, Radia	no email
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Sheridan, Jim	jsherida@ibm.com
Showalter, Jim	gamma@mintaka.dca.mil
Shue, Chi	chi@osf.org
Sklower, Keith	sklower@okeeffe.berkeley.edu
Stine, Robert	stine@sparta.com
Stursa, Scott	xndmis14@servax.bitnet
Vaudreuil, Greg	gvaudre@nri.reston.va.us
Wilder, Rick	rick@gateway.mitre.org
Winkler, Linda	b32357@anlvm.ctd.anl.gov
Woodburn, Robert	woody@saic.com
Youssef, Mary	mary@ibm.com

OSI Connectionless Transport Over Internet UDP

A RFC Proposal

(A Companion Document to RFC 1006)

February 7, 1990

Chi Shue, Bill Haggerty, Kurt Dobbins

OSI Connectionless Standards

7498/AD1: OSI - Basic Reference Model- Addendum 1
Connectionless-Mode Transmission

Layer	Service	Protocol	Status
Application	8649/AD2	10035	IS (June, 1990)
Presentation	8822/AD1	9576	IS (June, 1990)
Session	8326/AD3	9548	IS
Transport	8072/AD1	8602	IS
Network	8348/AD1	8473	IS

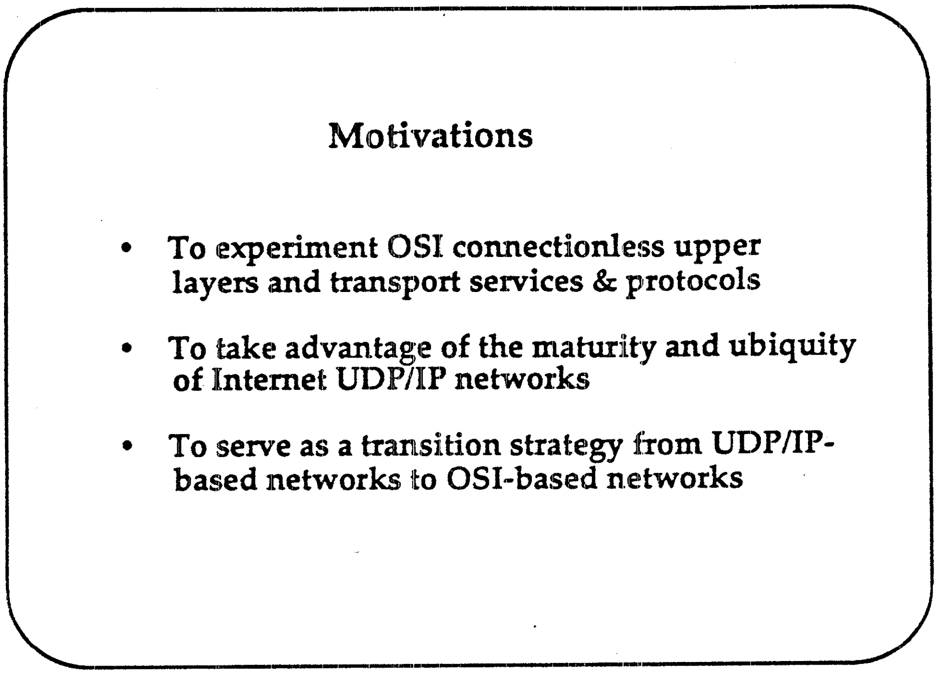

Motivations

- To experiment OSI connectionless upper layers and transport services & protocols
- To take advantage of the maturity and ubiquity of Internet UDP/IP networks
- To serve as a transition strategy from UDP/IP-based networks to OSI-based networks

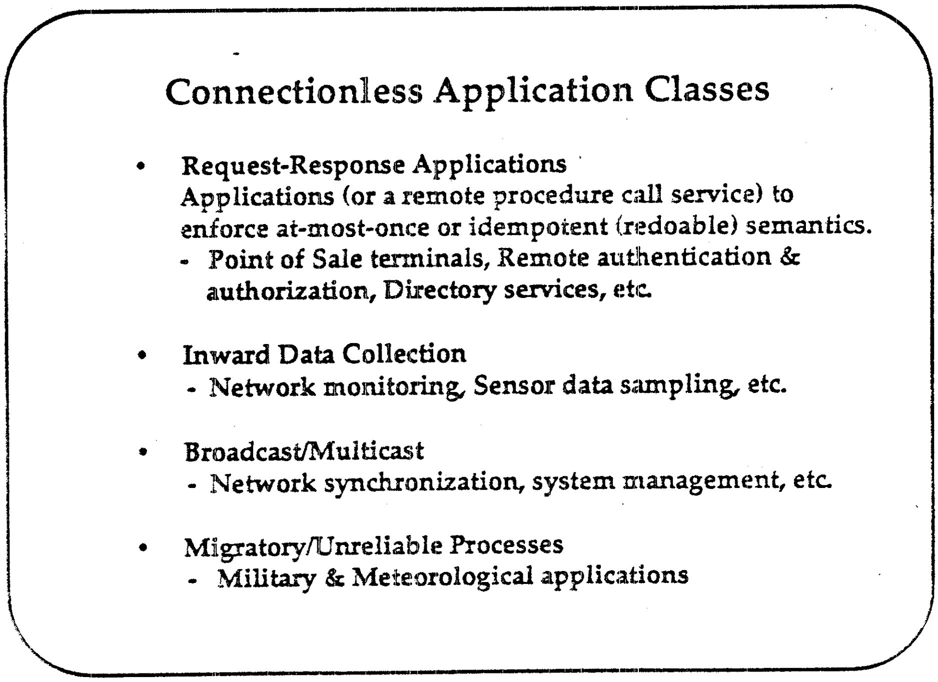

Connectionless Application Classes

- Request-Response Applications
  Applications (or a remote procedure call service) to enforce at-most-once or idempotent (redoable) semantics.
  - Point of Sale terminals, Remote authentication & authorization, Directory services, etc.
- Inward Data Collection
  - Network monitoring, Sensor data sampling, etc.
- Broadcast/Multicast
  - Network synchronization, system management, etc.
- Migratory/Unreliable Processes
  - Military & Meteorological applications

The Model

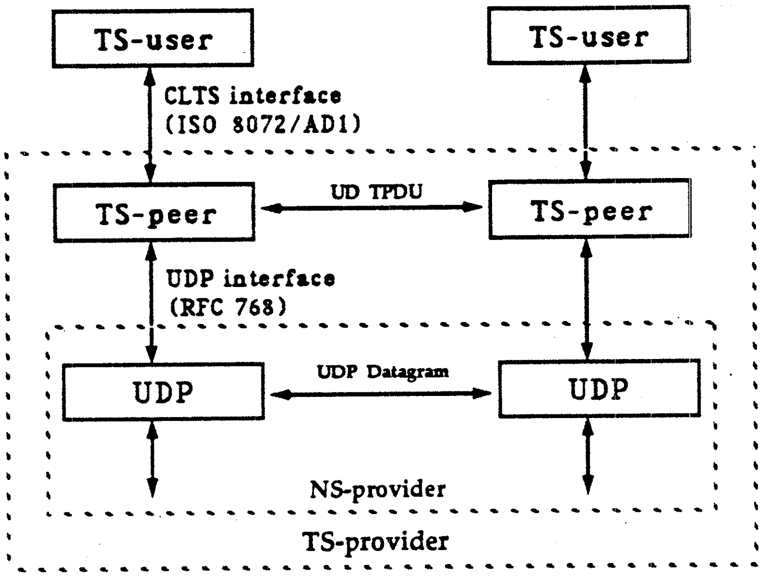

TS-user

TS-user

CLTS interface
(ISO 8072/AD1)

TS-peer

UD TPDU

TS-peer

UDP interface
(RFC 768)

UDP

UDP Datagram

UDP

NS-provider

TS-provider

Mapping Between Connectionless Network Service and UDP
Service Primitives

<u>CLNS</u>	<u>UDP</u>
N-UNIT-DATA.REQUEST	SEND DATAGRAM
N-UNIT-DATA.INDICATION	READ DATAGRAM

Note: UDP port 102 will be reserved for the use of this RFC

Mapping of Service Parameters

CLNS	UDP
Source address	Source IP address from calling TS-address
Destination address	Destination IP address from called TS-address
Quality of service	(ignored)
NS-user data	UD TPDU constructed from T-UNIT-DATA

Protocol Format

UD TPDU structure encapsulated in the UDP data field

1	2	3 m	m+1 n
LI	UD 01000000	Variable Part	User Data

LI - Length of the header including parameters

UD - Unit Data TPDU type

Variable Part - Source T-Selector ID
Destination T-Selector ID

User Data - TSDU

Implementation Status

- Complete OSI Connectionless Transport and Upper layers were successfully implemented in December, 1988
- Design was based on ISODE version 4.0
- Better response time and less code size were achieved than that of the connection-oriented ISODE
- Design and source code will be contributed to the future ISODE release
- Connectionless ROS was proposed to CCITT/VII DAF Rapporteurs Group

3.5.3 OSI X.400 Working Group (osix400)

CHARTER

Chairperson: Robert Hagens/UWisc., hagens@cs.wisc.edu

Mailing Lists:

ietf-osi-x.400@cs.wisc.edu
ietf-osi-x.400-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.

Specific Objectives:

1. Develop a memo describing known issues and problems.
2. Develop a scheme to alleviate the need for static RFC 987 mapping tables.
3. Develop a scheme to support X.400 routing.
4. Consider ways in which directory services may be utilized in order to hide the details of RFC 822 and X.400 addressing.

Estimated Timeframe for Completion:

The timeframe is being reconsidered by the working group chair.

CURRENT MEETING REPORT

Reported by Robert Hagens/University of Wisconsin

MINUTES

The meeting was convened by chairman Rob Hagens. An attendance list will be published with the Proceedings of the IETF.

The group (much smaller than the last meeting) began the meeting by discussing the sort of address structure required when X.400/MHS is introduced into the Internet. The group concluded that there is no need to design or define an interm/transition address structure. RFC 987/1138 defines an acceptable transition structure: the RFC-822 domain defined attribute. There was some discussion about the lack of widespread support for DDAs. It appears that most ADMD providers support DDAs in their MTAs, however there may be user agents that do not support the entry of DDAs. The general feeling was that this was a shortcoming of the specific user agents; it could be corrected. It should not effect the organization of the NREN X.400/MHS service.

After this, the group discussed the need for a PRMD authority for the NREN. Note: the group's definition of NREN is the US portion of the IP-connected Internet.

Several points were discussed:

- An NREN PRMD is a long term solution for certain organizations.
- Organizations may "grow out" of the NREN PRMD and register elsewhere.
- As X.400 software is deployed within the NREN, we need to organize a coherent MTS before chaos decends.
- We must provide cheap and quick registration services for the NREN PRMD.
- The NREN PRMD may negotiate with US ADMDs for international service. Such negotiations must be on the basis of originator keeps all revenue.
- The NREN PRMD may relay traffic for other PRMDs
- An NREN PRMD must provide a registration service as well as manage operational connectivity (i.e., routing).
- The 'ole ADMD field question: there is a general desire to keep the ADMD field blank. Many ADMD providers require the ADMD field to contain the name of the ADMD. Should the NREN PRMD automatically fill in the ADMD field? One suggestion was that within the US, the ADMD field should be kept blank, but international traffic would have the ADMD field set as the message leaves the country. Some discussion of munging P1.originator vs. P2.originator ensued. Is that a protocol violation? How does it effect security? How does it

effect a "reply"?

- The question was raised as to whether the NREN could register itself as an ADMD. The answer was not known.
- Creation of 987/1138 mapping tables for members of the NREN PRMD was considered a good thing. Non-members could relay traffic into and out of the PRMD via DDAs.
- The operation of the NREN PRMD will not be free. There is a need to fund a few people who will organize and operate the PRMD.

The group agreed that a document must be written which describes the responsibilities and operational aspects of the NREN PRMD. A tentative title for this document is "Transition and long term strategy for operation of X.400/MHS in the NREN". I hope to have a preliminary draft of this document by the end of March. James Galvin (of TIS), offered to draft an outline.

At the end of the meeting, Professor Kirstein from UCL described his interest in promoting and experimenting with ODA. He has access to implementations of ODA that may be utilized by an Internet-ODA-X.400 experiment. There are no proposed experiments at this time. Any one interested in any Internet-ODA-X.400 experiments should contact Professor Kirstein or an OSI area director.

ATTENDEES

Bennett Derek	xndmis14@servax.bitnet
Colella, Richard	colella@osi3.ncsl.nist.gov
Galvin, James M.	galvin@tis.com
Gardner, Ella	epg@gateway.mitre.org
Goguely, Herve	rvg@bridge2.3com.com
Gross, Martin	martin@protolaba.dca.mil
Hagens, Robert	hagens@cs.wisc.edu
Jensen, Phil	jensen@fsu1.cc.fsu.edu
Kellen, Daniel	kellen@eglin.af.mil
Kirstein, Peter	kirstein@cs.ucl.as.uk
Lazear, Walter	lazear@gateway.mitre.org
Shue, Chi	chi@osf.org
Stursa, Scott	xndmis14@servax.bitnet
Sturtevant, Allen	sturtevant@ccc.nmfecc.gov
Winkler, Linda	b32357@anlvm.ctd.anl.gov

3.6 Operation Area

Iterim Director: Phill Gross/NRI

The Interconnectivity WG (chaired by Guy Almes) will conquer by dividing. IWG has had two main activities in recent meetings - BGP, and operational routing and topology management. We have decided to create a new WG, Topology Engineering (tewg), to focus specifically on the second issue. Scott Brim (Cornell Theory Center) will chair the new TEWG.

TEWG will have a specific goal of coordinating among the various relevant operational routing and topology management groups in the Internet. This includes regional networks, FARNET, national backbones, etc. Guy Almes will continue to chair IWG, which will now take BGP as its single focus. Please see the charters for IWG and TEWG, or contact the chairs, for additional information.

The Joint Monitoring for Adjacent NSFnet Networks WG (JoMANN) has undergone a minor transformation. Sue Hares (Merit) organized JoMANN, at least partly, to assist Merit in interacting with the regional networks attached to NSFnet. JoMANN proved useful enough that we have decided to establish it as a mainstay of the new Operations Area. The WG will be renamed Network Joint Monitoring (NJM) to emphasize that the new focus will be monitoring issues beyond simply networks adjacent to NSFnet. Gene Hastings will chair this important continuing effort.

There is some other preliminary activity in the Operations Area. We held a meeting of the reporters from the major national backbones (NSFnet, ESnet, NSI, DCA/DARPA) in an attempt to make the network status reports a more regular and standardized feature of all IETF meetings. We also had an ad hoc meeting of folks interested in developing standard ways of collecting and reporting network data. We hope to bring these two efforts together, if possible.

3.6.1 Benchmarking Methodology Working Group (bmwg)

CHARTER

Chairpersons: Scott Bradner/Harvard, sob@harvard.harvard.edu
Mick Scully, mcs@ub.com

Mailing List: bmwg@harvisr.harvard.edu

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 standalone 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 standalone devices has 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.

Specific Objectives:

1. Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.
2. The document will also define various classes of standalone 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.
3. 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.

In addition, this document will make specific recommendations on a common reporting structure for benchmark results.

The document will be organized such that each section::

(a) Defines a device class.
(b) Defines the performance characteristics important to this class of device.
(c) Recommend a specific benchmark suite (FLINTSTONES) for this class of device.
(d) Define a common reporting format for the results of the benchmark suite.

Estimated Timeframe for Completion:

We plan to issue a draft document for Objective No. 1 by late December 1989. A document for Objective No. 2 is planned for the end of February 1990 concentrating on a selected set of device classes. The effort will continue on Objective No. 2 and No. 3 with final reports available in the late 1990 time frame.

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

MINUTES

We reviewed and edited the draft version of the Benchmarking Terminology memo. Consensus was reached on a number of changes. A final version has now been edited and will be submitted as soon as an appendix consisting of a set of mathematically precise formal definitions has been completed.

The next meeting time was set for a video conference on Feb 23rd.

ATTENDEES

Fred Baker	baker@vitalink.com
Chet Birger	cbirger@bbn.com
Scott Bradner	sob@harvard.harvard.edu
Jim Forster	forster@cisco.com
Der-Hwa Gan	dhg@bridge2.3com.com
Stev Knowles	stev@ftp.com
Robert M. Enger	enger@sccgate.scc.com
Gary Malkin	gmalkin@proteon.com
K.K. Ramakrishnan	rama@dsmail.dec.com
Joel Replogle	replogle@ncsa.uiuc.edu
Chi Shue	chi@osf.org
Frank Solensky	solensky@interlan.com
Steven Willis	swillis@wellfleet.com
Mary Youssef	mary@ibm.com

3.6.2 Topology Engineering Working Group (tewg)

CHARTER

Chairperson: Scott Brim, swb@devvax.tn.cornell.edu

Mailing Lists:

tewg@devvax.tn.cornell.edu
tewg-request@devvax.tn.cornell.edu

Description of Working Group:

The Topology Engineering Working Group monitors and coordinates connections between networks, particularly routing relationships.

Specific Objectives:

1. Monitor interconnectivity among national and international backbones and mid-level networks.
2. Monitor interconnection policies with a view of moving toward a common scheme for managing interconnectivity.
3. 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.

Estimated Timeframe for Completion:

1. 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.
2. An immediate project is to produce an RFC which will help mid-level networks when changing their interconnectivity.
3. This is an operational and liason WG and, as such, has an indefinite lifetime.

CURRENT MEETING REPORT

Reported by Scott Brim/Cornell University

AGENDA

Report on Europe by Mats Brunell Introduction to TEWG by Scott Brim Decisions on initial action items

MINUTES

RIPE connects to 13,000 IP hosts and 95 organizations in Europe. RIPE agreement with RARE last week removes political obsticals to IP in Europe. RIPE is only a coordination activity; not a service provider.

They need a European "root server".

They are setting up databases at KTH, CWI, and INRIA.

Four Task Forces

- Connectivity and Routing
- Network Management and Operations
- Domain Name Systems
- Formal Coordination

Four Hubs

- NORDUnet; KTH in Stockholm
- EUnet: CWI in Amsterdam
- CERN in Switzerland
- INRIA

Trans-Atlantic Links

T1 from Ithaca, NY to CERN due in spring. Will connect directly into NSFNET. 56kb/s from JvNC to NORDUnet 56kb/s from SURAnet to CWI 9.6kb/s from NYSERnet to Karlsruhe 56kb/s from JvNC to INRIA

Introduction to TEWG by Scott Brim

There are a number of groups engineering interconnectivity among components of the Internet (e.g., the FRICC Engineering Planning Group, the FARnet Technical Committee, and groups involved with geographic areas such as California or the northeast). TEWG will not try to replace these groups as a single forum where all such decisions should be made; on the contrary it will depend on their work, since there is far too much to do in one group or mailing list. Instead it will serve as the point of coordination for all of *them* inasmuch as they affect one another.

TEWG will serve as a clearinghouse for interconnectivity issues which cannot be handled in any more specialized group, for example interactions between private and government-funded nets. Another example is international connections. We are about to get our first Internet loop around the world, involving federally-funded general infrastructure nets, federally-funded mission-oriented nets, commercial internets, and mixed private and public internets – and at least three countries. TEWG will act as a forum for coordinating such situations.

TEWG will be both reactive and proactive in dealing with these problems. We will also gather and share knowledge in the form of RFCs.

Discussion

Discussion centered around what work the group should be doing and what needed to be done soon.

Action Items, to be done before next meeting:

- Explore interactions between VMNET logical topology and Internet topology. – Scott Brim
- Write CSnet plans for using multiple connections to NSFNET. – Dan Long
- Write requirements for useful database on inter-AS connectivity (initial step before taking a survey). Start collecting sample maps, tools, and data. – Paul Tsuchiya
- Write CA*NET plans for using multiple connections. – Dennis Ferguson.
- Start a generic routing policy paper for mid-level networks. Will have at least two sections: "General Principles" and "Rules of Thumb" (including things to watch out for when making changes) – Kent England, Dave O'Leary, Gene Hastings, Vince Fuller, and Matt Mathis
- Explore match/mismatch of BGP with real inter-AS needs – Guy Almes and Matt Mathis
- Write an RFC on "What is an AS?" – Guy Almes

- Liaison with ORWG – Paul Tsuchiya
- Liaison with RIPE – to be determined; Scott Brim will follow up with Phill Gross, RIPE Connectivity and Routing WG.

ATTENDEES

Almes, Guy	almes@rice.edu
Birger, Chet	cbirger@bbn.com
Brunell, Mats	mats.brunell@sics.se
Burgan, Jeffrey	jeff@nsipo.nasa.gov
Coltun, Rob	rcoltun@trantor.umd.edu
England, Kent	kwe@bu.edu
Ferguson, Dennis	dennis@gw.ccie.utoronto.ca
Feridun, Metin	mferidun@bbn.com
Fidler, Mike	ts0026@ohstvma.ircc.ohio-state.edu
Fuller, Vince	fuller@jessica.stanford.edu
Gerich, Elise	epg@merit.edu
Hahn, Jack	hahn@umd5.umd.edu
Hain, Tony	hain@nmfecc.arpa
Hastings, Gene	hastings@psc.edu
Hays, Ken	hays@scri1.scri.fsu.edu
Honig, Jeffrey	jch@tcgould.tn.cornell.edu
Jordt, Dan	danj@cac.washington.edu
Katz, Dave	dkatz@merit.edu
Long, Dan	long@bbn.com
Mathis, Matt	mathis@pele.psc.edu
Morris, Dennis	morrisd@imo-uvax.dca.mil
O'Leary, Dave	oleary@umd5.umd.edu
Oattes, Lee	oattes@utcs.utoronto.ca
Pace, Donald	pace@fsu1.cc.fsu.edu
Piscitello, Dave	dave@sabre.bellcore.com
Rekhter, Yakov	yakov@ibm.com
Tsuchiya, Paul	tsuchiya@thumper.bellcore.com
Veach, Ross	rrv@uiuc.edu
Ward, Carol	cward@spot.colorado.edu
Youssef, Mary	mary@ibm.com

WHAT'S THIS ALL ABOUT?

- TOPOLOGY ENGINEERING — TRY TO TUNE INTER-AS ROUTING FOR EVERYONE'S BENEFIT.
- DON'T REPLACE ANY OTHER GROUP, IN OR OUT OF IETF. DEPEND ON THEM, BUILD ON WORK.

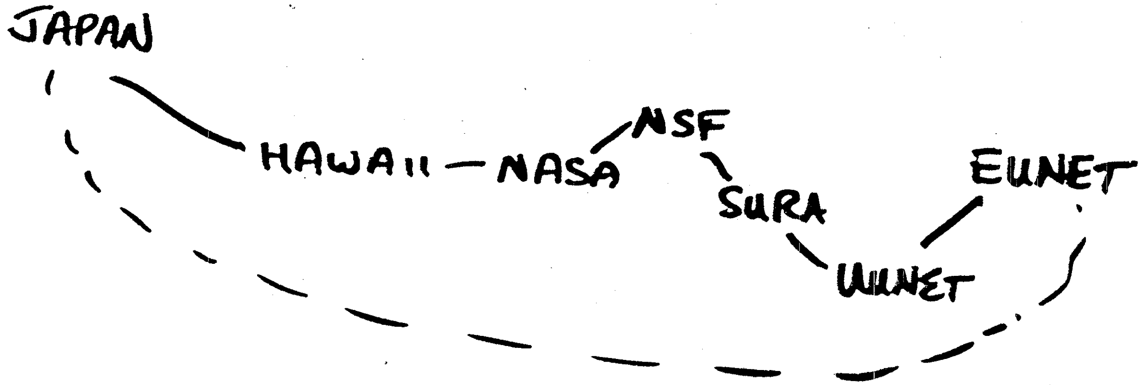

THESE LINKS ARE PRIVATE
REGIONAL
NSF (INFRASTRUCTURE)
MISSION AGENCY

+ INTERNATIONAL.

3.6.3 Network Joint Management (njm)

CHARTER

Chairperson:

Gene Hastings, hastings@psc.edu

Mailing List:

njm@merit.edu
njm-request@merit.edu

Description of Working Group:

There is a need for many different kinds of effort 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 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.

Specific Objectives:

- Examine known problems (continuing and transitory) and publish case studies
- Recommend and publish solutions to problems in terms of:
  - Communication procedures, problem tracking, and problem resolution procedures between NOCs
  - NOC Tools
  - Network Engineering in the areas of:
    - Inter-Administrative Domains
    - Intra-Administrative Domains
    - Routing Domains
  - Software fixes
- Publish User Reports (test drives) on NOC Tools
- Publish Tricks of the Trade.

Estimated Timeframe for Completion:

This is an operational and liason WG and, as such, has an indefinite lifetime.

CURRENT MEETING REPORT

Reported by Philip Almquist/Consultant

MINUTES

1. SNMP community names
   - all routers should support "monitor"
   - routers under the sole control of the regional NOC should support the NSFNET backbone community name
   - if neither of the above work to contact some gateway, try "public"
   - NSI "agrees in principle" to support community names that they will make available to regional NOC's
   - ditto for ESNET
   - regular polling of routers belonging to other organizations is a no-no, except that routers connecting two routing domains may be monitored by both NOC's (and should probably send traps to both NOC's).
   - the above restrictions on "regular polling" do not preclude sending queries to any router while actively debugging a problem
2. Network maps
   - Merit is 90regional maps which are accessible via anonymous FTP
   - regionals which have maps available via anonymous FTP should send pointers to them to the njm list; Merit will treat this as an implicit request to regularly retrieve copies of the map
   - all maps should include a creation date
3. NSFNET <-> BBN core interactions
   - MERIT and DCA have been working on coordinating responses to mailbridge problems at the FIX locations
4. BITNET II
   - Scott Brim expressed concern that BITNET II is being designed by people who do not understand the Internet topology. Thus, the substantial new load it will place on the Internet may occur in inappropriate places. Scott will investigate further.
5. Traceroute
   - several reported that third party traceroute is a real win, and hoped that other routers would support it soon
6. Appropriate us of the "status-reports" mailing list
   - the list is appropriate only for reports of current or very recent events, such as
     - "X will be down from ___ until ___"
     - "X is down"
     - "X was be down from ___ until ___"

    - Summary data can be interesting, but should be posted elsewhere

7. FARNET Report (by Guy Almes)
    - FARNET wants increased FARNET¡-¿IETF cooperation. Regionals should send people to IETF meetings; these people should report back to the regional operators and planners
    - periodic reports of usage/uptimes/etc. are useful (eg, the NSFNET and CERFNET monthly reports). People interested in helping to devise common reporting measures should send mail to Guy.
    - is application throughput commensurate with theoretical path bandwidths (ie, is performance as good as it ought to be)? This is an important question for assessing whether we run networks well and for justifying expensive, high-speed paths. Can we develop a "Dow Jones" average of network performance? Would this measure anything useful, or are most problems just broken TCP's that we have no control over? Interested parties should contact Guy about starting a joint IETF¡-¿FARNET project in this area.
8. NSFNET information files
    - there was a request to the NSF NIC to provide a file of responsible persons indexed by network number
    - other ideas for similar useful files should be sent to nsfnet-info
9. NREN planning
    - Steve Goldstein of NSF wants input on how NIC's and NOC's should be organized in the NREN
    - Gene will send his ideas to the njm list; others may respond
10. Whois service
    - NREN will use an X.500-based whois equivalent
    - some suggested that (in the shorter term) the existing NIC whois should be replicated on additional machines (this may not be practical)

ATTENDEES

Almes, Guy	almes@rice.edu
Almquist, Philip	almquist@jessica.stanford.edu
Aronson, Cathy	cja@merit.edu
Bradner, Scott	osb@harvard.harvard.edu
Brim, Scott	swb@devvax.tn.cornell.edu
Burgan, Jeffrey	jeff@nsipo.nasa.gov
Easterday, Tom	tom@nisca.ircc.ohio-state.edu
England, Kent	kwe@bu.edu
Ferguson, Dennis	dennis@gw.ccie.utoronto.ca
Feridun, Metin	mferidun@bbn.com

Fidler, Mike	ts0026@ohstvma.ircc.ohio-state.edu
Finkelson, Dale	dmf@westie.unl.edu
Fuller, Vince	vaf@stanford.edu
Gerich, Elise	epg@merit.edu
Goldstein, Steve	goldstein@note.nsf.gov
Hahn, Jack	hahn@umd5.umd.edu
Hain, Tony	hain@ccc.nmfecc.gov
Hallgren, Martyne	martyne@tcgould.tn.cornell.edu
Hastings, Gene	hastings@psc.edu
Hunter, Steven	hunter@ccc.nmfecc.gov
Jordt, Dan	janj@cac.washington.edu
Long, Dan	long@bbn.com
Love, Paul	loveep@sds.sdsc.edu
Lynn, Charles	clynn@bbn.com
Mathis, Matt	mathis@psc.edu
Medin, Milo	medin@nsipo.nasa.gov
Morris, Don	morris@ucar.edu
O'Leary, Dave	oleary@umd5.umd.edu
Oattes, Lee	oattes@utcs.utoronto.ca
Perkins, Dave	dave_perkins@3com.com
Pokorney, Dave	poke@nervm.nerdc.ufl.edu
Sheridan, Jim	jsherida@ibm.com
Steinberg, Louis	louiss@ibm.com
Streeter, Roxanne	streeter@nsipo.nasa
Sturtevant, Allen	sturtevant@ccc.nmfecc.gov
Veach, Ross	rrv@uiuc.edu
Ward, Carol	cward@spot.colorado.edu
Wintringham, Dan	danw@osc.edu
Woodburn, Robert	woody@saic.com

3.7 Routing Area

Director: Robert Hinden/BBN

MULTICAST OSPF W.G.

This WG met for the first time at the February IETF. Twenty two people attended the meeting, with the following topics being covered: introduction to IP multicast, overview of the IGMP protocol, survey of current multicast routing strategies, and proposed modifications / additions (algorithms and data) that will be necessary to support multicast routing in OSPF. Most the discussion centered on a desire for performance characteristics of multicast routing (e.g., how dynamic will host group membership be, how often will the cache entries be calculated).

OPEN ROUTING W.G.

The inter-domain policy routing architecture document became an Internet Draft at the beginning of February.

Martha Steenstrup gave a presentation to the IETF plenary outlining the important ideas in the document. The working group meet at IETF and discussed the details of how the architecture works. Work is progressing on the protocols for the initial version of inter-domain policy routing. The group is scheduling a video conference in Mid-March to discuss the proposed protocols.

The ORWG is now open. Send mail to msteenstrup@bbn.com if you would like to be put on the mailing list.

OSPF W.G.

John Moy gave a presentation to the IETF plenary describing OSPF, together with a comparison to the dual IS-IS. Also at the February IETF, there was a meeting of OSPF implementors (led by Rob Coltun and Jeff Honig). The main topic of this meeting was the incorporation of the University of Maryland's OSPF code into the "gated" program. Finally, field testing of OSPF in selected NSF regionals (and other Autonomous Systems) has begun. A new mailing list, ospf-tests@seka.cso.uiuc.edu has been formed to support this effort.

INTERCONNECTIVITY W.G.

The IWG meet several times at the IETF meeting and worked on a new versions of the BGP protocol and an accompanying usage document. New versions of these documents will be released in March.

As part of the reorganization of the IWG, the old BGP mailing list has been merged with the IWG list. The new list is iwg@rice.edu. Please send messages concerning IWG/BGP issues to the merged list.

IS-IS Working Group

Radia Perlman presented a talk to the IETF plenary on the IS-IS routing protocol and IP extensions. The working group meet several times at IETF to further refine the IP extensions and develop plans for several implementations.

OSPF / IS-IS Debate

There was much debate at the FSU IETF meeting on the merits of the OSPF v.s. IS-IS for routing IP traffic. The intention is to pick one as the recommended standard IGP for IP to allow for multivendor routing in a single autonomous system. The discussion was loud and heated, but no blood was shed.

I believe that the only conclusion that was reached is that we need real operational experience with these protocols before one can be selected as the "recommended standard IGP".

3.7.1 IS-IS for IP Internets Working Group (isis)

CHARTER

Chairperson: Ross Callon, callon@erlang.dec.com

Mailing Lists:

isis@merit.edu
isis-request@merit.edu

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.

Specific Objectives:

1. 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.
2. Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
3. Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

Estimated Timeframe for Completion:

We intend to have completed objectives 1 and 2 by February, 1990.

CURRENT MEETING REPORT

The timeframe is being reconsidered by the working group chair.

CURRENT MEETING REPORT

Reported by Steven Willis/Wellfleet

MINUTES

Dave Oran updated the status of the ANSI IS-IS document. It had now reached Draft proposal stage and had been assigned the ISO number 10589. He discussed some of the recent modifications:

- The LOC-AREA portion of the address was now gone.
- Partition repair is now optional.
- Two multicast addresses had been assigned for level 1 and level 2 intermediate systems.

The working group addressed a number of issues outstanding from the previous meeting in Ann Arbor. In particular :

- Someone brought up the point that it may be limiting for an IP router to only encapsulate over 802.3. (i.e., doing IP IS-IS over HDLC). Do we want to consider changing encapsulation from 802.3 to IP to allow for the additional link-layer flexibility?
- The IP L1 partition repair is a bit flaky. Do we want to just say that this is not allowed for IP or do we want to fix it? (Not a straight-forward task).
- Dave Oran has corrections for IP routing exchange authentication. We didn't resolve what to do with an authentication mismatch. Drop the packet and what management information?
- Presently, the Integrated IP spec says that IP external links can just be generated by L2 routers. This limits the topology (it cannot have a pocket of rip routers anywhere in an area that aren't connected to a L2 router) and will make it hard to transition from another IGP to IP IS-IS. The suggestion was made to be able to generate IP external information at L1 - this is a good idea but potentially creates a problem when AS border gateways are at L1 but are in different areas. Since AS external information is not flooded into areas, L1 routers in different areas will not hear another AS's border router's external information - thus t, Cris BGP won't work in the transit AS case - We were not sure how this will affect other EGPs.
- There was some debate about the merits of forwarding based upon default metric if there is no path using desired metric. Forwarding as we currently allow may result in violation of policy. Dropping packet isn't friendly. No conclusions. Tony Lauck pointed out there were three solutions to handling

datagrams with TOS not supported by an intermediate system :

1. Let the packet disappear into a black hole
2. Map the TOS into another TOS supported by the IS
3. Drop the packet and generate an ICMP message back to source.

Solution 3 appeared to be the best answer, but it was unclear as what the ICMP type code should be.

- We discussed the possibility of running a partial Dykstra if only the leaves of the tree had changed. Dave Oran pointed out that this was an implementation issue and suggested that an Implementor's Hints Annex be added to the IS-IS specification to address these issues.

ATTENDEES

Almquist, Philip	almquist@jessica.stanford.edu
Bagnall, Doug	bagnall_d@apollo.hp.com
Baker, Fred	baker@vitalink.com
Bare, Ballard	bare%hprnd@hplabs.hp.com
Brunner, Ted	tob@thumper.bellcore.com
Bryant, Stewart	bryant@janus.enet.dec.com
Cerf, Vint	vcerf@nri.reston.va.us
Chapin, Lyman	lyman-chapin@dgc.mceo.dgcom
Chatterjee, Samir	samir@nynexst.com
Colella, Richard	colella@osi3.ncsl.nist.gov
Deboo, Farokh	sun!iruucp!ntrlink!fjd
Farinacci, Dino	dino@bridge2.3com.com
Ferguson, Dennis	dennis@gw.ccie.utoronto.ca
Froyd, Stan	sfroyd@salt.acc.com
Gan, Der-Hwa	no email
Goguely, Herve	rvg@bridge2.3com.com
Hagens, Rob	hagens@cs.wisc.edu
Hain, Tony	hain@nmfecc.arpa
Heinanen, Juha	jh@funet.fi
Hytry, Tom	tlh@iwlcs.att.com
Karels, Mike	karels@berkeley.edu
Katz, Dave	dkatz@merit.edu
Kellen, Daniel	kellen@eglin.af.mil
Lauck, Tony	lauck@dsmail.dec.com
Little, Mike	little@saic.com
Marcinkevicz, Mike	mdm@gumby.dsd.trw.com
Minshall, Greg	minshall@kinetics.kinetics.com

Mogul, Jeff	mogul@decwrl.dec.com
Morris, Dennis	morrisd@imo-uvax.dca.mil
O'Leary, Dave	oleary@umd5.umd.edu
Oran, David	oran@oran.dec.com
Perlman, Radia	no email
Reilly, Michael	reilly@nsl.dec.com
Rekhter, Yakov	yakov@ibm.com
Senum, Steve	sjs@network.com
Sheridan, Jim	jsherida@ibm.com
Sklower, Keith	sklower@okeeffe.berkeley.edu
Solensky, Frank	solensky@interlan.interlan.com
Staw, Tony	staw@marvin.enet.dec.com
Sturtevant, Allen	sturtevant@ccc.nmfecc.gov
Tsuchiya, Paul	tsuchiya@thumper.bellcore.com
Willis, Steven	swillis@wellfleet.com
Winkler, Linda	b32357@anlvm.ctd.anl.gov
Woodburn, Robert	woody@saic.com

3.7.2 Interconnectivity Working Group (iwg)

CHARTER

Chairperson: Guy Almes, almes@rice.edu

Mailing Lists: iwg@rice.edu

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.

Specific Objectives:

1. Continue development of the Border Gateway Protocol (BGP).
2. Continue development of a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
3. 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.

Estimated Timeframe for Completion:

1. We have a draft of a General BGP Usage document. We hope to have an initial draft-RFC version of that by Spring 1990. This General Usage document will be hard work.
2. We will write a Stub BGP Usage document. We hope to have this document in hand and use it in test deployments of BGP alongside the current EGP- based Inter-AS routing structures. Experience gained from this test deployment will guide the evolution of both the BGP protocol and the General Usage document.

CURRENT MEETING REPORT

Reported by Guy Almes/Rice University

MINUTES

At the IETF meeting at the University of Hawaii, we spent some time discussing whether to split the Interconnectivity Working Group into two parts:

- one group of limited duration that would work on the BGP protocol and its usage and deployment, and
- another ongoing group more concerned with the operational aspects of interconnectivity, especially with tuning the routing at major connection points in the Internet.

The result of these discussions is that IWG will split, and the protocol development half will retain the name Interconnectivity Working Group (and fall within the Routing Area of the IETF), while the other half will be called the Topology Engineering Working Group and fall within the Operations Area of the IETF. At the coming IETF at Florida State, IWG will meet all of Tuesday, and TEWG will take the slot originally allocated to IWG on Wednesday morning.

ATTENDEES

Almes, Guy	almes@rice.edu
Birger, Chet	cbirger@bbn.com
Curci, Raymond	curci@gw.scri.fsu.edu
England, Kent	kwe@bu.edu
Farinacci, Dino	dino@bridge2.3com.com
Ferguson, Donald	dennis@gw.ccie.utoronto.ca
Fernandez, Louis	1fernandez@bbn.com
Fidler, Mike	ts0026@ohstvma.ircc.ohio-state.edu
Finkelson, Dale	dmf@westie.unl.edu
Fuller, Vince	fuller@jessica.stanford.edu
Hahn, Jack	hahn@umd5.umd.edu
Hastings, Gene	hastings@psc.edu
Hays, Ken	hays@scri1.scri.fsu.edu
Honig, Jeffrey C.	sch@tcgould.tn.cornell.edu
Jordt, Dan	danj@cac.washington.edu
Katz, Dave	dkatz@merit.edu
Long, Dan	long@bbn.com

Mathis, Matt	mathis@pele.psc.edu
O'Leary, Dave	oleary@umd5.umd.edu
Pace, Donald	pace@fsu1.cc.fsu.edu
Rekhter, Yakov	yakov@ibm.com
Sheridan, Jim	jsherida@ibm.com
Tsuchiya, Paul	tsuchiya@thumper.bellcore.com
Veach, Ross	rrv@uiuc.edu
Winkler, Linda	b32357@anlvm.ctd.anl.gov
Wintringham, Dan	danw@igloo.osc.edu

3.7.3 Multicast routing OSPF Working Group (mospf)

CHARTER

Chairperson Steve Deering, deering@pescadero.stanford.edu

Mailing Lists:

mospf@devvax.tn.cornell.edu mospf-request@devvax.tn.cornell.edu

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).

Specific Objectives:

The new multicast routing version of OSPF will be documented in an RFC, and at least two independent implementations will be developed to demonstrate the new protocol's viability.

The working group will be of short duration, lasting only for (hopefully) three IETF meetings. Alot of the work will be done between meetings using electronic mail and the teleconferencing facilities.

Milestones:

- Tallahassee: Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular looking at 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 (such as extent of backward compatibility, and whether all routers in an AS need be capable of multicast routing).
- 2nd meeting: 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 University of Maryland as an example.
- 3rd meeting: 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.

CURRENT MEETING REPORT

Reported by John Moy/Proteon

MINUTES

This was the initial meeting of the MOSPF working group. John Moy presented a number of slides to introduce the subject of multicast routing . The slides also attempted to discuss the main areas where the OSPF protocol needs to be changed/extended in order to provide multicast support.

The slide discussion was broken up into the following areas:

- There was a short introduction into IP multicasting, including the Internet Group Management Protocol (IGMP, RFC 1112). IGMP is responsible for maintaining host group membership. Multicast routers must support IGMP (although in multicast OSPF only the Designated router will be actively sending/receiving IGMP messages). Through using IGMP, a multicast router knows which multicast destinations are active on its attached LANs, but does not need to keep track of which particular hosts are requesting them (this is a considerable savings).
- A brief (and very incomplete) survey of multicast routing was attempted. This was done through examination of Steve Deering's "Multicast routing in Internetworks and Extended LANs" paper. This paper discusses a number of multicast routing methodologies: an extension to the learning bridges to better support multicast, four separate Bellman-Ford multicast routing algorithms (arranged in order of increasing functionality and complexity) and a link-state multicast algorithm. Finally, a mix of these approaches is explored for internet-wide multicast (and the wild-card multicast group is introduced).
  The most functional Bellman-Ford multicast algorithm in Steve's paper has been implemnted for BSD UNIX and is documented in RFC 1075.
  It is expected that the multicast OSPF extensions will closely follow the link-state multicast algorithm in Steve's paper.
- The basic mechanism behind link-state multicasting was explored. A shortest-path tree is built having as root the multicast datagram's source. Those branches not containing the specified multicast destination are then pruned. This tree then yields the multicast datagram's path.
  At each hop, the multicast datagram is sent as a link-level multicast (or broadcast). For this reason, multicast routers receive all multicast packets (i.e., must open up their multicast filters).
  The above trees are built on demand, and the results are cached (see below).

- Cache entries are used to store the results of the above SPF calculation. For each tree built (there is potentially a different tree for each [source net, multicast destination] pair), a cache entry is formed. The cache entry specifies the interface expected to receive the datagram, and the set of interfaces that the datagram should be forwarded out.
  Note that this proposed cache structure is slightly different than the one in Steve's paper. First, it skips caching whole subtrees (implementation experience with OSPF shows that subtree caching is a lot of work), and it simplifies things by ignoring TTL.
  A separate tree (and a separate cache entry) will also possibly be calculated for each TOS value. The entire cache must be cleared on topology changes. When a group's membership changes, only cache entries pertaining to that destination need be flushed.
- The changes and additions to the OSPF protocol are expected to be slight. This is because OSPF already maintains a complete topological map of the routing domain, enabling the multicast tree calculation. The expected changes to OSPF include the following: 1) During the Dijkstra calculation, routing table entries (e.g., for net X) will be marked with their corresponding "transit node". This node will be the root for multicast tree calculations having net X as datagram source.
- Expected additions to OSPF include the following: 1) OSPF Designated Routers (DRs) will need to speak IGMP. 2) There will be a new link state advertisement (type 6) that routers will originate when there are multicast destinations on attached LANs. There will be separate advertisements for each multicast destination. These advertisements will be originated by DRs only. This additional link state information will enable tree pruning. 3) There will probably be a new bit in the router links advertisement indicating "multicast-capable". This allows some routers in the AS to decline multicast support. 4) The multicast tree calculation must be made deterministic; i.e., all routers must calculate the same tree in the presence of multiple equal-cost routes.
- Inter-area multicast will be a little more complicated. This is because when the datagram source is in another area, the local topology surrounding the source is hidden, inhibiting the multicast tree calculation. In this case, we propose forming a tree for each of the other areas. Each tree can be thought of as still being rooted at the datagram source; there will be a link from the datagram's source to each of the area's border routers. The cost of these links will be that advertised by the area border routers in their summary link advertisements (note reverse direction). The rest of the tree will (as in the intra-area case) deals only with the area's own topology.

```
 datagram source
        :
 ----------------
 :    :    :    :
 :    :    :    :
ABR  ABR  ABR  ABR
 |    |         |
             -------
             |
```

> A multicast datagram enters the area through the border border routers. The above scheme works only if area border routers are "wild-card" multicast receivers (i.e., receive all muticast packets).
> The logic for when the source of the datagram is exterior to the AS should be similar to the inter-area case.

The following questions were raised by the slides. Bob Hinden asked whether there were any estimates for the CPU usage consumed by the potentially large number of tree calculations. In a related question, Chuck Davin wondered about the expected distribution of host group memberships (e.g., number per LAN and lifetime). These questions were put off, hopefully for Steve Deering to answer.

Scott Brim questioned the behavior of the above inter-area multicasting scheme in the presence of asymmetric paths. He thought that there might be the possibility of packet looping. This issue should be looked into further.

Besides the above, the following issues were brought up:

- Should the multicast specification be written as a separate document, or should it simply depend on RFC 1131 (the OSPF specification)?
- If we do not require all of the routers to be multicast-capable, is the possibility of reduced functionality acceptable?
- How much backward compatibility should there be with the present OSPF protocol?
- Should we try to be more efficient in inter-area multicasting, and drop the requirement that border routers be wild-card multicast receivers?

FUTURE MEETINGS

We intend to begin writing the specification for OSPF multicast extensions. This will be done primarily through communications on the mospf mailing list (mospf@devvax.tn.cornell.ed

There will be a MOSPF WG meeting at the next IETF (Pittsburgh). Also, if enough progress is made between meeting, we will attempt to schedule the teleconferencing facilities.

ATTENDEES

John Cavanaugh	john.cavanaugh@stpaul.ncr.com
Rob Coltun	rcoltun@trantor.umd.edu
Samir K. Chatterjee	samir@nynexst.com
George Clapp	meritec!clapp@bellcore.bellcore.com
Daniel Kellen	kellen@odixie.enet.dec.com
Stan Froyd	sfroyd@salt.acc.com
James R. Davin	jrd@ptt.lcs.mit.edu
Douglas Bagnall	bagnall-d@apollo.hp.com
Metin Feridun	mferidun@bbn.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Ballard Bare	bare%hprnd@hplabs.hp.com
Cathy Aronson	cja@merit.edu
Bob Hinden	hinden@bbn.com
Frank Solensky	solensky@interlna.com
Dave Miller	dtm@mitre.org
Tim Seaver	tas@mcnc.org
Joel Replogle	replogel@ncsa.uiuc.edu
Tony Mason	mason@transarc.com
Farokh Deboo	fjd@interlink.com
Dino Farinacci	dino@bridge2.3com.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Dave O'Leary	oleary@noc.sura.net

3.7.4 Open Distance Vector IGP Working Group (odv)

CHARTER

Chairperson: Charles Hedrick/Rutgers University, hedrick@cs.rutgers.edu

Mailing Lists:

odv@rutgers.edu
odv-request@rutgers.edu

Description of Working Group:

The Open Distance Vector Working Group is chartered to sponsor working on distance vector based routing protocols, and related work.

Specific Objectives:

1. Produce RFC describing IGRP. Should be ready by spring 90.
2. Sponsor and review work comparing distance vector and SPF algorithms. Timing depends upon actions of funding agencies. This is probably at least a one-year task.
3. Design a new distance vector protocol. This is a long-term goal.

CURRENT MEETING REPORT

Did not meet.

3.7.5 Open Systems Routing Working Group (orwg)

CHARTER

Chairperson: Marianne Lepp/BBN, mlepp@bbn.com

Mailing List: open-rout-interest@bbn.com

Description of Working Group:

The Open Systems Routing Working Group is chartered to develop a policy-based AS-AS routing protocol that will accommodate large size and general topology.

Specific Objectives and Milestones:

- Architecture
- Draft Protocol Specification of key elements of the protocol

Estimated Timeframe for Completion:

May 1990

CURRENT MEETING REPORT

Reported by Marianne Lepp/BBN

MINUTES

The Open Routing Working Group met for a full day at the recent IETF meeting. The previous day, we gave a formal presentation of the architecture, which was meant to give the audience a general overview of inter-domain policy routing. During the morning ORWG session, we continued where the formal presentation left off, delving into the details of the architecture as presented in the Internet Draft. In particular, we discussed the path setup procedure, the dissemination of routing information throughout the Internet, and the methods used to reduce the amount of routing information that each router must maintain. The information reduction methods include routing at the administrative domain (AD) level and leaving the IGP to route within a domain, the virtual gateway abstraction, the "super" AD, and the hierarchical organization of route servers, where the position in the hierarchy indicates the scope of the information it maintains. The one big question still open at this point is exactly how we are going to represent addresses when there are super ADs, i.e. a hierarchy of ADs.

During the afternoon session, we discussed some of the differences between IDPR and BGP, but unfortunately no BGP experts were in attendance. We were and still are hoping for critical review of the inter-domain policy routing architecture by some BGP experts. In fact, anyone with comments on the Internet Draft, please mail them to msteenst@bbn.com.

ATTENDEES

Chatterjee, Samir	samir@nynexst.com
Clapp, George	meritec!clapp@bellcore.bellcore.com
Farinacci, Dino	dino@bridge2.3com.com
Gross, Phill	pgross@nri.reston.va.us
Hinden, Bob	hinden@bbn.com
Jacobson, Van	van@helios.ee.lbl.gov
Little, Mike	little@saic.com
McKenney, Paul E.	mckenney@sri.com
Medin, Milo	medin@nsipo.nasa.gov
Moy, John	jmoy@proteon.com
Su, Zaw-Sing	zsu@tsca.istc.sri.com
Wintringham, Dan	danw@igloo.osc.edu

3.7.6 PDN Routing Working Group (pdnrout)

CHARTER

Chairperson: Carl-Herbert Rokitansky/Fern University of Hagen
roki@DHAFEU52.BITNET or roki@ISI.EDU

Mailing Lists:

- pdn-wg@BBN.COM: For internal discussions and information exchange between members of the PDN Routing working group.
- pdn-interest@BBN.COM: For information about:
  - Status report and proceedings of the PDN Routing WG
  - Draft proposals of documents and papers
  - Documents and papers published by PDN WG members
  - Important discussion on PDN Routing issues.
- pdn-request@BBN.COM: For people interested in being put on the "pdn-interest" mailing list.

Description of Working Group:

The DoD INTERNET TCP/IP protocol suite has developed into de facto industry standard for heterogenous packet switching computer networks. In the US, several hundreds of INTERNET networks are connected together; however the situation is completely different in Europe: The only network which could be used as a backbone to allow interoperation between the many local area networks in Europe, now subscribing to the DoD INTERNET TCP/IP protocol suite, would be the system of Public Data Networks (PDN). However, so far, no algorithms have been provided to dynamically route INTERNET datagrams through X.25 public data networks. Therefore, the goals of the Public Data Network Routing working group are the development, definition and specification of required routing and gateway algorithms for an improved routing of INTERNET datagrams through the system of X.25 Public Data Networks (PDN) to allow worldwide interoperation between TCP/IP networks in various countries. In addition, the application and/or modification of the developed algorithms to interconnect local TCP/IP networks via ISDN (Integrated Services Digital Network) will be considered.

Specific Objectives and Estimated Timeframe for Completion:

1. Application of the INTERNET Cluster Addressing Scheme to Public Data Net-

works. (Already done, see produced documents)

2. Development of hierarchical VAN-gateway algorithms for worldwide INTERNET network reachability information exchange between VAN-gateways (Already done, see produced documents)
3. Assignment of INTERNET/PDN-cluster network numbers to national public data networks. (Mapping between INTERNET network numbers and X.121 Data Network Identification Codes (DNICs) (Already done, see produced documents)
4. Assignment of INTERNET/PDN-cluster addresses to PDN-hosts and VAN-gateways according to the developed hierarchical VAN-gateway algorithms (Almost done, see produced documents)
5. Definition of the PDN-cluster addressing scheme as an Internet standard (Already done, [earlier than expected - a case that happens very seldom!] see produced documents)
6. Specification of an X.121 Address resolution protocol (RFC-Draft, expected to be completed by October '89)
7. Specification of an X.25 Call Setup and Charging Determination Protocol (RFC-Draft, expected to be completed by Fall '89)
8. Specification of an X.25 Access and Forwarding Control Scheme (to be written up as an RFC-Draft by Fall '89 or later)
9. Specification of routing metrics taking X.25 charges into account (to be written up as an RFC-Draft by Fall '89 or later)
10. Delayed TCP/IP header compression by VAN-gateways and PDN-hosts (new objective, will be considered Fall '89 or later)
11. Provide a testbed for worldwide interoperability between local TCP/IP networks via the system of X.25 public data networks (PDN) (starting June '89)
12. Implementation of the required algorithms and protocols in a VAN-BoX (Test version towards End '89)
13. Interoperability between ISO/OSI hosts on TCP/IP networks through PDN (1989/90)
14. Consideration of INTERNET Route Servers (1990)
15. Interoperability between local TCP/IP networks via ISDN (1990)
16. Development of Internetwork Management Protocols for worldwide cooperation and coordination of network control and network information centers (starting 1990).

CURRENT MEETING REPORT

Did not meet.

3.8 Security Area

Director: Steve Crocker/TIS

The new Security Policy WG, chaired by Rich Pethia, met at the IETF meeting in Florida. There was a considerable interest. The WG will propose ideas for an Internet-wide security policy. A mailing list has been established. Send requests to:

- spwg-request@nri.reston.va.us

A number of messages have already been sent on this list, and the ideas are flowing rapidly.

The SNMP Authentication portion of the Authentication WG met in Florida and discussed a trio of documents. These documents will continue to undergo further review, but have been released for general distribution with the intention of becoming a proposed standard (elective). The three documents are:

- "Authentication and Privacy in the SNMP"
- "Administration of SNMP Communities"
- "Experimental Definitions of Managed Objects for Administration of SNMP Communities"

Keith McCloghrie, Chuck Davin and Jim Galvin are to be congratulated for pushing through these documents.

The IP authentication portion of the Authentication WG did not meet, but its document is complete and will be submitted to the RFC editor for advancement to Proposed Standard (Elective).

Some security related topics have come up that are being pursued in other areas. This is expected to happen reasonably frequently, and our intent is to leave the primary responsibility with the other area and coordinate as needed. Specific topics being coordinated at the moment are:

- User profile, under development by the User Services WG, chaired by Joyce Reynolds
- Telnet encryption and authentication, under development by the Telnet WG, chaired by Dave Borman.
- Privacy Enhanced Mail, under development by the Privacy and Security Work-

ing Group in the IRTF, chaired by Steve Kent.

3.8.1 IP Authentication Working Group (ipauth)

CHARTER

Chairperson: Jeffrey Schiller/MIT, jis@bitsy.mit.edu

Mailing List: awg@bitsy.mit.edu

Description of Working Group:

To brainstorm issues relating to providing for the security and integrity of information on the Internet, with emphasis on those protocols used to operate and control the network. To propose open standard solutions to problems in network authentication.

Specific Objectives:

1. RFC specifying an authentication format which supports multiple authentication systems.
2. Document discussing the cost/benefit tradeoffs of various generic approaches to solving the authentication problem in the Internet context.
3. Document to act as a protocol designers guide to authentication.
4. RFC proposing A Key Distribution System (emphasis on "A" as opposed to "THE"). MIT's Kerberos seems the most likely candidate here.

Estimated Timeframe for Completion:

This working group will hopefully complete its current objectives within one year. At this point the group will either disband or will move on to other related problems/issues.

CURRENT MEETING REPORT

Did not meet.

3.8.2 Security Policy Working Group (spwg)

CHARTER

Chairperson: Rich Pethia, rdp@sei.cmu.edu

Mailing lists:

spwg@nri.reston.va.us
sp̊wg-request@nri.reston.va.us

Description of Working Group:

The Security Policy Working Group 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 adminstrative issues related to security, including integrity, authentication and confidentiality controls, and administration of hosts and networks.

Objectives and Milestones:

Among the issues to be considered in this working group are:

- Responsibilities and obligations of users, data base 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 Interenet components.

Specific steps that will be taken are:

1. First Meeting: 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 to be reviewed at the May IETF.
2. May IETF Meeting: revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.
3. In the July 1990 timeframe, prepare first draft policy recommendation for work-

ing group review and modification.

4. In the early September 1990 timeframe, finalize draft policy and initiate review following standard RFC procedure.

3.8.3 SNMP Authentication Working Group (snmpauth)

CHARTER

Chairperson: Jeffrey Schiller/MIT, jis@bitsy.mit.edu

Mailing List:

awg@bitsy.mit.edu
awg-request@bitsey.mit.edu

Description of Working Group:

To define a standard mechanism for authentication within the SNMP.

Specific Objective:

To write an RFC specifying procedures and formats for providing standardized authentication within the SNMP.

Estimated Timeframe for Completion:

May 1990

CURRENT MEETING REPORT

Reported by James Davin/MIT

AGENDA

The business of the meeting was the consideration of three documents describing proposed mechanisms for authenticating SNMP management operations:

1. Galvin, McCloghrie, and Davin. Authentication and Privacy in the SNMP.
2. Davin, Galvin, and McCloghrie. Administration of SNMP Communities.
3. McCloghrie, Davin, and Galvin. Experimental Definitions of Managed Objects for Administration of SNMP Communities..

MINUTES

The first portion of the meeting was devoted to presentations by Jim Galvin and Keith McCloghrie that summarized the substance of the three documents.

These presentations were followed by a lively discussion of relevant issues:

1. Timeliness — The single issue that elicited the most discussion was the problem of ensuring the "timeliness" of messages exchanged in the protocol. Concerns were voiced about several aspects of this problem:
   (a) The implications of the described timeliness mechanisms with respect to authentication communities of more than one management station need to be clarified.
   (b) Concern was voiced about potential problems involved with the setting of community clock values as described in the administration document. The idea was expressed that this mechanism for clock synchronization may be suboptimal in terms of both the state required in an agent and vulnerability to denial of service attack.
   (c) The possibility that subnet duplication of protocol messages could entail reversal of a community clock was suggested.
   (d) The implications of clock drift for the protocol were discussed. Although some concern remains on this topic, many felt that signficant problems are associated only with clock drifts several orders of magnitude larger than those typically experienced.
2. Key Distribution Options — The limitations of using SNMP as a key distribution mechanism were recognized, and the possibility of exploring other mechanisms was suggested. In particular, the role of multiple management stations

in the key distribution process needs clarification.

3. Algorithm Correctness — The desirability of algorithms with either sound formal foundations or reputations based on broad review and experience was noted. The desirability of citing relevant literature was also noted.
4. Liability Issues — Concern was raised about the legal liabilities that may accrue to a promulgating standards body by its choice of an algorithm.
5. Coordination — the desirability of coordinating this effort with other relevant efforts was acknowledged.

The meeting concluded with a consensus that the three documents should be introduced into the IETF process for consideration as possible standards with Elective status.

Action Items:

1. Chuck agreed to prepare minutes and be responsible for augmenting the AWG mailing list (awg@bitsy.mit.edu) to reflect any newcomers to the effort.
2. The document authors agreed to revise their documents to reflect the concerns raised at this meeting and to (re-)introduce them into the IETF Drafts repository for further review.

ATTENDEES

Doug Bagnall	bagnall_d@apollo.hp.com
Scott Bradner	sob@harvisr.harvard.edu
Ted Brunner	tob@thumper.bellcore.com
Jeff Case	case@utkcs.cs.utk.edu
Steve Crocker	crocker@tis.com
James R. Davin	jrd@ptt.lcs.mit.edu
Stan Froyd	sfroyd@salt.acc.com
James M. Galvin	galvin@tis.com
Steven Hunter	hunter@ccc.nmfecc.gov
Phil Jensen	jensen@fsu1.cc.fsu.edu
Tony Lauck	lauck@tl.enet.dec.com
Walt Lazear	lazear@gateway.mitre.org
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Greg Minshall	minshall@kinetics.com
Jeff Mogul	mogul@decwrl.dec.com
Dave Monaebello	dave@pluto.dss.com
Oscar Newkerk	newkerk@decwet.enet.dec.com
Dave Perkins	dave_perkins@3com.com

Jim Robertson	jar@esd.3com.com
Jon Saperia	saperia@tcpjon.enet.dec.com
Tom Seaver	tas@mcnc.org
Frank Solensky	solensky@interlan.com
Mike St. Johns	stjohns@umd5.umd.edu
Dean Throop	throop@dg-rtp.dg.com
Sudhanshu Verma	verma@hpindbu.hp.com
Steve Waldbusser	waldbusser@andrew.cmu.edu
Brian Yasaki	bky@twg.com

3.8.4 Site Security Policy Handbook Working Group (ssphwg)

CHARTER

Chairpersons: Paul Holbrook/CERT ph@SEI.CMU.EDU
Joyce K. Reynolds/USC-ISI 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.

Objectives and Milestones:

- After the group is announced and interested people are on the list, Holbrook will distribute current ideas about the handbook and the outline.
- First IETF Meeting (May 1990 - PSC): 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 a SSPHWG "editorial board" for future writing assignments for the first draft of document.
- Around the June USENIX in California: Finalize outline and organization of handbook. Partition out pieces to interested parties and SSPHWG editorial board members.
- Second IETF Meeting (August 1990 - UBC): In the early August 1990 timeframe, pull together a first draft handbook for working group review and modification.
- In the October 1990 timeframe, finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor for publication.

Chapter 4

Network Status Briefings

4.1 "State of the Internet"

Presentation by Chet Birger/BBN

STATE OF THE INTERNET

Chet Birger

February 7, 1990

BBN Communications Corporation

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH

INTERNET GROWTH SUMMARY

- 1202 Networks advertised
- 2203 Networks registered

NUMBER OF NETWORKS
December 1983-January 1990

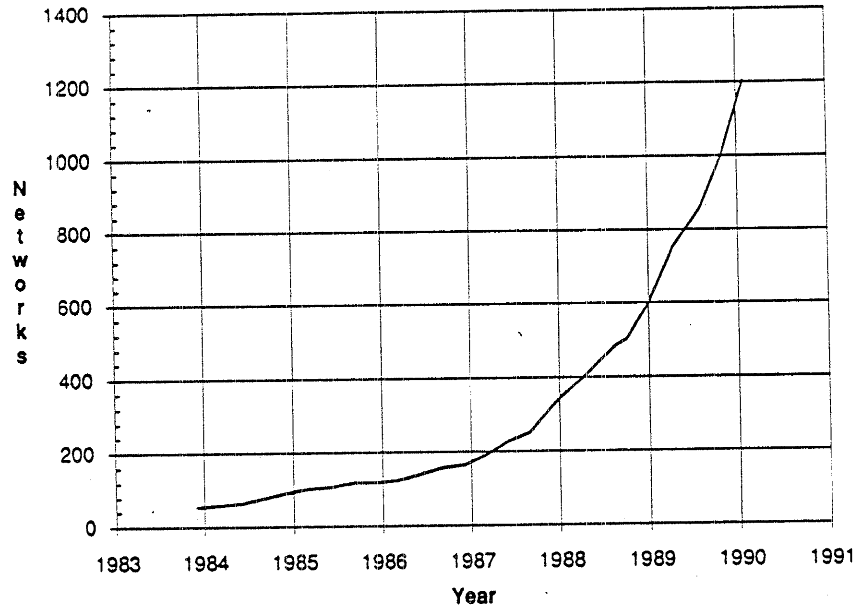

NUMBER OF NETWORKS
December 1983-January 1990

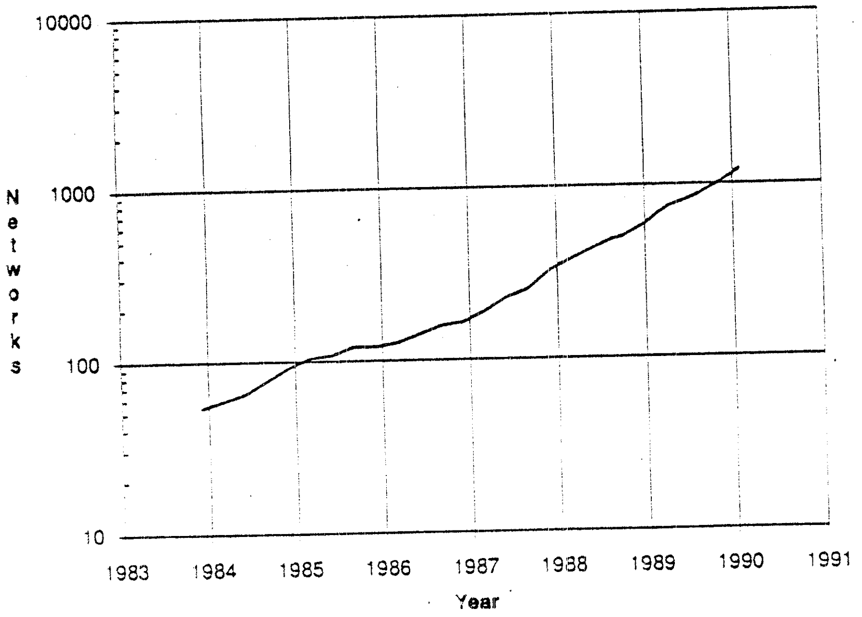

ADVERTISED NETWORKS

```
130.33.0.0     130.37.0.0     130.38.0.0     130.39.0.0
130.41.0.0     130.42.0.0     130.43.0.0     130.44.0.0
130.46.0.0     130.49.0.0     130.50.0.0     130.53.0.0
130.54.0.0     130.56.0.0     130.57.0.0     130.58.0.0
130.62.0.0     130.63.0.0     130.68.0.0     130.69.0.0
130.70.0.0     130.71.0.0     130.74.0.0     130.84.0.0
130.85.0.0     130.86.0.0     130.87.0.0     130.90.0.0
130.91.0.0     130.93.0.0     130.94.0.0     130.95.0.0
130.101.0.0    130.105.0.0    130.108.0.0    130.113.0.0
130.114.0.0    130.116.0.0    130.117.0.0    130.118.0.0
130.123.0.0    130.126.0.0    130.127.0.0    130.132.0.0
130.134.0.0    130.136.0.0    130.137.0.0    130.150.0.0
130.152.0.0    130.157.0.0    130.160.0.0    130.163.0.0
130.166.0.0    130.167.0.0    130.168.0.0    130.180.0.0
130.182.0.0    130.184.0.0    130.185.0.0    130.186.0.0
130.188.0.0    130.189.0.0    130.191.0.0    130.192.0.0
130.194.0.0    130.195.0.0    130.199.0.0    130.202.0.0
130.203.0.0    130.204.0.0    130.207.0.0    130.208.0.0
130.212.0.0    130.215.0.0    130.217.0.0    130.219.0.0
130.220.0.0    130.221.0.0    130.225.0.0    130.232.0.0
130.233.0.0    130.235.0.0    130.236.0.0    130.237.0.0
130.238.0.0    130.239.0.0    130.240.0.0    130.245.0.0
130.251.0.0    130.252.0.0    130.253.0.0    131.1.0.0
131.2.0.0      131.4.0.0      131.5.0.0      131.6.0.0
131.7.0.0      131.8.0.0      131.9.0.0      131.10.0.0
131.11.0.0     131.12.0.0     131.13.0.0     131.14.0.0
131.15.0.0     131.17.0.0     131.18.0.0     131.19.0.0
131.20.0.0     131.21.0.0     131.22.0.0     131.23.0.0
131.24.0.0     131.25.0.0     131.26.0.0     131.27.0.0
131.28.0.0     131.29.0.0     131.31.0.0     131.32.0.0
131.33.0.0     131.34.0.0     131.35.0.0     131.36.0.0
131.37.0.0     130.236.0.0    130.237.0.0    130.238.0.0
130.239.0.0    130.240.0.0    130.245.0.0    130.251.0.0
130.252.0.0    130.253.0.0    131.1.0.0      131.2.0.0
131.4.0.0      131.5.0.0      131.6.0.0      131.7.0.0
131.8.0.0      131.9.0.0      131.10.0.0     131.11.0.0
131.12.0.0     131.13.0.0     131.14.0.0     131.15.0.0
131.17.0.0     131.18.0.0     131.19.0.0     131.20.0.0
131.21.0.0     131.22.0.0     131.23.0.0     131.24.0.0
131.25.0.0     131.26.0.0     131.27.0.0     131.28.0.0
```

```
131.29.0.0     131.31.0.0     131.32.0.0     131.33.0.0
131.34.0.0     131.35.0.0     131.36.0.0     131.37.0.0
131.38.0.0     131.39.0.0     131.40.0.0     131.41.0.0
131.42.0.0     131.43.0.0     131.44.0.0     131.45.0.0
131.46.0.0     131.47.0.0     131.49.0.0     131.50.0.0
131.51.0.0     131.52.0.0     131.53.0.0     131.55.0.0
131.56.0.0     131.57.0.0     131.58.0.0     131.59.0.0
131.60.0.0     131.61.0.0     131.62.0.0     131.63.0.0
131.65.0.0     131.70.0.0     131.72.0.0     131.74.0.0
131.76.0.0     131.77.0.0     131.78.0.0     131.79.0.0
131.80.0.0     131.81.0.0     131.83.0.0     131.84.0.0
131.85.0.0     131.87.0.0     131.92.0.0     131.93.0.0
131.95.0.0     131.103.0.0    131.105.0.0    131.106.0.0
131.108.0.0    131.112.0.0    131.113.0.0    131.114.0.0
131.118.0.0    131.119.0.0    131.120.0.0    131.123.0.0
131.128.0.0    131.131.0.0    131.132.0.0    131.133.0.0
131.145.0.0    131.146.0.0    131.151.0.0    131.154.0.0
131.155.0.0    131.156.0.0    131.158.0.0    131.167.0.0
131.170.0.0    131.171.0.0    131.172.0.0    131.177.0.0
131.178.0.0    131.179.0.0    131.182.0.0    131.183.0.0
131.186.0.0    131.187.0.0    131.188.0.0    131.192.0.0
131.193.0.0    131.196.0.0    131.199.0.0    131.203.0.0
131.204.0.0    131.210.0.0    131.211.0.0    131.212.0.0
131.214.0.0    131.215.0.0    131.216.0.0    131.218.0.0
131.225.0.0    131.226.0.0    131.230.0.0    131.236.0.0
131.239.0.0    131.240.0.0    131.243.0.0    131.246.0.0
131.247.0.0    131.249.0.0    131.250.0.0    131.254.0.0
132.1.0.0      132.2.0.0      132.3.0.0      132.4.0.0
132.5.0.0      132.6.0.0      132.7.0.0      132.8.0.0
132.9.0.0      132.10.0.0     132.12.0.0     132.13.0.0
132.15.0.0     132.16.0.0     132.17.0.0     132.18.0.0
132.19.0.0     132.21.0.0     132.23.0.0     132.25.0.0
132.27.0.0     132.29.0.0     132.30.0.0     132.31.0.0
132.32.0.0     132.35.0.0     132.36.0.0     132.37.0.0
132.39.0.0     132.40.0.0     132.41.0.0     132.42.0.0
132.44.0.0     132.45.0.0     132.46.0.0     132.47.0.0
132.48.0.0     132.49.0.0     132.50.0.0     132.51.0.0
132.52.0.0     132.53.0.0     132.54.0.0     132.56.0.0
132.57.0.0     132.58.0.0     132.60.0.0     132.61.0.0
132.62.0.0     132.63.0.0     132.68.0.0     132.144.0.0
```

ADVERTISED NETWORKS

```
5.0.0.0        7.0.0.0        10.0.0.0       13.0.0.0
14.0.0.0       15.0.0.0       16.0.0.0       18.0.0.0
26.0.0.0       28.0.0.0       31.0.0.0       35.0.0.0
36.0.0.0       46.0.0.0       90.0.0.0       128.2.0.0
128.3.0.0      128.4.0.0      128.5.0.0      128.6.0.0
128.7.0.0      128.8.0.0      128.9.0.0      128.10.0.0
128.11.0.0     128.16.0.0     128.18.0.0     128.19.0.0
128.20.0.0     128.29.0.0     128.32.0.0     128.35.0.0
128.36.0.0     128.38.0.0     128.39.0.0     128.41.0.0
128.42.0.0     128.43.0.0     128.44.0.0     128.45.0.0
128.46.0.0     128.47.0.0     128.49.0.0     128.52.0.0
128.54.0.0     128.55.0.0     128.58.0.0     128.59.0.0
128.60.0.0     128.61.0.0     128.62.0.0     128.63.0.0
128.81.0.0     128.82.0.0     128.83.0.0     128.84.0.0
128.86.0.0     128.89.0.0     128.91.0.0     128.92.0.0
128.93.0.0     128.95.0.0     128.96.0.0     128.97.0.0
128.99.0.0     128.100.0.0    128.101.0.0    128.102.0.0
128.103.0.0    128.104.0.0    128.105.0.0    128.109.0.0
128.110.0.0    128.111.0.0    128.112.0.0    128.113.0.0
128.114.0.0    128.115.0.0    128.116.0.0    128.117.0.0
128.118.0.0    128.119.0.0    128.120.0.0    128.121.0.0
128.122.0.0    128.123.0.0    128.124.0.0    128.125.0.0
128.126.0.0    128.127.0.0    128.128.0.0    128.133.0.0
128.135.0.0    128.136.0.0    128.138.0.0    128.139.0.0
128.140.0.0    128.141.0.0    128.143.0.0    128.144.0.0
128.145.0.0    128.146.0.0    128.147.0.0    128.148.0.0
128.149.0.0    128.150.0.0    128.151.0.0    128.152.0.0
128.153.0.0    128.154.0.0    128.155.0.0    128.156.0.0
128.157.0.0    128.158.0.0    128.159.0.0    128.160.0.0
128.162.0.0    128.163.0.0    128.164.0.0    128.165.0.0
128.167.0.0    128.168.0.0    128.169.0.0    128.170.0.0
128.171.0.0    128.172.0.0    128.173.0.0    128.174.0.0
128.175.0.0    128.180.0.0    128.181.0.0    128.182.0.0
128.183.0.0    128.185.0.0    128.186.0.0    128.187.0.0
128.188.0.0    128.189.0.0    128.190.0.0    128.191.0.0
128.193.0.0    128.194.0.0    128.195.0.0    128.196.0.0
128.197.0.0    128.200.0.0    128.204.0.0    128.205.0.0
128.206.0.0    128.209.0.0    128.210.0.0    128.211.0.0
128.212.0.0    128.213.0.0    128.214.0.0    128.215.0.0
128.217.0.0    128.218.0.0    128.219.0.0    128.220.0.0
```

```
128.221.0.0    128.222.0.0    128.223.0.0    128.226.0.0
128.227.0.0    128.228.0.0    128.229.0.0    128.230.0.0
128.231.0.0    128.235.0.0    128.236.0.0    128.237.0.0
128.238.0.0    128.239.0.0    128.241.0.0    128.242.0.0
128.244.0.0    128.245.0.0    128.247.0.0    128.248.0.0
128.249.0.0    128.250.0.0    128.252.0.0    128.253.0.0
128.255.0.0    129.1.0.0      129.3.0.0      129.4.0.0
129.5.0.0      129.6.0.0      129.7.0.0      129.8.0.0
129.10.0.0     129.13.0.0     129.15.0.0     129.16.0.0
129.18.0.0     129.19.0.0     129.20.0.0     129.22.0.0
129.24.0.0     129.25.0.0     129.26.0.0     129.29.0.0
129.30.0.0     129.32.0.0     129.33.0.0     129.34.0.0
129.43.0.0     129.46.0.0     129.48.0.0     129.49.0.0
129.55.0.0     129.57.0.0     129.59.0.0     129.60.0.0
129.61.0.0     129.62.0.0     129.63.0.0     129.64.0.0
129.65.0.0     129.66.0.0     129.69.0.0     129.71.0.0
129.72.0.0     129.73.0.0     129.74.0.0     129.75.0.0
129.77.0.0     129.79.0.0     129.81.0.0     129.82.0.0
129.83.0.0     129.84.0.0     129.85.0.0     129.87.0.0
129.88.0.0     129.89.0.0     129.91.0.0     129.92.0.0
129.93.0.0     129.95.0.0     129.96.0.0     129.97.0.0
129.99.0.0     129.100.0.0    129.101.0.0    129.104.0.0
129.105.0.0    129.106.0.0    129.107.0.0    129.109.0.0
129.110.0.0    129.111.0.0    129.112.0.0    129.114.0.0
129.116.0.0    129.117.0.0    129.118.0.0    129.119.0.0
129.120.0.0    129.121.0.0    129.122.0.0    129.123.0.0
129.125.0.0    129.126.0.0    129.127.0.0    129.128.0.0
129.130.0.0    129.131.0.0    129.133.0.0    129.137.0.0
129.138.0.0    129.139.0.0    129.140.0.0    129.142.0.0
129.143.0.0    129.162.0.0    129.170.0.0    129.172.0.0
129.174.0.0    129.175.0.0    129.176.0.0    129.177.0.0
129.179.0.0    129.186.0.0    129.188.0.0    129.189.0.0
129.190.0.0    129.191.0.0    129.192.0.0    129.193.0.0
129.199.0.0    129.206.0.0    129.207.0.0    129.209.0.0
129.213.0.0    129.216.0.0    129.217.0.0    129.219.0.0
129.229.0.0    129.235.0.0    129.236.0.0    129.237.0.0
129.238.0.0    129.240.0.0    129.241.0.0    129.243.0.0
129.244.0.0    129.246.0.0    129.248.0.0    129.250.0.0
129.252.0.0    130.11.0.0     130.13.0.0     130.14.0.0
130.15.0.0     130.17.0.0     130.18.0.0     130.20.0.0
```

ADVERTISED NETWORKS

```
132.151.0.0    132.154.0.0    132.158.0.0    132.159.0.0        192.5.65.0     192.5.82.0     192.5.88.0     192.5.104.0
132.160.0.0    132.162.0.0    132.163.0.0    132.174.0.0        192.5.105.0    192.5.109.0    192.5.131.0    192.5.144.0
132.175.0.0    132.178.0.0    132.183.0.0    132.192.0.0        192.5.146.0    192.5.148.0    192.5.157.0    192.5.171.0
132.193.0.0    132.194.0.0    132.197.0.0    132.198.0.0        192.5.172.0    192.5.174.0    192.5.175.0    192.5.178.0
132.202.0.0    132.203.0.0    132.204.0.0    132.205.0.0        192.5.192.0    192.5.195.0    192.5.196.0    192.5.197.0
132.206.0.0    132.207.0.0    132.208.0.0    132.209.0.0        192.5.198.0    192.5.199.0    192.5.200.0    192.5.[illegible].0
132.211.0.0    132.230.0.0    132.235.0.0    132.236.0.0        192.5.204.0    192.5.210.0    192.5.211.0    192.5.213.0
132.238.0.0    132.239.0.0    132.241.0.0    132.245.0.0        192.5.214.0    192.5.216.0    192.5.218.0    192.5.219.0
132.246.0.0    132.249.0.0    132.250.0.0    132.251.0.0        192.5.237.0    192.5.238.0    192.5.241.0    192.5.[illegible].0
132.254.0.0    133.1.0.0      133.2.0.0      133.4.0.0          192.9.9.0      192.12.5.0     192.12.7.0     192.12.8.0
133.5.0.0      133.6.0.0      133.11.0.0     133.137.0.0        192.12.9.0     192.12.10.0    192.12.12.0    192.12.17.0
133.138.0.0    134.3.0.0      134.4.0.0      134.9.0.0          192.12.18.0    192.12.19.0    192.12.20.0    192.12.30.0
134.12.0.0     134.20.0.0     134.22.0.0     134.24.0.0         192.12.33.0    192.12.51.0    192.12.62.0    192.12.63.0
134.48.0.0     134.50.0.0     134.53.0.0     134.55.0.0         192.12.64.0    192.12.65.0    192.12.66.0    192.12.67.0
134.57.0.0     134.60.0.0     134.62.0.0     134.63.0.0         192.12.68.0    192.12.69.0    192.12.87.0    192.12.97.0
134.68.0.0     134.69.0.0     134.74.0.0     134.78.0.0         192.12.98.0    192.12.100.0   192.12.119.0   192.12.122.0
134.79.0.0     134.87.0.0     134.95.0.0     134.114.0.0        192.12.124.0   192.12.125.0   192.12.126.0   192.12.141.0
134.117.0.0    134.121.0.0    134.126.0.0    134.129.0.0        192.12.171.0   192.12.184.0   192.12.192.0   192.12.193.0
134.131.0.0    134.132.0.0    134.139.0.0    134.141.0.0        192.12.195.0   192.12.199.0   192.12.206.0   192.12.215.0
134.149.0.0    134.160.0.0    134.165.0.0    134.172.0.0        192.12.216.0   192.12.236.0   192.12.245.0   192.12.249.0
134.173.0.0    134.174.0.0    134.177.0.0    134.185.0.0        192.12.250.0   192.16.13.0    192.16.72.0    192.16.123.0
134.192.0.0    134.197.0.0    134.202.0.0    134.205.0.0        192.16.165.0   192.16.166.0   192.16.169.0   192.16.173.0
134.207.0.0    134.228.0.0    134.231.0.0    134.240.0.0        192.16.174.0   192.16.175.0   192.16.183.0   192.16.184.0
136.145.0.0    136.149.0.0    136.160.0.0    136.161.0.0        192.16.186.0   192.16.188.0   192.16.192.0   192.16.194.0
136.176.0.0    136.177.0.0    136.200.0.0    136.205.0.0        192.16.199.0   192.16.201.0   192.16.202.0   192.16.204.0
136.229.0.0    136.242.0.0    136.246.0.0    136.247.0.0        192.16.205.0   192.16.206.0   192.16.207.0   192.16.208.0
137.2.0.0      137.3.0.0      137.4.0.0      137.22.0.0         192.24.225.0   192.24.239.0   192.26.8.0     192.26.10.0
137.29.0.0     137.39.0.0     137.41.0.0     137.80.0.0         192.26.12.0    192.26.18.0    192.26.20.0    192.26.25.0
137.83.0.0     137.90.0.0     137.95.0.0     137.103.0.0        192.26.27.0    192.26.83.0    192.26.86.0    192.26.87.0
137.152.0.0    137.192.0.0    137.209.0.0    137.210.0.0        192.26.88.0    192.26.89.0    192.26.91.0    192.26.93.0
137.211.0.0    192.1.4.0      192.1.5.0      192.1.6.0          192.26.103.0   192.26.147.0   192.26.148.0   192.26.210.0
192.1.7.0      192.1.8.0      192.1.10.0     192.1.17.0         192.31.3.0     192.31.7.0     192.31.17.0    192.31.21.0
192.1.19.0     192.1.20.0     192.1.22.0     192.1.23.0         192.31.24.0    192.31.27.0    192.31.30.0    192.31.39.0
192.1.24.0     192.1.25.0     192.4.13.0     192.5.8.0          192.31.44.0    192.31.45.0    192.31.63.0    192.31.66.0
192.5.9.0      192.5.10.0     192.5.11.0     192.5.14.0         192.31.68.0    192.31.70.0    192.31.71.0    192.31.75.0
192.5.18.0     192.5.21.0     192.5.22.0     192.5.23.0         192.31.82.0    192.31.83.0    192.31.85.0    192.31.87.0
192.5.24.0     192.5.25.0     192.5.28.0     192.5.29.0         192.31.88.0    192.31.95.0    192.31.100.0   192.31.101.0
192.5.39.0     192.5.41.0     192.5.46.0     192.5.47.0         192.31.103.0   192.31.104.0   192.31.106.0   192.31.111.0
192.5.49.0     192.5.53.0     192.5.54.0     192.5.56.0         192.31.112.0   192.31.146.0   192.31.147.0   192.31.152.0
192.5.57.0     192.5.58.0     192.5.59.0     192.5.60.0         192.31.153.0   192.31.154.0   192.31.163.0   192.31.164.0
```

ADVERTISED NETWORKS

```
192.31.165.0   192.31.172.0   192.31.173.0   192.31.174.0       192.48.105.0   192.48.125.0   192.48.139.0   192.48.143.0
192.31.177.0   192.31.178.0   192.31.180.0   192.31.181.0       192.48.212.0   192.48.214.0   192.48.215.0   192.48.217.0
192.31.192.0   192.31.197.0   192.31.210.0   192.31.211.0       192.48.219.0   192.48.223.0   192.52.61.0    192.52.64.0
192.31.214.0   192.31.215.0   192.31.222.0   192.31.223.0       192.52.65.0    192.52.66.0    192.52.78.0    192.52.[illegible].0
192.31.225.0   192.31.230.0   192.31.231.0   192.31.238.0       192.52.106.0   192.52.107.0   192.52.111.0   192.52.112.0
192.31.239.0   192.31.242.0   192.31.245.0   192.31.246.0       192.52.117.0   192.52.154.0   192.52.156.0   192.52.179.0
192.31.253.0   192.31.254.0   192.33.4.0     192.33.5.0         192.52.180.0   192.52.194.0   192.52.195.0   192.52.219.0
192.33.13.0    192.33.14.0    192.33.19.0    192.33.33.0        192.52.232.0   192.52.236.0   192.54.33.0    192.54.81.0
192.33.36.0    192.33.112.0   192.33.115.0   192.33.116.0       192.54.104.0   192.54.105.0   192.54.106.0   192.54.109.0
192.33.140.0   192.33.141.0   192.33.144.0   192.33.145.0       192.54.129.0   192.54.138.0   192.54.226.0   192.54.240.0
192.33.146.0   192.33.148.0   192.33.149.0   192.33.153.0       192.54.252.0   192.55.81.0    192.55.90.0    192.55.97.0
192.33.156.0   192.33.159.0   192.33.167.0   192.33.168.0       192.55.98.0    192.55.103.0   192.55.109.0   192.55.117.0
192.33.170.0   192.33.178.0   192.33.179.0   192.33.181.0       192.55.120.0   192.55.134.0   192.55.187.0   192.55.207.0
192.33.182.0   192.33.185.0   192.35.44.0    192.35.49.0        192.55.214.0   192.55.245.0   192.55.246.0   192.58.91.0
192.35.62.0    192.35.74.0    192.35.75.0    192.35.76.0        192.58.107.0   192.58.109.0   192.58.127.0   192.58.181.0
192.35.78.0    192.35.79.0    192.35.82.0    192.35.86.0        192.58.194.0   192.58.199.0   192.58.204.0   192.58.206.0
192.35.96.0    192.35.97.0    192.35.98.0    192.35.140.0       192.58.218.0   192.58.222.0   192.58.223.0   192.58.224.0
192.35.148.0   192.35.154.0   192.35.162.0   192.35.163.0       192.58.225.0   192.58.244.0   192.65.78.0    192.65.131.0
192.35.169.0   192.35.170.0   192.35.171.0   192.35.180.0       192.65.137.0   192.65.143.0   192.65.147.0   192.65.175.0
192.35.196.0   192.35.200.0   192.35.208.0   192.35.213.0       192.65.183.0   192.67.6.0     192.67.53.0    192.67.62.0
192.35.226.0   192.35.229.0   192.36.125.0   192.36.148.0       192.67.80.0    192.67.92.0
192.39.2.0     192.39.11.0    192.39.12.0    192.39.13.0
192.39.14.0    192.40.51.0    192.41.140.0   192.41.146.0
192.41.177.0   192.41.197.0   192.41.200.0   192.41.202.0
192.41.204.0   192.41.211.0   192.41.217.0   192.41.228.0
192.41.237.0   192.41.245.0   192.41.246.0   192.41.249.0
192.42.2.0     192.42.4.0     192.42.8.0     192.42.41.0
192.42.61.0    192.42.62.0    192.42.70.0    192.42.80.0
192.42.81.0    192.42.82.0    192.42.98.0    192.42.108.0
192.42.110.0   192.42.114.0   192.42.141.0   192.42.142.0
192.42.144.0   192.42.152.0   192.42.155.0   192.42.239.0
192.42.244.0   192.42.245.0   192.42.246.0   192.42.247.0
192.42.248.0   192.43.152.0   192.43.188.0   192.43.197.0
192.43.199.0   192.43.204.0   192.43.205.0   192.43.207.0
192.43.239.0   192.43.244.0   192.43.250.0   192.43.252.0
192.44.1.0     192.44.82.0    192.44.83.0    192.44.84.0
192.44.85.0    192.44.216.0   192.44.217.0   192.44.220.0
192.44.221.0   192.44.222.0   192.44.223.0   192.44.224.0
192.44.225.0   192.44.226.0   192.44.237.0   192.44.253.0
192.48.33.0    192.48.80.0    192.48.95.0    192.48.100.0
```

DDN MAILBRIDGES

CURRENT STATUS

- Six DDN Butterfly Mailbridges operational
- 265 EGP neighbors
- Ethernet interfaces added to Mitre and Ames mailbridges
  - 192.52.194-NSFTRANSIT 5
  - 192.52.195-NSFTRANSIT 6
- BMILLBL has only one interface
  - ARPANET interface eliminated
  - Provides EGP server function on MILNET

TRAFFIC SUMMARY

- ~ 13,000,000 packets/day forwarded
- 0.6-1.9% packets dropped
- Average Bytes per packet

BMILAMES:	159	BMILISI:	221
BMILLBL:	359	BMILBBN:	149
BMILMTR:	169	BMILDCEC:	153

BMILDCEC DAILY THROUGHPUT

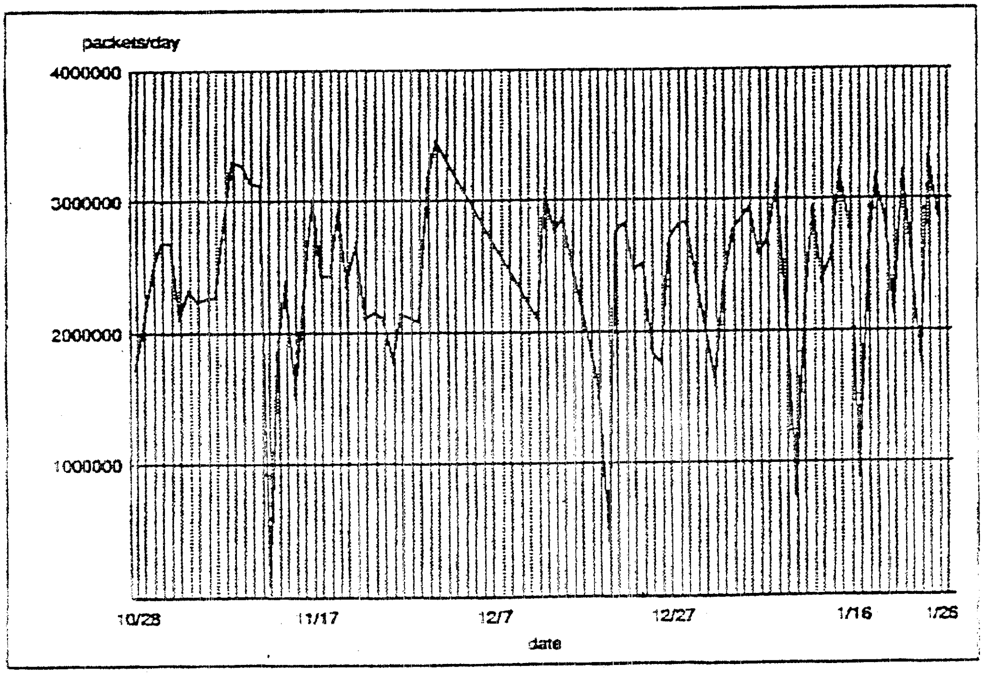

BMILAMES DAILY THROUGHPUT

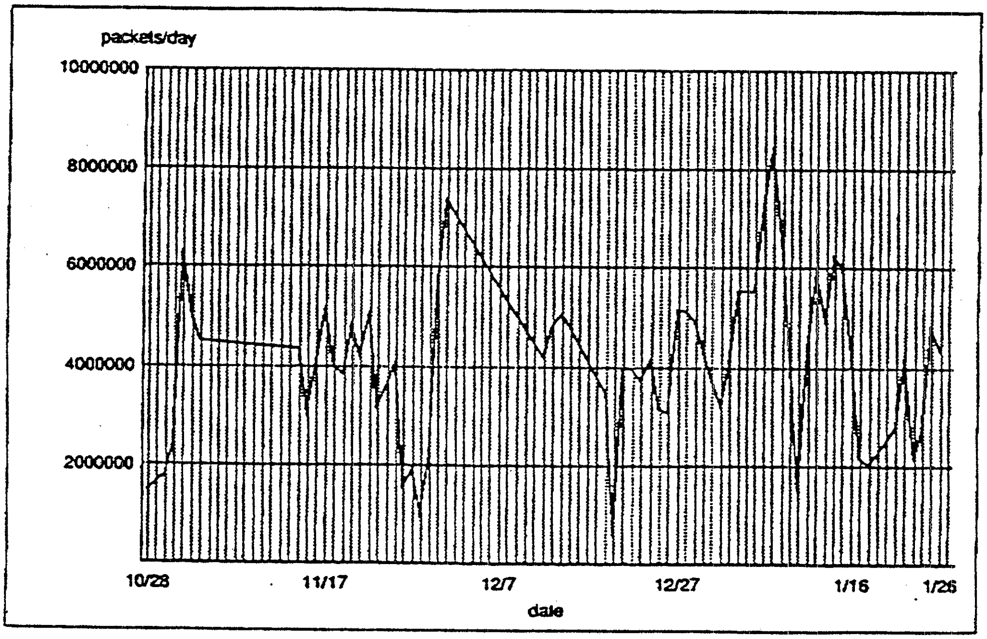

BMILMTR DAILY THROUGHPUT

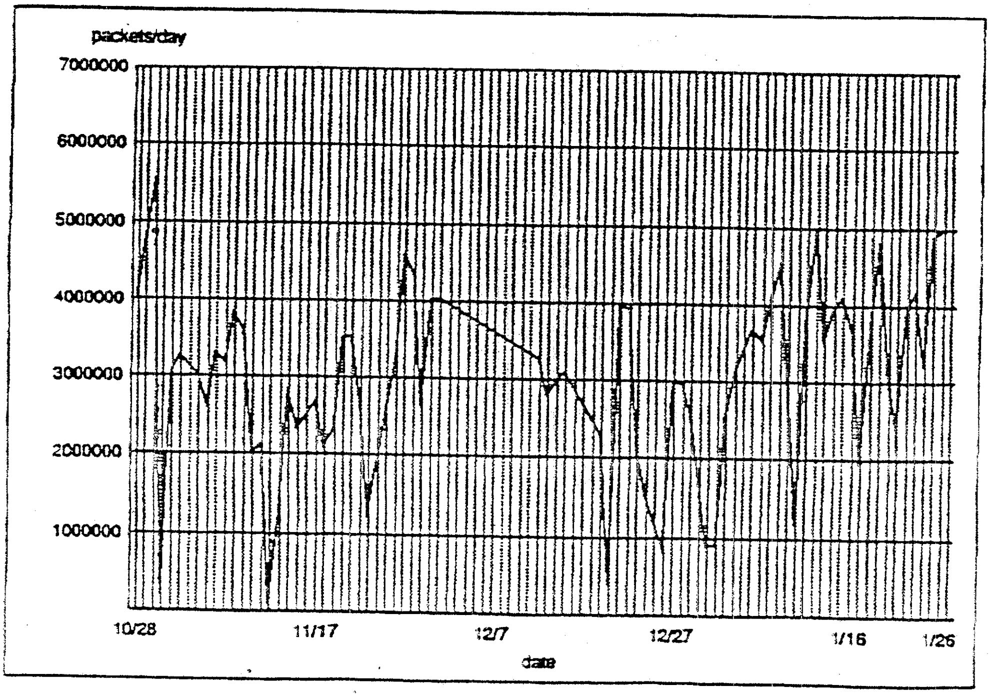

BMILISI DAILY THROUGHPUT

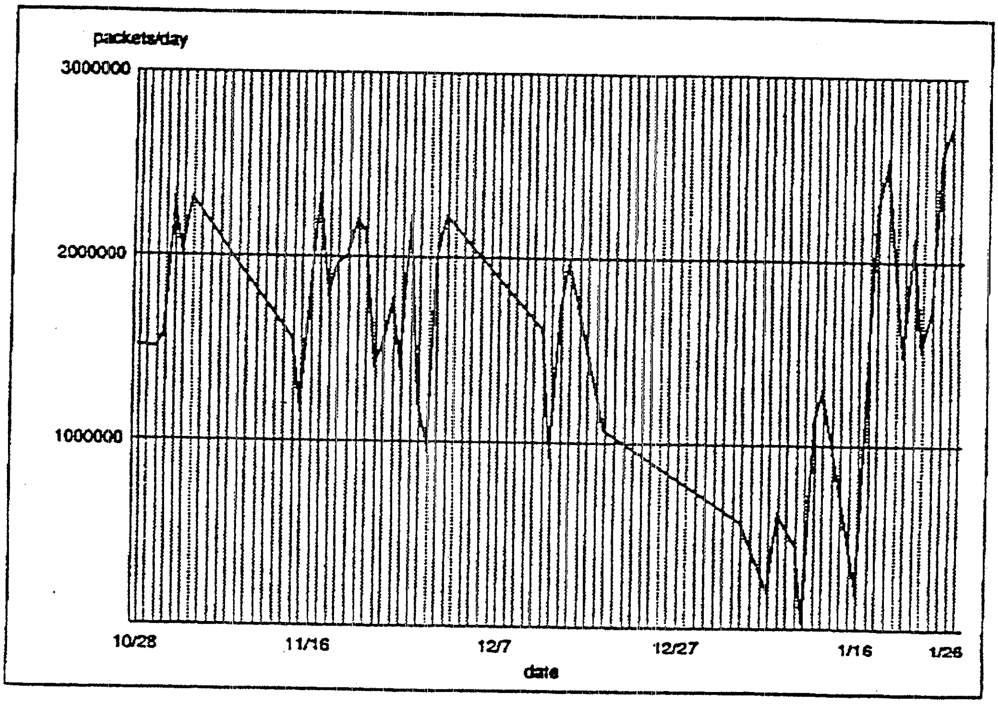

BMILBBN DAILY THROUGHPUT

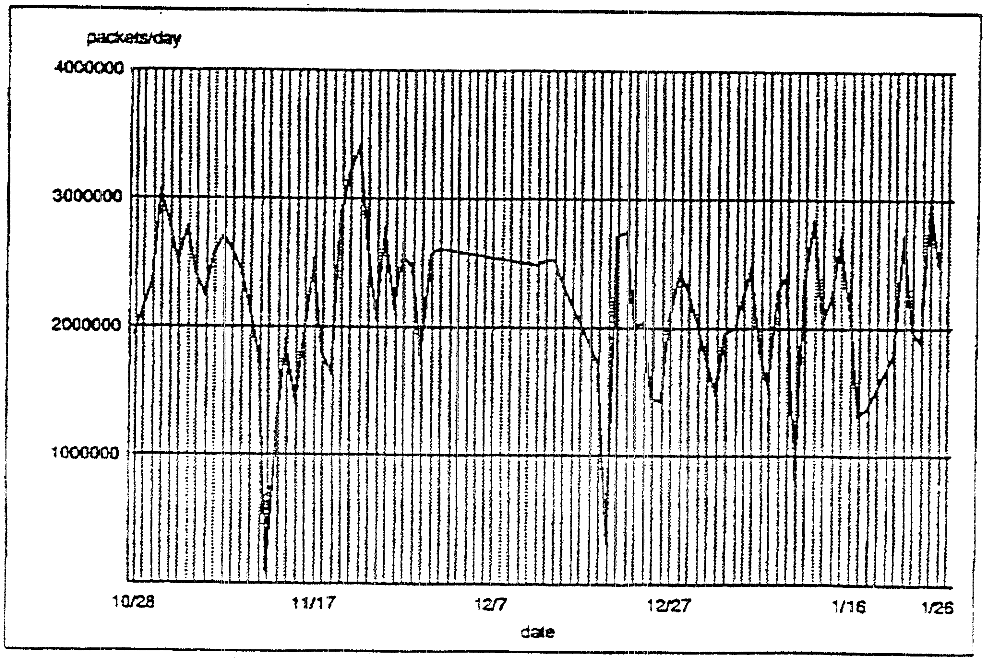

BMILLBL DAILY THROUGHPUT

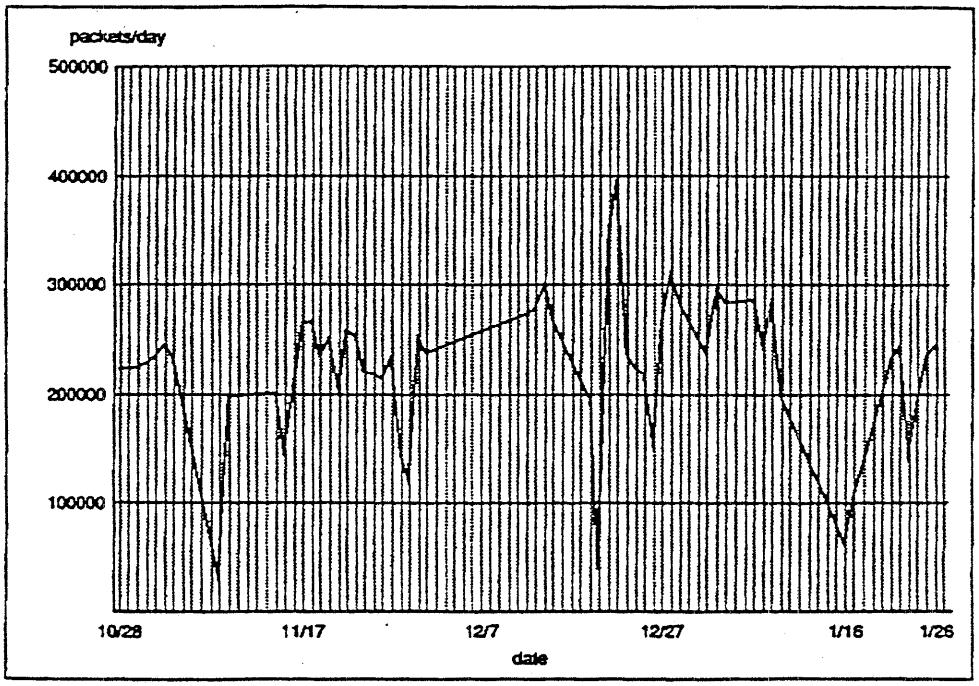

4.2 "ESnet Report"

Presentation by Tony Hain/ DOE

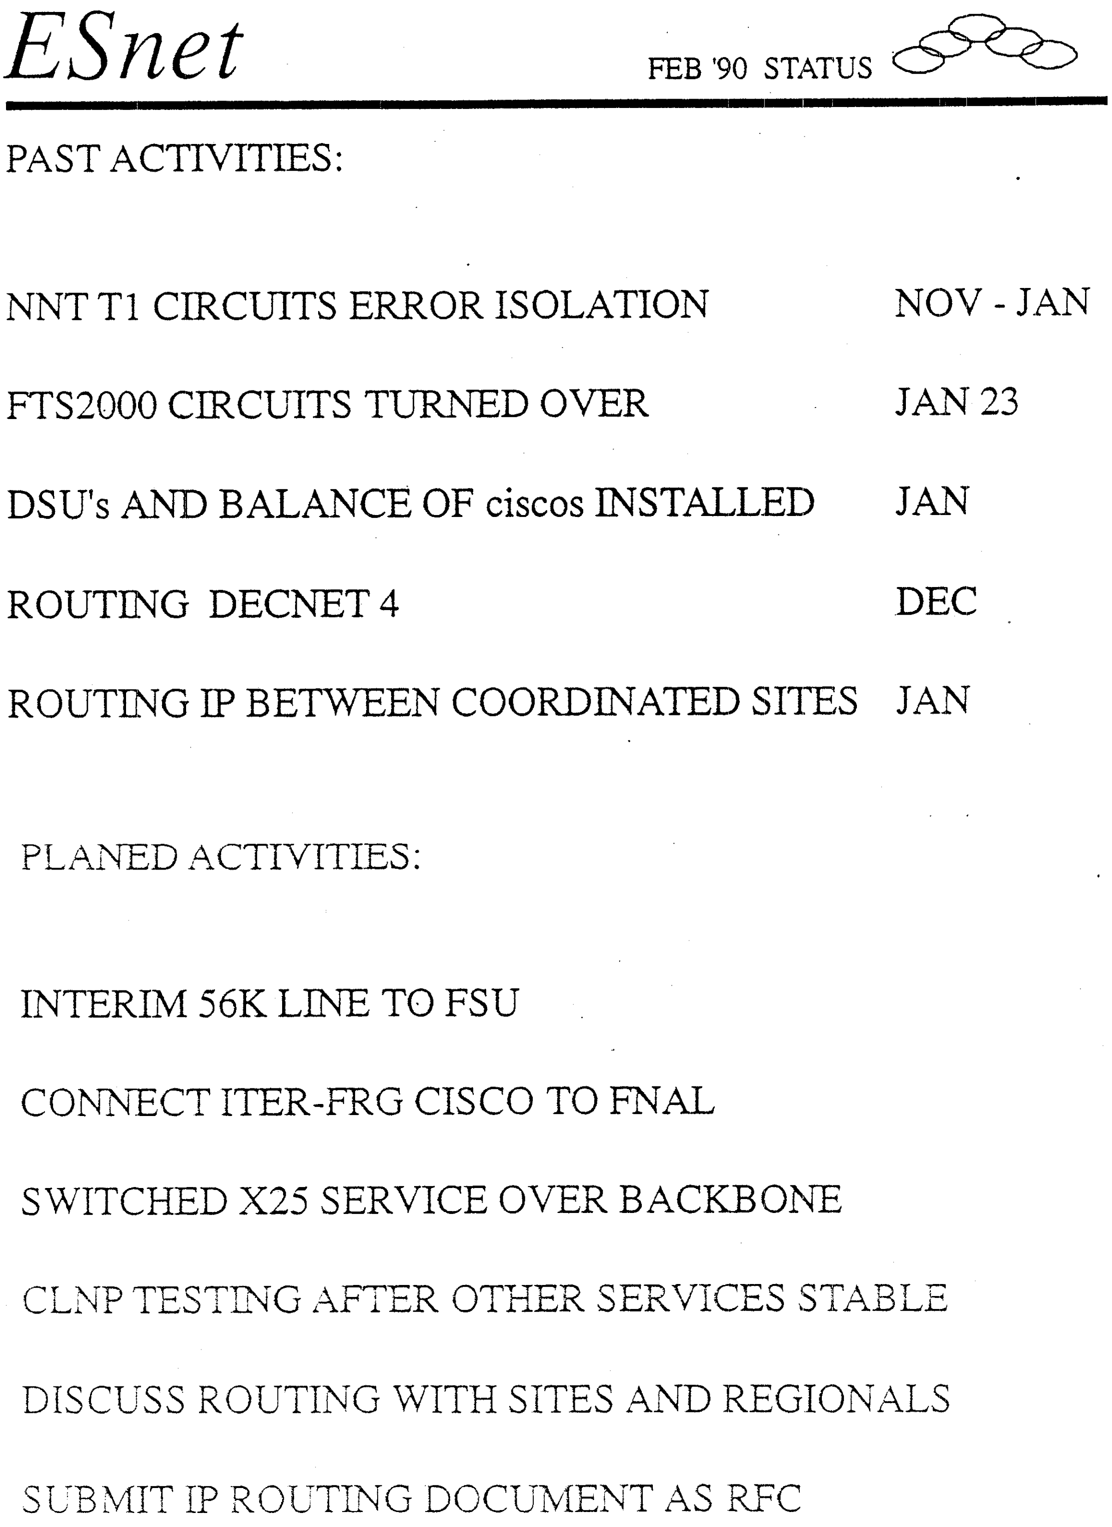

ESnet FEB '90 STATUS

PAST ACTIVITIES:

NNT T1 CIRCUITS ERROR ISOLATION	NOV - JAN
FTS2000 CIRCUITS TURNED OVER	JAN 23
DSU's AND BALANCE OF ciscos INSTALLED	JAN
ROUTING DECNET 4	DEC
ROUTING IP BETWEEN COORDINATED SITES	JAN

PLANED ACTIVITIES:

INTERIM 56K LINE TO FSU

CONNECT ITER-FRG CISCO TO FNAL

SWITCHED X25 SERVICE OVER BACKBONE

CLNP TESTING AFTER OTHER SERVICES STABLE

DISCUSS ROUTING WITH SITES AND REGIONALS

SUBMIT IP ROUTING DOCUMENT AS RFC

ESNET BACKBONE - 1990

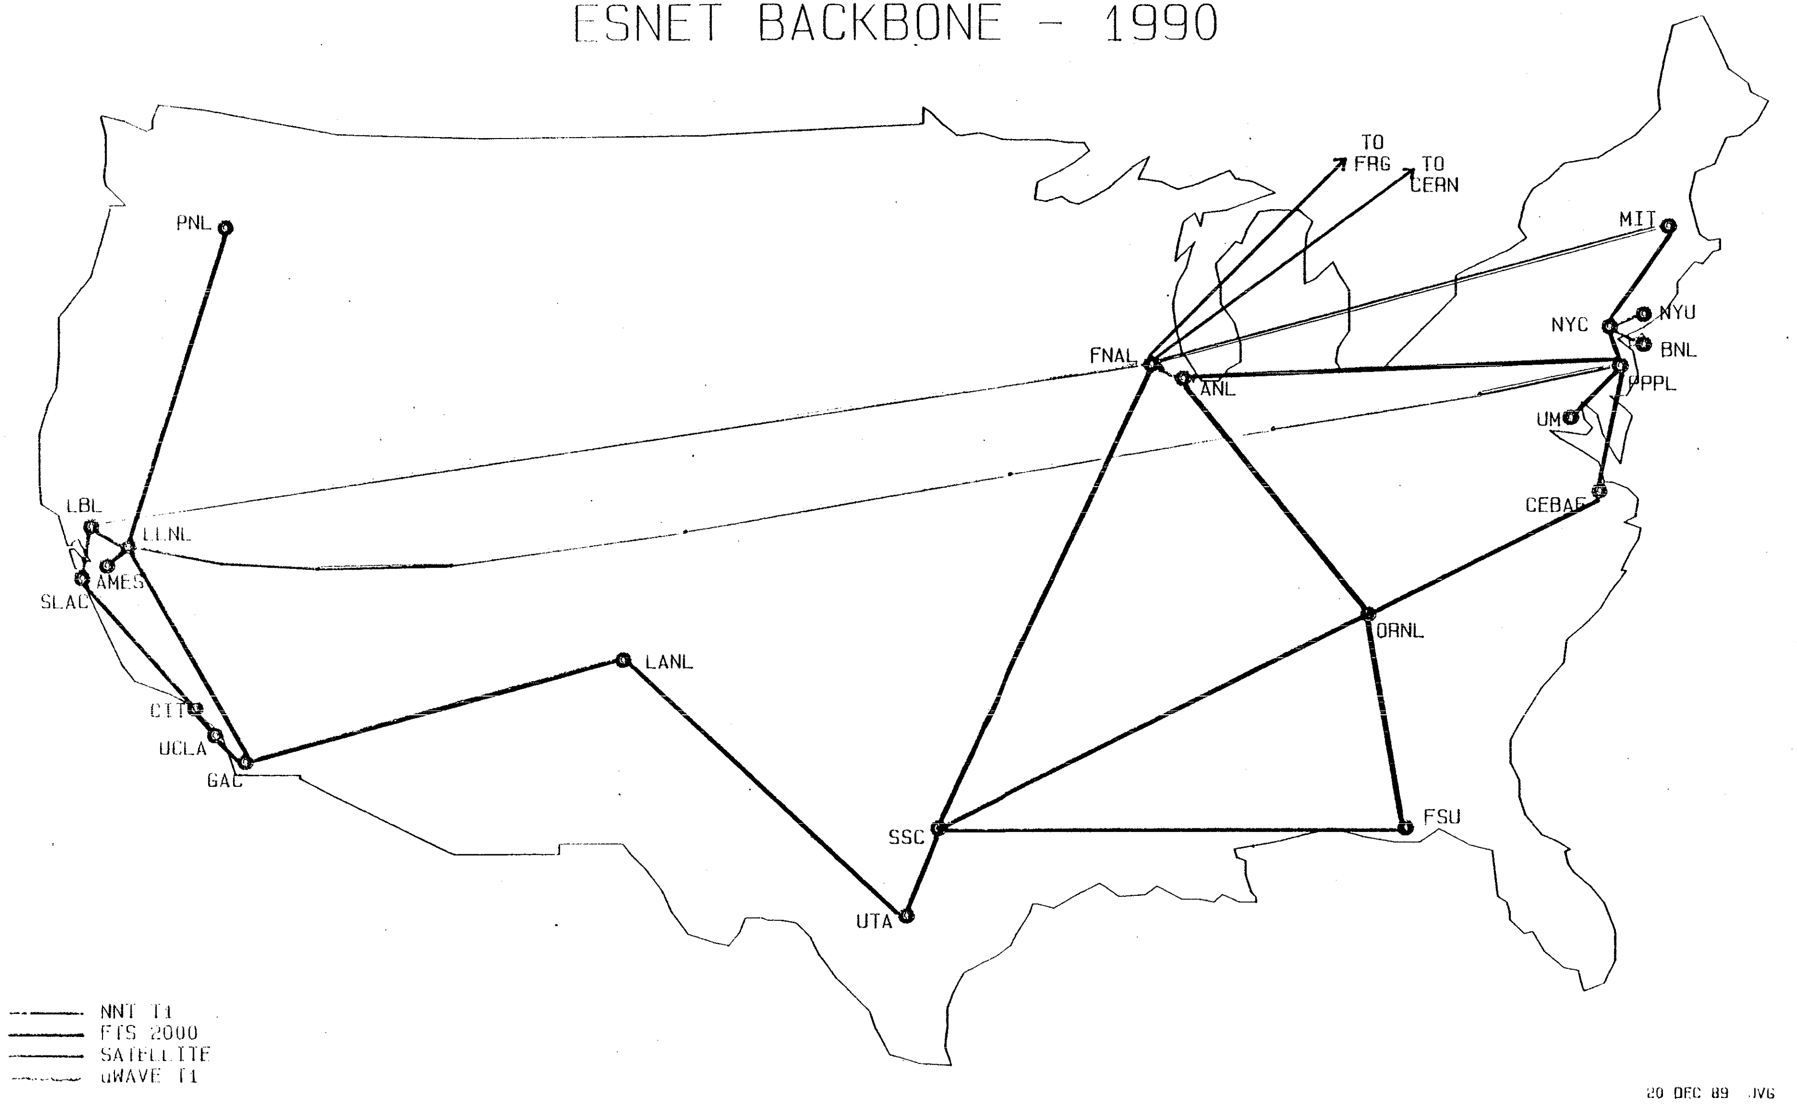

4.3 "NSF Report"

Presentation by Elise Gerich/ MERIT

NSFNET T-1 Data Network
New Topology

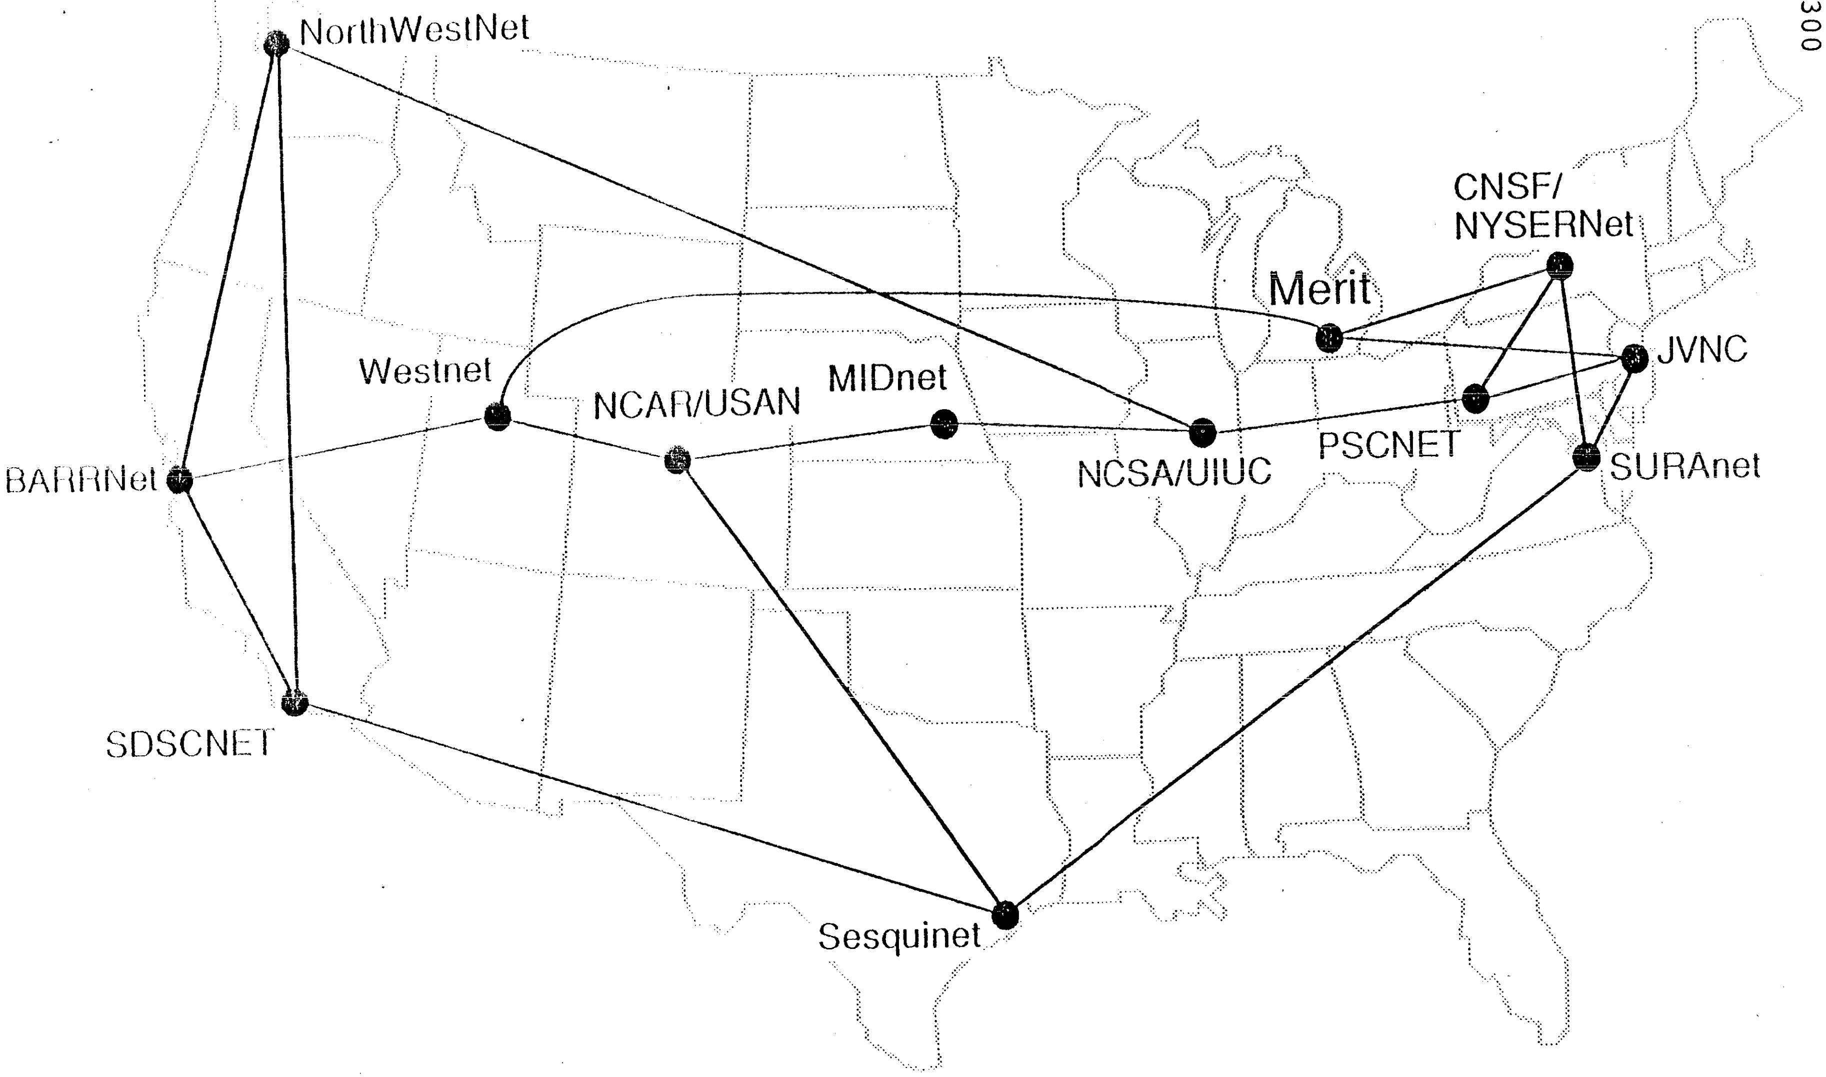

The Merit Computer Network

LK

2 8 1989

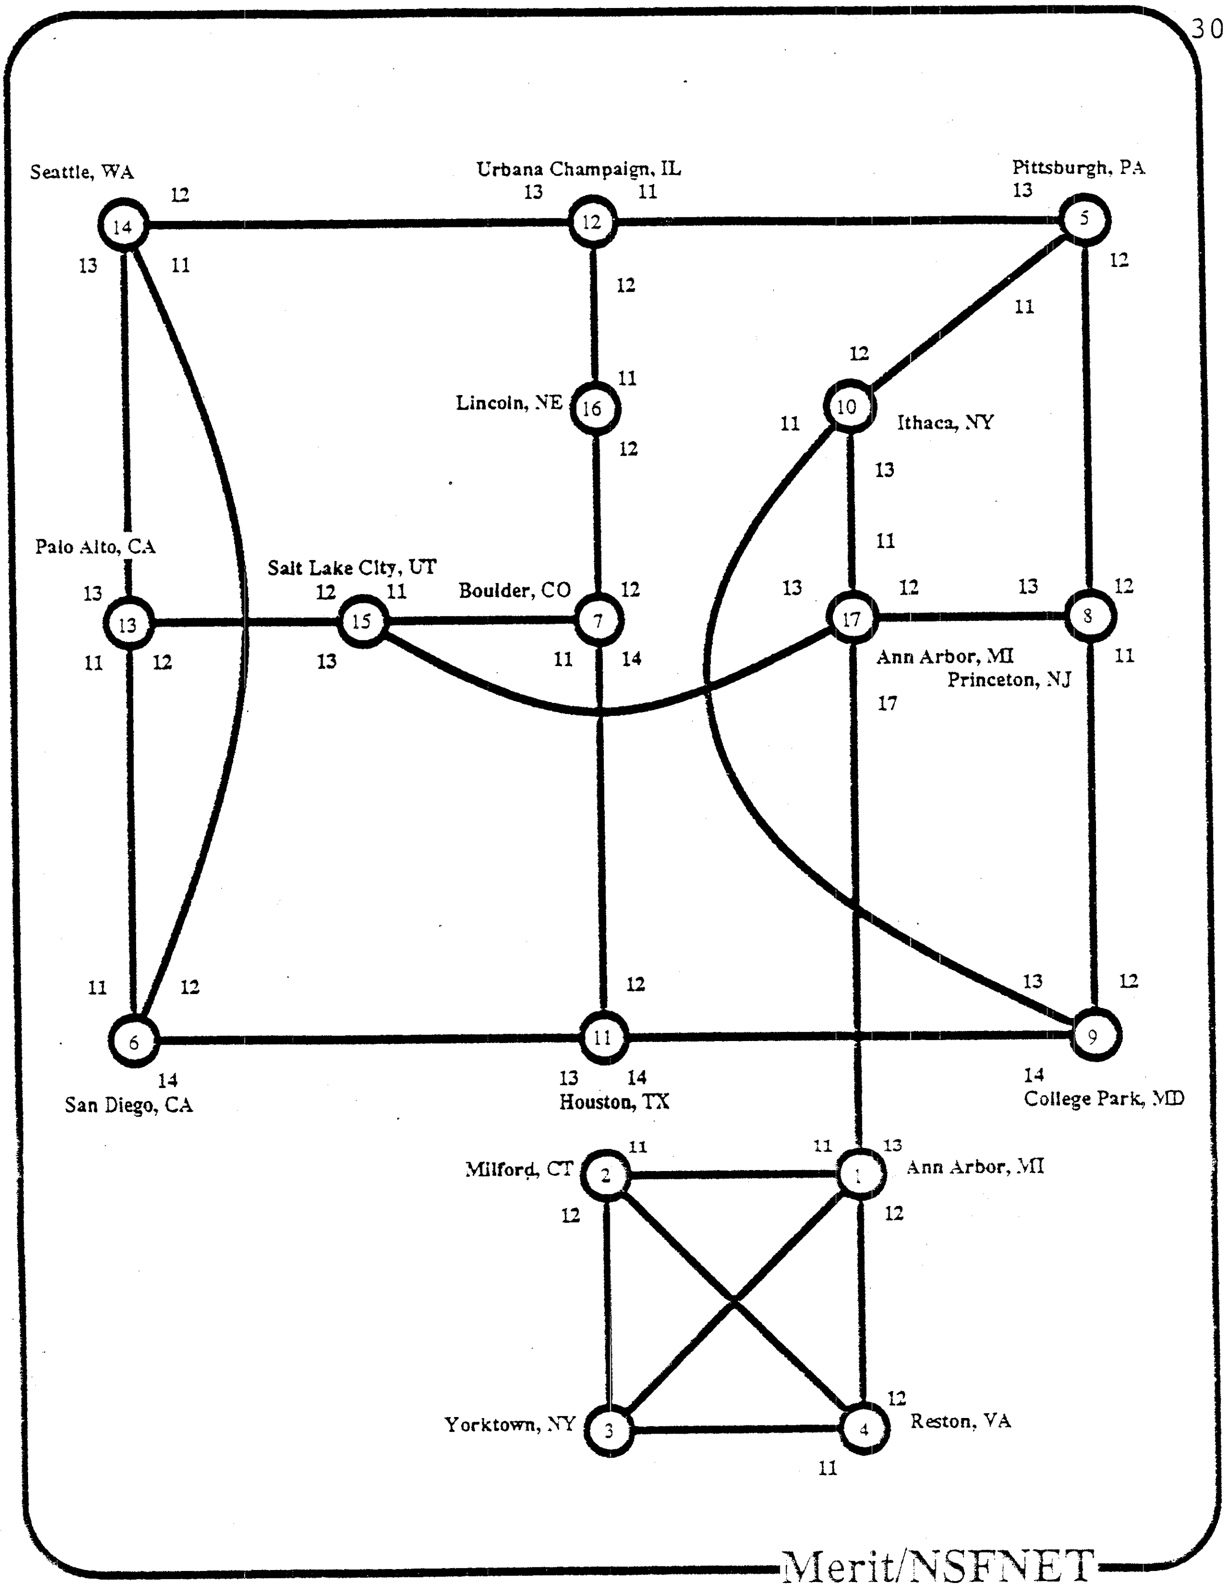

Merit/NSFNET

18 August 1989, HWB/BAC

COMPARISON OF NSFNET TRAFFIC 1988 and 1989

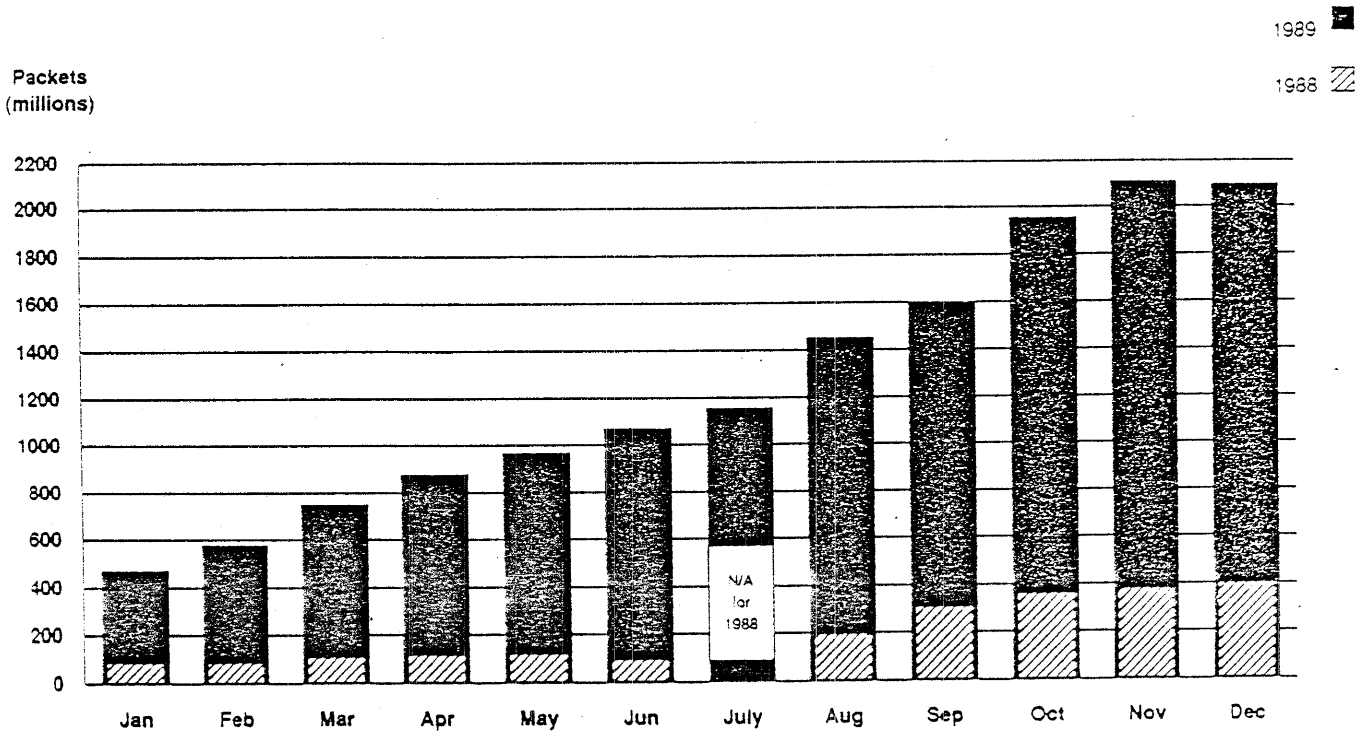

February 2, 1990

TOTAL NUMBER OF PACKETS BY NSS
December 1989

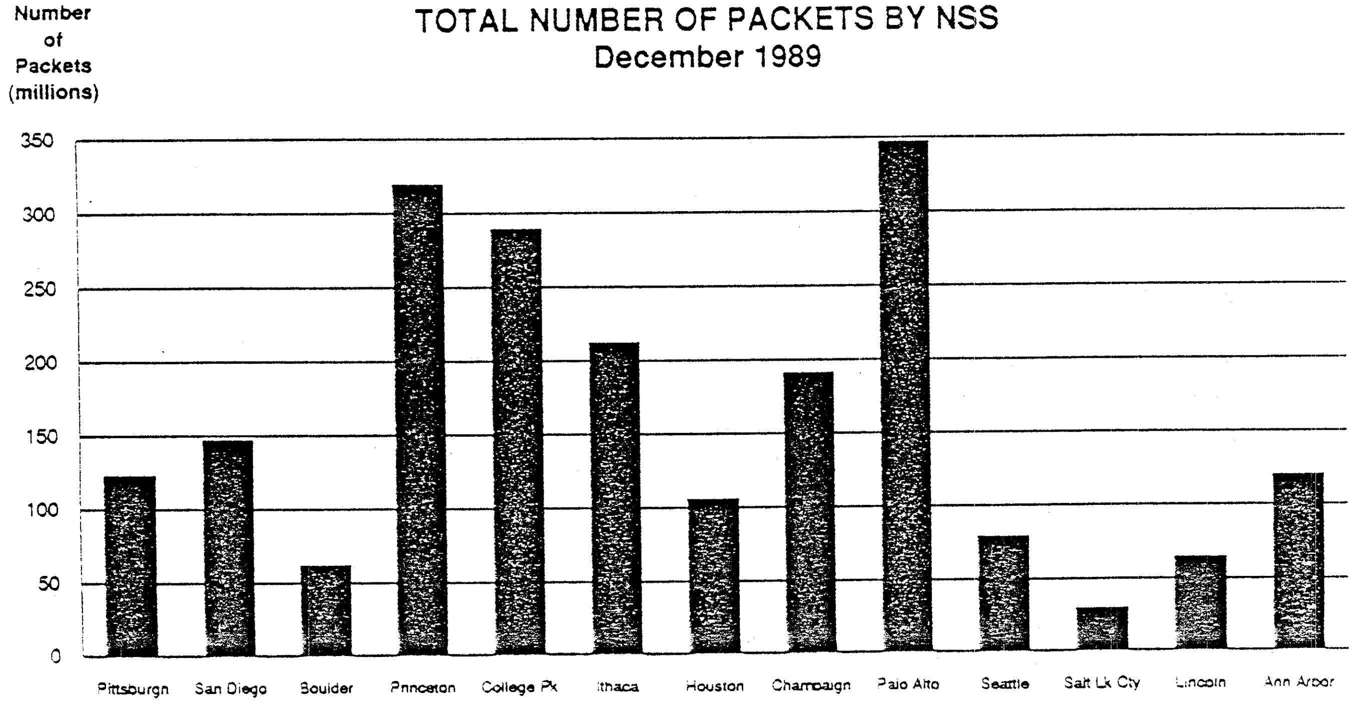

Packets into the NSFNET from NSS 5

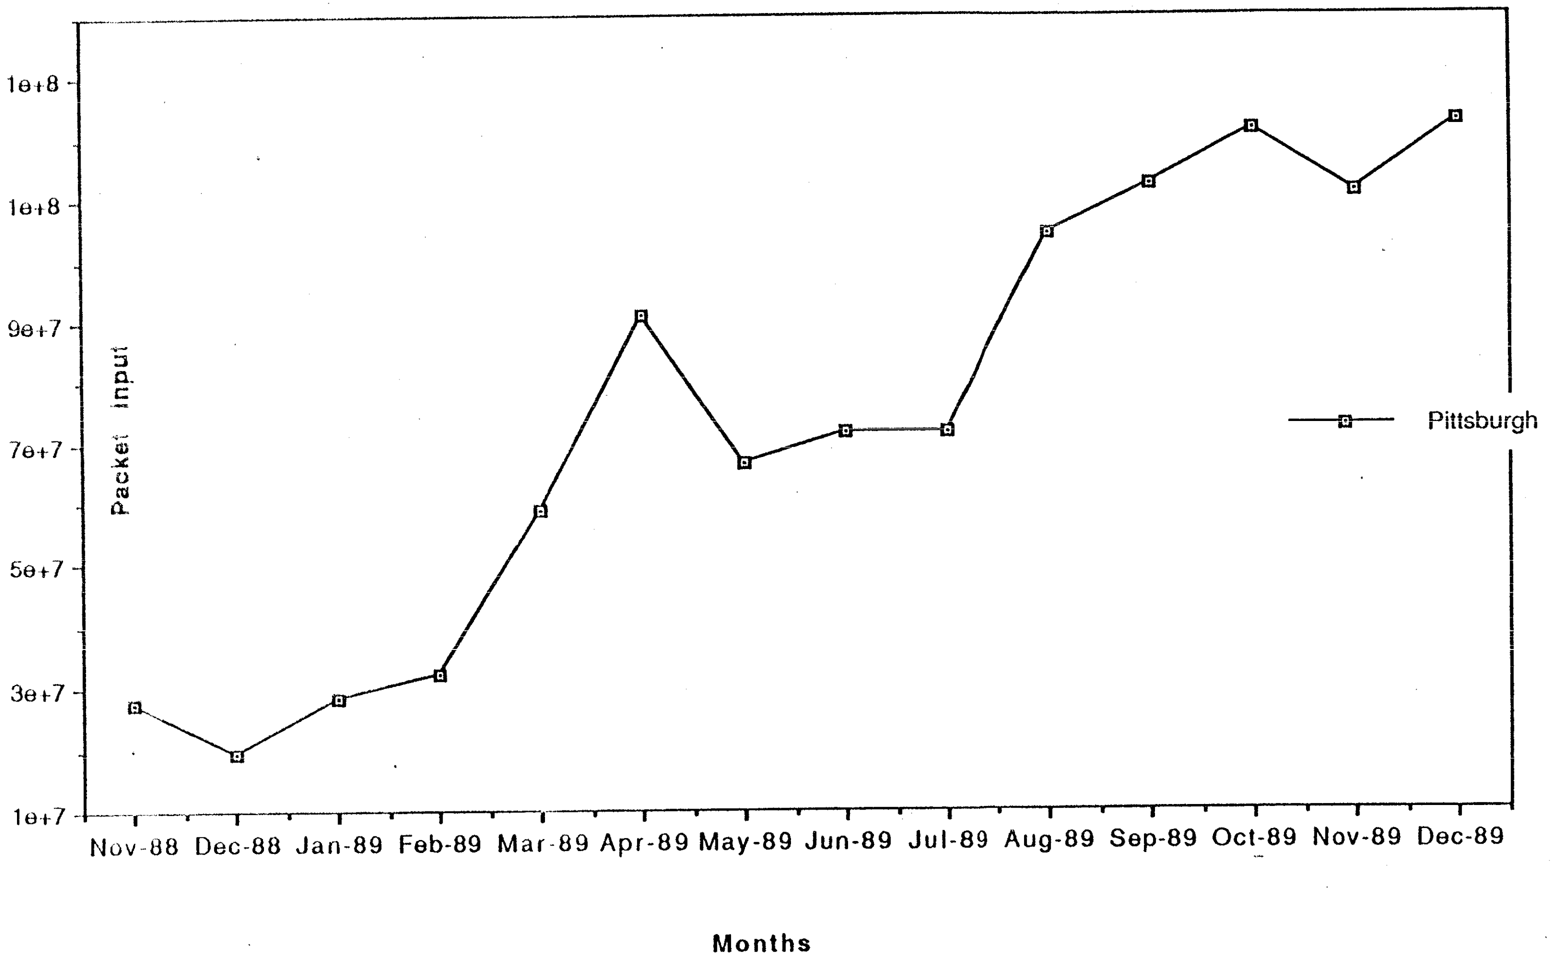

Number of Packets offered by 6 Nodes

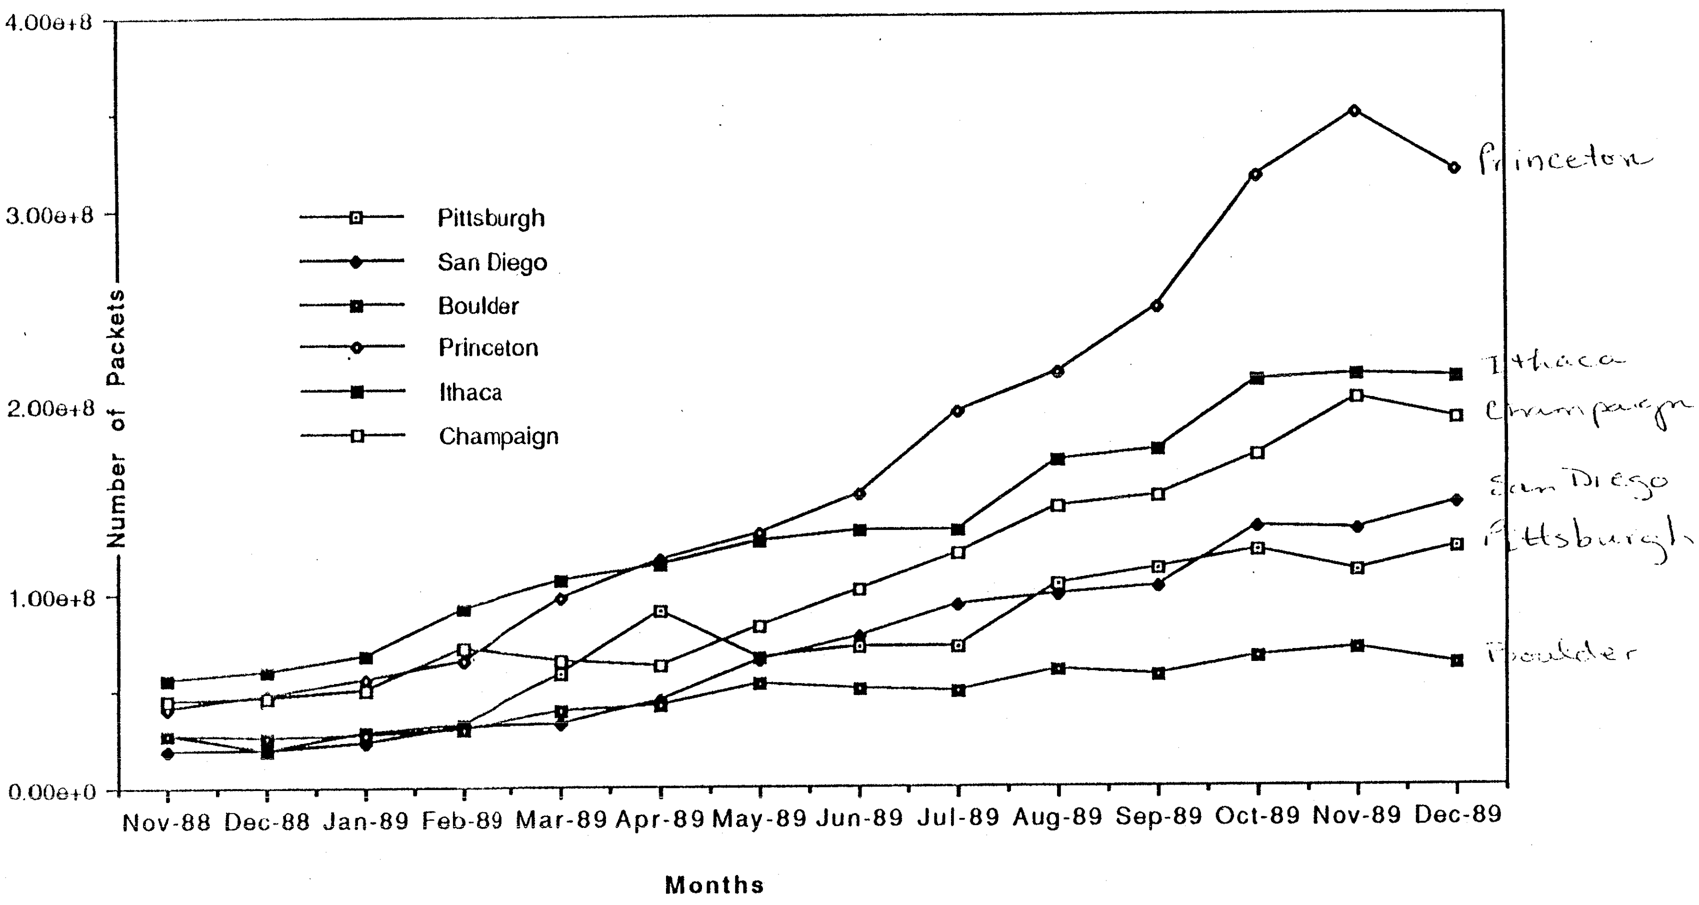

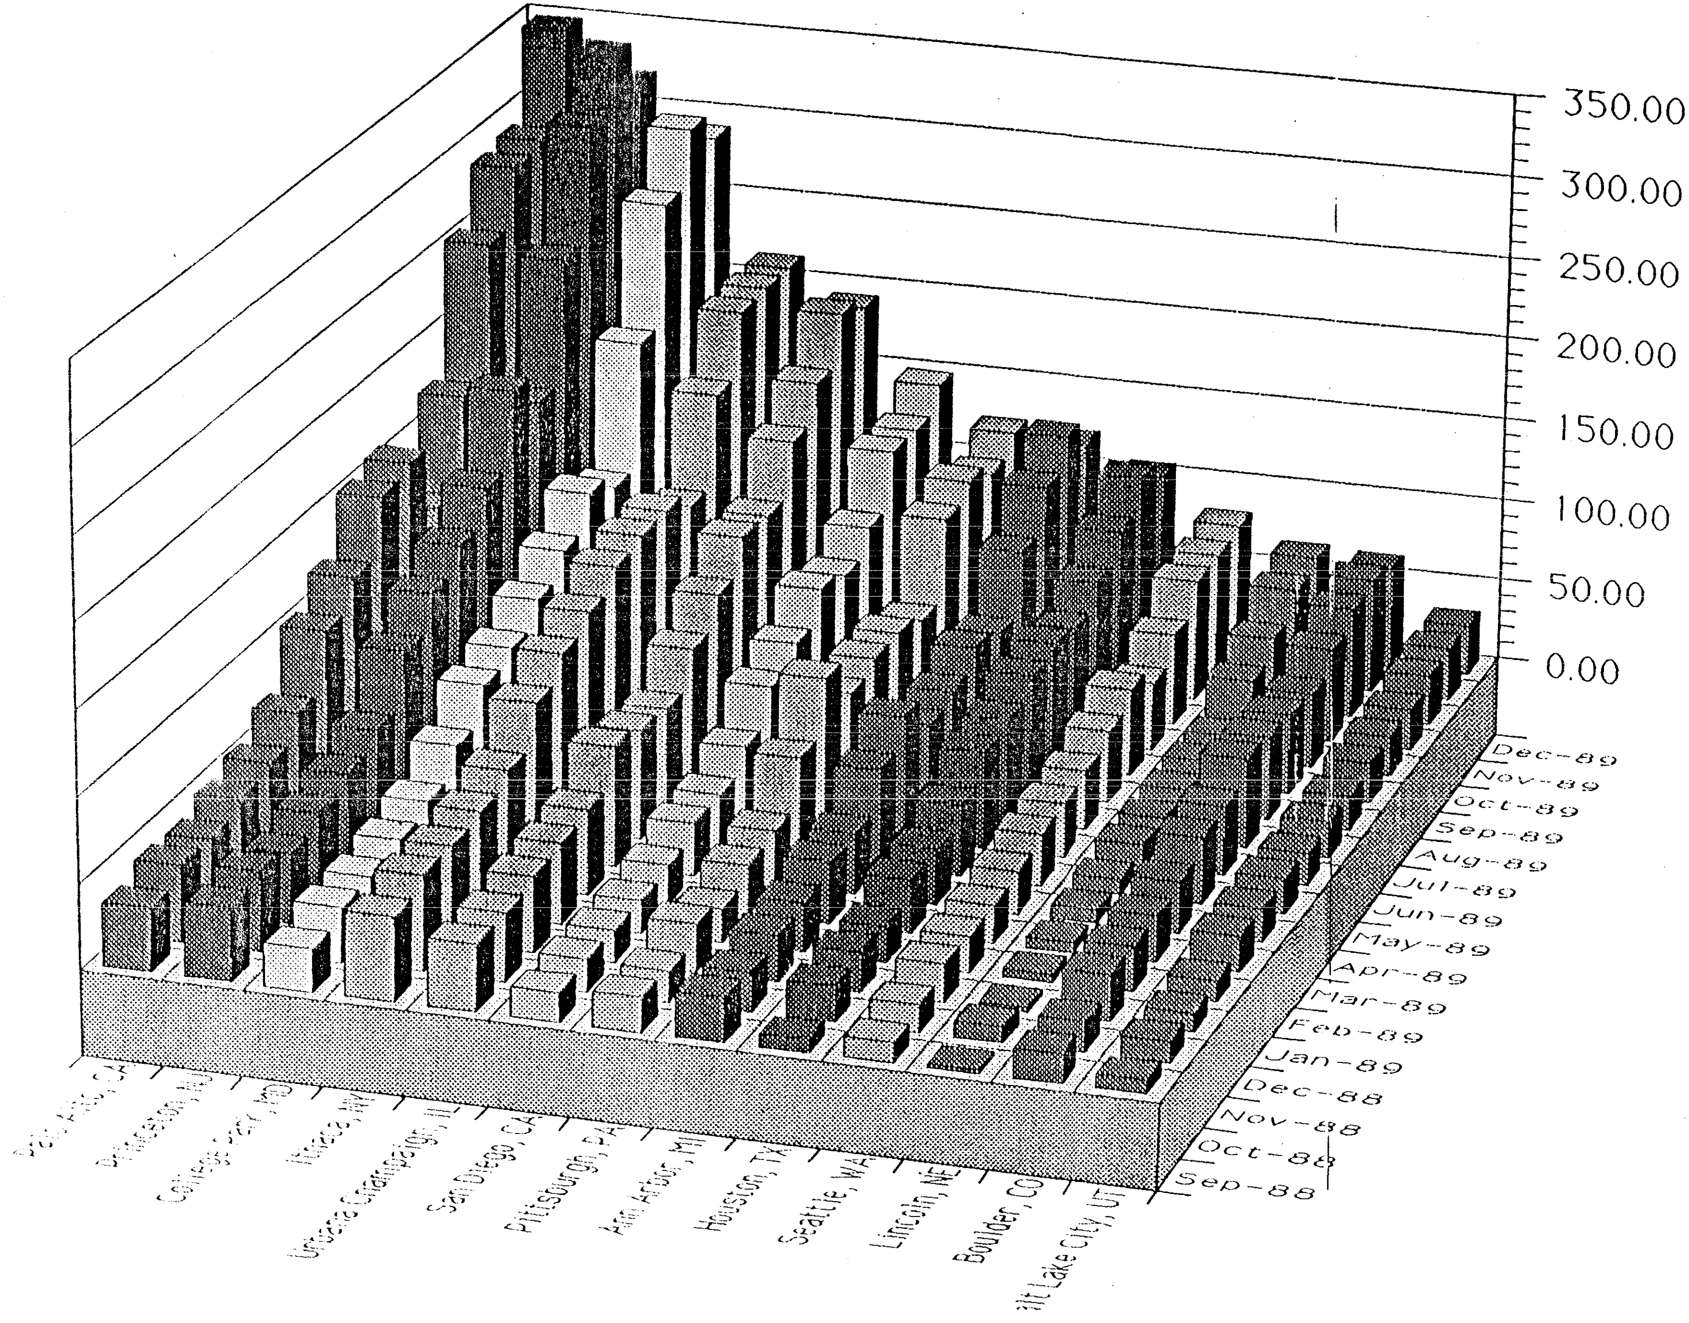
350.00
300.00
250.00
200.00
150.00
100.00
50.00
0.00
Dec-89
Nov-89
Oct-89
Sep-89
Aug-89
Jul-89
Jun-89
May-89
Apr-89
Mar-89
Feb-89
Jan-89
Dec-88
Nov-88
Oct-88
Sep-88
Palo Alto, CA
Princeton, NJ
College Park, MD
Ithaca, NY
Urbana Champaign, IL
San Diego, CA
Pittsburgh, PA
Ann Arbor, MI
Houston, TX
Seattle, WA
Lincoln, NE
Boulder, CO
ilt Lake City, UT

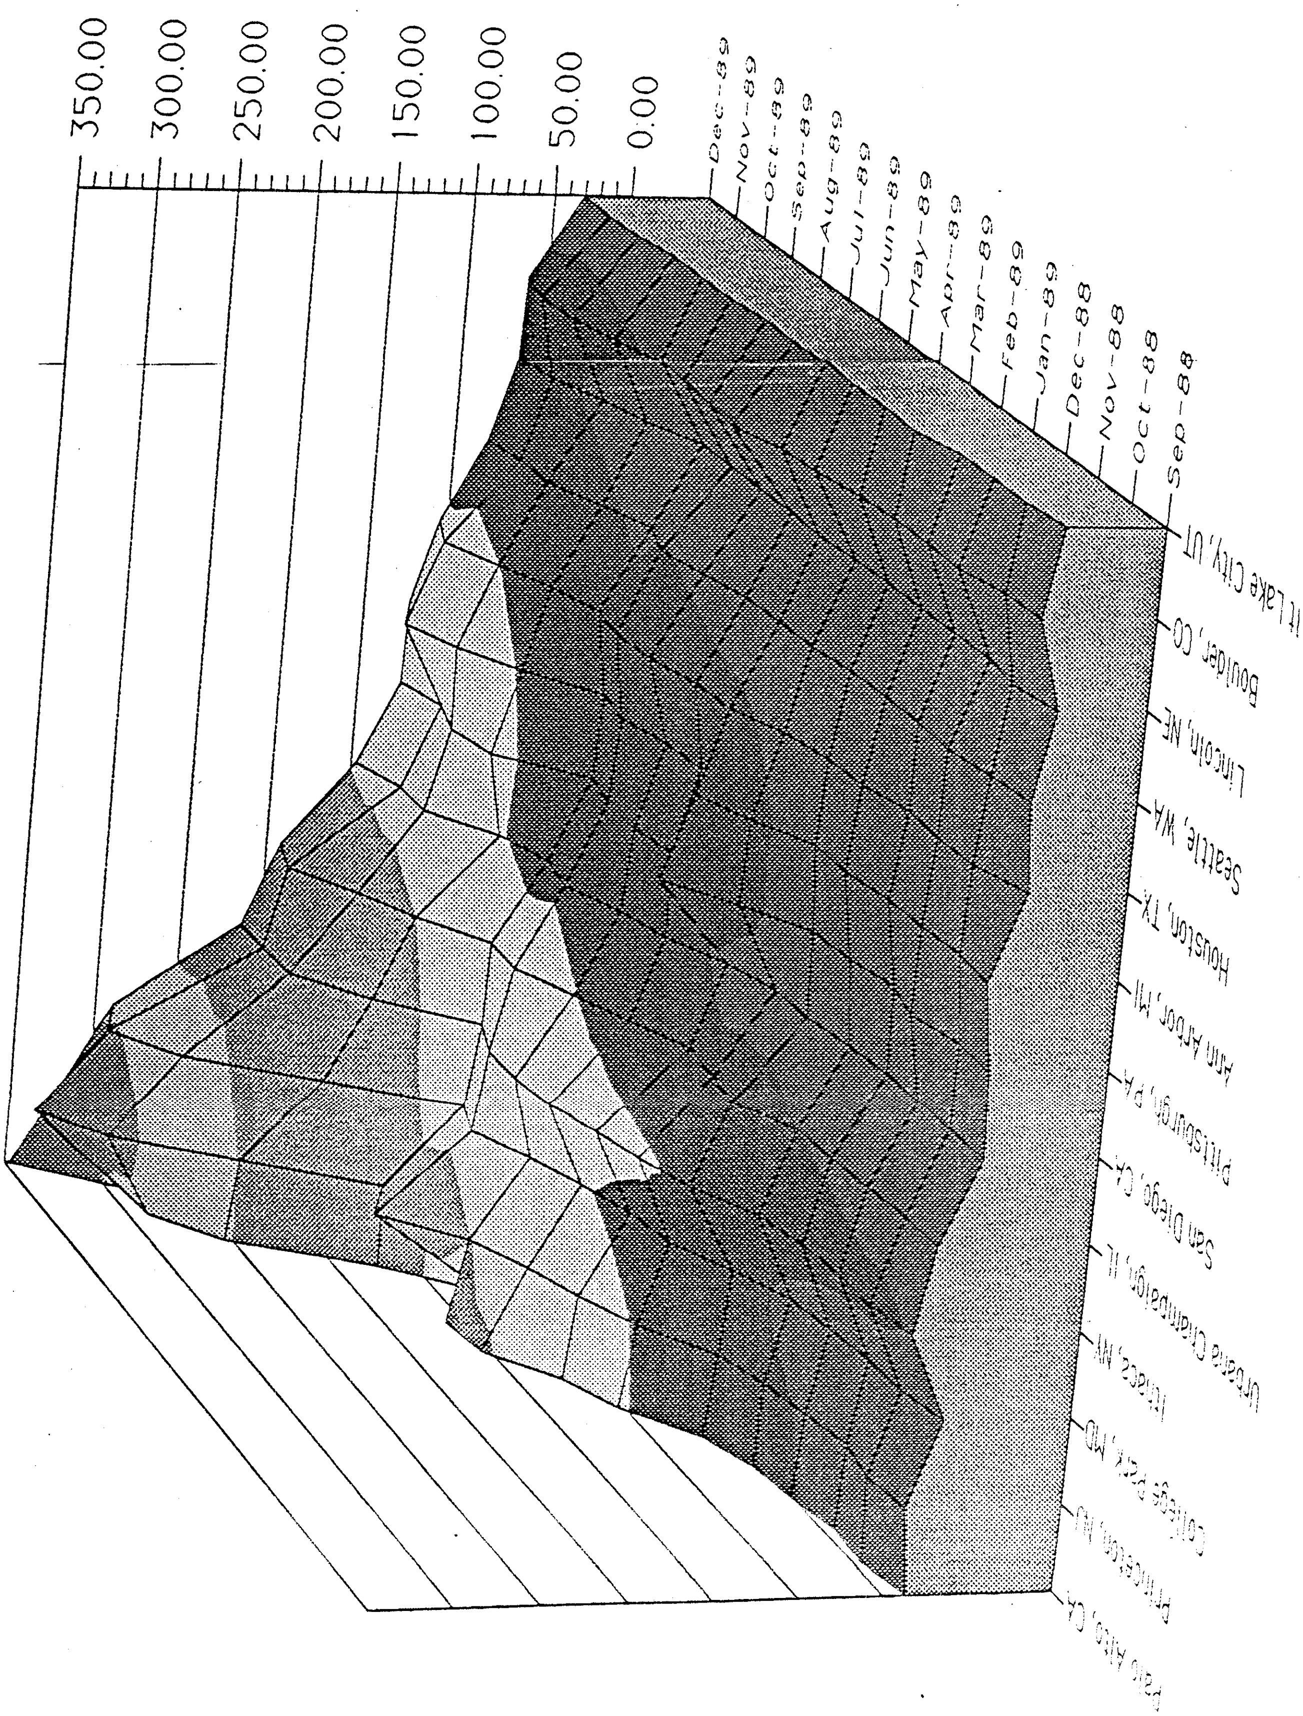
350.00
300.00
250.00
200.00
150.00
100.00
50.00
0.00
Dec-89
Nov-89
Oct-89
Sep-89
Aug-89
Jul-89
Jun-89
May-89
Apr-89
Mar-89
Feb-89
Jan-89
Dec-88
Nov-88
Oct-88
Sep-88
lt Lake City, UT
Boulder, CO
Lincoln, NE
Seattle, WA
Houston, TX
Ann Arbor, MI
Pittsburgh, PA
San Diego, CA
Urbana Champaign, IL
Ithaca, NY
College Park, MD
Princeton, NJ
Palo Alto, CA

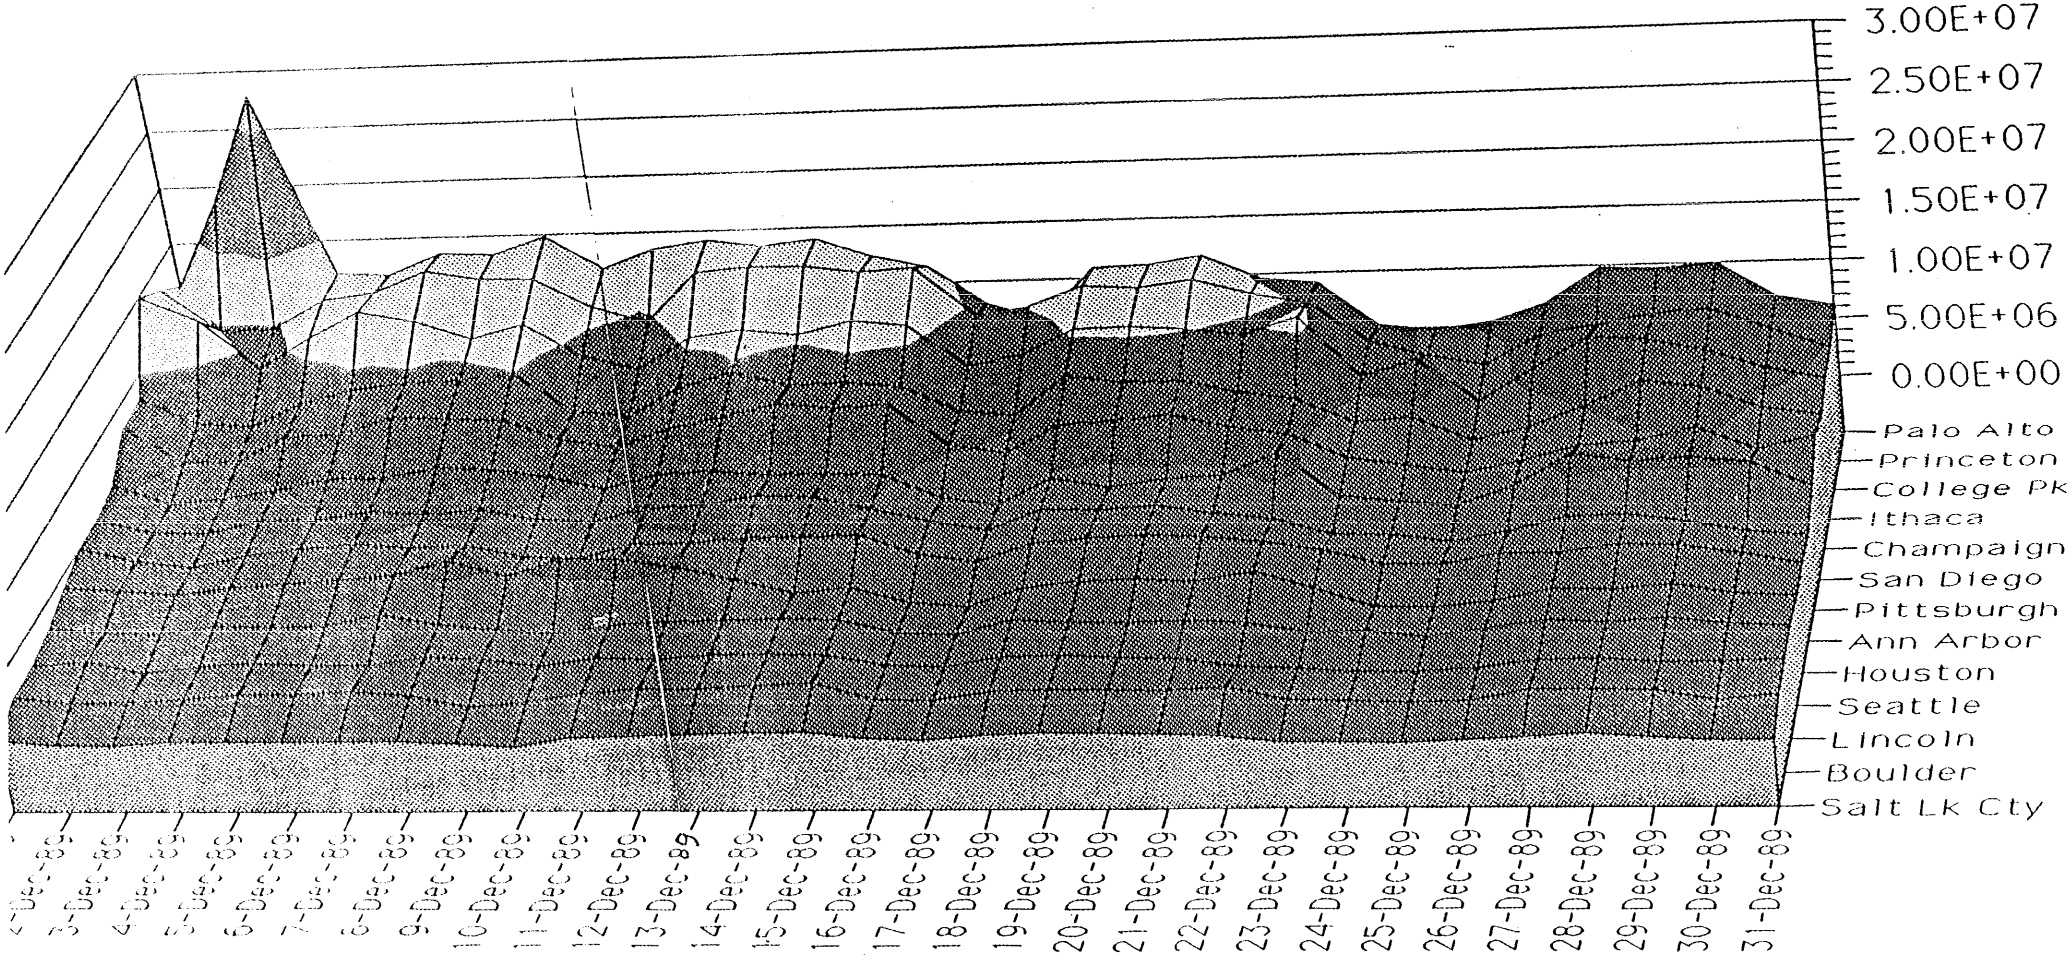
3.00E+07
2.50E+07
2.00E+07
1.50E+07
1.00E+07
5.00E+06
0.00E+00
Palo Alto
Princeton
College Pk
Ithaca
Champaign
San Diego
Pittsburgh
Ann Arbor
Houston
Seattle
Lincoln
Boulder
Salt Lk Cty
2-Dec-89
3-Dec-89
4-Dec-89
5-Dec-89
6-Dec-89
7-Dec-89
8-Dec-89
9-Dec-89
10-Dec-89
11-Dec-89
12-Dec-89
13-Dec-89
14-Dec-89
15-Dec-89
16-Dec-89
17-Dec-89
18-Dec-89
19-Dec-89
20-Dec-89
21-Dec-89
22-Dec-89
23-Dec-89
24-Dec-89
25-Dec-89
26-Dec-89
27-Dec-89
28-Dec-89
29-Dec-89
30-Dec-89
31-Dec-89

Offered Load Traffic Distribution

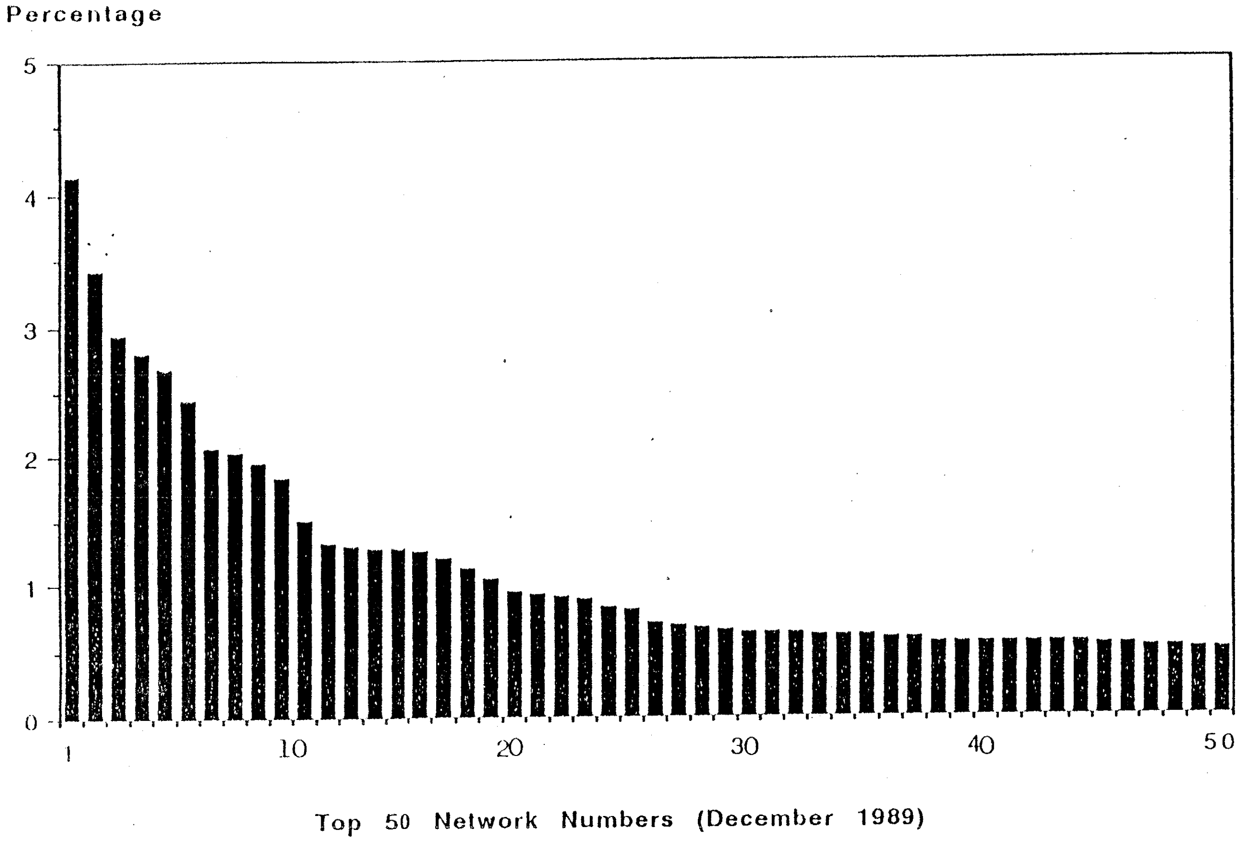

Peak of the Averages

Link Utilization

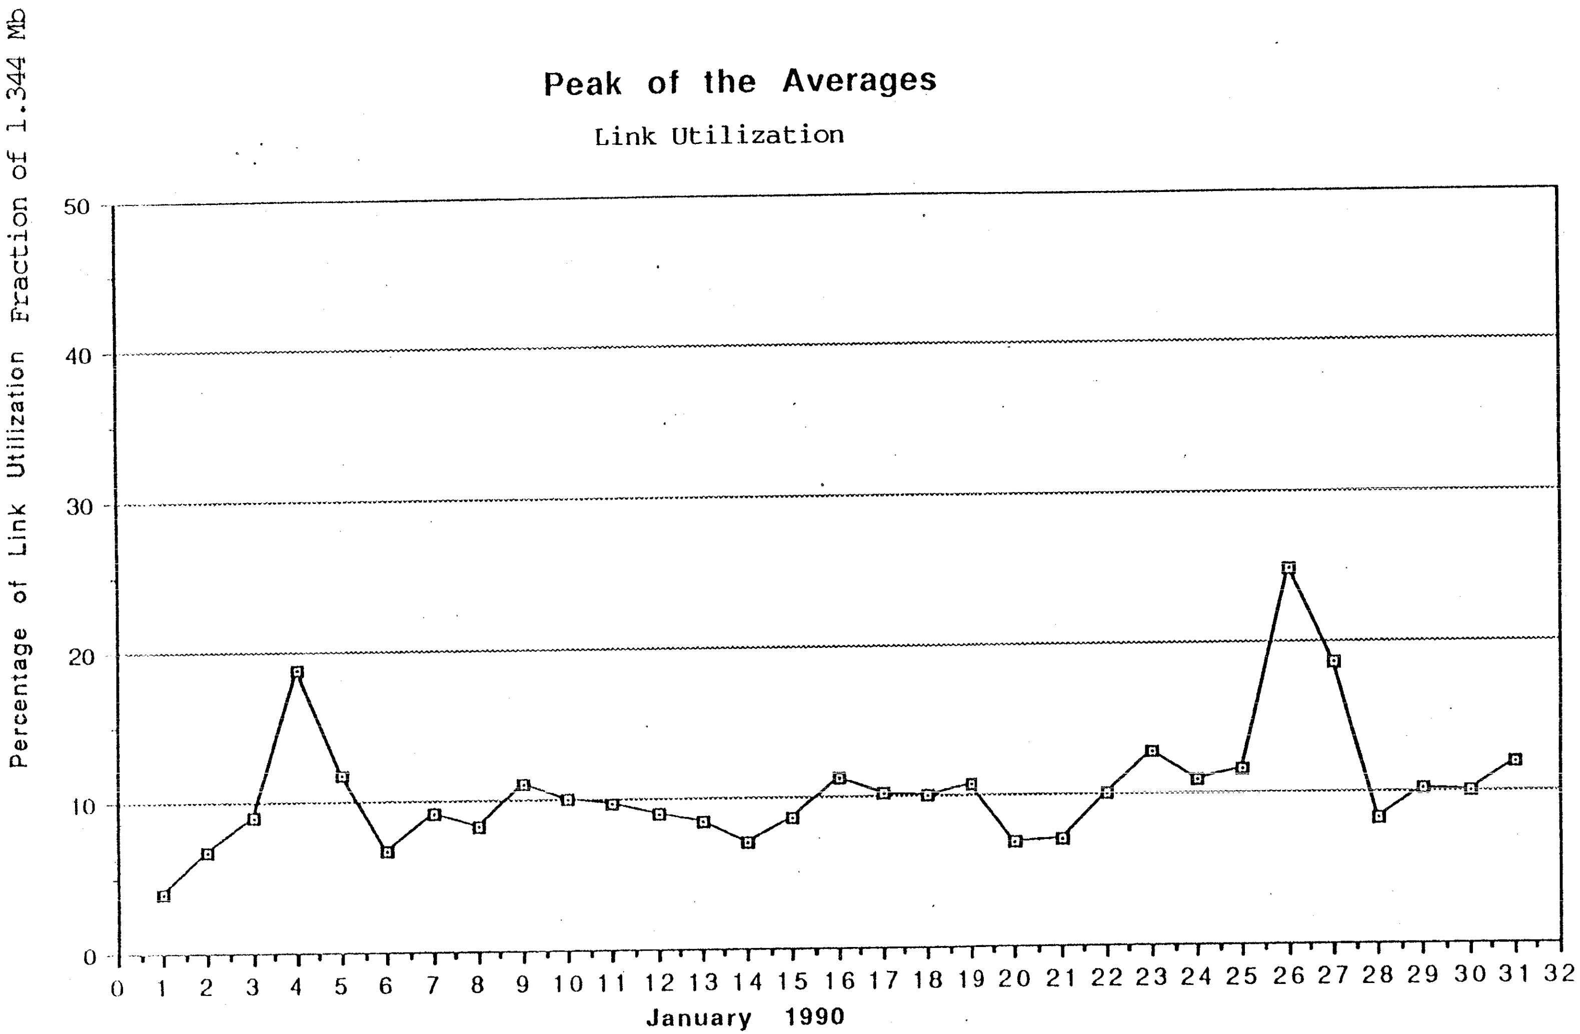

1	12	-	5
2	5	-	8
3	5	-	8
4	8	-	5
5	9	-	11
6	13	-	15
7	8	-	5
8	12	-	5
9	8	-	5
10	5	-	8
11	9	-	11
12	12	-	5
13	5	-	8
14	12	-	5
15	5	-	8
16	5	-	8
17	5	-	8
18	5	-	8
19	8	-	5
20	8	-	5
21	5	-	8
22	12	-	5
23	12	-	5
24	8	-	5
25	12	-	5
26	12	-	5
27	12	-	5
28	8	-	5
29	11	-	9
30	8	-	5
31	5	-	8

February 2, 1990

Peak of the Peaks

Link Utilization

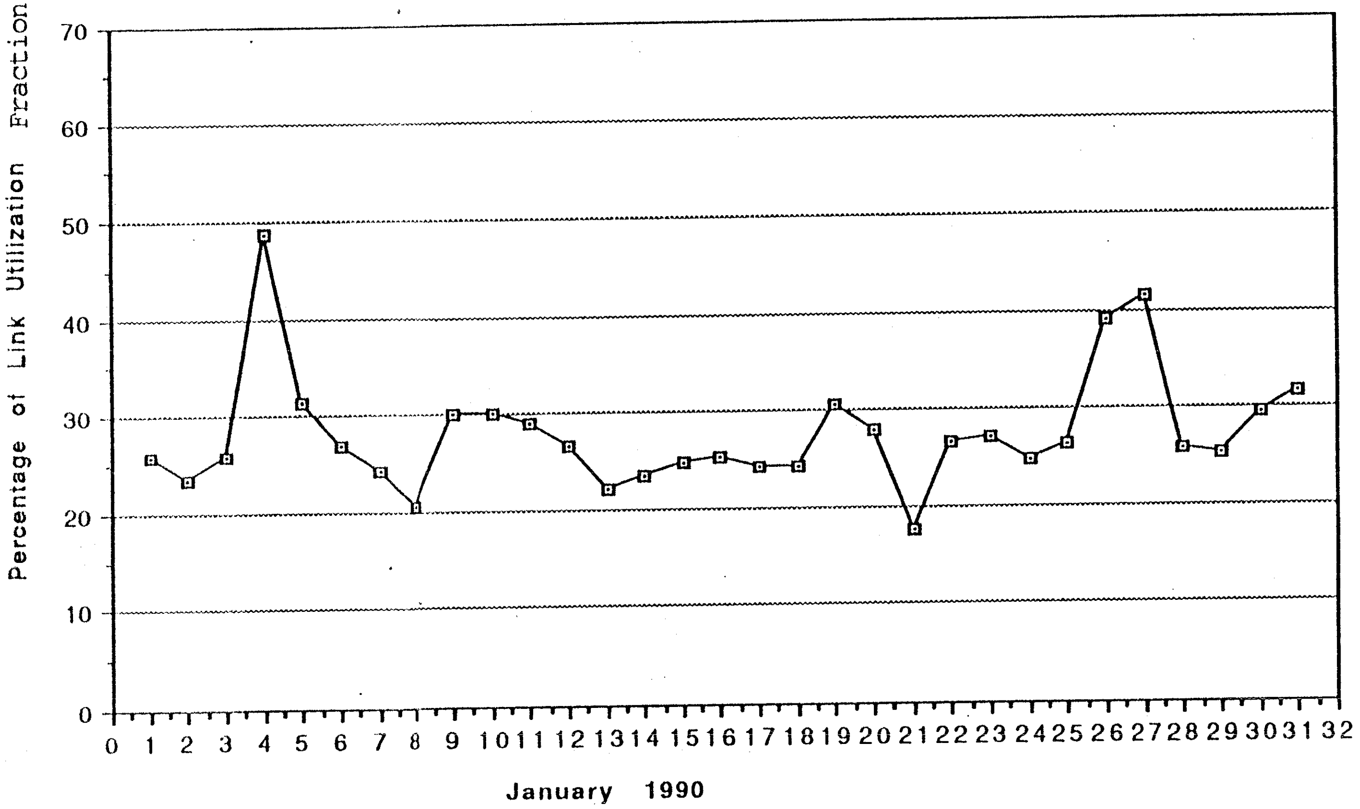

1	6	-	11
2	6	-	11
3	8	-	5
4	9	-	11
5	9	-	8
6	17	-	15
7	8	-	5
8	12	-	16
9	12	-	5
10	12	-	5
11	5	-	10
12	16	-	12
13	17	-	10
14	5	-	8
15	9	-	11
16	9	-	8
17	13	-	6
18	8	-	9
19	14	-	12
20	13	-	6
21	9	-	8
22	5	-	12
23	5	-	12
24	13	-	15
25	8	-	5
26	12	-	5
27	13	-	14
28	5	-	8
29	5	-	8
30	14	-	12
31	13	-	14

February 2, 1990

Peak and Average Link Utilization for Jan 4 '90

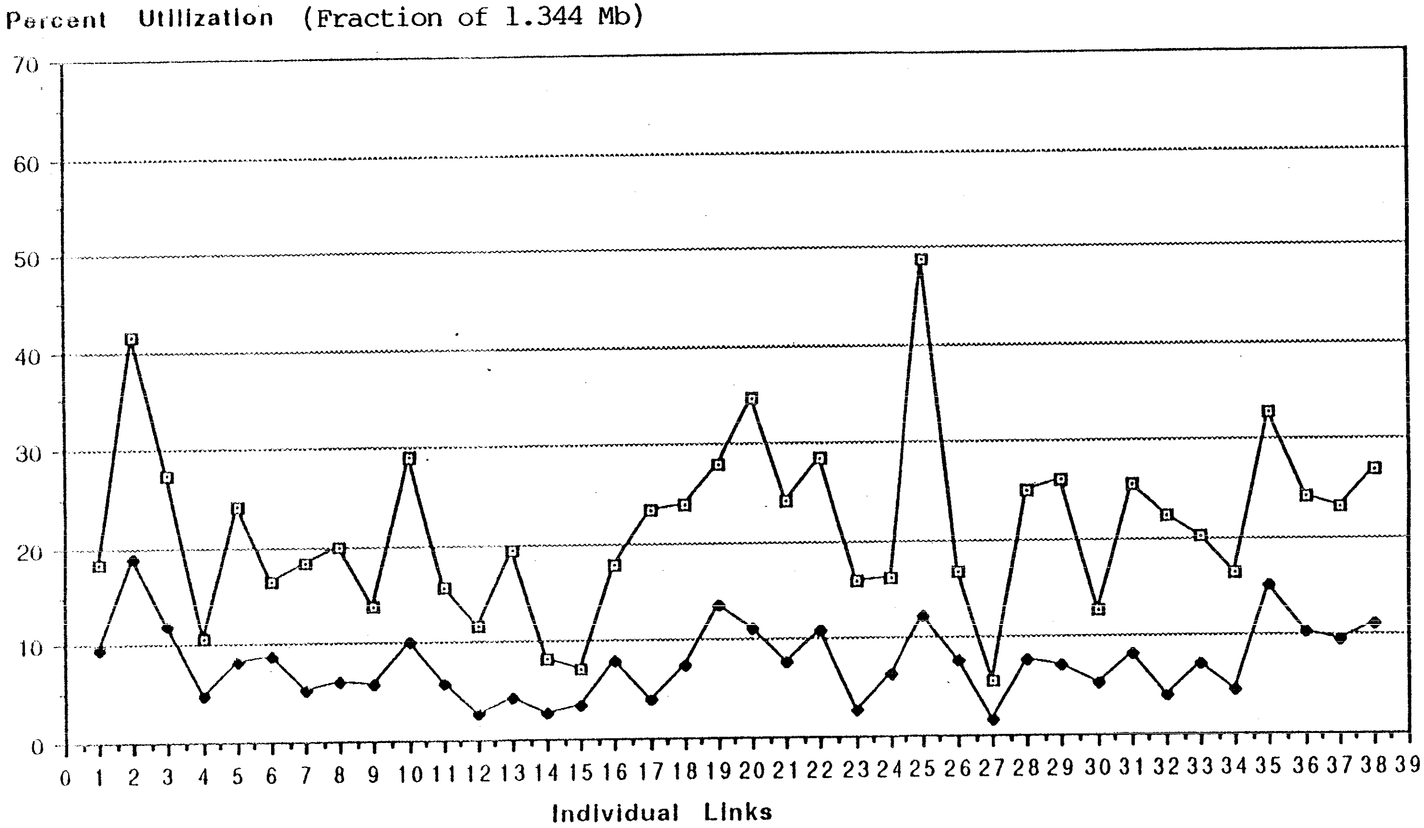

1	5	-	8
2	8	-	5
3	5	-	10
4	10	-	5
5	5	-	12
6	12	-	5
7	6	-	11
8	11	-	6
9	6	-	13
10	13	-	6
11	6	-	14
12	14	-	6
13	7	-	11
14	11	-	7
15	7	-	15
16	15	-	7
17	7	-	16
18	16	-	7
19	8	-	9
20	9	-	8
21	8	-	17
22	17	-	8
23	9	-	10
24	10	-	9
25	9	-	11
26	11	-	9
27	10	-	17
28	17	-	10
29	12	-	14
30	14	-	12
31	12	-	16
32	16	-	12
33	13	-	14
34	14	-	13
35	13	-	15
36	15	-	13
37	15	-	17
38	17	-	15

February 2, 1990

Chapter 5

Technical Presentations

5.1 "Explaining the Role of GOSIP"

Presentation by Phill Gross/ NRI

Background

The Government OSI Profile (GOSIP), issued as FIPS 146 by the National Institute of Standards and Technology (NIST), specifies the details of OSI for use in the U.S. Government.

OSI "Profiles" are important because OSI standards allow many potential options and choices. Without careful specification and prior agreement, different vendor products might very well conform to the OSI standards but not interoperate with each other. Therefore, a major goal of FIPS 146 is to insure that the U.S. government be able to buy interoperable OSI products from different commercial vendors.

The first version of GOSIP was published in August 1988 following a comment period beginning in early 1987. GOSIP was adopted as FIPS 146 in February 1989 and will become a Federal procurement requirement in August 1990 [1]. GOSIP was written by an inter-agency group and continues to evolve under the guidance of the GOSIP Advanced Requirements Group. A second version of GOSIP will become a FIPS in the summer of 1990 and will then become a Federal procurement requirement 18 months later [2].

There is an additional publication called the GOSIP Users' Guide which provides an expanded explanation of GOSIP including tutorials, interpretation, integration planning advice, and information on registration[3]. The GOSIP Users' Guide will be updated and re-released in coordination with each version of GOSIP.

The Internet Activities Board (IAB) and the Internet Engineering Task Force (IETF) are fully committed to integrate OSI into the Internet. In particular, one of the eight technical areas of concentration in the IETF is devoted to OSI integration, and the IETF is represented on the GOSIP Advanced Requirements Group.

Source of confusion?

GOSIP is an important tool for planning OSI integration. However, as the August 1990 requirement date for GOSIP compliance approaches, there has also been an increasing amount of concern as to how GOSIP should be applied to near-term network planning.

In particular, there appears to be a common misunderstanding that GOSIP mandates a transition to OSI beginning in August 1990.

For example, in the January 1990 IEEE Spectrum (Technology '90), there is the following quote in the section on Data Communications (page 35-36):

> "OSI protocols are viewed as a long-term answer to the problem. But, the scarcity of products on the market hinders devising a network strategy around OSI. ....
>
> Many vendors are still pinning their hopes for OSI on FIPS 146 (GOSIP), which requires that Federal agencies start using OSI products after August."

GOSIP does not "require" that Federal agencies start using OSI products after August 1990. GOSIP is a procurement specification. GOSIP does not mandate, or even explicitly address, the issue of protocol transition.

Some clarifying points about GOSIP

As a procurement specification, GOSIP does not apply to existing installed equipment. It applies to new network procurements.and major upgrades to existing networks. "Major upgrade" does not necessarily apply to increasing the number of components in existing non-GOSIP networks.

When GOSIP does apply, it is not exclusionary. That is, other protocol families can continue to be procured and used. When GOSIP does apply, waivers are allowed in consideration of specific agency requirements. When GOSIP does not apply, no waiver may be necessary.

Agencies have the responsibility for developing their own waiver process, and for determining the applicability of GOSIP to any specific procurement. NIST does not have an enforcement role regarding GOSIP. In general, agencies are responsible developing their own agency-wide plans for GOSIP compliance in their network procurements,

Summary

The large existing installed base of TCP/IP and other protocol users, the limited availability of commercial OSI products, and the still incomplete development of OSI standards (e.g., for routing, network management, and directory services) combine to make a near-term transition to a ubiquitous OSI environment in the Internet unrealistic.

GOSIP is an important tool for procuring interoperable OSI products for the U.S. government. However, GOSIP does not mandate, or even explicitly address, the issue of OSI transition.

This description was published in the March 1990 edition of "ConneXions, the Interoperability Newsletter".

The points made in this article will be given in more detail in a forthcoming RFC [4] issued by Vint Cerf (NRI, chair of the IAB) and Kevin Mills (NIST).

References

[1] U.S. Government Open Systems Interconnection Profile. August 1988. U.S. Federal Information Processing Standards Publication 146. Version 1.

[2] U.S. Government Open Systems Interconnection Profile. April 1989. U.S. Federal Information Processing Standards Publication 146-1.. Draft Version 2.

[3] Tim Boland. Government Open Systems Interconnection Profile Users' Guide. August 1989. NIST Special Publication 500-163.

[4] "Explaining the Role of GOSIP", RFC (to be published), Vint Cerf (IAB) and Kevin Mills (NIST).

5.2 "OSPF Update"

Presentation by John Moy/ Proteon

OVERVIEW

OSPF is a link-state or SPF-based IGP (Interior Gateway Protocol). In a link-state based protocol, each router maintains a complete topological map of the routing domain (this topological map is called the link state database). Each router synchronizes this database with its neighbors, and calculates a routing table by performing a shortest-path calculation on the database. Other examples of link-state based routing algorithms include the current Arpanet routing algorithm and the IS-IS protocol.

OSPF was designed by an IETF working group. This working group was co-chaired by Mike Petry of the University of Maryland and myself. The working group published the final OSPF specification in October of 1989 as RFC 1131. Other documents produced by the working group included a requirements document and a comparison to the IS-IS protocol. These documents, as well as intermediate versions of the OSPF protocol specification, were widely distributed throughout the community (including being sent to several foreign countries, ANSI meetings, etc.).

There are currently two independent, interoperable implementations of OSPF, one by Proteon and one by the University of Maryland. Also, the protocol has generated alot of enthusiasm among the IP user community.

Next, we review the OSPF design goals. In some respect, this is a defense of the OSPF working group, since a recent criticism of the WG is that we should have simply adopted the IS-IS protocol (the ISO routing protocol that was being developed by ANSI committee). For this reason, the design goals listed in the slides have been divided into two categories: those goals that were reasons not to simply adopt the IS-IS protocol, and other general design goals.

This paragraph lists design goals that were inconsistent with the adoption of IS-IS. First, we wanted a native IP protocol: one that would pass native IP addresses, would run directly over IP and would operate inside the normal IP framework. We also wanted a self-contained protocol specification; this specification should be easily understood by the IP community. The protocol should be designed for efficient IP processing: fragmentation should be avoided whenever possible (this problem has since been fixed in the IS-IS) and packet fields should be aligned in the usual IP fashion (n-byte quantities on n-byte boundaries). We wanted a multi-level routing hierarchy, with internal routes always preferred over external routes. We wanted an area routing scheme in which areas would be modeled after subnetted networks; appropriate information collapsing should be performed at area boundaries. Also, we wanted to start with the proven base of the BBN Arpanet SPF algorithm (which has been in operation for 10 years): in particular, using explicit acknowledgements in the flooding scheme and forcing a router to synchronize its database before it joins the routing domain.

Other design goals included the following. We wanted to support the extension to IP subnet addressing referred to as variable length subnets. We also wanted to take advantage of other work that had been done for link-state based routing: in particular, using some of the broadcast network support that had been developed by DEC and some of the improvements to the original BBN SPF algorithm that had been proposed by Radia Perlman.

COMPARISON TO DUAL IS-IS

I should mention a few things before I proceed with the comparison between OSPF and IS-IS. I am only comparing the two in terms of IP protocol functionality. Also, since I am one of the developers of OSPF, I have a definite OSPF bias and make no pretense of providing a balanced presentation. Lastly, I am comparing OSPF to the dual IS-IS draft that was published shortly before the February meeting, and the IS-IS document as it stood before the recent Paris editing session. Any changes in these two document since then is not reflected here. In particular, the dual IS-IS draft seems in places incomplete, forcing me in places to read between the lines (possibly incorrectly).

The protocol comparison is broken into several sections. First I address problems with the dual approach (running a single routing protocol to support multiple, independent networking stacks). These problems are separated into two cases: when running the dual IS-IS in an IP-only domain, and when running the dual IS-IS in a mixed IP/ISO domain. All problems are illustrated with examples taken from the dual IS-IS proposal. Then I present a list of miscellaneous differences between OSPF and IS-IS.

PROBLEMS WITH DUAL APPROACH (IP-ONLY)

The next few paragraphs detail problems encountered with the dual IS-IS approach when running in an IP-only domain.

In this case, IP routing is forced into an ISO framework, forcing the development of new capabilities and concepts that are unnecessary in an IP environment. Example: The dual IS-IS runs directly over the ISO link-level encapsulation. This encapsulation need not presently be supported by IP routers, and will in many cases be the first time that an IP router need deal with an odd-length link level header. Another example: The dual IS-IS manipulates MAC addresses, in order to establish adjacencies. This will be the first time that IP routing code will need to deal with MAC addresses (this is done by ARP in the IP world), and the ISO service primitives will need to be implemented in order to pass the link level addresses up to the IP routing layer.

Also, the dual IS-IS would carry extra baggage that is unnecessary in an IP environment. Example: As discussed above, it will carry MAC addresses which, since this is handled by the ARP protocol, is a duplication of effort. Another example: IS-IS packets will be difficult to parse, since their fields are not byte aligned.

Since IS-IS was developed as an ISO routing protocol, there may be some implicit assumptions made in the protocol that are invalid for the IP environment (since the IP and ISO environments are quite different; e.g. ISO routes to hosts instead of networks). Also, some IP features cannot be fitted into the ISO framework of IS-IS. Example: Authentication cannot be handled by the dual IS-IS, since it is not part of the IS-IS specification.

Lastly, the dual IS-IS lacks a clear specification for use as an IP-only routing protocol. This is because it is based on and refers extensively to an ISO document, which requires a good deal of ISO knowledge in order to be understood.

PROBLEMS WITH DUAL APPROACH (MIXED DOMAIN)

The next few paragraphs deal with problems encountered when using the dual approach in a mixed ISO/IP domain. The major problem here is the loss of flexibility, caused by the artificial tying together of two separate protocol stacks through the use of a single routing protocol. This is not surprising; since the main argument in favor of dual routing is conservation of resources, this conservation must have some cost.

Using the dual IS-IS, the following flexibility is lost: 1) IP and ISO areas are forced to be the same. Since areas select addressing, this ties the IP and ISO addressing schemes together (so you also lose addressing flexibility). 2) IP and ISO are forced to run over the same set of router interfaces, i.e., you cannot run only IP traffic on one router interface while running a mix of IP and ISO traffic over a router's other interfaces. 3) The link costs are forced to be the same for both IP and ISO, forcing the cost structure to be the same for both ISO and IP routing. 4) Not all IP TOS values can be supported. The IP TOS values must be mapped into the four ISO supported values before they are used, forcing a reduction of information.

When using the dual IS-IS, there are unexpected dependencies introduced between IP and ISO, which can make problems hard to debug. For example, the dual IS-IS requires (but does not enforce) all areas to be either IP-only, ISO-only, or all mixed routers. However, if for example a router inside a mixed domain is configured instead to be ISO-only, all IP routing in the domain will appear to be working, but those IP packets that are routed through the ISO-only router will unexpectedly be dropped. This is counter to intuition; you would instead expect IP traffic to automatically be routed around the ISO-only router (as it would using the SIN approach).

When using the dual IS-IS, there will be places where the IP support and ISO support will diverge. For example, the dual IS-IS wants to prefer internal IP routes over external (a concept that IS-IS does not have), so it has proposed splitting the Dijkstra calculation into two parts (as is done in OSPF), and the second part will be different for IP and ISO. Also, the dual IS-IS wanted area summaries (again as in OSPF), which are dynamic level 2 routes whose existence depends on the Level 1 Dijkstra calculation. Again, this is something not present in the ISO IS-IS that will cause ISO and IP support to diverge. Such divergences can be confusing, and will add to the difficulty of developing and maintaining the dual IS-IS protocol. For example, Proteon supports three XNS style network protocols: XNS, IPX and Apollo Domain. We try to support this with a single XNS code base, but sometimes all the special cases introduced makes this approach seem hardly worthwhile.

Lastly, I would point out that commercial multiprotocol routers on the market today use the SIN approach (which I would rather call the parallel stack approach) instead of the dual approach. In particular, Proteon supports seven different networking protocols simultaneously, each with its own separate routing protocol. Proteon has found that the parallel stack approach works well, and is both easy to analyze and debug.

OTHER DIFFERENCES

The following describes a collection of miscellaneous differences between OSPF and the dual IS-IS.

An OSPF router synchronizes its link state database with its neighbors before it joins the routing domain (unlike IS-IS), minimizing the possibility of temporary routing loop formation. OSPF supports authentication, while IS-IS does not. OSPF has no limit on path cost, where IS-IS limits the cost of any particular link to 63 and the total path cost to 1023 (this limitation can either restrict the choice of metric or restrict the size of the routing domain). OSPF supports non-broadcast multi-access networks (like X.25) where the IS-IS at best treats them as collections of point-to-point networks or stubs at the edge of the routing domain. OSPF also provides efficient inter-area routing, allowing for the intelligent calculation of exit routers. In comparison, the IS-IS simply routes to the nearest level 2 router, which may be in the opposite direction from the packet's true destination.

We believe that OSPF has a more robust and efficient flooding scheme than IS-IS (OSPF positively acknowledges all updates, and OSPF routers on a LAN synchronize with the Designated Router (and Backup) only, instead of trying to synchronize all router pairs (as in IS-IS)). OSPF specifies an incremental routing table update procedure for changes in external routes (IS-IS does not). OSPF also has two levels of external routes, allowing for an extended routing trust model and for the avoidance of metric conversion.

Finally, OSPF has a more flexible area routing scheme. In OSPF, the backbone area need not be physically connected (as required in IS-IS). Instead, OSPF uses virtual links to establish and maintain backbone connectivity. Also, OSPF routers can attach to multiple areas (in the IS-IS, a router can attach to only a single area). Finally, the OSPF area scheme was designed so that areas model subnetted networks; enough flexibility is provided by OSPF so that existing IP subnet topologies can be easily supported.

OSPF status

The OSPF specification was published in October 1989 as RFC 1131. There are two independent, interoperable OSPF implementations: one by Proteon and the other by the University of Maryland. OSPF is scheduled for inclusion into the "gated" program. There is OSPF MIB support in progress, and there is also network analyzer support available (e.g., OSPF packet parsing tools in several LAN monitors).

Rob Coltun has tested his UMD OSPF implementation in the Mitre testbed, at Stanford University and at University of Maryland (with the Proteon implementation). The Proteon implementation has been running in the Proteon internet (12-15 routers) for 4 months, is currently being deployed in SURANET (5 routers so far), and has been tested at NASA Ames and several NSF regional networks.

OSPF

Update

Topics of this Presentation

- o Overview and history of OSPF
- o Comparison to dual IS-IS
- o OSPF Status

OSPF Introduction

- Link-state (SPF) based IGP
- Designed by IETF WG
- Open Specification: RFC 1131 (10/89)
- Other WG documents:
  - Requirements doc
  - IS-IS comparison
- Two independent, interoperable implementations
- Broad user support

OSPF Design Goals

- Native IP protocol *
- Protocol specification that is :
  - Self-contained *
  - Easily understood by IP community *
- Efficient IP processing
  - No fragmentation **
  - Field alignment *
- Multi-level routing hierarchy *

* = reason not to go with IS-IS
**= original IS-IS problem that was fixed

OSPF Design Goals (continued)

- Integrated support for IP subnetting
  - Information collapsing *
  - VL subnet masks
- Start from BBN SPF base *
- Good support for broadcast nets (DEC)
- Use improvements to BBN proposed by Perlman

Comparison to Dual IS-IS

Problems with Dual Approach (IP- only)

- Force IP into ISO Framework
  - ISO encapsulation
  - MAC address manipulation
- Extra baggage
  - Packet parsing
  - MAC addresses
- Some ISO assumptions not valid for IP
  - ~~Database synchronization on serial lines assumes reliable data link~~
- Lack of clear IP specification

Problems with Dual Approach (running as dual)

- Tying ISO & IP together loses flexibility in:
  - Area configuration
  - Interface mix
  - Link cost assignments
  - TOS
- Unexpected dependencies hard to debug
  - e.g., misconfiguration of areas
- Divergences can be confusing

Other Differences:

OSPF

- Synchronizes before joining routing domain
- Supports authentication
- Has no limit on path cost
- Supports NBMA nets (e.g. x.25)
- Has efficient inter-area routing

Other Differences (continued)

OSPF

- Uses positive acknowledgement in flooding
- Synchronizes O(n) instead of O(n**2)
- Supports incremental routing table calculations
- Has two levels of external routers
- Has more flexible area scheme
    - Virtual links

OSPF'S FLEXIBLE AREA CONFIGURATIONS

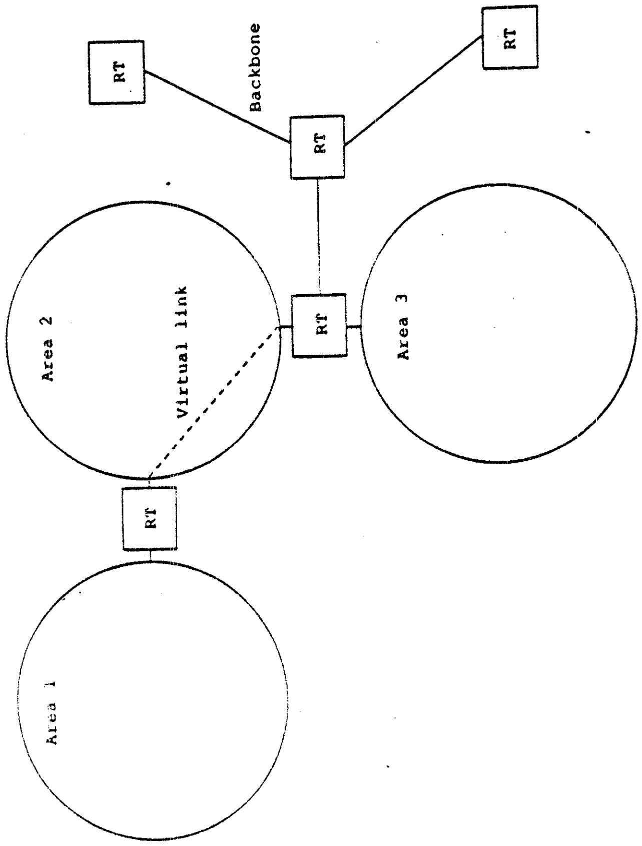

- Virtual links maintain backbone connectivity
- Routers can attach to multiple areas
- Mirrors present internet topology

OSPF Status

- RFC published (10/89)
- Two independent, interoperable implementations
  - Proteon release 8.2
  - UMD for BSD UNIX
- Schedule for gated
- MIB support in progress
- Network analyzer support available

OSPF Field Experience

- UMD Implementation tested at
  - Mitre
  - Stanford
  - UMD w/ Proteon
- Proteon Implementation
  - Running in Proteon internet (12-15 nodes) for 4 months
  - Being deployed in SURANET (5 routers so far)
  - Tested at NASA Ames and [redacted] other regionals

5.3 "Open Routing"

Presentation by Martha Steenstrup/BBN

We outlined the architecture of inter-domain policy routing, described in our Internet Draft (accessible from the online directories of the NIC and NNSC with the filename: draft-ietf-orwg-architecture-01.ps). The main features of the algorithm are the ability of a source to request and select routes according to policy requirements, the ability of an administrative domain to control access to its resources, and the ability to accommodate an Internet consisting of many administrative domains (ADs).

In order to give a source sufficient control over its routes, we have chosen link-state route generation and source routes, specified as a sequence of ADs and policy conditions. The architecture permits public and private policies and supports the notion of AD communities. To handle large Internets, the architecture allows the formation of super ADs as collections of individual ADs, uses a hierarchy of route servers dispensing routes, and reduces the essence of an AD to virtual gateways and virtual links.

A virtual gateway is the connecting fabric between two adjacent ADs. In its simplest configuration a virtual gateway consists of a pair of directly connected policy gateways, one in each AD. However, a virtual gateway may contain several directly connected policy gateways in each AD of the pair. A virtual link connects a pair of virtual gateways within an AD and carries policy restrictions imposed by the AD. More than one physical path may comprise a single virtual link. Virtual gateways and virtual links reduce the amount of information required for routing and their inherent redundancy increases fault tolerance and the ability to load share.

The principle routing functions are collection and distribution of AD virtual gateway and virtual link status, routing database maintenance, route synthesis and selection, route setup, packet forwarding, and route repair. We present a subset of the full set of protocols suggested in the Internet Draft for implementation in the initial version of the algorithm.

Inter-Domain Policy Routing

M. Steenstrup and M. Lepp

ORWG members: D. Estrin, M. Little, D. Clark, L. Zhang, N. Chiappa, P. Tsuchiya, Z-S. Su, L. Breslau, I. Castineyra, S. Resheff, P. Clark, R. Callon, B. Braden

Outline

- Why Policy Routing?
- The Internet Environment
- Overview of the Routing Architecture
- Preliminary Protocol Set

Policy Routing

- Allows users to specify routes that meet traffic requirements
- Allows networks to control how their resources are used

BBN Communications Corporation

Structure of the Internet

- Not strict hierarchy
- High-speed backbones
- Back doors
- Thousands of administrative domains (ADs)
- Many different service requirements

BBN Communications Corporation

Internet

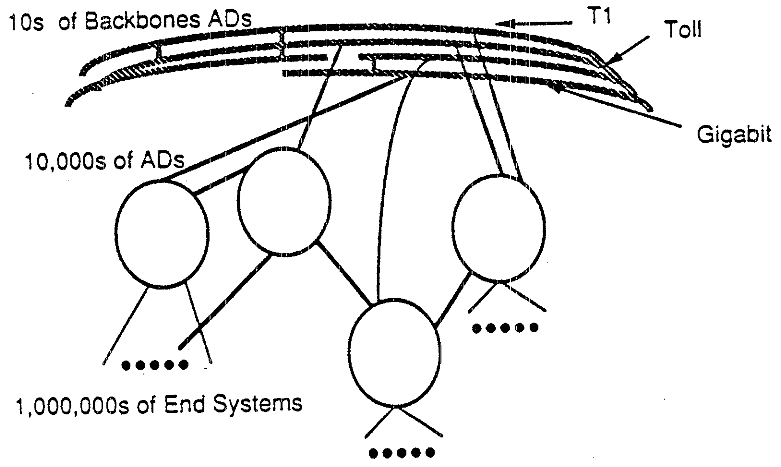

Architecture Features

- Route synthesis according to requested policy conditions
- Source routes specified as sequence of ADs and policy conditions
- Information reduction through abstraction and hierarchical organization
- Handles for security

Policies and Administrative Domains

- Service access restrictions, quality, and cost
- AD sets policy restrictions for transit traffic
- AD sets policy requirements for local users
- Public and private policies

AD Structure

- ADs connected by virtual gateways, consisting of at least two directly connected policy gateways
- Virtual gateways connected by virtual links, consisting of at least one path with a given set of policy conditions
- Superior ADs
- AD communities

Virtual Gateway

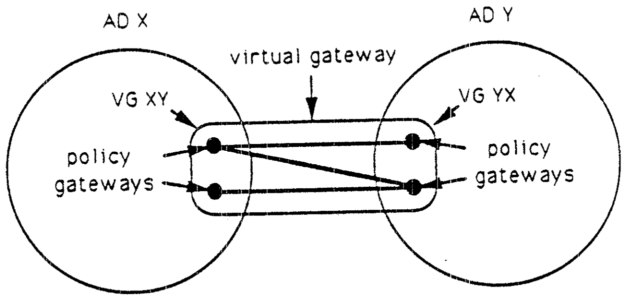

Virtual Links

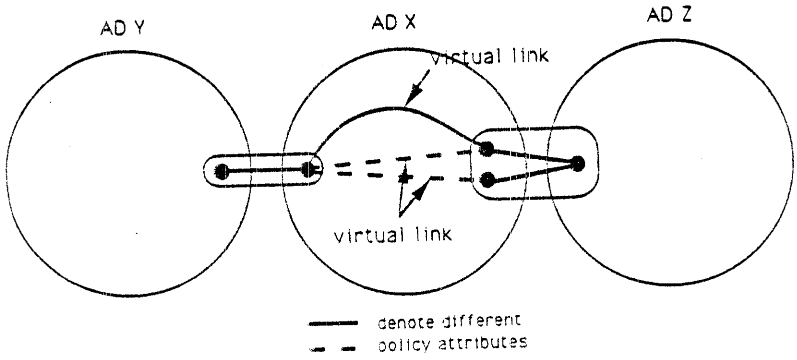

Routing Functions

- Collection and distribution of AD topology and policy information
- Route synthesis and selection
- Route setup and teardown
- Packet forwarding
- Route repair
- Database maintenance and query/response

Routing-Related Databases

- Name / address mapping
- Routes
- Global AD topology and policy information
- Local AD policies

Protocols for First Version

- Routing update
- Route setup
- Database query/response
- Virtual gateway and virtual link protocols
- Route synthesis algorithm

What First Version Gives You

- All basic routing functions
- Access restrictions based on source, destination, and time of day
- Quality of service specification based cost and bandwidth
- Synthesis of minimum hop routes that respect access restrictions and that account for cost and bandwidth requirements
- Minimum host participation required

5.4 "Use of OSI IS-IS in IP and Dual Environments"

Presentation by Radia Perlman/DEC

Overview of OSI IS-IS for Routing in TCP/IP and Dual Environments

These rough notes outline an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments, and dual environments. For further detail, see the Internet Draft "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments" (January 1990).

Protocol Overview

The TCP/IP protocol suite has been growing in importance as a multi-vendor communications architecture. With the anticipated emergence of OSI, we expect coexistence of TCP/IP and OSI to continue for an extended period of time. There is a critical need for routers to support both IP traffic and OSI traffic in parallel.

There are two main methods that are available for routing protocols to support dual OSI and IP routers. One method, known as "Ships in the Night", makes use of completely independent routing protocols for each of the two protocol suites. These notes describe an alternate approach, which makes use of a single integrated protocol for routing both protocol suites.

By supporting both IP and OSI traffic, this integrated protocol design supports traffic to IP hosts, OSI end systems, and dual end systems. This approach is "integrated" in the sense that the IS-IS protocol can be used to support pure-IP environments, pure-OSI environments, and dual environments. In addition, this approach allows interconnection of dual (IP and OSI) routing domains with other dual domains, with IP-only domains, and with OSI-only domains.

The protocol described here is based on the work of the IETF IS-IS working group.

What the Integrated IS-IS offers

The integrated IS-IS provides a single routing protocol which will simultaneously provide efficient routing for TCP/IP, and for OSI. This design makes use of the OSI IS-IS routing protocol, augmented with IP-specific information. This design provides explicit support for IP subnetting, variable subnet masks, TOS-based routing, and external routing. There is provision for authentication information, although the precise form of authentication to be used is outside of the scope of this document. IP reachability information (i.e., information specifying which IP addresses are reachable by each router) is carried independently from OSI reachability information, allowing independent address assignment for each protocol suite. Similarly, the external routing information (information about routes external to the routing domain) is carried independently for the two suites.

Both OSI and IP packets are routed "as is" – i.e., they are transmitted directly over the underlying link layer services without the need for mutual encapsulation. The integrated IS-IS is a dynamic routing protocol, based on the SPF (Dijkstra) routing algorithm.

The Integrated IS-IS Protocol allows for mixing of IP-only, OSI-only, and dual (IP and OSI) routers, as defined below.

An IP-only IS-IS router (or "IP-only" router) is defined to be a router which: (i) Uses the IS-IS protocol for routing IP packets, as specified in this report; and (ii) Does not otherwise support OSI protocols. For example, such routers would not be able to forward OSI CLNP packets.

An OSI-only router is defined to be a router which uses the IS-IS protocol for routing OSI packets. Generally, OSI-only routers may be expected to conform to OSI standards, and may be implemented independent of this specification.

A dual IS-IS router (or "dual" router) is defined to be a router which uses the IS-IS protocol for routing both IP and OSI packets, as specified in this report.

This approach does not change the way that IP packets are handled. IP-only and dual routers are required to conform to the requirements of Internet Gateways. The integrated IS-IS protocol described in this report outlines an Interior Gateway Protocol (IGP) which will provide routing within a TCP/IP routing domain (i.e., autonomous system). Other aspects of router performance (e.g., operation of ICMP, ARP, EGP, etc.) are not affected by this proposal.

Similarly, this approach does not change the way that OSI packets are handled. There will be no change at all to the contents nor to the handling of ISO 8473 Data packets and Error Reports, nor to ISO 9542 Redirects, ES Hellos, and IS Hellos. Other OSI packets (specifically those involved in the IS to IS intra-domain routing protocol) remain unchanged except for the addition of IP-related information.

This approach makes use of the existing IS-IS packets, with IP-specific fields added. Specifically: (i) authentication information may be added to all IS-IS packets; (ii) the protocols supported by each router, as well as each router's IP addresses, are specified in IS-IS Hello and LSP packets; (iii) internally reachable IP addresses are specified in all LSP packets; and (iv) externally reachable IP addresses, and external routing protocol information, may be specified in level 2 LSP packets.

The protocol described in this report may be used to provide routing in an IP-only routing domain, in which all routers are IP-only. Similarly, this protocol may be used to provide routing in a pure dual domain, in which all routers are dual. Finally, this protocol may be used to provide routing in a mixed domain, in which some routers are IP-only, some routers are OSI-only, and some routers are dual. The specific topological restrictions which apply in this latter case are described below (under "Support of Mixed Routing Domains"). The use of IS-IS for support of pure OSI domains is specified in the associates OSI specification.

The Integrated IS-IS protocol specification does not constrain which network management protocol(s) may be used to manage IS-IS-based routers. Management information bases (MIBs) for managing IP-only, OSI-only, and dual routers, compatible with CMIP, CMOT, and/or SNMP, are the subject of a separate, companion document.

Overview of the ISO IS-IS Protocol

The IS-IS Routing Protocol has been developed in ISO to provide routing for pure OSI environments. In particular, IS-IS is designed to work in conjunction with ISO 8473 (The ISO Connectionless Network Layer Protocol), and ISO 9542 (The ISO End System to Intermediate System Protocol). This section briefly describes the manner in which IS-IS is used to support pure OSI environments. Enhancements for support of IP and dual environments are described elsewhere in this note.

In IS-IS, the network is partitioned into "routing domains". The boundaries of routing domains are defined by network management, by setting some links to be "exterior links". If a link is marked as "exterior", no IS-IS routing messages are sent on that link.

Currently, ISO does not have a standard for inter-domain routing (i.e., for routing between separate autonomous routing domains). Instead, manual configuration is used. The link is statically configured with the set of address prefixes reachable via that link, and with the method by which they can be reached (such as the DTE address to be dialed to reach that address, or the fact that the DTE address should be extracted from the IDP portion of the ISO address).

OSI IS-IS routing makes use of two-level hierarchical routing. A routing domain is partitioned into "areas". Level 1 routers know the topology in their area, including all routers and end systems in their area. However, level 1 routers do not know the identity of routers or destinations outside of their area. Level 1 routers forward all traffic for destinations outside of their area to a level 2 router in their area. Similarly, level 2 routers know the level 2 topology, and know which addresses are reachable via each level 2 router. However, level 2 routers do not need to know the topology within any level 1 area, except to the extent that a level 2 router may also be a level 1 router within a single area. Only level 2 routers can exchange data packets or routing information directly with external routers located outside of the routing domains.

ISO addresses are subdivided into the Initial Domain Part (IDP), and the Domain Specific Part (DSP). The IDP is the part which is standardized by ISO, and specifies the format and authority responsible for assigning the rest of the address. The DSP is assigned by whatever addressing authority is specified by the IDP. The DSP is further subdivided into a "High Order Part of DSP" (HO-DSP), a Local Area (LOC-AREA), a system identifier (ID), and an NSAP selector (SEL). The HO-DSP may use any format desired by the authority which is identified by the IDP. Together, the IDP and the HO-DSP identify the routing domain. The LOC-area identifies the area within the routing domain.

In some cases, a single routing domain may use more than one (IDP,HO-DSP) combination, and LOC-AREA may be assigned independently for each (IDP,HO-DSP) combination. However, in all cases the combination of the (IDP, HO-DSP, and LOC-AREA) will identify the area. This combination may therefore be referred to as the "Area Address". Usually, all nodes in an area have the same area address. However, sometimes an area might have multiple addresses. Motivations for allowing this are:

- It might be desirable to change the address of an area. The most graceful way of changing an area from having address A to having address B is to first allow it to have both addresses A and B, and then after all nodes in the area have been modified to recognize both addresses, then one by one the nodes can be modified to forget address A.
- It might be desirable to merge areas A and B into one area. The method for accomplishing this is to, one by one, add knowledge of address B into the A partition, and similarly add knowledge of address A into the B partition.
- It might be desirable to partition an area C into two areas, A and B (where "A" might equal "C", in which case this example becomes one of removing a portion of an area). This would be accomplished by first introducing knowledge of address A into the appropriate nodes (those destined to become area A), and knowledge of address B into the appropriate nodes, and then one by one

removing knowledge of address C.

Since OSI addressing explicitly identifies the area, it is very easy for level 1 routers to identify packets going to destinations outside of their area, which need to be forwarded to level 2 routers.

In IS-IS, there are two types of routers:

- Level 1 intermediate systems – these nodes route based on the ID portion of the ISO address. They route within an area. They recognize, based on the destination address in a packet, whether the destination is within the area. If so, they route towards the destination. If not, they route to the nearest level 2 router.
- Level 2 intermediate systems – these nodes route based on the area portion of the ISO address. They route towards areas, without regard to the internal structure of an area. A level 2 IS is also a level 1 IS in one area.

A level 1 router will have the area portion of its address manually configured. It will refuse to become a neighbor with a node whose area addresses do not overlap its area addresses. However, if level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, then the level 1 router will accept the other node as a neighbor.

A level 2 router will accept another level 2 router as a neighbor, regardless of area address. However, if the area addresses do not overlap, the link would be considered by both routers to be "level 2 only", and only level 2 LSPs would flow on the link. External links (to other routing domains) must be from level 2 routers.

IS-IS provides an optional partition repair function. In the unlikely case that a level 1 area become partitioned, this function, if implemented, allows the partition to be repaired via use of level 2 routes.

IS-IS requires that the set of level 2 routers be connected. Should the level 2 backbone become partitioned, there is no provision for use of level 1 links to repair a level 2 partition.

In unusual cases, a single level 2 router may lose connectivity to the level 2 backbone. In this case the level 2 router will indicate in its level 1 LSPs that it is not "attached", thereby allowing level 1 routers in the area to route traffic for outside of the domain to a different level 2 router. Level 1 routers therefore route traffic to destinations outside of their area only to level 2 routers which indicate in their level 1 LSPs that they are "attached".

An end system may autoconfigure the area portion of its address by extracting the area portion of a neighboring router's address. If this is the case, then an endnode will always accept a router as a neighbor. Since the standard does not specify that the end system MUST autoconfigure its area address, an end system may be configured with an area address. In this case the end system would ignore router neighbors with non-matching area addresses.

The IS-IS provides for optional Quality of Service (QOS) routing, based on throughput (the default metric), delay, expense, or residual error probability.

Overview of the Integrated IS-IS

The integrated IS-IS allows a single routing protocol to be used to route both IP and OSI packets. This implies that the same two-level hierarchy will be used for both IP and OSI routing. Each area will be specified to be either IP-only (only IP traffic can be routed in that particular area), OSI-only (only OSI traffic can be routed in that area), or dual (both IP and OSI traffic can be routed in the area). This proposal does not allow for partial overlap of OSI and IP areas.

Similarly, within an IP-only or dual area, the amount of knowledge maintained by routers about specific IP destinations will be as similar as possible as for OSI. For example, IP-capable level 1 routers will maintain the topology within the area, and will be able to route directly to IP destinations within the area. However, IP-capable level 1 routers will not maintain information about destinations outside of the area. Just as in normal OSI routing, traffic to destinations outside of the area will be forwarded to the nearest level 2 router. Since IP routes to subnets, rather than to specific end systems, IP routers will not need to keep nor distribute lists of IP host identifiers.

The IP address structure allows networks to be partitioned into subnets, and allows subnets to be recursively subdivided into smaller subnets. However, it is undesirable to require any specific relationship between IP subnets and IS-IS areas. For example, in many cases, the dual routers may be installed into existing environments, which already have assigned IP and/or OSI addresses. In addition, even if IP addresses are not already pre-assigned, the address limitations of IP constrain what addresses may be assigned. We therefore will not require any specific relationship between IP addresses and the area structure. Reachability information (i.e., information about which addresses are reachable by each router and area) will be carried independently for the two protocol suites. Somewhat greater efficiency and scaling of the routing algorithm can be achieved if there is some correspondence between the IP address assignment structure and the area structure.

Within an area, level 1 routers exchange link state packets which identify the IP addresses reachable by each router. Specifically, zero or more IP address, subnet mask, metric combinations may be included in each LSP packet. Each level 1 router is manually configured with the IP address, subnet mask, metric combinations which are reachable on each interface. A level 1 router routes as follows:

- If a specified destination address matches an IP address, subnet mask, metric reachable within the area, the packet is routed via level 1 routing. If the specified destination address matches more than one IP address, subnet mask pair reachable within the area, the more specific address is the one routed towards (the one with more "1" bits in the mask – this is also know as "best match" routing).
- If a specified destination address does not match any IP address, subnet mask, metric combination listed as reachable within the area, the packet is routed towards the nearest level 2 router.

Flexible use of the limited IP address space is important in order to cope with the anticipated growth of IP environments. Thus an area (and by implication a routing domain) may simultaneously make use of a variety of different address masks for different subnets in the area (or domain).

Level 2 routers include in their level 2 LSPs a complete list of IP address, subnet mask, metric specifying all IP addresses reachable in their area. This information may be obtained from a combination of the level 1 LSPs (obtained from level 1 routers in the same area), and/or by manual configuration. In addition, Level 2 routers may report external reachability information, corresponding to addresses which can be reached via routers in other routing domains (autonomous systems).

Default routes may be announced by use of a subnet mask containing all zeroes. Default routes should be used with great care, since they can result in "black holes". Announcement of a default route by a level 1 router (in its level 1 LSP) will prevent routing of packets from that area via level 2 routing.

The integrated IS-IS provides Type of Service (TOS) routing, through use of the QOS feature from IS-IS.

Support of Mixed Routing Domains

The integrated IS-IS proposal specifically allows for three types of routing domains:

- Pure IP
- Pure OSI

- Dual

In a pure IP routing domain, all routers must be IP-capable. IP-only routers may be freely mixed with dual routers. Some fields specifically related to OSI operation may be included by dual routers, and will be ignored by IP-only routers. Only IP traffic will be routed in an pure IP domain. Any OSI traffic may be discarded (except for the IS-IS packets necessary for operation of the routing protocol).

In a pure OSI routing domain, all routers must be OSI-capable. OSI-only routers may be freely mixed with dual routers. Some fields specifically related to IP operation may be included by dual routers, and will be ignored by OSI-only routers. Only OSI traffic will be routed in a pure OSI domain. Any IP traffic may be discarded.

In a dual routing domain, IP-only, OSI-only, and dual routers may be mixed on a per-area basis. Specifically, each area may itself be defined to be pure IP, pure OSI, or dual.

In a pure IP area within a dual domain, IP-only and dual routers may be freely mixed. Only IP traffic can be routed by level 1 routing within a pure-IP area.

In a pure-OSI area within a dual domain, OSI-only and dual routers may be freely mixed. Only OSI traffic can be routed by level 1 routing within a pure OSI area.

In a dual area within a dual routing domain, only dual routers may be used. Both IP and OSI traffic may be routed within a dual area.

Within a dual domain, if both IP and OSI traffic are to be routed between areas, then all level 2 routers must be dual.

The integrated IS-IS protocol does not provide for encapsulation of OSI packets within IP packets, nor of IP packets within OSI packets. It is therefore not possible to route IP packets through OSI-only routers, nor to route OSI packets through IP-only routers. At some point in the future, optional mechanisms may be defined to allow encapsulation for this purpose. However, such mechanisms will be optional, and dual routers will not be required to provide encapsulation.

Advantages of Using Integrated IS-IS

Use of the integrated IS-IS protocol, as a single protocol for routing both IP and OSI packets in a dual environment, has significant advantages over using separate protocols for independently routing IP and OSI traffic.

An alternative approach is known as "Ships In the Night" (S.I.N.). With the S.I.N. approach, completely separate routing protocols are used for IP and for OSI. For example, may be used for routing IP traffic, and IS-IS may be used for routing OSI traffic. With S.I.N., the two routing protocols operate more or less independently. However, dual routers will need to implement both routing protocols, and therefore there will be some degree of competition for resources.

Note that S.I.N. and the integrated IS-IS approach are not really completely separate options. In particular, if the integrated IS-IS is used within a routing domain for routing of IP and OSI traffic, it is still possible to use other independent routing protocols for routing other protocol suites. In the future, optional extensions to IS-IS may be defined for routing other common protocol suites. However, such future options are outside of the scope of this document. This section will compare integrated IS-IS and S.I.N. for routing of IP and OSI only.

A primary advantage of the integrated IS-IS is that, since it requires only one routing protocol, it uses fewer resources. In particular, less implementation resources are needed (since only one protocol needs to be implemented), less CPU and memory resources are used in the router (since only one protocol needs to be run), and less network resources are used (since only one set of routing packets need to be transmitted). Primarily this translates into a financial savings, since each of these three types of resources cost money. This implies that dual routers based on the integrated IS-IS should be less expensive to purchase and operate than dual routers based on S.I.N.

Another advantage of the integrated IS-IS relates to the network management effort required. Since the integrated IS- IS requires only one protocol, there is less information for the operator to configure.

Note that the operation of two routing protocols with the S.I.N. approach are not really independent, since they must share common resources. For example, if one routing protocol becomes unstable and starts to use excessive resources, the other protocol is likely to suffer. A bug in one protocol could crash the other. However, with the integrated IS-IS, the interactions are explicit, whereas with S.I.N., the interactions are implicit.

The use of a single integrated routing protocol similarly reduces the likely frequency of software upgrades. Specifically, if you have two different routing protocols in your router, then you have to upgrade the software any time EITHER of the protocols change. If you make use of a single integrated routing protocol, then software changes are still likely to be needed, but less frequently.

Finally, routing protocols have significant real time requirements. In IS-IS, these real time requirements have been explicitly specified. In other routing protocols, these requirements are implicit. However, in all routing protocols, there are real time guarantees which must be met in order to ensure correct operation. In general, it is difficult enough to ensure compliance with real time requirements in the implementation of a single real time system. With S.I.N., implementation of two semi-independent real-time protocols in a single device makes this more difficult.

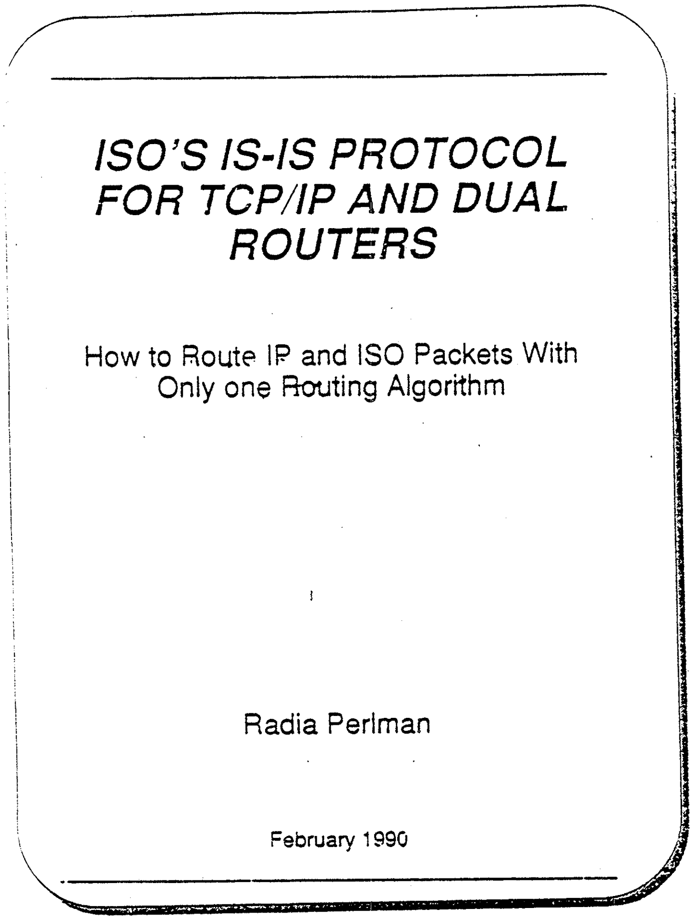

ISO'S IS-IS PROTOCOL FOR TCP/IP AND DUAL ROUTERS

How to Route IP and ISO Packets With Only one Routing Algorithm

Radia Perlman

February 1990

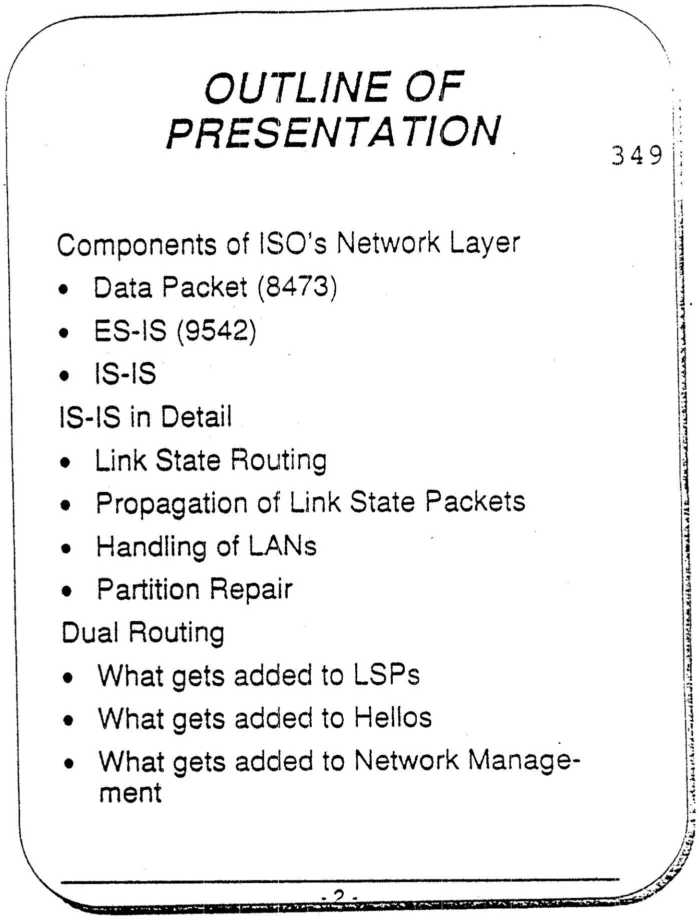

OUTLINE OF PRESENTATION

Components of ISO's Network Layer

- Data Packet (8473)
- ES-IS (9542)
- IS-IS

IS-IS in Detail

- Link State Routing
- Propagation of Link State Packets
- Handling of LANs
- Partition Repair

Dual Routing

- What gets added to LSPs
- What gets added to Hellos
- What gets added to Network Management

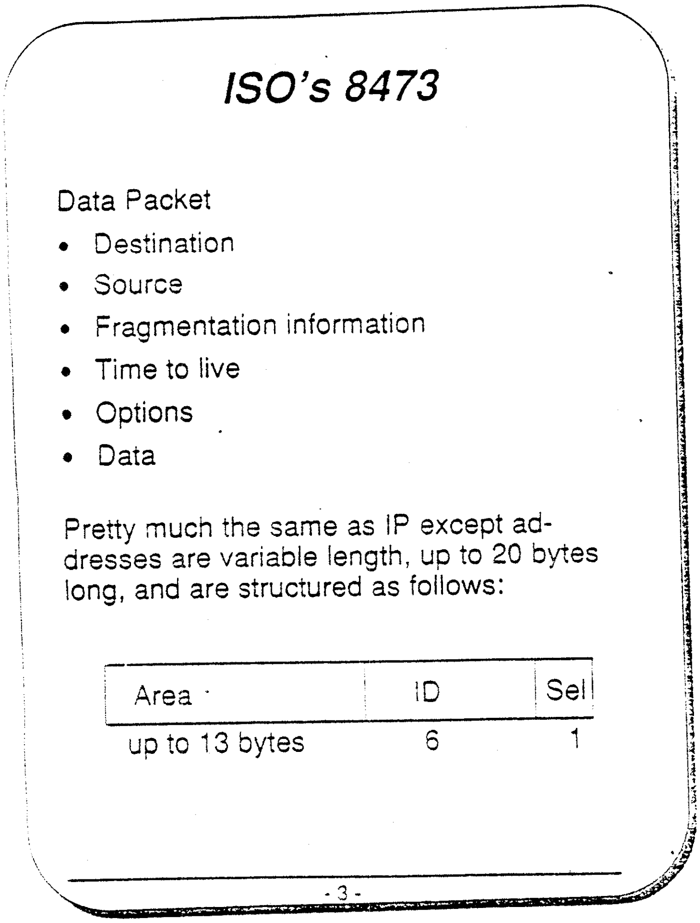

ISO's 8473

Data Packet

- Destination
- Source
- Fragmentation information
- Time to live
- Options
- Data

Pretty much the same as IP except addresses are variable length, up to 20 bytes long, and are structured as follows:

Area	ID	Sel
up to 13 bytes	6	1

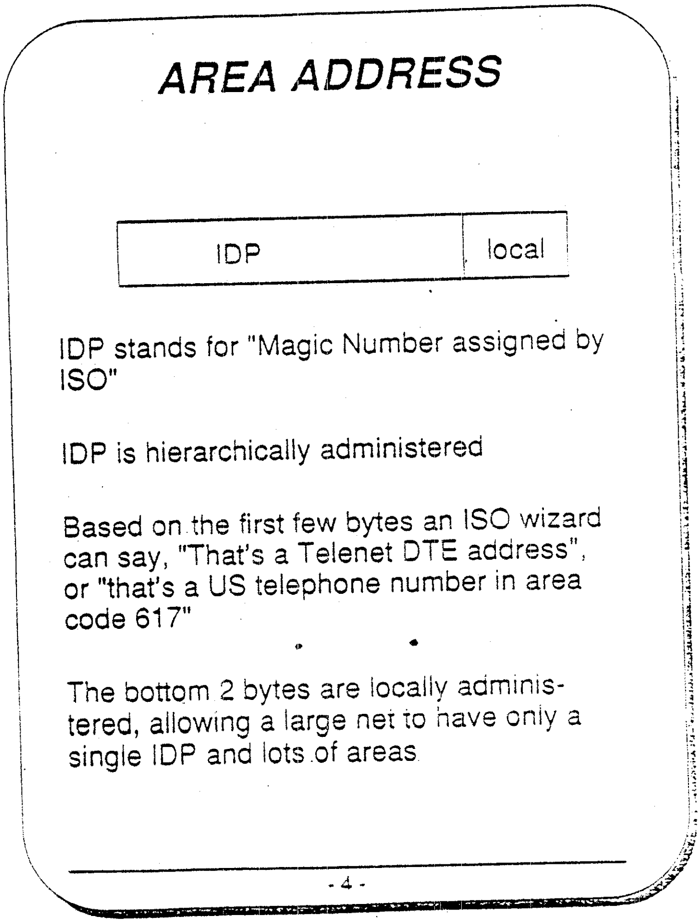

AREA ADDRESS

IDP	local

IDP stands for "Magic Number assigned by ISO"

IDP is hierarchically administered

Based on the first few bytes an ISO wizard can say, "That's a Telenet DTE address", or "that's a US telephone number in area code 617"

The bottom 2 bytes are locally administered, allowing a large net to have only a single IDP and lots of areas

EXAMPLE

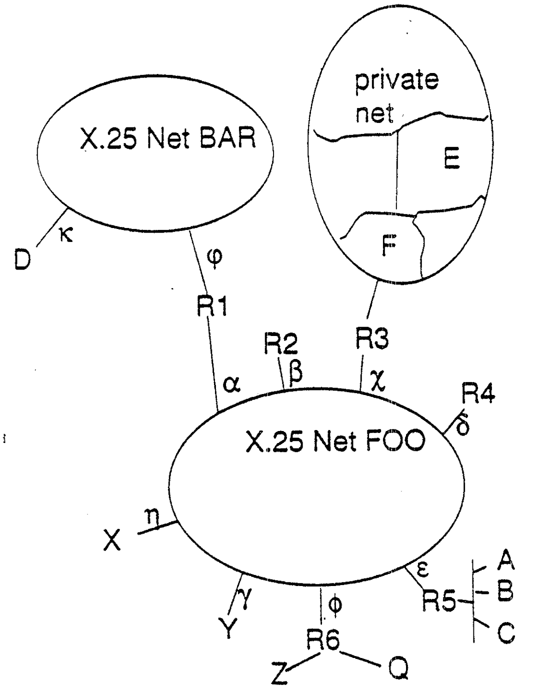

ISO's ES-IS

ES = End System (Endnode, Host)

IS = Intermediate System (Router, Gateway)

IS's Find Out about ES's

- ES's periodically issue Hellos
- Sent to "All IS's" Multicast Address
- Routers learn Network Layer Address
- Learn Data Link/Network Layer Address correspondence

ES's Find Out about IS's

- IS's periodically issue Hellos to "All ES's"
- IS's issue Redirects, either to another router or to the destination -- Redirect is to a Data Link Layer address

EQUIVALENT IN IP

Routers do not exchange information about individual ES's

LAN number part of address

Only nodes on the LAN keep track of individuals on the LAN, and keep cache of Data Link Layer address through ARP protocol

Need some method for endnodes to know about routers

IS-IS

Link State Routing

- Computation of Routes
- Dealing with LANs
- Propagation of Link State Packets

Extra goodies

- (Optional) multiple routing metrics with Type of Service routing
- Partition repair

LINK STATE ROUTING

Each IS discovers neighbors, constructs a Link State Packet (LSP)

Each IS broadcasts its LSP to all other ISs

Each IS keeps the most recently generated LSP from every other IS

Each IS computes routes (Dijkstra's algorithm)

Similar to the "New ARPANET" routing algorithm

Also called SPF Routing

HIERARCHICAL ROUTING

IF IN AREA, ROUTE DIRECT

ELSE ROUTE TO NEAREST LEVEL 2 ROUTER

LEVEL 1 ROUTERS DO NOT KNOW (OR CARE) WHICH LEVEL 2 ROUTER GIVES OPTIMAL ROUTE

SEQUENCE NUMBERS

Large linear space

Top reached only when errors have occurred

When top of sequence number reached in LSP for IS "R", R purges the old LSP and needs to wait for a few minutes before a new LSP (with low sequence number) will be believed by the network

Linear space is simpler, and more robust than earlier schemes (circular, lollipop)

DEALING WITH LANS

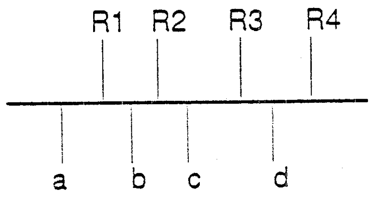

One IS gets elected "Designated Router", based on ID (and priority)

That IS names the LAN -- its 6 byte ID plus an extra byte, in case it is DR on multiple LANs (e.g., R4.5)

Each router just claims in its LSP that it has neighbor R4.5

Router R4 issues an extra LSP with source "R4.5" that lists all the routers and endnodes

LAN LSP PROPAGATION

Instead of explicit acknowledgements, the DR periodically (O(10 seconds)) issues a "Complete Sequence Numbers Packet" (CSNP)

The CSNP summarizes the LSP database, by including a list of LSP source/sequence number correspondence

If the DR has failed to receive an LSP, the router(s) that have that LSP will respond to the CSNP by transmitting the LSP

If another router has failed to receive an LSP, it will send a request to the DR, which will transmit the missing LSP

LARGE LSP PROPAGATION

Suppose LSP is too large to send in one packet (for instance, endnode membership of large LAN)

- "Source" of LSP actually contains an extra field "fragment number"
- Each fragment is propagated by the LSP Propagation Algorithm independently
- Each fragment has an independent sequence number and age
- Only the Routing Computation Algorithm is aware that there is any connection between fragments
- If only one fragment changes, then only it needs to be rebroadcast

LARGE CSNP

What if CSNP is too large to fit into one packet?

CSNP contains field "Address range".

CSNP can be broken into multiple CSNPs, each with a different address range

Each CSNP "fragment" can be acted on independently -- no need to get all fragments

PARTITION REPAIR

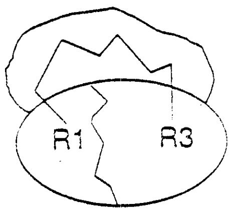

- Based on ID, a Level 2 router is named the "Partition Designated Level 2 Router"
- Level 2 LSPs give the attached area number and the Designated L2 Router
- If L2 routers both claim same area, but different Designated L2 Router, then there's a partition
- Glue the pieces together with virtual L1 link
- Level 1 routers never aware there was a partition
- This feature is optional: repair if partition capable L2 router in partition

OTHER GOODIES

PERFORMANCE REQUIREMENTS CLEARLY SPECIFIED

CORRECT OPERATION ENSURED UNDER MEMORY OVERLOAD CONDITIONS

COMPLETE MIB DEFINED

DUAL ROUTING

One routing protocol based on IS-IS augmented with IP information

Compatible with pure ISO routers

Conforming with ~~OSI~~ IP standards

Conforming with Internet standards

Three types of routers

- OSI only -- will not understand the additional IP information, will not understand IP data packets
- IP only -- will not forward 8473 Data Packets. DOES NOT NEED 9542
- Dual -- will handle both data packet formats

MIXING OF ROUTERS

An Area (or the level 2 backbone) will be either:

- Pure OSI -- all routers are OSI only or dual, and all data packets in the area are 8473
- Pure IP -- all routers are IP only or dual and all data packets in the area are IP
- Dual -- all routers in the area are dual, and both types of data packets are routed

In the future, optional encapsulation mechanisms may be defined to allow mixing of dual, IP only, and OSI only routers within an area

IP ADDRESS SUMMARY

- Each router (within an area) is configured with the set of {IP address, subnet mask, metric(s)} reachable over each interface
- These addresses are included in the level 1 LSP, and level 1 routing computes routes to {IP address, subnet mask} pairs
- Each level 2 router may additionally be configured with a set of {IP address, subnet mask, metric(s)}, representing the addresses reachable within the area
- The level 2 router announces a configured address provided at least one matching IP address is reachable via level 1 routing
- The level 2 router announces any other reachable {IP address, subnet mask} pairs that are reachable, and not already included in the summary info

IP SPECIFIC INFORMATION

IP INTERNAL REACHABILITY INFORMATION

- Included in L1 LSPs
- Included in L2 LSPs
- Contains a list of zero or more of
  - -- IP address
  - -- subnet mask
  - -- metrics
    - > must contain default
    - > may contain also delay, expense, and/or error
- L1 -- addresses directly reachable
- L2 -- addresses directly reachable or reachable via level 1 routing

IP SPECIFIC INFORMATION

IP EXTERNAL REACHABILITY INFORMATION

- Included only in Level 2 LSPs
- Similar in content to IP INTERNAL REACHABILITY INFORMATION
- Includes entries discovered through a direct link to an external router (e.g. EGP)

IS-IS REQUIRES ROUTERS TO HAVE A UNIQUE IDENTIFIER

ALL ROUTERS IN AN AREA MUST HAVE THE SAME AREA ADDRESS

AREA CAN BE OSI OBTAINED, OR DERIVED FROM "AS" BY ADDING MAGIC CONSTANT

ID CAN BE IEEE 48 BIT STATION ID OR TCP/IP ADDRESS + CONSTANT

IP SPECIFIC INFORMATION

- Protocols supported
  - -- in Hellos allows routers to know if neighbor is compatible
  - -- in LSPs allows routers to know if another router in the area (or level 2 backbone) is compatible
    - > easier to catch misconfiguration errors
    - > maybe useful in the future for encapsulation
  - -- if absent, indicates pure ISO
  - -- one value for IP
  - -- one value for ~~OSI~~ CLNP
  - -- multiple values can appear
  - -- might include other protocols in the future

IP SPECIFIC INFORMATION

IP INTERFACE ADDRESS

- In Hello Packets, because ICMP Redirect protocol requires IP Layer address of router Redirection is toward
- In LSPs, allows other routers to know a set of IP Network Layer addresses to use to send packets to that router
  - -- encapsulation in future
  - -- partition repair

AUTHENTICATION INFORMATION

- Optional in all PDU's. Contents outside the scope of this specification

IP SPECIFIC INFORMATION

INTERDOMAIN ROUTING PROTOCOL INFORMATION

- In Level 2 LSPs
- For the convenience of external routing protocol -- for instance, allows external border gateways to find each other
- Contents of field outside the scope of this specification

CONCLUSION

We've defined a single routing protocol which can support

- pure IP environments
- pure OSI environments
- dual environments

In dual environments, mixing can be independently in each area

For IP this supports IP subnetting, variable subnet masks, type of service routing, external routing, authentication, and partition repair

For IP this defines an IGP -- remainder (ARP, ICMP, ...) unchanged

Good solution for IP only environment, but particularly efficient for dual environment

STATUS

IS-IS

DP STATUS

FIRST DRAFT TO ANSI 4/87
ISO 6/88

PRECURSOR
RUNNING FOR 2 YEARS NSFNET
≈ 20 ROUTERS

AT DEC RUNNING IN NETWORKS
≈ 20 ROUTERS

DUAL ROUTING

INTERNET DRAFT

PUBLIC DOMAIN IMPLEMENTATION
BEING BUILT BY U OF WISC

ALPHA TEST CODE BY SHANER

PRODUCTION QUALITY BY
END OF YEAR

5.5 "From Smart Drop to Congestion Control"

Presentation by Martha Steenstrup/BBN

Congestion – the phenomenon of increased delay and reduced throughput in response to high offered load – is a potential problem in any packet-switching network where mismatched transmission rates or convergent flows are common, and in particular, is a phenomenon well-known to the Internet community. We propose a rate-based congestion control algorithm for the Internet, in which routers play the principal role in determining flow rates. Each router periodically computes maximum acceptable flow rates, called rations, for each of its associated backbone links and access networks, and for its processors. Routers refuse packets from flows that attempt to use more than their ration. Participating hosts collect ration information from the routers and use these ration values to set the rates at which they submit traffic to the Internet. Thus, both routers and cooperating hosts exercise control of flow rates.

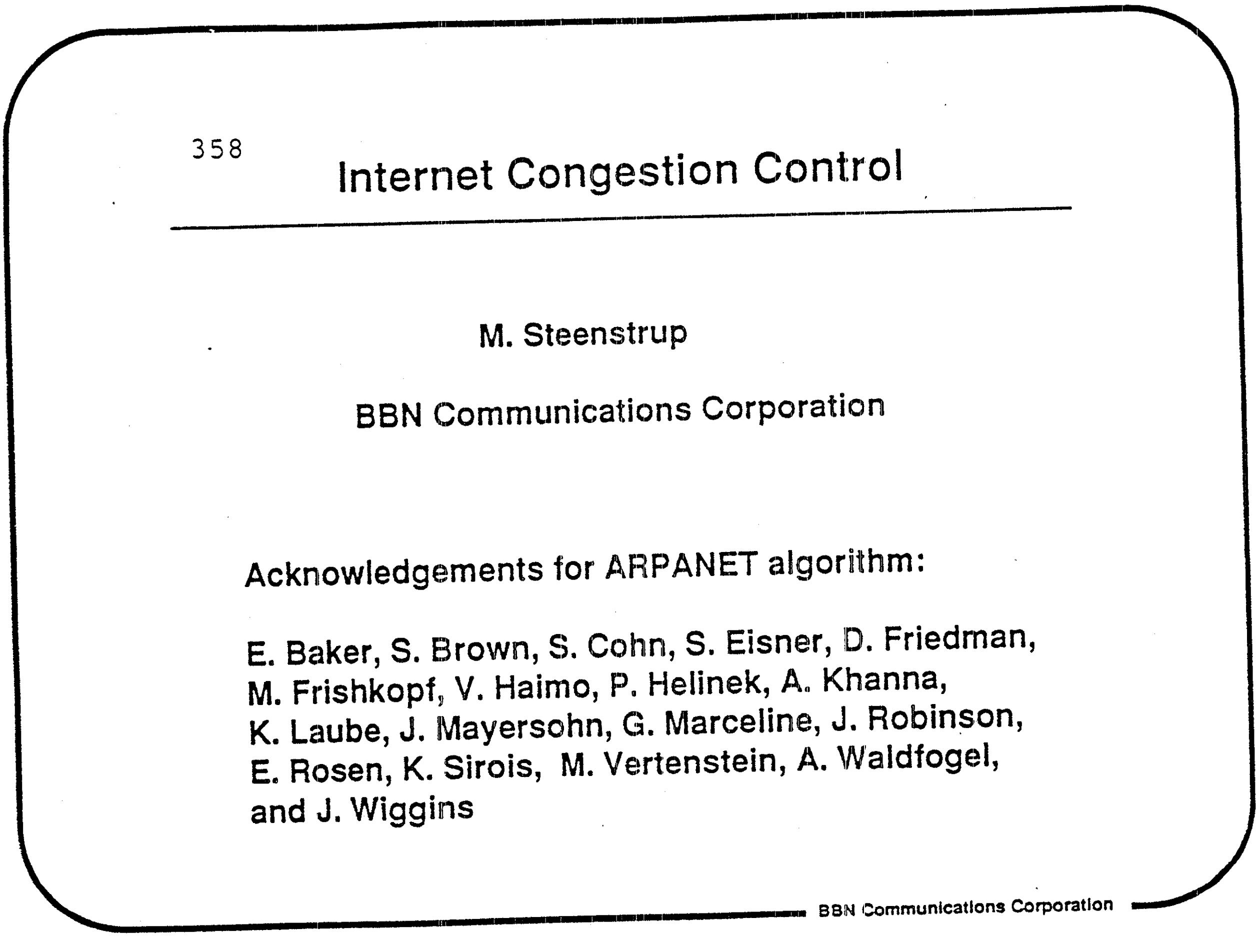

Internet Congestion Control

M. Steenstrup

BBN Communications Corporation

Acknowledgements for ARPANET algorithm:

E. Baker, S. Brown, S. Cohn, S. Eisner, D. Friedman, M. Frishkopf, V. Haimo, P. Helinek, A. Khanna, K. Laube, J. Mayersohn, G. Marceline, J. Robinson, E. Rosen, K. Sirois, M. Vertenstein, A. Waldfogel, and J. Wiggins

BBN Communications Corporation

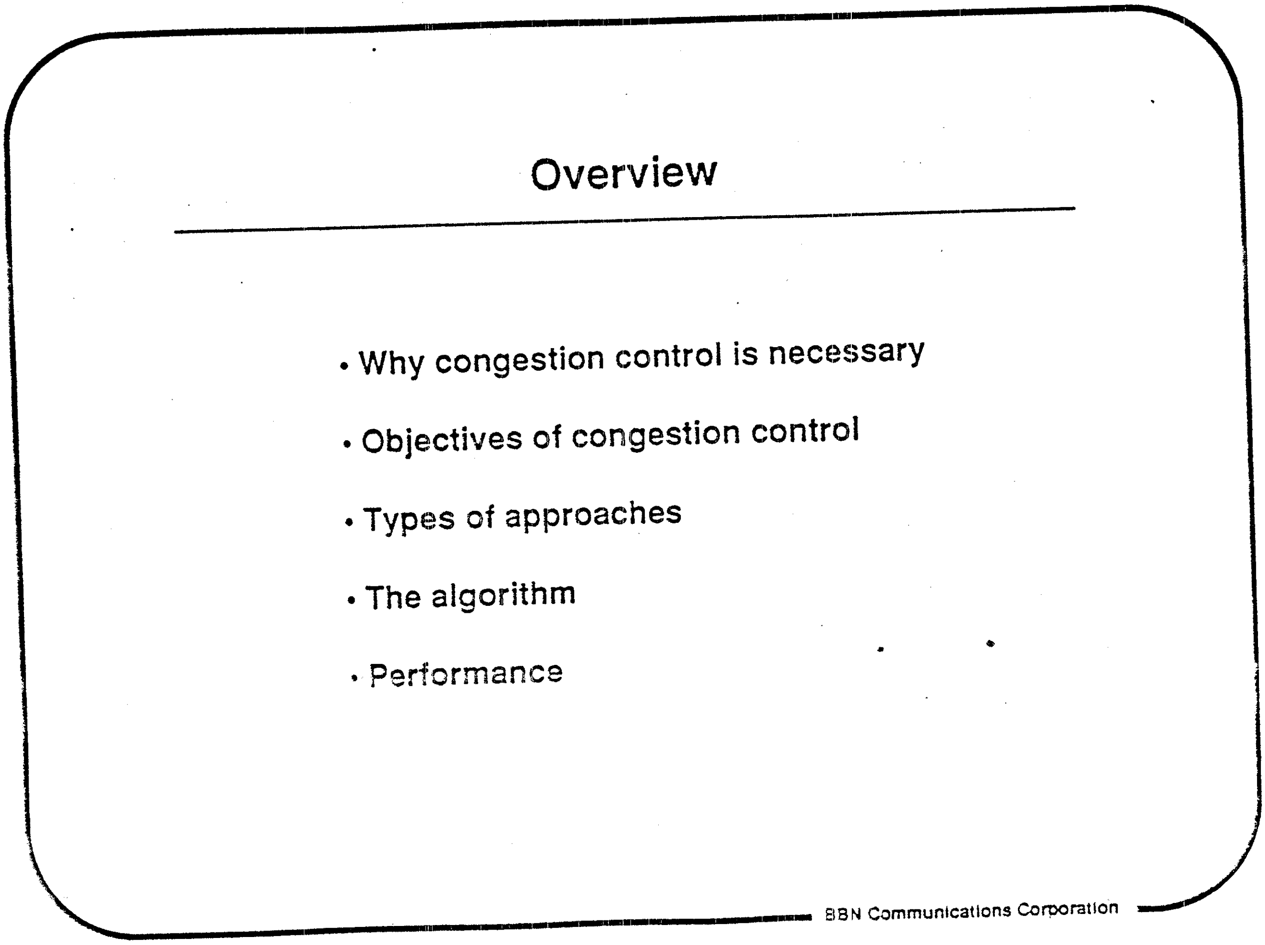

Overview

- Why congestion control is necessary
- Objectives of congestion control
- Types of approaches
- The algorithm
- Performance

BBN Communications Corporation

Congestion

- Increase in delay and reduction in throughput in response to higher offered load
- Aggravated by buffering and retransmissions

Why is it Necessary?

- Mismatched transmission speeds
- Convergent flows
- Static network design
- Component outages

Internet-Specific Issues

- Higher data rates
- More flows per resource
- Wider range of applications
- More packet switches
- Many different vendors
- No global administration

Objectives

- Minimize congestion
- Maximize throughput
- Fairness
- Stability
- Responsiveness
- Efficiency
- Synergy with routing
- Easy to implement

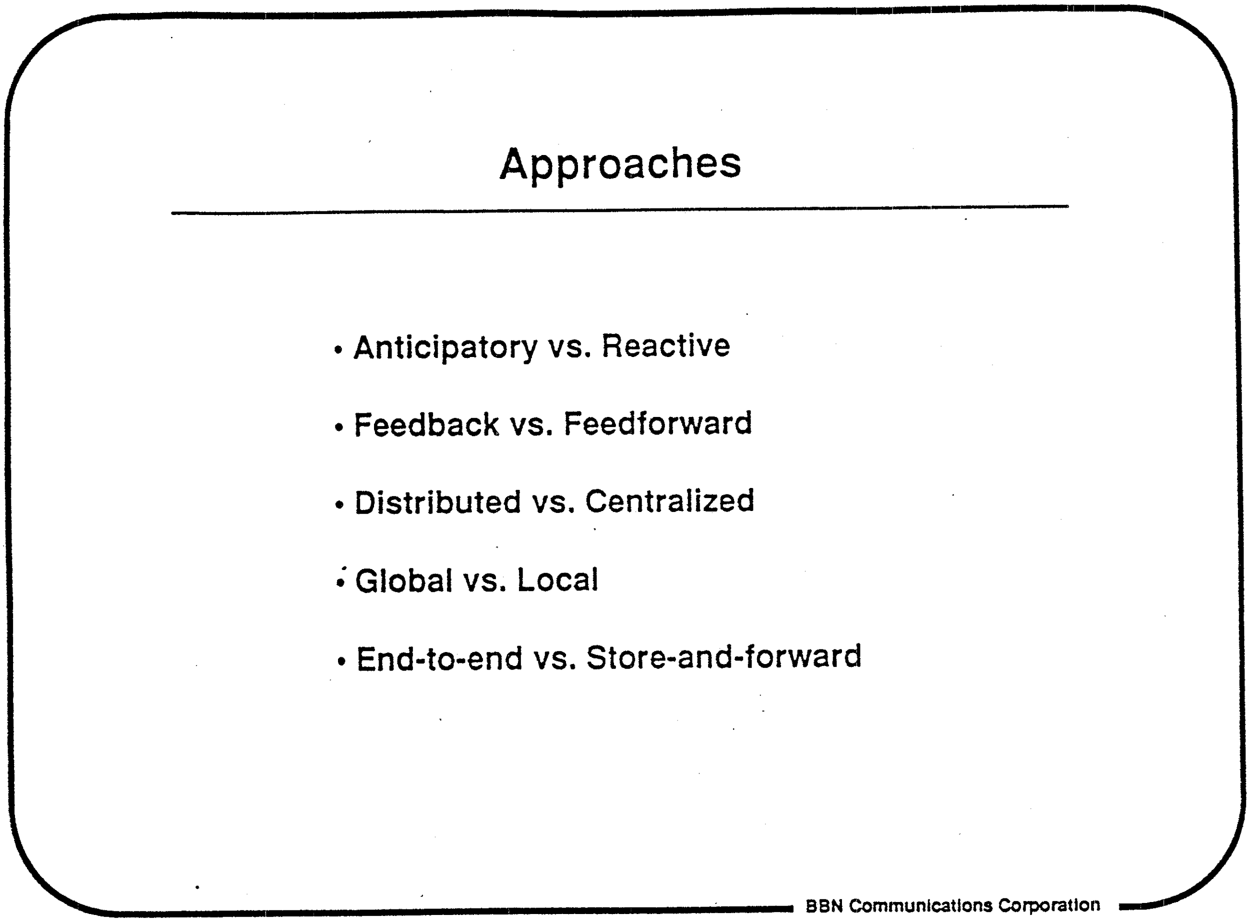

Approaches

- Anticipatory vs. Reactive
- Feedback vs. Feedforward
- Distributed vs. Centralized
- Global vs. Local
- End-to-end vs. Store-and-forward

BBN Communications Corporation

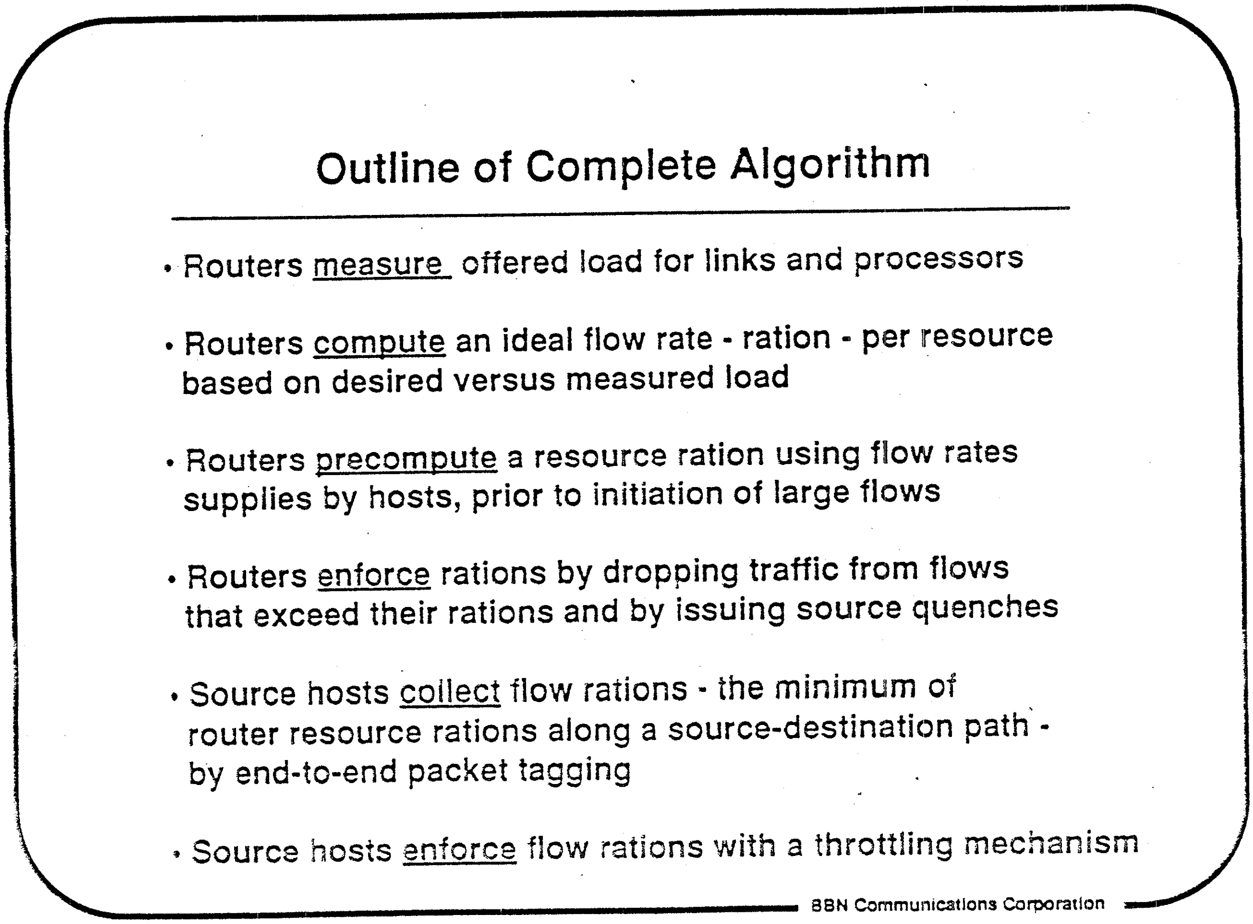

Outline of Complete Algorithm

- Routers measure offered load for links and processors
- Routers compute an ideal flow rate - ration - per resource based on desired versus measured load
- Routers precompute a resource ration using flow rates supplies by hosts, prior to initiation of large flows
- Routers enforce rations by dropping traffic from flows that exceed their rations and by issuing source quenches
- Source hosts collect flow rations - the minimum of router resource rations along a source-destination path - by end-to-end packet tagging
- Source hosts enforce flow rations with a throttling mechanism

BBN Communications Corporation

Resource Load

- Offered load measured for links and processors measured over 4-second interval
  Links in bytes
  Processors in busyness

- Recency-weighted average of measurements:

$$\text{load}(0) = \text{sample}(0)$$

$$\text{load}(t+1) = \text{load}(t) + a(\text{sample}(t+1) - \text{load}(t)),\ 0<a<1$$

Resource Rations

- Resource rations computed as:

$$\text{ration}(0) = \text{target}$$

$$\text{ration}(t+1) = \min\ \{\text{target},\ \text{target/load ration}(t)\}$$

- Large variance in packet size necessitates two separate rations:
  Links in bytes per second
  Processors in packets per second

- Must measure capacity of network links:
  M/M/1 number of customers in system
  bytes sent ÷ bytes sent/average queue length

Convergence

- Best Case:
  All flows greedy, one iteration

- Worst Case:
  All flows non-greedy, order log n iterations, where n is number of flows using the resource

Stability

- May overshoot because of buffer induced feedback delay

- Relies on routes remaining viable over several tens of seconds

Fairness

- Facilitated by <u>single</u> ration per resource
- At source-destination <u>flow</u> level

Smart Drop

- Hosts are <u>not required</u> to participate in congestion control
- Routers maintain <u>flow</u> database
  Routes
  Rations
  Flow use of resources
- Database accessed by <u>hashing</u> on source-destination addresses
- When flow exceeds ration, <u>drop</u> packet and issue a source quench

Feedforward Control I

```
n = number of flows using resource

d = target

ration = d/n

Repeat until all flows accounted for:

  If no flows are greedy with respect to ration, stop

  Other wise,

    d = d - the contributions of all nongreedy flows
    n = the number of greedy flows
    ration = d/n
```

Feedforward Control II

```
If flow setup

  If flow > or = ration

    ration = g/(g+1) ration
    g = g + 1

If flow teardown

  If flow > or = ration

    ration = g/(g-1) ration
    g = g - 1
```

Network Topology

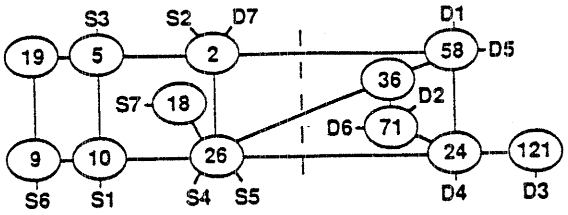

Usage of Cutlink

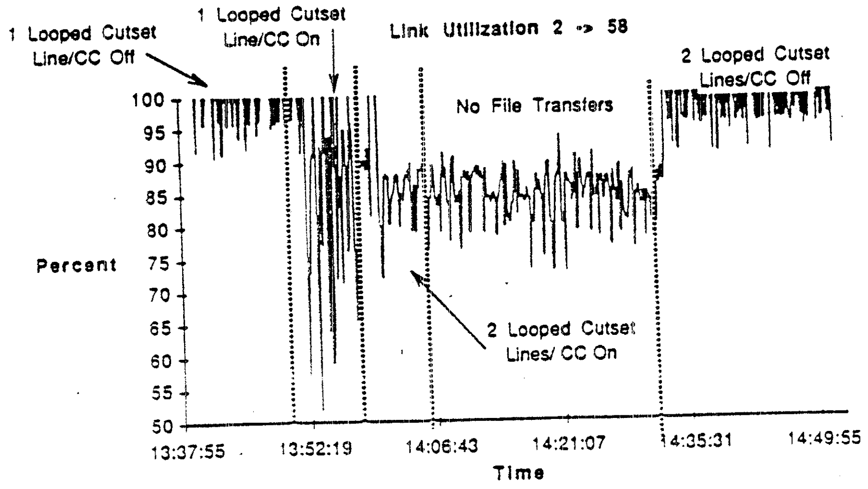
1 Looped Cutset Line/CC Off

1 Looped Cutset Line/CC On

Link Utilization 2 -> 58

2 Looped Cutset Lines/CC Off

No File Transfers

2 Looped Cutset Lines/ CC On

Rations on Cutlink

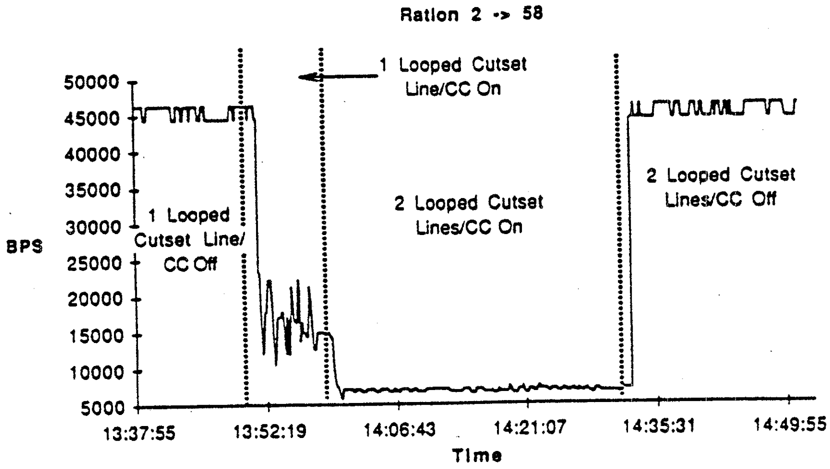

Usage of Feeder Link

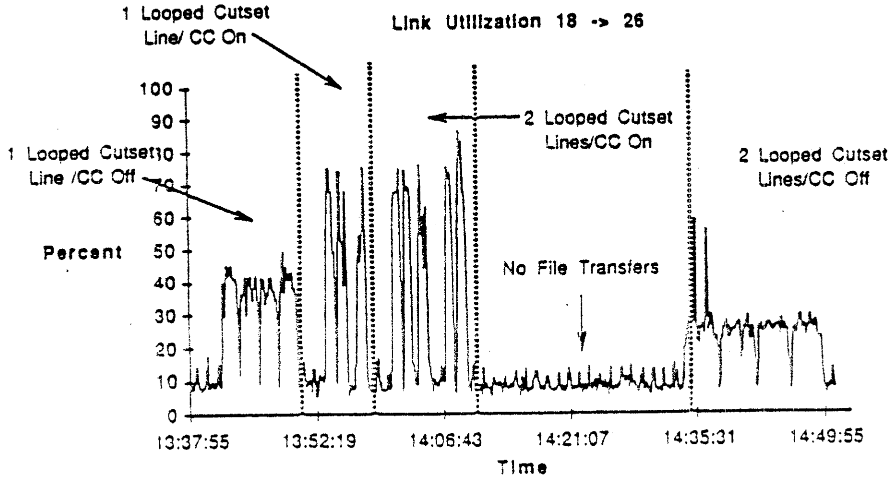

5.6 "NORDUNET"

Presentation by Mats Brunnell/NORDUNET

NORDUNET is a networking program in the Nordic countries (Denmark, Finland, Iceland, Norway and Sweden). The program is financed by the Nordic council of ministers, and runs from 1985 - 1991. The total budget is approximately 13,5 MSEK equal to 2 MUSD. The activities are focused at harmonizing networking in the R&D sector. As a result the NORDUnet network has been established.

NORDUNET Is the Inter-Nordic R&D which provides international network services to the National Nordic R&D networks, DENet in Denmark, FUNET in Finland, SURIS on Iceland, UNINETT in Norway and SUNET in Sweden. NORDUnet has made service agreements with other international networks like NSFnet, EUnet, EARN, HEPnet/SPAN.

NORDUNET is a multiprotocol "internet". The protocols supported are:

- Internet IP
- IBM/NJE - RSCS
- DECnet
- X.25

RIPE During the last year Europe has seen a growing number of IP networks with a coverage that goes beyond that of a typical LAN. Networks of regional, national and international importance have come into operation. Though these networks are operated under the separate executive authority of the various organizations that own them, a growing tendency has been observed to interconnect them on an ad hoc basis.

In order to facilitate the interconnections of separate IP networks, a coordination body has been formed recently by most of the organizations running IP services in Europe today. Under the name RIPE (Reseaux IP Europeens), a framework has been set up within which a growing number of IP network service providers will coordinate the inter-network aspects of their services.

370

NORDUNET and the NORDUnet – an Overview

Latest status on RIPE – the activities

by

Mats Brunell
NORDUNET/SICS

Mats.Brunell@sics.se

The presentation

NORDUNET & NORDUnet

- About the NORDUNET program
- Background to the NORDUnet project
- The NORDUnet backbone
  - technical
  - organisational
- Some conclusions
- About the future

Update on RIPE activites

- Status on European IP networking (yesterday...)
- RIPE Task Forces
- RIPE Current activities

The NORDUNET program

- funded by the Nordic Council of Ministers
- runs from 1986 - 1991
- total budget is appr. ~~6~~ 2 MUSD
- participating organisations in NORDUNET are:
  - Denmark/DENET
  - Finland/FUNET
  - Iceland/SURIS
  - Norway/UNINETT
  - Sweden/SUNET

Plus individuals!

The goals

- harmonised network services

  In cooperation with the national networks

- secondary goals:

  – knowledge in networking

  – establish good inter-nordic relations

  – establish international relations

The activities

- provision of services today:
  - The NORDUnet Backbone
  - Ensure Harmonised mailservices, e g application level gateways RFC987 style
- planning for a introduction of OSI based services
  - OSI-pilots and experiments

Working group and project activities

NORDUNET WG:s:

- Directories X.500 WG
- Message Handling Systems X.400 WG

NORDUNET projects:

- NORDUnet implementation
- Gateway, evaluation of Mailway RFC987 MHS/SMTP application level gateway software
- NORDUNET EAN/X.400 software distribution
  - "RFC-style addressing
  - Standard attribute style of addressing
- NORDUnet network layer addressing
  - X.25 '80 and '84, ISO addressing
  - ISO CONS and CLNS NSAP allocation scheme
- NORDIC Mail Harmonisation
  - Nordic mail Harmonisation addressing
  - RFC987 Gateway mapping
  - Service organisation
  - Service quality monitoring

External activties

RARE, European umbrella organisation for R&D networking harmonisation

- WG1, MHS, Ex Chairman: Alf Hansen ELAB-RUNIT
- WG2, FTAM, Participation: Einar Løvdal UiO
- WG3, Directories, Participation: Juha Heinänen and national
- WG4, "X.25+" National representation
- WG6, National representation
- WG8, Application Management, Chairman: Mats Brunell
- RARE/COSINE IXI project participation

Other activities

- RIPE, European IP coordination group
- HEPnet Technical Committee
- CCIRN, Coordination Committee for Intercontinental Research Networks
- IANW conferences, national and NORDUNET

Organisation

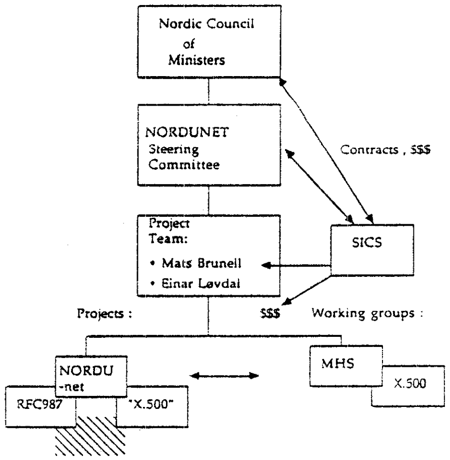

Background to the NORDUnet project

- IBM funding for EARN lines ended 1987
- Cost effectiveness needed for more than NJE/RSCS type of traffic
- New protocol types of interest:
  – DARPA IP
  – DECnet
  – ISO OSI - X.25 based services
- Cost sharing for international connections to the US and Major European Networks

The NORDUnet project

- Implementation NORDUNET coordinated project
- Project initiated in September 1987
- Project plan agreed in NORDUNET in Jan -88 Sent to SUNET, UNINETT, FUNET,UNI-C and SURIS for comments and decisions
- Decision to start ready in May -88
- Project start May -88
- Contract signed for long term responsibility for funding etc, by the National networks in June -89
- Official opening in October -89

NORDUnet - A TRIBUTE TO NORDIC COOPERATION!

The NORDUnet plug (or socket)

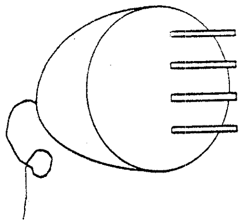

DARPA IP

DECnet

X.25 1980 & 1984

EARN

Basic transport level

Uses a mix of Vitalink TranLan III CISCO and X.25, both private and public

Availability of lines (down time statistics)

	May-89	*August-89*
SE - NO	*96,6%*	*97,4%*
SE - FI	*97,3%*	*97,6%*
SE - DK	*96,0%*	*86,0%*
SE - NL (CWI)	*99,1%*	*99,5%*
SE - US (JvNC)	*96,1%*	*94,0%*

Only KTH-CWI has acceptable availability (<1% down time)

Transmission delay across NORDUnet

Inter Nordic:	ca 30 - 45 ms
Towards EUnet (CWI):	ca 60 ms
Towards NSFnet (JvNC):	ca 450 ms (Satteilite)

Optimal IP performance - DECnet performs less efficiently

NORDUnet Configuration

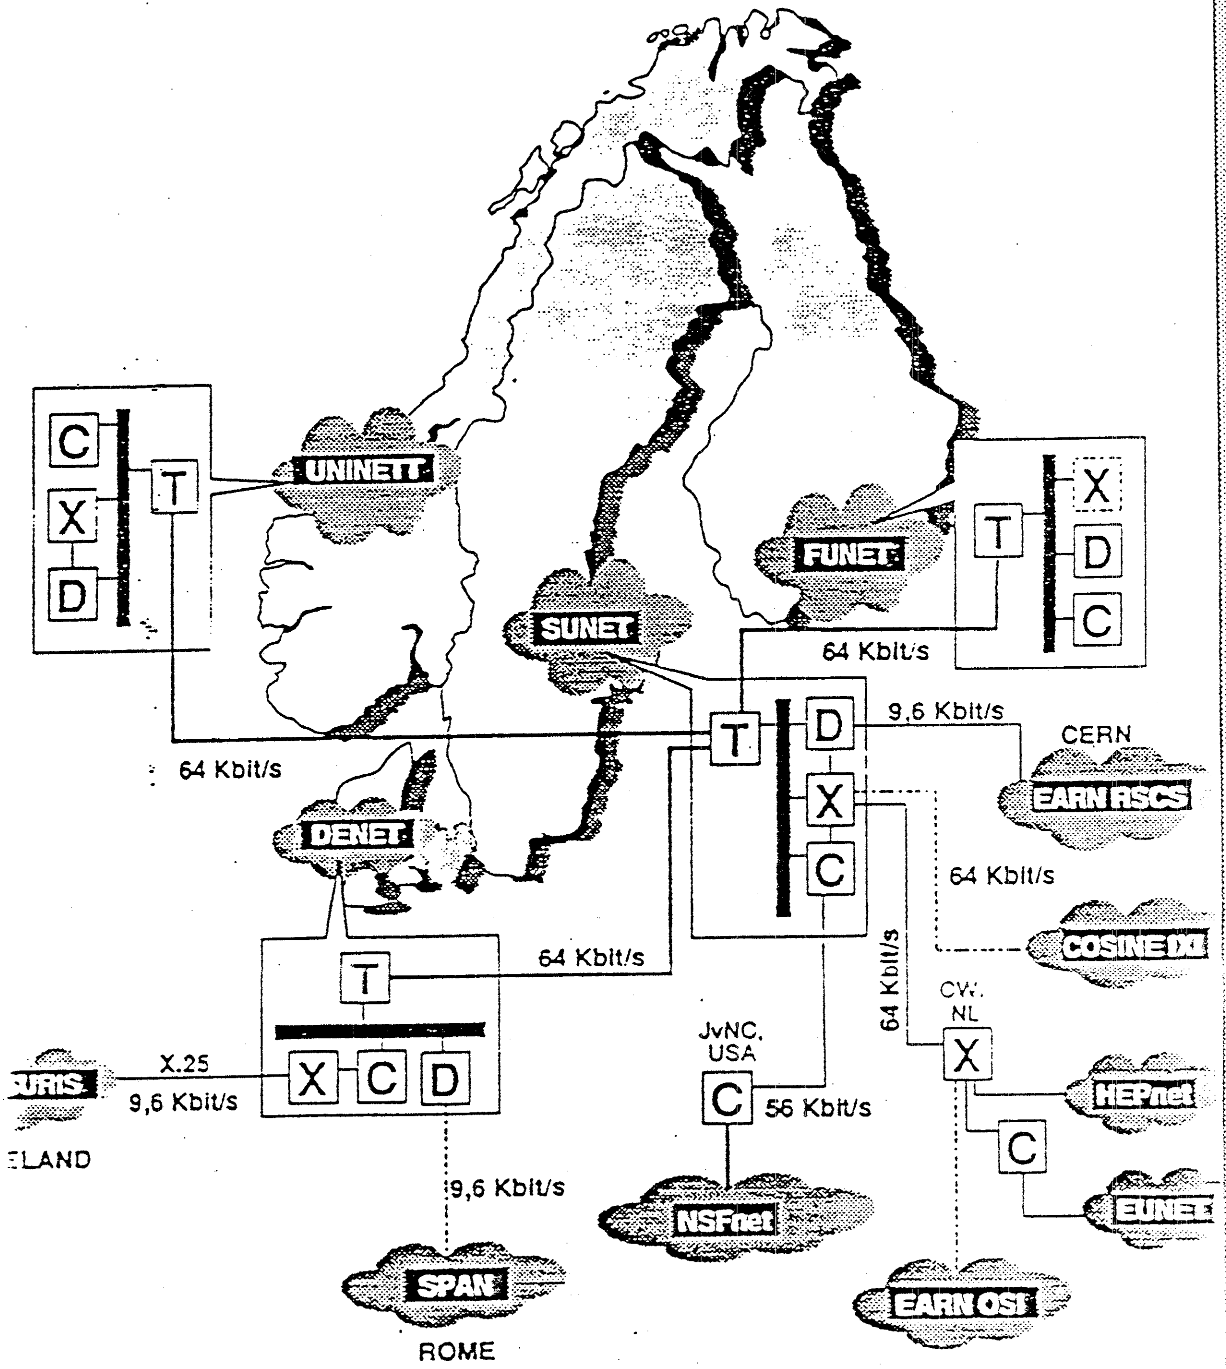

NORDUnet IP sub-network

- more than 20 networks
- One adminstrative routing domain
  - primary route via JvNC
  - secondary via CWI/SURANet (backup link) uses EGP to JvNC/NSS
  - announcement of "our nets" at relevant gateways using "administrative distance"
- Each contry runs own DNS with secondaries in US and at CWI
- Fake rootserver is run at KTH (permission by SRI-NIC) (backup at Lund University, Sweden)
- initially 2.500 hosts (after 6 months 5.000 and now after one year 6.500 hosts
  - all major universities
- Internet/NSFnet connected via JvNC
- EUnet connected via CWI, Amsterdam
- HEPnet/IP (CERN) via CWI/NIKHEF

Uses CISCO routers

SNMP for management purposes

NORDUnet DECnet sub-network

- organised according to HEPnet/SPAN addressing scheme:
  - national DECnets use DECnet areas 47-62
  - NORDUnet routers plus HEP and SPAN nodes on area 21

  implying:

  - all DECnet services available between Nordic nodes
  - all DECnet services available between Nordic area 21 nodes and non-Nordic SPAN and HEP nodes
  - only DECnet file copy and mail available via socalled "Poor Mans Routing" between Nordic ordinary nodes and SPAN/HEPnet
- Non standard DECnet router technology used on NORDUnet
  - some performance problems
- 3000 - 4000 DECnet nodes (including PC's using PCSA)
- Coordination of node names planned

Uses uVAX 3600 "EARN gboxes "as routers

NORDUnet X.25

- Stable NORDUnet X.25 access point established at KTH
- Is connected to HEPnet/EUnet X.25 sub-net, EARN X.25 sub-net
- Addressing recommendation (both X.121 and ISO NSAP scheme) worked out
- important experience gained running X.25 over Ethernet

National entrypoints only!

Will give access to COSINE IXI pilot

Uses Satelcomms switches

NORDUnet X.25 cont'd

But:

- X.25 over NORDUnet Ether is still a pilot, not a service

X.25 Conclusions:

- Running X.25 PLP using LLC1 is not recommendable, LLC2 required!!
- Current bridge technology is not very well adapted to running connection oriented link and network layer services
- Extending X.25 access to NORDUnet and European X.25 networks throughout the national Nordic networks still remains to be solved

NORDUnet EARN

- Uses a mix of solutions today!

  - RSCS/BSC
  - RSCS/DECnet
  - RSCS/NJE/OSI
  - RSCS/IP VMNET

Evaluation will determine NORDUnet solution

Decision not yet taken in EARN on the OSI transition program!

- about 50 RSCS nodes

NORDUnet organisation

- Contract between the participating organisations *No legal body yet!*
- Policymaking body is a *separate* body, the NORDUnet Board repr from the owners/contractors
- Contracts to operational sites and coordinators

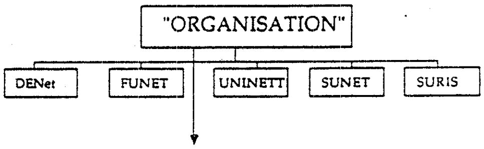

Contracts:

- KTH: IP & Basic LAN coord
- UNI•C: DECnet coord
- RUNIT-D: X.25
- HUT: EARN
- JvNC
- STS
- Service agreements et c

NORDUnet Costs overview, approximate

Initial costs:	800 kUSD
NORDUNET contribution :	500 kUSD
Operational costs:	
Estimated 1990:	725 kUSD
• Lines	350 kUSD
• Personell, travel et c	320 kUSD
• Maintenace	15 kUSD
• Equipment	40 kUSD

External contributions trough contracts:

- 40 - 120 kUSD (NSF funding for US connection pending)

Experiences, contract and decisionmaking

- Formal/legal organisation needed to sign certain contracts!
- Long term commitments needed – but hard to get!
- Commit each party to pay if he drops out!!

One plan, containing

- Technical solution for indentified services
- Implementation costs timeplan etc
- *Operational cost estimate*
- *Organisational solution drafted*

376

Some Conclusions

- Stable and well performing NORDUnet IP service established
- Stable DECnet service
- X.25 over the NORDUnet backbone still at a pilot stage
- EARN - solution for EARN operational service between the Nordic countries and towards EARN/Europe still to be decided
- We have a backbone which can live for about 3-5 years
- It can utilize 2 Mb speeds from end-LAN to end-LAN today
- It can carry existing and ISO protocols

NORDUnet is now a reality - at the service of theNordic Research Community

Short term activities

- COSINE /IXI
- EMBnet consultancy
- IBM EASI discussions
- Directory projects
- ISO IP/DECnet Phase V project preparation

Longer term future

- High-speed networking 100 Mbps +
- "New protocol standardisation methods" needed!?
- Secure long term funding (beyond this technical solution)

Key questions?

How well can Political intitives and pragmatic networking cooperate?

How can we solve the hard problems which lies ahead of us?

RIPE Status

Status on European IP networking (yesterday...)

- About 13.000 nodes and 95 organisations
- Support from most major R&D networks in Europe
- Will be in "partnership" with RARE and thus bring in the "last countries" to RIPE

Contacts:

RIPE-request@nic.EU.net

Main Intern. IP L.Lines

(incl. ordered lines)

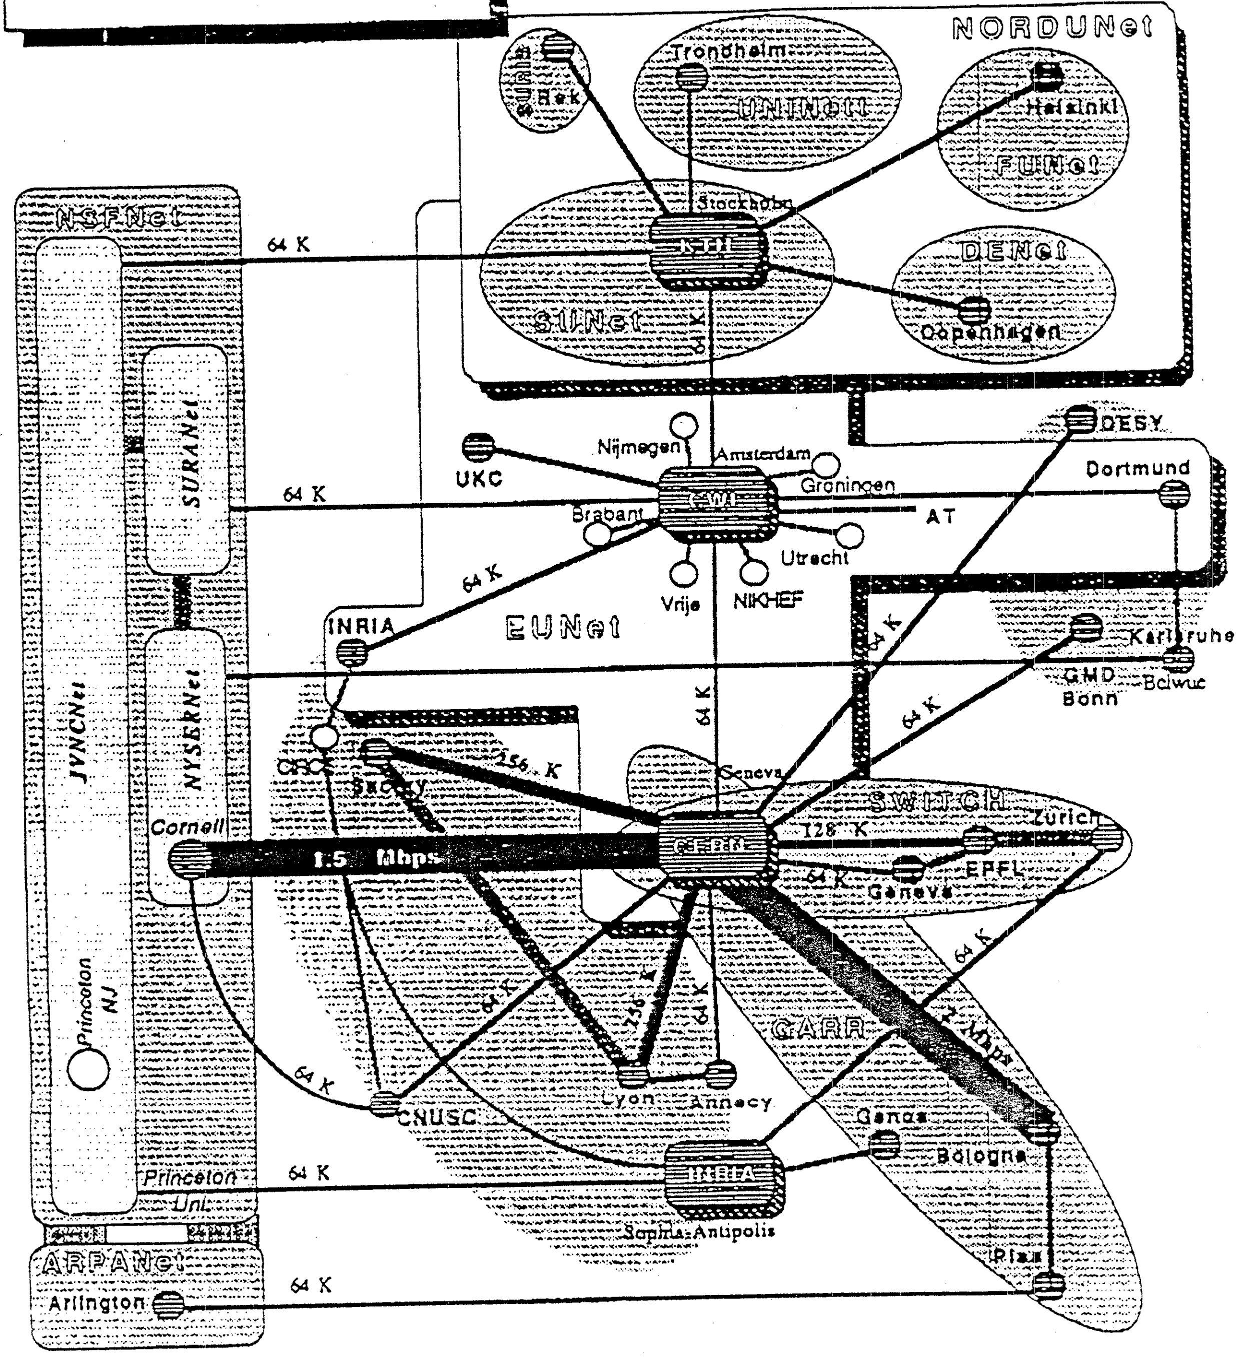

▬	International Leased Line
—	National connection (not all are represented)
○	National IP Net.
○	International IP Net.
▤	Multi-International IP Node

CERN-DD/CSMap05
January 1989

(Prepared for CERN internal use. Available to outside from F.G.de Bilio CERN-DD/CS)

RIPE Task Forces

- Task Force 1 "Connectivity and Routing"
  Headed by Thomas Lenggenhager SWITCH
- Task Force 2 " Network Management and Operations"
  Headed by Daniel Karrenberg EUnet/CWI
- Task Force 3 "Domain Name System"
  Headed by Francis Dupont Inria
- Task Force 4 "Formal coordination"
  Headed by Rob Blokzijl NIKHEF

RIPE Current activities

- "Setting it all up"
  - assisting with info et c national and local people on IP networking
- Setting upp DB-servers for info "whois" et c
  - CWI: "nic.EU.net"
  - Inria
  - KTH
- Inventory of IP nets and links

Next Phase:

- European root server!
- Proposal on topology and routing
- Set of bi- and multilateral agreements

5.7 "Report of the Open Software Foundation"

Presentation by Brad Johnson/OSF

The Open Software Foundation is currently conducting an in-depth technical evaluation of distributed computing technologies that were submitted to OSF in response to a Request For Technology (RFT). Technologies being evaluated include Remote Procedure Call and Presentation Services, Naming/Directory Services, Authentication, and Distributed File Services.

In February 1990, the IETF begins a review and evaluation of host services technologies to be considered for standardization. OSF believes that both the Internet community and OSF could benefit greatly by coordinating these parallel evaluation efforts.

This presentation introduces OSF, the RFT process, and specifically addresses the Distributed Computing Environment RFT. We will describe the goals of the project and its relationship to the industry in general and the potential role of Internet community in the process. Finally, We will describe the technical evaluation process, the key findings of the evaluation to date, and outline the plans for announcement and release of the selected technologies.

Report of the
Open Software Foundation

Distributed Computing Environment
Request For Technology

Brad C. Johnson

How the hell did I get here?

- David Clark
- Paul Mockapteris
- Craig Partridge
- David Crocker

Presentation

What is OSF?

Update on the DCE RFT

How does the RFT process work?

What is OSF?

Key organizations:

- Research Institute
- Operations
  - human resources
  - finance
- Development
  - OS
  - Motif
  - DCE
  - ANDF

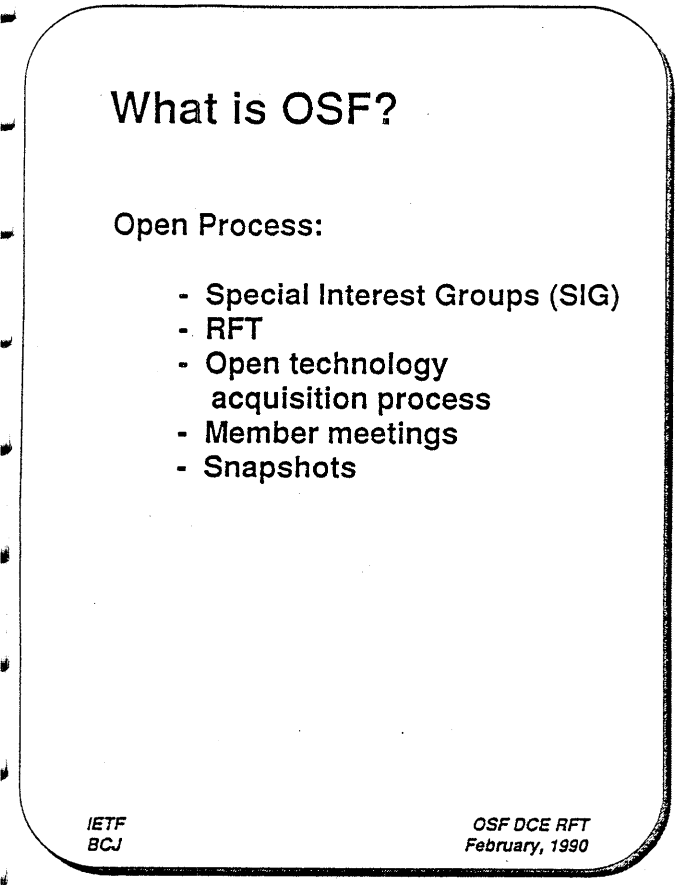

What is OSF?

Open Process:

- Special Interest Groups (SIG)
- RFT
- Open technology acquisition process
- Member meetings
- Snapshots

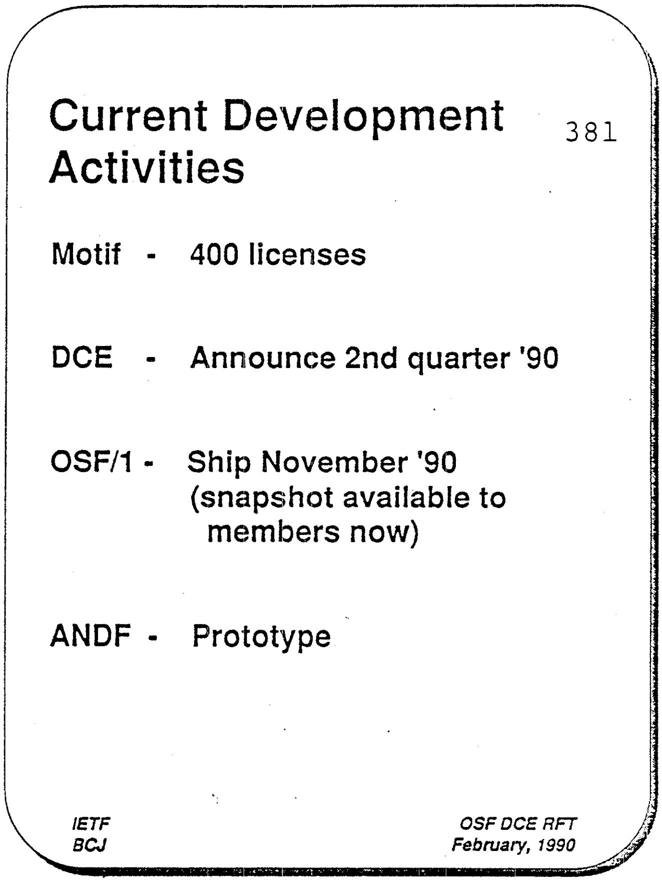

Current Development Activities

Motif -	400 licenses
DCE -	Announce 2nd quarter '90
OSF/1 -	Ship November '90 (snapshot available to members now)
ANDF -	Prototype

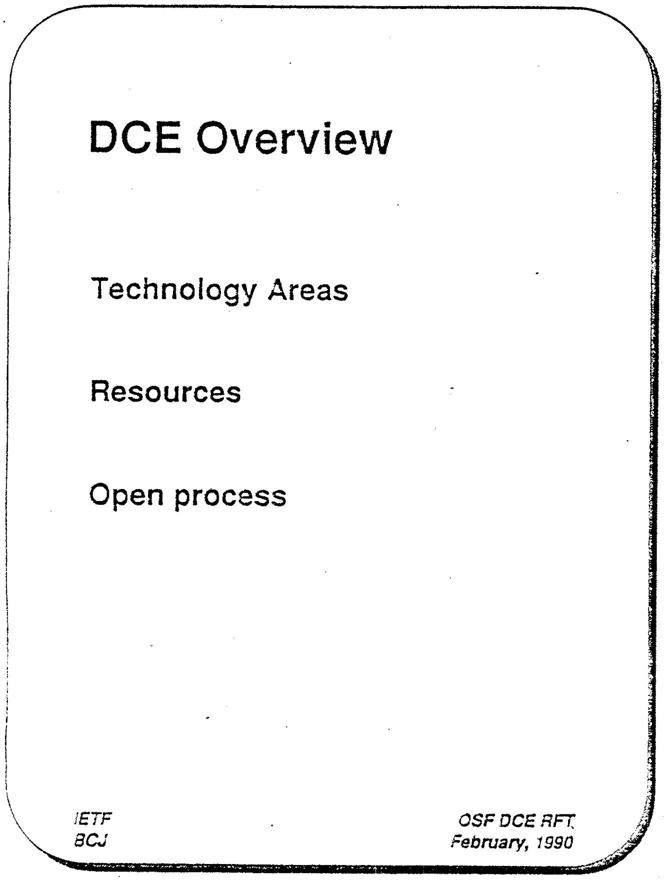

DCE Overview

Technology Areas

Resources

Open process

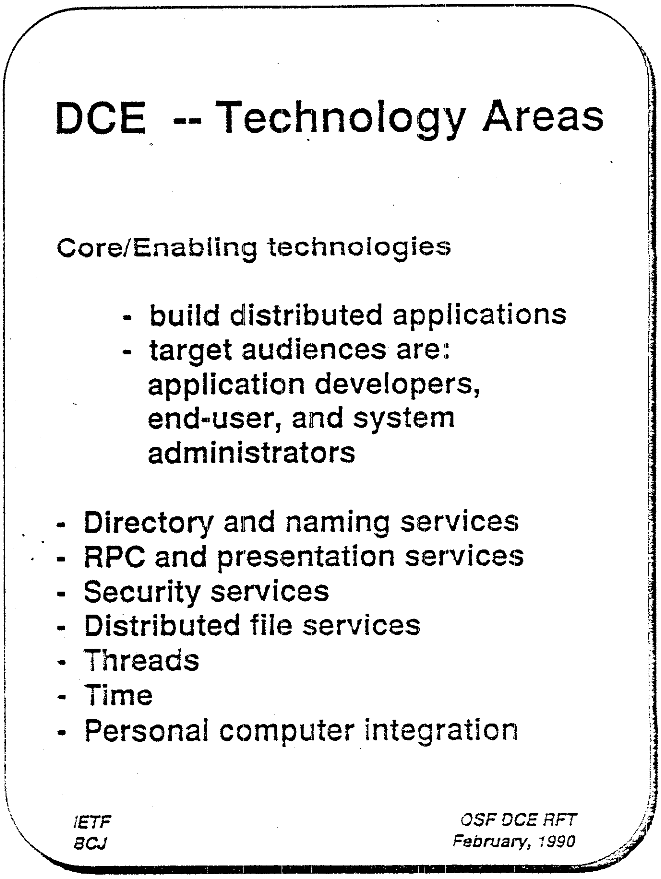

DCE -- Technology Areas

Core/Enabling technologies

- build distributed applications
- target audiences are: application developers, end-user, and system administrators

- Directory and naming services
- RPC and presentation services
- Security services
- Distributed file services
- Threads
- Time
- Personal computer integration

DCE Architecture

Properties of distributed environments

- physical separation
- administrative autonomy
- heterogeneity
- scalability
- extensibility

Expectations of distributed environments

- resource sharing
- availability
- performance
- cost effective

DCE Architecture, cont.

Technology areas beyond the scope of this RFT

- distributed applications
- distributed databases
- distributed development tools
- distributed operating systems
- event notification
- interprocess communication
- mainframe integration
- message handling
- object oriented environments
- spooling technologies
- system management
- terminal oriented technologies

DCE Architecture, cont.

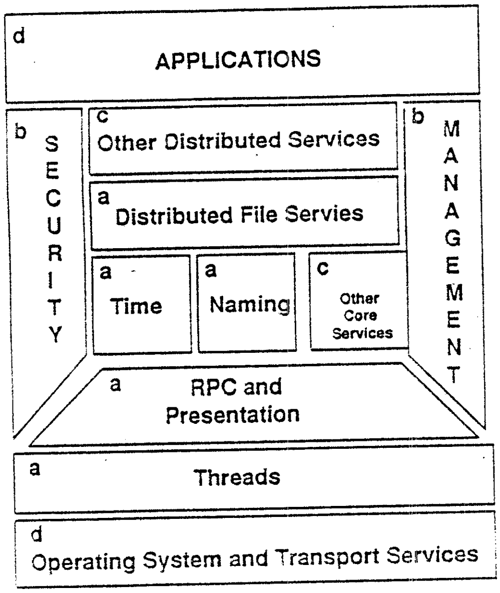

d APPLICATIONS

b SECURITY

c Other Distributed Services

a Distributed File Servies

a Time

a Naming

c Other Core Services

b MANAGEMENT

a RPC and Presentation

a Threads

d Operating System and Transport Services

DCE -- Resources

Members and
Submitters (members)

- member meetings
- open door policy
- surveys
- lab period
- references

SIG

- as needed basis
- concentration of technical expertise

DCE -- Resources, cont.

Consultants:

- Andrew Birrell, DEC
- Heinz Juergen Burkhardt, GMD
- David Cheriton, Stanford
- Paul Mockapetris, USC
- Sape Mullender, CWI
- Roger Needham, U. Cambridge
- William Pigott, DHL
- Russel Sandberg, Legato
- Peter Schay, Gartner Group
- Walter Ulrich, ADL
- Peter Weinberger, AT&T

DCE -- Resources, cont.

StAnDaRdS

- ANSI
- CCITT
- Internet
- IEEE
- NIST
- X/Open
- ISO

Two-way communication

DCE -- Resources, cont.

OSF evaluation team

- 9 technical contributors
- 3 management and operations

Other OSF expertise

- OS overlap
- business directions
- research directions

How does the RFT process work?

Solicit proposals

Solicit technology

Solicit feedback

Offer software

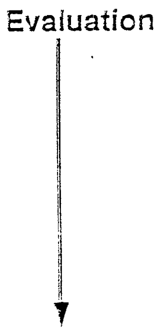

Evaluation

DCE -- Open Process

Solicit Proposals

- OSF survey: Fall '88
- SIG: January, March, April, May '89
- RFT: June '89

Solicit Technology

- Letters of intent: July '89
- SIG: August '89
- Refinement: September '89
- Submissions: October '89

DCE -- Open Process cont.

Solicit Feedback

- November member meeting
  - submission presentations
  - demonstrations
  - panel discussions
  - consultant discussions
  - surveys
- January member meeting
  - framework presentations
  - working group discussions
  - consultant discussions
  - participation of key standards groups

On-going process of communication

DCE -- Open Process, cont.

Evaluate Technologies

- Letters of intent
- Full submission review
- November meeting Q&A
- DCE framework/architecture
- January meeting
- Lab period

DCE -- Open Process, cont.

Evaluate Technologies cont.

Vertical characteristics

- Framework document
- Evaluation criteria
- Key issues
  - technology
  - submission

Horizontal characteristics

- standards
- ease-of-use
- heterogeneity
- scalability
- performance
- internationalization
- serviceability
- portability

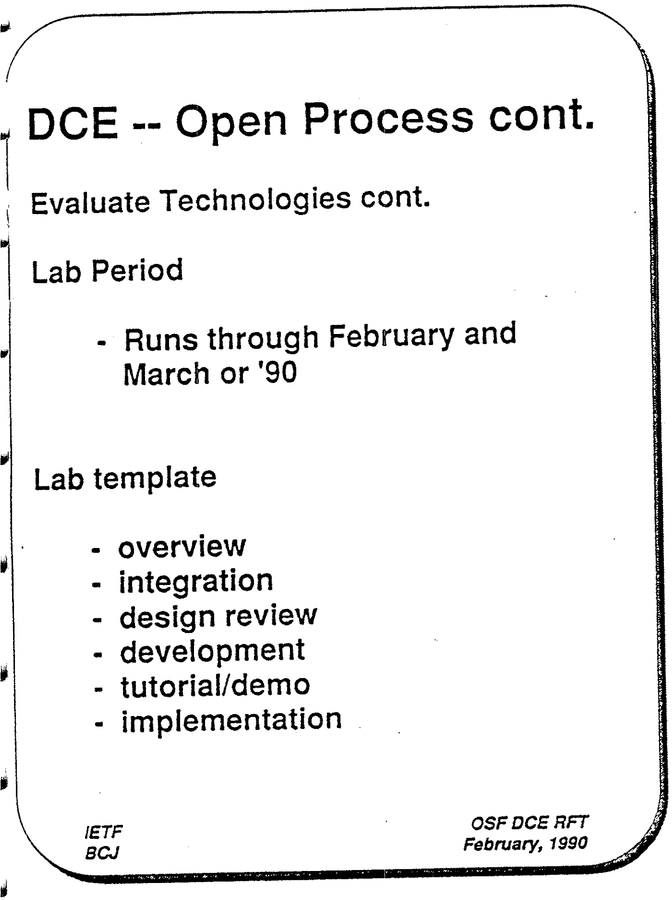

DCE -- Open Process cont.

Evaluate Technologies cont.

Lab Period

- Runs through February and March or '90

Lab template

- overview
- integration
- design review
- development
- tutorial/demo
- implementation

IETF
BCJ

OSF DCE RFT
February, 1990

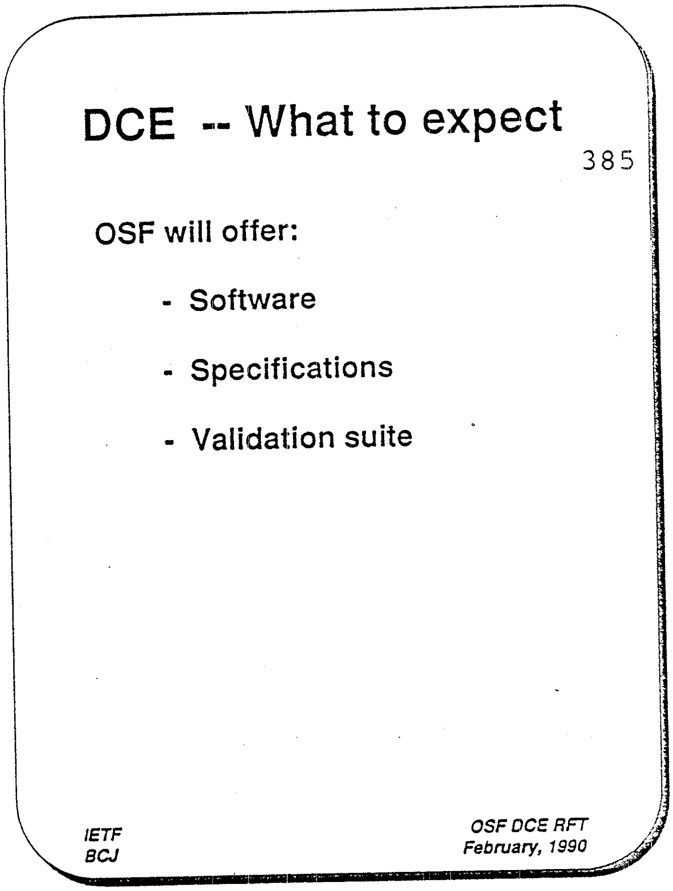

DCE -- What to expect

OSF will offer:

- Software
- Specifications
- Validation suite

IETF
BCJ

OSF DCE RFT
February, 1990

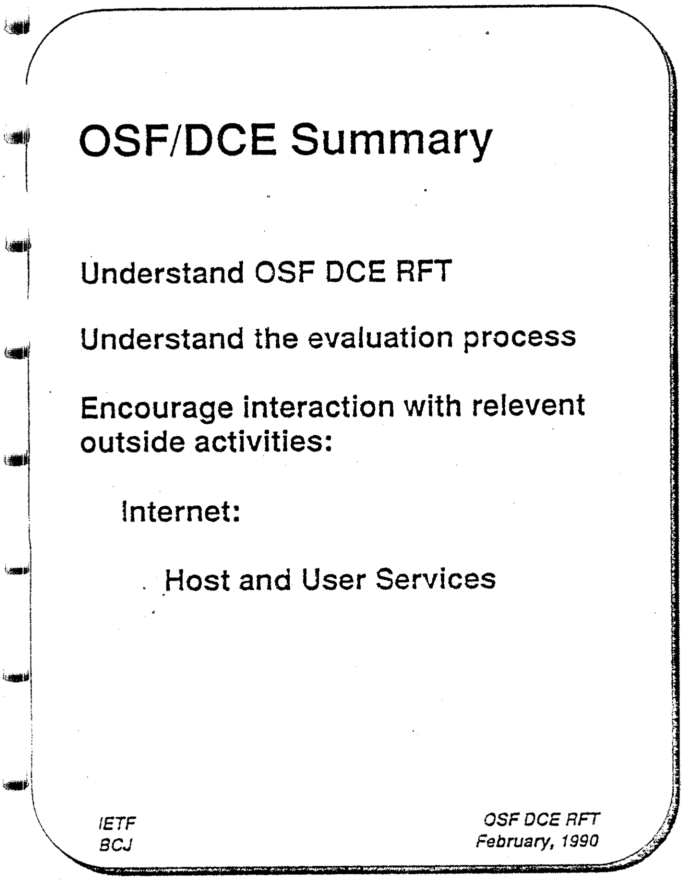

OSF/DCE Summary

Understand OSF DCE RFT

Understand the evaluation process

Encourage interaction with relevent outside activities:

Internet:

Host and User Services

IETF
BCJ

OSF DCE RFT
February, 1990

5.8 "The INTEROP 89 Network: Design, Problems, and Lessons Learned"

Presentation by Phillip Almquist/Stanford

SYNOPSIS:

This talk described our experiences designing and operating a large multi-vendor demonstration network at INTEROP 89. The most important purpose of the talk was to allow the Internet community to learn from those experiences. Because the Internet policy makers have stated their intent to concurrently support OSI and TCP/IP in the reasonably near future, the talk particularly emphasized what we learned about the practical realities of building and operating a non-trivial OSI (CLNS) network; I figured that I was hardly the last TCP/IP network designer who would be faced with having to build an OSI network.

The talk began with a brief description of the network, primarily for the benefit of those who had not been to INTEROP 89. The network included approximately 600 hosts and routers interconnected with nearly 6 miles of cabling. There were two separate T-1 connections to the Internet, including our own NSFNet node. A microwave link provided Internet access from the hotel which most of the attendees were staying at (an idea we should consider for future IETF meetings). We also had a working multivendor FDDI ring, tons of SNMP, a 4.4BSD machine, the first (?) public demonstration of PPP, and lots of other interesting stuff. Maps of the network are included in the slides that follow this summary.

Basic Network Design

The meat of the talk consisted of an extended description of the network's design and the rationale behind the choices we made. In cases where those choices did not work well in practice, I attempted to include whatever wisdom hindsight seemed to offer.

The environment in which the network would be built and operated dictated a very conservative design. Few enough days were allowed for installation that the network pretty much had to work as built, without extensive debugging. Because we had little effective control over what was connected to the network or how it was configured, the design had to assume that hosts could not be trusted to be well-behaved. And although the network only had to last for three days, the design had to ensure that if any failures did occur, they could be isolated and repaired with the utmost rapidity.

We chose a topology that consisted of a large number of small subnets. Each subnet was connected to one of several backbone routers, and the backbone routers were in turn connected together by a short backbone Ethernet (see the diagram in the slides). This resulted in a tree topology. Vendors could either connect their hosts directly to our subnets or, if they preferred, could build their own subnets and connect them to one of ours with a router. The tree topology was chosen because it is very simple to debug, and precludes the sorts of finger-pointing that could have resulted had we allowed back-door paths that we didn't control. In general, the topology was designed based on what we knew worked well for IP networks, since we didn't have a good feel for what would work well for OSI networks. However, what we did seemed to work reasonably well for OSI.

TCP/IP

The part of the network design which was specific to supporting the TCP/IP protocols is discussed at length in the slides, but was skimmed over only very briefly during the talk because it was pretty much just a scaled up version of what we did at INTEROP 88 (which I talked about at the Ann Arbor IETF). However, I will briefly mention several things:

> We had routers from several vendors, and the only IGP that they had in common was RIP. Because RIP doesn't have any mechanisms for authentication, a misconfigured router or host running routed can easily create black holes. Also, we had had problems at INTEROP 88 with smaller systems pausing noticeably every 30 seconds because the routes from NSFNET made the RIP updates very large. As a result, we were forced to use static routing except on the short Ethernet that interconnected the backbone routers. Hopefully, by time you read this, there will be an Internet standard IGP which will be usable in networks such as INTEROP's.
>
> Most of the problems we observed were due to basic errors in host configurations, such as failing to set the subnet mask or the default gateway correctly. I take this to be a sign of the maturity of TCP/IP: most vendors feel sufficiently confident about their products that the engineers are delegating the system configuration tasks to the marketing staff.
>
> In the absence of protocol police, no amount of rules will give you a "clean" network. Fortunately, a well-designed IP network will survive most sorts of abuse by hosts.

In general, The TCP/IP aspects of the network went quite smoothly, in part because we had learned from our mistakes during INTEROP 88 and in part because The NSFNET people had considerably improved the reliability of inter-AS routing in the Internet.

OSI

Needless to say, supporting OSI was much more of an adventure than supporting TCP/IP, because we didn't have much practical experience to guide our decisions. Fortunately it worked pretty well anyway.

The first major decision we had to make was how much of the network would support OSI. Because of our inexperience and because none of the vendors of our backbone routers had released their OSI support, we decided that the better part of valor was to turn on the OSI code in only one brand of backbone router.

The second major decision we faced was the choice of an address format for network layer (NSAP) addresses. Because we would be connecting to other OSI networks, we needed to choose a format owned by some group that would be willing to assign us an address range. We wanted a format that would simplify routing as much as possible. And although in theory hosts and routers can handle completely arbitrary address formats, we weren't sure we believed that practice followed theory. The OSINET format seemed the safest choice because the network would be connected to OSINET and because we knew many of the exhibitors had tested their implementations on OSINET, using OSINET addresses. The OSINET format (pictured in the slides) also allowed us two bytes of local routing information, which was quite adequate for our purposes.

In hindsight, the only major problem we had with this format was one that we would have had with any widely used format: hosts which recognize the address format sometimes attempt to do clever things. At least one vendor's hosts came up with incorrect addresses because they though that they knew how to guess them based on the contents of IS Hello packets. Conversations I've had since then suggest that this unfortunate practice is expected to be widely implemented, so OSI network designers are well advised to be aware of these heuristics when assigning addresses.

In the future, the choice of format will become much simpler, since anyone who wants to use the IS-IS routing protocol will have to use the GOSIP format. However, none of our vendors implemented any routing protocol for OSI, so we had to use static routing. The implementations did support hierarchical static routing, which was crucial since we didn't want to maintain a static route to every host in every router.

Through the clever use of the two bytes of suborganization id in the OSINET format, we kept the number of static routes to under half a dozen per router. We used the second byte for the cable number (i.e. the IP subnet number), and the first byte for a number identifying which of the backbone routers was closest to that cable. A default route was used to get to the OSINET.

The routing worked pretty well, except for two minor problems. Although we were able to keep the static routes to a manageable number, typing them in was error-prone because of the length of OSI addresses. We also had routing loops when packets were sent to non-existent hosts, for the following interesting reason: in IP, addressing is tightly coupled to topology, so a router can easily tell whether or not it can send a packet directly to the destination by simply examining the destination address. In OSI, this tight coupling between topology and addressing is not required to exist; a router looks in the equivalent of its ARP table to determine if the destination address is directly reachable. Thus, if a host didn't exist, the final hop router would forward the packet to its default gateway, which would send it back, etc. The router vendor has since solved this problem by allowing static routes which cause matching packets to be discarded.

The problem for which we could come up with no good solution was name service. This is supposed to be handled by X.500, and in fact Marshall Rose's White Pages was demonstrated at the show. However, OSI applications in the real world today don't use X.500 to look up protocol addresses. And because X.500 is the ultimate answer, there is no OSI-standard (or even de facto standard) OSI host table format. To make matters even worse, in the OSI world you have a lot more addresses you need to keep track of, since you need a separate address for each host/application pair, rather than just for each host like we're used to in the TCP/IP world. Our solution, such as it was, was to invent our own host table format and make it available via anonymous FTP (not anonymous FTAM).

One of my goals for the network design was to provide the same level of support for the OSI protocols that we did for the TCP/IP protocols. I didn't entirely succeed, and didn't expect to, but it was an interesting to see what was possible. Most of the network could forward CLNP datagrams, and (modulo the learning curve) the routing wasn't much more painful than the (mostly static) routing we had to do for IP. We got our CLNP network connected to two national networks (the NSFNET test network and OSINET). ES-IS (the OSI ARP/ICMP redirect equivalent) worked well, except that one vendor didn't implement it. There were no real OSI network meltdowns. The only real hole was name service.

What We Learned

In this section I spent a couple of minutes trying to cull out some less technical lessons. Despite the fact that this was our second try at doing an INTEROP net, managing such an endeavor was hard. The economics of trade shows force a schedule for constructing the network that doesn't allow much time for things to go wrong. We did our best to coordinate all the people and plan all the details, but in the end it worked because the (volunteer) crew was good enough that they could solve the problems that we didn't plan for.

It was gratifying to note that we really had learned from our experiences the previous year, resulting in improvements ranging from cabling practices to booting of routers to coordination with MERIT and the NIC. It was less gratifying to note that we still hadn't learned to do some other things well, including network operation and management.

Credits

The talk concluded with credits for those responsible for the success of the INTEROP 89 network. I am very grateful to Peter de Vries, who helped with the network's design and managed its construction; and to Rob Hagens, Dave Katz, Jim Forster, and Marshall Rose, who share most of the responsibility for both the success of the OSI component of the network and for whatever understanding of OSI I have managed to achieve. The network was built and operated by Karl Auerbach, Dave Bridgeham, Eric Brunner, Jeff Burgan, Mario Castro, Shelley de Vries, Stev Knowles, Steve Larbig, Lisa Robertson, and John Romkey.

392

The INTEROP 89 Network: Design, Problems, and Lessons Learned

Philip Almquist

almquist@Jessica.Stanford.EDU
214 Cole Street, Suite #2, San Francisco, CA 94117-1916

I. Introduction

The INTEROP 89 Network

- 2-3 times the size of INTEROP 88 network
  - Nearly 100 exhibitors
  - Nearly 600 hosts and routers
  - Nearly 6 miles of cabling (various media)
- Both TCP/IP and OSI
- Massive demo of SNMP (25 vendors)
- Working FDDI ring (15 vendors)
- Microwave link to the hotel
- PPP demo
- Designed by Peter de Vries and myself

Outline

- The network design
  - Topology
  - TCP/IP specifics
  - OSI specifics (and tutorial)
- What we learned
- Credits

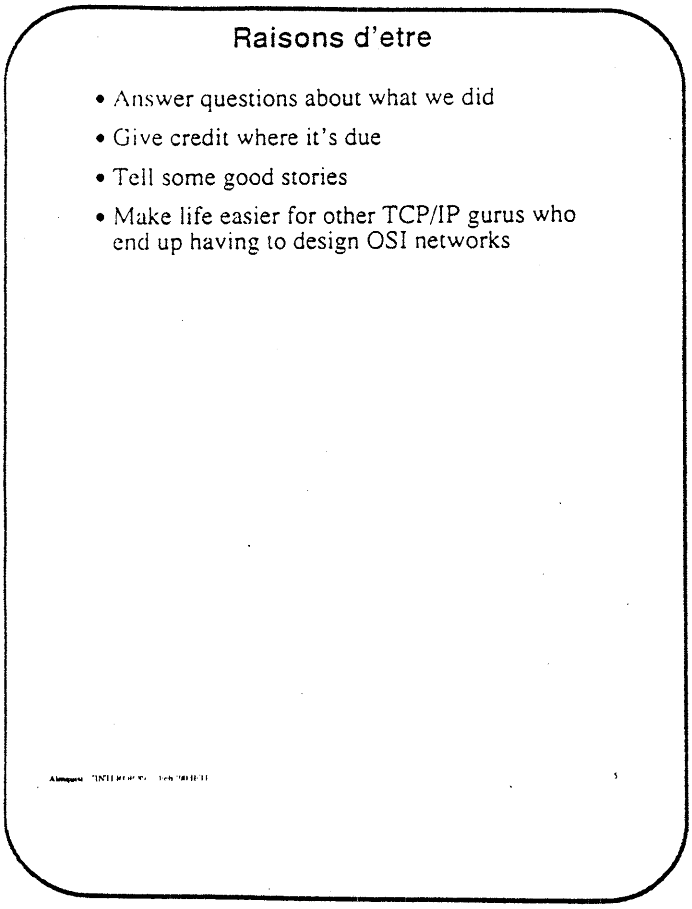

Raisons d'etre

- Answer questions about what we did
- Give credit where it's due
- Tell some good stories
- Make life easier for other TCP/IP gurus who end up having to design OSI networks

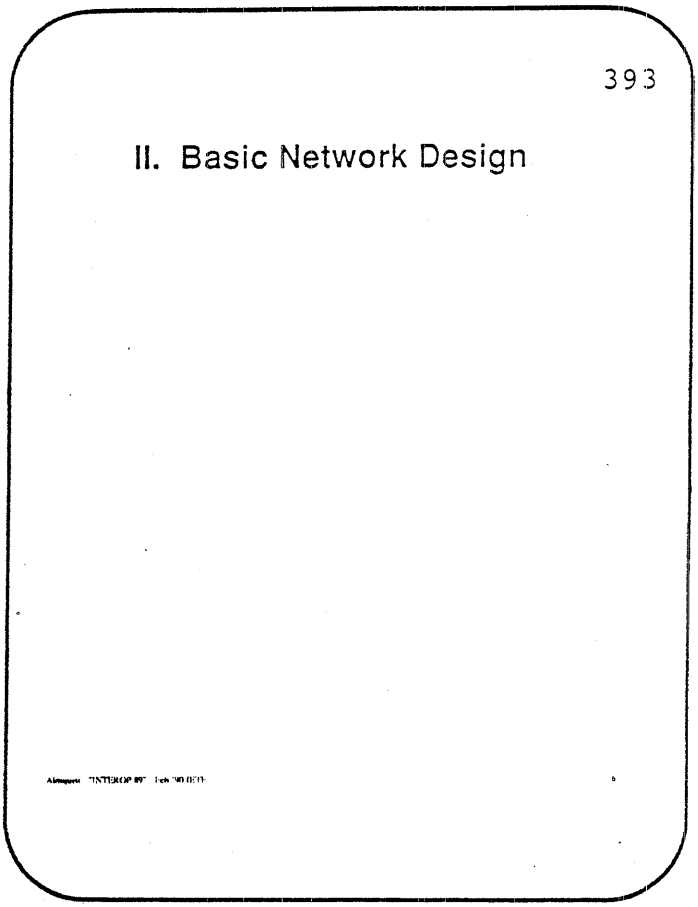

II. Basic Network Design

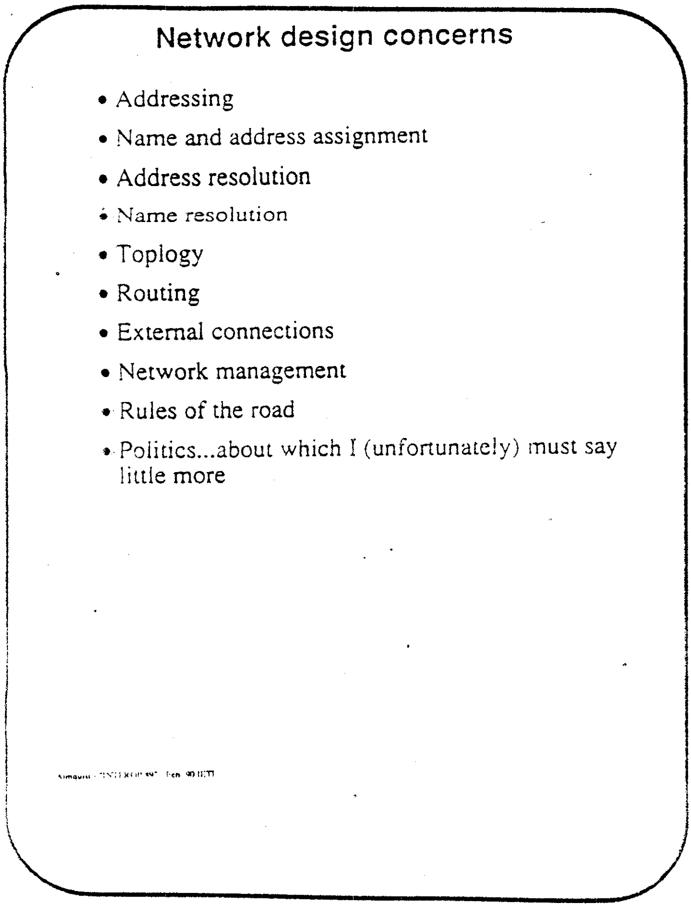

Network design concerns

- Addressing
- Name and address assignment
- Address resolution
- Name resolution
- Toplogy
- Routing
- External connections
- Network management
- Rules of the road
- Politics...about which I (unfortunately) must say little more

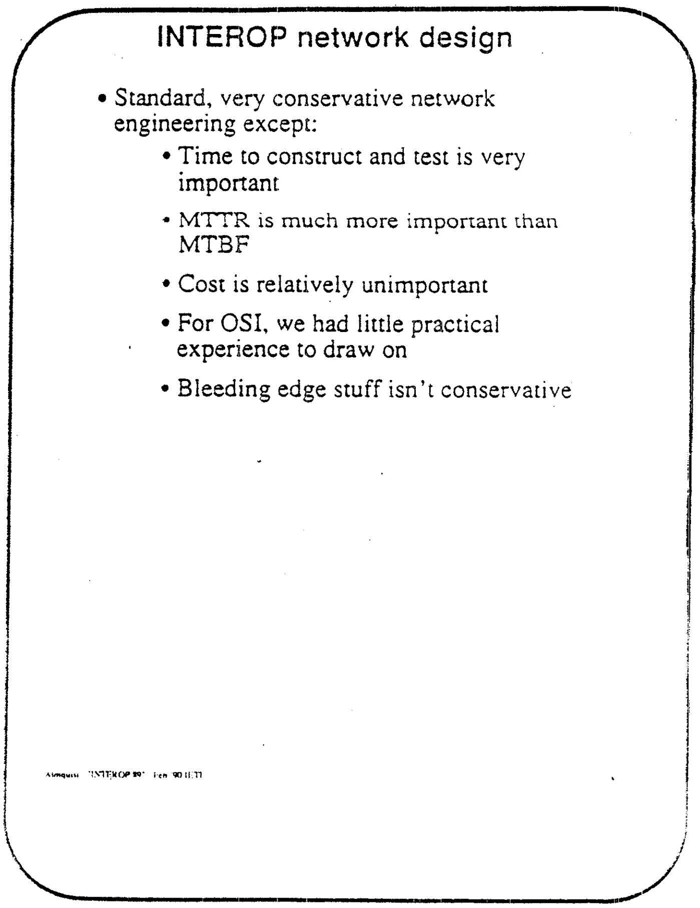

INTEROP network design

- Standard, very conservative network engineering except:
    - Time to construct and test is very important
    - MTTR is much more important than MTBF
    - Cost is relatively unimportant
    - For OSI, we had little practical experience to draw on
    - Bleeding edge stuff isn't conservative

Topology

- Logically a tree - no alternate routes
- Physically closer to a star

INTEROP BACKBONE SEGMENTS

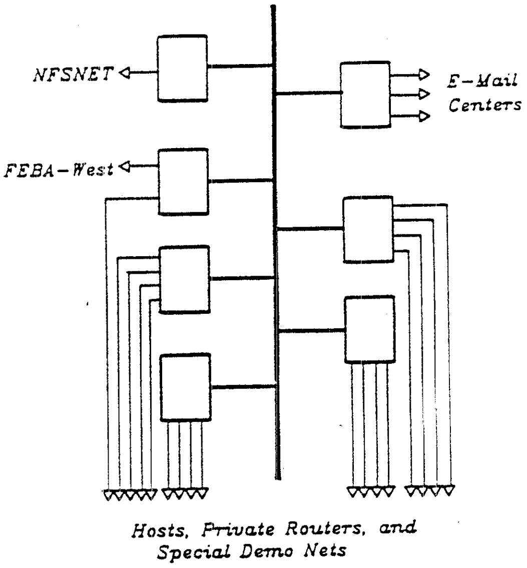

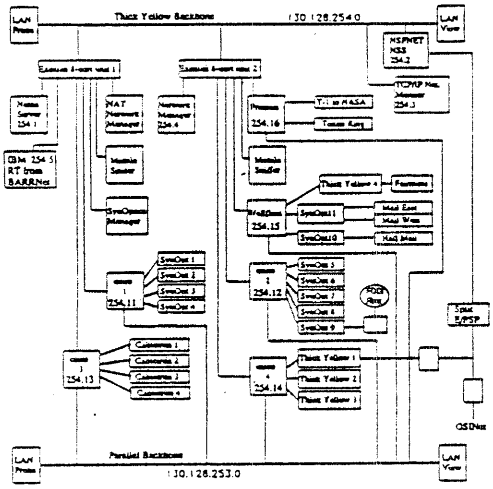

COURTESY PETER DE VRIES

Design criteria

- Subnets should be small and simple
- Balancing number of hosts and number of exhibitors per subnet
- Vendor requests
- Special protocol requirements (LAT, DECNET, Novell)
- Geographical locality
- Reducing single points of failure

III. TCP/IP

Addressing

- Subnetted class B network
- Subnetted class C network for FDDI ring
- Broadcast address format per HRRFC draft

Name and address assignment

- NIC allocated us a domain and IP network numbers
- Exhibitors chose host names of the form *companyname-hostname.ShowNet.COM*
- Exhibitors filled out forms specifying host name, type, ...
- We assigned IP addresses based on the forms

Address resolution

- ARP
- Proxy ARP disabled because it obscures a multitude of sins

Name resolution

- Host information entered into ShowNet.COM zone file
- IN-ADDR.ARPA zone file and HOSTS.TXT automatically generated
- Two primary name servers (1 local, 1 off-site)
- HOSTS.TXT and /etc/hosts via anonymous FTP for primitive hosts

Routing

- Routers from multiple vendors
  --> RIP and/or static routing
- Default gateway was host 1 on all subnets
- Subnets containing hosts used static routing
- Router interconnect subnet used RIP
- Improved technology allowed us not to send RIP updates on subnets where it wasn't used

396 External connections

- NSFNET
  - NSS connected directly to router interconnect subnet
  - Most backbone routers EGP'd with the NSS
  - EGP routes (minus BARRNet, NSI, and BBN Core routes) were dumped into RIP
- FEBA-West
  - T-1 link to NASA-AMES
  - RIP default route
- Corporate external nets not advertised at all

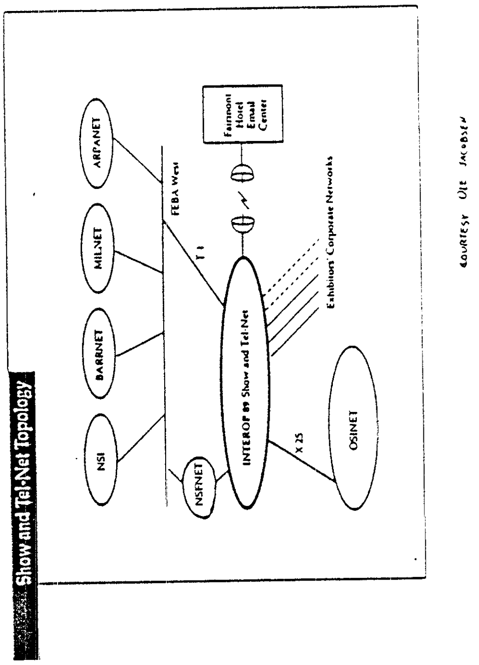

Show and Tel-Net Topology

ARPANET
MILNET
BARRNET
NSI
NSFNET
FEBA West
T 1
Fairmont Hotel Email Center
INTEROP 89 Show and Tel-Net
Exhibitors' Corporate Networks
X 25
OSINET

COURTESY OLE JACOBSEN

Network management

- Many high-tech tools, most of which the NOC personnel had no time to learn to use
- Packet-watchers on router interconnect net
- Patch panel so packet watchers could easily be attached to any backbone segment were spec'd but not as-built
- UTP vendor managed the UTP hubs

Rules of the road

- Broadcast restrictions
- Gatewaying restrictions
- Recommended HRRFC compliance
- Required RFC1009 compliance for routers
- "We can disconnect you if you break somebody else..."

IV. OSI

Overview

- 13+ vendors
- routers from cisco, CMC, Proteon, MERIT
- OSINET connection (national X.25 test net based on ACCUNET)
- Experimental NSFNET connection
- OSI support initially planned for only a few segments, but grew to cover most of the network

Basic OSI-speak

- end system (ES): host
- intermediate system (IS): router
- ES-IS: equivalent of ARP + ICMP redirect
- IS-IS: IGP
- subnet point of attachment (SNPA): link layer (e.g. Ethernet MAC) address

Protocols

- CONS vs CLNS
- 802.2 (not Ethernet) at the link layer
- CLNP, similar to IP
- TP4, similar to TCP
- no widely-implemented UDP equivalent yet
- session and presentation protocols
- VT for remote login, FTAM for file transfer, X.400 for email, X.500 for white pages

Layers and addresses

- "If we had only consulted the ancient mystics, we would have seen immediately that seven layers were required."

  -- Steve Crocker, *RFC1000*
- More ISO-speak: SAP addresses, selectors, and entity titles
- All layers use addresses; unlike in TCP/IP, a network layer address + some well-known magic numbers is not sufficient

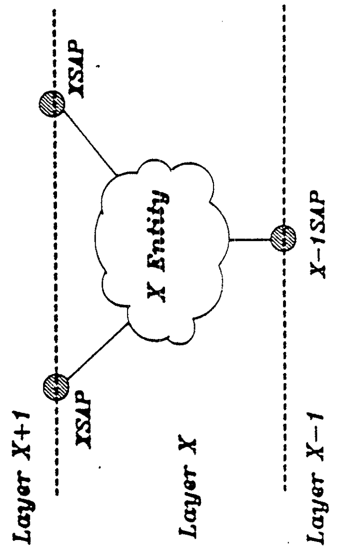

Addressing

- OSINet NSAP address format (though primative) seemed safest choice
- High order byte of suborg-id used to simplify routing routing
- Low order byte of suborg-id indicated corresponding IP subnet (OSI not used on class C net)

```
47 0004 0046 RR SS NNNNNNNNNNNN SS
(a) (b)  (c) (d) (e)    (f)      (g)
```

(a) Authority and Format Identifier (AFI). 47 means "ISO International Code Designator with Binary Abstract Syntax."

(b) ICD 0004 = OSINET

(c) OSINET Organizational ID 0046 = INTEROP Show Network

(d) and (e) OSINET Suborganization ID. This was assigned such that the first octet signified a router number within the show backbone, and the second octet selected a particular subnetwork (LAN).

(f) SNPA (MAC) address

(g) NSAP Selector

COURTESY DAVE KATZ

Name and address assignment

- Flat name space
- Exhibitor-chosen host names identical to IP host names without the domain part
- Address assignment handled the same as IP address assignment

Address resolution -- ES-IS

- Hosts multicast their existence
- Routers remember SNPA addresses of "local" hosts
- Routers multicast their existence
- Hosts remember one or more router SNPA addresses
- Routers can send redirects (including case where source and destination are on the same cable)
- Routers used static tables to remember SNPA addresses of hosts which could not support ES-IS

Name resolution

- Recall: applications (rather than hosts) have addresses
- Name service (X.500) is generally not implemented
- Machines use host tables, but there is no standard format
- (Non-standard) format host table provided via anonymous FTP

Toplogy

- Subset of the IP topology
- Attempts to limit OSI to one or two backbone routers didn't work

Routing

- Entirely static (IS-IS doesn't exist yet)
- Static routes are route prefixes; router uses longest match algorithm
- High-order byte of suborg-id used to route to backbone router closest to destination
- Low-order byte used to route from there to private router (if necessary)
- SNPA tables used to route to hosts on directly-attached segments

External connections

- NSFNET
  - We assigned addresses to NSFNET end systems, so no special routing was required
- OSINET
  - OSI default route

Network management

- Most packet sniffers understand OSI
- NSFNET and cisco agreed on Wisconsin-compatible ping
- Network operators had an OSI-capable host
- CMIP applications not available to us

Rules of the road

- NIST Stable Implementors' Agreements
- Basically, not yet known

V. What We Learned

What we learned

- We need to make sure that exhibitors know
  - The subnet mask
  - That bridging subnets together is very bad
- We should have protocol police
- Non-essential pieces don't get built
- Internet routing works much better than it did in 1988

What we re-learned

- Fancy net mgmt tools are not very useful (or used) if the staff doesn't have time beforehand to configure and learn to use the tools
- Planning is vital, and takes far longer than predicted
- Competitors really can work together
- Static routing is a b*tch
- Questionnaire writing is a fine art
- Operations staffing during the show should be preplanned
- We don't know how to convey to attendees the power and complexity of the network
- The best laid plans will sometimes fail; somehow, the talented volunteers pull it together anyway

What we did right

- Many small things, since this was our second attempt
- No DNS problems this year
- Backboards prevented cabling problems
- Internet routing worked reliably
- The network really did pass OSI packets

VI. Credits

What they did, and why

- Setup time allowed
  - 5.25 days total
  - 12.5 minutes per host
  - 15 seconds per foot of backbone cabling
- "I was invited to join about 10 other people to put together, play with, and take apart an arbitrarily complex network. A lot of people take this very seriously, and ignore that this is in reality a challenging goal. During the brief time it is up, it is a pretty amazing toy. The fact it was kept running in a near flawless condition was due to our pride in its magnificence. Our grand toy should appear nothing less than awe-inspiring to other people. As for our motivation for putting forth all this effort, if you have to ask, you wouldn't understand the answer." -- Stev Knowles

The core group

- Peter de Vries (project manager)
- Karl Auerbach
- Dave Bridgeham
- Eric Brunner
- Jeff Burgan (FEBA-West connection)
- Mario Castro
- Shelley de Vries
- Stev Knowles
- Steve Larbig
- Lisa Robertson
- John Romkey (hostmaster)

Greatly assisted by

- Karen Auerbach
- Geoff Baehr
- John Bashinski
- David Burdelski
- John Burruss
- James Davidson
- Carl Feil
- Elise Gerich
- Sue Hares
- Ronnie Hueter
- Ole Jacobsen
- Sandy Lerner
- Chris Lynch
- Milo Medin
- Dave Noble
- Ceal Preston
- Dave Preston
- Sue Romano
- Bill Rust
- Anthony Scampavia
- Cris Schudler
- Jim Sheridan
- Mary Stahl
- Geof Stone
- Mark Strangio
- James van Bokkelen
- Bruce Van Nice
- John Veizades
- Dave Vereeke
- Joe Vermeulen
- Denis Yaro
- and undoubtedly others...

The OSI gurus

- Rob Hagens (INTEROP OSI consultant)
- Jim Forster
- Dave Katz
- Marshall Rose/*The Open Book*

Those who helped with this talk

- Peter de Vries
- Dave Katz
- Allen Penny
- Ole Jacobsen/*ConneXions -- The Interoperability Report*

Appendix A

Attendees

Guy Almes
Rice University
PO Box 1892
Dept of Computer Science
Houston TX, 77251-1892
713-527-6038
almes@rice.edu

Philip Almquist
Consultant
214 Cole Street
San Francisco CA, 94117
415-752-2427
almquist@jessica.stanford.edu

Stan Ames
Mitre Corporation
Burlington Road
K302; M/S D115
Bedford MA, 01730
617-271-3182
sra@mbunix.mitre.org

Karen Armstrong
San.Diego SCC
CERFNet
PO Box 85608
San Diego CA, 92138-5608
619-534-5077
armstrong@sds.sdsc.edu

Cathy Aronson
Merit Computer Network
CICNet
1075 Beal Avenue
Ann Arbor MI, 48109-2112
313-936-2090
313-763-3000
cja@merit.edu

Douglas Bagnall
Hewlett-Packard
Apollo Division
330 Billerica Road; M/S CHR 03 DC
Chelmsford MA, 01824
508-256-6600
bagnall_d@apollo.hp.com
mit_eddie!apollo!bagnall_d

Fred Baker
Vitalink
6607 Kaiser Drive
Fremont CA, 94555
415-794-1100
baker@vitalink.com

Ballard Bare
Hewlett-Packard
8000 Foothills Blvd.
Roseville CA, 95678
916-785-5608
bare%hprnd@hplabs.hp.com

Richard Basch
MIT
1 Amherst Street
Cambridge MA, 02139
617-253-0100
probe@mit.edu

Derek Bennett
Florida State University
Education Center
F.I.R.N., B1-14
Tallahassee FL, 32399
904-487-0911
xndmis14@servax.bitnet

Art Berggreen
Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara CA, 93101
805-963-9431
art@sage.acc.com

Chet Birger
BBN Communications
150 CambridgePark Drive
Cambridge MA, 02140
617-873-2676
cbirger@bbn.com

Dave Borman
Cray Research
1440 Northland Dr.
Mendota Heights MN, 55120
612-681-3398
dab@cray.com
dab%oliver.cray.com@uc.msc.umn.edu

Karen Bowers
Corp. for Natl. Research Initiatives
1895 Preston White Drive
Suite 100
Reston VA, 22091
703-620-8990
(703) 582-9236
kbowers@nri.reston.va.us

Scott Bradner
Harvard U
William James Hall 1232
33 Kirkland Street
Cambridge MA, 02138
617-495-3864
sob@harvard.harvard.edu

Scott Brim
Cornell University
265 Olin Hall
Theory Center
Ithaca NY, 14853
607-255-8686
swb@devvax.tn.cornell.edu

Ronald Broersma
Naval Ocean Systems Center
Code 914
San Diego CA, 92152-5000
619-553-2293
ron@nosc.mil

Mats Brunell
SICS
Distributed Systems Lab
PO Box 1263
S-164 28 Kista ,
46 8 752 15 63
mats.brunell@sics.se

Theodore Brunner
Bell Communications Research
445 South Street
Morristown NJ, 07960
201-829-4678
tob@thumper.bellcore.com

Stewart Bryant
Digital Equipment
PO Box 121, Worton Grange
Imperial Way; RE02/GH2
Reading BERKS, RG2 OTU
0734 854682
bryant@janus.enet.dec.com

Jeffrey Burgan
NASA Ames Research Center
Moffett Field CA, 94035
415-604-5705
jeff@nsipo.nasa.gov

Jeffrey Case
U of Tennessee
Dept of Computer Science
107 Ayres Hall
Knoxville TN, 37996
615-974-0822
615-573-1434
case@utkux1.utk.edu

Stephen Casner
University of Southern California
4676 Admiralty Way
ISI
Marina del Rey CA, 90292
213-822-1511
casner@isi.edu

John Cavanaugh
NCR Comten
2700 Snelling Avenue N.
St. Paul MN, 55113
612-638-2822
john.cavanaugh@stpaul.ncr.com

Vinton Cerf
Corp. for Natl. Research Initiatives
1895 Preston White Drive
Suite 100
Reston VA, 22091
703-620-8990
703-573-1965
cerf@isi.edu
vcerf@nri.reston.va.us

A. Lyman Chapin
Data General
4400 Computer Drive
Westborough MA, 01580
508-870-6056
lyman-chapin@dgc.mceo.dgcom

Samir Chatterjee
NYNEX
500 Westchester Avenue
White Plains NY, 10604
914-683-2344
samir@nynexst.com

J. Noel Chiappa
IESG
708 E. Woodland
Grafton VA, 23692
804-898-7663
804-898-8183
617-898-2800
jnc@lcs.mit.edu

George Clapp
Ameritech Services
Gould Center, Building 40
2850 Golf Road
Rolling Meadows IL, 60008
708-806-8318
meritec!clapp@bellcore.bellcore.com

Richard Colella
Nat Inst of Standards & Tech
Building 225
Room B217
Gaithersburg MD, 20899
301-975-3627
colella@osi3.ncsl.nist.gov

Rob Coltun
U of Maryland
Computer Science Dept.
College Park MD, 20742-2411
301-454-2946
rcoltun@trantor.umd.edu

John Cook
Chipcom
195 Bear Hill Rd.
Waltham MA, 02154
617-890-6844
cook@chipcom.com

Dave Crocker
Digital Equipment Corporation
Network Systems Lab
100 Hamilton Ave.
Palo Alto CA, 94301 USA
415-688-6820
dcrocker@nsl.dec.com

Steve Crocker
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood MD, 21738
301-854-6889
crocker@tis.com

Raymond Curci
Florida State University
SCRI/442 SCL
Tallahassee FL, 32306
904-644-4587
curci@gw.scri.fsu.edu

James Davin
MIT LCS
Computer Science Lab, NE43-507
545 Technology Square
Cambridge MA, 02139
617-253-6020
jrd@ptt.lcs.mit.edu

Philip DeMar
Fermi National Accelerator Lab
PO Box 500
Batavia IL, 60510
708-840-3678
demar@fnalb.fnal.gov

Farokh Deboo
Interlink Computer Science
47370 Fremont Blvd.
Fremont CA, 94538
415-657-9800
fjd@interlink.com

Ralph Droms
Bucknell University
Computer Science Department
323 Dana Engineering
Lewisburg PA, 17837
717-524-1145
droms@sol.bucknell.edu

Tom Easterday
Ohio State University
1971 Neil Avenue
Room 406
Columbus OH, 43210
614-292-4027
tom@nisca.ircc.ohio-state.edu

Robert Enger
Contel Federal Systems
P O Box 10814
Chantilly VA, 22021-0814
703-818-5555
enger@sccgate.scc.com

Hunaid Engineer
Cray Research
1400 Northland Drive
Mendota Heights MN, 55120
612-681-3015
hunaid@opus.cray.com

Kent England
Boston U
Information Technology
111 Cummington Street; M/S IT
Boston MA, 02215
617-353-2780
kwe@bu.edu

Dino Farinacci
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
415-940-7661
dino@bridge2.3com.com

Dennis Ferguson
University of Toronto
5 King's College Road
Toronto ONTARIO, M5S 1A4
416-978-2455
dennis@gw.ccie.utoronto.ca

Metin Feridun
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
617-873-1870
mferidun@bbn.com

Louis Fernandez
BBN Communications
150 CambridgePark Drive
Cambridge MA, 02140
617-873-2781
lfernandez@bbn.com

Michael Fidler
Ohio State University
1971 Neil Avenue, Rm 406
IRCC
Columbus OH, 43210-1210
614-292-4843
ts0026@ohstvma.ircc.ohio-state.edu

Dale Finkelson
U of Nebraska-Lincoln
29 Walter Scott Engr Ctr
MIDnet
Lincoln NE, 68588-0534
402-472-5032
dmf@westie.unl.edu
dmf@westie.unl.edu

James Forster
cisco Systems
1350 Willow Road
Menlo Park CA, 94025
415-326-1941
forster@cisco.com

Richard Fox
Hughes LAN
950 Linden Ave. #208
Sunnyvale CA, 94086
415-966-7924
sytek!rfox@sun.com

Stanley Froyd
Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara CA, 93101
805-963-9431
sfroyd@salt.acc.com

Vince Fuller
Stanford University
Networking Systems
115 Pine Hall
Stanford CA, 94305
415-723-6860
fuller@jessica.stanford.edu

James Galvin
Trusted Information Systems
3060 Washington Road
Glenwood MD, 21738
301-854-6889
galvin@tis.com

Der-Hwa Gan
3Com
2081 North Shoreline Blvd.
Mountain View CA, 94043
415-969-4400

Ella Gardner
Mitre Corporation
7525 Colshire Drive
McLean VA, 22102-3481
703-883-5826
epg@gateway.mitre.org

Elise Gerich
University of Michigan
1075 Beal Avenue
Ann Arbor MI, 48109-2112
313-936-3000
epg@merit.edu

Adrianne Glappa
NEC America
110 Rio Robles
San Jose CA, 95134
408-922-3862

Herve Goguely
3Com
2081 North Shoreline Blvd
Mountain View CA, 94043
415-940-7645
rvg@bridge2.3com.com

Steven Goldstein
National Science Foundation
1800 G Street NW
Rm 416
Washington DC, 20550
202-357-9717
goldstein@note.nsf.gov

Martin Gross
DCA
1860 Wiehle Avenue
Code R640
Reston VA, 22090-5500
703-437-2165
martin@protolaba.dca.mil

Phill Gross
Corp. for Natl. Research Initiatives
1895 Preston White Drive
Suite 100
Reston VA, 22091
703-620-8990
pgross@nri.reston.va.us

Robert Hagens
U of Wisconsin-Madison
Computer Science Dept.
1210 West Dayton Street
Madison WI, 53706
608-262-1017
hagens@cs.wisc.edu

Jack Hahn
University of Maryland
Computer Science Center
College Park MD, 20742
301-454-5434
hahn@umd5.umd.edu

Tony Hain
Lawrence Livermore Natl Lab
PO Box 5509
Livermore CA, 94550
415-422-4017
hain@nmfecc.arpa

Tom Halcin
Hewlett-Packard
19420 Homestead Rd.
Cupertino CA, 95014
408-447-2480
halcin%hpinddm@hplabs.hp.com

Martyne Hallgren
Cornell University
265 Olin Hall
Theory Center
Ithaca NY, 14853-5210
607-255-8686
martyne@tcgould.tn.cornell.edu

Brian Handspicker
Digital Equipment Corporation
550 King Street
LKG1-2/E19
Littleton MA, 01460
508-486-7894
bd@vines.dec.com

Gene Hastings
Pittsburgh Supercomputer Center
4400 5th Avenue
Pittsburgh PA, 15213
412-268-4960
hastings@psc.edu

Kenneth Hays
Florida State University
Super. Computations Res. Inst.
400 Science Center Library
Tallahassee FL, 32306-4052
904-644-7053
hays@scri1.scri.fsu.edu

Juha Heinanen
Tampere University of Technology
Software Systems Laboratory
PO Box 527
SF-33101 Tampere ,
358 31 162578
jh@funet.fi

Christine Hemrick
Bell Communications Research
331Newman Springs Road
Red Bank NJ, 07701
201-758-2754
cfh@sabre.bellcore.com

Robert Hinden
BBN Communications
50 Moulton Street
Cambridge MA, 02138 USA
617-873-3757
hinden@bbn.com

Russell Hobby
University of California
Computing Services
Surge II - Room 1400
Davis CA, 95616
916-752-0236
rdhobby@ucdavis.edu

Peter Honeyman
U of Michigan
IFS Project
535 West William
Ann Arbor MI, 48103-4943
313-763-4403
honey@citi.umich.edu

Jeffrey Honig
Cornell University
265 Olin Hall
Theory Center
Ithaca NY, 14853-5201
607-255-8686
jch@tcgould.tn.cornell.edu

Binh Hua
IBM
Neighborhood Road
80SA/700
Kingston NY, 12401
914-385-3365

Doug Hunt
Prime Computer
500 Old Connecticut Path
Framingham MA, 01701
508-879-2960
dhunt@enr.prime.com

Steven Hunter
Lawrence Livermore Natl Lab
PO Box 808
Livermore CA, 94550
415-423-2219
hunter@ccc.mfecc.arpa

Tom Hytry
AT&T Bell Laboratories
1100 E Warrenville Road
Naperville IL, 60566
312-979-7313
tlh@iwlcs.att.com

Joel Jacobs
Mitre Corporation
Burlington Road
Bedford MA, 01821
617-271-7373
jdj@mitre.org

Ole Jacobsen
Interop, Inc.
806 Coleman Avenue
#9
Menlo Park CA, 94025
415-325-9542
ole@csli.stanford.edu

Van Jacobson
Lawrence Berkeley Lab
One Cycloctron Road
Berkeley CA, 94720
415-486-6411
van@helios.ee.lbl.gov

Ron Jacoby
Silicon Graphics
2011 N. Shoreline Blvd.
PO Box 7311
Mountain View CA, 94039-7311
415-960-1980
rj@sgi.com

B.V. Jagadeesh
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
bvj@chamundi.esd.3com.com

Phil Jensen
Florida State University
Systems Group-Computing Ctr
200 Sliger Building
Tallahassee FL, 32306-3042

Brad Johnson
Open Software Foundation
11 Cambridge Center
Canbridge MA, 02142
617-621-8849
bradcj@osf.org

Dan Jordt
U of Washington
170 Academic Computer Center
3737 Brooklyn Avenue NE
Seattle WA, 98105
206-543-7352
danj@cac.washington.edu

Mary K.
SRI International
Net Info Sys Ctr
333 Ravenswood Ave. Rm EJ 296
Menlo Park CA, 94025
415- 859-4775
stahl@nisc.sri.com

Michael Karels
U of California
CSRG Computer Science Div. EECS
457 Evans Hall
Berkeley CA, 94720
415-642-4948
karels@berkeley.edu

Dave Katz
Merit Computer Network
1075 Beal Avenue
Ann Arbor MI, 48103
313-763-4898
dkatz@merit.edu

David Kaufman
Proteon
Two Technology Dr.
Westborough MA, 01581-5008
508-898-2800
dek@proteon.com

Daniel Kellen
Digital Equipment
5401 Corporate Woods Dr., #850
Pensacola FL, 32561
904-882-5498
kellen@eglin.af.mil

Peter Kirstein
University College London
Department of Computer Science
31 Bancroft Ave.
London , N2 OAR
44-1-380-7286
kirstein@cs.ucl.ac.uk

Stev Knowles
FTP Software
26 Princess Street
Wakefield MA, 01880-3004
617-246-0900 x270
stev@ftp.com

Lee LaBarre
Mitre Corporation
Burlington Rd
M/S E066
Bedford MA, 01730
617-271-8507
cel@mbunix.mitre.org

Tony Lauck
Digital Equipment
550 King Street
LKG 1-2/A19
Littleton MA, 01460-1289
508-486-7644
lauck@dsmail.dec.com

Walter Lazear
Mitre Corporation
7525 Colshire Drive
McLean VA, 22102
703-883-6515
lazear@gateway.mitre.org

Mike Little
SAIC
8619 Westwood Center Dr.
Vienna VA, 22182
703-749-5360
little@saic.com

Dan Long
BBN Systems & Technologies
NEARNet
10 Moulton Street
Cambridge MA, 02138
617-873-2766
long@bbn.com

E. Paul Love
San Diego Supercomputer Center
PO Box 85608
San Diego CA, 92138
619-534-5043
loveep@sds.sdsc.edu

Dan Lynch
Interop, Inc.
21370 Vai Avenue
Cupertino CA, 95014
415-941-3399
408-996-1108
lynch@isi.edu

Charles Lynn
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
617-873-3367
clynn@bbn.com

Gary Malkin
Proteon
2 Technology Drive
Westborough MA, 01581
508-898-2800
gmalkin@proteon.com

Louis Mamakos
University of Maryland
Computer Science Center - Syst
College Park MD, 20742
301 454-2943
louie@trantor.umd.edu

Mike Marcinkevicz
TRW
Space and Defense Sector
One Space Park; M/S: R3/2089
Redondo Beach CA, 90278
213-812-2161
mdm@gumby.dsd.trw.com

Tony Mason
Transarc Corporation
The Gulf Tower
707 Grant Street
Pittsburgh PA, 15219
412-338-4400
mason@transarc.com

Matt Mathis
Pittsburgh Supercomputer Center
4400 5th Ave.
Pittsburgh PA, 15213
412-268-3319
mathis@pele.psc.edu

Jim McCabe
Computer Sciences Corporation
NASA Ames Research Center
Moffett Field CA, 94035
415-604-4425
jmccabe@orville.nas.nasa.gov

Keith McCloghrie
Hughes LAN Systems
1225 Charleston Road
Mountain View CA, 94043
415-966-7934
sytek!kzm@hplabs.hp.com

Paul McKenney
SRI International
333 Ravenswood Ave.
EK 360
Menlo Park CA, 94025
415-859-4910
mckenney@sri.com

Leo McLaughlin
Wollongong Group
1129 San Antonio Road
Palo Alto CA, 94303
415-962-7100
ljm@twg.com

Milo Medin
NASA Ames Research Center
NASA Science Internet Project Office
Moffett Field CA, 94035
415-490-9157
medin@nsipo.nasa.gov

Bill Melohn
Sun Microsystems
2550 Garcia Avenue
Mountain View CA, 94043
415-336-2941
melohn@sun.com

Donald Merritt
USA Ballistic Research Lab
Attn: AMXBR-SECAD
Aberdeen Proving Ground MD, 21005-5066

301-278-6808
don@brl.mil

David Miller
Mitre Corporation
Burlington Road
Bedford MA, 01730
617-271-3993
dtm@mitre.org

Cyndi Mills
BBN Communications
150 CambridgePark Drive
Cambridge MA, 02140
617-873-4143
cmills@bbn.com

Greg Minshall
Novell
1340 Treat Blvd.
Suite 500
Walnut Creek CA, 94596
415-947-0998
415-938-2562
minshall@kinetics.kinetics.com

Paul Mockapetris
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey CA, 90292-6695 USA
(213) 822-1511
pvm@isi.edu

Jeffrey Mogul
Digital Equipment
Western Research Labs
100 Hamilton Ave.
Palo Alto CA, 94301
415-853-6643
mogul@decwrl.dec.com

Dave Monachello
Datability Systems
322 8th
New York NY, 10001
212-807-7800
dave@pluto.dss.com

Berlin Moore
PREPnet
530 N. Neville Street
Pittsburgh PA, 15213
412-268-7873
prepnet@andrew.cmu.edu

Dennis Morris
Defense Communication Agency
DDN Program Office
Code DDEI
Washington DC, 20305-2000
703-285-5221
morrisd@imo-uvax.dca.mil

Donald Morris
NCAR
PO Box 3000
Scientific Computing Division
Boulder CO, 80307
303-497-1282
morris@ucar.edu

John Moy
Proteon
Two Technology Drive
Westborough MA, 01581-5008
508-898-2800
jmoy@proteon.com

Oscar Newkerk
Digital Equipment
14475 NE 24th St.
Bellevue WA, 98007
206-865-8913
newkerk@decwet.dec.com

Rebecca Nitzan
NASA
600 Maryland Ave., SW
Suite 200 East
Washington DC, 20024
202-554-8677
703-765-6115
nitzan@nsipo.nasa.gov

David O'Leary
University of Maryland
SURAnet
Computer Science Center
College Park MD, 20742
301-454-8055
301-454-5434
oleary@umd5.umd.edu

Lee Oattes
University of Toronto
Computer Services
4 Bancroft Avenue, Rm. 102
Toronto ON, M5S 1C1
416-978-5448
oattes@utcs.utoronto.ca

Dave Oran
Digital Equipment
216 Lakewood Dr.
Bloomington IN, 47401
812-332-0237
oran@oran.dec.com

Donald Pace
Florida State University
Computing Center
Tallahassee FL, 32306
904-644-2591
pace@fsu1.cc.fsu.edu

Craig Partridge
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
617-873-2459
craig@bbn.com
craig@nnsc.nsf.net

Gurudatta Parulkar
Washington University
Campus Box 1045
St. Louis MO, 63130
314-889-4621
guru@flora.wustl.edu

David Perkins
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
415-694-2808
dave_perkins@3com.com

Drew Perkins
Inter Stream
824 Lilac Street
Pittsburgh PA, 15217
412-422-9828
ddp@andrew.cmu.edu

Radia Perlman
Digital Equipment
550 King St
LKG 1-2/A19
Littleton MA, 01460-1289
508-486-7648
508-263-6730
NAC::Perlman
Perlman%BERGIL.DEC@DECWRL.DEC.COM

Richard Pethia
Carnegie Mellon U
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh PA, 15213-3890
412-268-7739
rdp@sei.cmu.edu

Michael Petry
University of Maryland
Computer Science Center
Room 3339
College Park MD, 20742
301-454-2943
petry@trantor.umd.edu

David Piscitello
Bell Communications Research
331 Newman Springs Road
Red Bank NJ, 07701
201-758-2286
dave@sabre.bellcore.com

Dave Pokorney
U of Florida
Space Science Research Bldg
Room 112
Gainesville FL, 32611
904-392-4601
poke@nervm.nerdc.ufl.edu

Nuggehalli Pradeep
Computer Sciences Corporation
Nasa Ames Research Center
Moffett Field CA, 94035
415-604-4514
pradeep@orville.nas.nasa.gov

K.K. Ramakrishnan
Digital Equipment
LKG 1-2/A19
550 King Street
Littleton MA, 01460-1289
508-486-7267
rama%erlang.dec.com@decwrl.dec.com

Michael Reilly
Digital Equipment
Network Systems Lab
100 Hamilton Ave.; M/S UCO-4
Palo Alto CA, 94301
415-853-6593
reilly@nsl.dec.com

Yakov Rekhter
IBM
TJ Watson Research
Route 134, PO Box 218
Yorktown Heights NY, 10598
914-945-3896
yakov@ibm.com

Joel Replogle
University of Illinois
National Center Supercomputing Applications
605 East Springfield Avenue
Champaign IL, 61820
217-244-0636
replogle@ncsa.uiuc.edu

Joyce Reynolds
University of Southern California
4676 Admiralty Wy #1001
ISI
Marina del Rey CA, 90292-6695
213-822-1511
jkrey@venera.isi.edu

Jim Robertson
3Com
2081 N Shoreline Blvd
Mountain View CA, 94043
415-940-2683
jar@bridge2.3com.com

Marshall Rose
Performance Systems International, Inc.
(PSI)
California Office
420 Whisman Court
Mountain View CA, 94043-2112 USA
(415) 961-3380
mrose@psi.com

Mark Rosenstein
MIT
Amherst St.
Room E40-3311
Cambridge MA, 02139
617-253-1530
mar@athena.mit.edu

Karen Roubicek
BBN Systems & Technologies
10 Moulton Street
NSF Network Service Ctr
Cambridge MA, 02138
617-873-3361
roubicek@nnsc.nsf.net

Jonathan Saperia
Digital Equipment
550 King Street
M/S LKG 1-2/B13
Littleton MA, 01460-1289
508-486-5542
saperia%tcpjon@decwrl.dec.com

Martin Schoffstall
Performance Systems International
PO Box 3850
Reston VA, 22091
schoff@psi.com
info@psipost.com

Tim Seaver
Microelectronics Center of North Carolina

PO Box 12889
Research Triangle Park NC, 27709
919-248-1973
tas@mcnc.org

Steve Senum
Network Systems Corporation
7600 Boone Avenue North
Minneapolis MN, 55428
612-424-4888
sjs@network.com

Jim Sheridan
IBM
166 East Shore Drive
PO Box 334
Whitmore Lake MI, 48189
313-393-6537
jsherida@ibm.com

Steven Shibuyama
Unisys
Network and Communications Group
5151 Camino Ruiz; M/S 02I203
Camarillo CA, 93010
805-987-9481

Jim Showalter
DCA
1860 Wiehle Avenue
Reston VA, 22090-5500
703-437-2580
gamma@mintaka.dca.mil

Chi Shue
Open Software Foundation
11 Cambridge Center
Cambridge MA, 02142
617-621-8972
chi@osf.org

Keith Sklower
University of California
Computer Science Dept.
570 Evans Hall
Berkeley CA, 94720
415-642-9587
sklower@okeeffe.berkeley.edu

Pat Smith
Merit Computer Network
1075 Beal-NDSB
Ann Arbor MI, 48109
800-66-Merit
psmith@merit.edu

Frank Solensky
Racal InterLAN
155 Swanson Rd
Boxborough MA, 01719
508-263-9929
solensky@interlan.interlan.com

Michael St Hohn
Department of Defense
Attn T41
9800 Savage Rd
Ft. Meade MD, 20755
301-688-6742
stjohns@umd5.umd.edu

Mary Stahl
Hewlett-Packard
Apollo Systems Division
330 Apollo Drive
Chelmsford MA, 01824
508-256-6600 x5963
strohl@apollo.hp.com

Tony Staw
Digital Equipment
Digital Park
Worton Grange, Imperial Way
Reading BERKS, RG2 OTE
011 44 734 868711
x3908
staw@marvin.enet.dec.com

Martha Steenstrup
BBN Communications
150 CambridgePark Dr.
Rm 20/665
Cambridge MA, 02140
617-873-3192
msteenst@bbn.com

Louis Steinberg
IBM
472 Wheelers Farms Rd
M/S 91
Milford CT, 06460
203-783-7175
louiss@ibm.com

Robert Stine
SPARTA
7926 Jones Branch Dr
Suite 1070
McLean VA, 22102
703-448-0210
stine@sparta.com

Roxanne Streeter
NASA Ames Research Center
Sterling Software
Moffett Field CA, 94035
415-694-4845
streeter@nsipo.arc.nasa.gov

Scott Stursa
Florida State University
Education Center
F.I.R.N., B1-14
Tallahassee FL, 32399
904-487-0911
xndmis14@servax.bitnet

Allen Sturtevant
Lawrence Livermore Natl Lab
PO Box 808
Livermore CA, 94550
415-422-8266
sturtevant@ccc.nmfecc.gov

Zaw-Sing Su
SRI International
333 Ravenswood Ave.
EJ280
Menlo Park CA, 94025
415-859-4576
zsu@tsca.istc.sri.com

Dean Throop
Data General
62 Alexander Dr.
Research Triangle Park NC, 27709
919-549-8421
throop@dg-rtp.dg.com

Claudio Topolcic
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
617-873-3874
topolcic@bbn.com

Paul Tsuchiya
Bell Communications Research
435 South Street
Morristown NJ, 07960
201-829-4484
tsuchiya@thumper.bellcore.com

James Van Bokkelen
FTP Software
26 Princess Street
Wakefield MA, 01880-3004
617-246-0900
jbvb@ftp.com

Kenneth Van Wyk
SEI/ CERT
4500 Fifth Avenue
Pittsburgh PA, 15213
412-268-6935
krvw@sei.cmu.edu

Gregory Vaudreuil
Corp for Natl Research Initiatives
1895 Preston White Dr
Suite 100
Reston VA, 22091 USA
703-620-8990
gvaudre@nri.reston.va.us

Ross Veach
University of Illinois
Computing Services Office
1304 West Springfield
Urbana IL, 61801
217-244-4274
rrv@uiuc.edu

John Veizades
Apple Computer
20525 Mariani Ave.
Cupertino CA, 95014
408-974-2672
veizades@apple.com

Sudhanshu Verma
Hewlett-Packard
19420 Homestead Road
Cupertino CA, 95014
408-447-3417
verma@hpindbu.hp.com

A. Lee Wade
NASA Ames Research Center
Moffett Field
Mountain View CA, 95045
415-604-4789
wade@orion.arc.nasa.gov

Steve Waldbusser
Carnegie-Mellon University
4910 Forbes Avenue
Pittsburgh PA, 15213
412-268-6628
sw01@andrew.cmu.edu

Carol Ward
U of Colorado
Westnet
3645 Marine Street
Boulder CO, 80309-0455
303-492-5860
cward@spot.colorado.edu

Rick Wilder
Mitre Corporation
1031 Poplar Drive
Falls Church VA, 22046
703-883-6174
rick@gateway.mitre.org

Steve Willis
Wellfleet Communications
12 DeAngelo Drive
Bedford MA, 01730
617-275-2400
swillis@wellfleet.com

Linda Winkler
Argonne National Laboratory
Building 221, B-251
Argonne IL, 60439
708-972-7236
b32357@anlvm.ctd.anl.gov

Dan Wintringham
Ohio Supercomputer Center
1224 Kinnear Road
Columbus OH, 43212
614-292-0901
danw@igloo.osc.edu

David Wittbrodt
cisco Systems
1350 Willow Road
Menlo Park CA, 94025
415-326-1941
dmw@cisco.com

John Wobus
Syracuse University
Computing and Network Services
Machinery Hall
Syracuse NY, 13244
315-443-4324
jmwobus@suvm.acs.syr.edu

Robert Woodburn
SAIC
CSEIC
8619 Westwood Center Drive
Vienna VA, 22182
703-734-9000
woody@saic.com

Brian Yasaki
Wollongong Group
7799 Leesburg Pike
Suite 1100, North Tower
Falls Church VA, 22043
703-847-6340
bky@twg.com

Raj Yavatkar
University of Kentucky
Department of Computer Science
Lexington KY, 40506-0027
(606) 257-6745
(606) 257-3961
(606) 257-4078
raj@ms.uky.edu

Mary Youssef
IBM
472 Wheelers Farms Rd.
M/S91
Milford CT, 06460
203-783-4338
mary@ibm.com

Aileen Yuan
Mitre Corporation
7525 Colshire Dr.
McLean VA, 22102
703-883-7023
aileen@gateway.mitre.org

David Zimmerman
Convex Computer
3000 Waterview Parkway
PO Box 833851
Richardson TX, 75083-3851
214-497-4164
dpz@convex.com