PROCEEDINGS OF THE THIRTEENTH INTERNET ENGINEERING TASK FORCE

APRIL 11-14, 1989

COCOA BEACH, FLORIDA

Compiled and Edited

by

Phill Gross

Karen L. Bowers

May 1989

Acknowledgements

The Thirteenth IETF Meeting was held at the Cocoa Beach Hilton and Towers at Cocoa Beach, Florida on April 11-14, 1989. On behalf of the IETF I would like to thank Susie Karlson (ACE) for making this a successful and smooth running session; she was faced with many a challenge in a somewhat "user hostile" hotel setting and came through with flying colors. A special thanks to you, Susie, for the great on site support!!

Though we were unable to make use of the conference facilities at the training center of the Kennedy Space Center, we certainly enjoyed being sponsored by the hospitable Dave Roberts (NASA EL-ESS-4 Office). Dave and his associates, Dale Pope, Dick Batton and Jack Martin personally treated us to a VIP tour of the KSC. Logistics were orchestrated by Mitch Barnes of the Public Affairs Office. Thank you all for the first class touch this tour placed on a very productive four day meeting.

The final thank you goes again to Monica Hart (NRI) whose relentless efforts have greatly improved the format and quality of these Proceedings.

Karen L. Bowers

Table of Contents

I. Chairman's Message

Phill Gross
NRI

Chairman's Message

An Important Anniversary

The April 11-14 1989 meeting of the IETF coincided closely with the anniversary of a significant event. On April 7, 1969, RFC 1 was issued. The title was "Host Software", and the author was Steve Crocker. It was interesting to read this paper from a twenty year vantage point.

Quoting from RFC 1:

****

"Introduction

The software for the ARPA network exists partly in the IMPs and partly in the respective HOSTs. BB&N has specified the software in the IMPs and it is the responsibility of the HOST groups to agree on the HOST software.

During the summer of 1968, representatives from the initial four sites met several times to discuss the HOST software and initial experiments on the network. There emerged from these meetings a working group ...

I present here some of the tentative agreements reached and some of the open questions encountered. VERY LITTLE OF WHAT IS HERE IS FIRM [emphasis added] and reactions are expected. ...

Some Requirements Upon the Host-to-Host Software

... As with any new facility, there will be a period of very light usage until the community of users ... begins to depend on it. ... It seems natural to provide the ability to use any remote HOST as if it had been dialed from a TTY (teletype) terminal. Additionally, we would like some ability to transmit a file in a somewhat different manner perhaps than simulating a teletype. ...

One of the inherent problems in the network is the fact that all responses from a remote HOST will require on the order of a half-second or so, no matter how simple. For teletype use, we could shift to a half-duplex local-echo arrangement, but this would destroy some of the usefulness of the network."

****

Working Groups, host requirements (unfirm host requirements, we should note!), and the expectation that a community of users will grow to depend on network communication -- there is much prophesy in this first RFC! We can also see the roots of TELNET and FTP.

In reading RFC 1, I was struck by how much has changed over 20 years, and yet how many of the challenges and fundamental problems remain today. Twenty years ago there were four IMP sites on the world's only packet switching network. Today there are 10's of 1000's of hosts and 100's of 1000's of users reachable on an international Internet of low-1000's of networks. The growth of the IETF has mirrored that pattern. The

group had its origins in a 15 person working party of government contractors in the mid-80's. Typical meetings are now nearly ten times that large with a large vendor and user constituency. When we first formed working groups, there were four; now there are 22 with several others in various stages of formation. Of these working groups, the IETF currently has a working group still attempting to firm up host rquirements and the Telnet WG is still dealing with the issue of a line oriented local-echo.

However, the continued need for such groups is more an indication of the incredible growth and change in the environment, rather than an indication of a lack of progress. We have always been victimized by the growth and demand created by the abundance of technical success, rather than the lack of success. So, challenges certainly remain. But judging from the particularly sharp technical progress in recent meetings, I believe we have the right to celebrate this notable anniversary in style.

(Thanks to Bob Braden (ISI) for noting the anniversary of RFC 1)

Another Significant Milestone

At the June 1988 IETF meeting in Annapolis, Mark Pullen of DARPA reported that the Arpanet was being decommissioned. Plans for this remarkable event have become firm, and the initial actions are being taken. The Arpanet Evaporation schedule, as provided by DARPA, was reported at the April 1989 meeting. The slides speak for themselves, and little needs to be added. However, Vint Cerf, considered by many to be the father of the Arpanet, has provided a touching perspective on the subject. Please enjoy.

The April 1989 Meeting

Fifteen Working Groups met at the Cocoa Beach IETF. Two new WGs met for the first time. These new groups are the NOC Tools catalogue WG and the Dynamic Configuration WG. Another experimental WG to examine a possible unified network management interface also met for the first time, but it has not yet been decided to continue that effort. Since April, two other groups have formed under the auspices of the User Services WG. These are the Network Information Services Infrastructure WG and the User Documents bibliography WG. Descriptions and reports on all these activities are contained in these Proceedings.

Three other WGs had progressed to the point of making detailed technical reports to the Plenary. These WGs are Interconnectivity, Open Routing, and Open OSF Routing. There was even a minority report to the IWG's proposal for a mid-term routing architecture.

Just prior to the April meeting, Cornell released "The Computer Worm" report resulting from an internal investigation. Jeff Schiller, who had previewed the report, was kind enough to give an unscheduled discussion of that report to the Plenary. The conclusion section of the Cornell report has been included in these Proceedings.

Other News

Responding to a vote of attendees at the January 1989 IETF meeting in Austin Texas, the April meeting was extended to 3.5 days. The first two days were devoted fully to Working Group sessions, the third day was devoted fully to technical presentations, and the concluding half day was devoted to reports from the Working Groups. This increased the number of sessions for working group meetings, where much of the technical work is pursued, from three to four.

At the April meeting in Cocoa Beach, we held a luncheon for the Working Group chairs to receive feedback on this and other actions (See Section II for more details). At this luncheon, a second refinement was suggested to give even more time for Working Groups activity. The suggested schedule was:

Days 1 and 2

9 am - 12 WG Morning session
1 pm - 4 pm WG Afternoon session
4 pm - 5:30 pm Technical Presentations (in Plenary)

Day 3

9 am - 12 WG Morning session
1 pm - 5:30 pm Technical Presentations (in Plenary)

Day 4

9 am - 12 WG Reports

This gives an additional period for WG sessions, making a total of five sessions, but retains the overall time available for technical Plenary presentations. This will reduce the number of overlapping WGs meetings. This format will be tried as an experiment at the next several meetings.

IETF Working Group Status

(April 1989)

Working Groups	RFC or Draft?	Met Apr 89?	Current Report?	Chair or POC (address)
ALERTMAN	-	Yes	Yes	Louis Steinberg (IBM) louiss@ibm.com
Authentication	Yes	Yes	Yes	Jeff Schiller (MIT) jis@athena.mit.edu Jon Rochlis (MIT) jon@athena.mit.edu
CMIP-over-TCP (CMOT)	Yes	No	-	Lee LaBarre (MITRE) cel@mitre.org
DNS (new)	Yes	Yes	Yes	Paul Mockapetris (ISI) pvm@isi.edu
Dyn. Host Config. (new)	-	Yes	Yes	Ralph Droms (Bucknell) droms@cs.purdue.edu
Host Requirements	Yes	No	-	Bob Braden (ISI) braden@isi.edu
Interconnectivity	Yes	Yes	Yes	Guy Almes (Rice) almes@rice.edu
Internet MIB	Yes	No	-	Craig Partridge (BBN) craig@nnsc.nsf.net
JoMann/NSFnet Req Mon.	Yes	Yes	Yes	Susan Hares (Merit) skh@merit.edu
LAN Mgr MIB	No	Yes	Yes	Amatzia Ben-Artzi (Stan) amatzi@spd.3mail.3com.com
NISI (new)	-	May 89	Yes	Karen Bowers (NRI) bowers@sccgate.scc.com Phill Gross (NRI) gross@sccgate.scc.com
NM Ser Interface	-	Yes	Yes	Jeff Case (UTK) case@utkcs2.cs.utk.edu
NOC Tools (new)	-	Yes	Yes	Bob Enger (Contel) enger@sccgate.scc.com

Working Groups	RFC or Draft?	Met Apr 89?	Current Report?	Chair or POC (address)
OSPF	Yes	Yes	Yes	Mike Petry (UMD) petry@trantor.umd.edu John Moy (Proteon) jmoy@proteon.com
Open Systems Routing	Yes	Mar 89	-	Marianne Lepp (BBN) mlepp@bbn.com
OSI Interoperability	No	Yes	Yes	Ross Callon (DEC) callon@erlang.dec.com Rob Hagens (UWISC) hagens@cs.wisc.edu
PDN Routing Group	No	Yes	YES	CH Rokitansky(Fern Univ) roki@isi.edu or roki@dhafeu52.bitnet
Performance and CC	Yes	Yes	Yes	Allison Mankin (MITRE) mankin@gateway.mitre.org
Pt-Pt Protocol	Yes	Yes	Yes	Drew Perkins (CMU) ddp@andrew.cmu.edu Russ Hobby (UC Davis) rdhobby@ucdavis.edu
ST and CO-IP	No	Yes	Yes	Claudio Topolcic (BBN) topolcic@bbn.com
TELNET Linemode	Yes	No	Yes	Dave Borman (Cray) dab@cray.com
User Documents (new)	-	Jun 89	Yes	Karen Roubicek (NSF) roubicek@nnsc.nsf.net Tracey LaQuey (UTexas) tracy@emx.utexas.edu
User Services	No	Yes	Yes	Karen Bowers (NRI) bowers@sccgate.scc.com

Future IETF Meeting Sites

25-28 July 1989	Stanford University
31 Oct - 3 Nov 1989	University of Hawaii
6-9 February 1990	Florida State University
1-4 May 1990	Pittsburgh Supercomputer Center
31 Jul - 3 Aug 1990	University of Washington
November 1990	Princeton
February 1991	OPEN - VOLUNTEERS encouraged!!

II. Chairman's Luncheon

Karen L. Bowers
NRI

The Corporation for National Research Initiatives has a cooperative agreement with the National Science Foundation to provide overall technical guidance to the IETF. This technical guidance includes IETF high-level planning and direction, identification of Internet issues, and organization and guidance to appropriate Working Groups to address these issues. In order to concentrate our efforts more fully on the primary goal of technical guidance, we embarked on a short-term effort to streamline such administrative necessities as Proceedings preparation, quarterly meeting planning, and working group reporting.

A luncheon for Working Group chairs was held on April 12, 1989. The purpose of this session was to provide all Working Group chairs with essential information on how these new procedures would facilitate their technical reporting and distribution of information.

The Chairman's Luncheon briefing outlined the current IETF and Internet-Draft Directory activities and contents; the procedure employed in preparation of the Proceedings and the associated formats for the WG Charter (Form 2), the Status Update (Form 3) and Current Meeting Report; and an overview of activity in progress between the quarterly IETF meetings (slides attached). Two distinct directories, the IETF Directory and the Internet Drafts Directory, will be maintained on line to better facilitate progress of the IETF and to provide public access to information important to IETF members and newcomers alike.

The IETF Directory (to be in place shortly) will consist of files containing: a general IETF description, the Working Group Matrix, meeting dates/locations, current meeting information, a READ ME file with a high level overview of the IETF Directory, and individual Working Group files. Each Working Group will have a file dedicated to its particular activities and will contain a Charter (Form 2), a Status Update (Form 3) and the most Current Meeting Report.

The Internet-Draft Directory (in place now) is a repository of working Internet-Drafts made available for review and comment,in preparation for transition to full RFC status, as appropriate. This directory will soon contain a READ ME file and Index-Abstract to aid the reader in locating files of interest and points of contact for each Internet-Draft. Internet-Drafts are installed as they are made available to the IETF Office and are done so in RFC format, to facilitate document submission to the RFC editor. Eleven Internet Drafts are currently installed inthe Internet Draft Directory; nine are queued up to be installed shortly.

The procedures for preparation of the quarterly IETF Proceedings are currently under revision. The goal is to better capture the accomplishments of the individual Working Groups as well as improve the quality and format of the document itself. This

includes timely distribution of the Proceedings to the Working Group Chairs and IETF meeting attendees. To assist in this process all Working Group Chairs and technical briefers are requested to provide their input as soon as possible upon their return home from the IETF meeting, preferably within the first two weeks. Their early submission will result in an expedited release of the Proceedings to the printers and in turn to WG Chairs, IETF attendees and the like.

Activity between the quarterly meetings is ongoing. WG Chairs hold interim meetings and video teleconferences; WG members continue their work on identified priorities; mid-term meeting reports are submitted to the IETF Chairman; new and revised Internet-Drafts are installed in the Internet-Draft Directory; the agenda for the next IETF quarterly meeting is drafted and finalized through joint participation of all the WG Chairs; and arrangements for the next quarterly IETF meeting are announced and finalized.

This is an iterative process, which when firmly in place, will greatly assist the IETF's primary technical mission.

Karen L. Bowers

The Internet Engineering Task Force

The IETF Directory
(maintaining better information to facilitate progress)

The Internet-Draft Directory
(finally!)

Schedule of Activity between IETF Meetings

Other Issues

IETF Directory Content

- General IETF Description (History, Organization, Goals, etc.)
- WG Matrix (list of current WGs w/status)
- Meeting Dates/Locations (to include previous meetings)
- Current Meeting Information (Agenda/Logistics)
- READ ME File (high-level overview of IETF Directory)
- Individual WG Files (see below)

WG Files

- Charter (Form 2)
- Status Update (Form 3) - NEW!
- Current Meeting Report

CHARTER

(Form 2)

1) Description of WG (in simple terms)
2) Specific Objectives
3) Estimated Timeframe for Completion (of objectives)

(EX: NOC-Tools)

STATUS UPDATE

(Form 3)

1) Chairman's Name(s)/E-Mail Address
2) Name of WG Mailing List(s)
3) Date/Site of Last Meeting
4) Date/Site of Next Meeting
5) Progress to Date (toward completion of objectives)
6) Pending or New Objectives
   (with estimate of timeframe for completion)
7) Documents Produced

CURRENT MEETING REPORT

(to be submitted by WG Chair for each meeting)

1) Agenda
2) List of Attendees
3) Minutes
4) Slides (if reported at Plenary)

CHARTER + STATUS UPDATE +
CURRENT REPORT = PROCEEDINGS ENTRY!

NOC-TOOLS WG

1) CHARTER:

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.

2) OBJECTIVES:

A) IDENTIFY TOOLS AVAILABLE TO ASSIST NETWORK MANAGERS IN DEBUGGING AND MAINTAINING THEIR NETWORKS.

B) PUBLISH A REFERENCE DOCUMENT LISTING WHAT TOOLS ARE AVAILABLE, WHAT THEY DO, AND WHERE THEY CAN BE OBTAINED.

C) ARRANGE FOR THE CENTRAL (OR MULTI-POINT) ARCHIVING OF THESE TOOLS IN ORDER TO INCREASE THEIR AVAILABILITY.

D) ESTABLISH PROCEDURES TO ENSURE THE ONGOING MAINTENANCE OF THE REFERENCE AND THE ARCHIVE, AND IDENTIFY AN ORGANIZATION WILLING TO DO IT.

E) IDENTIFY THE NEED FOR NEW OR IMPROVED TOOLS AS MAY BECOME APPARENT DURING THE COMPILATION OF THE REFERENCE DOCUMENT.

3) TIMEFRAME:

THE FIRST EDITION OF THE CATALOG WILL BE SUBMITTED FOR FINAL REVIEW AT THE OCTOBER-NOVEMBER IETF MEETING. PRELIMINARY VERSIONS WILL BE MADE AVAILABLE EARLIER.

Advertise new I.D.

INTERNET DRAFTS

- Directory Contents

  README
  Index-Abstract
  Individual Internet-Drafts
  — incl. POC for comments

- Rules

  RFC Format
  Naming: DRAFT-<TFNAME>-<WGNAME>-
  <ABBREVITITLE>-<REVNO>-.TXT
  Postscript (specify fonts)

IETF staff action item

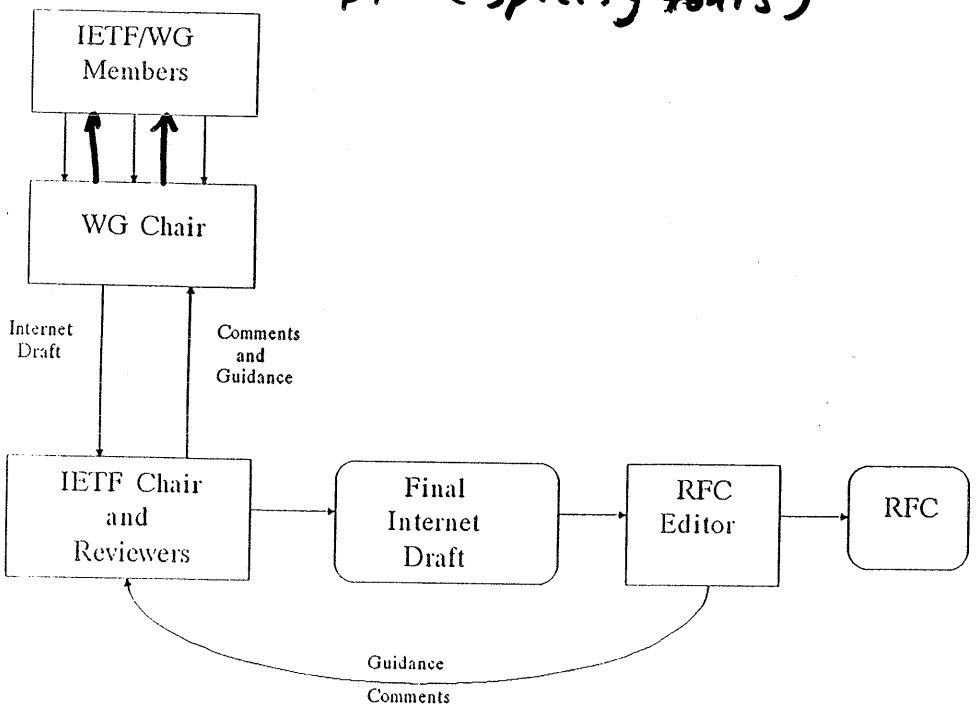

Status of IDEAS:

14 were deleted (either obsolete or already RFCs)

10 are installed in new directory
(some of which will become RFCs)

1 newly installed

9 new documents queued up to be installed shortly

Proceedings Schedule *

Day 1-9	Collect All WG Meeting Reports and TB Summaries-Begin Filling Proceeding Shell
Day 10-21	Proceedings Production
Day 22-28	Proceedings Finalized/Chairman's Review
Day 29-35	Proceeding at Printers (5 days)
Day 36	Proceedings Mailed to Attendees and VIPs (eg. IAB, WG Chairs)

*Compressed schedule

Overview: Activity between 1/4ly IETF Meetings (General Guide)

WK1
- o Request/Receive Wg Meeting Reports and TB Summaries from Recent IETF Plenary
- o Begin Assembling Proceedings

WK2
- o Final Receipt of WG Reports/TB Summaries
- o Prepare Mailing Labels/Attendee Rosters

WK3
- o Proceedings Production
- o Assist WG Chairs w/set up of VTs and mid-term meetings
- o Announce Meeting Info/Logistics/RSVP Requirements for next Plenary

WK4
- o Proceedings Finalized/Chairman's Review
- o IETF Directory Updated w/New WG Reports and Forms 3
- o Assist WG Chairs w/set up of VTs and mid-term meetings

check out PAST

WK5
- o Proceedings Delivered to Printers
- o Request Agenda for next Plenary (for scheduling of WG sessions and technical briefings)

WK6
- o Proceedings Mailed
- o Request Individual WG Agendas for next Plenary
- o Mid-term WG Meetings Taking Place

WK7
- o Announce Preliminary Agenda for next Plenary
- o Collect Mid-term WG Meeting Reports (to be included in next Proceedings)

WK8
- o Install New Internet Drafts or Revisions (as required)

WK9
- o Collect All Individual WG Agendas for next Plenary
- o Release Final Agenda for next Plenary

WK10
- o Release Individual Agendas to IETF Mailing List
- o Consolidate (RSVP) Attendee Roster for next Plenary

WK11
- o Prepare for Upcoming IETF Plenary
- o Finalize All Meeting/Logistics Requirements

WK12
- o Next IETF Plenary

Other Issues

Inquiry: Technical Details for Configuring Mailers?

Future Meetings: Individual WG Agendas will be Requested in Advance

Proceedings: Courtesy Copies vs. $35 Non-Attendee Document Charge

Directories: Need a Location Volunteered to Shadow IETF and Internet-Draft Directories
(East Coast Preferred)

Coffee Breaks - THE GONG?

Discuss - 1 day WG
1 day TB
1 day WG
1/2 day WG Reports

AV Requirements

If Chairs need other than an OH projector for their meetings please let ACE know in advance

Rosters

ACE will always provide blank rosters to be signed by WG attendees and returned to ACE staff for processing. In turn, ACE will provide within 1-2 weeks both a master attendee list (name, affiliation, address, phone, email) and a copy of the roster from their respective WG (name, email)

III. IETF Attendees

Mohammad Alaghebandan

3Com Corp.
PO Box 832
Cupertino, CA 95015

415-940-7679 mra@bridge2.3com.com

Guy Almes

Rice University
PO Box 1892
Dept of Computer Science
Houston, TX 77251-1892

713-527-6038 almes@rice.edu

Philip Almquist

Stanford University
285 Clinton Park
San Francisco, CA 94103

415-552-0588 almquist@jessica.stanford.edu

Cathy Aronson

Merit/ CICNet
1075 Beal Avenue
Ann Arbor, MI 48109-2112

313-936-2090 cja@merit.edu

Amatzia Ben-Artzi

3Com Corp.
3165 Kifer Road
Santa Clara, CA 95052

408-733-4823 amatzia@spd.3mail.3com.com

Art Berggreen

ACC
720 Santa Barbara St
Santa Barbara, CA 93101

805-963-9431 art@sage.acc.com

April 1989 - IETF Meeting Attendee List

Rick Boivie

IBM
472 Wheelers Farms Road
Milford, CT 06460

203-783-7284 rboivie@ibm.com

Dave Borman

Cray Research
1440 Northland Dr.
Mendota Heights, MN 55120

612-681-3398 dab%oliver.cray.com@uc.msc.umn.edu

Leonard Bosack

cisco Systems
1360 Willow Road
Menlo Park, CA 94025

415-326-1941 bosack@mathom.cisco.com

Karen Bowers

National Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091

703-620-8990 bowers@sccgate.scc.com

David Bridgham

FTP Software
26 Princess Street
Wakefield, MA 01880

617-246-0900 dab@ftp.com

Scott Brim

Cornell University
265 Olin Hall
Cornell Theory Center
Ithaca, NY 14853

607-255-8686 swb@devvax.tn.cornell.edu

Jeffrey Burgan

NASA Ames Research Center
M/S 233-18
Moffet Field, CA 94035

415-694-6440 jeff@nsipo.nasa.com

Ross Callon

Digital Equipment Corporation
550 King Street
LKG1-2/A19
Littleton, MA 01460-1289

508-486-5009 callon@erlang.dec.com

Jeffrey Case

University of Tennessee
Computing Center
200 Stokely Management Center
Knoxville, TN 37996-0520

615-974-6721 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 Chao

BBN Communications Corp.
50 Moulton Street
Cambridge, MA 02138

617-873-3888 jchao@bbn.com

J. Noel Chiappa

MIT/Proteon
Two Technology Drive
Westborough, MA 01581-5008

617-898-2800 jnc@lcs.mit.edu

April 1989 - IETF Meeting Attendee List

Joseph Choy

NCAR
PO Box 3000
Boulder, CO 80307-3000

303-497-1222 choy@ncar.ucar.edu

Danny Cohen

University of Southern California
4676 Admiralty Way
ISI
Marina del Rey, CA 90272

213-822-1511 cohen@isi.edu

Michael Collins

Lawrence Livermore Natl Lab
PO Box 5509 L-561
Livermore, CA 94550

415-422-4018 collins@nmfecc.arpa

Rob Coltun

University of Maryland
Computer Science Dept.
College Park, MD 20742-2411

301-454-2946 rcoltun@trantor.umd.edu

John L. Cook

Chipcom Corp.
195 Bear Hill Rd.
Waltham, MA 02154

617-890-6844 cook@chipcom.com

Jerry Cronin

USAF
1842 EEG/EEMSI
Scott AFB, IL 62225-6348

618-256-2658 1842eeg-intg2@afcc-oal.arpa

James Davin

MIT LCS
Computer Science Lab, NE43-507
545 Technology Square
Cambridge, MA 02139

617-253-6020 jrd@ptt.lcs.mit.edu

Farokh Deboo

3Com
2081 North Shoreline Blvd.
Mountain View, CA 94043

415-940-7677 ..!sun!bridge2!fjd

Phil Draughon

Northwestern University
Academic Computing & Network Services
2129 Sheridan Road
Evanston, IL 60208

312-491-4070 jpd@accuvax.nwu.edu

Ralph Droms

Bucknell University
323 Dana Engineering
Lewisburg, PA 17837

717-524-1145 droms@cs.purdue.edu

Charles Eldridge

SPARTA, Inc.
7926 Jones Branch Drive
Suite 1070
McLean, VA 22102

703-448-1683 eldridge@sparta.com

Robert Enger

CONTEL Federal Systems
1300 Quince Orchard Blvd
Gaithersburg, MD 20878-4199

301-840-4040 enger@sccgate.scc.com

April 1989 - IETF Meeting Attendee List

Hunaid Engineer

Cray Research, Inc.
1400 Northland Drive
Mendota Heights, MN 55344

612-681-3015 hunaid@cray.com

Dino Farinacci

3Com
2081 N. Shoreline Blvd.
Mountain View, CA 94043

415-940-7661 dino@bridge2.3com.com

Mark Fedor

NYSERNet Inc
Rensselear Technology Park
165 Jordan Road
Troy, NY 12180

518-283-8860 fedor@nisc.nyser.net

Craig Fox

Network Systems Corp
7600 Boone Avenue North
Brooklyn Park, MN 55428

612-424-4888 foxcj@nsco.network.com

Jose Garcia-Luna

SRI International
Telecommunications Sciences Ct
333 Ravenswood Avenue
Menlo Park, CA 94025

415-859-5647 garcia@sri.com

Lionel Geretz

ACC
720 Santa Barbara Street
Santa Barbara, CA 93101

805-963-9431 lionel@salt.acc.com

Elise Gerich

Merit Computer Network
University of Michigan
1075 Beal Avenue
Ann Arbor, MI 48109-2112

1-800-666-3748 epg@merit.edu

Chuck Gerlach

AT&T
1100 East Warrenville Rd
Naperville, IL 60566-7044

312-979-7325 cag@iwlcs.att.com

Robert Gilligan

Sun Microsystems
2550 Garcia Ave.
Mountain View, CA 94043

415-336-1012

Jim Greuel

Hewlett-Packard
3404 E. Harmony Road
Fort Collins, CO 80525

303-229-2493 jimg%hpcndpc@hplabs.hp.com

Martin Gross

DCA/DCEC
1861 Wiehle Avenue
Reston, VA 22090

703-437-2165 martin@protolaba.dca.mil

Phill Gross

National Research Initiatives
1895 Preston White Dr. Ste.100
Reston, VA 22091

703-620-8990 gross@sccgate.scc.com

April 1989 - IETF Meeting Attendee List

Rob Hagens

University of Wisconsin
Computer Science Dept.
1210 W. Dayton Street
Madison, WI 53706

608-262-1017 hagens@cs.wisc.edu

Tony Hain

Lawrence Livermore Natl Lab
PO Box 5509, L-561
Livermore, CA 94550

415-422-4200 hain@nmfecc.arpa

Brian Handspicker

Digital Equipment Corp.
550 King St.
Littleton, MA 01460

508-486-7894 bd@vines.dec.com

Bob Harris

BBN Communications
50 Moulton St.
Cambridge, MA 02238

617-873-1817 bharris@bbn.com

Robert T. Harris

SPARTA, Inc.
7926 Jones Branch Dr.
McLean, VA 22102

703-448-0210 harris@sparta.com

Gene Hastings

Pittsburgh SCC
5000 Forbes Ave.
Pittsburgh, PA 15213

412-268-4960 hastings@morgul.psc.edu

Russell Hobby

University of California
Computing Services
Surge II - Room 1400
Davis, CA 95616

916-752-0236 rdhobby@ucdavis.edu

Jeffrey Honig

Cornell Theory Center
Theory Center
265 Olin Hall
Ithaca, NY 14853-5201

607-255-8686 jch@tcgould.tn.cornell.edu

Steven Hunter

Lawrence Livermore Natl Lab
PO Box 5509, L-561
NMFECC
Livermore, CA 94550

415-423-2219 hunter@nmfecc.arpa

Tom Hytry

AT&T Bell Labs
1F-313
1100 E. Warrenville
Naperville, IL 60566

312-983-5058 att!iwles!tlh

Ole-Jørgen Jacobsen

Connexions, ACE
480 San Antonio Rd.
Suite 100
Mountain View, CA 94040

415-941-3399 ole@csli.stanford.edu

Van Jacobson

Lawrence Berkeley Lab
One Cycloctron Road
M/S 46A
Berkeley, CA 94720

415-486-6411 van@helios.ee.lbl.gov

April 1989 - IETF Meeting Attendee List

Michael Karels

University of California
1107 Liberty Street-EECS
CSRG, CS Division
Berkeley, CA 94720

415-642-4948 karels@berkeley.edu

Steve Knight

Cray Research, Inc.
1440 Northland Drive
Mendota Heights, MN 55120-1095

612-681-3124 knight@cray.com

Lee LaBarre

Mitre Corporation
Burlington Rd
M/S E066
Bedford, MA 01730

617-271-8507 cel@mitre.org

Tracy LaQuey

University of Texas
Computation Center
Austin, TX 78712

512-471-3241 tracy@emx.utexas.edu

John Lekashman

NASA Ames Research Center
MS 258-6
Moffett Field, CA 94035

415-694-4359 lekash@orville.nas.nasa.gov

Marianne Lepp

BBN Communications
50 Moulton St.
Cambridge, MA 02133

617-873-2458 mlepp@bbn.com

Norbert Leser

Open Software Foundation
11 Cambridge Center
Cambridge, MA 02142

617-621-8715 nl@osf.org

Mike Little

SAIC
8619 Westwood Center Dr.
Vienna, VA 22182

703-749-5360 little@saic.com

Mark Lottor

SRI International
333 Ravenswood Ave
EJ282
Menlo Park, CA 94025

415-859-2652 mkl@sri-nic.arpa

Paul Love

San Diego Supercomputer Center
PO Box 85608
San Diego, CA 92138-5608

619-534-5043 loveep@sds.sdsc.edu

Charles Lynn, Jr.

BBN STC
10 Moulton Street
Cambridge, MA 02238

617-873-3367 clynn@bbn.com

April 1989 - IETF Meeting Attendee List

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

Allison Mankin

Mitre Corporation
7525 Colshire Drive
McLean, VA 22102

703-883-7907 mankin@gateway.mitre.org

Matt Mathis

Pittsburgh SCC
4400 5th Ave.
Pittsburgh, PA 15213

412-268-3319 mathis@fornax.ece.cmu.edu

Keith McCloghrie

Wollongong Group
1129 San Antonio Rd
Palo Alto, CA 94303

415-962-7160 kzm@twg.com

Milo Medin

Sterling Software
1121 San Antonio Road
Palo Alto, CA 943036

415-694-6440 medin@nsipo.nasa.gov

Donald Merritt

Ballistic Research Laboratory
Attn: AMXBR-SECAD
Aberdeen Proving Grounds, MD 21005-5066

301-278-6808 merritt@brl.mil

Donald Morris

NCAR
POB 3000
Scientific Computing Division
Boulder, CO 80307

303-497-1000 morris@ncar.ucar.edu

John Moy

Proteon, Inc.
Two Technology Drive
Westborough, MA 01581-5008

508-898-2800 jmoy@proteon.com

Russ Mundy

DCA
Code B602
Washington, DC 20305-2000

703-285-5481 mundy@beast.ddn.mil

Ronald Natalie

Rutgers University
CCIS PO Box 879
Busch Campus-Hill Center
Piscataway, NJ 08855-0879

201-214-0832 ron@rutgers.edu

Rebecca Nitzan

Lawrence Livermore Natl Lab
L-561
PO Box 5301
Livermore, CA 94550

415-422-9775 nitzan@nmfecc.llnl.gov

Bill Norton

Merit Computer Network
1075 Beal Ave.
Ann Arbor, MI 48104

313-764-9423 wbn@merit.edu

Bill Nowicki

Sun Microsystems
2550 Garcia Avenue
M/S 14-49
Mountain View, CA 94043

415-960-1300 nowicki@sun.com

Zbigniew Opalka

BBN Communications
50 Moulton St.
Cambridge, MA 02238

617-873-2888 zopalka@bbn.com

Philippe Park

BBN STC
10 Moulton Street
Cambridge, MA 02238

617-873-2892 ppark@wilma.bbn.com

Drew Perkins

Carnegie Mellon University
4910 Forbes Ave
Pittsburgh, PA 15213

412-268-8576 ddp@andrew.cmu.edu

Michael Petry

University of Maryland
Computer Science Center
Room 3339
College Park, MD 20742

301-454-2943 petry@trantor.umd.edu

Paul Pomes

University of Illinois
1304 W. Springfield Ave.
Urbana, IL 61801-2987

217-333-6262 paul@uxc.cso.uiuc.edu

Rex Pugh

Hewlett-Packard
8000 Foothill Blvd
R3NF3
Roseville, CA 95678

916-785-4471 pugh%hprnd@hplabs.hp.com

K.K. Ramakrishnan

Digital Equipment Corporation
LKG 1-2/A19
550 King Street
Littleton, MA 01460-1289

508-486-7267 rama%erlang.dec.com@decwrl.dec.com

Jacob Rekhter

IBM
TJ Watson Research Center
Route 134 PO Box 218
Yorktown Heights, NY 10598

914-945-3896 yakov@ibm.com

Joel Replogle

NCSA
605 E Springfield
Champaign, IL 61820

217-244-0636 jr@ncsa.uiuc.edu

April 1989 - IETF Meeting Attendee List

Robert Reschly

Ballistic Research Laboratory
ATTN: SLCBR-SF (Reschley)
Aberdeen Proving Ground, MD 21005-5066

301-278-6678 reschly@brl.mil

Joyce Reynolds

University of Southern California
4676 Admiralty Wy #1001
ISI
Marina del Rey, CA 90292-6695

213-822-1511 jkrey@venera.isi.edu

Carl Rokitansky

Fern Univ. of Hagen,FB ET/DVT
Frauenstuhlweg 31
D-5860 ISERLOHN, WEST GERMANY

roki@dhafeu52.bitnet or roki@a.isi.edu
roki%dhafeu52.bitnet@cunyvm.cuny.edu

Milt Roselinsky

Communication Machinery Corporation
125 Cremona Drive
Santa Barbara, CA 93117

805-562-3194 cmcvax!milt@hub.ucsb.edu

Karen Roubicek

NSF Network Service Center
10 Moulton St.
Cambridge, MA 02138

617-873-3361 roubicek@nnsc.nsf.net

Greg Satz

cisco Systems
1360 Willow Road
Suite 201
Menlo Park, CA 94025

415-326-1941 satz@cisco.com

Jeffrey Schiller

MIT
1 Amherst Street
Room E40-311
Cambridge, MA 02139

617-253-4101 jis@bitsy.mit.edu

Bruce Schofield

DCEC
1860 Wiehle Avenue
Reston, VA 22090

703-437-2556 schofield@edn-unix.arpa

John Scott

Data General
62 TW Alexander Drive
Research Triangle Park, NC 27709

919-248-5995 scott@dg-rtp.dg.com

Dallas Scott

MITRE Corporation
7525 Colshire Drive
M/S W420
McLean, VA 22102

703-883-6700 dscott@gateway.mitre.org

Rajeev Seth

Hewlett-Packard
19420 Homestead Rd
MS 43LH
Cupertino, CA 95014

408-447-3573 rajs%hpindbu@hplabs.hp.com

April 1989 - IETF Meeting Attendee List

Jim Sheridan

IBM
PO Box 334
Whitmore Lake, MI 48189

313-393-6537 jsherida@ibm.com

Michael St. Johns

DCA
Code B612
Washington, DC 20305-2000

703-285-5133 stjohns@beast.ddn.mil

Mary Stahl

SRI-NIC
Net Info Sys Ctr
333 Ravenswood Ave. Rm EJ 296
Menlo Park, CA 94025

415- 859-4775 stahl@sri-nic.arpa

Louis Steinberg

IBM
472 Wheelers Farms Rd
M/S 91
Milford, CT 06460

203-783-7175 louiss@ibm.com

Robert Stine

SPARTA, Inc
7926 Jones Branch Dr.
Suite 1070
McLean, VA 22102

703-448-0210 stine@sparta.com ~~guest@edn-unix.dca.mil~~

Zaw-Sing Su

SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025

415-859-4576 zsu@tsca.istc.sri.com

IV. Final Agenda

Agenda for the April 11-14 IETF Meeting

TUESDAY, APRIL 11th

9:00 am Opening Plenary, Introductions and Local Arrangements
Phill Gross (NRI)

9:15 am Morning Working Group Sessions

- OSPFIGP (Petry, UMD and Moy, Proteon)
- Network Management Services Interface (Case, UTK and McCloghrie, TWG)
- OSI Interoperation (Callon, DEC and Hagens, UWisc)
- Performance and Congestion Control, TCP Subgroup (Mankin, Mitre)
- Point-Point Protocol (Perkins, CMU and Hobby, (UCDavis)
- User Services (Bowers, NRI)

12:30 pm Lunch Break

1:30 pm Afternoon Working Group Sessions

- Authentication (Schiller, MIT and Rochlis, MIT)
- LANMAN (Ben-Artzi, 3Com)
- OSI Interoperation (Callon, DEC and Hagens, UWisc)
- Performance and Congestion Control (Mankin, Mitre) (Open Meeting, but attendees are expected to have reviewed, and prepared comments on, the draft Gateway Congestion Control paper. Send to mankin@gateway.mitre.org for a copy.)
- Point-Point Protocol (Perkins, CMU and Hobby, (UCDavis)
- User Services (Bowers, NRI)

5:00 pm Recess (for the those not attending the DMS WG session)

5:00 pm Domain Name System WG (convened by Drew Perkins, CMU) - 6:30 pm

WEDNESDAY, APRIL 12th

9:00 am Opening Plenary
9:15 am Morning Working Group Sessions

- o NOC Tools (Enger, Contel and Stine, Sparta)
- o Joint Interconnectivity and Open Routing WGs (Almes, Rice and Lepp, BBN)
- o Public Data Network Routing (Rokitanski, FERN) (Open meeting)
- o Performance and Congestion Control (Mankin, Mitre) (Editing session for WG members only)
- o ST and Connection IP (Topolcic, BBN)
- o ALERTMAN (Louis Steinberg, IBM)

12:15 pm Lunch Break

12:30 pm Working Lunch of the WG Chairs

- o IETF Office Update (Bowers, NRI)

1:30 pm Afternoon Working Group Sessions

- o Host Dynamic Configuration (Droms, Bucknell and P.Gross, NRI)
- o Interconnectivity (Almes, Rice)
- o Public Data Network Routing (Rokitanski, FERN) (Members Only)
- o Performance and Congestion Control (Mankin, Mitre) (Editing session for WG members only)
- o ST and Connection IP (Topolcic, BBN)

5:00 pm Recess

7:30 pm
- o Joint Monitoring Access for (NSFNET) Adjacent Networks (Gerich, Merit)

THURSDAY, APRIL 13th

9:00 am Opening Plenary

9:10 am Everything You Ever Wanted to Know about OSPFIGP (including how to pronounce it) (Moy, Proteon)

10:00 am The Open Routing Architecture (Lepp, BBN)
10:45 am Break
11:00 am Report on the NASA Science Internet (Medin, Ames)
11:50 am State of the Internet (Opalka, BBN)
12:10 pm Growth of the Internet (St. Johns, B600)

12:30 noon Lunch Break

1:45 pm Report on the DOE Energy Science Network (ESNET) (Hain, LBL)
2:00 pm An Interim Routing Architecture (Mundy, DCA B600)
2:15 pm NSFNET Report
- o Architectural Changes to NSFNET (Gerich, MERIT)
- o Nifty NSFNET Stats, using NNStat (Gerich, MERIT)

2:45 pm Interim Routing Architecture (an Alternative View) (Rekhter, IBM)
2:55 pm Arpanet Evaporation Timetable and An Overview of FRICC Initiatives (eg, the NNT, RIB, and RIG) (P.Gross, NRI)

3:05 pm Mailbridge Access Control (Lepp, BBN)
3:15 pm Authentication WG Report (Schiller, MIT)
3:30 pm Break
3:45 pm Cornell Worm Report Commentary (Schiller, MIT)
3:55 pm The DCA TCP/IP Certification Program (M. Gross, DCA-DCEC)
4:15 pm Header Compression for TCP/IP Datagrams (Van Jacobson, LBL)
5:00 pm Recess

FRIDAY, APRIL 14th

9:00 am Opening Plenary
9:15 am Working Group Reports and Discussion

- Network Management Services Interface (Case, UTK and McCloghrie, TWG)
- LANMAN (Ben-Artzi, 3Com) (Hares, Merit)
- OSI Interoperation (Callon, DEC and Hagens, UWisc)
- Joint Monitoring Access for (NSFNET) Adjacent Networks
- User Services (Bowers, NRI)
- Performance and Congestion Control (Mankin, Mitre)
- Point-Point Protocol (Perkins, CMU and Hobby, UCDavis)
- ST and Connection IP (Topolcic, BBN)

10:30 am Break

10:45 am Working Group Reports and Discussion

- ALERTMAN (Steinberg, IBM)
- Host Dynamic Configuration (Droms, Bucknell and P.Gross, NRI)
- NOC Tools (Enger, Contel and Stine, Sparta)
- Domain Name System (Perkins, CMU)
- Public Data Network Routing (Rokitanski, FERN)

11:30 pm Concluding Plenary Remarks and Group Discussion

12:00 pm Adjourn

V. Working Group Summaries

- Charters
- Status Updates
- Current Meeting Reports
- Slides Presented at IETF (if available)

Alert Management Working Group
Chairperson: Louis Steinberg/IBM

CHARTER

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 IETF meeting, with final review expected at the October IETF meeting. The second RFC draft will be submitted for initial review at the October IETF meeting. A date for final review of this document has not yet been determined.

Alert Management Working Group
Chairperson: Louis Steinberg/IBM

STATUS UPDATE

1. Chairperson: Lou Steinberg, louiss@ibm.com

2. WG Mailing list: alert-man@merit.edu and
alert-man-request@merit.edu

3. Last Meeting (and first): Cocoa Beach, FL, April 12, 1989

4. Next Meeting: Stanford, July 25-28, 1989

5. Progress to date: Initial review of topics, defining what we are attempting to accomplish.

5. Pending or new objectives: see objectives in Charter

6. Progress to date (e.g., documents produced):

   - o initial review of topics and defining what we are attempting to accomplish

Alert Management Working Group
Chairperson: Louis Steinberg/IBM

CURRENT MEETING REPORT
Reported by Louis Steinberg

AGENDA

a. Introduction
b. Discussion of group's charter and goals
a. Chair's action items

1) establish mailing list
2) write first draft of "information flow management" document

ATTENDEES

Cathy Aronson	cja@merit.edu
Amatzia Ben-Artzi	amatzia@spd.3+.3com.com
Jeff Case	case@utkux1.utk.edu
John Chao	jchao@bbn.com
John Cook	cook@chipcom.com
Chuck Davin	jrd@ptt.lcs.mit.edu
Mark Fedor	fedor@nisc.nyser.net
Lionel Geretz	lionel@salt.acc.com
Bob Harris	bharris@bbn.com
Steven Hunter	hunter@nmfecc.llnl.gov
Tom Hytry	tlh@iwles@att
Lee Labarre	cel@mbunix.mitre.org
Charles Lynn	clynn@bbn.com
Keith McCloghrie	kzm@twg.com
Bill Norton	wbn@merit.edu
Joel Replogle	jr@ncsa.uiuc.edu
Greg Satz	satz@cisco.com
Bruce J. Schofield	schofield@edn-vax.dca.mil
John Scott	scott@dg-rtp.dg.com
Jim Sheridan	jsherida@ibm.com
Robert Stine	stine@sparta.com
Steve Waldbusser	sw01@andrew.cmu.edu
Dan Wintringham	danw@igloo.osc.edu

MINUTES

The first meeting of the Alert Management Working Group began with an introduction from the Chairman (Lou Steinberg).

A discussion of the goals of this group then followed. It was decided that the output of this group must take great care to not impact the letter or spirit of either the SNMP or CMOT RFCs. Each document produced will demonstrate the use of proposed alert management techniques in a manner conferment with both SNMP and CMOT.

Several divisions were proposed for the work to be done. Lou discussed his interpretation, in which he focused on (1) the flow of previously generated alerts and (2) managed MIB objects to generate them. Examples of each were briefly cited, as Lou has already coded and tested these ideas.

Jeff Case expressed strong concern that the Working Group not focus on managed MIB objects to generate alerts. The feeling of many SNMP implementors was that this would violate the philosophy of SNMP, which uses the protocol (rather than managed objects) to generate traps. While some SNMP users may be using such objects in the experimental space of the MIB, it is inconsistent to define such variables. Doing so (even with "optional" MIB objects), would make the MIB appear to be slanted towards use with CMOT.

Lee LaBarre presented the CMIP view of Events, and the areas that ISO looks at to manage them. This was basically a superset of Lou's view.

The desire to develop techniques fully compatible with both SNMP and CMOT led to a decision that two RFCs would be submitted. The first would deal with managing the flow of information caused by asynch. generated alerts. The second (with a requested status of "optional") would discuss techniques for generating such alerts that are self-limiting; those that do not allow an excess of alerts to be generated.

Lou agreed to set up a mailing list for the group. He did not think that ibm.com was available, and agreed to look for an alternative (SNMP @ nisc.nyser.net was suggested for a start). Lou also took the action item to write up a draft of the first ("flow") document, describing some of the systems he has used. This will be posted to the list for initial review.

Authentication Working Group
Chairperson: Jeffrey Schiller/MIT

CHARTER

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 with either disband or will move on to other related problems/issues.

Authentication Working Group
Chairperson: Jeffrey Schiller/MIT

STATUS UPDATE

1. Chairperson: Jeffrey Schiller, jis@bitsy.mit.edu

2. WG Mailing list: AWG@BITSY.MIT.EDU

3. Last Meeting: Cocoa Beach April 1989

4. Next Meeting: To Be Scheduled

5. Pending or New Objectives:

6. Progress to Date (e.g., documents produced):

   A draft RFC was circulated at the last meeting to proposing a standard authentication format for multiple protocols (addresses object [1] above).

   A draft "Authentication Requirements" document was also circulated. This is the beginning of an effort that will lead to writing a protocol designers guide to authentication (addresses objective [3] above).

   A draft RFC for the Kerberos Authentication system was circulated as well (addresses objective [4] above).

Authentication Working Group
Chairperson: Jeffrey Schiller/MIT

CURRENT MEETING REPORT
Reported by Jeffrey Schiller

ATTENDEES

Danny Cohen	cohen@isi.edu
John Cook	cook@chipcom.com
Charles Eldridge	eldridge@sparta.com
Hunaid Engineer	hunaid@hall.cray.com
Phill Gross	gross@sccgate.scc.com
Mike Karels	karels@berkeley.edu
Steve Knight	knight@baldmt.cray.com
Lee LaBarre	cel@mbunix.mitre.org
Norbert Leser	nl@osf.org
Louis Mamakos	louie@trantor.umd.edu
Don Merritt	merritt@brl.mil
John Moy	jmoy@proteon.com
Russ Mundy	mundy @beast.ddn.mil
Jeff Schiller	jis@bitsy.mit.edu
Mike St. Johns	stjohns@beast.ddn.mil
Ross Veach	rrv@uxc.cso.uiuc.edu
Steve Waldbusser	sw01@andrew.cmu.edu

MINUTES

Three handouts were distributed at the beginning of the meeting.

1. First Draft of an Authentication Requirements document currently being authored by Jon Rochlis.
2. A copy of a Draft RFC for Version 4 of the Kerberos Authentication System authored by Jennifer Steiner.
3. A copy of a Draft RFC for a new IP option for IP level authentication by Jeffrey Schiller.

Most of the discussion at the meeting was on the IP option draft paper. This paper proposes the creation of a new IP option for carrying a cryptographic checksum of selected portions of the packet's IP header and the entire data contents of the packet. Its primary use would be as a mechanism to permit the addition of authentication to already existing protocols that currently have no provision for carrying authentication information within the protocol's own data contents.

The members of the working group discussed the document and proposed several modifications. The meeting came to a consensus on the modifications.

Action Items:

Jeffrey Schiller will put together another version of the IP option document and will distribute it to the members for consideration.

Authentication Working Group

IP option for network auth code.

* need to add authentication to existing protocols
* no change to protocol specific packet format
* handles different types of checksum computations, variable data length.

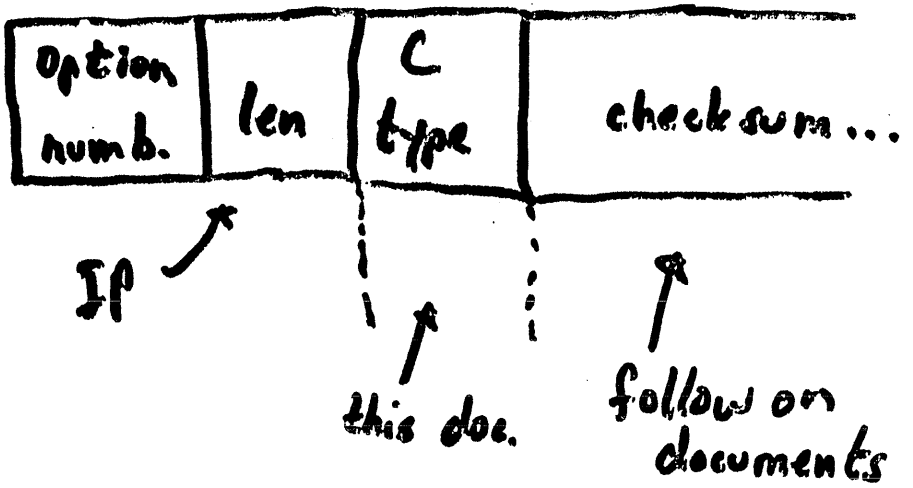

* protects packet data and part of IP header

source addr.
dst addr.
IP protocol.

what we do <u>not</u> do:

* privacy: packet is not encrypted
* replay protection: requires more state, may not always be oppropriate.

CMIP-over-TCP (CMOT) Working Group
Chairperson: Lee LaBarre/Mitre

CHARTER

Description of Working Group:

- Develop a long term approach to management of the Internet based on the OSI Network Management Framework and the Common Management Information Protocol (CMIP).

- Provide input to the OSI standards process based on experience in the Internet, and thereby influence the final form of OSI International Standards on network management, in particular CMIS/P.

Specific Objectives:

a) Develop prototype implementors agreements on CMIP over TCP.
b) Develop prototype implementations based on the CMOT agreements and IETF SMI and MIB agreements.
c) Experiment with CMOT and extensions to the SMI and MIB.
d) Develop final implementors agreements for CMOT.
e) Promote development of products based on CMOT.
f) Provide input to the OSI Network Management standards process in time to effect the International Standards.

Estimated Timeframe for Completion:

The group's work should be completed by June 1989.

CMIP-over-TCP (CMOT) Working Group
Chairperson: Lee LaBarre/Mitre

STATUS UPDATE

1. Chairperson: Lee LaBarre, cel@mitre.org.com

2. WG Mailing List: netman@gateway.mitre.org

3. Last meeting: 19 January, 1989, Austin, Texas

4. Next Meeting: TBD as required

5. Pending or New Objectives:

   The remaining tasks for the group include:

   - updating the specification when the OSI standards reach international standard (IS) status,
   - specification of event generation and event report control mechanisms.

   The latter task has moved to a subgroup of the MIB WG. However,if it is decided that generic event generation and report control mechanisms are not desired, then this group will address the problem.

6. Progress to date (e.g., documents produced):

   - RFC1095, "The Common Management Information Services and Protocol over TCP/IP (CMOT", edited by U. Warrier and L. Besaw

   - The group has completed a major portion of its charter to develop a long term approach to network management, namely an specification of an architecture and protocol that is consistent with OSI and will facilitate management of future networks containing TCP/IP and OSI components. That specification, contained in RFC1095, is based on the DIS version of CMIP, and on Internet RFCs. The RFC1095 and the new SNMP RFC1098 have been given equal status by the IAB.

CURRENT MEETING REPORT

None

Domain Working Group
Chairperson: Paul Mockapetris/USC/ISI

CHARTER

Description of Working Group:

The goal of the Domain Working Group is to advise on the administration of the top levels of the DNS ("the root servers"), consider proposed extensions and additions to the DNS structure and data types, and resolve operational problems as they occur.

Specific Objectives:

The specific short-term objectives are:

1. Adding load balancing capability to the DNS.
2. Adding DNS variables to the MIB.
3. Implementation catalog for DNS software.
4. Responsible Person Record.
5. Adding network naming capability to the DNS.
6. Evaluate short term measures to improve, or at least describe the security of the DNS.

Estimated Timeframe for Completion (for above objectives):

1. The preferred method for Load Balancing was decided upon at the April '89 IETF meeting at Cocoa Beach. A short RFC will be written before the next meeting in July '89.

2. End of 1989

3. Questionaire sent, responses data being organized, summary and detail to appear. (PVM)

4. July '89.

5. RFC issued April 89, implementations to follow.

Domain Working Group
Chairperson: Paul Mockapetris/USC/ISI

STATUS UPDATE

1. Chairpersons: Permanent - Paul Mockapetris (pvm@isi.edu)
   Temporary - Drew Perkins (ddp@andrew.cmu.edu)

2. WG Mailing Lists(s): namedroppers@sri-nic.arpa

3. Date of Last Meeting: April '89, Cocoa Beach

4. Date of Next Meeting: July '89, Stanford University

5. Pending or New Objectives: see Charter

6. Progess to Date (e.g., documents produced):

   - RFC 1101 - on Network Name Mapping
   - Advice to Internet Host Requirements Editor

Domain Working Group
Chairperson: Paul Mockapetris/USC/ISI

CURRENT MEETING REPORT
Reported by Drew Perkins

AGENDA

1. What should the DWG suggest to the Host Requirements WG.

2. How do DNS processes appear in the MIB.

3. Policy on load balancing.

4. Addition of dynamic add and delete to the DNS.

5. Firm up the rules for defining new types and classes, and the interpretation of wildcards.

6. Implementation catalog for DNS software.

7. A test/validation suite for the DNS.

8. Enhancements to the DNS in general.

ATTENDEES

Momhammad Alaghebandan	mra@bridge2.3com.com
Philip Almquist	almquist@jessica.stanford.edu
Cathy Aronson	cja@merit.edu
Dave Borman	dab@cray.com
Mike Collins	collins@ccc.mfecc.llnl.gov
Mark Fedor	fedor@nisc.nyser.net
Jose Garcia-Luna	garcia@sri.com
Elise Gerich	epg@merit.edu
Mike Karels	karels@berkeley.edu
Steve Knight	knight@baldmt.cray.com
Tracy LaQuey	tracy@emx.utexas.edu
Mark Lotter	mlk@sri-nic.arpa
Paul Love	loveep@sds.sdsc.edu
Russ Mundy	mundy@beast.ddn.mil
Bill Norton	wbn@merit.edu
Bill Nowicki	nowicki@sun.com
Drew Perkins	ddp@andrew.cmu.edu
Rex Pugh	pugh%hprnd@hplabs.hp.com
Mary Stahl	stahl@sri-nic.arpa
Mike St. Johns	stjohns@beast.ddn.mil
Zaw-Sing Su	zsu@sri.com
Paul Tsuchiya	tsuchiya@gateway.mitre.org

MINUTES

The Domain WG met at the April '89 IETF in Cocoa Beach, Fla. Since Paul Mockapetris (ISI), the permanent chairman of the Domain WG, was unable to attend the meeting, Drew Perkins (CMU) filled in as temporary chairman.

The group quickly decided that no one in attendance had any strong desires to talk about any of the issues with the exception of Load Balancing. This happened to be the main concern of the temporary chairman, so the rest of the meeting was spent discussing it.

The goal of "load balancing" is to use the DNS as a tool to dynamically balance the load across some number of servers. The particular situation is as follows. There are a number of server machines HOST1, HOST2, HOST3, etc. Each of these machines is connected to a network file system and appear identical to users (with the exception of the host name of course). Users would like to simply say "TELNET HOST" and be connected to the least loaded host. It was decided that there were a number of ways of accomplishing this, most of them using the DNS.

1. The TELNET application could use some protocol to find out the load across all servers whenever a user wanted to connect to a server. It could then pick the least loaded system. This of course has the disadvantage that every TELNET application must be modified. Therefore this possibility was rejected.

2. A new DNS Load Balancing Resource Record (LB) could be defined. This record could be similar to the MX record for mail. There would be one record for each system. These records would be continuously given new preference values, and would always have a small TTL, preferably zero. Again, this choice would require modifying every implementation of TELNET, so it was rejected.

3. A single CNAME RR could be used to dynamically alias HOST to HOSTn. This RR would always have a small TTL (zero) and would be changed dynamically to reflect the least loaded machine. For example, at first HOST1 may be the least loaded, so there would be an RR "HOST 0 IN CNAME HOST1". If the load on HOST1 increased so that HOST2 became the least loaded, then this RR would be be removed and a new RR would be added: "HOST 0 IN CNAME HOST2".

4. A dynamically sorted list of A RRs could be used. The domain "HOST" could include the same A RRs as HOST1, HOST2, etc. These RRs would be sent with small TTLs and would be resorted as the load on each machine changed.

The group decided that option 3 was the most preferable and was the easiest to implement. This brought up the issue of zero TTLs in the DNS. RFC 1034 is somewhat ambiguous with respect to zero TTLs. However, since the meeting it has been pointed out that RFC 1035 is not ambiguous. Zero TTLs mean that an RR cannot be cached, but that it can be used for the transaction in progress, no matter how long it takes to resolve.

Domain WG Report 4/14/89

Dynamic Load Balancing

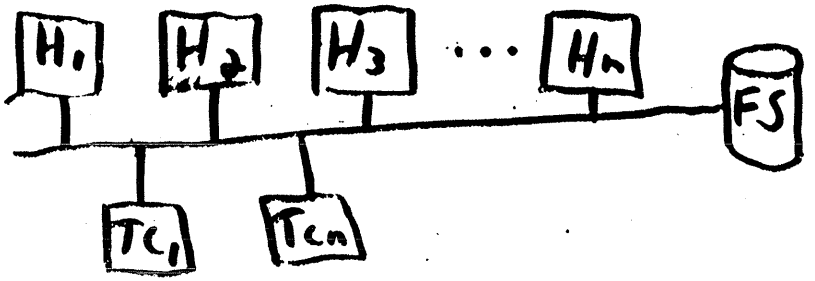

3 ways to resolve "host" ⇒ "hostn"

1. HOST ∅ CNAME Hostn

2. Host ∅ A ipaddr(Hostn)
   Host ∅ A ipaddr(Hostm)
   ⋮
   Host ∅ A ipaddr(Hostx)
   } sorted RR's

3. Host ∅ LB Hostn
   ∅ LB Hostm
   ∅ LB ⋮ Hostx
   } sorted RR's

3. is NEW RR requiring changes to all resolvers

Decided to go with #1

Dealing with zero ttl

zero means "cannot cache"

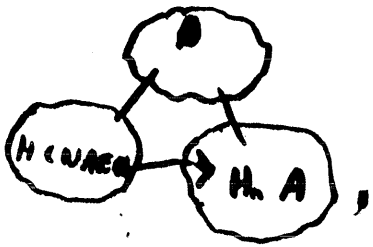

presents problem...

decided that you can cache for lifetime of query

Future issues

1. Responsible Person
2. Dynamic Add & Delete
3. Use of DNS for AS# for routing
4. Other enhancements to DNS & BIND

Dynamic Host Configuration Working Group
Chairpersons: Ralph Droms/UMD and Phill Gross/NRI

CHARTER

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).

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: (to be determined)

Dynamic Host Configuration Working Group
Chairpersons: Ralph Droms/UMD and Phill Gross/NRI

STATUS UPDATE

1. Chairpersons: Ralph Droms, droms@sol.bucknell.edu and Phill Gross, gross@sccgate.scc.com.

2. WG Mailing List: host-conf@rutgers.edu

3. Last Meeting: Cocoa Beach, April 1989

4. Next Meeting: Videoconference about June 12, or July IETF meeting in Palo Alto.

5. Pending or New Objectives: see Charter

6. Progress to Date (e.g., Documents Produced)

   Organizational meeting at Cocoa Beach: agreed on Charter and began discussion of Objective 1.

Dynamic Host Configuration Working Group
Chairpersons: Ralph Droms/UMD and Phill Gross/NRI

CURRENT MEETING REPORT
Reported by Ralph Droms and Phill Gross

AGENDA

a) Discuss Charter and Objectives
b) Set date and agenda for next meeting
c) Mailing list

ATTENDEES

Philip Almquist	almquist@jessica.stanford.edu
Dave Borman	dab@cray.com
David Bridgham	dab@ftp.com
Farokh Deboo	..!sun!bridge2@fjd
Ralph Droms	droms@sol.bucknell.edu
Hunaid Engineer	hunaid@cray.com
Bob Gilligan	gilligan@sun.com
Phill Gross	gross@sccgate.scc.com
John Lekashman	lekash@orville.nas.nasa.gov
Norbert Leser	nl@osf.org
Mark Lottor	mkl@sri-nic.arpa
Louis Mamakos	louie@trantor.umd.edu
Ron Natalie	ron@rutgers.edu
Bill Nowicki	nowicki@sun.com
Drew Perkins	ddp@andrew.cmu.edu
Mike Petry	petry@trantor.umd.edu
Robert Reschly	reschly@brl.mil
Carl-Herbert Rokitansky	roki@dhafeu52.bitnet
Rajeev Seth	rajs%hpindbu@hp-sde.sde.hp.com
Mike St. Johns	stjohns@beast.ddn.mil
Lance Travis	cmt@appollo.com
Bill Westfield	billw@cisco.com

MINUTES

This meeting kicked off the Dynamic Host Configuration Working Group. The WG was formed to study the automatic management of network configuration. Phill Gross characterized the problem by referring to the TCP/IP protocol suite documents written by Chuck Hedricks. Phill pointed out the introductory document is 25 pages long, while the management guide is 48 pages long. What we hope to do in this working group is reduce the amount of hand configuration and "wizardry" required to manage TCP/IP networks.

The group began by considering the charter and a list of objectives. The charter and objectives met with general approval. The WG modified the objectives to include the writing of an RFC based on the results of Objective 1, and/or more RFCs

based on the results of Objective 3. The WG felt that the mechanisms to be enumerated for Objective 2 represented a simple rehash of information published elsewhere, and did not warrant the publishing of a new RFC.

After agreeing on the charter and the objectives, the WG dove straight into a discussion of Objective 1. We quickly decided to limit the scope of our discussion to "Internet participants" with only a single interface. This decision allowed us to avoid the "host versus gateway" and "multi-homed host" religious wars...

Next, we talked about several configuration scenarios:

1. Virgin host
2. Rebooted host
3. Moved host
4. Replaced host
5. X Window System terminal

We launched this discussion with Objective 1 in mind - but soon discovered that we needed to solve the "host identification" problem before we could address Objective 1.

We developed the following model of host identification. There are several data items that might be used to identify a host:

1. Interface (hardware) address
2. Machine identifier (e.g., serial number)
3. IP address
4. Domain name

In all of the scenarios we considered, identifying a host involves fixing (at least) one of the above data items and then developing the other data items from existing or new bindings. The bindings may be stored in the host, stored elsewhere in the net or assigned dynamically.

For example, consider replacing a user's broken workstation. What remains fixed is the host's domain name, and the remaining information must be found from existing bindings (e.g., the Domain Name system for IP address) or existing bindings must be updated (e.g., the IP to interface address in the RARP server must be updated).

Administrivia: Ron Natalie volunteered to manage the host-conf@rutgers.edu mailing list.

DYNAMIC HOST CONFIGURATION

CHARTER

THE PURPOSE OF THIS WORKING GROUP IS THE INVESTIGATION OF HOST CONFIGURATION AND RECONFIGURATION MANAGEMENT. WE WILL DETERMINE THOSE FUNCTIONS THAT CAN BE AUTOMATED (E.G., INTERNET ADDRESS ASSIGNMENT [??]) AND THOSE THAT CANNOT (I.E., THOSE THAT MUST BE MANAGED BY SYSTEM ADMINISTRATORS). WE WILL SUMMARIZE EXISTING AUTOMATIC CONFIGURATION MECHANISMS AND PROPOSE NEW MECHANISMS.

OBJECTIVES

1) IDENTIFY THE INFORMATION REQUIRED FOR HOSTS TO:

   a) EXCHANGE PACKETS WITH OTHER HOSTS
   b) OBTAIN PACKET ROUTING INFORMATION
   c) ACCESS THE DOMAIN NAME SERVICE
   d) ACCESS OTHER LOCAL AND REMOTE SERVICES

   WORKING PAPER?

2) SUMMARIZE THOSE MECHANISMS ALREADY IN PLACE FOR MANAGING THE INFORMATION IDENTIFIED IN OBJECTIVE 1

⇒ 3) SUGGEST NEW MANAGEMENT MECHANISMS

?? 4) EXAMINE SPECIFIC SCENARIOS AND SHOW HOW THOSE SCENARIOS CAN BE RESOLVED USING THE MECHANISMS FROM 2 AND 3.

IMPL+ EXPER.

Host Requirements Working Group
Chairperson: Robert Braden/ISI

CHARTER

Description of Working Group:

The Host Requirements Working Group has the goal of producing an RFC defining the official requirements for the software on a host which is to be part of the Internet.

Specific Objectives:

1. Produce a document that is the host equivalent of RFC-1009, "Requirements for Internet Gateways", providing guidance for vendors, implementors, and users of host software for internet applications.

2. Enumerate the protocols required, referencing the RFC's and other documents describing them in detail.

3. Provide further clarification, discussion, and guidance in those areas of the referenced specifications that contain ambiguous or incomplete information.

4. Define the current architecture as completely and carefully as possible, don't invent new architecture.

5. As a secondary task, provide a forum for discussing particular solutions to pressing host problems.

Estimated Timeframe for Completion:

Our objective is to publish the document early in 1989.

Host Requirements Working Group
Chairperson: Robert Braden/ISI

STATUS REPORT

1. Chairperson: Bob Braden, braden@isi.edu, (213) 822-1511

2. WG Mailing List: ietf-hosts@NNSC.NSF.NET

3. Last Meeting: Austin IETF, January 1989

4. Next Meeting: Stanford IETF, July 1989

5. Pending or New Objectives: see Charter

6. Progress to date (e.g., documents produced):

   The document has grown to 190 pages, and a number of sets of very extensive comments have been considered and incorporated when appropriate. The decisions made at the last meeting have been incorporated, and some further changes suggested by email discussion have been made in the April 17 version, that will be installed as an Internet-Draft document.

   There are a few hard issues still to be resolved, and a number of areas in which further investigation would be desirable. Recent discussions have concerned the multihoming rules for a non-gateway host, the rules for addressing and source routing in SMTP, immediate vs. deferred processing of SMTP RCPT TO: commands, and IP source routing.

   The size of the document has been a cause for concern. It has been suggested that it would be better to split it into two documents, one for the application layer and support programs, and one for the transport layer and below. No decision has been taken on this.

CURRENT MEETING REPORT

None

Interconnectivity Working Group
Chairperson: Guy Almes/Rice and Scott Brim/Cornell

CHARTER

Description of Working Group

We aim to improve practical inter-autonomous system routing in the Internet.

Specific Objectives:

Produce a practical system for Inter-Autonomous System routing that is (a) significantly better than the current system based on EGP-2 and the Stub Model, and (b) significantly more timely than we expect the outcome of the Open Routing Working Group to be. We hope to produce:

* a Mid-Term Inter-AS Routing Architecture, and
* a Border Gateway Protocol both implemented and deployed.

Estimated Timeframe for Completion:

April 1990

Interconnectivity Working Group
Chairperson: Guy Almes/Rice and Scott Brim/Cornell

STATUS UPDATE

1. Chairpersons: Guy Almes, almes@rice.edu
   Scott Brim, swb@chumley.tn.cornell.edu

2. WG Mailing List(s): IWG@rice.edu

3. Date of Last Meeting: April, 89 at the Cocoa Beach IETF

4. Date/Site of Next Meeting: At the summer IETF meeting.

5. Pending or New Objectives:

   Revision of the draft RFCs by the summer IETF meeting.

6. Progress to Date (e.g., documents produced):

   We have draft RFCs of both the MIRA architecture and the BGP protocol.

   Draft RFC on MIRA and draft RFC on BGP; both internal documents at this stage.

Interconnectivity Working Group
Chairperson: Guy Almes/Rice and Scott Brim/Cornell

CURRENT MEETING REPORT
Reported by Guy Almes

AGENDA

Morning session: Joint meeting with the Open Routing Working Group. Afternoon session: Closed meeting to work on problems raised that morning.

ATTENDEES

Guy Almes	almes@rice.edu
Philip Almquist	almquist@jessica.stanford.edu
Rick Boivie	rboivie@ibm.com
Scott Brim	swb@devvax.tn.cornell.edu
Jeffrey Burgan	jeff@nsipo.nasa.gov
Noel Chiappa	jnc@lcs.mit.edu
Joe Choy	choy @ncar.ucar.edu
Mike Collins	collins@ccc.mfecc.llnl.gov
Dino Farinacci	dino@bridge2.3com.com
Jose Garcia-Luna	garcia@sri.com
Elise Gerich	epg@merit.edu
Tony Hain	hain@ccc.mfecc.llnl.gov
Jeffrey Honig	jch@sonne.tn.cornell.edu
Mike Karels	karels@berkeley.edu
Steve Knight	knight@cray.com
Mike Little	little@saic.com
Paul Love	loveep@sds.sdsc.edu
Marianne Lepp	mlepp@bbn.com
Charles Lynn	clynn@bbn.com
Matt Mathis	mathis@faraday.ece.cmu.edu
Milo Medin	medin@nsipo.nasa.gov
Don Merritt	merritt@brl.mil
Russ Mundy	mundy@beast.ddn.mil
Rebecca Nitzan	nitzan@nmfecc.llnl.gov
Yakov Rekhter	yakov@ibm.com
Milt Roselinsky	cmcvax!milt@hub.ucsb.edu
Bruce J. Schofield	schofield@edn-vax.dca.mil
Dallas A. Scott	dscott@gateway.mitre.org
Mike St. Johns	stjohns@beast.ddn.mil
Paul Tsuchiya	tsuchiya@gateway.mitre.org
Ross Veach	rrv@uxc.cso.uiuc.edu

MINUTES

The morning session was spent briefing the Open Routing Working Group on the Mid-Term Inter-AS Routing Architecture (MIRA) and the Border Gateway Protocol (BGP). Our assumptions about the timing of the ORWG were essentially confirmed. We presented the essential ideas of MIRA and BGP.

The afternoon session was spent going over remaining problems in MIRA and BGP. We spent the most time discussing the neighbor acquisition and neighbor reachability problems in the context of MIRA. A number of solutions were posed and discussed. Several work, but few have a reasonable combination of elegence and reliability.

Internet MIB Working Group
Chairperson: Craig Partridge/BBN

CHARTER

Description of Working Group:

As defined in RFC 1052, the original purpose was to devise an Internet Management Information Base (MIB) and Structure of the Management Information (SMI).

Specific Objectives:

After finishing version 1 of the MIB and SMI in the summer of 1988, the group continued meeting to discuss questions of upgrading and enhancing the MIB.

Estimated Timeframe for Completion:

The group currently plans to release a revised Internet MIB sometime in 1989.

Internet MIB Working Group
Chairperson: Craig Partridge/BBN

STATUS UPDATE

1. Chairperson: Craig Partridge, craig@bbn.com

2. WG Mailing List: mib-wg@nnsc.nsf.net (to join, mail to mib-wg-request@nnsc.nsf.net).

3. Last Meeting: Austin, TX, January 1989

4. Next Meeting: May 18th in Boston where it will consider a draft for a second version of the MIB.

5. Pending or New Objectives: see Charter

6. Progress to Date (e.g., documents produced):

CURRENT MEETING REPORT

none

JOMANN Working Group
Chairperson: Susan Hares/MERIT

CHARTER

Description of Working Group:

This "Joint Monitoring Access for Adjacent Networks focusing on the NSFNET Community" Working Group will:

- o discuss how to identify problems in the next hop network
- o create a list of existing tools which can solve these problems (We will discuss to see if NOC-Tools Working Group can take over this. NSFNET will archive a list of these tools.)
- o create a list of routing topology maps of regionals (possibly prepare a MAP Internet-Draft)

Specific Objectives:

See above

Estimated Timeframe for Completion:

6-9 months (August 31, 1989)

JOMANN Working Group
Chairperson: Susan Hares/MERIT

STATUS UPDATE

1. Chairperson: Susan Hares (Merit), skh@merit.edu

2. WG Mailing List: njm@merit.edu (Regional or National Net NOC people)
njm-interest@merit.edu (anyone interested)
njm-request@merit.edu

3. Date of Last Meeting: Cocoa Beach, April 11-14, 1989

4. Date of Next Meeting: Stanford, July 25-26, 1989

5. Pending or New Objectives: to be determined

6. Progress to Date (e.g., documents produced):

- Common SNMP monitor session
- Policies discussed; agreement on error reporting, outage reporting, and virus reporting
- MAPs collected
- Tools list - no progress. Possible project for NOC-Tools WG

WHO SHOULD ATTEND:

Technical representatives from mid-level or peer networks. In the future we may want to extend this to technical representatives from campus networks. However, in interest of getting a lot of work done quickly the initial working group will be limited.

JOMANN Working Group
Chairperson: Susan Hares/MERIT

CURRENT MEETING REPORT
(Chaired by Elise Gerich in the absence of Susan Hares)
Reported by Elise Gerich/MERIT

AGENDA

1. SNMP Community Names
2. Traffic Statistics Request
3. ARPANET

SNMP Community Names
==========

The discussion centered around the distribution and changing of the global SNMP community names.

It was mentioned that there should be a policy for distribution along with the name.

Merit asked if anyone minded changing the name. All parties, except for NYSERNet, did not feel strongly about the periodic changes of the SNMP community name. It was felt that this was not needed.

A vote was taken and by majority (not unanimously), it was decided to not change the SNMP community name at a constant interval. The name will stay the same until further notice.

As a side note, it was mentioned that MERIT/IBM now has SNMP deployed in the backbone and is giving out the community name to any regional who wants to interoperability test. There has been very little interoperability testing as of the time of the meeting. Anyone who was interested in testing the implementation was asked to talk with Elise after the meeting.

Mark Oros's Questions
==========

There were some questions for the group posed by Mark Oros. These were discussed. As follows:

How many people use SNMP/SGMP?

Almost all of the regionals currently use it or are in the process of installing it.

Vendor specific traps:

Mark proposed that it would be useful to have a Vendor Specific SNMP trap for the Proteon Gateway which would be used to indicate someone logging into the gateway. When someone successfully logged into the gateway, the Proteon would generate a trap.

After some discussion, the group felt that this was a good idea. Proteon will be notified and it was suggested that all concerned parties call their friendly Proteon sales rep.

Traffic Statistics Request

Merit has asked the regionals to provide them with traffic statistics/data on the traffic flowing from the regional into NSFNET. For example, the amount of packets/bytes sent from the regional to the NSS.

Ideally Merit would like:

Packet Size Distribution
Distribution of traffic
Packets
Bytes

If you don't have the data, then you obviously cannot provide it. The time-frames Merit would like to see include: 24 hours, weekly, hourly. Raw data is also preferred.

In summary, Merit will take whatever you have in terms of data collected and reports generated.

ARPANET

PSCNET is now the primary route to ARPANET. Much thanks to Matt Mathis and others for dealing with this at the spur of the moment.

There is a planned NSFNET connection to the MILNET at Univ. of Maryland. There is also a connection to the MILNET planned at NASA AMES.

FUTURE OF JOMANN

This topic was left up to the mailing list. An E-mail discussion is planned.

Lan Manager Working Group
Chairperson: Amatzia Ben-Artzi/3Com

CHARTER

Description of Working Group:

To define 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 four basic areas:

- Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
- Define 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.
- Define the mechanisms of exchange of management information between clients and servers so that proxies can be developed.

Estimated Timeframe for Completion:

LAN Manager Working Group
Chairperson: Amatzia Ben-Artzi/3Com

STATUS UPDATE

1. Chairperson: Amatzia Ben-Artzi, amatzia@spd.3mail.3com.com

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

3. Date of Last Meeting: Cocoa Beach April 1989

4. Date of Next Meeting: May 1989

5. Pending or New Objectives: to be defined

6. Progress to Date (e.g., documents produced):

   - Draft Proposal for MIB Objects
   - Discussed 2nd draft at the meeting
   - Working on 3rd draft. To be posted in mailing list this week. Hope to bring it before the "large" MIB-WG at their 5/18 meeting.

LAN Manager Working Group
Chairperson: Amatzia Ben-Artzi/3Com

CURRENT MEETING REPORTS
Reported by Amatzia Ben-Artzi

FEBRUARY INTERIM MEETING

The Lan Manager MIB working group met Wednesday, 2/22 in Sunnyvale, CA for our first meeting. The meeting was very productive and generated a long list of output and action items. Below is a summery of the meeting and major decisions reached.

The minutes cover seven topics:

1. Introduction
2. The Group's Objectives
3. To Proxy or Not?
4. Initial Documents and Editors
5. Relationship to the MIB WG
6. Mailing List
7. Timetable

1. Introduction to the Lan Manager.

Amatzia gave a short introduction on the Lan Manager. The emphasis is on the management interoperability issues between the Lan Manager as a standard in the workgroup environment and the standards being developed in the "enterprise" environment. As both are based on the TCP/IP it is important that they can cooperate. Microsoft offered to provide a more extensive overview of the LanManager if people will find it useful. Please send me feedback on this one!!

2. Group's Objective

The objective of the group is to define 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).

We found that it translated into four basic areas:

- Define a set of management information out of the existing Lan Manager objects to allow for useful management from a TCP/IP based manager.

- Define extensions to the TCP/SMI where appropriate.

- Develop requirements for additional network management information, as needed, and work to extend the Lan Manager interfaces to support such information.

- Define the mechanisms of exchange of management information between clients and servers so that proxies can be developed.

3. Proxy or Not?

We concluded that the manager need to have access to the management information in every client in a direct way. It amounts to the following:

- A client may have a Management/TCP stack.

- A client may be supported by a server that acts as "proxy" on his behalf.

- The manager need not be aware of which one of the techniques above is being used in the workgroup.

However, we recognized that in the proxy mode, the server may have two types of objects: server resident, and client resident. For the second type, a server-client mechanism has to be developed to allow the implementation of a proxy.

4. Initial Documents and Editors

The following documents have been identified as needed. Initial editors were also selected.

- SMI Extensions: Pranati Kapadia, HP

- MIB Objects: Jim Greuel, HP

- NDIS Extensions (not assigned)

- Transport APIs for NM Information access (not assigned)

- Client-Server management protocol

5. Relationship to the MIB group

A real objective of this MIB group is to work under the "BIG" MIB group. One implication was that the MIB specification should follow the 1066 RFC (specifying all attributes as "objects") with an appendix that actually describe the containment relationship (Same technique that was used in the CMOT RFC to re-state the supported MIB)

A big question mark is SMI. Can we live with the guideline of "no SMI extensions" ? We shall address it when the first required extension shows. We do know, however, that EVENTS or alerts, or alarms) are a big issue, but we where not sure if this was an SMI issue or what.

We also feel very strongly that the recommendation of the previous MIB group should be followed: Lan Manager should be assigned a number of the MIB (Like TCP, IP, or CMOT) and define it objects under this branch. Then, bring them forward to the larger group for discussion and approval. ONLY EXPERIMENTAL OBJECTS SHOULD BE PLACED UNDER THE EXPERIMENTAL BRANCH.

The whole branch is OPTIONAL, so people who don't implement it do not have to worry about conformance. It seems like we would want, for simplicity of conformance, at least initially, to say: the branch is optional, but if you implement it, it is ALL mandatory. The target is roughly 20 - 30 objects initially.

6. Mailing list:

Welcome a new mailing list: LanManWG@spam.istc.sri.com

As usual, there is also a LanManWG-request.

Initial membership:

aguilar@istc.sri.com	SRI
jimg@hpcnd.hp.com	H-P
...!ucbvax!mtxinu!excelan!ramesh	Excelan
...microsoft!henrysa	Microsoft
geo@ub.com	Ungerman-Bass
kapadia@hpda.hp.com	H-P
davep@esd.3com.com	3Com
jcham@mbunix.mitre.org	Mitre
Hunter@nmfecc.arpa	Laurence Livermore Lab
...!ucbvax!mtxinu!excelan!pramod	Excellan
jonab@cam.unisys.com	Unisys
kzm@twg.com	The Wollongong Group
amatzia@spd.3com.com	3Com

7. Timetable

March 17: First draft proposal for MIB objects in the mail.
March 31: Comments are due back to the editor
April 11-14: IETF meeting. We shall meet sometime there to discuss the proposed draft (currently planned as a half-day meeting)

Thanks to all the participants for a very effective meeting.

APRIL IETF MEETING

ATTENDEES

Amatzia Ben-Artzi	amatzia@spd.3+.3com.com
Jeff Case	case@utkuxl.utk.edu
Jim Greuel	jimg%hpcndpc@hplabs.hp.com
Steven Hunter	hunter@ccc.mfecc.llnl.gov
Lee LaBarre	cel@mbunix.mitre.org
Rajevv Seth	rajs%hpindbu@hplabs.hp.com
Jim Sheridan	jsherida@ibm.com
Louis Steinberg	louiss@ibm.com

MINUTES

Key actions/directions:

- Group agreed that lanman-mib should be part of the standard mib version 2.
- Events will be introduced on a limited basis and enter as "experimental" to the MIB.
- Experimental events should have examples in the standard how they fit into CMOT and SNMP.
- The event "thresholds" (i.e., boundery conditions) should be defined as standard objects using the standard SMI (RFC 1065). They will, however, be defined in tables to preserve their internal dependency/relation.

Next activities:

- Review 3rd draft over the mailing list.
- Present to MIB-WG at 5/18.
- Next meeting to finalize the RFC or start the next phase [depending on progress with the MIB-WG] at IETF meeting in July.

LANMAN - WG

WHAT IS LanManager

- ▷ Network Operating System
  - standard
  - Provides
    - File sharing
    - Print sharing
    - Distributed Apps.
  - Based on NetBios (Rfc 1001, 1002)
- ▷ APIs to many of its functions
- ▷ Works ON:
  - DOS & OS/2 (Now)
  - UNIX, Mac (soon)
  - VMS (Later)

LANMAN - WG

What is LanMan-WG

1. SUB-Group of the MIB-WG
2. Defines NM Objects that can be used by TCP NM protocols
3. Extends the scope of TCP Manager to workgroups.

LANMAN-W6

ARCHITECTURE

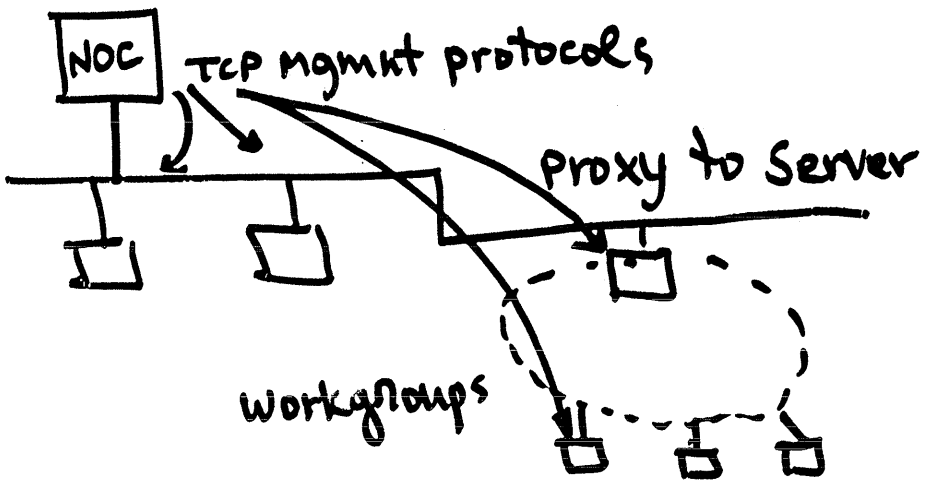

- Noc doesn't care

THE MIB

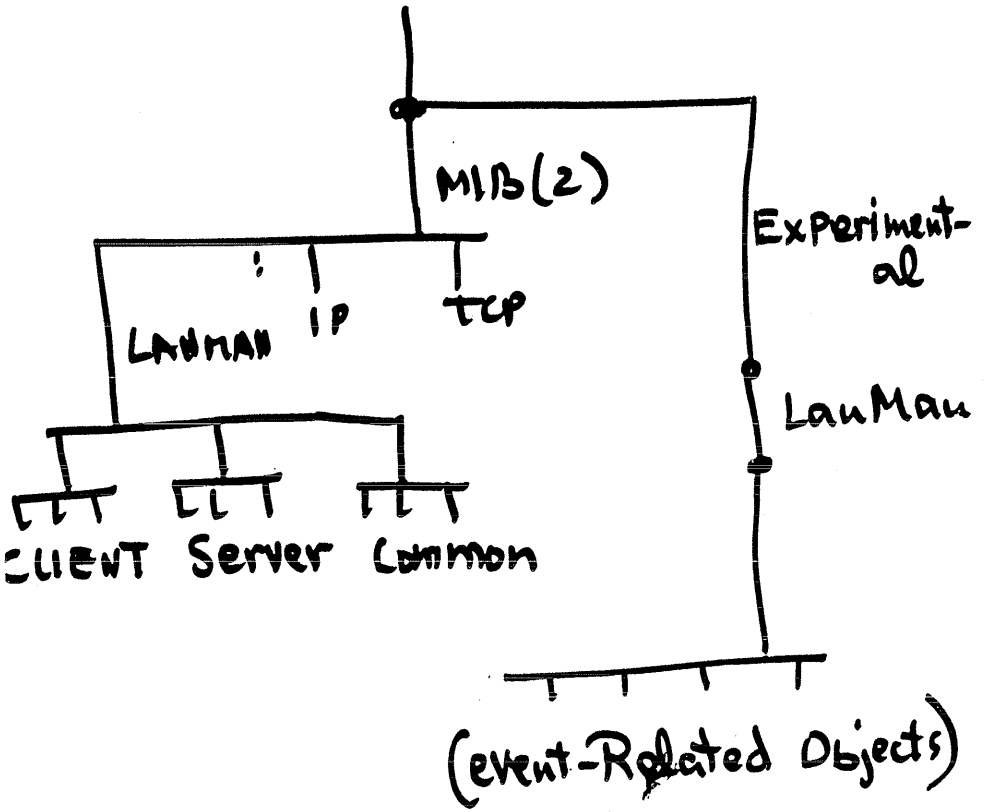

LanMan-WG

PHASES

I.
- Map only existing LanMan objects
- Read only in Main Section.
- Experimental Section on events — using only objects and Conforming to 1065

II. (?) Control

(?) Events conforming to the output of the new "Event Mgmnt" TF.

LanMan-WG

Status/Actions

- Second draft discussed at the meeting this week.
- The "event" Section to be revised. Mapping onto CMOT and SNMP to be included
- Will be placed on MIB-WG ML. (soon)
- Will Request review/approval in next MIB, that submitted as "Draft RFC".

LanMan-Wg

- LanManWg@spam.istc.SRI.Com

(-Request)

- Next Meeting

(3) As part of the MIB-WG
Boston, May 18

NISI Working Group
Chairpersons: Karen Bowers/NRI and Phill Gross/NRI

CHARTER

Description of Working Group:

The NISI WG will explore the requirements for common, shared internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. This effort will be a sub-group of the User Services WG and will coordinate closely with other IAB and FRICC efforts in the area of Directory Services.

Specific Objectives:

1) Write a short white paper to serve as a starting point for discussions on an Internet-wide information services infrastructure.

2) Develop a more detailed statement of required information services as currently supplied by a typical network information service organization. This will initially take the form of an annotated outline of services, suitable to be expanded into a full Requirements Document.

3) Define candidate pilot projects for consideration by this or other groups to implement. Initial candidates include:

- Define common user interface for information retrieval by electronic mail.

- Define common user interfaces for other information services (e.g., white pages)

- Define the minimally required information content for an Internet-wide use registration database and begin to collect such information.

Estimated Timeframe for Completion:

Objective 1 -- Draft to be distributed by email to the USWG mailing list by June 30, 1989.

Objective 2 -- Annotated outline, ready for volunteer writing assignments, to be distributed by email to the USWG and the NISI mailing lists by June 30, 1989

Objective 3 -- To be determined

NISI Working Group
Chairpersons: Karen Bowers/NRI and Phill Gross/NRI

STATUS UPDATE

1. Chairpersons: Karen L. Bowers / bowers@sccgate.scc.com
   Phill Gross / gross@sccgate.scc.com

2. WG mailing list: NISI@MERIT.EDU

3. Date of Last Meeting: 04 May 89 / NRI

4. Date of Next Meeting: 25-28 Jul 89 / Stanford IETF

5. Pending Objectives:
   Two Immediate Actions: Network Information Services Requirements Document

   White Paper: Network Information Services Infrastructure (NISI)

   Both drafts to be distributed prior to the July IETF Meeting. These drafts will be further expanded/edited by the NISI WG and then presented to the full body of the USWG for review.

6. Progress to Date (e.g., documents produced):

   First Meeting Held

NISI Working Group
Chairpersons: Karen Bowers/NRI and Phill Gross/NRI

CURRENT MEETING REPORT
Reported by Karen Bowers and Phill Gross

AGENDA

- Individual Information Services Briefings

    * NSFnet, NIS (MERIT) - Jim Sweeton
    * NSFnet, NNSC (BBN) - Craig Partridge
    * CSNET, CIC (UCAR) - Laura Breeden
    * BITNET, BITNIC (EDUCOM/BITNET) - Jim Conklin
    * DDN, NIC (SRI) - Mary Stahl
    * PREPnet (PA Research and Economic Partnership) - Thomas Bajzek

- Strategy for Preparation of the Network Information Services Requirements Document : purpose of document, how best to prepare

- Discuss Related Projects: IAB White Pages Workshop
  FRICC DECNET/Internet Directory Services (Phill Gross)

- Discuss/Modify Strawman Agenda:

    * Pilot Project Selection

    * Technical Approach to be Employed
        Explore Methods for Providing Selected Services
        Define Support Requirements
        Define the Common, Shared User-Interface
        Define Applications to Interface Current Internet Services

    * Follow-on Activities
        How to Implement
        Issues for an Internet-wide Information Services Infrastructure (NISI)

- Determine Critical Actions to be Undertaken Immediately by the NISI WG

ATTENDEES: Thomas Bajzek (PREPNET), Karen L. Bowers (NRI), Laura Breeden (CSNET), Jim Conklin (BITNET), Jose Garcia-Luna (DDN/SRI), Phill Gross (NRI), Craig Partridge (NNSC/BBN) Mike Roberts (EDUCOM), Mary Stahl (DDN/SRI), Jim Sweeton (NSI/MERIT). Participants unable to attend this session but who have been invited to join this closed forum are:
Waverly Williams (SPAN) and Karen Roubicek (NNSC/BBN) (FARNET representation was provided by both Laura Breeden and Thomas Bajzek. Guy Almes was unable to attend.)

MINUTES

The Network Information Services Infrastructure (NISI) WG held its first meeting on May 4, 1989, 12:00pm - 5:45pm at the Corporation for National Research Initiatives, Reston, VA.

The goal of this meeting was to explore the technical and non-technical issues of providing common, shared Internet-wide information services.

As a basis for understanding the current, general information services requirements of the Internet community, each attendee who currently is responsible for providing network information services was asked to provide a short presentation summarizing the type and quality of services they provide, and the services they see a need for but currently do not provide. This exercise enabled the WG to develop a baseline of what types of services exist, overlap and differ, how those services are implemented, and the value-added services for which there could be a future demand. This information will be assembled into a "Network Information Services Requirements Document" by the NISI WG and presented to the body of the USWG for review and comment.

Phill Gross provided an update on the results of Dave Clark's Workshop on Internet White Pages. Last week a draft Plan for Internet Directory Services was released by Karen Sollins to all the workshop participants for comment/review. It is purported this document will be released shortly to the public, at which time it will be available to the NISI WG. In summary, X.500 was identified as the most likely candidate for directory service. However, the consensus of the workshop is to propose a field trial(s) to include experiments with at least an X.500 implementation, Profile to explore a non-hierarchical structure and possibly DECnet's Network Architecture Naming Service (DNANS). Another related directory services project underway is the current indepth examination of DECnet's Network Architecture Naming Service (DNANS) by the FRICC, in anticipation of implementing this service in the large scale DECnet Internet. This will be an important issue in the upcoming transition from DECnet Phase IV to DECnet Phase V.

A discussion of the short and mid-term agenda for the NISI WG ensued. The WG members decided on two immediate actions:

1) to write a Network Information Services Requirements Document that identifies information services currently available, those that overlap and differ, value-added services that are or should be made available, and how these services are currently implemented. In addition, this document could include "lessons learned" on how to better provide these services.

2) to prepare a white paper defining the concept of a Network Information Services Infrastructure, to be made available as a reference document for any organization proposing/developing a new or follow on network information service or network information service infrastructure effort.

Actions for follow-on meetings include:

Explore the preliminary steps in developing a Common User Interface(s) for email information services and for white pages and determine if these are pilot projects the USWG should undertake. Perhaps at a minimum, the NISI WG could "define" what those Common User Interfaces should be.

Examine the requirements for and issues associated with common/standard User Registration Database Attributes and accomplish a draft definition of those attributes. An associated task would be to define tools for database management as well.

Coordinate with the FRICC/IAB on our mutual concerns and related activities.

The next meeting will be in Stanford, 25-28 July 1989. In the interim, initial drafts of the NIS Requirements Document and the NISI white paper will be prepared by Phill Gross and Karen Bowers and provided by email to the NISI WG members for their reciprocal input prior to the July IETF meeting.

Network Management Services Interface Working Group
Chairpersons: Jeff Case/UTK and Keith McCloghrie/TWG

CHARTER

Description of Working Group:

The objective of the Network Management Services Interface Working Group is to define a management services interface by which network management applications may obtain access to a heterogeneous, multi-vendor, multi-protocol set of network manageable devices. While such an interface is desirable, the extent to which an implementation is feasible in real systems, is not clear.

The services interface is intended to support the network management protocols and strategies commonly found in networking devices today. As this is an Internet Engineering Task Force Working Group, the natural focus is on current and future network management strategies used in the Internet. However, the interface being defined is expected to be sufficiently flexible and extensible to also allow support for other protocols, at little or no extra cost. The anticipated list of supported strategies and protocols includes: the standard TCP/IP network management protocol, i.e., the Simple Network Management Protocol (SNMP); the standard OSI network management protocol, i.e., the Common Management Information Protocol (CMIP); an experimental TCP/IP network management protocol, i.e., Common Management Information Protocol over TCP/IP (CMOT); the Manufacturing Automation Protocol and Technical Office Protocol (MAP/TOP); as well as proprietary network management protocols.

Specific Objectives:

1) define an architectural framework for such a service interface,
2) advance an RFC which describes the interface,
3) implement one or more prototypes, and
4) evaluate the viability of the interface including recommendations for future work and the feasibility of implementation in real systems.

Estimated Timeframe for Completion:

Not known at this time

Network Management Services Interface Working Group
Chairperson: Jeff Case/UTK and Keith McCloghrie/TWG

STATUS UPDATE

1. Chairpersons: Jeff Case, case@UTKUX1.UTK.EDU
   Keith McCloghrie, kzm@TWG.COM

2. WG Mailing List(s): none implemented at this time

3. Date of Last Meeting: Cocoa Beach, FL April 11, 1989

4. Date of Next Meeting:

   May 15-17, 1989 in conjunction with the Boston IFIPS conference
   (Editor's note: this meeting will be postponed because expected contributions have not been received as expected.)

5. Pending or New Objectives:

   No changes in objectives resulted from this meeting.

6. Progress to Date (e.g., documents produced):

   - Initial meeting held. We are working through the process. Revisons to the initial draft document are underway.

   - There is a draft document which was discussed in the meeting. Revisions to the draft are underway.

Network Management Services Interface Working Group
Chairpersons: Jeff Case/UTK and Keith McCloghrie/TWG

CURRENT MEETING REPORT:
Reported by Jeff Case and Keith McCloghrie

AGENDA
1st Meeting

Introductions, registration, etc	(15 min)	KZM
Introduction to the problem and work in progress	(45 min)	JDC
Discussion of the draft RFC	(60 min)	JDC
Brainstorm about architecture required	(45 min)	KZM
Organizational details: mtgs, mailing list, etc.	(30 min)	KZM

ATTENDEES

Momhammad Alaghebandan	mra@bridge2.3com.com
Cathy Aronson	cja@merit.edu
David Bridgham	dab@ftp.com
Jeff Case	case@utkux1.utk.edu
John Chao	jchao@bbn.com
Danny Cohen	cohen@isi.edu
John Cook	cook@chipcom.com
Hunaid Engineer	hunaid@cray.com
Lionel Geretz	lionel@salt.acc.com
Jim Greuel	jimg%hpcndpc@hplabs.hp.com
Brian Handspicker	bd@vines.dec.com
Steven Hunter	hunter@ccc.mfecc.llnl.gov
Mark Kepke	mak%hpcndm@hplabs.hp.com
Lee LaBarre	cel@mbunix.mitre.org
Norbert Leser	nl@osf.org
Paul Love	loveep@sds.sdsc.edu
Keith McCloghrie	kzm@twg.com
Bill Norton	wbn@merit.edu
Rajeev Seth	rajs%hpindbu@hp-sde.sde.hp.com
Jim Sheridan	jsherida@ibm.com
Lance Travis	cmt@apollo.com
Steve Waldbusser	sw01@andrew.cmu.edu
Dan Wintringham	danw@igloo.osc.edu

MINUTES

An initial meeting of the Working group was held on April 11 at the Cocoa Beach IETF meeting. Jeff Case presented the problem and explained the work that had taken place before the meeting. A draft document was distributed and discussed. Several individual volunteered to prepare and forward alternate text for sections of the document before May 1 so that a new draft can be prepared and distributed before the next meeting. It was agreed to attempt to meet in Boston sometime between May 15-17 in conjunction with the IFIPS network management conference and the MIB working group meeting. (Editor's note: however, since expected contributions to the draft document have not materialized in a timely fashion, this meeting will be postponed.)

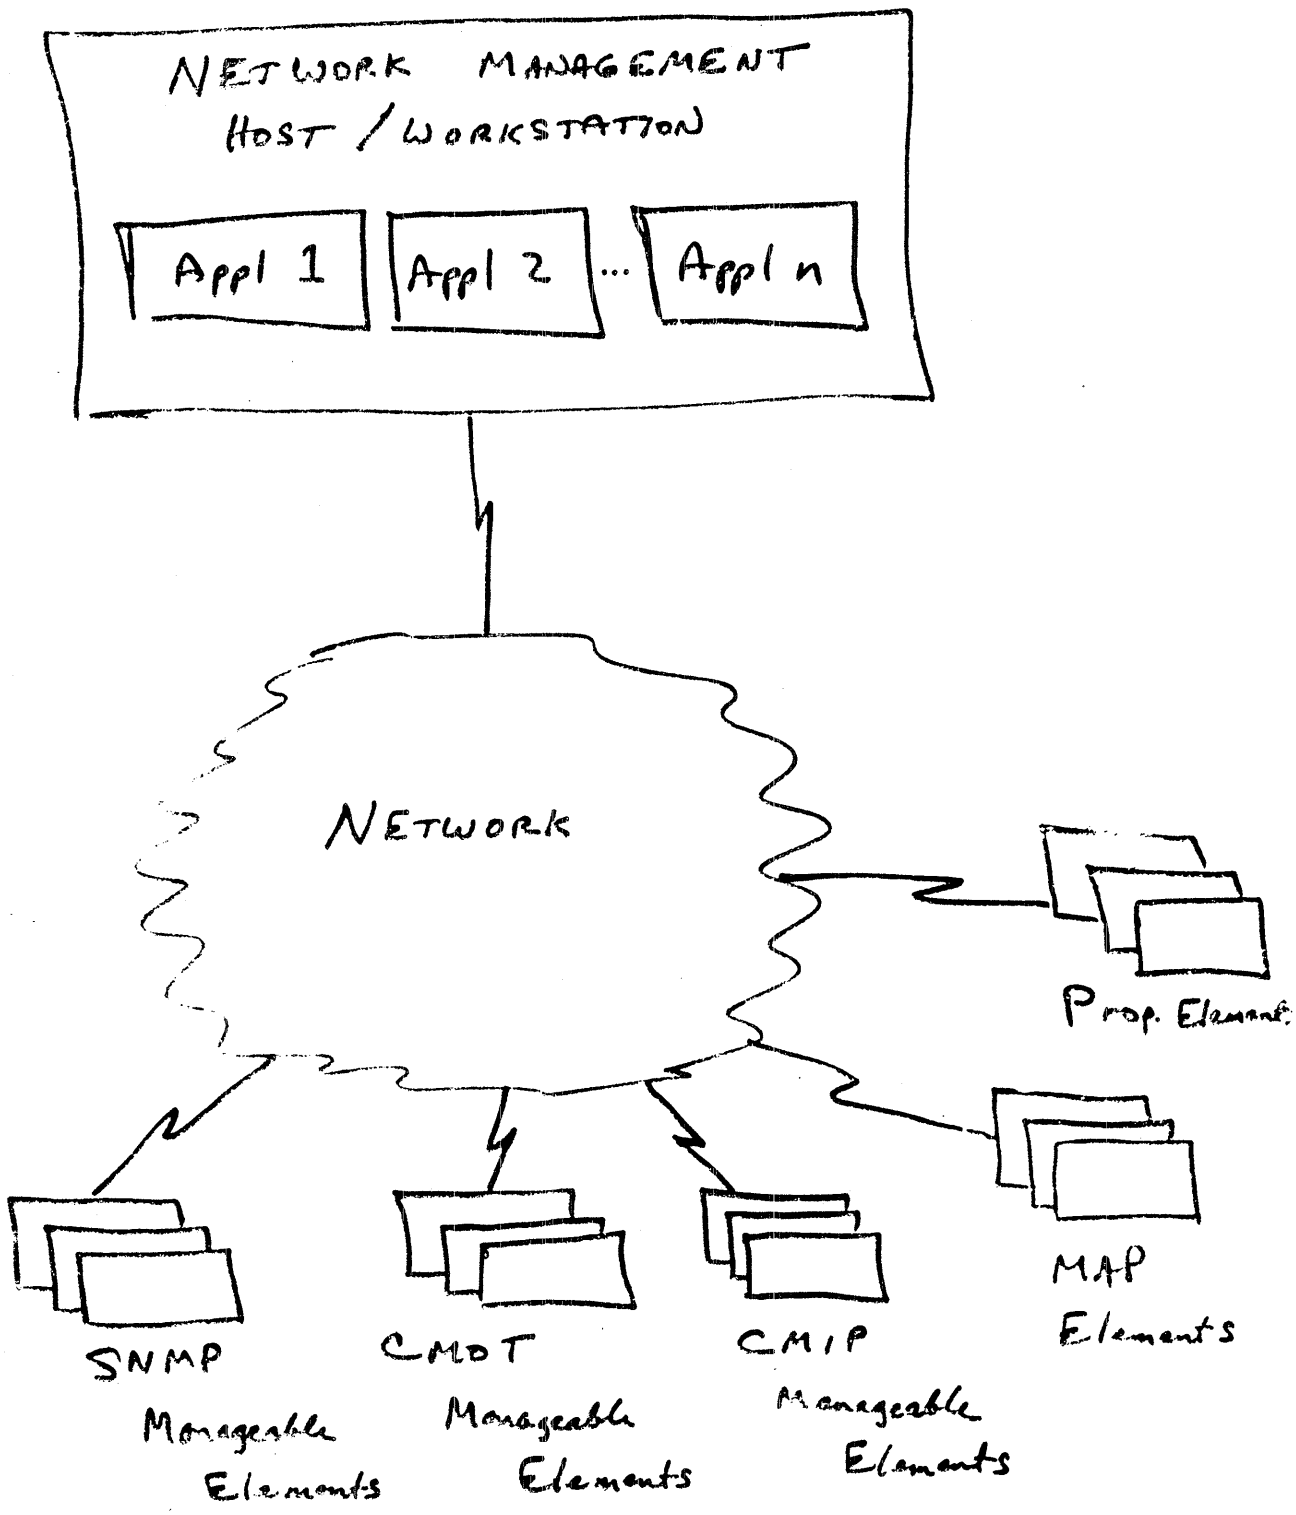
NETWORK MANAGEMENT
HOST / WORKSTATION
Appl 1
Appl 2
...
Appl n
NETWORK
Prop. Element
SNMP
CMOT
CMIP
MAP
Elements
Manageable
Elements
Manageable
Elements
Manageable
Elements

THE PROBLEM

NOC-Tools Working Group
Chairperson: Robert Enger/Contel

CHARTER

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-November IETF meeting. Preliminary versions will be made available earlier.

NOC- Tools Working Group
Chairperson: Robert Enger/Contel

STATUS UPDATE

1. Chairperson: Robert Enger, enger@sccgate.scc.com

2. WG Mailing List: noctools@merit.edu

3. Date of Last Meeting: Cocoa Beach, April 12, 1989

4. Date of Next Meeting: Stanford, July 25-28 1989

5. Pending or New Objectives: Final review of the catalog planned for October/November IETF Meeting.

6. Progress to Date (e.g., documents produced):

   Preliminary catalog outling and entry format have been developed. First pass made at enumerating the tools available.

NOC-Tools Working Group
Chairperson: Robert Enger/Contel

CURRENT MEETING REPORT
Reported by Robert Enger

ATTENDEES

Mohammad Alaghebandan	mra@bridge2.3com.com
Cathy Aronson	cja@merit.edu
Jeff Case	case@utkuxl.utk.edu
John Chao	jcha@bbn.com
John Cook	cook@chipcom.com
Robert Enger	enger@sccgate.scc.com
Hunaid Engineer	hunaid@cray.com
Mark Fedor	fedor@nisc.nyser.net
Steven Hunter	hunter@ccc.mfecc.llnl.gov
Gary Malkin	gmalkin@proteon.com
Don Morris	morris@ncar.ucar.edu
Bill Norton	wbn@merit.edu
Joel Replogle	jr@ncsa.uiuc.edu
Joyce Reynolds	jkrey@venera.isi.edu
Karen Roubicek	roubicek@nnsc.nsf.net
John Scott	scott@dg-rtp.dg.com
Rajeev Seth	rajs%hpindbu@hp-sde.sde.hp.com
Jim Sheridan	jsherida@ibm.com
Mary Stahl	stahl@sri-nic.arpa
Bob Stine	quest@edn-unix.dca.mil
Steve Waldbusser	sw01@andrew.cmu.edu
Bill Westfield	billw@cisco.com
Dan Wintringham	danw@igloo.osc.edu

MINUTES

The kick-off meeting was quite productive, and was a good start to meeting our group's goal of producing a document and disbanding by December of this year.

We began the meeting by reviewing the charter of the Noctools Working Group. Briefly, the purpose of the Noctools Working Group is to produce the first edition of a catalog of network management tools for TCP/IP internets, and to identify one or more sites to serve as tool and catalog repositories. Incredibly, we have already received volunteers for catalog and tool repositories: Phil Almquist, of Stanford, and Mark Fedor, of Nysernet, each volunteered his own site.

We then discussed the proposed structure of the tool catalog. The bulk of the document will be an alphabetic list of network management tools. The list of entries will be preceded by an intro describing the catlog's purpose and structure, and an index into the catlog. The index will group tools according to function or category; some tools will be listed in several functional groups. Developing a workable taxonomy is important

for indexing the tool catalog. Our goal is to allow the document's readers to quickly locate the tools that will meet their needs.

The need for a useful index was the motivation for a discussion on a taxonomy of network management tools. There was little consensus, other than that the "IEEE Five" (performance, fault, configuration, accounting, and security management) do not constitute a very useful taxonomy. From floor came the sound point that attempting to categorize at this stage is premature; we need to examine our data before attempting to describe them. We decided to defer the development of a taxonomy until we had the opportunity to review some number of entries.

We also discussed the scope of the catalog. In summary, even though we are cataloging tools for internet management, some LAN management tools and tools that run only under certain operating systems or hardware configurations are potentially eligible for inclusion in the catalog. We did not, nor do we need to, develop criteria for differentiating internet management tools from other management tools. There was, however, consensus that breakout boxes will not be included in the catalog.

As a basis for discussion, we developed a quick list of internet tools that probably will be included:

arp
ding
etherfind
internet rover
lan analyzer
mconnect
netspy
netspy
netwatch
nnstat
nslookup
nysernet snmp tools
overview
ping(s)
sniffer
tcpdump
tcptrace
tokenview
xnetmon
xup

After enumerating the tools, we then discussed the scope and format of individual tool descriptions. To avoid unfairness to products and legal battles with their vendors, descriptions will

be strictly objective. We agreed on a working format for the tool entries; each entry will have the following sections, with no more than a few paragraphs for each:

1. Tool name.
2. Key word list: for now, a best guess. We will have a uniform key-word list in the final catalog.
3. Abstract: a brief description of the tool's purpose and characteristics.
4. Mechanism: how the tool works.
5. Caveats: cautions about tool impact on system performance, etc.
6. Hardware requirements.
7. Software requirements.
8. Related analysis tools: post processors, data reduction tools, etc.
9. Availability/Support: How to obtain the tool, and whether it is public domain, copyrighted, or commercially available. This will NOT include pricing information.
10. Bugs/limitations.

The issue of accepting vendor inputs for the catlog came up. Our consensus was that we will welcome vendor submissions, though the Noctools Working Group will retain editorial control over the catalog.
>From the floor came the excellent suggestion to establish a review panel. This would provide some anonymity to reviewers, and could insulate them from vendor pressure.

Towards the meeting's end, we solicited volunteers to write draft entries for the tools we had listed. For now, entries will be submitted in ASCII, to the mailing list. As entries are collected, the draft of the tool catalog will reside in the ietf drafts directory.

In other business, co-chairmen Bob Stine and Bob Enger were tasked to develop a tool-use survey and forward it to NOCs, mailing lists, trade magazines, and a few of the internet old boys. Phill Gross was tasked to inform Jon Postel of our group's activities.

A mailing list, noctools@merit.edu, has been established for the working group. As usual, requests to join the list should be directed to noctools-request@merit.edu.

NOC-Tools WG

An effort to provide network managers with increased access to the tools available.

Co-Chairs: Bob Stine - Sparta
Bob Enger - Contel

Noc-Tools is a brief spin-off from the User Services Working Group.
Expected TTL: 9 months.

NOC-Tools is an **open** WG.
Your participation is encouraged.

Status

two organizational meetings were held prior to this conference. The charter was written, and a preliminary catalog outline and entry format were developed.

Noctools held it first functional meeting on Wednesday morning. There were 24 attendees, as well as a few cameo appearances.

Three individuals indicated that their site might be willing to act as one of repositories.

Discussions were held on the format of the catalog entry, as well as what would be the most usefull index organization. We also tried to enumerate all the tools available, BUT:

YOUR HELP IS NEEDED

If you know of a useful tool:
Let us know.

Send e-mail giving the tool's name, function, and how to go about getting it.

Mailing List: NOCTOOLS@merit.edu

To join the list: -request

ALSO

If you would like to write the Catalog entry for a tool: please drop us a note.

We will supply you with the format, and enough examples to dislodge even the most stubborn writer's block.

Open SPF-based IGP Working Group
Chairpersons: Mike Petry/UMD and John Moy/Proteon

CHARTER

Description of Working Group:

The OSPF working group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage muliple vendor implementations.

Specific Objectives:

1. Design the routing protocol, and write its specification.

2. Develop multiple implementations, and test against each other.

3. Obtain performance data for the protocol.

4. Make changes to the specification (if necessary) and publish the protocol as an RFC.

Estimated Timeframe for Completion:

We have a complete protocol specification. Implementation experience and performance data should be obtained during the summer of 1989. The specification should be ready for final review by the October-November IETF.

Open SPF-based IGP Working Group
Chairpersons: Mike Petry/UMD and John Moy/Proteon

STATUS UPDATE

1. Chairpersons: Mike Petry, petry@trantor.umd.edu
   John Moy, jmoy@proteon.com

2. WG Mailing List(s): ospfigp@trantor.umd.edu

3. Date of Last Meeting: April 1989, Cocoa Beach, FL

4. Date of Next Meeting: Teleconference scheduled tentatively for the beginning of June

5. Pending or New Objectives:

   - o New revision of specification to be submitted (5/89)
   - o Testing of implementations (summer 89)
   - o Final specification to be submitted (fall 89)

6. Progress to Date (e.g., documents produced):

   - o The OSPF Specification, first revision (1/89)
   - o First revision of the OSPF specification finished (1/89)
   - o Two OSPF implementations nearing completion
   - o OSPF presentation given during IETF plenary (4/89)

Open SPF-based IGP Working Group
Chairpersons: Mike Petry/UMD and John Moy/Proteon

CURRENT MEETING REPORT
Reported by John Moy

AGENDA

The OSPFIGP working group met for a half day on April 11th in Cocoa Beach.

ATTENDEES

Jeffrey Burgan	jeff@nsipo.nasa.gov
Rob Coltun	rcoltun@trantor.umd.edu
Dino Farinacci	dino@bridge2.3com.com
Elise Gerich	epg@merit.edu
Jeffrey Honig	jchesonne.tn.cornell.edu
Van Jacobson	van@helio.ee.lbl.gov
Steve Knight	knight@baldmt.cray.com
Mike Little	little@saic.com
Milo Medin	medin@nsipo.nasa.gov
John Moy	jmoy@proteon.com
Mike Petry	petry@trantor.umd.edu
Robert J. Reschly	reschly@brl.mil
Paul Tsuchiya	tsuchiya@gateway.mitre.org
Ross Veach	rrv@uxc.cso.uiuc.edu

MINUTES

1. It was decided to shorten the name of the protocol, and the name of the working group, from OSPFIGP to OSPF.

2. We talked for a while about variable length subnet masks. The present OSPF specification attempts to specify a minimum number of rules so that the set of subnet masks for any IP network is consistent. It was widely viewed that these rules, and their attendant explanation, were inadequate. Much more text is needed to explain valid methods for assigning subnet masks. It was agreed that this kind of discussion is outside the scope of the routing protocol, and so should be moved to a separate document.

3. The subject of authentication came up. OSPF has a 16-bit authentication type field, and a 64-bit field that can be used for things like encrypted checksums, in its standard packet header. Jeff Schiller explained his new proposal for an authentication IP option. This may alleviate the need for authentication support in the OSPF header. However, for the moment we decided to keep the authentication support unchanged.

4. We previewed the presentation that was presented to the IETF

5. The following changes to the OSPF specification were agreed upon. Small changes have been omitted. For more details, see the next draft of the OSPF specification.

   a. The D-bit was removed from the summary links advertisement. Router IDs now must be distinguishable from class A, B, and C network numbers.
   b. IP interface addresses are now specified in a router's router links advertisement.
   c. Cost of links to transit nodes must be > 0.
   d. Receiving a Database Description packet with the INIT bit set always resets the adjacency.
   e. The DR calculation procedure was modified slightly.
   f. A DR priority of 0 means that the router is ineligible to become DR.
   g. OSPF routing packets will have their IP precedence set to "Internetwork Control".
   h. A network mask has been added to the summary links advertisement.

6. There was discussion of the LS checksum field which resides in the link state advertisement header. This currently specifies the Fletcher (ISO) checksum. We could not find any reason why this checksum is stronger than the IP checksum, and it is harder to calculate. A decision to change to the IP checksum is pending. Van Jacobson offered references comparing the two checksums. In any case, the calculation of the LS checksum field is not optional. 0 will be an illegal value, regardless of which checksum is used.

7. The two current OSPF implementations reported their progress. John Moy's implementation is written and is in the debugging stage. Rob Coltun has made similar progress. The two implementations have not yet attempted to talk to each other. Performance data is expected at the next IETF.

Open Systems Routing Working Group
Chairperson: Marianne Lepp/BBN

CHARTER

Description of Working Group:

The Open Systems Routing Working Group is chartered to develop a policy-based AS-AS routing protocol that will accommodate size and general topology.

Specific Objectives:

- Architecture
- Functional Specification
- Draft Protocol Specification

Estimated Timeframe for Completion:

December 1989

Open Systems Routing Working Group
Chairperson: Marianne Lepp/BBN

STATUS UPDATE

1. Chairperson: Marianne Lepp, mlepp@bbn.com

2. WG Mailing List: open-rout-interest

3. Last meeting: teleconference at BBN and ISI March 23, 1989

4. Next Meeting: late May, to be announced

5. Pending or New Objectives:

6. Progress to date (e.g., Documents Produced):

   - IDEA 007 Requirements
   - Functional Specification
   - Architecture in draft

CURRENT MEETING REPORT

none

OSI Interoperability Working Group
Chairpersons: Ross Callon/DEC and Robert Hagens/Univ of Wisc

CHARTER

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:

Still being determined

OSI Interoperability Working Group
Chairpersons: Ross Callon/DEC and Robert Hagens/Univ of Wisc

STATUS UPDATE

1. Chairperson's: Ross Callon (DEC) callon@erlang.dec.com
   Rob Hagens (UWisc) hagens@cs.wisc.edu

2. WG Mailing List(s):

   ietf-osi@cs.wisc.edu - submissions to list
   ietf-osi-request@cs.wisc.edu - addition/deletions

3. Date of Last Meeting: Cocoa Beach, April 11-14, 1989

4. Date of Next Meeting: Stanford, July 1989

5. Pending or New Objectives:

   Write RFC for CLNP Echo. (draft by the next meeting)

   Propose mechanism for encapsulation/routing/network management of CLNP inside DoD IP for production purposes. (very rough draft of issues by the next meeting)

   Prepare the IETF-OSI "OSI documents to read" list. (ongoing)

   Prepare IETF-OSI comments on Gosip V2. (at the next meeting)

6. Progress to Date (e.g., documents produced):

   - RFC 1069 and 1070

   - We have made significant progress toward aligning the proposed Internet NSAP address format (from RFC 1069) with the Gosip NSAP address format, version 2.

   - We have started definition of an echo-request/ echo-reply function to propose as an RFC and possibly as an addendum to ISO 8473.

7. Documents Produced: RFC 1069, 1070

OSI Interoperability Working Group
Chairpersons: Ross Callon/DEC and Rob Hagens/Univ of Wisc

CURRENT MEETING REPORT
Reported by Ross Callon and Rob Hagens

AGENDA

Morning

1. Status:
   - Gosip, version 2
   - Berkeley 4.4 release

2. Echo function for CLNP (ISO 8473)

3. Alignment of OSI NSAP address proposals
   - RFC 1069
   - New proposal from NIST
   - New proposal from Boeing

Afternoon

4. Discussion: when do we really need OSI in the Internet?

5. Discussion: proposed interoperability strategies

6. Tentative plans for future WG meetings

ATTENDEES

David Borman	dab@cray.com
Ross Callon	callon@erlang.dec.com
Mike Collins	collins@ccc.mfecc.llnl.gov
Jerry Cronin	1842eeg-intg2@afcc-oal.arpa
Farokh Deboo	..!sun!bridge2!fjd
Dino Farinacci	dino@bridge2.3com.com
Lionel Geretz	lionel@salt.acc.com
Martin Gross	martin@protolaba.dca.mil
Rob Hagens	hagens@cs.wisc.edu
Tony Hain	haun@ccc.mfecc.llnl.gov
Bob Harris	harris@sparta.com
Gary Malkin	gmalkin@proteon.com
Keith McCloghrie	kzm@twg.com
Don Merritt	merritt@brl.mil
Russ Mundy	mundy@beast.ddn.mil
Zbigniew Opalka	zopalka@bbn.com
Rex Pugh	pugh%hprnd@hplabs.hp.com
Yakov Rekhter	yakov@ibm.com
Carl-herbert Rokitansky	roki@dhafeu52.bitnet
Milt Roselinsky	cmcvax!milt@hub.ucsb.edu
Dallas A. Scott	dscott@gateway.mitre.org
Jim Sheridan	jsherida@ibm.com
Lance Travis	cmt@apollo.com
Rick Wilder	rick@gateway.mitre.org

MINUTES

The meeting was convened by co-chairmen Ross Callon and Rob Hagens. An attendance list will be published with the Proceedings of the IETF. The major issues discussed at this meeting included: status reports on 4.4 BSD and GOSIP Version 2, an Echo function for CLNP, a timetable for needed OSI functionality in the Internet, and alignment of the NSAP address proposals.

Brief status reports on 4.4 BSD and GOSIP Version 2 were given by Rob Hagens.

4.4 BSD
TP4/CLNP runs over software loopback.
Next step will be testing with EON

GOSIP Version 2
A draft for public comment will be released in early May.
90 day comment period.

ECHO Function for CLNP
Rob Hagens presented his proposal for an Echo function for CLNP. The proposal was discussed and several changes to the proposal were recommended by the working group. An outline of the proposal along with the recommended changes is included below.

Mechanics
Echo and Echo Reply must be processed in the same manner as a DT PDU. EC and ECR PDUs are identical to the DT PDUs with the following exceptions: 1) type field EC 1E, ECR 1F. 2) segmentation part is always present. 3) the E/R flag is optional.

The E/R flag being optional was one of the changes recommended by the working group. The original proposal stated that the E/R flag never be set.

Receiving EC
EC processed like DT.
EC given to Echo entity
Echo entity issues an ECR with:

1. source and destination addresses reversed
2. TTL set to the normal value that the system uses when generating a DT
3. EC PDU included as data parameter of ECR
4. ER flag optionally set on ECR.

A note will be added stating that the E/R flag can be optionally set, but the usefulness of the ER pdu will depend upon the ability of the receiving system to process it.

PDU Options
All options in EC are copied in ECR.
RR option is copied but initialized since the EC RR is copied into the ECR data parameter.

The setting of the priority option was discussed and it was decided that it should not always be set to a default value of 0.

Source Routing
The group discussed whether source routing could be used as a substitution to the proposed echo function. The idea would be that to simulate echo, one could source route a packet through a remote peer and back to the originator. The first problem with source routing is that it has a bug in its specification whereby it will only work if all intermediate systems along the path implement source routing. The second problem is that the use of the source routing option will change the packet processing, thereby failing to meet the goal of treating echo packets as much as possible like normal data packets. In addition, it was pointed out that the number of source route entries is limited due to the PDU header size. In principle, it may not be possible to ping any system given the limited number of source route entries.

In addition, ECR pdus will never have a source route automatically generated.

When is OSI needed in the INTERNET?

The group discussed this issue which has been a repeated topic in several network forums. The group developed a timetable of when OSI functionality should be available on the internet.

t0: 4.4 BSD is released (beta).
CLNS subnets (isolated).

Encapsulation between isolated subnets by using smart gateways with fixed tables.

t1: 4.4 BSD is released.
t2: Intra Domain IS-IS Routing (Vendor's release).
Inter Domain IS-IS Routing (Static Tables).
Intra Domain IS-IS reaches DP status.

ALIGNMENT OF NSAP ADDRESS PROPOSALS

Ross Callon led the discussion of the three NSAP address proposals. The three proposals are outlined below.

RFC 1069

Goal
Consistent with ISO NSAP, ANSI/ISO Intra Domain Routing, Other

Intra Domain Routing, and Future Inter Domain Routing.
Flexibility.

Long Term Growth.
Result

Use ICD value assigned to DoD.
Specify high order fields.
Flexible 'IGP Specific' part.
Selector.
Pad to 20 octets.
Two IGP Specific formats specified; DoD, and ANSI/OSI.

Problems
Divergence (GOSIP and ANSI).
Too much flexibility (multiple IGP Specific Formats).

NIST PROPOSAL
Basically RFC 1069 with the following changes:

Expanded Global Area to 3 octets and renamed the field 'ADMIN AUTHORITY' (was 2 octets in RFC 1069).
Specifies one IGP Specific Format that is ANSI consistent (last 9 octets consistent with intra domain routing).
padding was moved.

BOEING PROPOSAL
AFI specifies ISO DCC for the IDI.

ie. for the US	AFI	DCC	IDI	
		39	84	0F

Not padded to 20 Octets
ORG ID 5 bytes

The working group discussed the Boeing proposal and realizes there should be more fields to meet hierarchical needs but the 5 byte ORG ID is excessive, therfore the group concluded that the proposal is not appropriate.

It was decided that the differences between the RFC 1069 and the NIST PROPOSAL would be resolved in the following manner:

1) Rewrite RFC 1069 with Global Area expanded to 3 octets and specifying one IGP Specific Format.

2) Suggest that NIST change their proposal so that the reserved section matches our padding (after routing domain).

This work should be completed by the next IETF.

AGENDA FOR NEXT IETF

The agenda for the next IETF was proposed as follows: one half day each for presentations and discussions on (i) Directory Services; (ii) Intra Domain Routing; (iii) Inter Domain Routing; and (iv) comments on Gosip version 2.

OSI WG Report

April, 1989

Status

- New RFC's
  - 1069 - NSAP Address Format
  - 1070 - EON (CLNP in DOD IP)
- BSD 4.4 Release
  - TP4/CLNP runs over loopback
  - EON testing next
- GOSIP Version 2
  - May - available for comments
  - Oct - published as new FIPS

CLNP ECHO

- Defined ECHO (EC) and ECHO-Reply (ECR) PDUs
- EC and ECR look like DT
- ECR will contain entire EC PDU
- Most Details worked out (ER)
- TODO:
  - write out procedures
  - try it
  - forword results to ANSI (as addendum to 8473)

NSAP Address Formats

- Modify RFC 1069
  - expand, rename field
  - define only 1 I6P specific part
- Propose change to GOSIP V2
  - move padding
- These changes will align <u>all</u> of the Internet and GOSIP V2 with 1 NSAP address format!
- TODO: align with ANSI?

OSI in the Internet

A Timeline

t_0 - BSD 4.4 released (β)

t_1 - Isolated CLNS subnets appear
- Fixed table routing
- CLNP encapsulation between subnets

t_2 - BSD 4.4 released

t_3 - Intra Domain IS-IS released by Vendors
- Fixed tables used for Inter-Domain IS-IS
- Dual GWs
- Inter Domain IS-IS becomes a DP

July Agenda

(Draft)

- 1½-2 Days
- Presentations on
  - Intra-Domain IS-IS
  - Inter-Domain IS-IS
  - Naming / Directory Services

PDN ROUTING WORKING GROUP
Chairperson: Carl-Herbert Rokitansky/Fern University of Hagen

CHARTER

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 were 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 Networks. (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) (almost done, RFC-Draft)

4) Assignment of INTERNET/PDN-cluster addresses to PDN-hosts and VAN-gateways according to the developed hierarchical VAN-gateway algorithms (RFC-Draft, will be completed by July '89)

5) Definition of the PDN-cluster addressing scheme as an Internet standard (to be written up as an RFC-Draft by Fall '89)

6) Specification of an X.121 Address resolution protocol (RFC-Draft, expected to be completed by July '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).

PDN Routing Working Group
Chairperson: Carl-Herbert Rokitansky/Fern University of Hagen

STATUS UPDATE

1. Chairperson: Dr. Carl-Herbert Rokitansky, Fern University of Hagen, D-5860 ISERLOHN, FRG

   E-Mail: roki@DHAFEU52.BITNET, roki@A.ISI.EDU;
   Tel: ++49/2371/566-235

2. WG Mailing List(s):

   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.

3. Date of Last Meeting: April IETF 1989, Cocoa Beach, FL

4. Date of Next Meeting: Oct 31 - Nov 2, 1989, IETF/Univ of Hawaii

   (Intensive information exchange via e-mail meanwhile)

5. Pending or New Objectives:

   PO6) Definition of the PDN-cluster addressing scheme as Internet standard: Expected to be written up as an RFC-Draft by Fall '89.

   PO7) Specification of X.25 Call Setup and Charging Determination Protocol: Functionality, data structure and state diagram have already been defined, will allow reverse charging on international (!) X.25 connections, is currently in the progress to be written up as an RFC-Draft, and is expected to be completed by Fall '89.

   PO8) Specification of an X.25 Access and Forwarding Control Scheme: Functionality and data structure have already been defined, might be included in the RFC-Draft above on X.25 Call Setup and Charging Determination Protocol, and is expected to be completed by late Fall '89.

   PO9) Specification of routing metrics taking X.25 charges into account: Proposal is expected to be written up as an RFC-Draft by Fall '89 or later.

PO10) Delayed TCP/IP header compression: As a result from Van Jacobsons presentation (IETF, April 13, 1989), a delayed version will be considered to be used on X.25 connections by VAN-gateways and PDN-hosts, as a new objective (Fall '89)

PO11) Provide a testbed for worldwide interoperability between local TCP/IP networks via the system of X.25 public data networks (PDN): Several institutes and research establishments in Europe, USA and Australia have already agreed to participate in these tests, which are expected to start in June '89.

PO12) Implementation of the required algorithms and protocols in a VAN-BoX: All the algorithms and protocols specified above are intended to be implemented in a VAN-BoX (or on a workstation) towards the end of 1989 (first test version).

PO13) Interoperability between ISO/OSI hosts on TCP/IP networks through PDN: Will be tested and demonstrated in connection with national and international PDN-tests (see below).

PO14) Consideration of Route Servers: Already discussed, but no detailed specification so far; will be considered with regard to results from international PDN-tests (see below).

PO15) Interoperability between local TCP/IP networks via ISDN: First discussions and proposals were already made, will be considered in detail in 1990.

PO16) Internetwork Management Protocols (cooperation of NOCs and NICs): Will be considered with regard to results from international PDN-tests (see above).

6. Progress to Date (e.g., documents produced):

1. PDN-cluster addressing scheme:
Rokitansky, C.-H., "Internet Cluster Addressing Scheme and Its Application to Public Data Networks", in Proceedings of the 9th International Conference on Computer Communication (ICCC'88), pp. 482-491, Editor: J.Raviv, Tel Aviv, Israel, Oct 30 - Nov 4, 1988.

2. Hierarchical VAN-Gateway Algorithms:
Rokitansky, C.-H., "Hierarchical VAN-Gateway Algorithms and PDN-Cluster Addressing Scheme for Worldwide Interoperation Between Local TCP/IP Networks Via X.25 Networks", in Proceedings of ITG/GI Conference on "Communication in Distributed Systems" (Informatik Fachberichte 205, Kommunikation in verteilten Systemen, ITG/GI Fachtagung, Stuttgart, P.J. Kuehn (Hrsg.), ISBN 3-540-50893-7

Springer-Verlag Berlin Heidelberg New York, ISBN 0-387-50893-7 Springer-Verlag New York Berlin Heidelberg), pp. 758-774, Stuttgart, Feb 22-24, 1989.

3. Assignment of INTERNET/PDN-cluster network numbers to DNICs: Rokitansky C.-H., Fern University of Hagen, "Assignment and Reservation of INTERNET Network Numbers for the PDN-Cluster", RFC-Draft, Feb 1989.

4. Assignment of default INTERNET/PDN-cluster addresses to VAN-gateways: Currently in the progress of being written up as an RFC-Draft, expected to be completed by July '89, (might be included into the RFC-Draft above).

6. X.121 Address Resolution Protocol (first version has already been written up as an RFC-Draft to be discussed between members of the PDN Routing WG, and is expected to be completed by July '89).

PDN Routing Working Group
Chairperson: Carl-Herbet Rokitansky/Fern University of Hagen

CURRENT MEETING REPORT
Reported by Carl-Herbert Rokitansky

AGENDA

- Introduction
- Background information (European situation, X.25 Research Network)
- Report from CeBIT MultiNET TCP/IP demonstration on Hannover Fair, March 8-15, 1989 (Carl-H. Rokitansky, FernUni)
- Status report on BBN-VAN-GATEWAY (butterfly replacement, EGP, etc.) (Mike Brescia, BBN) -Discussion
- Status and technical discussion of short term goals
- Assignment and reservation of PDN-cluster network numbers to national X.25 public data networks (DNICs), RFC Draft (Roki)
- Assignment of INTERNET IP addresses to VAN-gateways according to the developed hierarchical VAN-gateway algorithms (Roki)
- X.121 address resolution protocol, RFC Draft (Mike Brescia, Roki)
- Access control and reverse charging on international X.25 connections
- Discussion on assignment of an autonomous system number to PDN
- Discussion on a modified EGP and routing metrics to be used between VAN-gateways
- Discussion of methods and requirements involving route servers
- Discussion on the application of the INTERNET cluster addressing scheme and developed gateway algorithms to Integrated Services Digital Networks (ISDN) to provide interoperability between local TCP/IP networks through ISDN (C. Rokitansky, FernUni)
- Coordination of PDN Routing performance tests
- Discussion on documents to be published by members of the PDN Routing WG
- Assignment of action items
- Miscellaneous (mailing lists, etc.)

ATTENDEES

Jerry Cronin	2eeg-intg2@AFCC-OA1.ARPA
Farokh Deboo	..!sun!bridge2!fjd
Lionel Geretz	lionel@SALT.ACC.COM
Martin Gross	martin@PROTOLABA.DCA.MIL
John Lekashman	lekash@ORVILLE.NAS.NASA.GOV
John Moy	jmoy@PROTEON.COM
Bill Nowicki	nowicki@SUN.COM
Zbigniew Opalka	zopalka@BBN.COM
Carl-Herbert Rokitansky	roki@DHAFEU52.BITNET
Mary Stahl	stahl@SRI-NIC.ARPA
Rick Wilder	rick@GATEWAY.MITRE.ORG

MINUTES

X.25 Research Network (reported by C.-H. Rokitansky (Roki), Fern University):
Towards the end of this year an "X.25 Research Network" will be installed and operated by the German PTT. A large number of German universities and research institutes will be connected to this X.25 Research Network at fixed costs. A gateway to the German DATEX-P network will allow interoperation with the worldwide system of X.25 Public Data Networks (PDN).

Since most universities do have local TCP/IP networks, they are especially interested in exchanging TCP/IP datagrams with each other through this X.25 research network. An advantage of this kind of X.25 network is the fact, that the exchange of network reachability information between hosts/gateways attached to the X.25 research network will NOT BE COST SENSITIVE, although we should try, of course, to limit the amount of such messages.

Similar X.25 research networks will be made operational in other European countries, and they will probably be interconnected to an European X.25 research network one day.

Report from CeBIT MultiNET (Mar '89) TCP/IP Demo on Hannover Fair (by Roki):
"The CeBIT MultiNET '89 demonstration of system-independent networks was a further development of the two preceding MultiNET shows. MultiNET's main object has remained unchanged: the joint presentation of international communication standards with products already available by a variety of EDP suppliers." (see slides presented at the Plenary) "Like before, CeBIT MultiNET '89 demonstrates the current state of communications technology. ... After a careful analysis of the actual needs of many network users, the following essential topics have been choosen for the CeBIT MultiNET '89: 1. Coexistence and migration from TCP/IP to "ISO"; 2. Network Management; 3. Network Applications. ... The clearly defined range of functions together with the popular program interfaces and the broad use of the TCP/IP protocol family will surely, continue to stimulate the porting of network applications onto these communcication protocols, just as it has done up to now. Recently, a new algorithm for TCP/IP has been published, which offers considerable performance advantages for the protocol handling and which shows that TCP/IP is still being further developed. Some of the newer implementations are already based on this new algorithm." (Extract from "CeBIT-MultiNET", MultiNET Services Gesellschaft fuer Tele- und Datenkommunikation mbH, Venloer Strasse 131, D-5024 Pulheim, FRG.)

As can be seen from the slides, 32 suppliers demonstrated TCP/IP interoperability and ISO/OSI coexistence. The relatively great number of suppliers offering TCP/IP via X.25 (16), IP Router (17), and Gateways (17), is of special interest for PDN Routing WG, although,of course, no dynamic routing of TCP/IP datagrams through the system of

X.25 Public Data Networks is provided. Most suppliers expect an even increasing demand for the TCP/IP protocol suite within the next years.

Status Report on BBN-VAN-GATEWAY (Zbigniew Opalka, BBN):

The current LSI-11/23 BBN-VAN-GATEWAY, will be replaced by a butterfly gateway by late May '89. EGP will be available by then.

Hierarchical VAN-Gateway Algorithms (Roki):

The concept of hierarchical VAN-gateway algorithms, which allow the distribution of worldwide INTERNET network reachability information within a few number of hops in a very efficient way, were discussed in detail. A four level hierarchy is defined so far: 1. LOCAL-VAN; 2. DATA-NETWORK- VAN; 3. COUNTRY-VAN; 4. ZONE-VAN. An additional level (REGION-VAN) might be specified between COUNTRY-VAN and ZONE-VAN. An advantage of the proposed algorithms is, that each VAN-gateway might have several direct neighbors, but there is ONE (!) only, which it has to call ACTIVELY for exchanging worldwide network reachability information. It is important to understand, that the whole data traffic will not necessarily be routed via the higher level VAN-gateways (level 2 to 4, which might be regarded as route servers in this case), but can be exchanged via a direct X.25 connection (SVC) between LOCAL-VANs and/or PDN-hosts. The concept also provides for stupid LOCAL-VANs, which do not maintain worldwide network reachability tables, but use their DATA-NETWORK-VAN as a default gateway. If this gateway knows a better route towards the destination network through the PDN, it sends an ICMP-Redirect message (similar to an ICMP Host Redirect message) to the calling LOCAL-VAN, specifying the INTERNET address of the next hop VAN-gateway, which can be any address in the PDN-cluster, (as an significant advantage of the INTERNET cluster addressing scheme). John Moy, Proteon, suggested to call this message "ICMP Gateway Redirect Message".

Assignment and Reservation of PDN-Cluster Network Numbers to DNICs (Roki):
This RFC draft contains a proposal for the assignment of INTERNET network numbers to existing X.25 national public data networks according to the proposed PDN-cluster addressing scheme, taking the structure of the X.121 international numbering plan for public data networks (6 zones worldwide) into account. The reservation (1024 class B network numbers for the PDN-cluster) and the assignment is based on the expected growth of national public data networks in the various countries, and was done with regard to a specification of subclusters of the PDN-cluster (Europe-cluster, North America-cluster, North-Europe-cluster, etc.). Where possible, adjacent countries were assigned adjacent networks or subclusters, having in mind some kind of subcluster oriented (cartesian) routing algorithms for future use. A direct mapping between the INTERNET

network number and the corresponding DNIC can be performed by an algorithm for all US PDN-cluster network numbers within the North-American subcluster.

The following Internet network numbers are reserved for the PDN cluster:

188.000 - 188.255	Europe	(Zone 2),	assigned: 41
189.000 - 189.255	North America	(Zone 3),	assigned: 37
190.000 - 190.127	Asia	(Zone 4),	assigned: 40
190.128 - 190.191	<reserved>		
190.192 - 190.255	Pacific	(Zone 5),	assigned: 20
191.000 - 191.127	South America	(Zone 7),	assigned: 40
191.128 - 191.191	<reserved>		
191.192 - 191.255	Africa	(Zone 6),	assigned: 8

A list of already assigned PDN-cluster network numbers is shown in the slides presented in the report at the Plenary.

Assignment of INTERNET IP Addresses to VAN-Gateways and PDN-Hosts:

According to the hierarchical VAN-gateway algorithms and the PDN-cluster addressing scheme, a proposal for the assignment of INTERNET addresses to PDN-hosts and VAN-gateways was presented by Roki and discussed in detail. This scheme allows to address a maximum number of 32.512 PDN-hosts [p.n.1.h] - [p.n.127.h] and 32.512 LOCAL-VANs [p.n.128.v] - [p.n.254.v] in each national public data network to which an INTERNET (class B) PDN-cluster network number [p.n.r.r] has been assigned. In addition, 255 INTERNET addresses in each PDN network are reserved for higher level VAN-gateways:

- DATA-NETWORK VANs [p.n.255.1] - [p.n.255.63] [p.n.255.64] - [p.n.255.127] reserved for future use
- COUNTRY VANs [p.n.255.128] - [p.n.255.159] [p.n.255.160] - [p.n.255.191] reserved for future use
- REGION VANs [p.n.255.192] - [p.n.255.207] reserved for future use [p.n.255.208] - [p.n.255.223] reserved for future use
- ZONE VANs [p.n.255.224] - [p.n.255.231] and
- (ZONE VANs) [p.0.255.224] - [p.0.255.231] "escape code" [p.n.255.232] - [p.n.255.239] reserved for future use
- WORLD VANs [p.n.255.240] - [p.n.255.243] reserved for future use [p.n.255.244] - [p.n.255.254] reserved for future use [p.n.255.255] <reserved> Advantages of this hierarchical addressing scheme are:

- a bitmask can be used to distinguish between a PDN-host and a LOCAL-VAN

- a bitmask can be used to distinguish between PDN-hosts/LOCAL VANs and higher level VANs (level 2 VANs and up = [p.n.255.v])

- a bitmask can be used to determine the functionality of each VAN-gateway

- each PDN-host or LOCAL VAN can determine automatically its default DATA NETWORK VAN (next higher level) [p.n.255.1]

- each (higher level) VAN-Gateway (and PDN-host) can determine automatically its default top-level (ZONE VAN) by an "escape code". [p.0.255.224]

Access Control and Reverse Charging on International X.25 Connections:

An access control scheme, proposed by Roki, has been discussed in detail with the group, and it was agreed that some access control in connection with charging determination is required for the routing of TCP/IP datagrams through the PDN, if we provide worldwide interoperability. As a result from this discussion, a modified and more flexible "X.25 Access Control and Forwarding Scheme" was worked out after the meeting, and was presented in the PDN Routing WG report at the plenary (see slides).

Discussion on a Modified EGP and Routing Metrics to be Used Between VANs:

Since there was not enough time to discuss this issue during the PDN Routing WG meeting, usage of either EGP2 or EGP3 was discussed with Marianne Lepp, Zbigniew Opalka, Roki and others at a "working lunch": Depending on the support for EGP3, either a modified version of EGP2 or EGP3 will be considered to be used by VAN-gateways to exchange worldwide network reachability information (eventually on an event driven basis).

Implementation of the Proposed Algorithms in a VAN-BoX (or on a Workstation):

Fortunately, within the last year, the IETF-PDN Routing working group has developed most of the required PDN addressing schemes and gateway algorithms to allow a dynamic routing of TCP/IP datagrams through the worldwide system of X.25 Public Data Networks (PDN). The required algorithms and protocols include:

- PDN-cluster addressing scheme: published in ICCC'88 proceedings;
- Hierarchical VAN-gateway algorithms: published in ITG/GI '89 Proc.

- Assign. áRes. of PDN-cluster net #: draft, ready to be published as RFC
- Assign. áRes. of PDN-cluster addr.: currently being written up as an RFC
- X.121 Address Resolution Protocol: draft, will be published as an RFC
- X.25 Call Setup and Charging Determ.: draft, currently being specified
- X.25 Access and Forwarding Control: draft, currently being specified
- Modified EGP2 or EGP3 between VANs: currently in progress to be defined
- Delayed TCP/IP header compression: will be considered (new objective)

By putting all these pieces together, it is intended to implement these algorithms, with support of the gateway companies (BBN, Proteon, SUN, 3COM, ACC, cisco, etc.), in a small "VAN-BoX" (and on a workstation) with an Ethernet and an X.25 interface. By placing this "VAN-BoX" between a local TCP/IP network and an X.25 public data network, the implemented gateway algorithms will automatically exchange network reachability information to provide worldwide INTERNET interoperability between local TCP/IP networks through X.25 Public Data Networks.

Application of the INTERNET Cluster Addressing Scheme to ISDN (Roki):

The application of the INTERNET cluster addressing scheme and the developed gateway algorithms to Integrated Services Digital Networks (ISDN) to provide interoperability between local TCP/IP networks through ISDN has been discussed briefly.

E.164 specifies the numbering plan for the ISDN era. According to this numbering plan, the International ISDN Number (max. 15 digits) consists of the Country Code (CC), the National Destination Code (NDC) and the Subscriber Number (SN). NDC and SN form the National Significant Number (NSN):

<International ISDN Number> ::= <CC><NSN>; <NSN> ::= <NDC><SN>; or <International ISDN Number> ::= <CC><NDC><SN>.

The INTERNET cluster addressing scheme could be applied to the ISDN by mapping cluster nets to CC or NDC.

Internetwork scenarios for PDN-ISDN-PDN and ISDN-PDN-ISDN using X.121 or E.164 escape codes were discussed. Also the long-term PDN to/from ISDN solution using numbering identifiers was sketched.

Coordination of International PDN Routing Performance Tests:

The developed PDN addressing schemes and VAN-gateway algorithms will be tested with participating sites in the following countries:

Zone 2 (Europe):

Germany:	Fern University of Hagen (all VAN-gateway levels)
	GMD, St. Augustin (DFN-Gateway)
	University of Dortmund (UUCP-Gateway)
	University of Karlsruhe (BELWUE)
	University of Stuttgart (BELWUE)
Austria:	University of Salzburg
Finland:	University of Helsinki (NORDUNET)
Italy:	CNUCE, Pisa *
Norway:	NTARE, Oslo, (NORDUNET) *
Sweden:	SICS, Stockholm (NORDUNET)
UK:	Portsmouth Polytechnic
	University College London (INTERNET Gateway) *

Zone 3 (North America):

USA:	BBN, Cambridge, MA
	CISCO, Menlo Park, CA *
	PROTEON *
	SRI, Menlo Park, CA *
	SUN, Mountain View, CA *

Zone 4 (Asia): Israel ?, Japan ?

Zone 5 (Pacific):
Australia: CSIRO
Indonesia: LAPAN

Zone 6 (Africa): Egypt ?

Zone 7 (South America): Argentina ?, Brazil ?

(* ... intended, but not yet agreed)
(? ... these countries will be contacted for participation, to have at least one representative site for each zone).

First tests have already started within Germany. International PDN-tests are expected to start in June '89 between BBN and sites in Europe and Australia.

Assignment of action items:

Stahl: Check assignment and specification of INTERNET/PDN-cluster network numbers for US national public data networks for correctness (03, June '89)

Roki: Submit Internet Draft "Assignment and Reservation of the INTERNET Network Numbers for the PDN-Cluster" to IETF Chair and Reviewers (03, July '89).

Roki: Finish Internet Draft "Addressing Scheme for the Assignment of INTERNET/PDN-Cluster Addresses to VAN-Gateways and PDN-Hosts" for submission to the IETF Chair and Reviewers (O4, July '89).

Opalka/ Finish Internet Draft "X.121 Address Resolution Protocol",
Roki: for submission to the IETF Chair and Reviewers (O6, July '89).

Roki: Write up Internet Drafts "INTERNET Cluster Addressing Scheme", and "Application of the INTERNET Cluster Addressing Scheme to X.25 Public Data Networks" for submission to the IETF Chair and Reviewers (O5, Fall '89).

Roki: Write up an Internet Draft "Hierarchical VAN-Gateway Standards for Worldwide INTERNET Interoperability" for submission to the IETF Chair and Reviewers (O5, Fall '89 or later).

Roki: Continue the specification of an Internet Draftt "X.25 Call Setup and Charging Determination Protocol" (O7, expected to be completed by Fall '89)

Roki: Continue the specification of an Internet Draft "X.25 Access and Forwarding Control Scheme" (O8, expected to be completed by Fall '89 or later)

Opalka/ Perform international PDN-tests according to the developed PDN- Roki: cluster addressing scheme and hierarchical VAN-gateway algorithms between USA (BBN) and sites in Europe (Fern University of Hagen, University of Dortmund, University of Salzburg, etc.), starting June '89 (O11).

A closed PDN Routing WG meeting was held during the afternoon session of April 12, where some of the PDN Routing objectives and action items were discussed in more detail.

Agenda of PDN Routing WG Meeting, Apr'89

- German/European Networking Situation
- Report from CeBIT MultiNET TCP/IP Demo
- Status Report on BBN-VAN-GW (B-fly, EGP)
- Status of short term goals
- Assignment of PDN-cluster net numbers
- Assignment of PDN cluster IP addresses to VAN-GWs according to hierarchical schemes
- X.121 address resolution protocol
- Access control and reverse charging (X.25)
- Autonomous system number for PDN
- Modified EGP to be used between VAN-GWs
- Requiremnts involving route servers
- Application of the cluster addressing scheme to ISDN for TCP/IP interoperation
- PDN Routing performance tests
- Documents to be published by PDN WG
- Assignment of action items

GERMAN X.25 RESEARCH NETWORK

- OPERATIONAL 1989/90
- FOR MORE THAN HUNDRED
  UNIVERSITIES
  RESEARCH ESTABLISHMENTS
- AT FIXED COSTS
  64 kBit/s ~ US $ 3.000.-/month
  9.6 kBit/s ~ US $ 900.-/month
- HUNDREDS OF LOCAL TCP/IP NETWORKS CAN BE INTERCONNECTED
- INTERCONNECTION WITH OTHER EUROPEAN X.25 RESEARCH NETWORKS AND PDN
- PDN ROUTING AND VAN-GATEWAY ALGORITHMS REQUIRED !

Produktübersicht der am MultiNET angeschlossenen Firmen

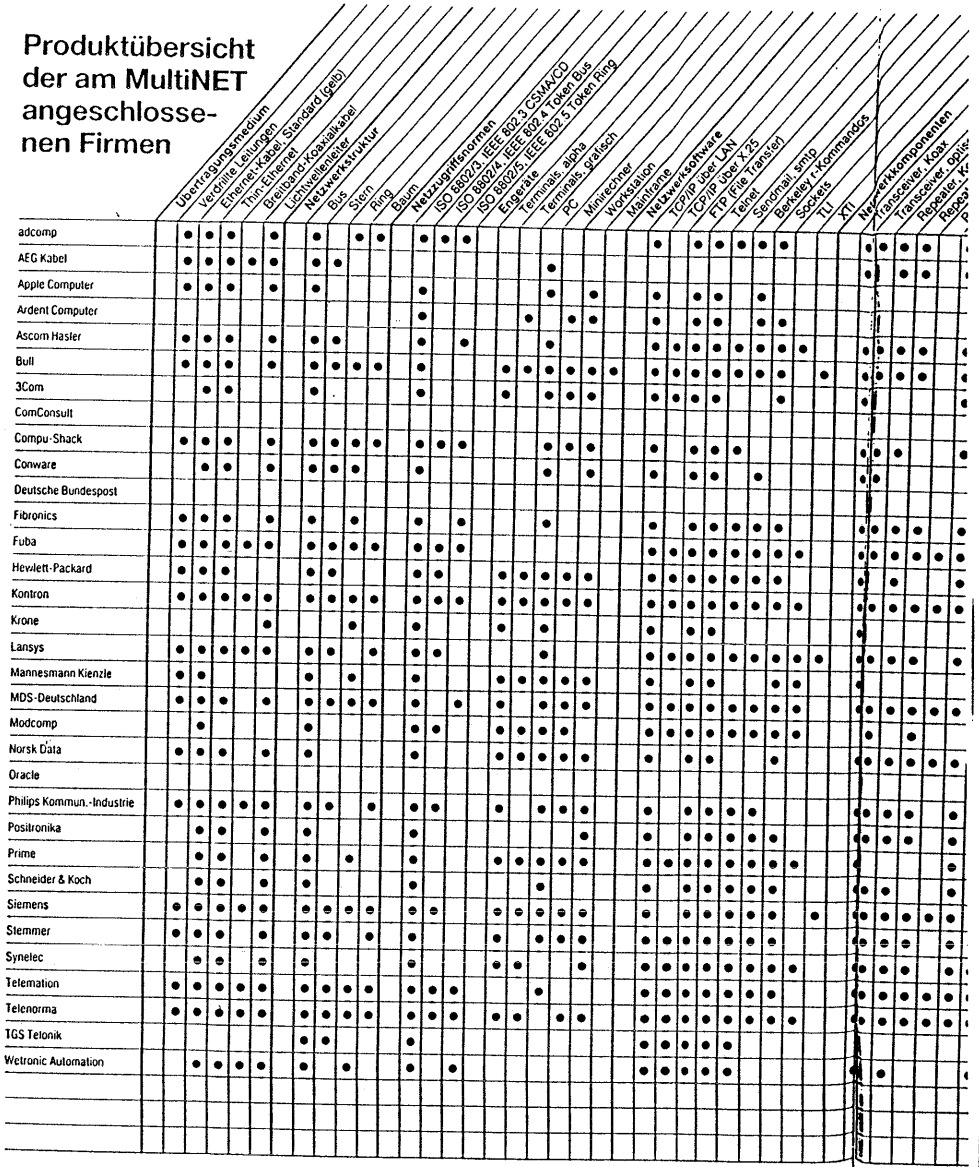

EXTRACTED FROM CeBit - MULTINET March '89

EXTRACTED FROM CeBit - MULTINET March '89

WORLDWIDE NETWORK REACHABILITY INFORMATION EXCHANGE BETWEEN VAN-GATEWAYS

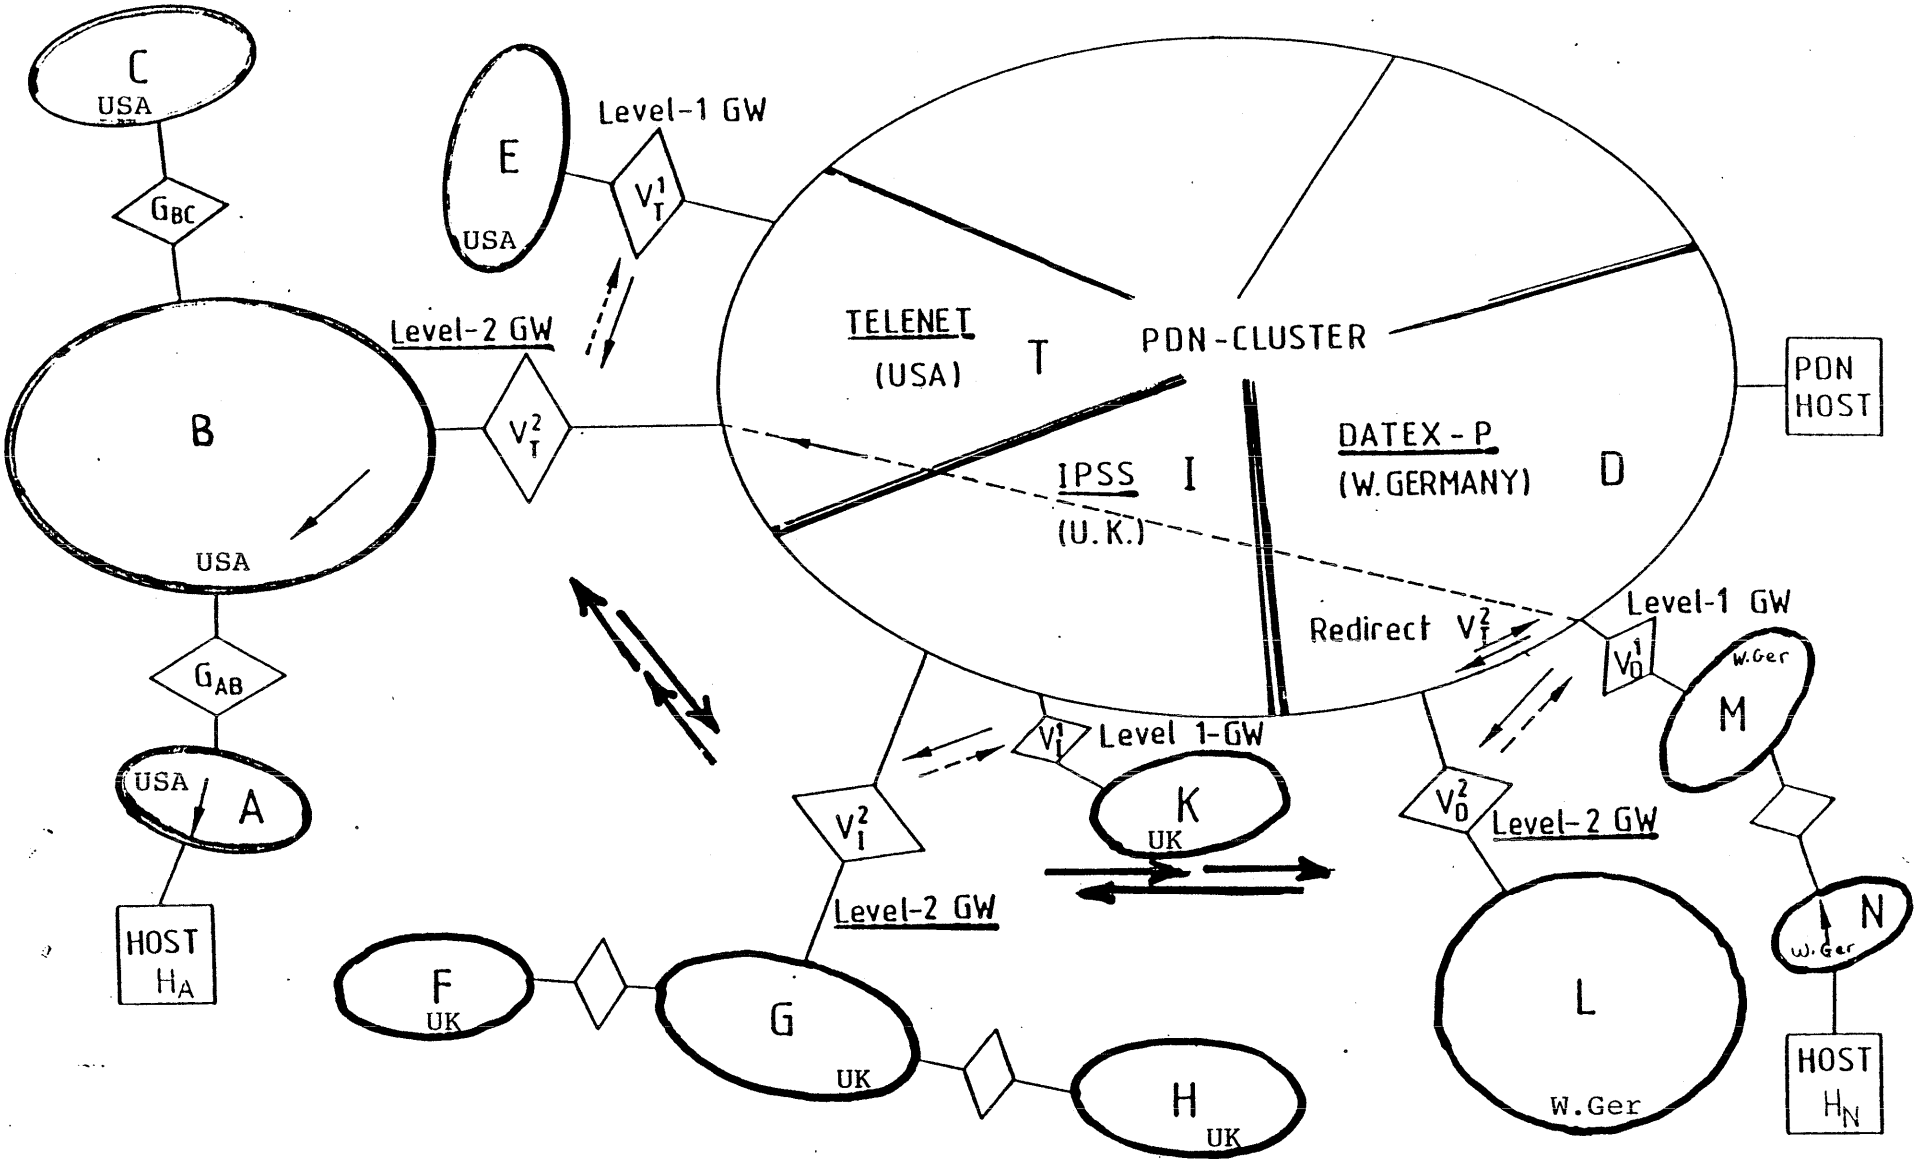

Assigned DNICs and PDN-Cluster Reservation

Zone	Area	assign.	reserve	spare
2	(Europe) :	49	256	207
3	(North America):	39	256	217
4	(Asia) :	40	192	152
5	(Pacific) :	20	64	44
6	(Africa) :	8	64	56
7	(South America):	44	192	148
Total		200	1024	824

PDN-Cluster Net Number Assignment to National Public Data Networks (DNICs)

PDN-Net	DCC	DNIC	Public Data Network	Country
188.001	262	2624	DATEX-P	Germany
188.016	270	2704	LUXPAC	Luxemburg
188.027	206	2062	DCS	Belgium
188.032	204	2041	DATANET1	Netherlands
188.033	204	2044	DABAS	Netherlands
188.042	238	2382	DATAPAK INT	Denmark
188.043	238	2383	DATAPAK NAT	Denmark
188.048	242	2422	DATAPAK	Norway
188.049	242	2427	DATAPAK	Norway
188.059	240	2402	DATAPAK	Sweden
188.064	234	2341	IPSS	Great Britain
188.065	234	2342	PSS	Great Britain
188.075	235	2352	TELEMATIC	Great Britain
188.080	272	2724	EIRPAC	Eire
188.088	274	2740	ICEPAC	Iceland
188.091	290	2901	KANUPAC	Greenland
188.096	268	2680	TELEPAC	Portugal
188.097	268	2682	DATACESS	Portugal
188.106	214	2141	NIDA	Spain
188.107	214	2145	IBERPAC	Spain
188.112	208	2081	NTI	France
188.138	232	2322	DATEX-P	Austria
188.139	232	2329	RADAU	Austria
188.144	222	2222	ITAPAC	Italy
188.145	222	2227	ITAPAC	Italy
188.155	292	2922	X-NET SMP	San Marino
188.160	202	2022	HELPAK	Greece
188.192	220	2200	YUPAC	Yugoslavia
188.201	220	2209	YUPAC	Yugoslavia
188.208	216	2160	NEDIX	Hungary
188.240	250	2502	IASNET	USSR
188.250	244	2441	Teletex	Finland
188.251	244	2442	DATAPAK	Finland
189.004	302	3020	DATAPAC	Canada
189.005	302	3025	GLOBEDAT-P	Canada
189.004	302	3028	INFOGRAM	Canada
189.004	302	3029	INFOSWITCH	Canada
189.033	310	3101	PTN-1 or WUTCO ?	USA
189.034	310	3102	MCI-DATA-TRANSPORT	USA
189.035	310	3103	ITT-UDTS II	USA
189.036	310	3104	MCI-MAIL	USA
189.038	310	3106	TYMNET	USA
189.039	310	3107	ITT-UDTS I	USA
189.042	311	3110	TELENET or UNINET ?	USA
189.045	311	3113	RCA-LSDS	USA
189.046	311	3114	UNINET ?	USA
189.051	311	3119	TRT DATAPAK	USA
189.056	312	3124	PSTS (FTCC)	USA
189.057	312	3125	UNINET ?	USA
189.058	312	3126	ADP-AUTONET	USA
189.059	312	3127	TELENET ?	USA
189.064	313	3132	COMPUSERVE	USA
189.066	313	3134	AT&T ACCUNET	USA
189.068	313	3136	MARKNET	USA
189.069	313	3137	CSC-INFONET	USA
189.071	313	3139	NETEXPRESS	USA
189.072	314	3140	SNET	USA
189.074	314	3142	BELL SOUTH	USA
189.077	314	3145	PACIFIC BELL	USA

189.079	314	3147	DIGIPAC or US WEST ?	USA
189.082	315	3150	GLOBENET	USA
189.224	334	3340	TELEPAC	Mexico
189.231	364	3640	BATELCO	Bahamas
189.236	338	3380	JAMANTEL	Jamaica
189.240	370	3700	?	Dominic. Rep.
189.243	330	3300	UDTS	Puerto Rico
189.244	340	3400	DOMPAC	Fr.-Antillen/Martinique/Guadeloupe
189.249	374	3740	TEXTEL	Tobago
189.250	374	3745	DATANET	Tobago
189.252	350	3503	C&W	Curacao
190.016	460	4600	PKTELCOM	China
.				
190.025	460	4609	PKTELCOM	China
190.030	487	4872	PACNET	China (Taiwan)
190.031	487	4877	UDAS	China (Taiwan)
190.032	425	4250	ISRANET	Israel
.				
190.041	425	4259	ISRANET	Israel
190.044	424	4243	EMDAN	United Arabian Emirates
190.047	426	4263	BAHNET ?	Bahrain
190.050	419	4190		Kuwait
.				
190.059	419	4199		Kuwait
190.101	454	4542	INTELPAK	Hongkong
190.102	454	4544	PSDS	Hongkong
190.103	454	4545	DATAPAK	Hongkong
190.112	440	4401	DDX-P	Japan
190.113	440	4408	VENUS-P	Japan
190.120	450	4501	DACOMNET	Korea (Rep.)
190.192	520	5200	IDAR	Thailand
.				
190.201	520	5209	IDAR	Thailand
190.207	515	5156	EASTNET	Philippines
190.208	502	5021	MAYPAC	Malaysia
190.215	525	5252	TELEPAC	Singapur
190.224	510	5101	INDOSAT	Indonesia
190.232	505	5052	AUSTPAC	Australia
190.233	505	5053	DAS	Australia
190.234	505	5054	?	Australia
190.240	546	5460	TOMPAC-NC	New Caledonia
190.247	547	5470	TOMPAC-PF	French-Polynesia
190.248	530	5301	PACNET	New Zealand
191.016	708	7080	HONDUTEL	Honduras
191.024	712	7122	RACSA-DATOS	Costa Rica
191.032	714	7140	INTELPAQ	Panama
.				
191.041	714	7149	INTELPAQ	Panama
191.054	732	7320	DAPAQ	Columbia
.				
191.063	732	7320	DAPAQ	Columbia
191.072	716	7160	ENTEL	Peru
.				
191.081	716	7169	ENTEL	Peru
191.088	730	7302	E-COM	Chile
191.089	730	7303	CHILEPAC	Chile
191.090	730	7305	TOMNET	Chile
191.096	722	7220	ARPAC	Argentinia
191.097	722	7222	ARPAC	Argentinia
191.108	724	7240	INTERDATA	Brasilia
191.109	724	7241	RENPAC	Brasilia
191.123	742	7420	DOMPAC	French-Guayana
191.192	602	6020	ARENTO	Egypt
191.199	605	6050	?	Tunesia
191.216	612	6122	SYTRANPAC	Ivory Coast
191.224	628	6282	GABONPAC	Gabon
191.232	655	6550	SAPONET-P	South Africa
191.233	655	6559	SAPOPAC	South Africa
191.251	617	6170	MAURIDATA	Mauritius
191.252	647	6470	DOMPAC	Reunion

HIERARCHICAL VAN-GATEWAY ALGORITHMS FOR PDN-CLUSTER

Level	Gateway
4	ZONE - GW
3	COUNTRY-GW
2	DATA NETWORK-GW
1	LOCAL-VAN-GW

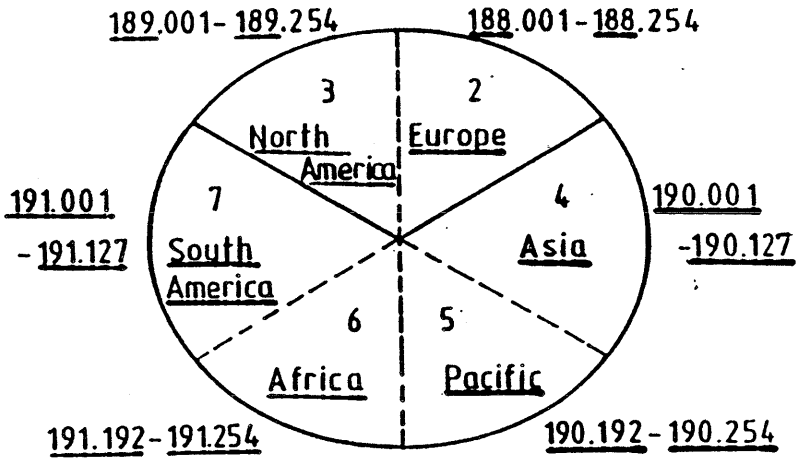

PDN-Cluster

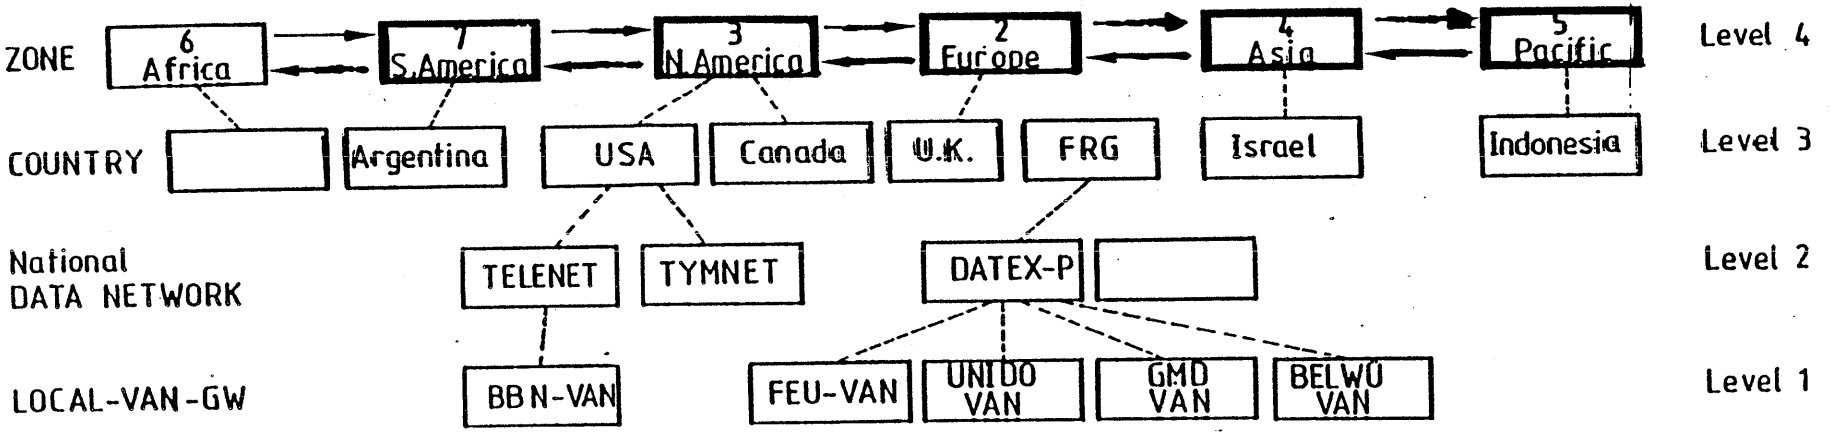

HIERARCHICAL PDN-HOST/VAN-GW IP ADDRESS ASSIGNMENT

IP ADDRESS	LEVEL	TYPE	MAX.#	MSB 31..16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0 LSB
p.n. 0.h		<reserved>		
p.n. 1.h : p.n.127.h	0	PDN-HOST	32.512	x..x 0 x.....x
p.n.128.v : p.n.254.v	1	LOCAL VAN	32.512	x..x 1 x.....x 0 x.....x
p.n.255.0		<reserved>		
p.n.255.1 : p.n.255.63	2	DATA NETWORK VAN	63	x..x 1111111 1 00 x.....x
p.n.255.64 : p.n.255.127		<reserved>		
p.n. : p.n.	3	COUNTRY VAN	32	x..x 1111111 1 00 x.....x
p.n.255.160 : p.n.255.191		<reserved>		
p.n.255.192 p.n.255.207	4	REGION VAN	16	x..x 1111111 1 10 x.....x
p.n.255.208 : p.n.255.223		<reserved>		
p.n.255.224 : p.n.255.231	5	ZONE VAN	8	x..x 1111111 1 1 1 0 x.....x
p.n.255.232 p.n.255.239		<reserved>		
p.n. : p.n.	6	. . .	4	x.x 1111111 1 1111110 xxx
p.n.255.244 p.n.255.255		<reserved>		

HIERARCHICAL VAN-GATEWAY IP ADDRESSES (EXAMPLE)

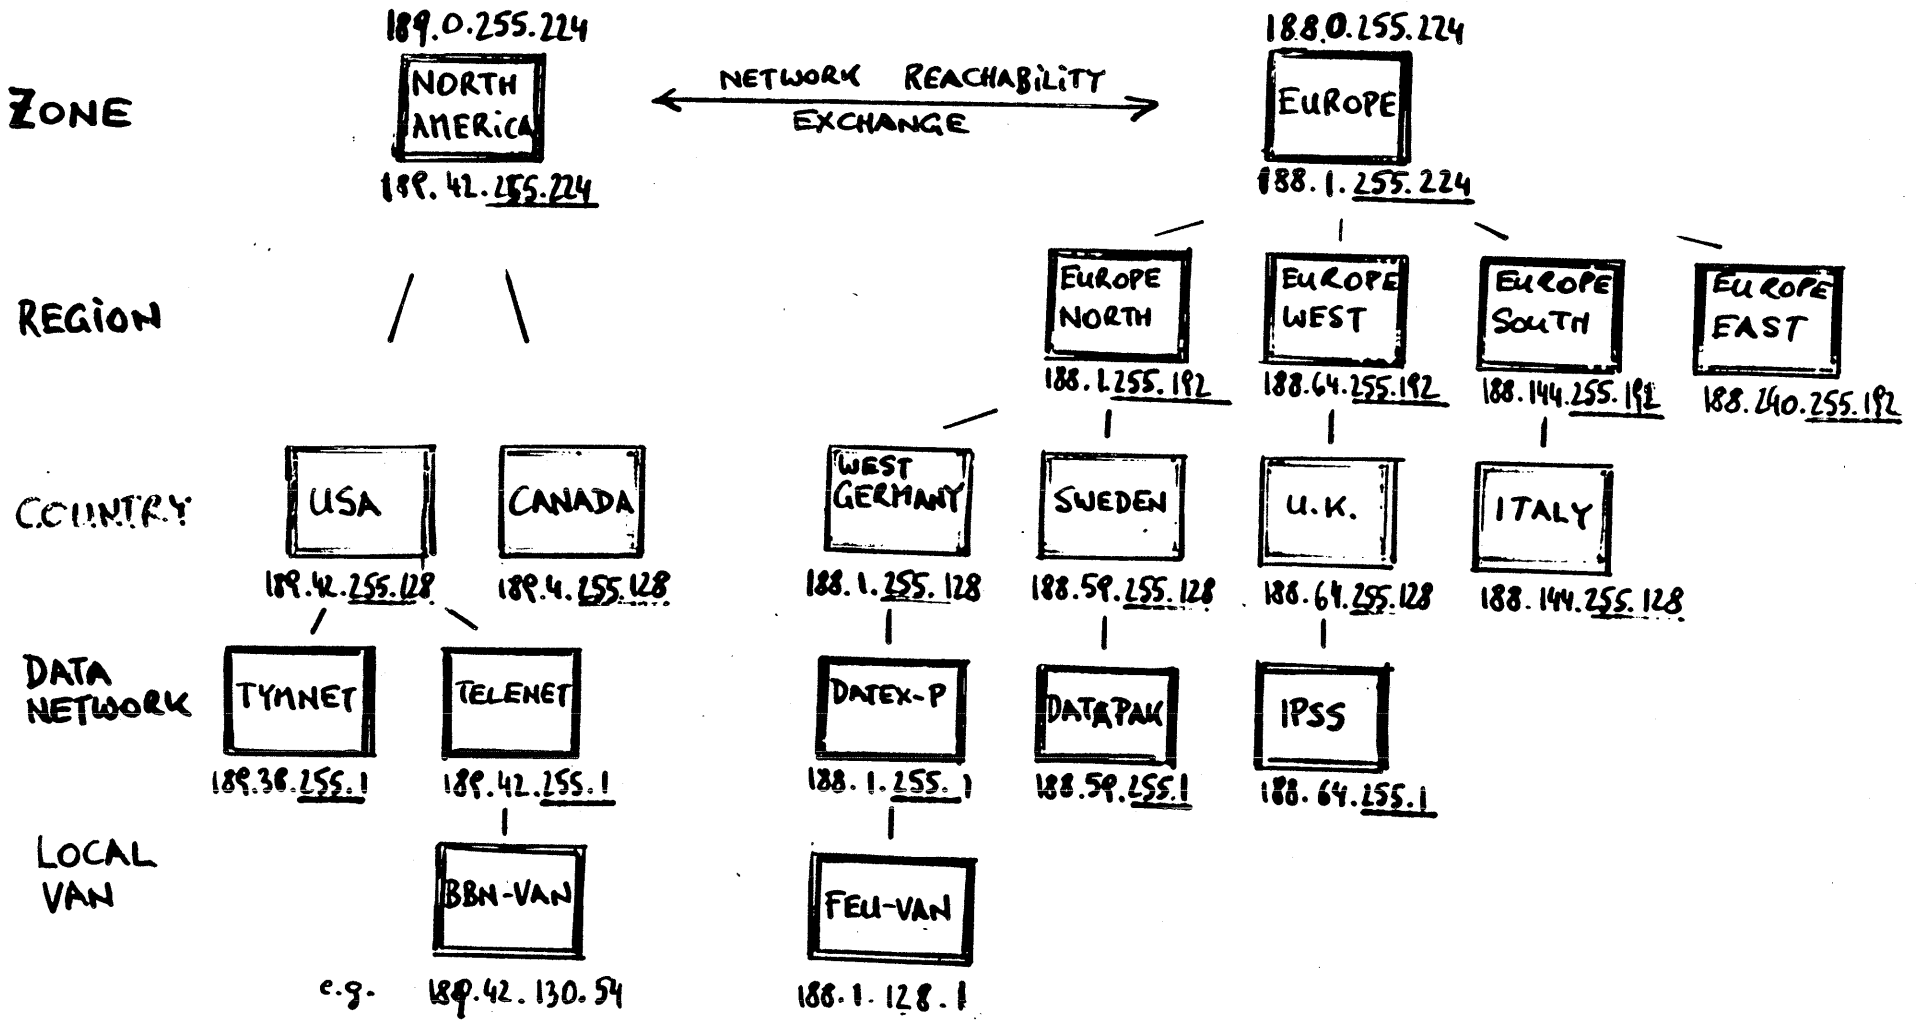

PDN X.25 ACCESS AND FORWARDING CONTROL SCHEME

CHECK:

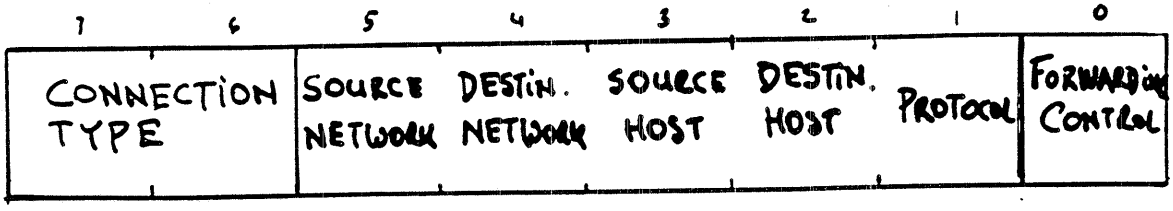

CONNECTION TYPES (Bit 6,7)

7 6

00 INCOMING CALL

01 OUTGOING CALL, REVERSE CHARGING REQUESTED

10 INCOMING CALL, REVERSE CHARGING REQUESTED

11 OUTGOING CALL,

↑ ↑

CHARGE BIT — INCOMING: 0, OUTGOING: 1

CHECK (Bit 1-5)

0 = DON'T CHECK

1 = CHECK

5 4 3 2 1

1 . . . . SOURCE NETWORK

. 1 . . . DESTINATION NETWORK

. . 1 . . SOURCE HOST

. . . 1 . DESTINATION HOST

. . . . 1 PROTOCOL

FORWARDING CONTROL (Bit 0)

0

0 . . . . CHECK ONLY WHEN ESTABLISHING X.25 CONNECTION

1 . . . . CONTINUE CHECKING ON FORWARDING PACKETS

X.25 ACCESS CONTROL & REVERSE CHARGING EXAMPLE

	NORMAL; check:	REVERSE CHARGING REQUESTED; check:
OUTGOING CALL	DESTINATION NETWORK SOURCE HOST 11011001	DESTINATION NETWORK SOURCE NETWORK 01110000
INCOMING CALL	SOURCE NETWORK 00100000	SOURCE HOST DESTINATION HOST PROTOCOL 10001111

PROVIDE SOURCE IP ADDRESS AND DESTINATION IP ADDRESS, PROTOCOL IN USER DATA FIELD OF CALL SETUP (Code CD).

International PDN-Tests:

Zone 2 (Europe):

Germany:	Fern University of Hagen
	GMD, St. Augustin (DFN-Gateway)
	Uni Dortmund (UUCP-Gateway)
	Uni Karlsruhe (BELWÜ)
	Uni Stuttgart (BELWÜ)
Austria:	University of Salzburg
Finland:	University of Helsinki (NORDUNET)
Italy:	CNUCE ?
Norway:	NTARE ? (NORDUNET)
Sweden:	SICS (NORDUNET)
UK:	Portsmouth Polytechnic
	UCL ?

Zone 3 (North America):

USA:	BBN, Cambridge, MA
	CISCO, Menlo Park, CA
	ISI ?, Los Angeles, CA
	PROTEON,
	SRI, Menlo Park, CA
	SUN, Mountain View, CA

Zone 4 (Asia):
Israel ?, Japan ?

Zone 5 (Pacific):
Australia: CSIRO
Indonesia: LAPAN

Zone 7 (South America):
Argentinia ?, Brasilia ?

Rokitansky/FernUni Hagen, FEB 1989

WORLDWIDE INTERNET / PDN - ROUTING

VAN - BOX

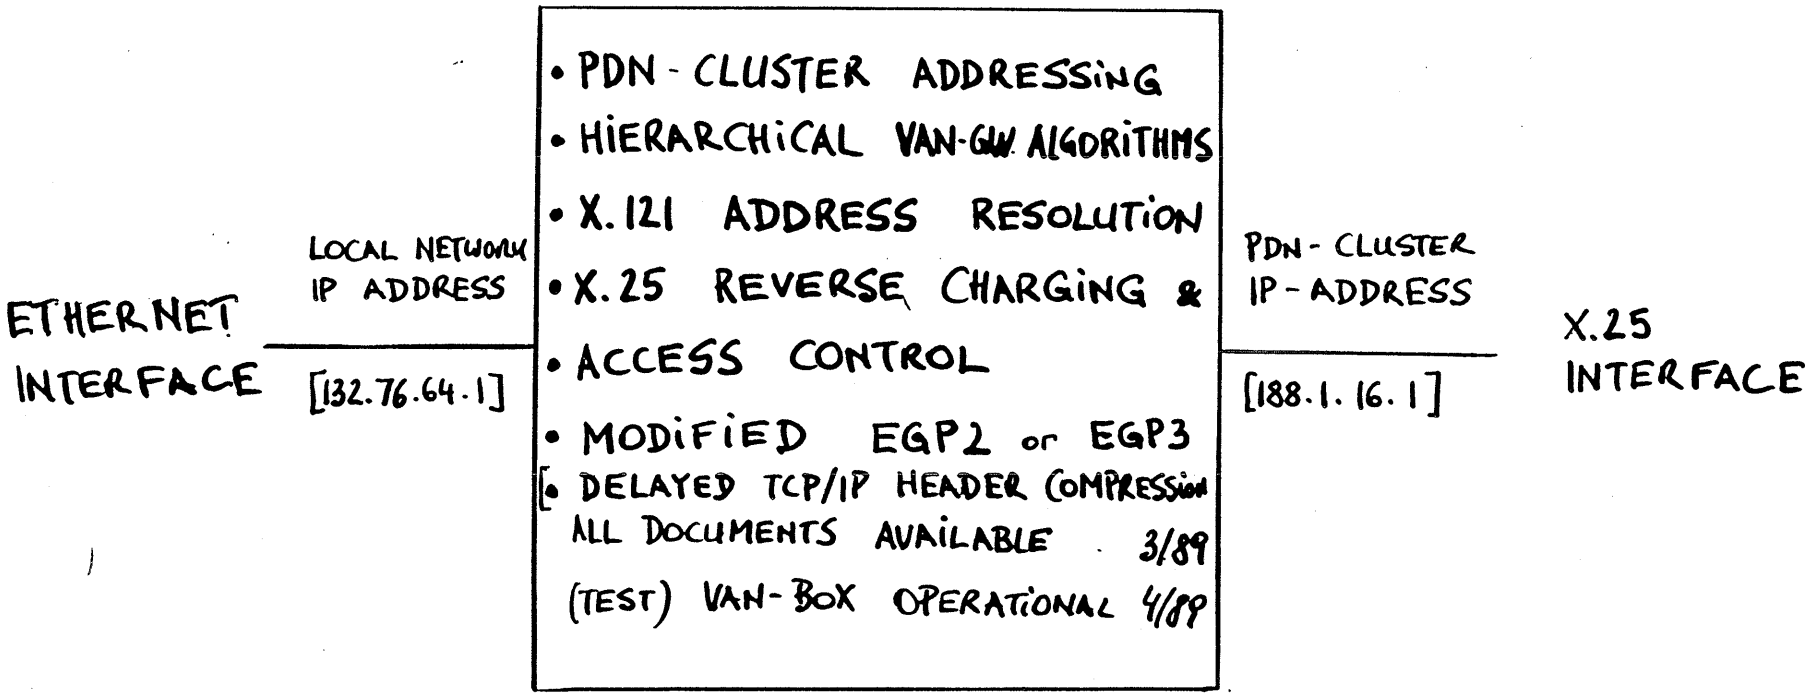

WHEN PLACED BETWEEN A LOCAL TCP/IP NETWORK AND AN X.25 PDN, THE VAN-BOX WILL AUTOMAGICALLY EXCHANGE WORLDWIDE NETWORK REACHABILITY INFORMATION WITH OTHER VAN-GWs (VAN-BOXES) ACCORDING TO THE HIERARCHICAL VAN - GATEWAY AND PDN ROUTING ALGORITHMS

Performance and Congestion Control
Chairperson: Allison Mankin/Mitre

CHARTER

Description of Working Group:

The charter of the IETF Performance and Congestion Control Working Group is to collect and develop short-term techniques for improving Internet performance, methods which like TCP Slow-start are retrofittable and inexpensive to implement. After a preliminary draft of a white paper documenting such performance enhancements for hosts and gateways, it was decided to sharpen the focus and divide the material into two papers.

One of the resulting papers is the RFC on gateway congestion control policies and algorithms. The intent of this paper is to present what is now known about the difficult problem of avoiding congestion in Internet gateways. It describes proposed policies such as Random Drop, Congestion Indication, and Fair Queuing, and sketches ground-rules for their adoption. An additional goal of the paper (achieved during the writing) is to generate dialogue on longer-term Internet gateway performance problems.

The other paper is an RFC on TCP performance. This describes TCP algorithms such as Retransmit Backoff, Slow-start, Nagle (Small-Packet Avoidance), and Delayed Ack, as well as their correct interaction. The scope is to expand the treatment of TCP performance found in the Host Requirements RFC.

Performance and Congestion Control Working Group
Chairperson: Allison Mankin/Mitre

STATUS UPDATE

1. Chairperson: Allison Mankin/mankin@gateway.mitre.org

2. Name of WG Mailing List(s):
ietf-perf(-request)@gateway.mitre.org

3. Date of Last Meeting: Cocoa Beach April 11-14, 1989

4. Date of Next Meeting: Stanford IETF in July

5. Pending or New Objectives:

The other objective of the group is the TCP Performance RFC. This is in rough draft state at this point, but the hope is to complete its good draft by the July meeting.

6. Progress to Date (e.g., documents produced):

- Gateway Congestion Control Policies
- We are closing in on completion of the RFC on gateway congestion control--the next revision (ready mid-May) will be placed in the IETF-DRAFTS directory for review by the IETF at large, and the contents will be presented in full at the July plenary.

Performance and Congestion Control
Chairperson: Allison Mankin/Mitre

CURRENT MEETING REPORT
Reported by Allison Mankin/Mitre

AGENDA

9-1 Tue: TCP Sub-Group
2-5 Tue: GW Draft (Open)
9-5 Wed: GW Draft (Members)

ATTENDEES

Art Berggreen	art@sage.acc.com
David Borman	dab@cray.com
Noel Chiappa	jnc@lcs.mit.edu
Mike Karels	karels@berkeley.edu
John Lekashman	lekash@orville.nas.nasa.gov
Charles Lynn	clynn@bbn.com
Allison Mankin	mankin@gateway.mitre.org
Matt Mathis	mathis@faraday.ecc.cmu.edu
Bill Nowicki	nowicki@sun.com
Bruce J. Schofield	schofield@edn-vax.dca.mil
John Scott	scott@dg-rtp.dg.com
Bill Westfield	billw@cisco.com

MINUTES

TCP Performance

The TCP Sub-group met in Cocoa Beach to get organized and agree on the scope and orientation of the TCP Performance paper. We agreed that the paper will expand on the Host Requirements RFC treatment of algorithms such as Retransmit Backoff, Slow-start, Nagle Small-Packet Avoidance, and Delayed Ack, and their correct interaction. It will be ancillary to the Host RFC; in particular, it will support the Host RFC's Musts, Shoulds and Mays.

A list of topics for scope had been generated by going through RFC-793 and the Host Draft. The group marked these according to our sense of the state of knowledge: Y if well understood, M if incompletely understood. We defined N for not understood, but only used it for Type of Service, because it seems orthogonal to TCP performance.

TCP Performance RFC Topics

	Certainty
Maximum Segment Size	M
TOS	N
Precedence	M
Connection Establishment	Y
Management of TCBs	Y

TCP Performance RFC Topics (chart con't from pg 1)

	Certainty
Connection Reuse	Y
PUSH	Y
Sender Silly Window Avoidance	Y
Nagle Small Packet Avoidance	Y
Slow-start Congestion Avoidance	Y
Relationship of Nagle S-S	Y
RTO Calculation	Y
Rxt Amount	Y
Rxt Backoff	Y
Response to Source Quench	Y
Window Upper Limits	M
Zero-window SWS Cant-Send State	M
Receiver Silly Window Avoidance	Y
Out-of-order Processing	Y
Delayed ACK	Y
Piggy-back ACKs	M
Application-TCP Interface	Y
Fairness Among Connections	Y
Interface of TCP to IP	M
Fair Processor	Y
Connection Instrumentation	Y
Extended Window Option	Y
High-Speed Implementation Techniques	M
Off-board Issue	M

Many of the attendees took action items to be responsible for topics. For quite a few of the topics, the members of the sub-group local to Washington, D.C. wrote drafts already, and these will be distributed.

Gateway Congestion Control

The rest of the working group time in Cocoa Beach was devoted to the Gateway Congestion Control Policies paper, now nearing release to the IETF. At the January meeting we had decided to recommend Random Drop in the paper. We revised the paper during the interim so that it detailed the RD policy more and gave it a fairly caveated recommendation (stating that experimentation was needed before any algorithm would be described).

The paper stimulated some comment. Lixia Zhang and Eman Hashmen (MIT) separately did simulation experiments on RD. Scott Shenker (Xerox PARC) did some analysis of RD as a sideline to a paper he is writing on game theory and gateway performance. All three have ongoing work on gateway congestion control algorithms that derive from Fair Queueing.

Chuck Davin gave an informal presentation of the MIT results on Tuesday afternoon. Scott's insights were offered in a long mail message, which was distributed to the working group.

A bottom-line summary of the studies:

RD as congestion control (choose random packet to drop instead of first or last on queue) is uncontroversially viewed as a win.

With RD, non-cooperating TCP gets excessive bandwidth at the expense of Slow-start connections.

RD (and other policies) give too much control feedback to connections whose RTTs (paths) are longer than others sharing the resource. Host-pair state as in FQ or SF (DEC Selective Feedback) is one cure.

Random Drop will remain at the center of the paper, but with explication of the performance problems it handles well and of its limitations. Its good properties for control (at queue overflow) will be made clearer. We will expand the sections on other policies and make a number of other changes suggested by the ten or so non-member reviewers who read the draft this time (and whom we thank wholeheartedly). Some of these changes are:

> Clarify distinction between congestion control and avoidance.
>
> Expand RD congestion control vs. RD congestion avoidance.
>
> Present the components of control system (congestion detection, feedback method, feedback selection).
>
> Survey congestion detection methods.
>
> In RD CA, constant and adaptive probability of drop.
>
> Survey the time constants identified for gateway performance, i.e. the interval used for averaging in congestion detection.

These changes are not as onerous as they sound. The paper will be redistributed in mid-May. At the same time, it will become an IETF-DRAFT. We will try for an interim meeting, perhaps, despite it not being great for everyone, on the day before INARC (May 31).

Some follow-on studies will be going on. One that came out of Cocoa Beach: at the open meeting on Tuesday, Rick Boivie and Yakov Rekhter offered to set an RT instrumented by Van into the NSSs to allow characterization of the time constants of backbone gateways.

PERF. + CONGESTION CONTROL WG

TASK 1: GW CONGESTION CONTROL DRAFT RFC

TASK 2: TCP PERFORMANCE "

1/ THIS MTG - REVIEWS BY GROUP AND INTERESTED OTHERS OF 2ND DRAFT OF GW PAPER

STATUS - MORE TEXT ON SOME POINTS, SOME CLARIFICATIONS, WILL BE COMPLETED BY MID-MAY

FULL PRESENTATION AT NEXT IETF

ietf-perf-request@gateway.mitre.org

PERF + CC WG

2/ TCP SUB-GROUP

MET (D.C. LOCALS) IN MARCH + WROTE A NEW ROUGH DRAFT BUT MISSED E-MAIL DEADLINE

THIS MTG - REVIEWED SCOPE AND TOPICS, AGREED ON ORIENTATION →

- SUPPLEMENT HOST RFC
- NOT SET REQUIREMENTS

Point-to-Point Protocol Working Group
Chairpersons: Drew Perkins/CMU and Russ Hobby/UC Davis

CHARTER

Description of Working Group:

The working group is defining the use of serial lines in data networks. While the main intent is to standardize the connection of IP networks over point-to-point links, the protocol is being designed to be extensible to other network protocols as well. The protocol will provide the capability of establishing the link parameters, authentication, link encryption, link testing, as well as control of the link while it is up. The protocol will also allow configuration and control of the higher level protocols such as IP, OSI, 802.3 bridging, and others.

Specific Objectives:

The main objective of the workgroup is to produce an RFC defining the protocol for the link and IP levels.

Estimated Timeframe for Completion:

The final draft of the RFC will be completed for the Fall 89 IETF Meeting.

Point-to-Point Protocol Working Group
Chairpersons: Drew Perkins/CMU and Russ Hobby/UC Davis

STATUS UPDATE

1. Chairpersons: Russ Hobby, University of California - Davis, rdhobby@ucdavis.edu
Drew Perkins, Carnegie Mellon University, dpp@andrew.cmu.edu

2. WG Mailing lists: ietf-ppp@ucdavis.edu - main mail list
ietf-ppp-request@ucdavis.edu - requests for addition to list

3. Last meeting: Cocoa Beach, April 1989

4. Next meeting: Stanford, July 23-24, 1989

5. Pending or New Objectives:

Produce RFC on protocol definition, final draft expected Fall 89, IETF meeting.

6. Progress to date (e.g., documents produced):

- Requirements for a Point-to-Point Protocol, Perkins September 1988.
- Complete protocol definition of link configuration and control. Definition of IP configuration and control being finished.

Point-to-point Protocol Working Group
Chairpersons: Russ Hobby/UC Davis and Drew Perkins/CMU

CURRENT MEETING REPORT
Reported by Russ Hobby

I. Introduction

The PPP WG met on April 11 and 12. The work group plans to have the current work written up and have one or two video conferences before the next IETF meeting.

Work was continued from where it was left at the video conference with the decision to handle configuration of the link and upper level protocols separately. The link must be up and ready before any of the upper level protocols can be configured and started. Once the line is up the upper level protocols may be brought up and taken down at any time. If the link goes down, the interface must inform the upper level protocols so that they may take appropriate action.

The group defined the protocols and procedures to bring the line up and work was started on the control protocol to bring up IP.

The link must be brought up by the following sequence:

1. Configuration exchange - this step is not completed until a Config Ack has been both received and sent. All configuration items are assumed to be at default values until configuration exchange is completed.

2. Authentication - authentication methods used are those agreed to in the configuration exchange if any. Authentication is accomplished using the PCP Authentication Protocol. A simple user/password authentication method is defined. Development of other methods is encouraged.

3. Encryption turned on - encryption methods used are those agreed to in configuration exchange if any. Only the data fields of PPP packets are encrypted and PCP packets are never encrypted insuring the control messages can get through even if encryption methods are out of sync. Currently no encryption methods are defined.

4. Line testing - check if line quality is sufficient to bring up upper level protocols. Suggested methods of line testing are being defined (Medin and Satz).

5. Line up - ready for upper level protocols.

   If any control packets are received that do not conform to the sequence, the equence is restarted. At any time if a data packet, as opposed to a control packet, is received

for an unknown protocol or a protocol that is not up, the packet is dropped.

II. PPP Control Protocol

As it was reported from the last IETF meeting, HDLC is the link packet format and in the data section of the packet is a field indicating the protocol in the remainder of the data. One the these protocols, number 33, is the PPP Control Protocol (PCP) used to configure and control the link and upper level protocols. PCP packets have a 16 bit protocol number field. After the protocol field the data can used as desired for configuration and control by each protocol.

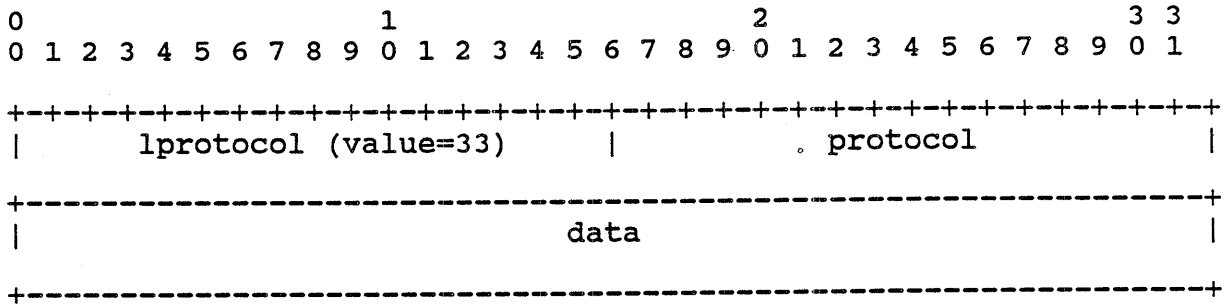

III. Link Control Protocol

Protocol 0 of the PCP, the Link Control Protocol (LCP) is used to configure and control the link itself. LCP includes functions for establishing the initial configuration, determining loopback, up/down control, circuit disconnect and other functions.

The PCP packet is as follows:

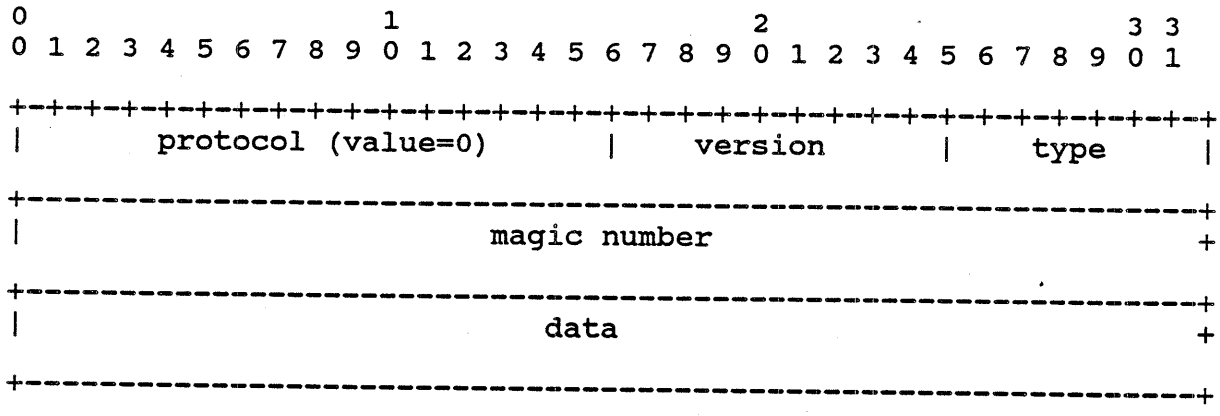

A. Version - The version of LCP supported.

B. Type - Type of LCP packet. Defined types are:

1 - Configuration	2 - Config Ack
3 - Config Huh?	4 - Config Nak
5 - Version Reject	6 - Type Reject
7 - Terminate Request	8 - Terminate Ack
9 - NOP	10- Keep Alive
11- Echo Request	12- Echo Reply
13- Protocol Unknown	

C. Magic Number - This pseudo-random number is used to uniquely identify an end of the point-to-point connection. This field is used to detect if a line is looped back to itself. Once a number is selected the same number is used for the duration for the connection. All LCP packets sent out must contain the senders magic number (See discussion on loopback detection)

D. Data - Additional data associated with the packet type.

LCP Packet Types

1. Configuration - This packet type is sent out the line to indicate pertinent configuration information and is used to establish a connection. Receipt of a Configuration packet means that the line is being reset and restarted. Exchange of Configuration packets can continue until both sides send Config Acks or one side gives up. If no response is received from the other side after a timeout period, the Configuration packet can be resent. Suggested timeout period is three seconds.

Configuration Items (CI) are placed in the data field of the PCP. Multiple CIs may be included in each packet. The format of a CI is:

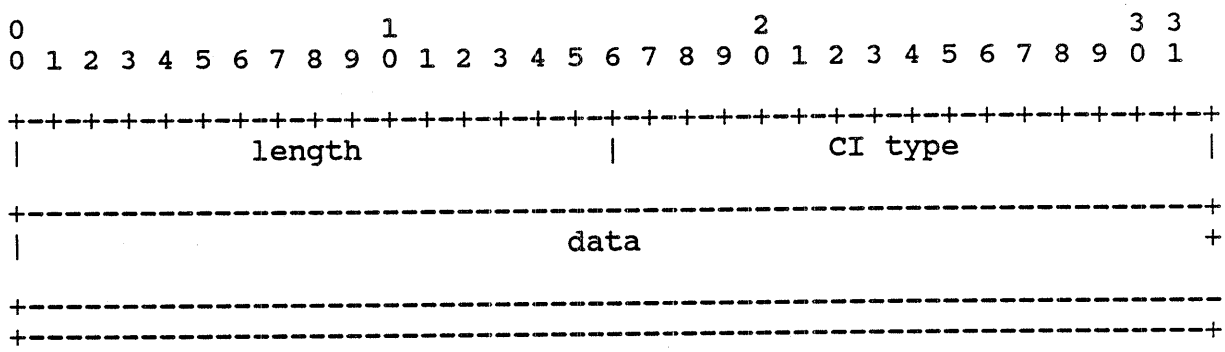

a. Length - inclusive length of the CI.

b. CI Type - Type number of the CI.

c. Data - value or other information for the CI.

CIs provide information on MRU, async character mapping, link authentication method and link encryption method. If a CI is not included in the config packet, the default is assumed. The end of the list of CIs is indicated by a zero length CI.

Config Item	Bits	Default
1 - Max Receive Unit	16	1024
2 - Async Control Char Map	32	all ones
3 - Authentication Type	16	none
4 - Encryption Type	16	none
5 - Keep Alive Parameters	??	none

Sync lines will accept any value for the Async Control Char Map.

2. Config Ack - This packet type is sent in response to a configuration packet and indicates acceptance of the other ends CIs. LCP data field contains CIs of accepted configuration.

3. Config Huh? - This packet type is sent in response to a configuration packet and indicates the configuration packet contained unknown CI type(s). The LCP data field will contain the CI entries of the unknown types.

4. Config Nak - This packet type is sent in response to a configuration packet and indicates the configuration packet contained unacceptable CI(s). The LCP data field will contain the CI entries of the unacceptable CI(s) and may contain suggested new values.

5. Version Reject - This packet type is sent in response to a LCP packet of an unacceptable version. The packet will return an acceptable version number.

6. Type Reject - This packet type is sent in response to a LCP packet of an unacceptable type. Any information in the LCP data field may be ignored.

7. Terminate Request - This packet type is sent to indicate the connection is going to be terminated. If possible wait for the Terminate Ack before breaking the connection. Any information in the LCP data field may be ignored.

8. Terminate Ack - This packet type is sent in response to a Terminate Request. Any information in the LCP data field may be ignored.

9. NOP - This packet type may be used to send non-LCP related data on the line, such as modem control information. When received the packet will be discarded.

10. Keep Alive Parameters - To be defined. Note: Both sides must agree on common parameters for keep alives.

11. Echo Request - This packet type is sent requesting that an echo reply packet be returned. Any information may be placed in the LCP data field.

12. Echo Response - This packet type is sent in response to the echo request packet. The LCP data field must be a copy the LCP data field of the request.

13. Protocol Unknown - This packet is sent in response to a PCP packet of an unknown or unimplimented protocol. The data field contains the 16 bit value of the unknown protocol.

IV. Authentication Protocol

Protocol 1 of PCP, the Authentication Protocol, is used to verify the entity on the other end of the link. The authentication method agreed to in the configuration exchange is the method an entity will use in verifying the other end. Each end may use a different method if agreed in the configuration phase.

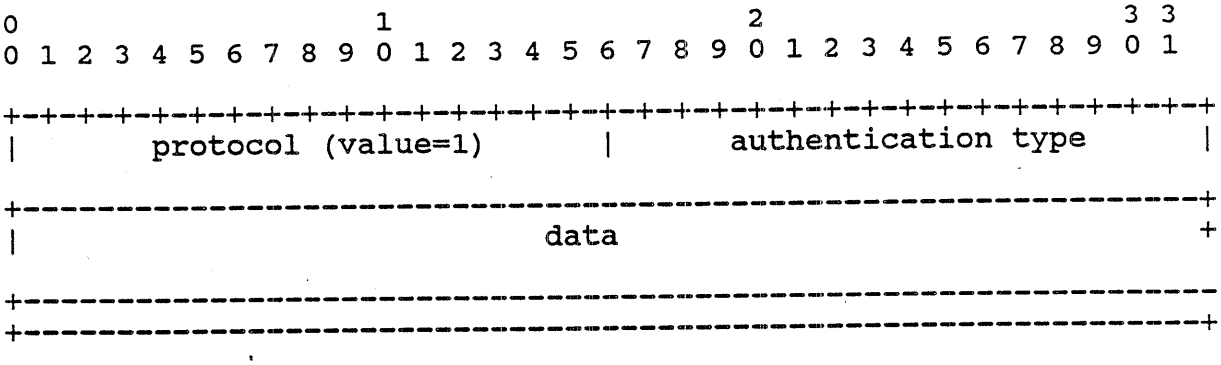

The data field may be used in any way defined by the authentication method.
A simple user/password (auth type = 1) method is defined here.

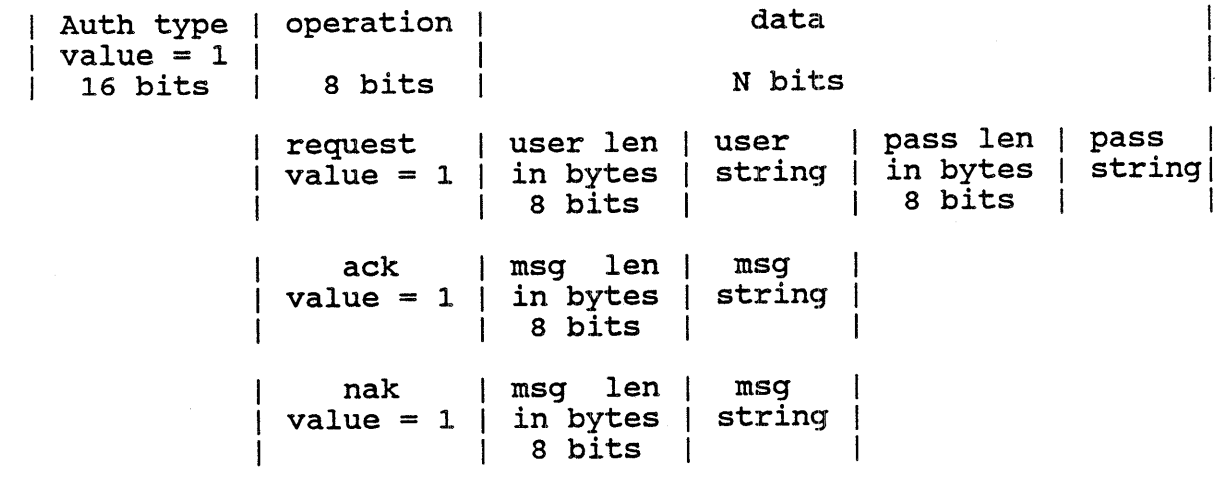

V. IP Control Protocol

Protocol 35 of the PCP, the IP Control Protocol (ICP) is used to configure and control the IP protocol. ICP includes functions for establishing the initial configuration, taking down the protocol and other functions. The ICP packet format is as follows:

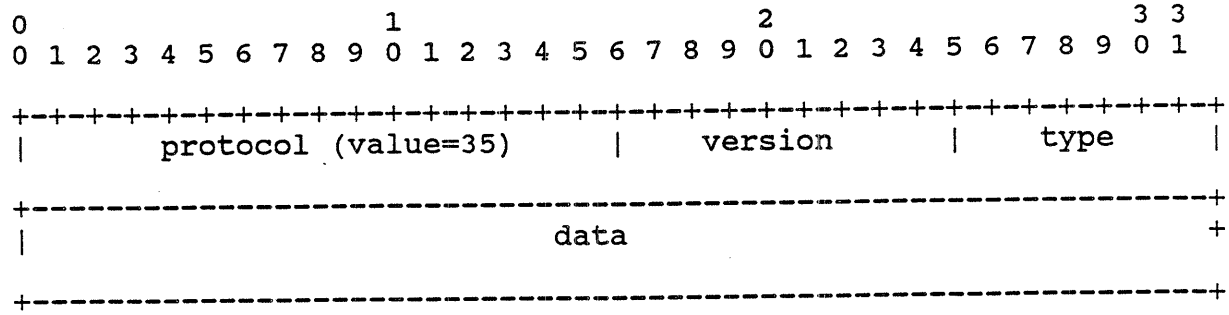

A. Version - The version of ICP supported.

B. Type - Type of ICP packet. Defined types are:

1 - Configuration	2 - Config Ack
3 - Config Huh?	4 - Config Nak
5 - Version Reject	6 - Type Reject
7 - Terminate Request	8 - Terminate Ack
9 - NOP	

The ICP functions above work the same as for their counterparts in the LCP section. The CIs for ICP are:

Configuration Item	Bits	Default
1. Addresses	??	none
2. Compression Method	??	none

IP Address Negotiation
On Init
  Send REQ w/ my IP addr or 0 if unknown.
  Receive ACK w/ both addresses (1) or
  Receive NACK. Can retry REQ w/ a different address.
On Recv REQ
Remote addr set if like remote addr then ACK w/both addresses else NACK
Remote addr not set pick an appropriate remote addr and ACK w/ both addresses.(2)

Note 1: If remote addr is 0 then ignore it. He'll be soon asking you to set it so remember it then.

Note 2: If have no idea what to pick (such as both ends ask each other end for its address) then give defaults, 127.0.0.X for one with smaller magic number and 127.0.0.X+1 for one with larger. There is discussion if the net 127 address is reasonable and what the value of X should be.

PPP Sequence

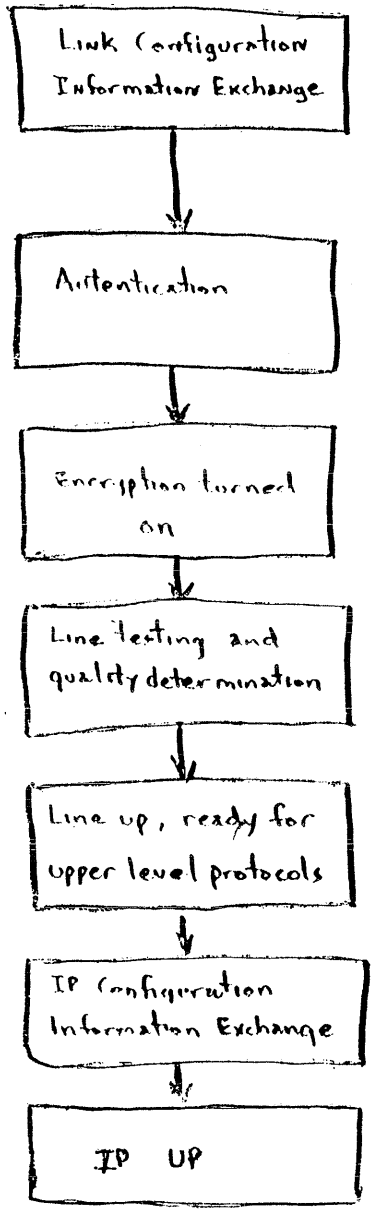

MTU.
Async Control Character Map.
Authentication Method.
Encryption Method.
Keep-alive Parameters.

Simple User/Password defined.
Development of other methods encouraged.

No Methods currently defined.
Encryption only on data fields of protocols.
Not used on Link Control Protocol.

Echo Request/Reply Packets.
Keep-alive Packets.
Testing method suggested.

Upper level protocols can be brought up and taken down at any time.

IP Address
Compression method used

HDLC

8	8	8	8 or 16	N	16	8
FLAG	ADDRESS	CONTROL	LPROTOCOL	DATA	FCS	FLAG

LPROTOCOL 33 – [redacted] CONTROL PROTOCOL

16
PROTOCOL

PROTOCOL ∅ – LINK CONTROL

8	8	16	N
VERSION	TYPE	MAGIC #	DATA

Types

1. config
2. Config Ack
3. Config Nak
4. Config Nak
5. Echo Request
6. Echo Reply
7. Version Rej
8. Type Rej
9. Terminate Req
10. Terminate Ack
11. Keep Alive
12. Nop
13. Protocol Unknown

Config Packet

16	16	N
LENGTH	TYPE	DATA

(MAY BE REPEATED)

Types	Default
1. MAXIMUM REC UNIT	1024
2. Async Control Char Map	ALL ONES
3. Authentication Method	NONE
4. Encrypting Method	NONE
5. KEEP-ALIVE PARAMETERS	NONE

PROTOCOL 33, PROTOCOL 1

AUTHENTICATION

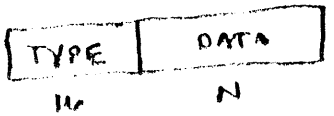

Type 1 - SIMPLE AUTHENTICATION

1 - REQUEST

2 - ACK

3 - NAK

REQUEST

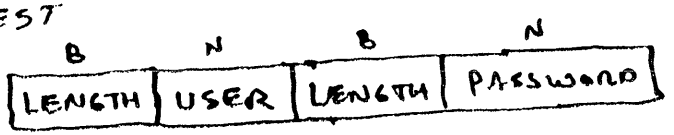

ACK, NAK

PROTOCOL 33, PROTOCOL 35

IP CONTROL

SIMILAR TO PROTOCOL 0 - (LINK CONTROL)

CONFIGURATION ITEMS

1. ADDRESSES
2. COMPRESSION METHOD

ST and Connection IP Working Group
Chairperson: Claudio Topolcic/BBN

CHARTER

Description of Working Group:

Define the next version of the ST protocol, explore future connection oriented internet protocol, use the former as a testbed to perform experiments in support of the latter.

Specific Objectives:

- Produce a new specification of ST
- Produce a specification of a next generation connection oriented protocol

Estimated Timeframe for Completion:

a) Produce a new specification of ST. (2-3 months)
b) Produce a specification of a connection oriented protocol. (6-12 months)

ST and Connection IP Working Group
Chairperson: Claudio Topolcic/BBN

STATUS UPDATE

1. Chairperson: Claudio Topolcic, BBN Labs, topolcic@bbn.com

2. Name of WG Mailing List(s): cip@bbn.com

3. Date of Last Meeting: April 12, Cocoa Beach Florida

4. Date of Next Meeting: July 25, Stanford

5. Pending or New Objectives: none

6. Progress to Date (e.g., documents produced)

   - Internal draft of ST Specification
   - Numerous e-mail messages describing issues in connection oriented protocols

St and Connection IP Working Group
Chairperson: Claudio Topolcic/BBN

CURRENT MEETING REPORT
Reported by Steve Casner/ISI and Claudio Topolcic/BBN

AGENDA

Characterize applications
Define functions of the internet layer
Identify implications on underlying networks
ST- discuss a number of issues we had not yet agreed to

ATTENDEES

Ross Callon	callon@erlang.dec.com
Steve Casner	casner@isi.edu
Danny Cohen	cohen@isi.edu
Phil Draughon	jpd@accuvax.nwu.edu
Phil Park	ppark@bbn.com
Zaw-Sing Su	zsu@sri.com
Claudio Topolcic	topolcic@bbn.com
Paul Tsuchiya	tsuchiya@gateway.mitre.org

MINUTES

The working group held two meetings, which correspond to the two tracks we are pursuing. The meeting held during the day of Wednesday 12 April covered the high level and long term issues of connection oriented internet protocols. A second meeting was held in the evening of 12 April. It covered a number of short term issues that need to be discussed to finalize the ST specification.

COIP Meeting of April 12, 1989:

Phil Park gave a presentation on Application Characterization, following the message he sent out. Applications are characterized by a list of parameters: packet size, packet rate, etc. We are calling these the "Quality of Service" (QOS) parameters.

Danny proposed an additional parameter indicating how long a connection is expected to last; e.g., for FTP, saying how long the file is, or for a teleconference, how long it will run. It's not clear just how this will be used. Perhaps we need only two values: short (transient, don't count it), and long. If only two values are necessary, then it was argued that the "transient" transfer is equivalent to datagram service and could be supported by a datagram internet layer rather than a connection oriented internet layer.

Although the problem of resource management in the network was not discussed, a possible use of this information is the following. We can imagine that the management might occur on several levels. Part of that management might be scheduling of resources on a time scale of tens of minutes to hours. If we want to establish a connection for a teleconference, and can specify that the teleconference is to last two hours, we might want the connection to be refused if my resources are going to be preempted by a higher-priority teleconference that has been scheduled to begin in five minutes. Or, we might want a different route to be chosen if such a conflict is encountered and an alternate route is possible. (If automatic rerouting with minimal disruption is possible, then maybe we don't have to care about this route choice initially.) This scenario presupposes a scheduling protocol which is likely to be outside the scope of the connection oriented internet layer as we envision it now. However, if scheduled connections are to be considered, then there must be some interaction between the scheduling mechanism and the connection establishment mechanism. Also a network's decision to accept a connection may be based on administrative policies. Those policies may allow use of a network as long as you don't hog it. The duration of a teleconference or the size of a file transfer would be one measure to be considered in such a policy.

There are some parameters that applications give to nets, and some vice versa. Steve asserts that it's best to have full information transfer both ways so that nets can make decisions based on the information when they care. There's not much penalty in providing too much information.

Some of the parameters that stimulated discussion are the following:

Bandwidth -- don't want a linear scale.
Packet size -- Danny proposed that the network should report back the maximum packet size available along the complete path, rather than having the application say what packet size it wants. BUT -- the Wideband Network, for example, needs to know the size packets to be sent in order to efficiently allocate a stream to match. It was agreed that once a connection was established, the maximum packet size along that route could be reported to the application to provide it some flexibility.
Reliability -- need only a few choices: perfect, a lot, a little, don't care. This is true unless you are really controlling something like coding, then you need enough bits to specify the range of control settings. Ross suggested we select levels that whose results we can understand, specifically:

1) The reliability supplied by TCP to an application.
2) The reliability over which TCP works efficiently.
3) The reliability across which voice works well.
4) The reliability over which TCP will collapse.
5) Don't care.

Delay -- this is a "softer" requirement than bandwidth. We couldn't figure out how an agent could guarantee a given delay, nor what it should do if it could not meet it.
Security -- may affect routing decisions, or processing load in agents if they have to decrypt headers, but for real security, applications don't get a choice, so there may be no point in including this parameter. Security should be supported by SDNS.
Burst factor -- this should be covered in the rate parameters if they are rich enough to describe the distribution. Later, we were not so sure, and considered a "leaky bucket" parameter, i.e., how big does the bucket (buffer) have to be to avoid overflowing given an output at the bottom with a flow at the average rate.

Danny suggested that we have shorthand codes for various traffic types. This was based on the observation that in NVP-II we had a full set of parameters but only ever used the "vocoder codes" that were a shorthand for the full set. Steve suggests the reason is that we never did vocoder experimentation over the net; we experimented in the lab, then used completed vocoders over the net. For application characterization, the communication of parameters is not between peers, but between layers. Steve suggests there will be a wider variety of applications, and that at least some parameters will require numerical adjustability if a shorthand is defined.

Danny suggested that each packet could carry a priority specification. Then the high priority packets would be processed first. Ross suggested an alternate approach, in which a connection would have a priority, and this priority would be used as input to the decision to accept, reject or pre-empt that connection. It would be the entire connection, not any given packet that would be deleted. This approach was more popular among the attendees. This discussion brought out two interesting questions. In times of congestion, should a connection be singled out and pre-empted, or should packets from a number of connections be dropped without pre-empting any given connection, and how would the connections be selected? Second, how it could be decided that the requirements of an established connection could no longer be met and that connection should be pre-empted? Neither of these issues were resolved.

Nevertheless, allowing packets in a connection to have different priorities supports a single application that carries data of different priorities. This would be the case if we are using a layered coding scheme where lower priority layers get discarded

if there is congestion. If this were done with separate connections each with its own fixed priority, that could cause problems with synchronization. We did not resolve this issue.

We discussed how the values of these parameters would be arrived at. There can easily be flexibility in what an application can accept. Possible reasons include a) in many applications packet size and packet rate can be traded off and so long as they represent the same bandwidth, since many applications only care about bandwidth, b) some applications can transmit comparable data at different bandwidths, such as when using multiple rate voice coders, or c) some applications, such as file transfer, can adapt their behavior to offer different bandwidths. Therefore, it is reasonable to assume that there can be a negotiation of these parameters between the application and the network layer.

In negotiating packet rate and packet size, there are two distributions to consider: 1) the range over which the application can operate (its adjustability), and 2) the range over which it varies as it operates (for variable-rate coding schemes). For (1), the application would want to give two values, what it needs (min) and what it wants (max). For (2), a variable rate application would want to specify the average rate (what it wants to pay for) and the peak rate it wants the network to handle. Specifying ranges of average and peak values won't work because the network can't tell what value of peak goes with a given choice of average (or vice versa). Instead it might be better to give a list of (avg,peak) pairs; each pair would specify a range of type (2), while the collection of pairs would specify the range of type (1). However, there is a problem, the packet rate and packet size are related; either might be fixed while the other changes (either for adjustability or variability). So, it might be necessary to specify a list of quadruples: (avg rate, peak rate, avg size, peak size) where avg=peak for at least one of rate or size.

The parameters list that we arrived at is the following:

Bandwidth - (avg pkt rate, peak pkt rate, avg pkt size, peak pkt size)
Delay - (maximum tolerable at some percentile, some indication of acceptable distribution, flag to allow discard)
Reliability - (some small number of values)
Discard option - (newest or oldest)
Total duration of transfer - (?)
Burst factor - (?)

We finally decided that we need to look at what the network will do with this information before we can say much more about the negotiation process. We decided that the next step should be to talk about the network layer and we should then come back and take another pass at characterizing the applications and defining the negotiation procedure.

ST protocol specification meeting of 12 April 1989

ATTENDEES

Steve Casner, Danny Cohen, Phil Draughon, Phil Park, and Claudio Topolcic

We met for a few hours the same evening. We reached a number of decisions and agreed not to decide on a number of other issues.

DECISIONS

We reviewed Steve Casner's proposal for doing encapsulation of ST packets in IP for the purpose of sending the ST packets across IP-only parts of the Internet. We approved it. In this technique, an IP header is placed in front of an ST packet and the IP destination is set to be the IP address of the next ST agent along the ST connection's route. This is used when the best route includes passing through gateways that do not support ST. Since the IP-only gateways and networks cannot perform resource management, we assume that this will only be the case when that IP-only part of the Internet is only lightly loaded.

Danny suggested using this IP encapsulation technique for all ST packets. Again, the IP destination would be the IP address of the next ST agent, even if that agent is the next gateway in the path. ST agents would still perform all the resource management functions they would if the packets were not so encapsulated. The motivation would be to take advantages of security implemented for IP as part of SDNS, or other services provided with IP. Everyone was intrigued, but we weren't convinced that this would be of enough benefit. We did not accept this view at this time. We agreed that it is not mandatory for the new ST specification to maintain the same interface with the next higher protocol layer as current ST. However, changes should be based on sound reasoning. Specifically, we may cause the next higher layer to need to have more information than current ST does. Phil Draughon offered to look into the requirements of the next protocol layer.

We agreed that we will need to write more specification documents than just the ST spec we are currently working on. We need to describe things like the interface between the ST layer and the next higher layer in a host, and the routing algorithm that will be used.

We agreed that adding new control messages to ST is acceptable, but only as long as the new ones have a different and distinct function.

We decided to get rid of the ST.DG bit.

We agreed that a source route option would be a good idea, and thought it would be relatively easy to implement, though we did not decide on a mechanism.

We agreed with Danny's proposal that security be provided by SDNS. IP encapsulation will be necessary to allow this.

NO DECISION REACHED

We talked about the impact of what we are doing with the conferencing part of ST on the point to point part of ST. We could try to not force any changes on the point to point ST, we could make only the obviously necessary changes, such as using a consistent Flow Spec, or at the other extreme, we could eliminate point to point ST altogether. Making point to point ST be simplex is equivalent to eliminating it. We did not make a decision because nobody had a particularly strong opinion.

We talked about the possibility of changing the route of an established connection, but we decided this was hard and we would postpone a decision.

We agreed that aggregation will be useful, and also agreed that it will be hard to specify and implement. We did not discuss how to do it.

We discussed whether a field would be needed that specifies the next protocol above ST. Such a field exists, it is the extension. However, we did not agree how that field should be interpreted. An obvious possibility is to partition the space and assign different parts to different protocols.

Somebody suggested adding a flag, and function, that causes the reverse path of a (conference) connection to be built automatically, in essence allowing conferencing connections to be full omniplex. We decided to table this idea for now.

We did not discuss routing.

ST and connection oriented internet protocol working group - Claudio Topolcic

- Two parallel tracks
  - Longer term connection oriented protocol issues
  - Short term ST specification

1. Long Term
   - Characterizing Applications with "Quality of Service" (QoS) parameters (first pass, not final)
     - Packet rate (avg & peak)
     - Packet size (avg & peak)
     - Delay (max tolerable ~~at~~ some percentile, some measure of acceptable distribution, flag to allow discard)
     - Reliability (small number of values)
     - Total duration of transfer (?)
     - Discard option (oldest or newest)
     - Burst factor (?)
   - Negotiating QOS parameters
     - Application should request operating points rather than ranges
     - This is hard to talk about until we think about what the nets will do with this information
   - Next step is to think about characterizing the networks

2. ST specification
   - Agreed
     - How to do IP encapsulation
     - Should have a way to do source routing
     - OK to change the interface to NVP
     - OK to expect more functions from NVP
     - We need to write more than one document
   - Further discussion
     - How to do aggregation
     - To change or eliminate point-to-point
     - When/if to change an existing route
     - Routing

TELNET Working Group
Chairperson: Dave Borman/Cray

CHARTER

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:

1. Either re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol, or issue a companion RFC to update and expand on fuzzy areas of RFC 854.

2. Create or update RFCs for TELNET options to clarify or fill in any missing voids in the current option set. (Most noteably, some method to allow automatic user authentication is needed).

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

4. When the above objectives have been met, go dormant, and will be re-activated as needed to fullfill the objective of being a clearing house for future extensions to the TELNET protocol.

Estimated Timeframe for Completion:

Estimates will be determined after the first meeting.

TELNET Working Group
Chairperson: David Borman/Cray

STATUS UPDATE

1. Chairperson: Dave Borman, dab@cray.com

2. WG Mailing List(s): telnet-ietf@cray.com
(Subject to change...)

3. Date of Last Meeting: New Group

4. Date of Next Meeting: July 1989 at Stanford

5. Pending or New Objectives: see Charter

6. Progress to Date (e.g., documents produced): none

CURRENT MEETING REPORT

none

USER-DOC Working Group
Chairpersons: Tracy LaQuey/Univ of Texas and Karen Roubicek/NSF

CHARTER

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:

- 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.)

- Draft bibliography will be prepared by mid-December 89.

USER-DOC Working Group
Chairpersons: Tracey LaQuey/Univ of Texas and Karen Roubicek/NSF

STATUS UPDATE

1. Chairpersons: Tracy LaQuey / tracy@emx.utexas.edu
   Karen Roubicek / roubicek@nnsc.nsf.net

2. WG Mailing Lists: us-wg@nnsc.nsf.net (temporary?)
   us-wg-request@nnsc.nsf.net

3. Date of Last Meeting: JVNC Supercomputer Center, Princeton NJ / 1 Jun 89

4. Date of Next Meeting: July IETF meeting/10:45am - 4:00pm, 25 July 89
   Stanford, CA.

5. Pending or New Objectives: see Current Meeting Report

6. Progress to Date (e.g., documents produced):

   1st formal meeting 1 Jun 89 / draft charter and objectives drawn up

USER-DOC Working Group
Chairpersons: Tracey LaQuey/Univ of Texas and Karen Roubicek/NSF

CURRENT MEETING REPORT
Reported by Karen Bowers

Several members of the USWG took the opportunity to convene a WG session during the FARNET meeting at the JVNC Supercomputer Center, Princeton, NJ on 1 June 1989. The purpose of this session was to discuss the formal formation of a distinct working group to assemble a bibliography of documents and user training tools useful to NICs, LAN managers and end users. The Agenda, outlined below, was very ambitious for the time allotted and consequently will extend into a follow-on WG session during the upcoming IETF meeting at Stanford University, 25-28 July 89.

AGENDA

- Form a distinct Working Group
- Write Charter and Objectives
- Select the Various Categories of Documents/Info to be Included
- Determine "Plan of Attack"
- Identify Existing Sources of Information
- Discuss in Detail Biblio Format to be Adopted

ATTENDEES

Karen L. Bowers	Karen Roubicek
Tracy LaQuey	Don Morris
Martyne Hallgren	Ed Krol
Joel Maloff	Tom Bajzek

MINUTES

Accomplishments:

- Karen Roubicek and Tracy LaQuey were asked to co-chair this effort as a WG (tiger team) under USWG and graciously accepted.

- A draft charter was drawn up and will be further revised and presented to the IETF Chairman for comment and approval prior to the Stanford IETF in July.

- Some basic requirements were identified as essential to this effort:

  Contacting individuals directly for their participation and biblio inputs is probably the most effective way of obtaining information, though mailing lists and bulletins will still be employed.

  Each listing should contain some or all of the following information:

  * date of document
  * shelf life
  * where to obtain and format
  * abstract
  * limitations/caveats
  * version #

Non-document sources of information could be included, such as videos, available user training workshops, etc:

* MCI Video on Internetting
* SIGUCCS/EDUCOM/MERIT Workshops
* England Study (Source: Jim Sweeton)

- Some documents have already been collected by Tracy. They are on emx.texas.edu. (The list of documents can be accessed by anonymous ftp. Cd to "user.wg/biblio" and the file is called "bibliography". The actual documents thus far collected are in "user.wg/documents".) These documents have been placed in the following "tentative" categories:

  * Introductions to TCP/IP and the Internet
  * Technical TCP/IP Tutorials
  * Network Administrators Tutorials
  * Electronic Mail Tutorials
  * Electronic Mail Configuration Tutorials and Reference Materials for Network Managers
  * Directory Services Documents
  * Reference Materials

Interim Activity Planned:

Within the next several weeks, Karen L. Bowers, Tracy LaQuey and Karen Roubicek will confer via a teleconference, finalize the Charter/Objectives, and outline the specific approach to be taken in assembling the bibliography. They will also discuss the official WG name to be adopted and determine if a mailing list separate from the us-wg mailing list is necessary or counter to active WG participation.

Next Meeting:

This "biblio" WG will convene at Stanford. It is essential the USWG, NISI and "biblio" WG are scheduled at times independent of one another to ensure essential participation in all three forums.

User Services Working Group
Chairperson: Karen Bowers/NRI

CHARTER

Description of Working Group:

The User Services Working Group will identify and address critical service requirements needed by "those people who help end users" (e.g. local net managers) and develop tools and materials to aid in the productivity of end users. The purpose is to answer the needs of the lower levels (*) within this hierarchy:

NATIONAL NETWORK
NET MANAGERS (NSF, DCA, ETC.)
NICs/NOCs
REGIONAL NET MANAGERS
LOCAL NET MANAGERS*
END USERS*

Specific Objectives:

1. Assemble a non-static cadre of interested experts within an open forum to exchange user services information, to share problem-solving techniques, and to select critical projects to be undertaken on behalf of the local net manager and end user.

2. Select projects based on production-oriented criteria. The Project(s)

   - must lend itself to accomplishment within a reasonable timeframe
   - must culminate in a measurable/quantifiable end result
   - must address user assistance needs = be user oriented
   - must yield products/tools designed to be both easily maintained and updated (with built in accountability)
   - must not duplicate efforts (This will be pre-empted by surveying exisitng resources.)

3. Determine the most appropriate approach to a respective project (s):

   - produce a totally new product
   - enhance/improve/influence an existing resource
   - table action for future consideration

4. Spin off various small WGs (tiger teams) to address very specific, short term projects (EX: NOC-Tools WG

and NISI WG). Once the respective project(s) is completed, members of the tiger team(s) will reassemble within the USWG to participate in the identification of the next project(s) to be undertaken.

Estimated Timeframe for Completion:

Selection and completion of projects will occur on a continuous basis, with timelines established for each individual tiger team formed.

User Services Working Group
Chairperson: Karen Bowers/NRI

STATUS UPDATE

1. Chairperson: Karen L. Bowers
bowers@sccgate.scc.com

2. WG Mailing List(s): US-WG@NNSC.NSF.NET and
US-WG-REQUEST@NNSC.NSF.NET

3. Date of Last Meeting: Cocoa Beach, Fl/April 11-12, 1989

4. Date of Next Meeting:

- o NISI WG formation meeting 4 May 89 and Plenary/WG meeting Stanford University/ 25-28 July 89
- o "Biblio" WG formation meeting 1 June 89 and Plenary/WG meeting Stanford University/ 25-28 July 89

5. Pending or New Objectives:

USWG Chairman to investigate Campus Awareness via meeting at NSF with D. Vanbelleghem, SIGUCCS et al and determine if USWG should get involved.

6. Progress to Date (e.g., documents produced):

Formation of three WGs (tiger teams): NOC-Tools, NISI WG, and a bibliography WG (not yet named); refer to the CURRENT MEETING REPORT attached

User Services Working Group
Chairperson: Karen Bowers/NRI

CURRENT MEETING REPORT
Reported by Karen Bowers

AGENDA

- Brief Intro of USWG (for new attendees)
- Planned USWG Organizational Structure
- Briefing on NOC-Tools WG (Bob Stine)
- Individual Briefings on AREAS FOR CONSIDERATION

*Network Resources Handbook (T. LaQuey and K. Roubicek)
*How to Set Up Campus NIC/NOC (T. LaQuey)
*Bibliography of Documents Every NIC Should Have
(M.Schoffstall and F. Perillo by e-mail)
*Mailing List Management: Listserv (J. Sweeton)

- Other: USWG Review of DRAFT Outline for U of Texas' Directory of Computer Networks (T. LaQuey and USWG)
- Selection of Projects to be Undertaken

ATTENDEES

Joyce K. Reynolds JKREYNOLDS@ISI.EDU	USC/ISI	213 822-1511
Jim Sweeton SWEETON@Merit.Edu	MERIT	313 936-3000
Karen Roubicek ROUBICEK@nnsc.nsf.net	NNSC/BBN	617 873-3361
Mary Stahl STAHL@SRI-NIC.ARPA	SRI/NIC	415 859-4775
Jose J. Garcia-Luna garcia@sri-com	SRI/NIC	415 859-5647
Robert Stine stine@sparta.com	SPARTA	703 448-0210
Ole Jacobsen ole@csli.stanford.edu	ConneXions/ ACE	415 941-3399
Tracy LaQuey tracy@emx.utexas.edu	UT Austin	512 471-3241
Don Morris morris@ncar.ucar.edu	NCAR	303 497-1282
Scott Brim swb@devvax.tn.cornell.edu	Cornell	607 255-8686
Philip Almquist almquist@jessica.stanford.edu	Stanford	415 723-2229
Elise Gerich epg@merit.edu	MERIT/NSFNET	313 936-3000
Paul Love LOVEEP@SDS.SDSC.EDU	SDSC	619 534-5043
Karen L. Bowers bowers@sccgate.scc.com	NRI	703 620-8990
Rebecca Nitzan nitzan@nmfecc.llnl.gov	LLNL	415 422-9775

MINUTES

After a brief review of our goals and charter the first item of discussion was the organizational structure of the USWG. Just as the members of the MIB WG have concluded, we too have determined that the USWG will be an umbrella organization from which smaller Working Groups (tigerteams) will be created to address short term projects. We have taken this one step further and have revised our Charter (provided above) to align with this structure and to reflect the project management aspects of our WG activities.

Bob Stine presented a briefing on the newly formed NOC-Tools WG, one such tiger team. The thrust of NOC-Tools is "to develop a catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools (both hardware and software) for TCP/IP internets. This WG, co-chaired by Bob Stine (SPARTA) and Bob Enger (Contel), convened a separate WG session on Wednesday, 12 April 89; minutes from that meeting have been prepared and are available.

Some of the AREAS FOR CONSIDERATION identified during our JAN/Austin,TX meeting (see Agenda) were researched during the spring quarter. Four of these were briefed by individual WG members and then were discussed at length by the USWG.

As the discussions progressed it was evident that several distinct requirements were surfacing. The most dramatic realization was the need to define what "network information services" are and how those services overlap/differ at each level of the networking hierarchy (national/ regional/campus). During a somewhat impromptu working lunch several USWG members concluded that what is needed is a requirements document addressing "network information services", to include requirements at each level; basic/advanced/elite/(perhaps, private) services that could be made available; and how these services could be successfully interconnected as an Internet-wide Network.

Information Service. The decision was made to hold a follow on exploratory meeting at NRI and invite a small group of NIC representatives. The purpose of the meeting is to define the general requirements for Internet information services and assemble that information into a network information services requirements document. Tentative plans are for this requirements document to be produced by a small USWG tiger team, with review to be accomplished by the entire USWG and individual reviewers selected by the IETF Chairman. (This requirements document could conceivably be the precursor to the design of a Network Information Services Infrastructure to provide internet-wide network information services.)

The second area identified for action is the creation of a bibliography useful to NICs, LAN managers, and end users. This will be an expansion of the work already done by

Marty Schoffstall and Francine Perillo and will include documents which answer such questions as "what is the Internet" and "what services are available (mailinglists/enhanced services)" and will provide references to basic information such as "how to ftp", "how to use email" and "how to set up a campus NIC/NOC". This will include a reference to the NOC-Tools bibliography andother pertinent glossaries. Bibliotext was suggested as a useful format for setting up this bibliography. The bibliography and any related on-line documents will be installed in a repository. Tracy LaQuey volunteered to act as a temporary repository until the bibliography has been produced and procedures established to maintain/update the bibliography/on-line documents. A Video Teleconference will be held on/about 2 June to further address this project and better define the "boundaries" of what should be included in this document.

Finally, the issue of "campus awareness" was raised and discussed. How can information be provided campus-wide to ensure colleges/universities are aware of 1) what connectivity currently exists on their respective campuses , and/or 2) how they can connect to the Internet and what connectivity provides them in terms of enhanced research capability and information exchange. This is an area currently being addressed by MERIT EDUCOM and SIGUCCS. The USWG needs to further investigate if we should/can play a role in this user education process; should we get involved in campus "road shows"? This will be discussed further at the next IETF plenary.

AREAS FOR CONSIDERATION tabled at this time (to be addressed in the future):

* Network Resources Directory
* How to Set Up a Campus NIC/NOC
* Mailing List Management
* Consolidated Common End User Questions/Answers

Other action items:

Ensure Ed Krol (U of IL) and Charlie Catlett (NCSA) are invited to participate in the bibliography. (K. Bowers)

Personally invite all key players (identified during our first meeting) to attend the Stanford meeting and/or be placed on the USWG mailing list. (J.K. Reynolds and K.Bowers)

AGENDA

- Brief Intro of USWG (for new attendees)
- Planned USWG Organizational Structure
- Briefing on NOC Tools WG
  (a newly established WG Under USWG) (B. Stein/SPARTA)
- Individual Briefings on Areas for Consideration
  - Network Resources Directory
    (Tracy LaQuey and Karen Roubicek)
  - How to Set Up Campus NIC/NOC
    (T. LaQuey)
  - Biblio of Docs Every NIC Should Have
    (M. Schoffstall/F. Perillo by email)
  - Mailing List Mgt/Listserv (Jim Sweeton)
  - Consolidated Common End User Questions/Answers
    (Sergio Heker - tabled)
- Review of Draft Outline: U of Texas' Directory of Computer Networks
- Selection of Projects to be Undertaken

April 11, 1989

Three Actions

April 11, 1989

I. Bibliography Working Group

- Bibliotext Format
- General Content
  - Title
  - Author
  - Citation: publishing house/date
  - Page Reference
  - Comments (abstract?)
  - (If on-line) Where found
- Addresses Documents Useful to NICs, LAN managers and end users

  1. Expands work done by Marty Schoffstall and Francine Perillo
  2. Adds:
     - What is the Internet?
     - What services are available?
       - mailing lists
       - enhanced services
     - How to ftp, email, set-up campus NIC/NOC, etc.
     - NOC-Tools Bibliography
     - Glossaries
- Repository
  - U of Texas - Tracy LaQuey/
    tracy@emx.utexas.edu

II. NISI Working Group (Network Information Services Infrastructure)

- Requirements document
  Internet->National->Regional->Campus
- Design Internet-wide information services?
- WG to be formed immediately?/Exploratory meeting to be held at NRI in May (NIC representation)

III. Investigate "Campus Awareness"

- SIGUCCS
- MERIT
- NSF (Dan VanBelleghem)
- USWG involvement?

VI. Network Status Briefings and Technical Presentations

Everything You Ever Wanted to Know About OSPFIGP
Presented by John Moy

This is a short introduction to the OSPF protocol. It has been developed during the past year in the OSPFIGP working group of the Internet Engineering Task Force.

OSPF is an IP routing protocol, intended to be used internal to an Autonomous System. In internet terminology, such a routing protocol is called an Internal Gateway Protocol. The OSPF utilizes SPF-based routing technology in order to find the set of best paths to each internet destination. The "O" in OSPF stands for open; it is hoped that multiple vendors will implement the protocol and interoperate.

OSPF has benefitted from the existing SPF routing technology. BBN first developed an SPF-based routing algorithm for the Arpanet in the 1970s. A paper by Radia Perlman ("Fault Tolerant Broadcast of Routing Information") introduced modifications to the SPF algorithm (e.g., the lollipop-shaped sequence space and the addition of a checksum field to links state advertisements) that enabled a reduction in the amount of routing traffic, and the removal of the link-up waiting time. DEC's IS-IS proposal introduced the concept of a Designated Router which generates a link state advertisement for transit networks. Finally, BBN did some work on area routing in an SPF-based system.

Based on this foundation, the OSPF working group was formed in the spring of 1988. The major features of OSPF are as follows. There is fast response to topology changes with a minimum of routing traffic. When multiple best paths are available to a destination, they are disovered and used. Separate sets of paths are possible for each IP Type of Service. A network mask is passed with each advertised destination, enabling "variable length subnet masks". Externally derived routes (e.g., EGP routes) are tagged and distributed independently from internal OSPF routes. All OSPF routing protocol packet exchanges are authenticated. Finally, OSPF protocol traffic uses IP multicast instead of broadcast.

OSPF also has an area routing scheme. This is very similar to the area routing developed by BBN. In OSPF area routing, routing inside any particular area is protected from outside interference. Also, the topology of the area is invisible from outside the area (similar to an IP subnetted network). Finally, the area ID is NOT encoded into the destination addresses.

Quickly, this is how OSPF works. Link state advertisements describe the local topology. Each router originates a link state advertisement, called a "router links advertisement". This indicates the type, cost, and state of each of the router's interfaces, together with what the interface attaches to (a Page

transit network, stub network, or to another router via a point-to-point connection). Each transit network has a "network links advertisement" originated for it by the Designated Router. This link state advertisement lists the set of routers attached to the transit network. Both of these link state advertisements are flooded throughout the routing domain. The collected set of advertisements forms the routing database. This database is identical in all nodes. From this database, each node calculates a shortest path tree, with itself as the root. This calculation yields the routing table.

The presence of areas modifies the above somewhat. The algorithm executes in each area as above, calculating all the intra-area routes. Area border routers (those attached to more than one area) learn routes to destinations in other areas, and transmit this learned information to their attached areas by means of "summary link advertisements". This third kind of link state advertisement is also flooded throughout a single area.

All area border routers must be attached to a single area: the "backbone". The backbone transmits the "inter-area" routing information (routes between areas). The backbone must be connected; all areas are attached to the backbone, forming a star topology with the backbone as hub. Areas may dedicate some of their resources to the' backbone; this enables the maintenance of backbone connectivity through the configuration of "virtual links".

External routes are described by a fourth kind of link state advertisement, "AS external link advertisements". These, unlike the previous three advertisement types, are flooded throughout the entire Autonomous System instead of just throughout a single area. These advertisments are considered last when building a routing table. The following comparisons can be made between OSPF and the DEC IS-IS proposal. Both protocols are SPF-based, and as such use many of the same mechanisms (flooding, the shortest path calculation, etc.), just as any two Ford algorithms use many of the same mechanisms (broadcast of routing tables, etc.). The main differences between OSPF and the DEC proposal cna be broken up into SPF differences, area routing differences, and special IP considerations.

The SPF differences include the following. OSPF ensures that a router's routing database is synchronized before the router is allowed to forward data traffic. This guards against packet looping. OSPF has made the following routing traffic reductions: on a transit networks database synchronization occurs only over OSPF adjacencies (an o(n) problem rather than o(n**2)), no attempt is made to sychronize link state advertisement ages, and external routes are specified each in a separate advertisement (allowing incremental updates). OSPF allows the specification of

two types of external metric (one comparable to the link state metric, and the other larger than any link state path). Finally, OSPF has no special "link state confusion logic"; link state checksum conflicts are treated the same as sequence number conflicts.

OSPF area routing looks much like the BBN area scheme, instead of the DEC IS-IS areas. The OSPF area ID is not part of the destination address. This allows the intelligent selection of exit/entry routers when routing to destination areas, and avoids the introduction of area partition repair logic (partitioned areas instead appear as two separate areas). The OSPF backbone is similar to the level two routing in the DEC IS-IS proposal. However, the OSPF backbone need not be physically connected, and may instead be connected by means of virtual links.

Finally, OSPF is an IP routing protocol, while the DEC IS-IS is an ISO routing protocol. OSPF passes around native IP addresses, and provides explicit IP subnetting support. OSPF packet formats have been designed so that they can be efficiently parsed in an IP environment. The packet formats have also been designed so that IP fragmentation and assembly should not be necessary. Finally, OSPF should provide some experience with IP multicast.

OSPF

An internet routing protocol

OSPF Topics

- Introduction and History
- List of features
- Protocol operation
- Comparison to DEC IS-IS

What is OSPF?

- A routing protocol
- Classified as an IGP
- Uses SPF technology
    - Distributed database
    - Replicated in all nodes
- The "O" stands for open

OSPF Roots

- BBN develops for Arpanet (1978)
- Enhancements by R.Perlman (1983)
    - Fault tolerant routing
    - Avoids line-up wait
- BBN work on area routing (1986)
- DEC modifications for broadcast (1987)
- OSPFIGP working group forms (1988)

OSPF features

- Resistant to routing loops
- Small amount of routing traffic
- Fast convergence when topology changes
- Equal-cost multipath
- Uses configurable 16-bit metric
- Separate metric for each TOS

OSPF features (cont.)

- Subnet masks attached to routes
- External routes are tagged
- Routing exchanges authenticated
- IP multicast used instead of broadcast
- Two levels of routing (areas)

Area Functionality

- Intra-area routing protected
- Information hiding at area boundaries
- Generalization of subnetted network

How OSPF works (single area)

- Link state advertisements describe local topology
- Routers originate advertisements
- Transit networks originate advertisements
    - via Designated Router
- Advertisements flooded throughout area
- Collected advertisements form routing database
- Dijkstra calculates routing table

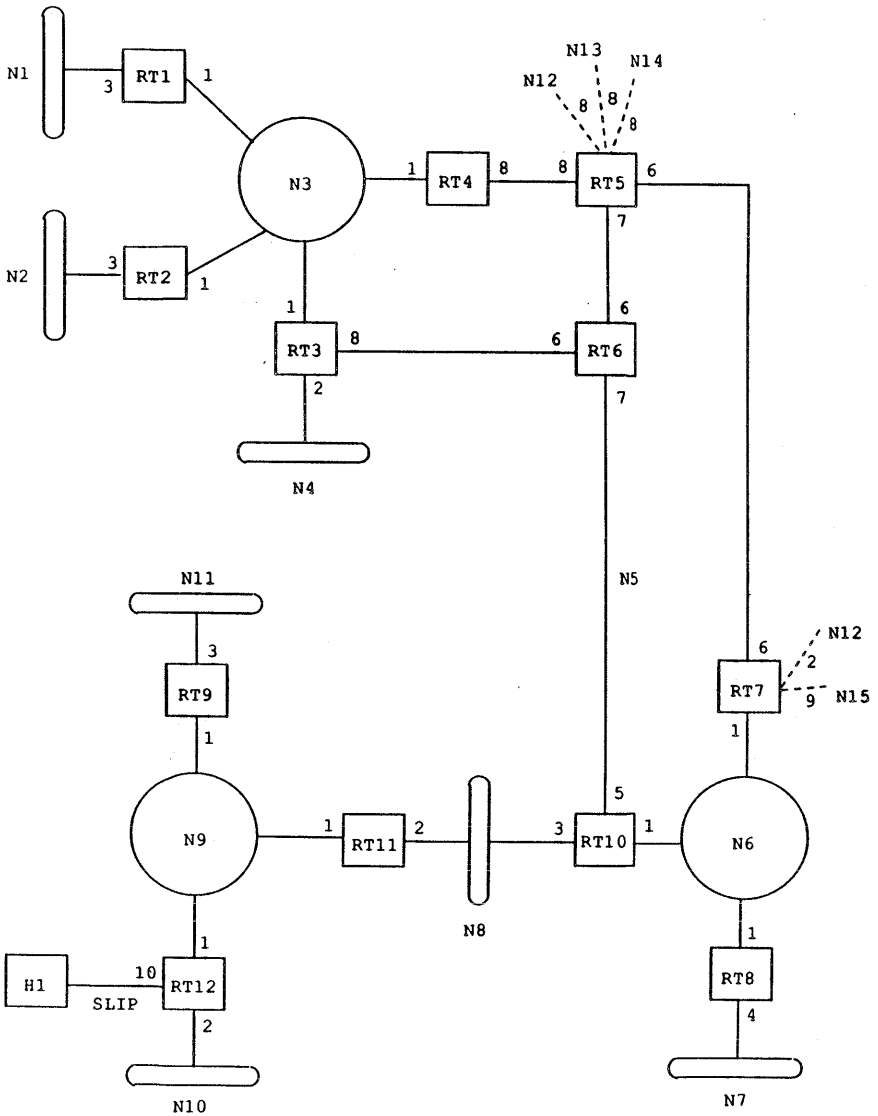

A Sample Autonomous System Configuration

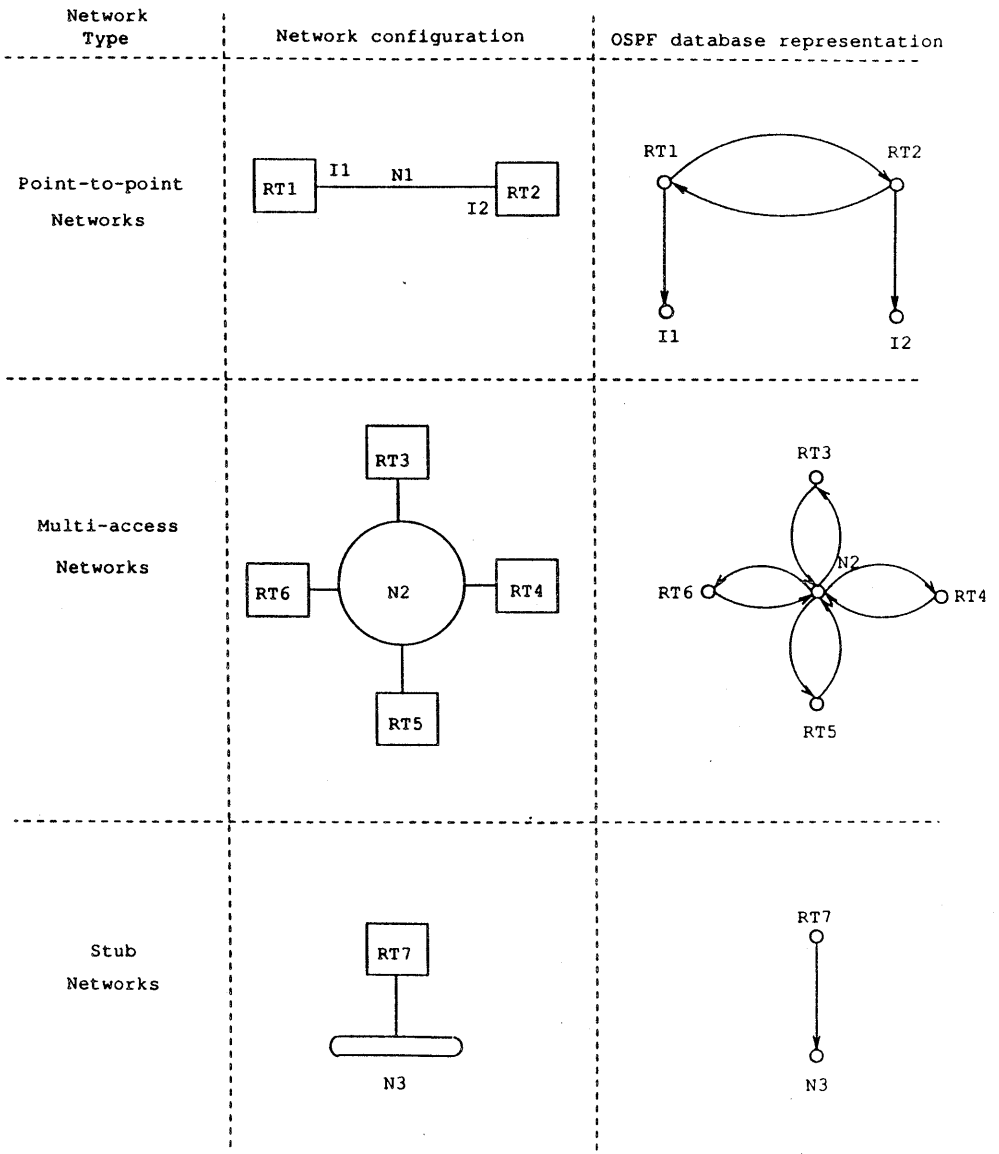

OSPF Routing Database Representation

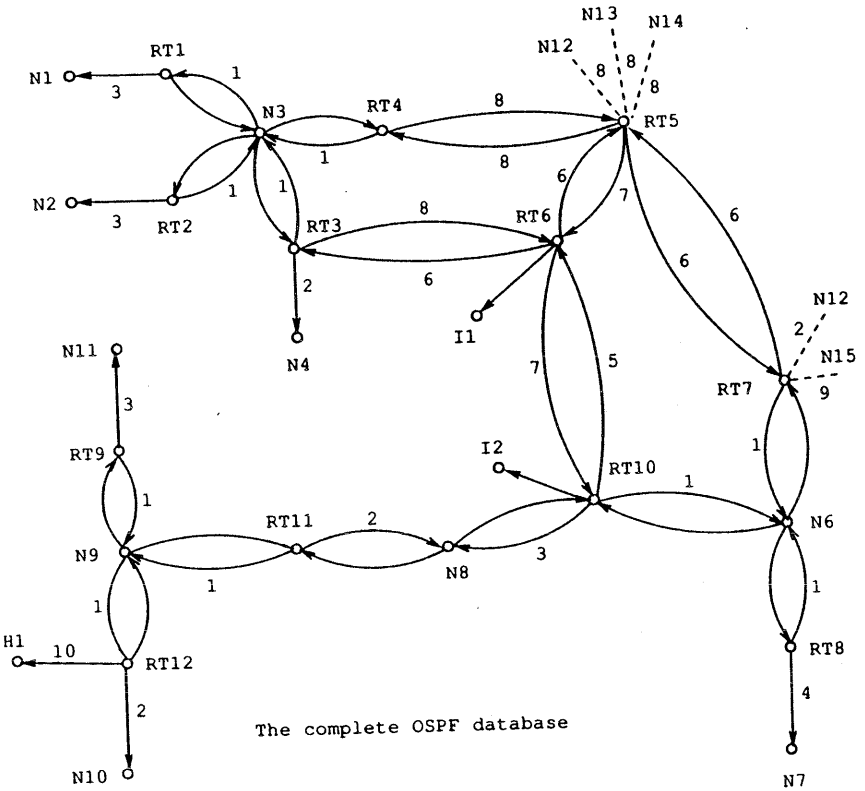

The complete OSPF database

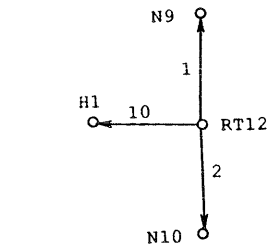

A router links advertisement

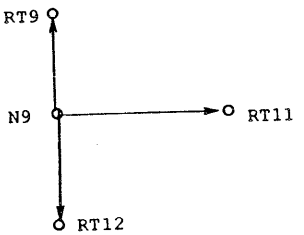

A network links advertisement

Individual link state components

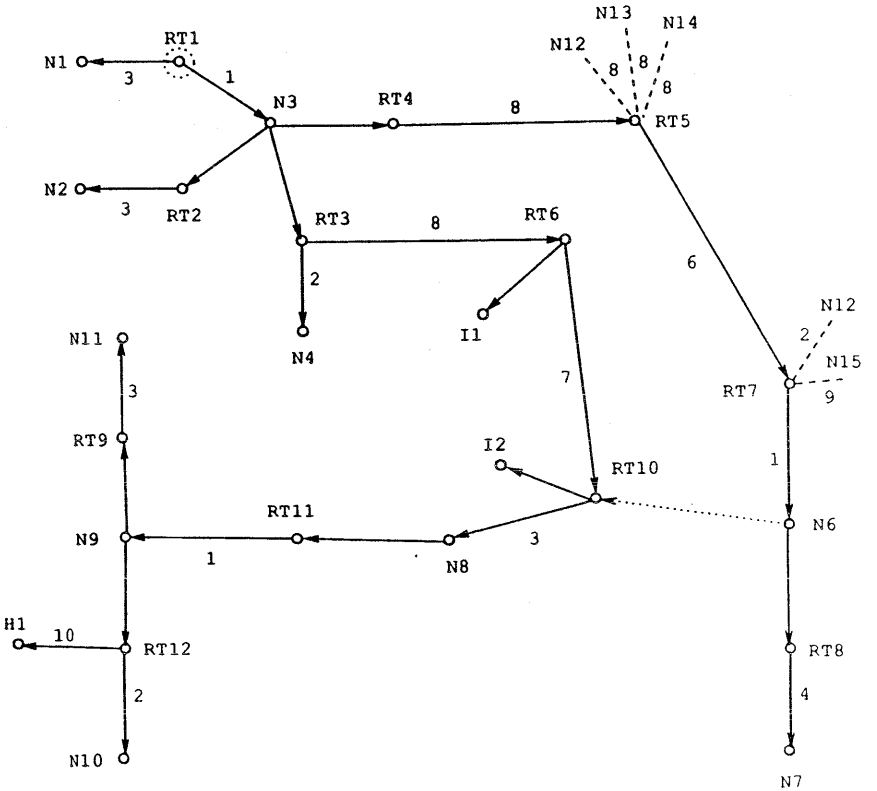

Resulting SPF tree (rooted at RT1)

Inter-area Routing

- Extra functionality in area border rtrs (ABR)
- ABRs describe routes to non-area dest.
    - Descriptions called summary links
- Summary links flooded throughout area
- Examined after Dijkstra procedure

The Backbone Area

- Connects area border routers
- Enables learning of inter-area routes
- Area routing forms star, with backbone as the hub
- Virtual links maintain backbone connectivity

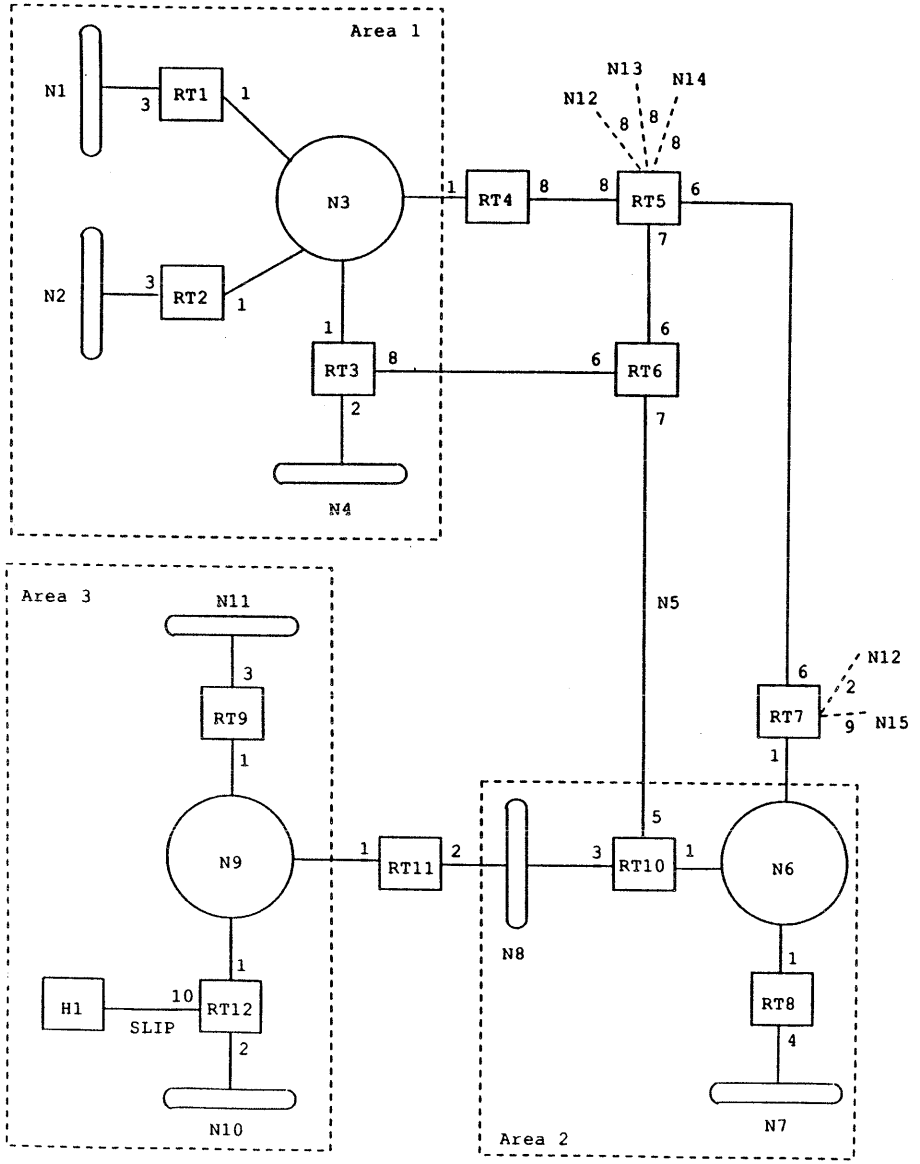

A Sample OSPF area configuration

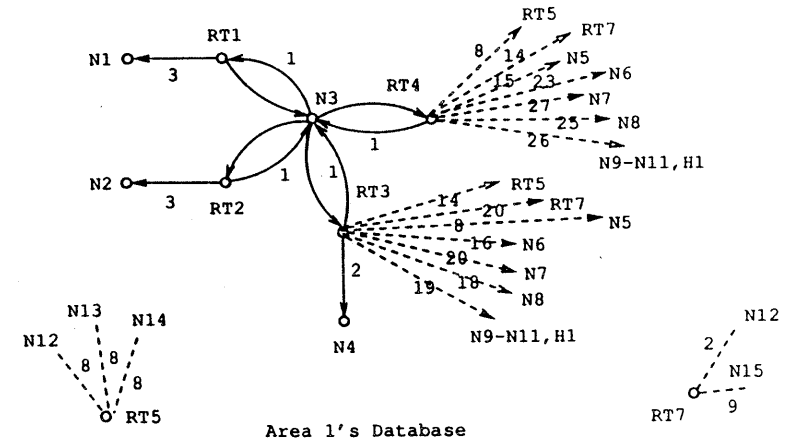

Area 1's Database

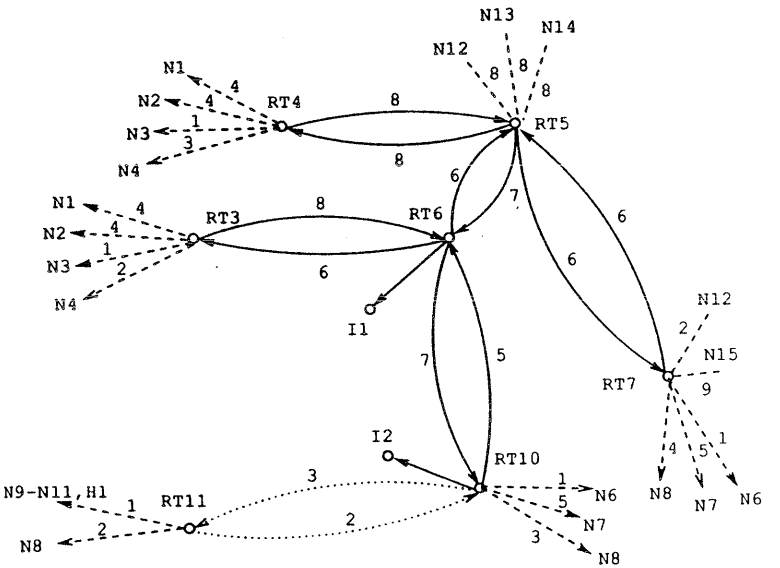

The backbone database

External routing info

- Described by "AS external link" adv.
- These are examined last in routing table build process
- Flooded throughout all areas
- Two available metric types:
  - Same as link state (type 1), or
  - Larger than any link state path (type 2)

OSPF Routing Hierarchy

- Intra-area routes
- Inter-area routes
- External type 1 routes
- External type 2 routes

More Details

- Hello Protocol
  - Discovers, maintains neighbors
  - Elects Designated Router
- Adjacencies established between selected routers
- For transit nets:

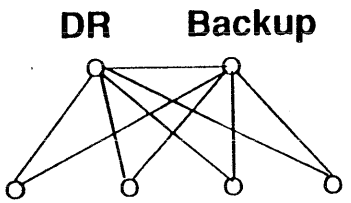

More Details (cont.)

- Full adjacencies
  - Require synch. databases
  - Are reflected in link state adv.
- Flooding done on (partial) adjacencies
  - o(N) pairs to sync, not o(N**2)
- Backup Designated Router
  - Hot standby for DR

Comparison to IS-IS

- SPF considerations
- Area considerations
- IP protocol considerations

SPF Considerations

- Database synchronization before joining routing domain
- Routing bandwidth reductions
  - Flood only over OSPF adjacencies
  - No age synchronization
  - External links in separate adv.
- Two types of external metric
- No special "link state confusion" logic

Area considerations

- Remove topological constraint on level two domain (backbone)
- Area ID not part of address
- No need for area partition repair

IP Considerations

- Advertise native IP addresses
- Keep flat IP address space
- Provide IP subnetting support
- Packets designed so that:
  - they can be efficiently parsed
  - IP fragmentation unnecessary
- Use IP multicast

The Open Routing Architecture
Presented by Marianne Lepp

The AS-AS routing architecture of the Open Routing Working Group has been designed for an Internet in which 10,000s of entities will participate in the routing. We expect that the bulk of the transit traffic will be carried by a small number of networks/ASs that are designated as fully transit. For ease of discussion, we call them "backbones". Other systems will carry a limited amount of transit traffic, dictated by policy agreements. Many others will act as stub systems carrying no transit traffic.

We are designing for a reasonably simple topology with back-doors and short-cuts. We expect to be connecting heterogenous systems, where heterogenaity include gateways, protocols, and network technologies.

The conceptual elements of the protocol are Routing Agents, Policy Agents, Forwarding Agents, User Agents, and Data Collection Agents. The Routing agents compute routes based on topology, policy, and type of service. They negotiate with the User Agent about whether a service can be provided or not. The Policy Agent maintains the policy database, including the validation and sanity checking functions. The Forwarding Agent is what we currently think of as an IP-router. It accepts packets for forwarding. The User Agent may reside in the host or the host's gateway and negotiates a route based on the application/user's policy credentials. The Data Collection Agent supports the dynamic features of the protocol.

The key features of the protocol are data reduction by recursively dividing the Internet into "areas", source routing, route set-up, and link attribute lists. The routing element is the way-station, which is the entrance point to Autonomous Regions. More information about the architecture itself can be found in the slides.

Open Routing Architecture

ORWG
Marianne Lepp, chair

April 13, 1989

Overview
Problem & Requirements
Architecture
Status and Timetable

Requests for interest mailing list → mlepp@bbn.com

Driving Requirements

- 10,000s routing entities
- General topology
- Complicated (but not all) policies
- Heterogeneity
  - boxes
  - protocols
  - network technologies
    - T3
    - circuit
    - :
- Limited cooperation
- Security
- Performance
  - CPU
  - Bandwidth
  - Robustness
- Conceptual Simplicity

Architecture

Conceptual Elements

- Routing Agent
- Policy Agent
- Forwarding Agent
- User Agent
- Data Collection Agent

Features

- Data Reduction
- Source Routing
- Route Set-up
- Link State
- Security

Problem(s)

- Full mesh routing protocol
- AS-AS
- usable in Internet today
- Will grow gracefully
- Policy-based

ORWG World View

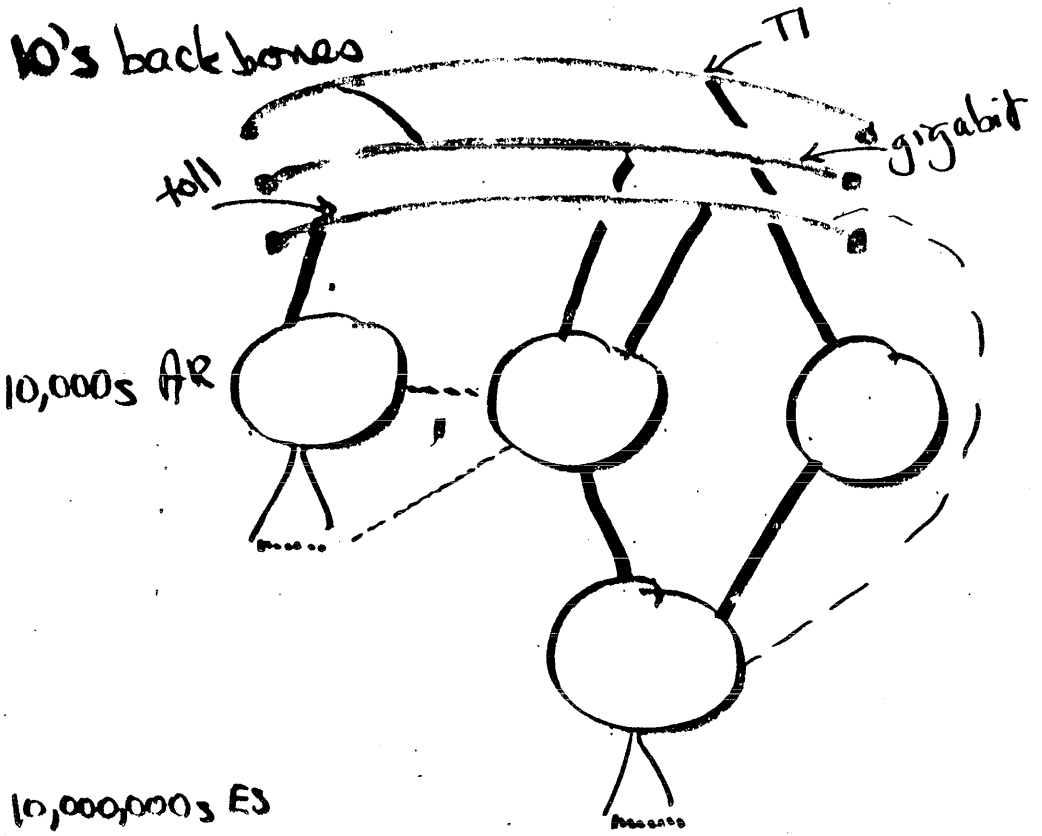

ARs can be: stub
limited transit
fully transit

(Problem): Universal Connectivity

Data Reduction

Why
- Scale
- Privacy
- Reduction of Information Exchange
- Graceful growth
- allow growth above + below AR level

How
- ARs → points
- AR may view itself as several points
- many APs can be viewed as a point

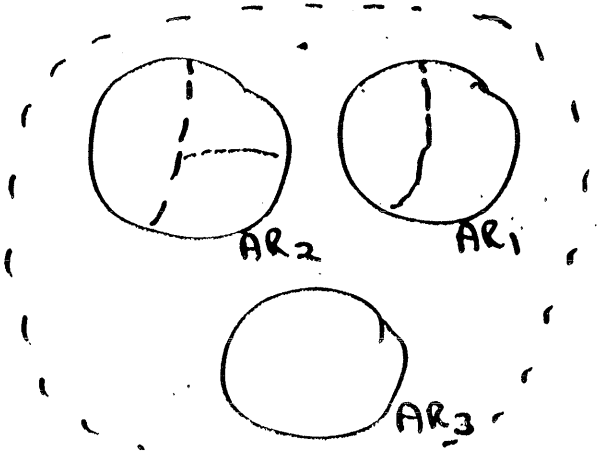

Data Reduction (cont.)

Way Stations
- Entry points to ARs
- Routing Element
- may be several physical BWS

"Areas" are configured

Partitions are repaired dynamically
- Policy - don't fix flag

Virtual Links
- Route tragically
- Kept up by way-stations

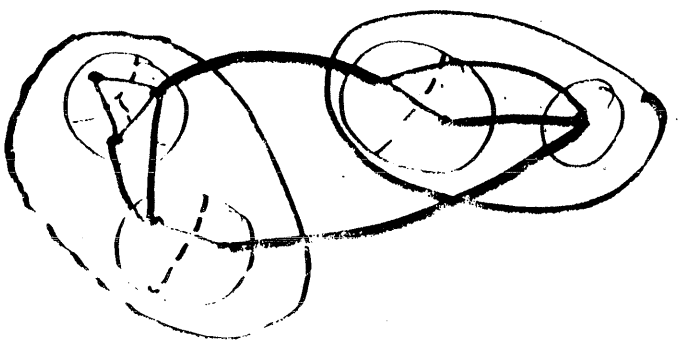

Source Routing

Why

- Select user policy
- Hide user policy
- Resolves routing consistency problems

How

- Way-station to way-station "tack" route to Internet
- way-stations route on their virtual links

Route Set-up

Why

- Reduce packet-level overhead
- Reduce processing overhead
- hardware forwarding functio easy
- resource reservation easy
- vehicle for enforcement of policy

How

- map interface, next hop to flow id
- data goes in set-up packet (can erase behind it)
- includes policy information

Route Set-up

Route repair

- local
- notify only as high up as needed to effect repair
- source not notified
- policy must be visible
- some policies will require source re-routing

Link State

- attributes
  - Policies
  - Physical properties
- areas and virtual links
- Dynamic changes on a very slow scale
- The higher the level, the slower the change
  - links
  - networks
  - ARs
  - :

Security

Protocol updates

Route / flow / policy authentication

Data authentication

Status & Timetable

Requirements	IDEA 007
Architecture	draft Sep 30
Functional Spec	July 1989
Draft protocol spec	December 1989
Test implementations	early 1990

NASA Science Internet

Milo S. Medin
Sterling Software Corporation
NASA Science Internet Project Office
NASA Ames Research Center
medin@nsipo.nasa.gov

NASA HEADQUARTERS ORGANIZATION CHART

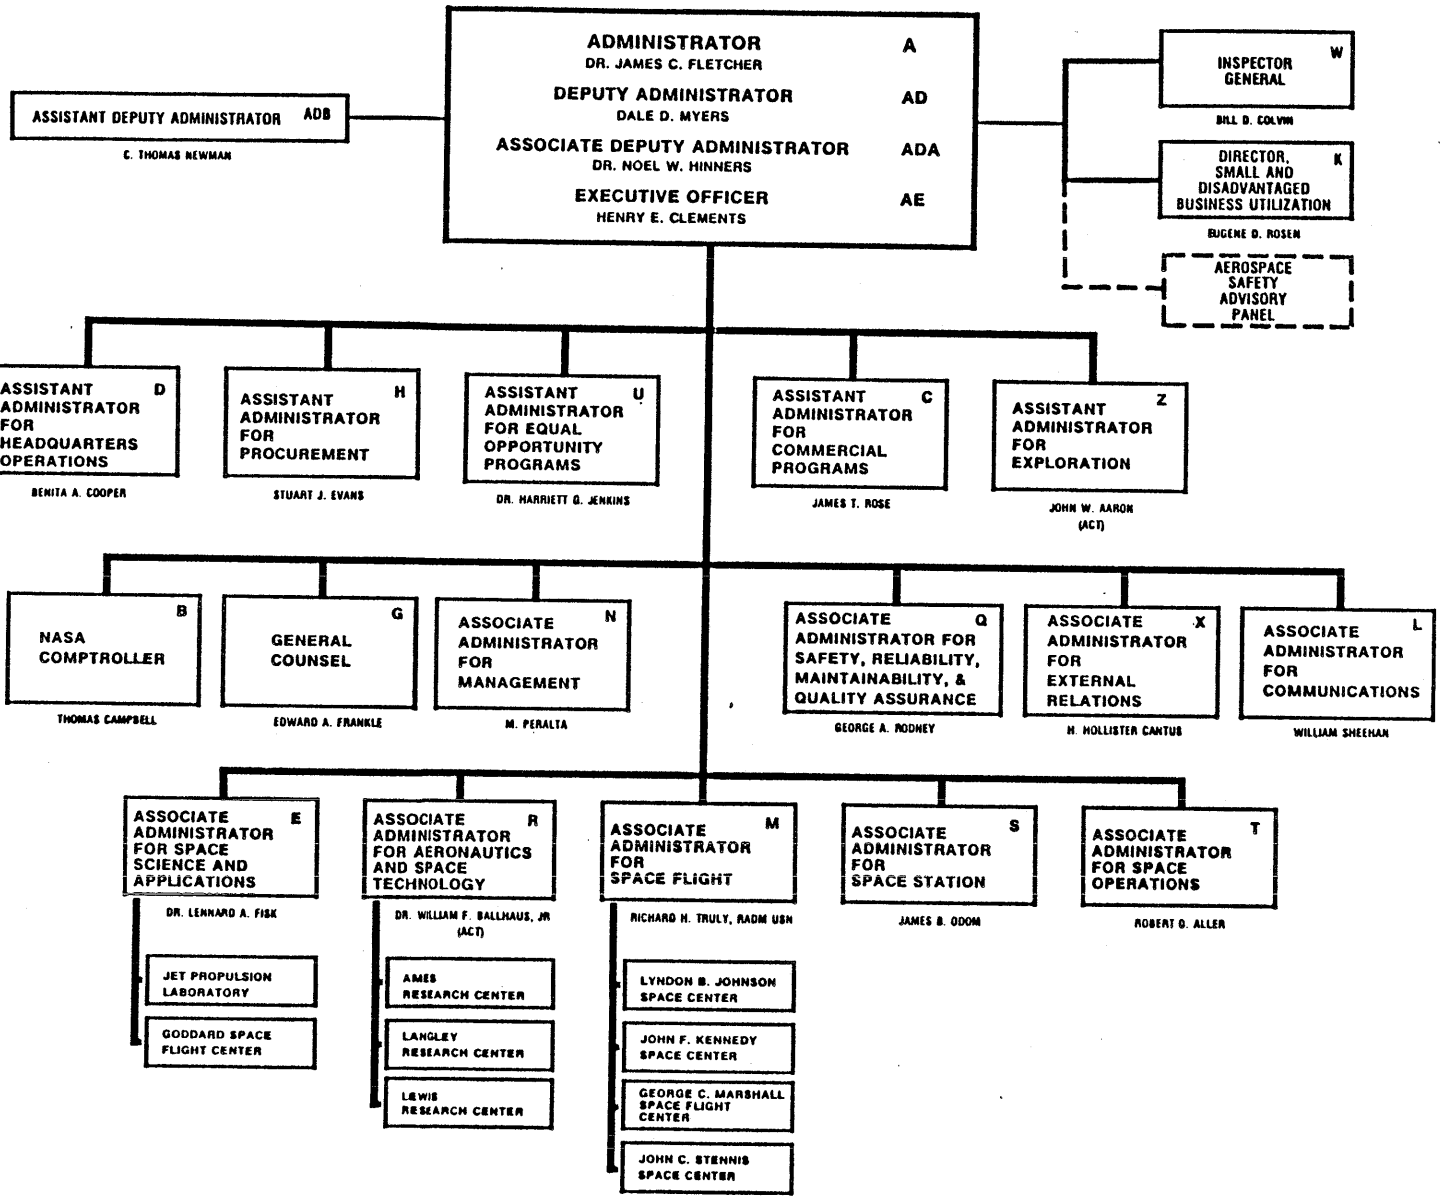

NOVEMBER 1, 1988

Who's on first?

Code E - Office of Space Science Applications (OSSA)

- sponsors NSI

Code R - Office of Aeronautics and Space Technology

- sponsors NAS program

Code T - Office of Space Operations

- provides communications support to other codes

NASA Science Internet Objectives

- Provide focus for OSSA networking
- Coordinate and consolidate diverse networking efforts
- Save money
- Increase performance
- Participate in federal networking efforts
- Provide better networking for the benefit of NASA science

NSI protocol requirements

- DoD IP
- DECnet Phase IV
- ISO OSI CLNP (8473 and friends) (future)

Multi-Protocol approaches

- Hardwired multiplexing
- Encapsulation
- Multi-protocol routers
- Convert the user base

Hardwired multiplexing

Simple +

Bypasses the real issue +/-

Inefficient use of lines -

Capital intensive -

Labor intensive -

Relatively inflexible -

DECnet encapsulation in IP (Multinet)

Efficient use of lines +

Ability to 'tunnel' through other IP networks

Capital intensive -

Labor intensive -

Added complexity -

Multi-protocol routers

Efficient use of lines +

Ability to interoperate with other IP nets +

Capital efficient +

Labor efficient +

Higher performance +

Complex -

Convert the user base

Efficient use of lines +

Capital efficient +

Labor efficient +

Interoperability with other nets of same protocol +

Best network engineering choice in most cases +

Completely impractical for a variety of reasons -

Space Physics Analysis Network (SPAN)

HQ at Goddard Space Flight Center (GSFC)

Carries exclusively DECnet

"Single Vendor Solution"

Large installed base due to age

Lines are mostly 9.6 Kbps

SPAN
North American Sites

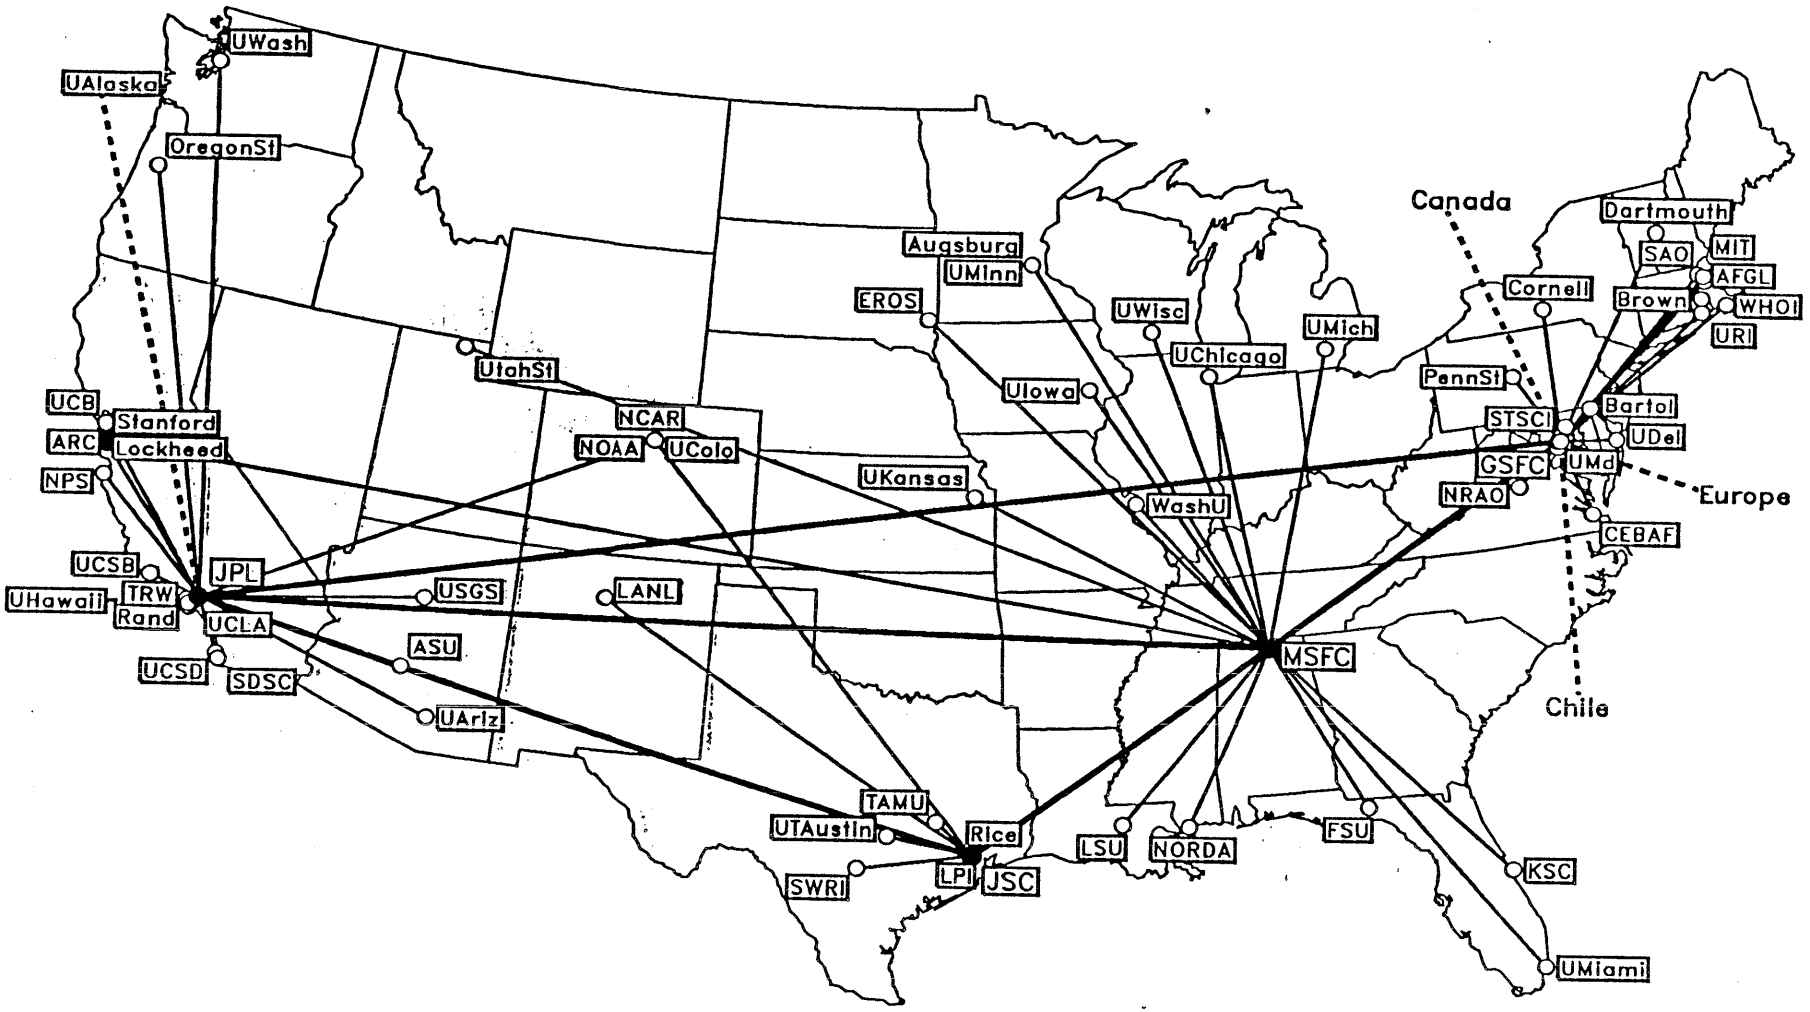

Prepared for: NASA Science Internet Project Office
by: Sterling Software, NAS2-11555
12 March 1988

NASA Science Network (NSN)

- HQ at Ames Research Center
- Carries both IP and DECnet
- Connects to NSFnet, DDN, regionals, sites, etc...
- Connects about 40-50 nets currently
- Uses Proteon p4200 routers
- Most links at 56 Kb, upgrades in works
- Major Internet connect points at ARC & GSFC

NASA Science Network
Gateway Interconnections

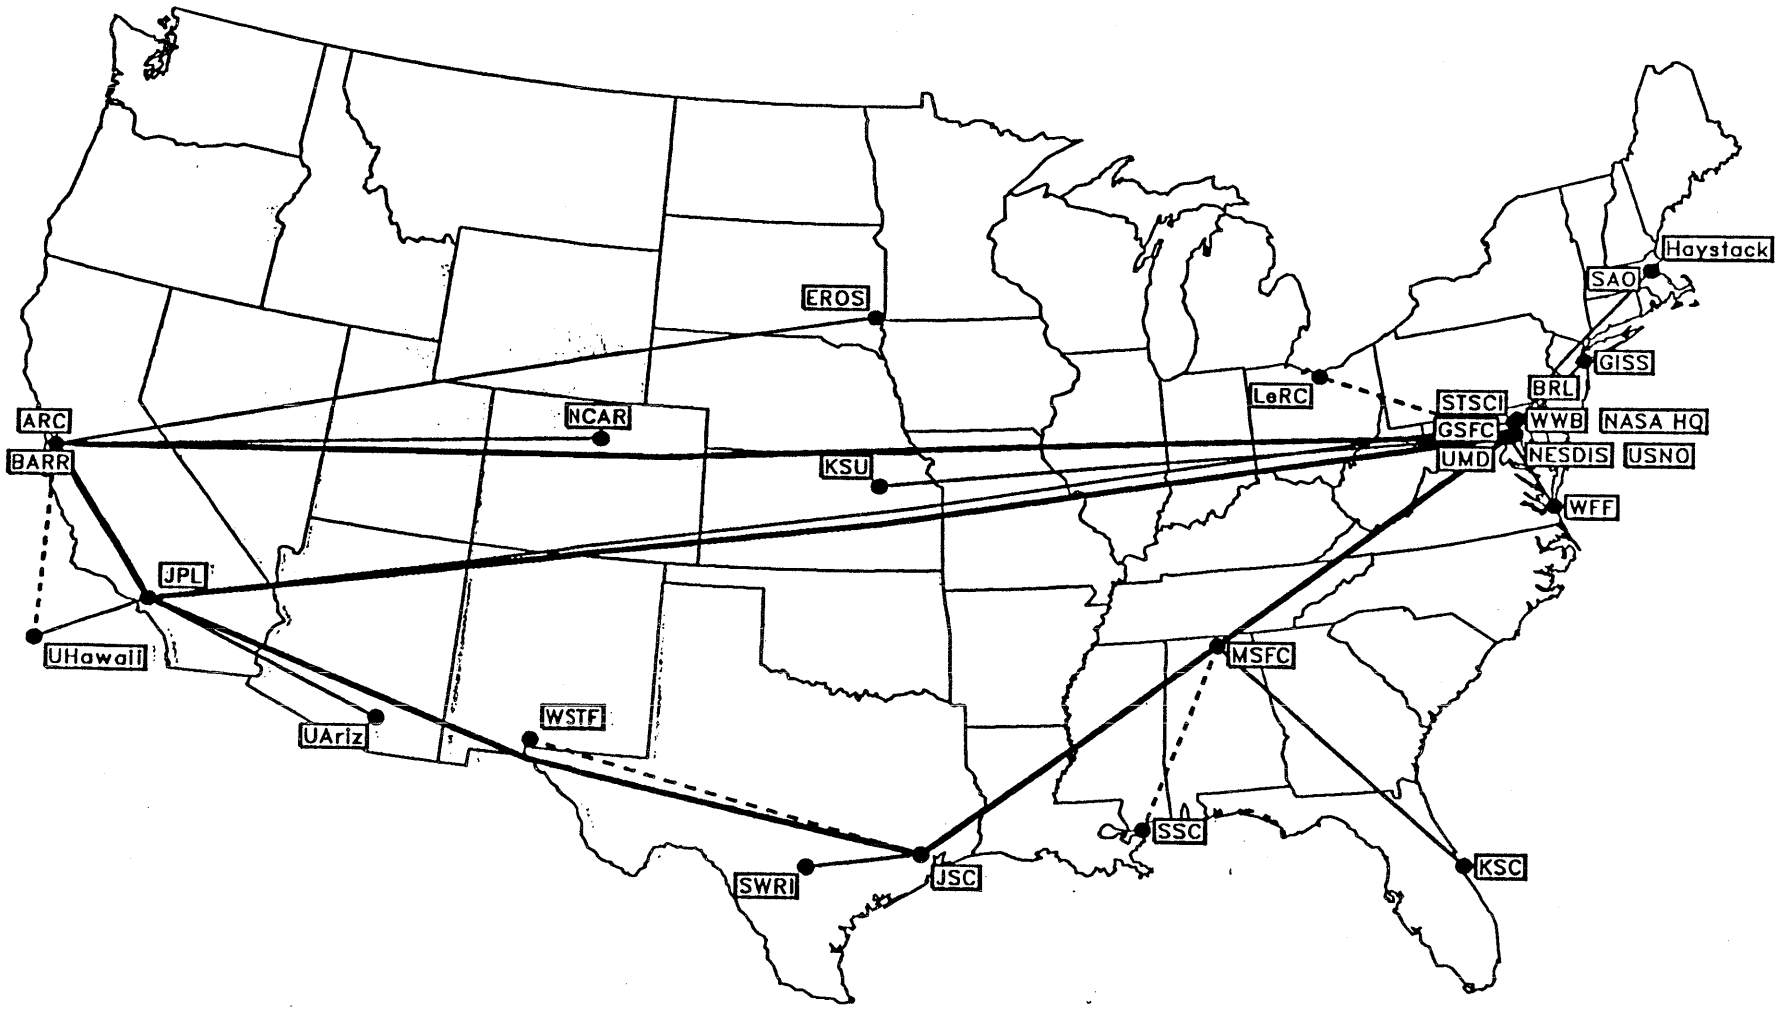

Prepared for: NASA Science Internet Project Office
by: Sterling Software, NAS2–11555
23 January 1989

NASA Science Internet Project Office (NSIPO)

- HQ at Ames Research Center
- Funds both NSN and SPAN activities
- Technical review of requirements and engineering
- Responsible to NASA HQ and science projects
- Primary driver behind interoperability efforts
- Primary interface to other codes and users

NSN & SPAN interoperability issues

- Mail interconnectivity now
- Login and FTP capabilities via 'Ultrix Dgate' shortly
- Security issues (YUK!)
- Other networks and protocols
- Multi-protocol forever ?

NSI communications capabilities

- Primary 'carrier' is Code T
- Lines come from a voice/data network called PSCN
- Switched 56 Kb service, and increments in DS-0 units
- Other 'carriers' are being vigorously pursued

NSN management support

- Primarily use SGMP/SNMP
- PC based status tools (Overview, NETMON)
- Unix (SUN) based statistics tools (using UCB code)
- Primary concern involves lines (status, congestion, etc)
- No statspy-type support deployed presently
- X.25 PAD and dial-up console port support

NSN circuit utilization

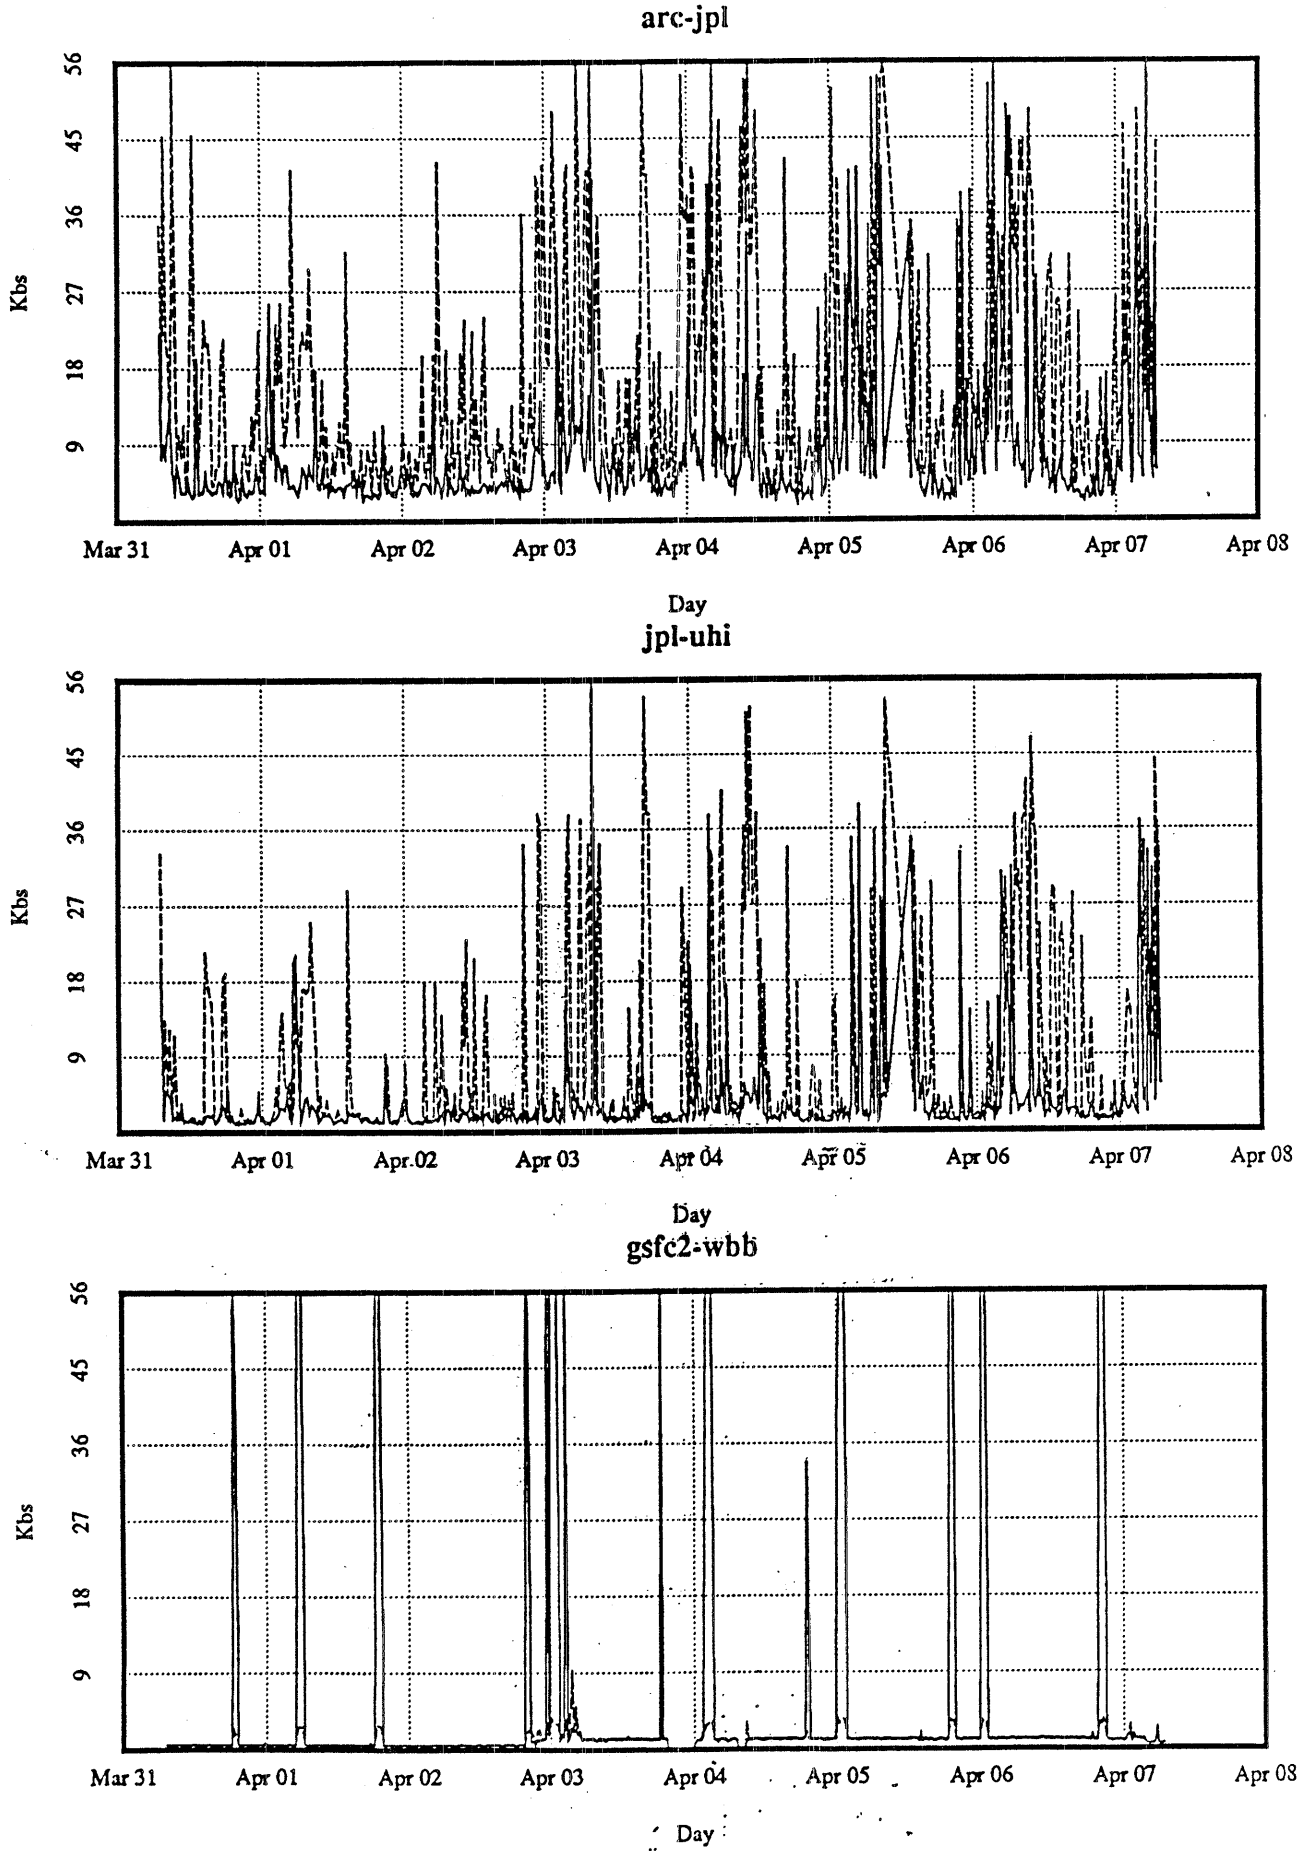

Future activities

- Better interoperability support
- More direct NSFnet interconnect via NSS EGP
- OSPFIGP integration and deployment
- Much stronger efforts to consolidate SPAN & NSN
- OSI CLNP support
- Higher bandwidth links
- Further expansion of sites and technologies

State of the Internet
Presented by Zbigniew Opalka

The number of networks in the Internet continues to grow. Since the beginning of the year the number of networks, as reported by the Butterfly Mailbridges, has risen from approximately 550 to 750 by April 1. If this trend continues, as is expected, the Internet will reach 1,000 networks by the third quarter of this year.

The Arpanet is steadily going away. Sites previously on the Arpanet are moving to various regional networks across the country. These regional nets are connected by the NSFNet; though connectivity to the Milnet is still provided by the Mailbridges.

A new experimental terrestrial Wideband network is planned to provide connectivity between some sites currently attached to the Arpanet. The old satellite Wideband is being replaced by a network connected by T1 trunks, using a modified IEEE 802.6-type technology. The services provided by this network include datagram service (unannounced traffic), stream messages (resources reserved across the network), as well as a multi-casting scheme using dynamic group addressing.

The terrestrial Wideband network is composed of Butterfly Wideband Packet Switches (WPS) connected by T1 trunks. Attached to the WPSs are Butterfly Internet Gateways and Stream (ST-protocol based) Gateways.

Mailbridges

All six of the Butterfly Mailbridges are operational. They are currently supporting 210 neighbors (combined on both the Arpanet and Milnet). The LSI-11 Mailbridges were decommissioned on March 6, though the LSI-11 EGP servers are still functional. It is highly recommended that anyone still using the LSI-11 EGP servers move over to the Butterflies as soon as possible.

The Mailbridges are passing around 8 million packets per day. This figure varies greatly, anywhere between 6 and 13 million packets are sent per day. The drop rate across the Mailbridges for queuing reasons (not counting ttl expiration and unreachability) is insignificant (a total of 675 packets or .00856 %). Average length of the packets passing through the Mailbridges is around 150 bytes.

As stated earlier, the Mailbridges provide the connectivity between the Arpanet and Milnet. To enhance Milnet connectivity with the rest of the Internet, Ethernet interfaces will be added to the Butterfly Mailbridges first at Ames (west coast) and Mitre (east coast), then to the other Mailbridges. The Ethernets, at Ames and Mitre, will also have attached to them an NSS, an NSFNet backbone switch.

Two bugs have been discovered in the Mailbridges. The first deals with decrementing the TTL field in the IP header and then dropping the packet is the TTL drops to 0. The Mailbridges were decrementing the TTL as required. They would drop a packet if its TTL was 0 when it first entered the Mailbridge. The problem occurs when a Mailbridge decrements a TTL whose value was 1 when it entered the Mailbridge, the Mailbridge would check for 0; decrement the 1 (to 0) and forward the packet. The fix will be deployed in the next major release of the Mailbridge code,

The second problem dealt with the EGP finite state machine (fsm) implementation. EGP, after neighbor acquisition, expects either Polls or Hellos to bring the EGP "link" up. The Mailbridges were expecting only Hellos after the initial neighbor acquisition, thereafter they would accept either Hellos or Polls. The code was fixed to accept both types of messages after neighbor acquisition and the fix has been deployed.

STATE OF THE INTERNET

Zbigniew Opalka

BBN Communications Corporation

TOPICS

- Internet Growth
- Arpanet
- Wideband Net
- DDN Mailbridges

INTERNET GROWTH

NUMBER OF NETWORKS
DECEMBER 1983 - APRIL 1989

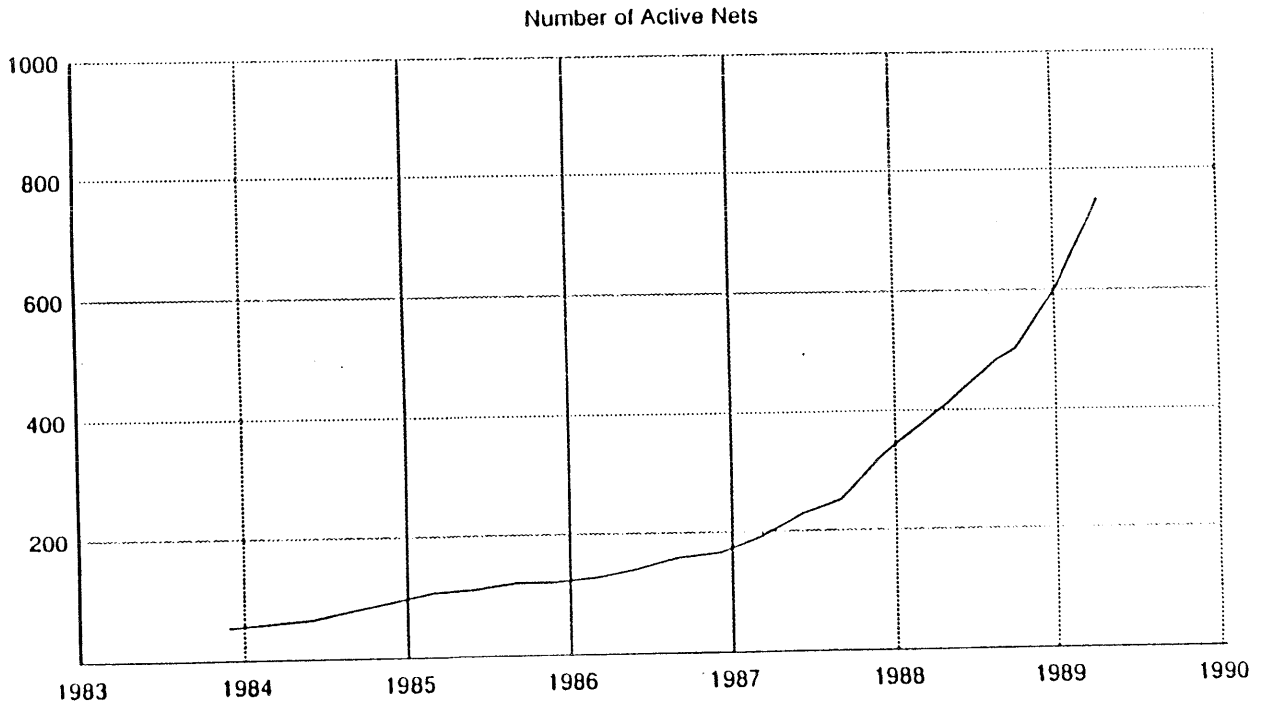

NUMBER OF NETWORKS
DECEMBER 1983 - APRIL 1989

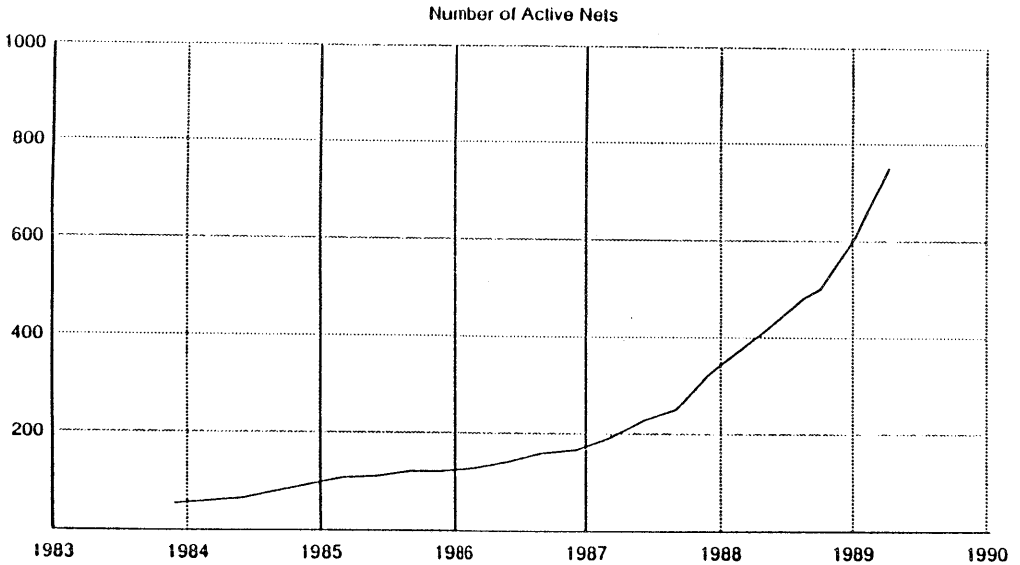

NUMBER OF NETWORKS
JANUARY 1989 - APRIL 1989

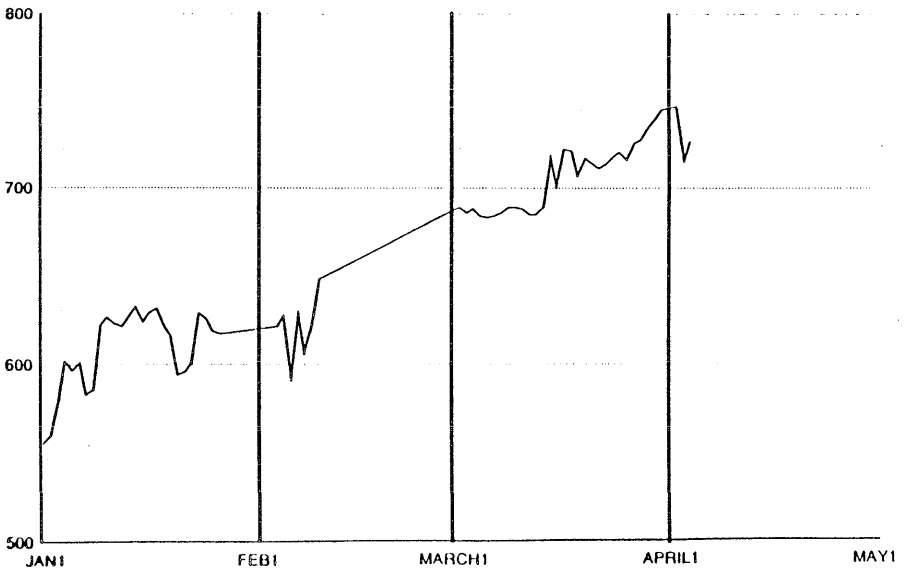

ESTIMATED GROWTH IN NUMBER OF NETS

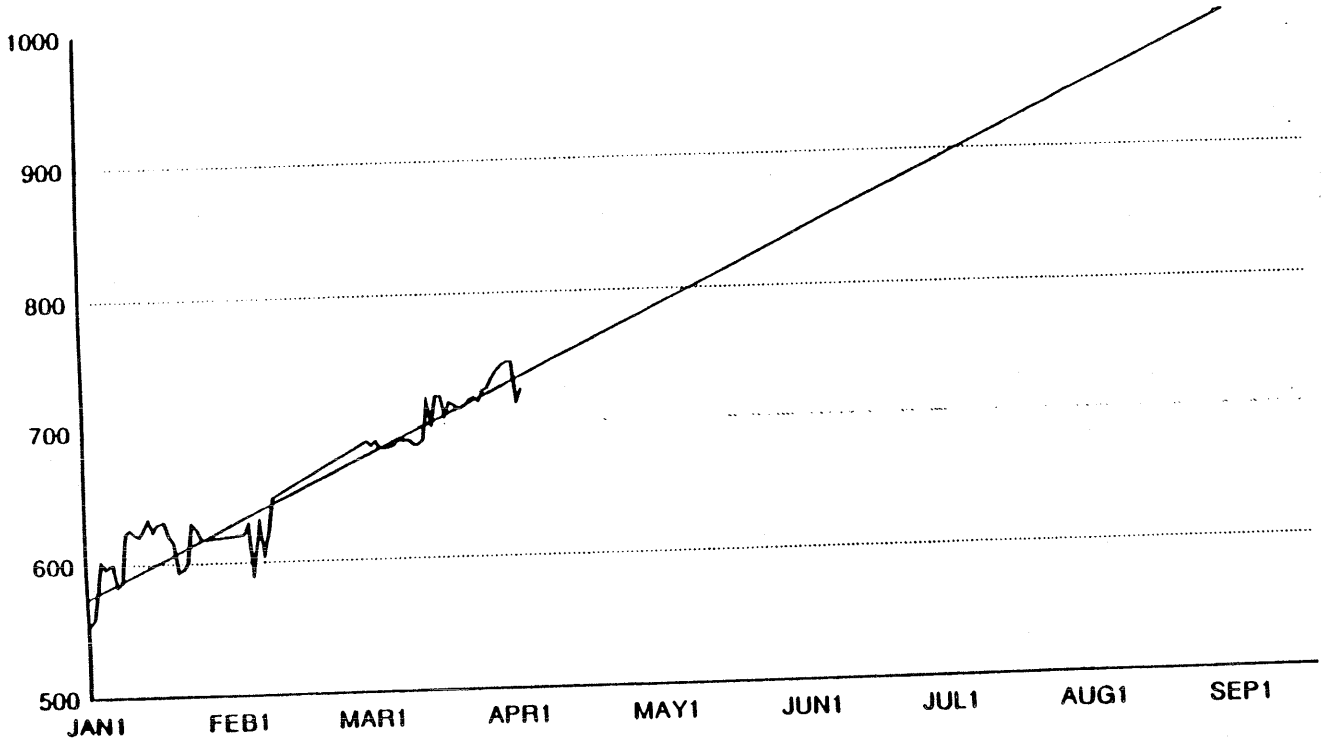

ARPANET

- Being Phased Out
- Sites Moving to Regional Networks
- Regional Nets Connected
  - NSFNET
  - Terrestrial Wideband
  - MILNET

WIDEBAND NET

- Current Satellite Net being phased Out
- Replaced with Experimental Wideband Network
- Composed of Modified IEEE 802.6 MAN Using T1 Links
- Terrestrial Wideband will take Load off Remaining Arpanet

TERRESTRIAL WIDEBAND SERVICES

- Datagrams (Unannounced Traffic)
- Stream Messages (Resources Reserved)
- Multicasting via Dynamic Group Addresses

T1 TERRESTRIAL NETWORK
INITIAL INSTALLATION

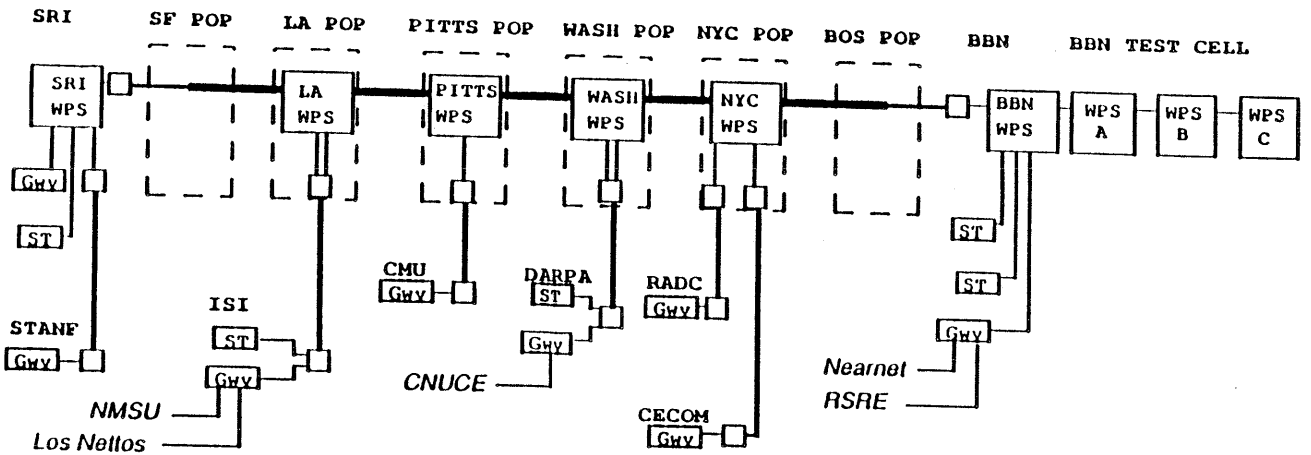

WPS = Wideband Packet Switch
GWY = IP Gateway
ST = ST Gateway
= T1 NNT Trunk
= T1 Tail
= Local Wire
= T1 DSU/CSU

MAILBRIDGE PROBLEMS

- TTL
  - "BUG", Test and Decrement TTL
  - Fix Implemented and will be distributed with next patch release
- EGP "State Machine" Fixed to Accept "Polls" in place of "Hello" after Initial "Acquire"
  - Fix deployed

CURRENT STATUS

- All 6 Butterfly Mailbridges Operational
- Current Number of Neighbors - 210
- LSI-11 Mailbridges Decommissioned on March 6
- LSI-11 EGP Servers Still Running
  - All Gateways should be Using New Mailbridges
  - Planned to be Decommissioned Soon

MAILBRIDGE CHANGES

- Ethernet Interfaces will be added to Butterfly Mailbridges
  - AMES
  - MITRE
  - Other to follow

MAILBRIDGE THROUGHPUT

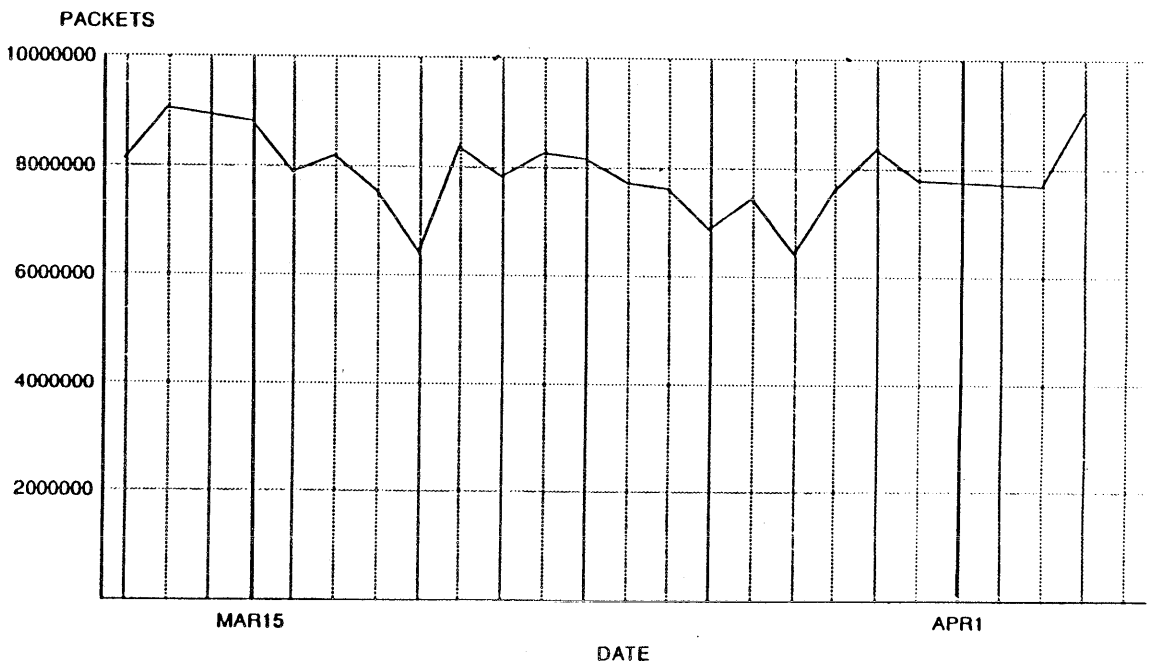

TRAFFIC SUMMARY

(AVERAGE - 21 DAYS)

- Packets/Day = 7,900,000
- Dropped/Day = 675
- % Dropped = .0085565%
- Average Packet Length = ~ 150 Bytes

Growth of the Internet
Presented by Mike St. Johns

Based on the accumulated 5 years of data from BBNCC regarding advertised networks, the growth of the internet appears to be exponential. Previous graphs of the growth of network numbers have been plotted on linear axises, and we've long suspected the growth was exponential. I regraphed the data on a semi-log graph (linear X is time, log Y is number of networks) and then did a fit on the data. The data line and the fit line appear to be pretty close. Unfortunately, I did not have access to tools which would have allowed a more formal statistical analysis of the data.

The doubling period of the data is approximately 13-14 months. We should reach 1000 networks by November of 1989, 8000 by March of 1993. If the trend continues, we could reach a million networks by sometime in 2000.

Based on the 5 year trend, I actually believe we could have as many as 8000 networks by 1993. I think its too early to believe the year 2000 prediction, but it is setting off some warning bells.

Advertised Nets
Nov 1986 - Mar 1989

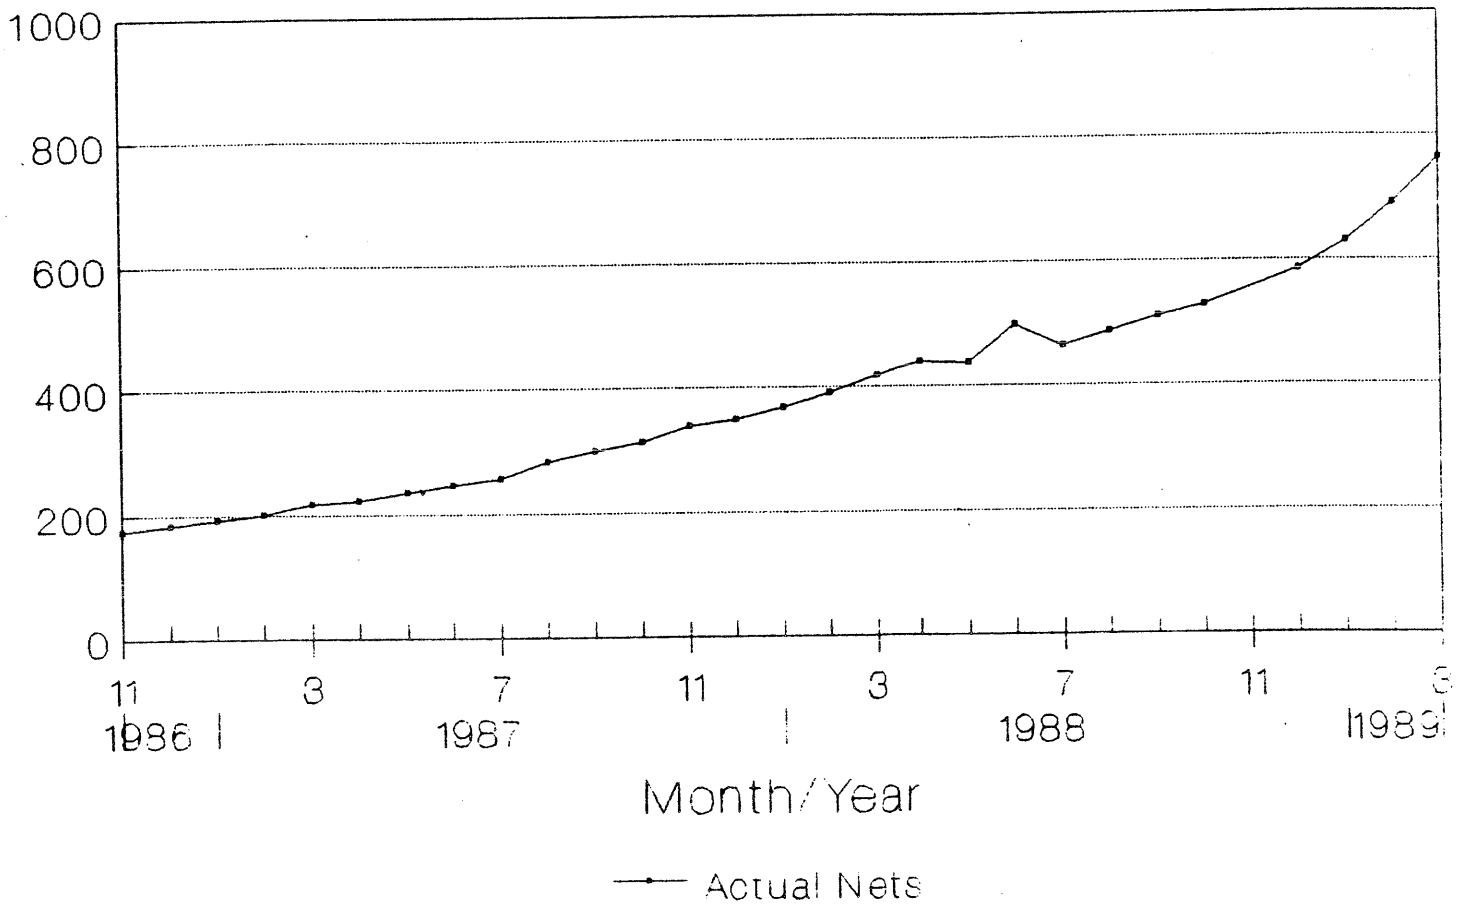

Data Courtesy BBNCC

Advertised Nets
Nov 1986 - Mar 1989

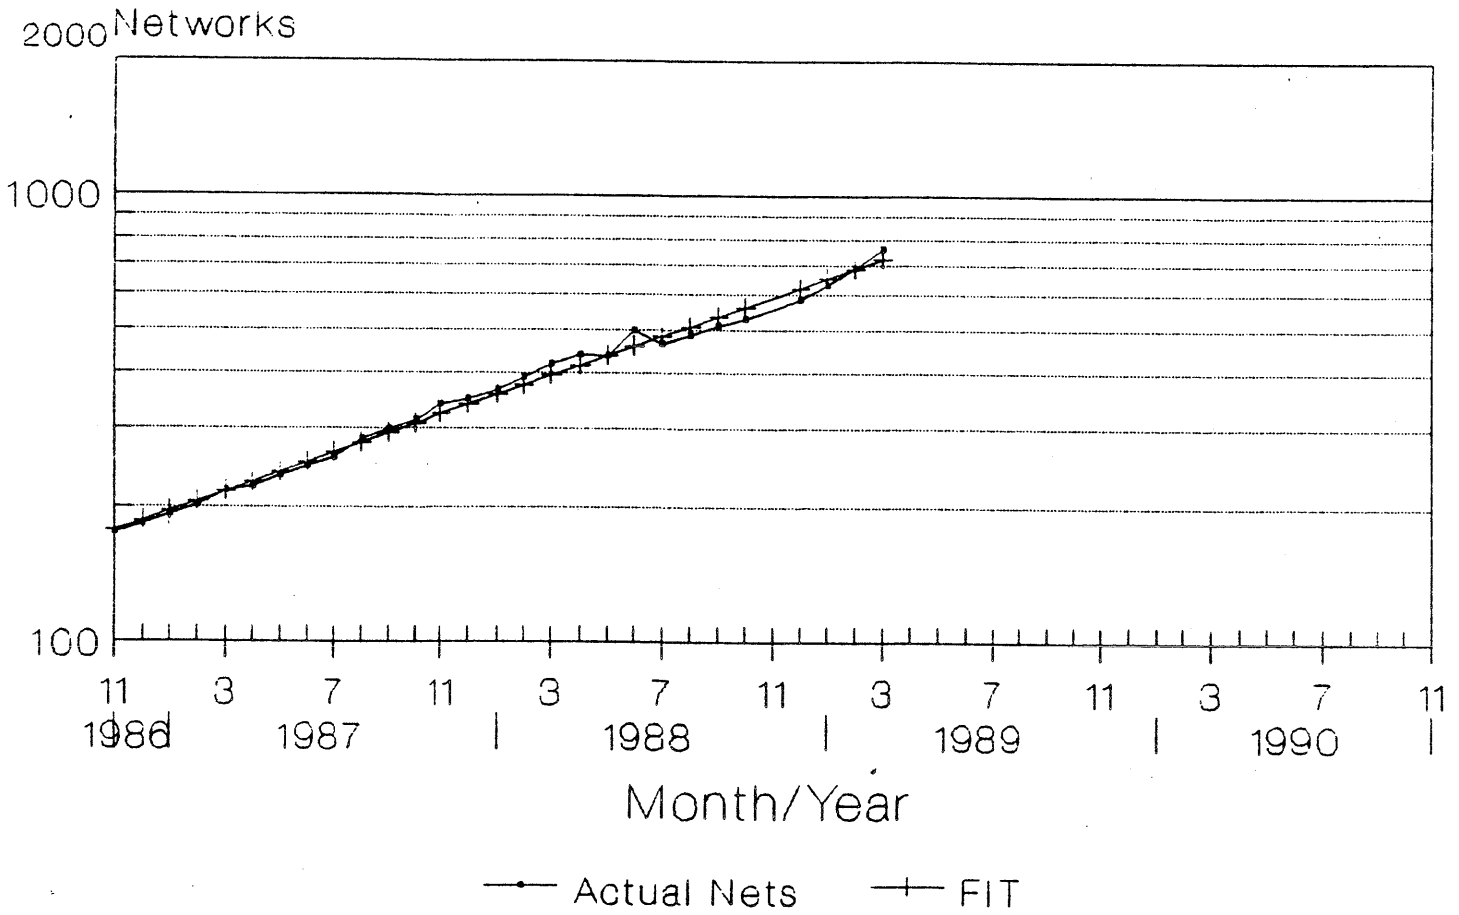

Data Courtesy BBNCC

Advertised Nets
Projected through Nov 1990

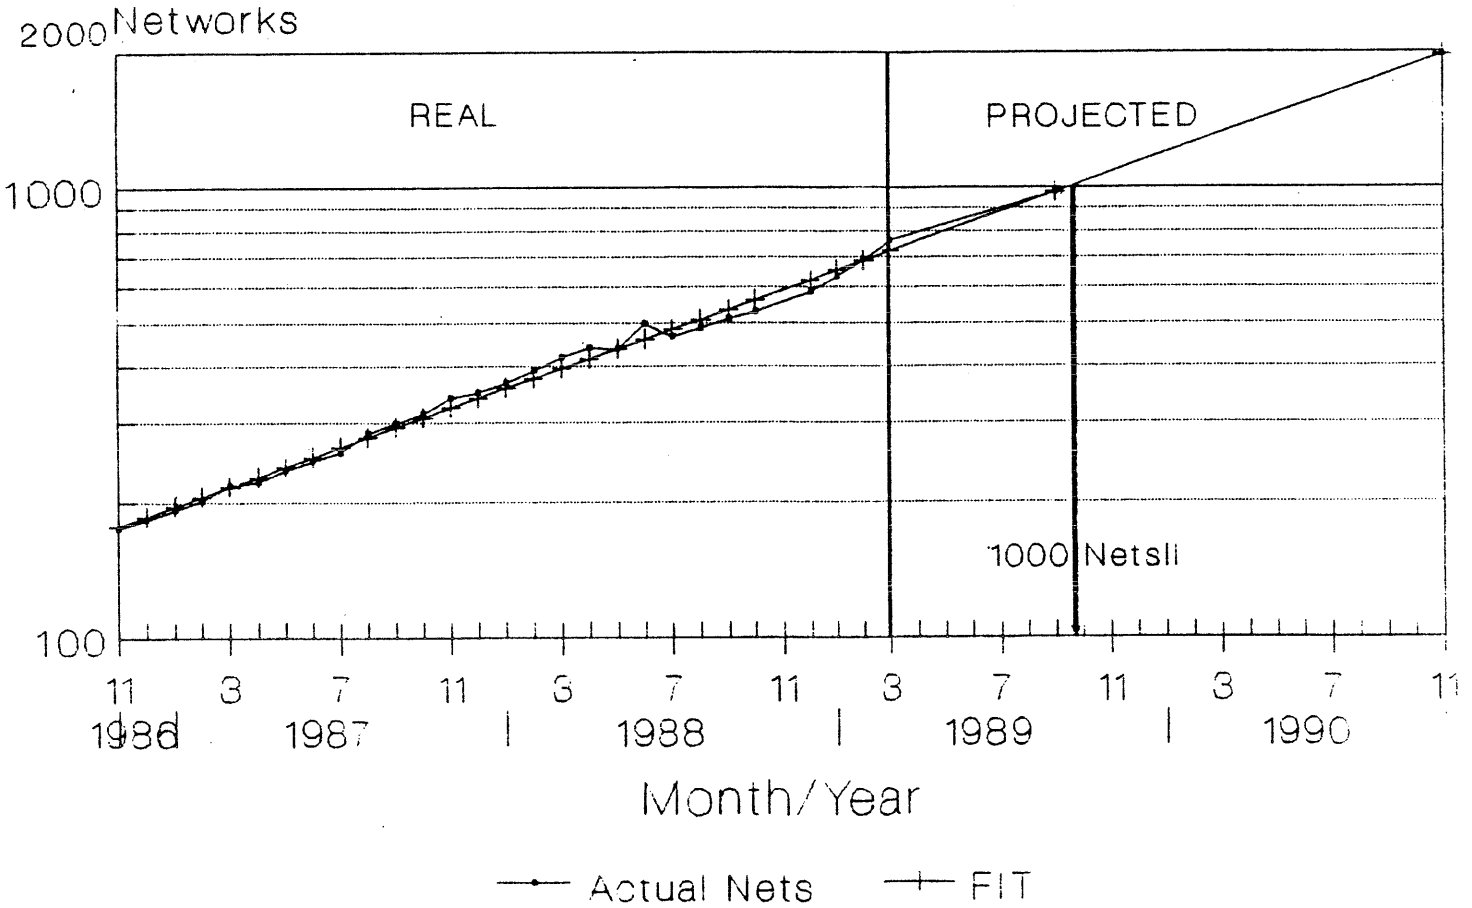

Network Predictions

2000 Dec 1990

8000 Mar 1993

64,000 Jul 1996

256,000 Sep 1998

1,024,000 Nov 2000

Period = 13.3 Month (Approx)
Doubles each period

DOE Energy Science Network
Presented by Tony Hain

ESnet is a service for DOE sponsored researchers providing enhanced communications facilities. Primary service is provided for the 5 Energy Research programs supported by the Office of Energy Research. Its goals are:

Enhance Inter-program communications
Enhance International collaborations
Reduce costs
Increase interconnectivity with other agency networks
Increase performance
Support OSI standards

Current projects include:

ITER program support for the US team
MFEnet II upgrade
Backbone Upgrades to T1
X.25 backbone services for international access
International and Interagency gateways

At the January IETF I reported that we were in the midst of an audit by the DOE/IG. The IG report has been filed recommending that we leave both HEPnet and MFEnet as they are, "all of the users are happy with the services they have". A formal rebuttal has been filed pointing out that current requirements cannot be met with existing services.

ESnet underwent a program review March 6,7 by an outside panel of network specialists. The final report has not been submitted, but preliminary recommendations included:

Centralize management of the DECnet with the IP backbone
Drop IP encapsulation (required to switch X25 and IP on the same lines) plans
Drop T1 multiplexing plans

ESnet has undergone several significant changes since the January meeting. A decision was made to drop X.25 as a native backbone service and allow DECnet phase IV in its place. This allowed procurement rather than development of the backbone routers. There are currently 19 sites identified for T1 service in calender '89. Dual protocol routers are being procured for the T1 backbone.

NASA and NSF backbone interconnects with ESnet are being planned.

Implementation discussions will be held in the near future at MERIT.

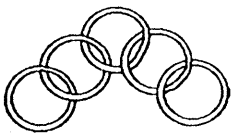

ESnet status report April '89

Tony Hain

ESnet

Why was it formed?

ESnet is intended to provide enhanced computer data communications for the programs funded by the U.S. Department of Energy's Office of Energy Research

Energy Programs Supported

- Applied Math Science
- Basic Energy Science
- Health and Environmental Research
- High-Energy and Nuclear Physics
- Magnetic Fusion Energy

Goals:

- Enhanced interprogram communications
- Enhanced international collaboration
- Cost savings:
  - Shared resources
  - Reduced redundant efforts and costs
  - Centralized network management
- Increased interconnectivity and interoperability with other agencies and networks
- Increased performance and functionality
- Foundation for support of OSI international standards

What are the current priorities?

NETWORKING PROJECTS

Status and current Issues

1. ITER
2. ESnet/MFEnet II (IP access)
3. Backbone T1 upgrades
4. X.25 backbone
5. ESnet X.25 access
6. European networking
7. Japanese networking
8. Interagency gateways

ESnet

How is it managed?

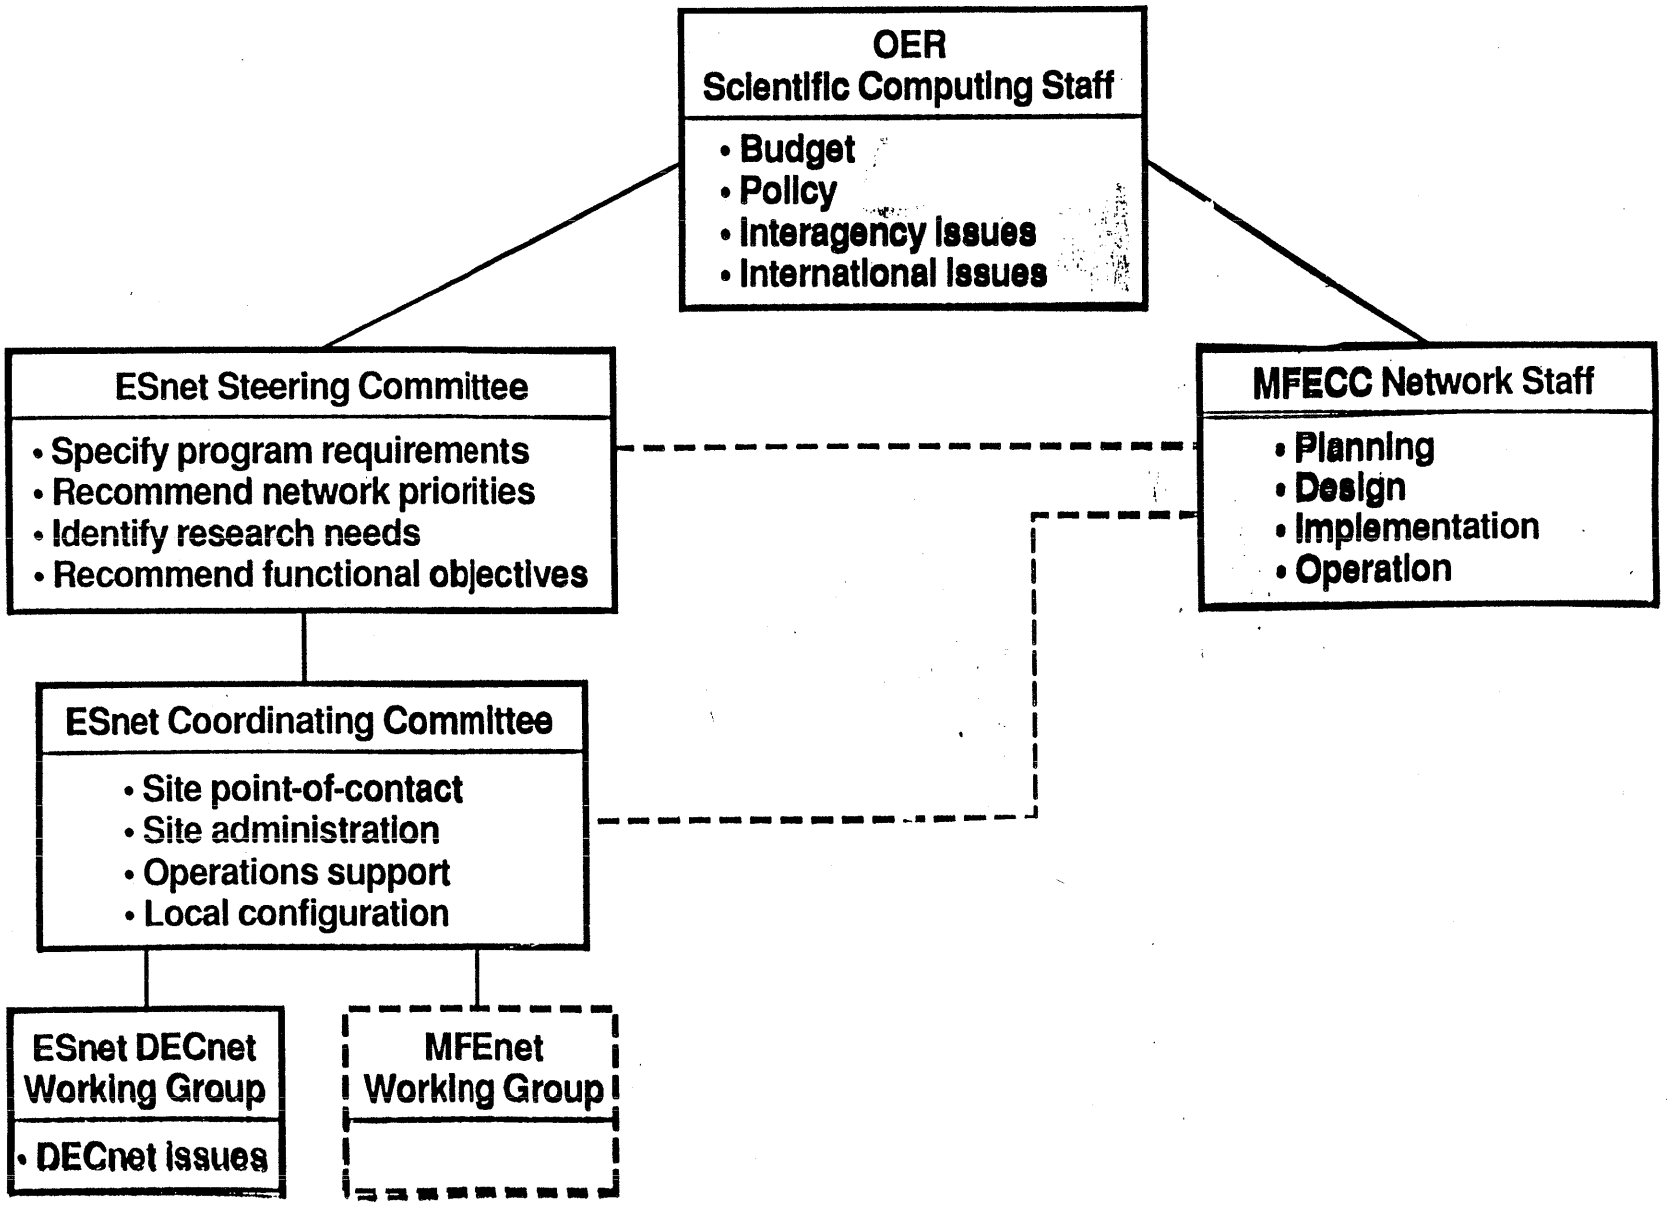

Networking and Engineering Groups

James F. Leighton — Manager, Networking and Engineering
Secretary — Angela Thompson

Network Host Systems Group Barry Howard	Data Communication Systems Group Paul Lund	Engineering Group Tim Voss
Host Network Access VAX System File Transport Network Servers Office Automation	Network Node Software Foreign Net Gateways Network Accnt. & Stats	Custom Hardware Network Installations Network Control
Bob Aiken Neal Mackanic Jackie Marr Jean Wolitzer	Mike Collins Bob Cooley Harry Massaro Rebecca Nitzan Lee Tennant	Steve Hunter Cliff Cordova Tony Hain Rick Schnetz Loyd Davis Linda Doyle
Others		
DEC Contract Phil Beck — VAX System Administrator Bill Horton — VAX System Analyst Meryl Anderson — Network-wide DEC Maint. TID Contract Lila Abrahamson — Documentation Consulting Open — Network support liaison		Tom Whitney Bill Callovini Jim Gagliardi Jon Hammond Jim Morton

ESnet What are plans for foreign connections?

- Japan (Nagoya, JAERI)
  - — Two low speed (9.6 Kbps) PDN "links" operational
  - — Consideration of 64 Kbps ESnet satellite link underway

- Switzerland (CERN)
  - — 64 Kbps X.25 satellite link operational
  - — Conversion to fiber-optic link underway

- Germany (Garching)
  - — Installation of 64 Kbps X.25 satellite link in progress
    - Will link to switching HUB for Germany
  - — Access being planned for ITER International Design Center

ESnet

APRIL '89 STATUS

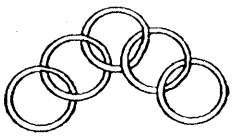

What ever happened to the IG report?

- Recommendation: Maintain existing networks as they are.
  - HEPnet users are satisfied with their 9.6k lines.
  - MFEnet users are very happy, leave them alone (have them time share addresses)
- Response: Formal rebuttal filed.
  - Current requirements cannot be met with existing networks.

ESnet

APRIL '89 STATUS

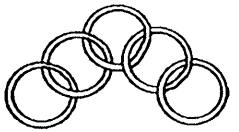

Project Review

- Conducted March 6 & 7, '89 by outside panel.
- Draft Recommendation Highlights:
  - Centralize management of DECnet with IP management.
  - Drop IP encapsulation plans.
  - Drop Multiplexing.

ESnet

APRIL '89 STATUS

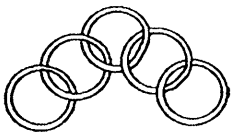

Current Activities

- 19 sites identified for T1 CY'89, expect first lines July.
- Backbone design constraints were relaxed, DECnet allowed as native protocol.
- Dual protocol routers being procured for use on the backbone. ~ Aug. 89
- Actively working several procurements (Routers, CSU/DSU's, lines...)

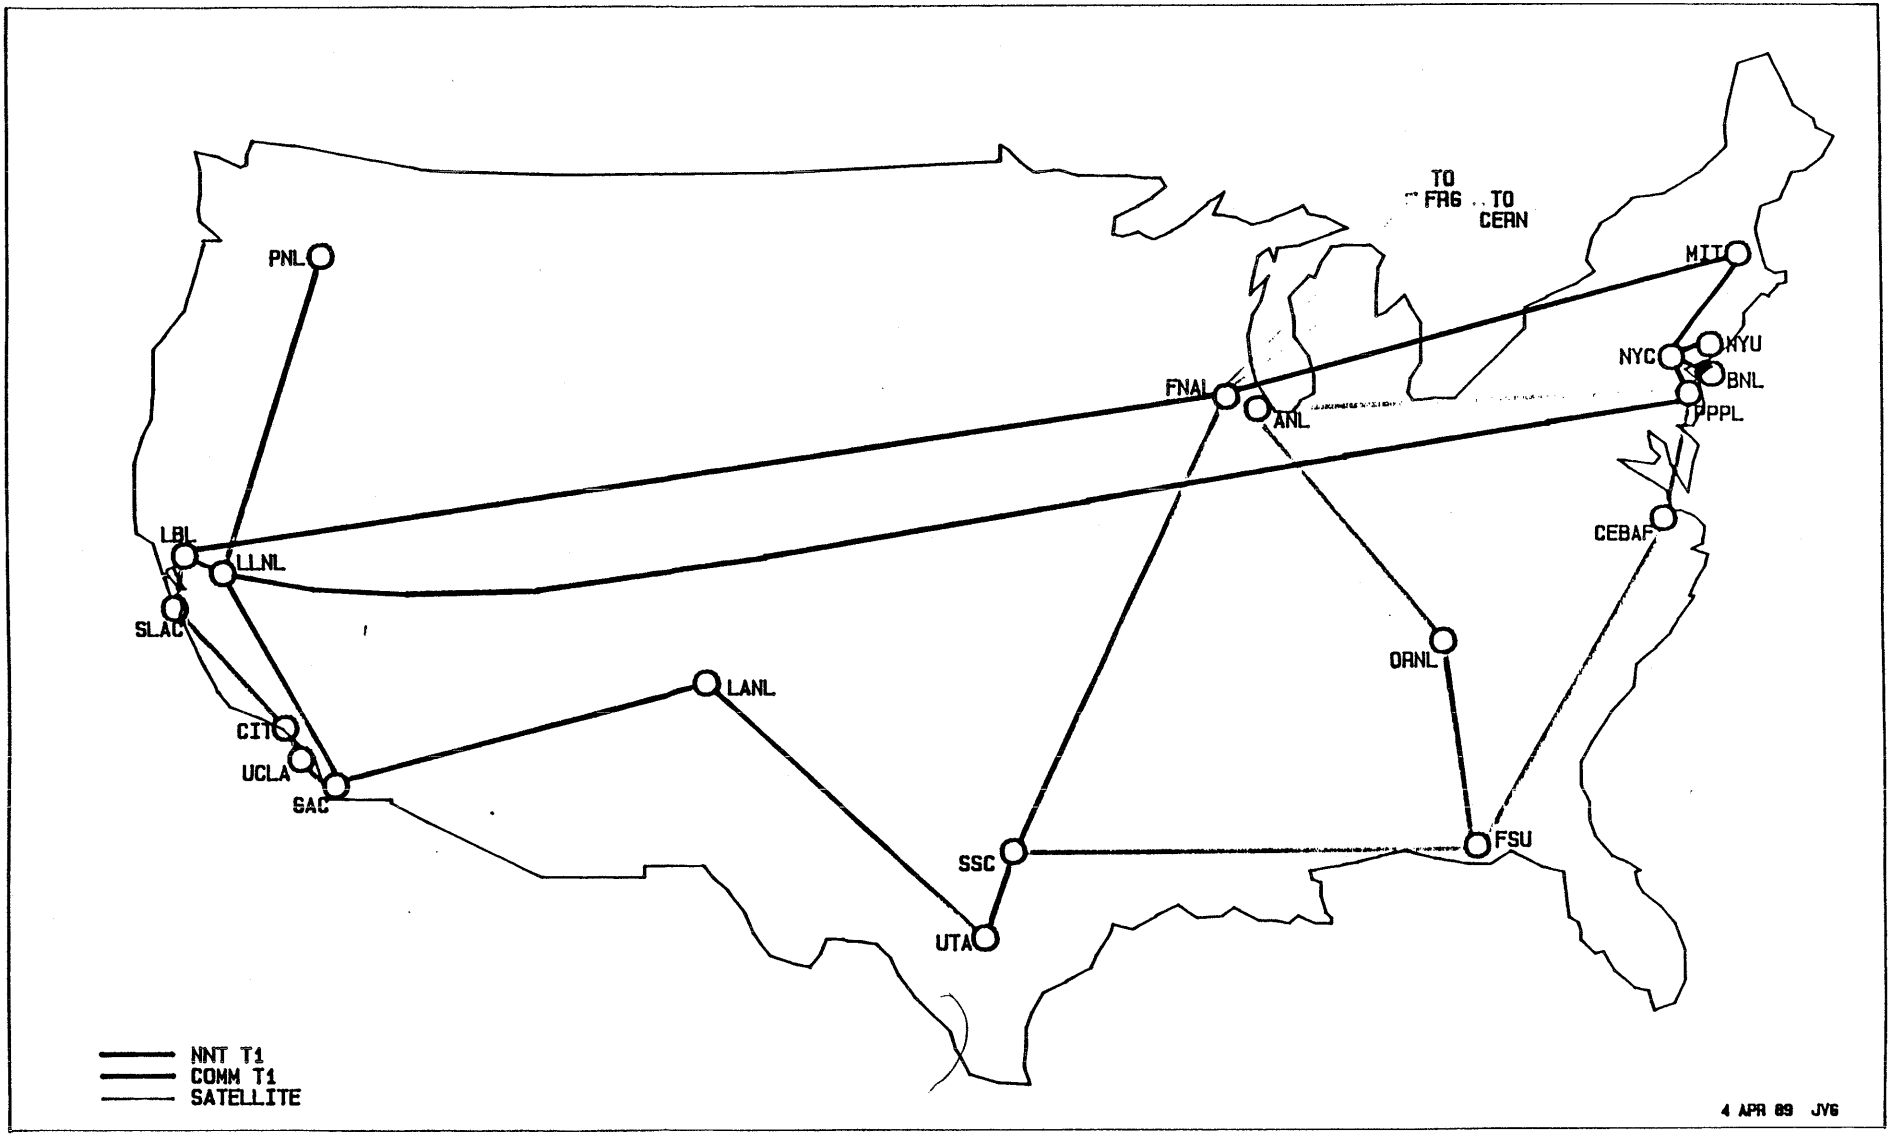
TO FRG
TO CERN
PNL
MIT
NYC
NYU
BNL
PPPL
FNAL
ANL
LBL
LLNL
SLAC
CEBAF
ORNL
LANL
CIT
UCLA
SAC
SSC
FSU
UTA
NNT T1
COMM T1
SATELLITE
4 APR 89 JV6

An Interim Routing Architecture
Presented by Russ Mundy

The purpose of the Interconnectivity Working Group (IWG) is to develop an interim Inter-Autonomous-System routing approach which will be available prior to the results of the work of Open Routing Working Group (ORWG). This is necessary since the current Exterior Gateway Protocol (EGP) is inadequate and implementation of new ORWG solution will not be available for some time.

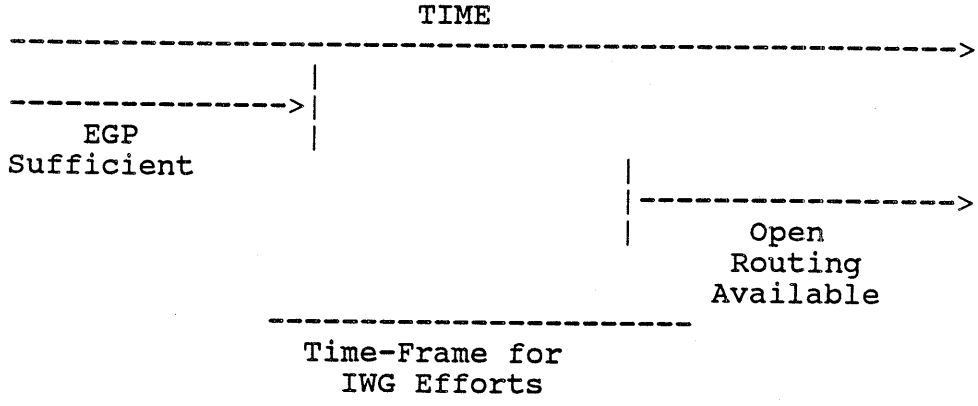

There have been several previous meetings of the IWG and significant mailing list activity by the working group. Prior to the Austin IETF meeting, the IWG participants had primarily been individuals involved with managing National and/or Regional networks. During the Austin IETF, vendor representatives were asked to provide their views on the IWG approach. During the open meeting, the vendor representatives made general comments but were not negative about the IWG approach. Subsequent to the meeting, each of the vendor representatives contacted the IWG chairman and expressed concern about the commercial viability of making frequent protocol changes in their products which was being inferred by the timing of the IWG and the ORWG efforts. Generally, their constraints include interoperability of the new protocol with the installed operational base and their cycle for major software releases (usually only one per year for major protocol changes).

As a result of the vendor inputs and the previous work of the IWG, the chairman drafted a Midterm Inter-Autonomous Routing Architecture (MIRA) that was distributed to the IWG mailing list prior to the Cocoa Beach IETF meeting. This draft paper became the basis for much of the discussion at this IETF meeting.

The ORWG participated in the first portion of the IWG meeting. This joint meeting provided ORWG participants the

opportunity to review and comment on the IWG efforts. The general impression is that MIRA is a different solution to many of the problems being addressed by the ORWG but there are some areas where the MIRA approach could provide useful lessons for the ORWG.

The approach described in the MIRA paper provides improved prevention of loops, better fault detection and better support for the extensive connectivity between Networks and Autonomous-Systems. MIRA also seeks to minimize implementation difficulties for both vendors and network providers/operators.

MIRA achieves many of these features by separation of functions. For example, route maintenance and packet forwarding is currently bundled in a single gateway. MIRA separates these functions into route server and border gateway. The separation also makes it feasible to provide a full Autonomous-System path that can be used to determine routes for packets. The route servers communicate with border gateways within its Autonomous-Systems as well as with route servers in other Autonomous-Systems.

There are several current problems/difficulties that MIRA has not yet solved. Some of these include: the method for the initial bootstrap of a system is not clear; a method of providing reliable communications between route servers needs to be defined; the method of achieving consistency between the route servers within an Autonomous-System is not defined; and the method(s) route servers use to choose between routes is not defined.

In addition, some of the ideas presented by the MIRA are being implemented by several participants of the IWG to gain experimental experience. Preliminary information from these implementations should be available for the July IETF meeting.

Intercommectivity
Working Group

Guy 'Almes
(Almes @ rice.edu)
Chairman

Russ Mundy
(Presenter)

Mundy @ beast .ddn. mil

Purpose:

Interim
Inter-AS
Routing
Approach

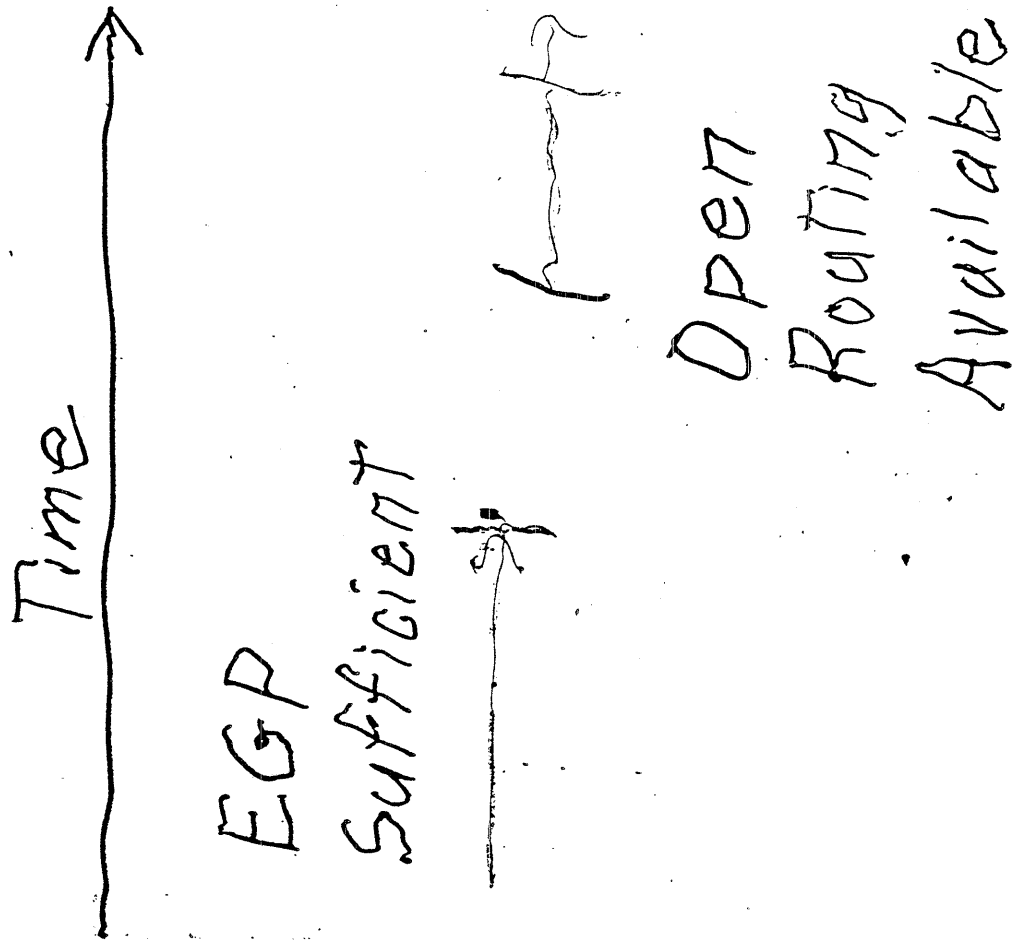
Time
EGP
Sufficient
Open
Routing
Available

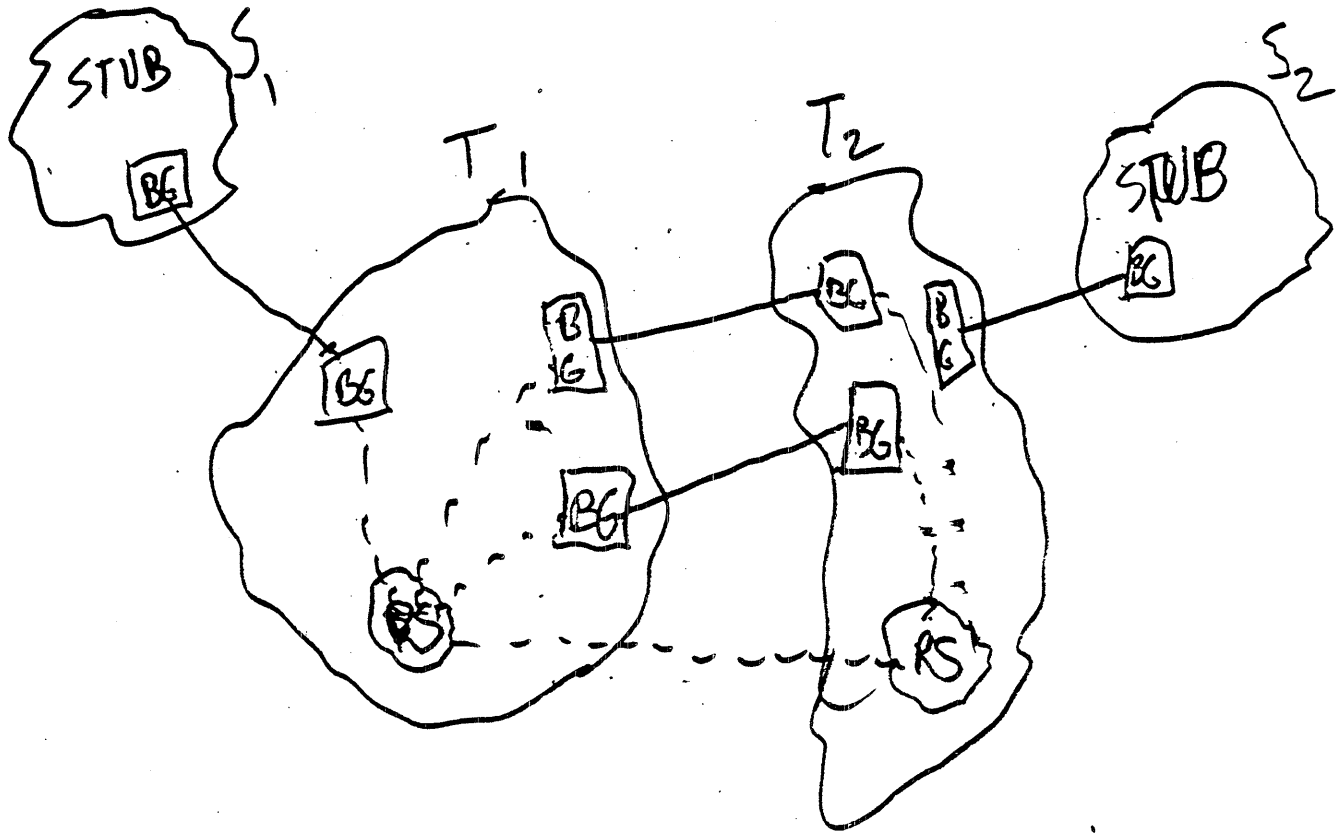
STUB
S1
T1
T2
S2
STUB
BG
RS

Architectural Changes to the NSFNET
Presented by Elise Gerich

In the nine months that Merit has managed and operated the NSFNET backbone, traffic on the network has grown substanitially as havethe number of networks announced by the backbone. In the lastmonth, March 1989, we saw an approximately 30% increase in thetraffic on the backbone, and have over 450 nets connected via mid-level networks.

Recognizing that the traffic was probably not going to decrease, and wanting to continue to offer good service to our users, MCI, IBM and Merit began to meet to plan an expansion of bandwidth for the backbone. As guidelines to our planning sessions, we established five primary objectives in redesigning the backbone.

These objectives were:

1) Eliminate the single tail circuits
2) Provide a network diameter of 3 or less
3) Provide a full T1 rate end-to-end bandwidth
4) Be financially feasible
5) Take initial steps toward DRS (Digital Reconfiguration Services)

The first objective, eliminating the tail circuits, was identified as the most pressing need to address, but we agreed it was desirable to fullfill all of the objectives. With these goals in mind, we came up with a new physical topology for the NSFNET backbone.

Instead of 14 circuits connecting the 13 nodes, six of which were single tail circuits, the new topology consists of 19 circuits with three circuits terminating at each node except for MIDnet which has two circuits coming into it. This redesigned topology meets all of our objectives.

The single tail circuits are eliminated.

The network diameter is 3.

Each node has circuits of full T1 rate bandwidth instead of the sub-T1 rate logical links that are currently provided on thebackbone.

The redesign fits within our budget.

And finally, this topology moves us toward DRS (Digital Reconfiguration Services). As you can see from the accompanying slide, we are moving toward an architecture where the circuits from each node feed into an MCI cloud. Within that cloud,

circuits may be allocated dynamically. We see this redesign as the first step toward implementing DRS.

Along with this redesign of the topology, MCI has been upgrading some of its digital radio circuits with fiber. As we put the new physical network in place, almost all of the circuits will be fiber, except for those circuits in the northwest and a few local loops.

The implementation phases of the new design are already under way. The circuits are ordered, additional hardware/software is being installed on the NSSs, and preparation by the sites is in process. There are four primary phases in the migration to the new topology.

Phase A which installs four additional circuits to the current 14 circuit topology is tentatively scheduled to be in place by the beginning of June '89. This phase addresses the need to eliminate the single tail circuits.

The following three phases, B thru D, consist of installing and disconnecting pairs of circuits. One of our objectives in order to stay within budget is to eliminate the need for redundant circuit terminations at the local site. For instance, SDSC currently has four circuits terminating at its site, but only two of those circuits will exist in the future backbone. So instead of installing an additional circuit at the local telco, we plan to roll one circuit out and another in. So at no time during the migration will any node exceed the final three local connections needed nor the existing number of connections. Furthermore, no nodes should become isolated during the installation.

Phases B thru D are expected to be in place in July '89. With the completion of those phases, we will have attained our objectives.

In addition to redesigning the backbone technology, Merit has been addressing the need for direct peer network connections. Toward this end, we have met with Milo Medin, of NASA, and Mike St. Johns,of DCA, to design direct connections between the NSFNET backbone and NSN and Milnet.

In the case of Milnet, we are in the process of deploying a configuration which we call a Split E-PSP (exterior packet switching processor). In this configuration,an NSS will have a RT colocated on an Ethernet with a Butterfly-Mailbridge. This Split E-PSP will establish an EGP session with the Mailbridge. Then, the Split E-PSP has a serial connection to the rest of the NSS, and this provides a direct connection between the backbone and the DDN.

The Split E-PSP configuration will be deployed from both NSS 13 (BARRNet) and NSS 9 (SURAnet), giving the NSFNET two connections to the DDN; one on the west coast and one on the east coast. We anticipate that at least one ofthe Mailbridges will be connected to the backbone by the end of May '89. This date hinges on the availablity of T1 circuits.

This Split E-PSP configuration at NSS 13 will also provide a direct connection to the NSN (NASA Science Net). Not only will the NSS EGP peer with the mailbridge, but also with a NSN-router. Milo Medin has also been negotiating with Jack Hahn at SURAnet to establish a SURAnet/NSN connection at NSS 9. If all goes well, NSFNET should have an east and west direct connection to both NSN and the Milnet.

The NSFNET backbone announces approximately 60 international networks. Most of these announcements are received via a standard NSS configuration, but we have implemented an EGP session via a serial link with CNUSC in Montpellier, France. This configuration is described on one of the accompanying slides. Currently the link between NSS 10(CNSF/NYSERnet) and CNUSC is a 56 Kbps Satellite link, but we are planning to upgrade this to T1 via TAT8.

Merit is pleased to announce the above architectural changes to the NSFNET backbone. However, this is just a beginning. We have recently deployed and made available SNMP. We continue to explore further changes and enhancements to NSFNET, such as ISO CLNP support, T3 upgrade, and possibly X.25 support. These are just a few of the next steps for NSFNET.

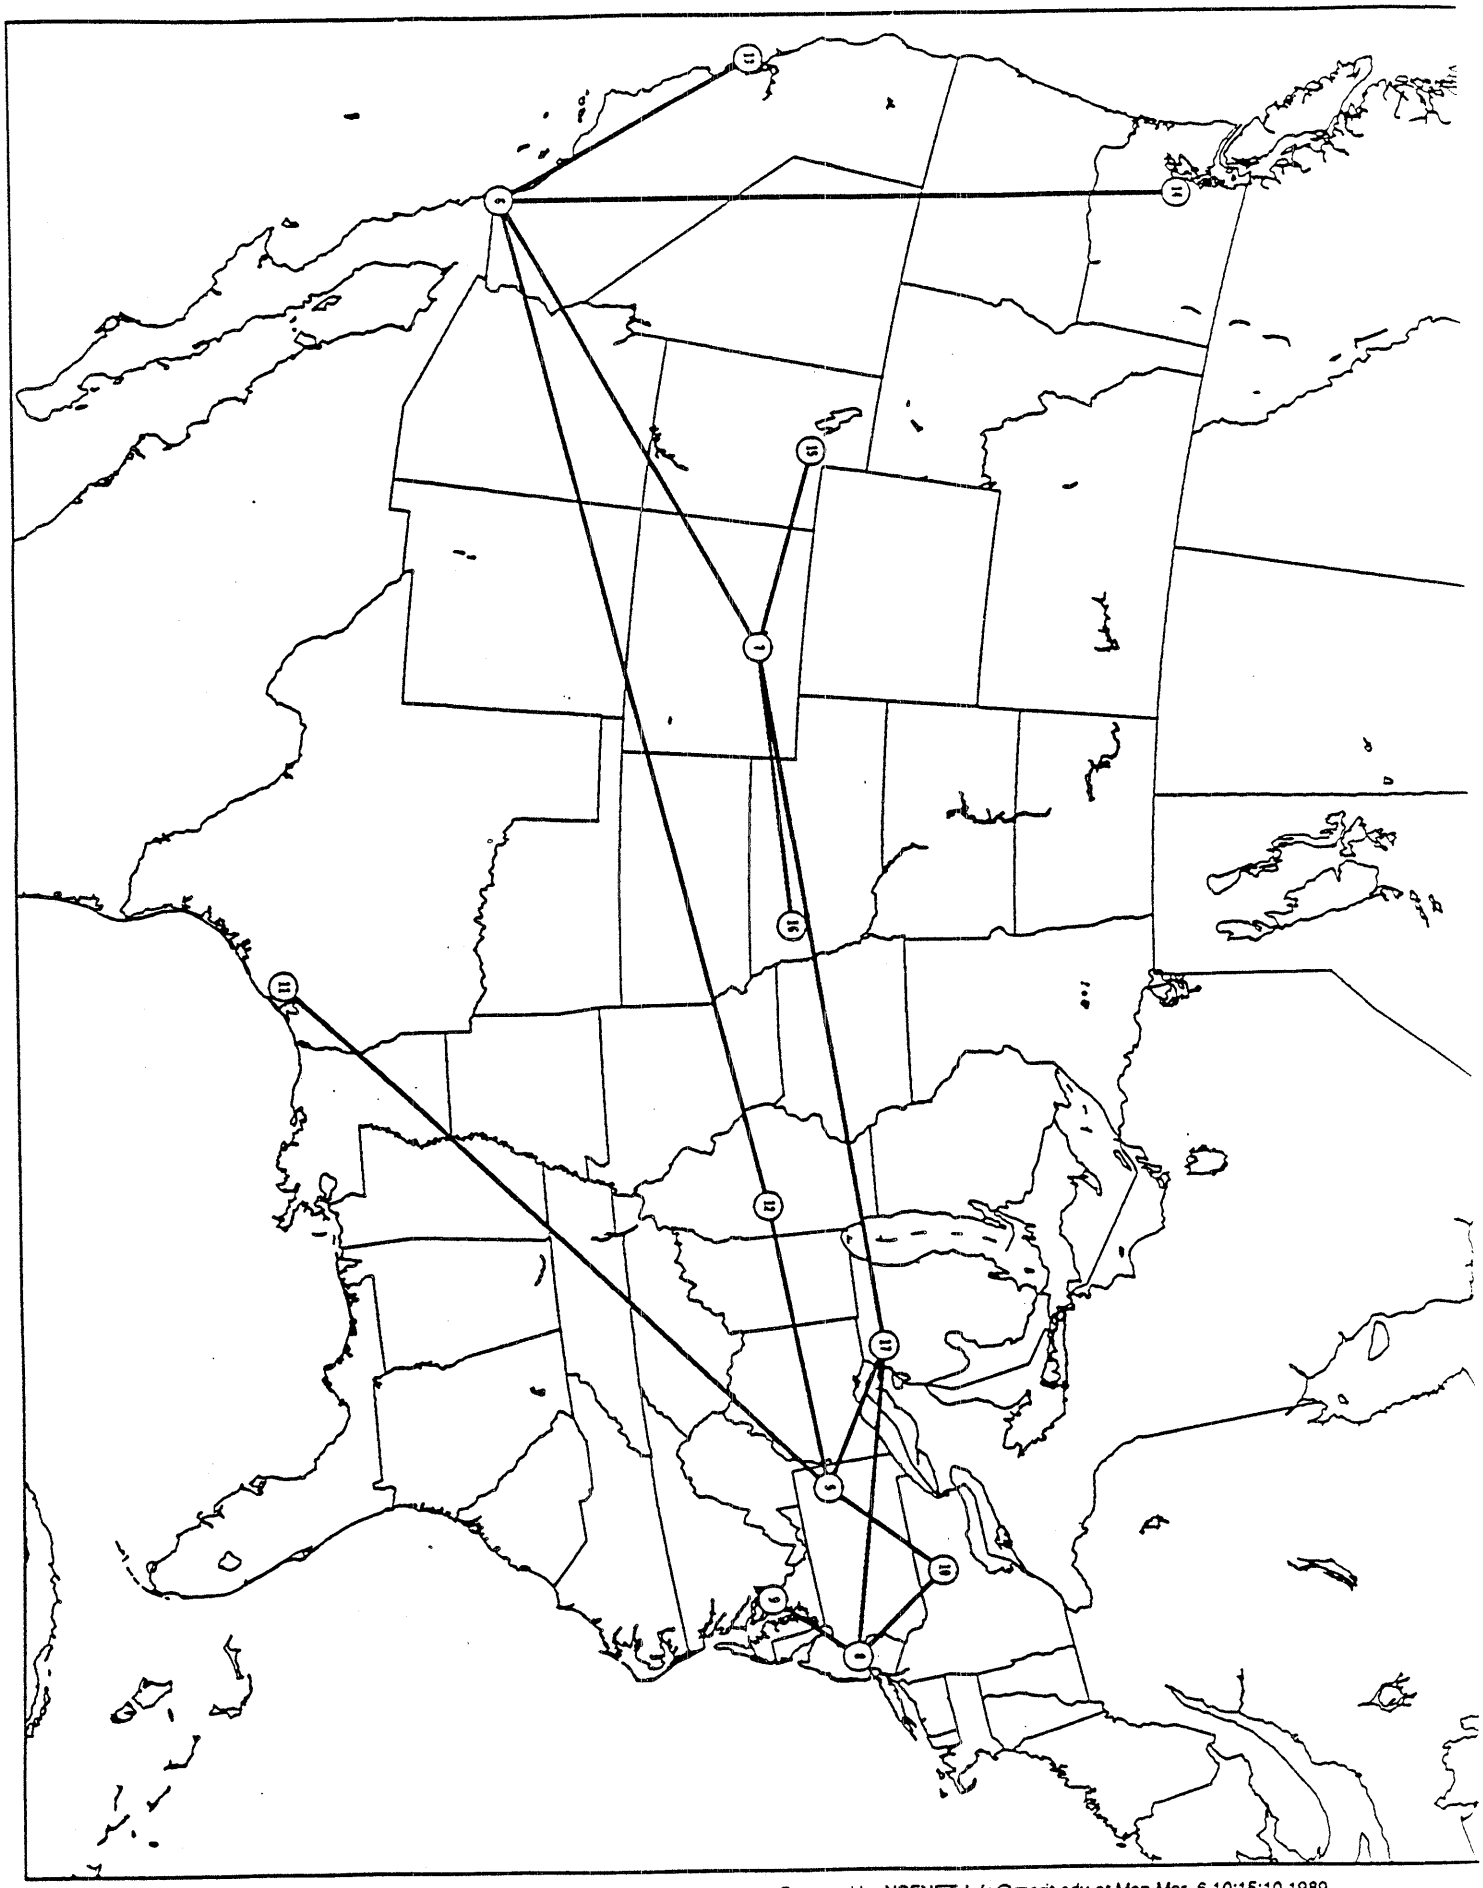

Physical NSFNET Topology
Circles contain NSS number

Prepared by NSFNET-Info@merit.edu at Mon Mar 6 10:15:10 1989
netmap-1.5 program by Brian Reid, map data from World Data Bank II
Lambert Conformal Projection (44°N,33°N), Map center: (40°N, 96° 30'W)
Image resolution 300/in., stroke limit 1 pixels

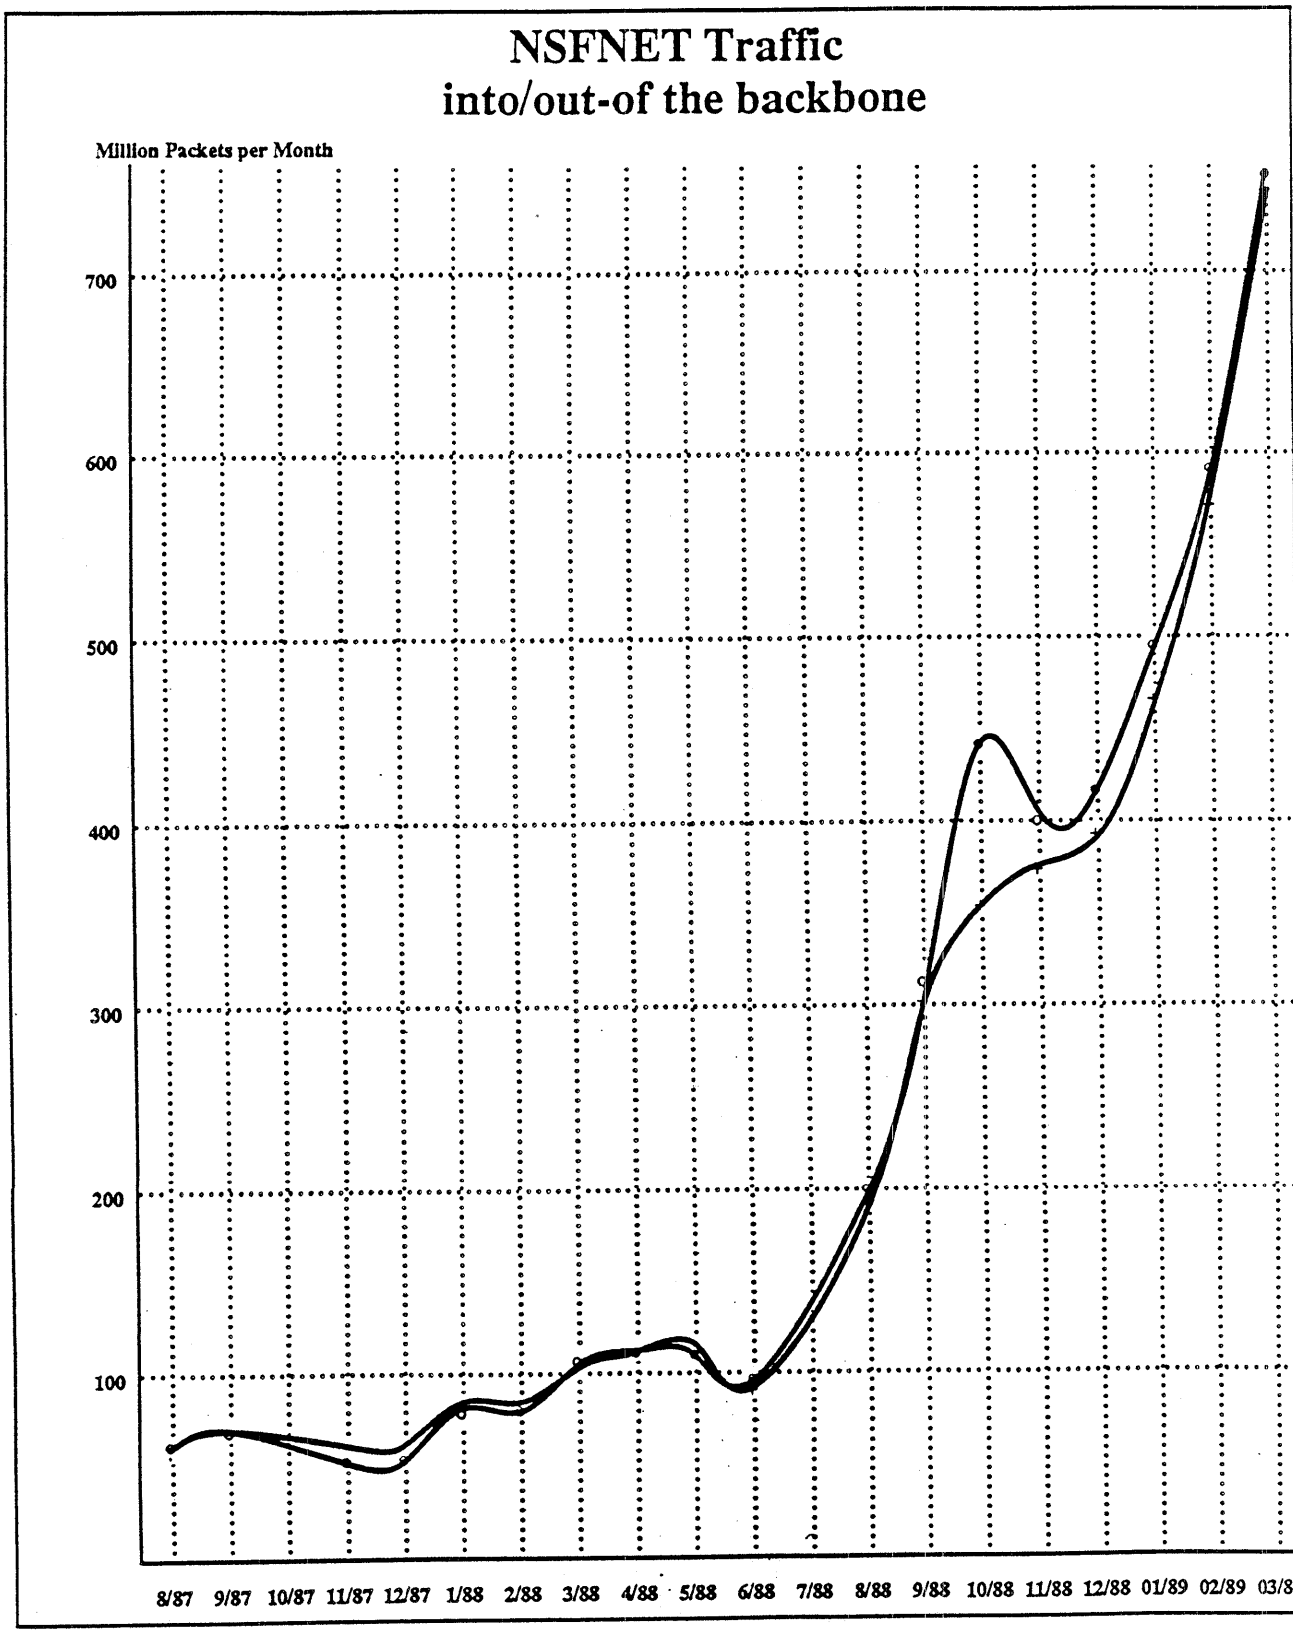
NSFNET Traffic
into/out-of the backbone
Million Packets per Month
700
600
500
400
300
200
100
8/87
9/87
10/87
11/87
12/87
1/88
2/88
3/88
4/88
5/88
6/88
7/88
8/88
9/88
10/88
11/88
12/88
01/89
02/89
03/89

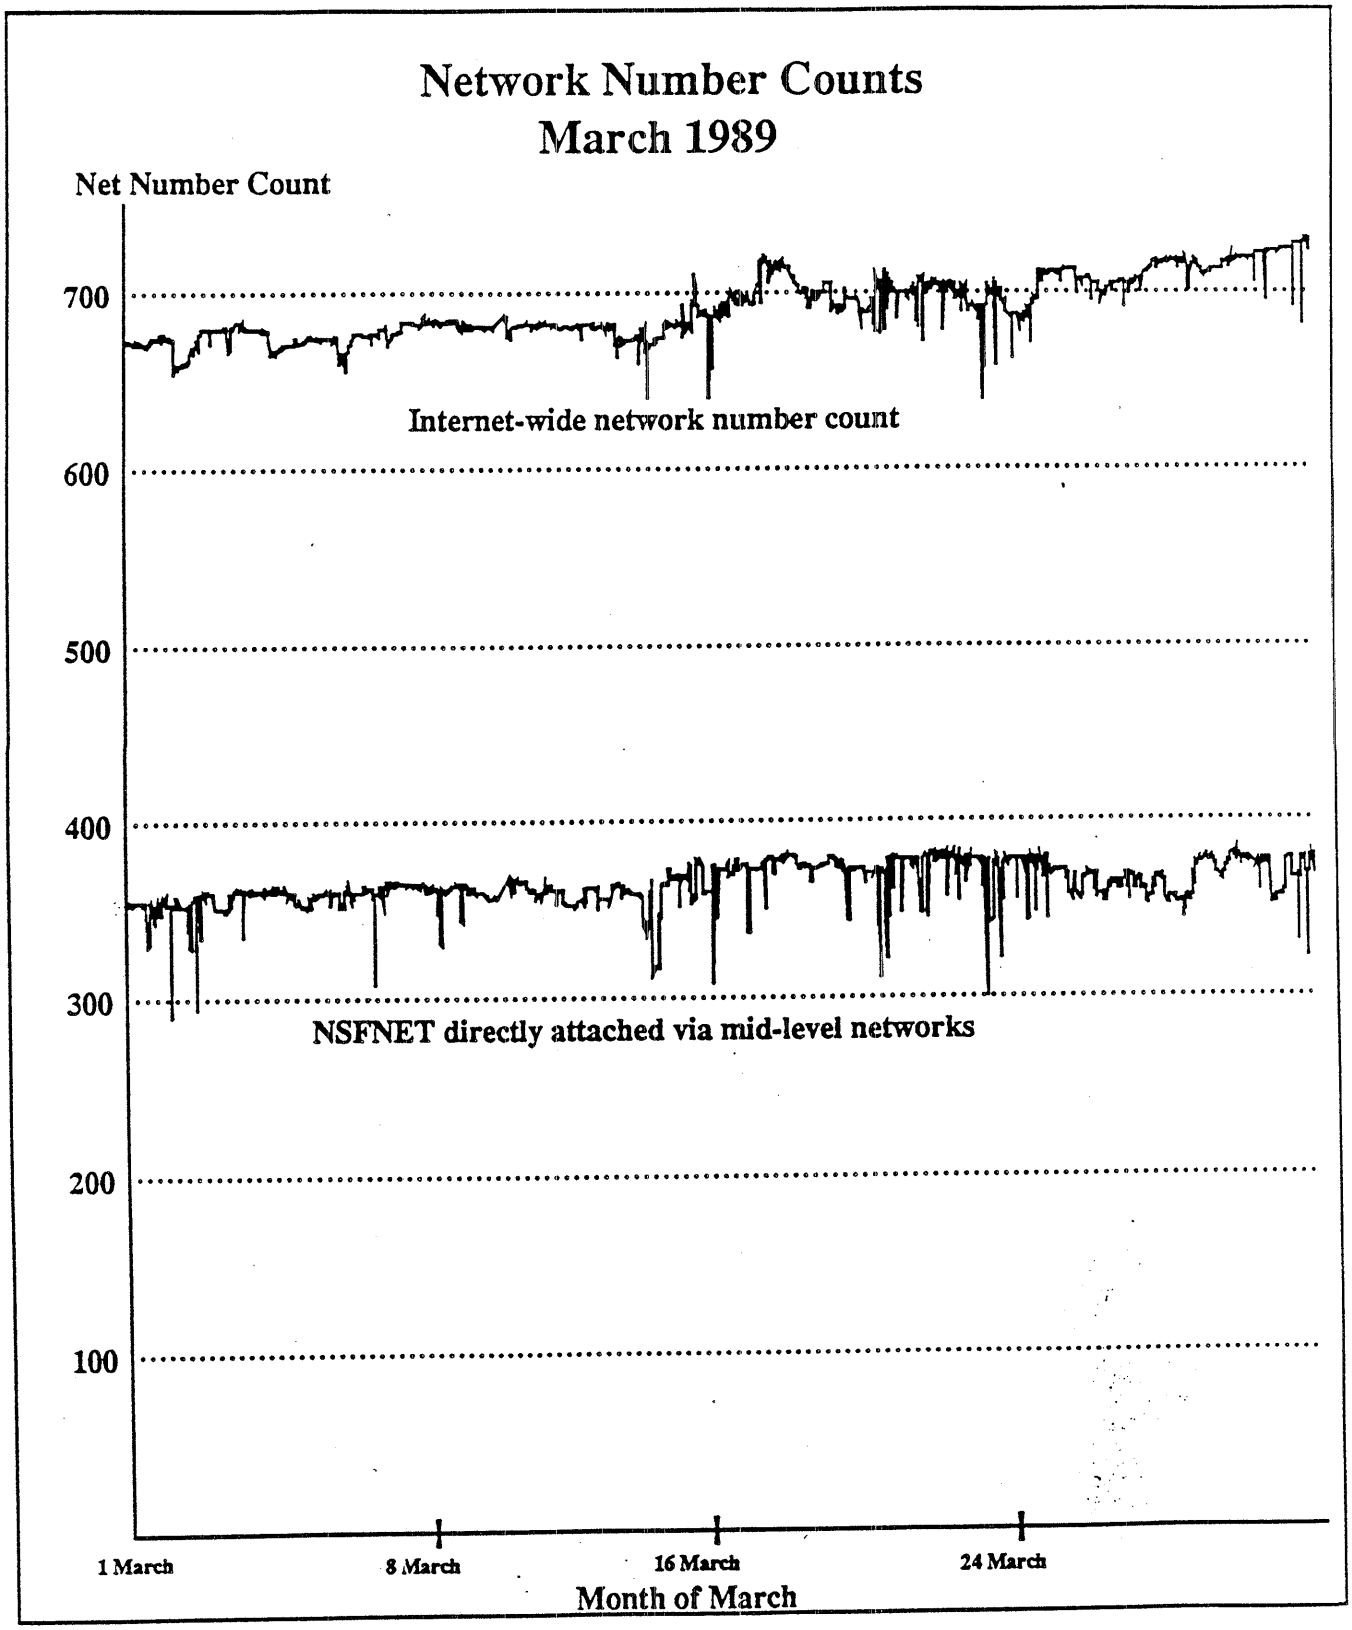
Network Number Counts
March 1989
Net Number Count
700
600
500
400
300
200
100
Internet-wide network number count
NSFNET directly attached via mid-level networks
1 March
8 March
16 March
24 March
Month of March

Primary Objectives

- Eliminate the tail circuits (Redundancy)
- Maintain network diameter of 3 or less
- Provide each node with bandwidth of full T1
- Be financially feasible
- Take initial steps toward DRS

1/18/89 epg

NSFNET DRS Application

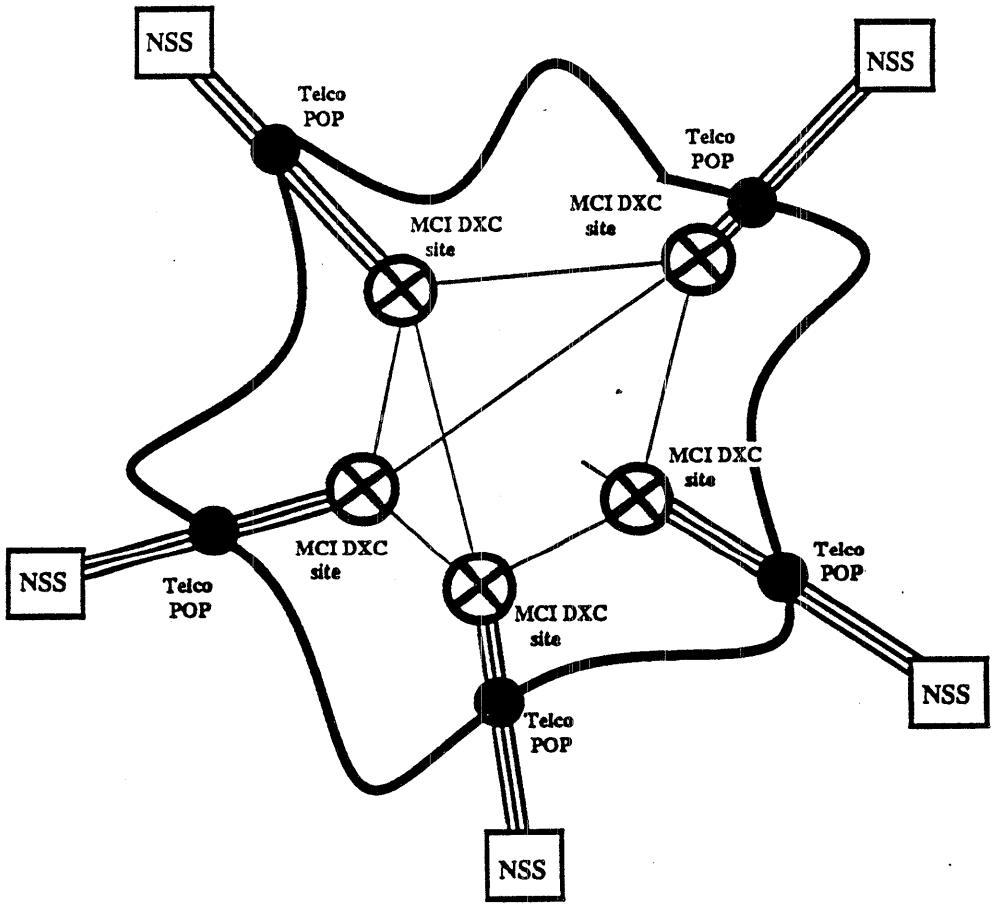

6 January 1989, HWB

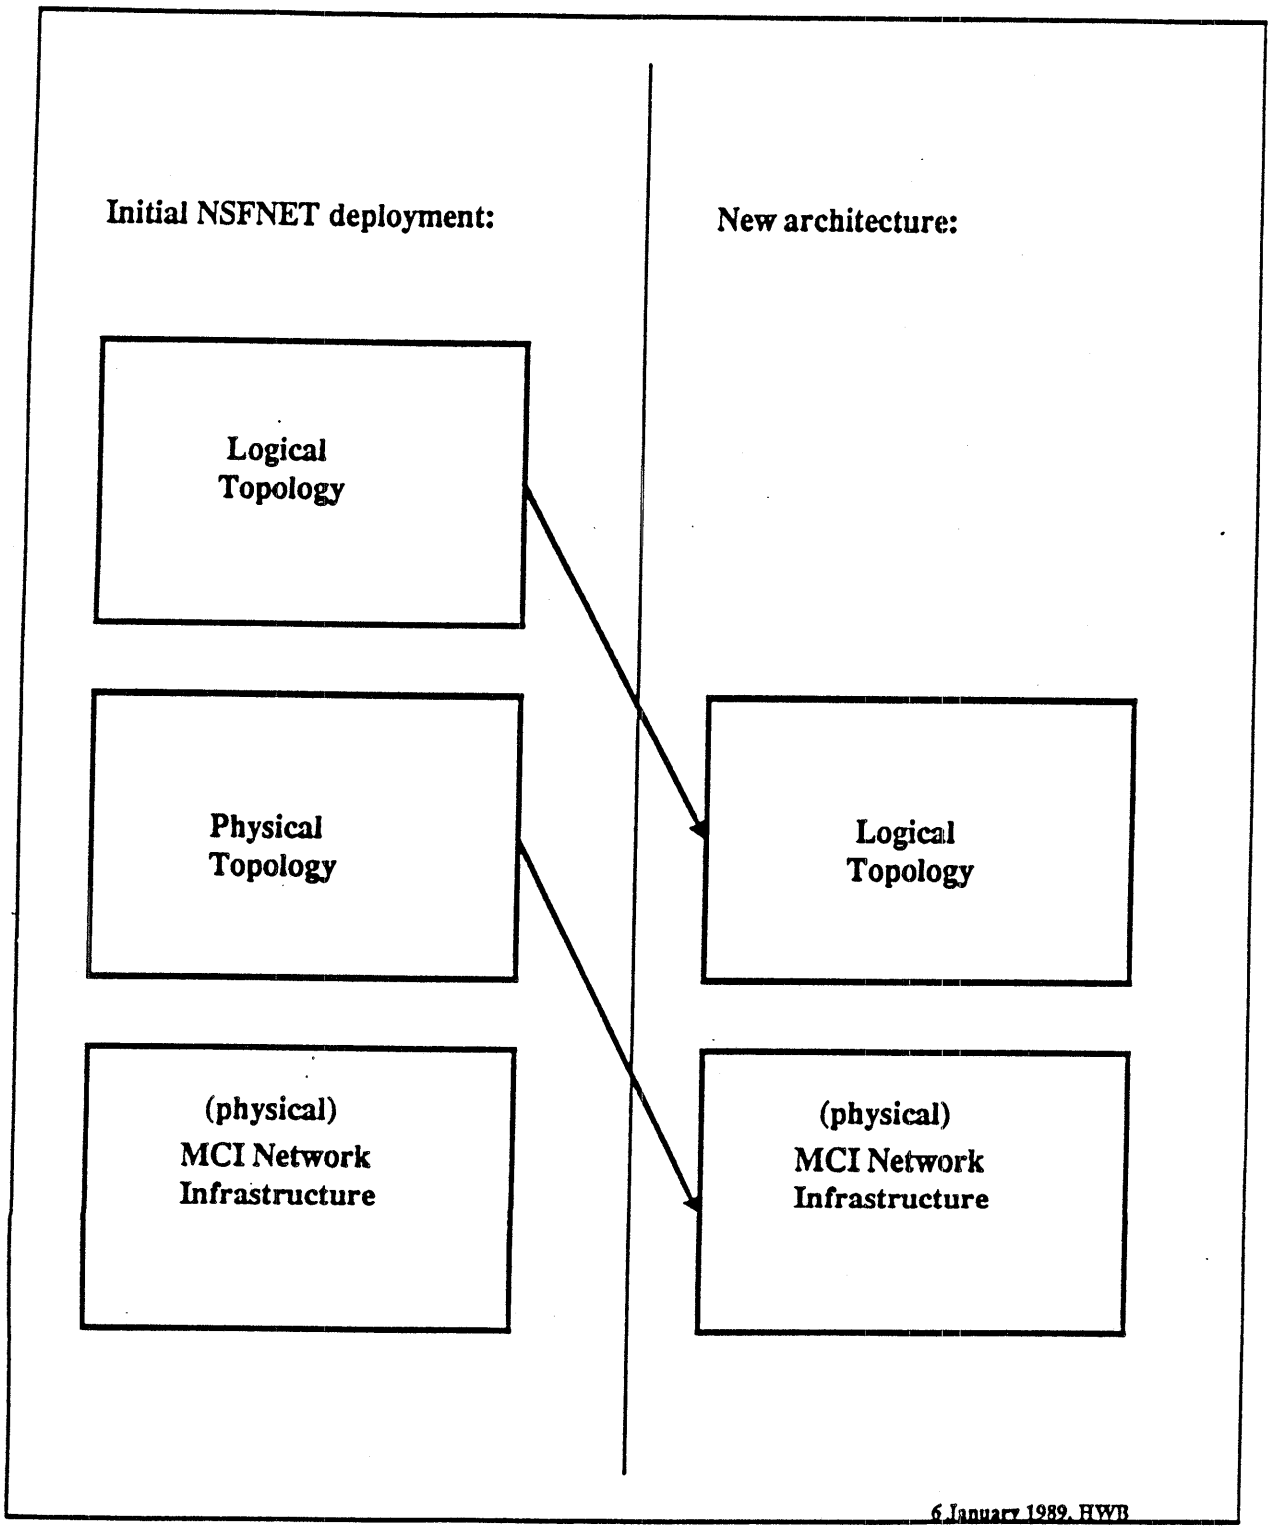
Initial NSFNET deployment:
New architecture:
Logical Topology
Physical Topology
(physical) MCI Network Infrastructure
Logical Topology
(physical) MCI Network Infrastructure
6 January 1989, HWB

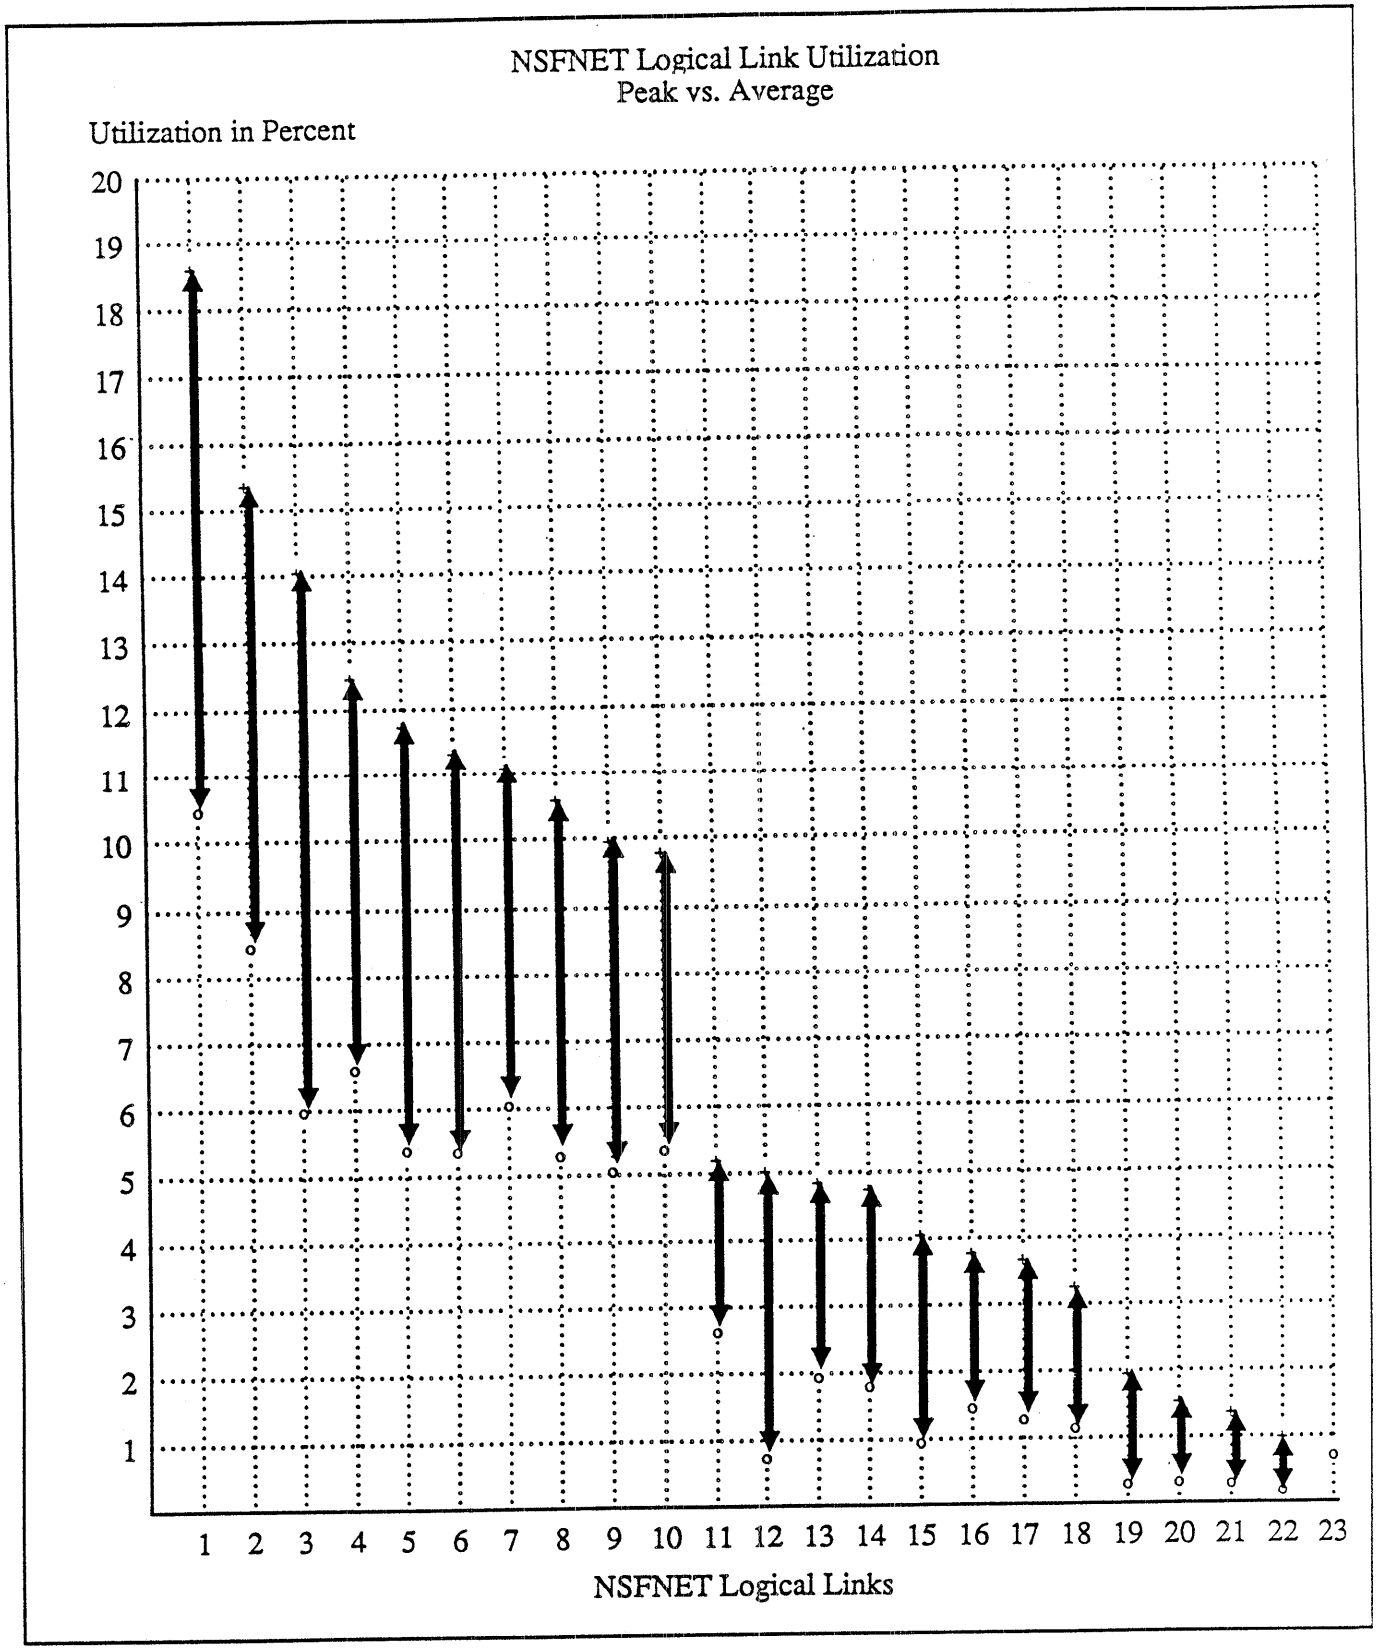
NSFNET Logical Link Utilization
Peak vs. Average
Utilization in Percent
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
NSFNET Logical Links

NATIONAL SCIENCE FOUNDATION

EXISTING T-1 DATA NETWORK

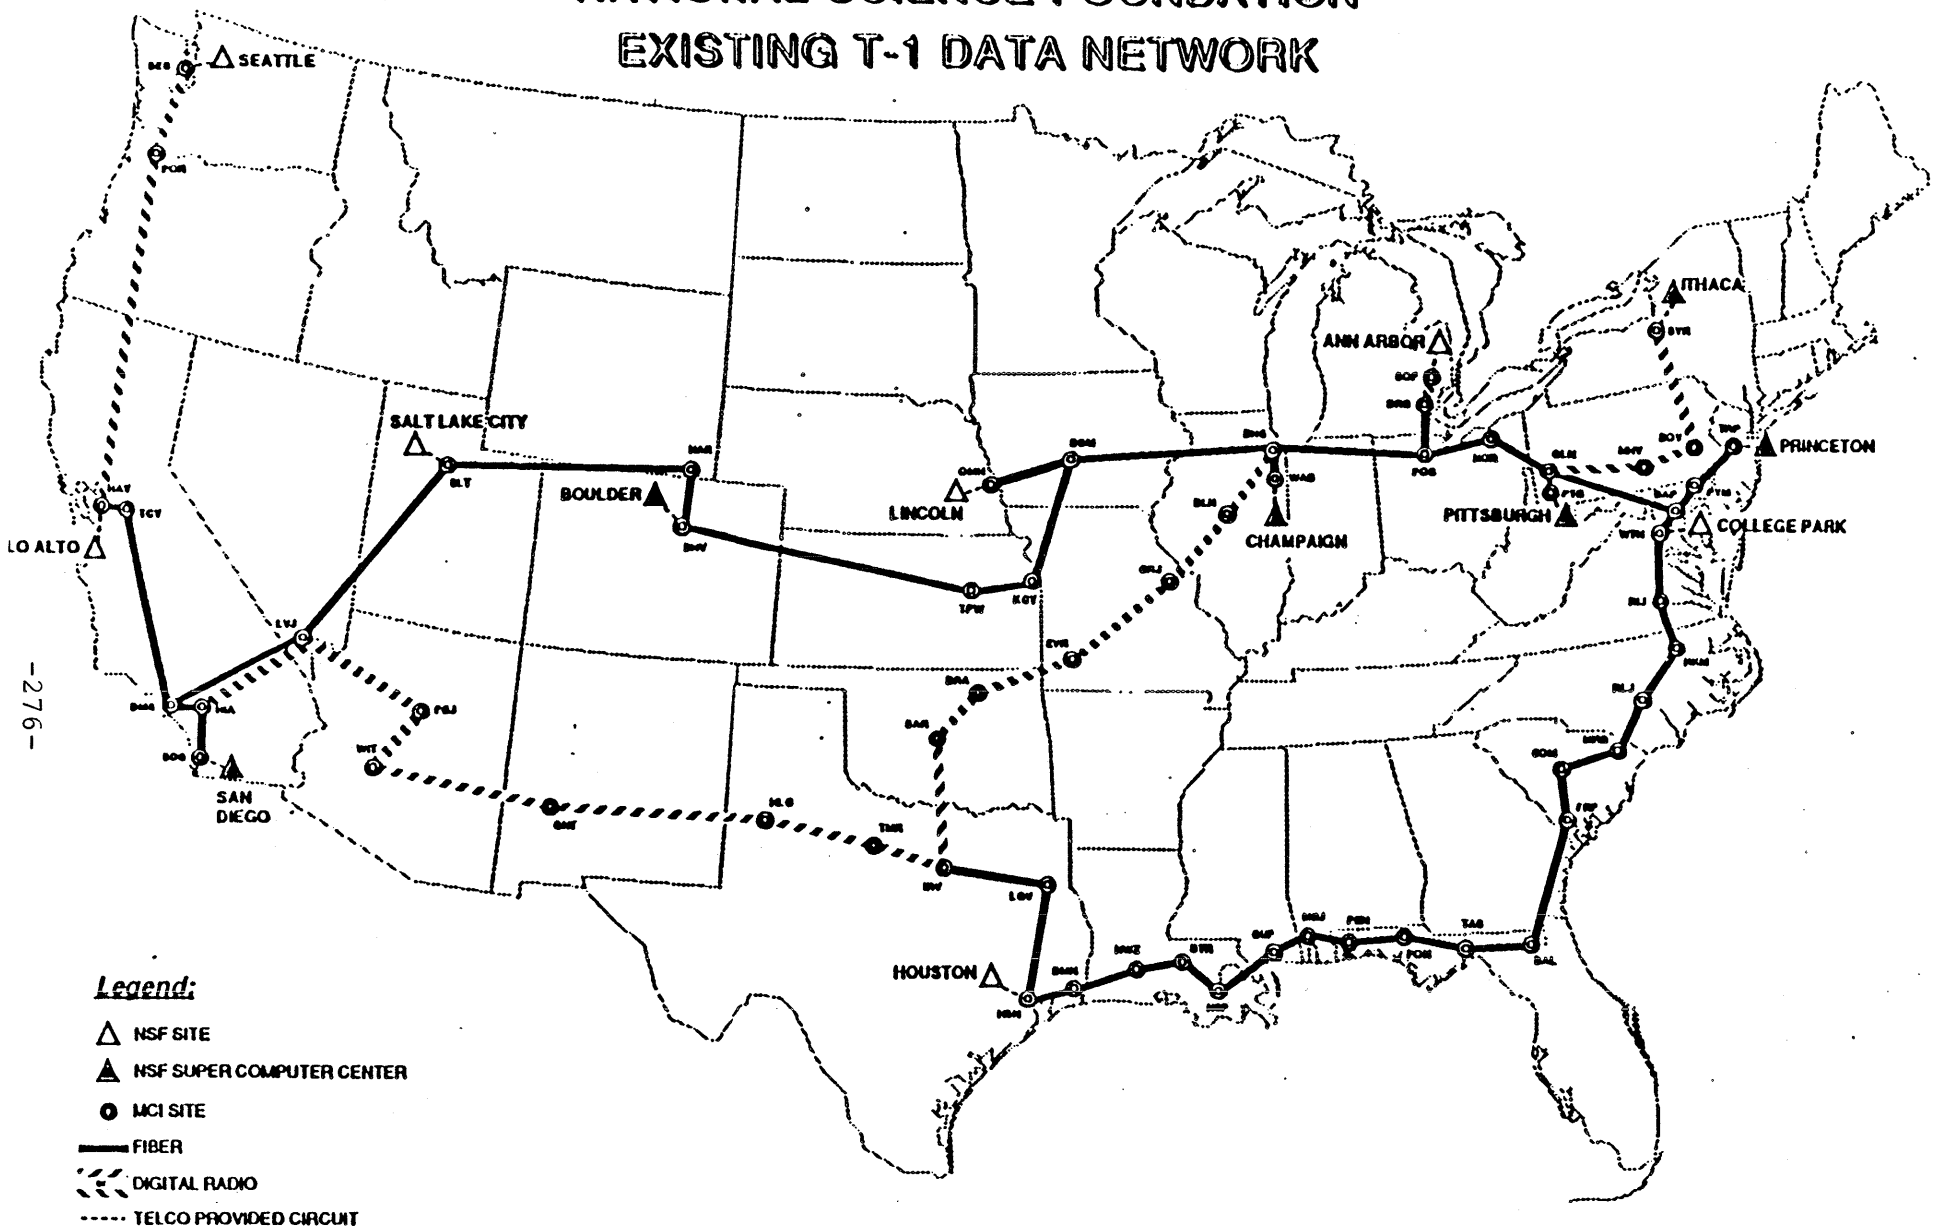

J. W. D., Jr.

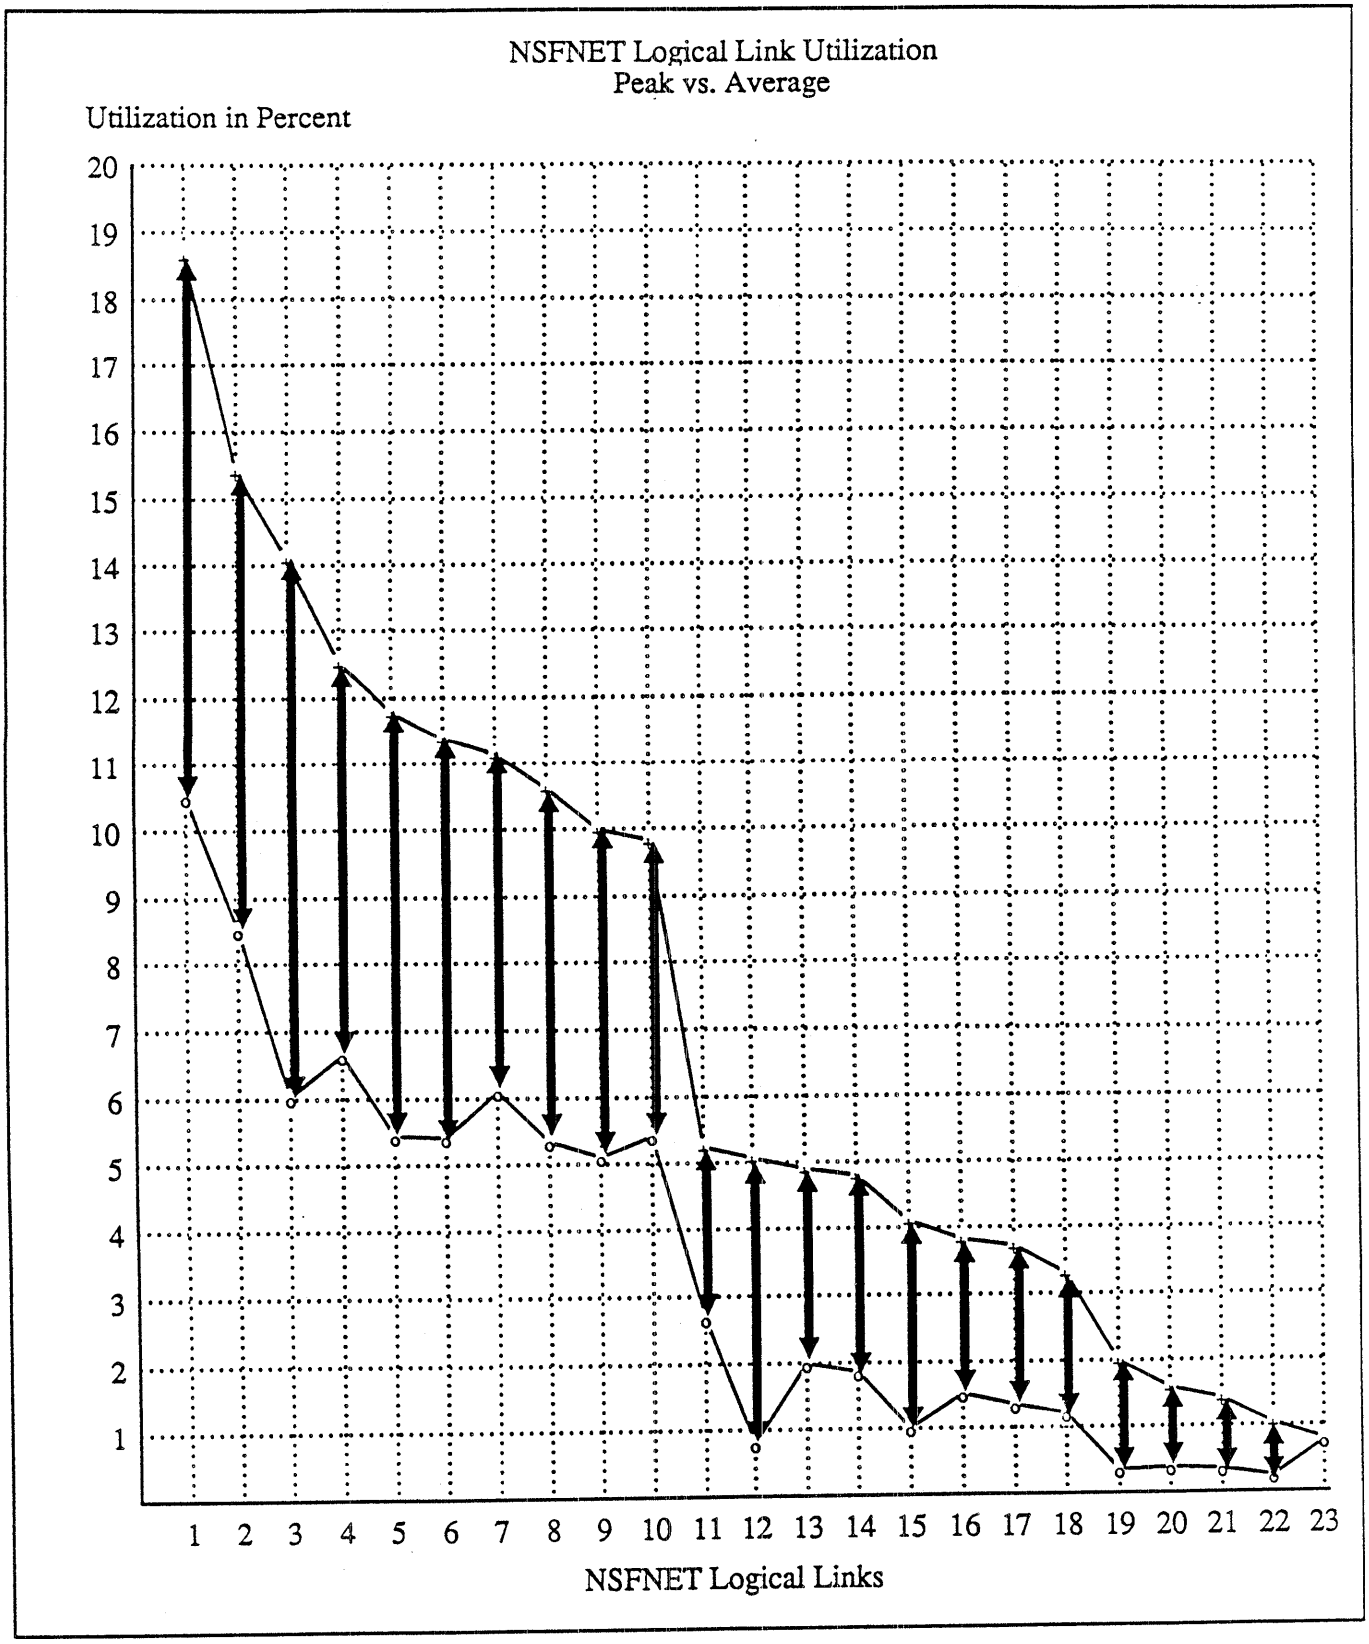
NSFNET Logical Link Utilization
Peak vs. Average
Utilization in Percent
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
NSFNET Logical Links

NATIONAL SCIENCE FOUNDATION
T-1 DATA NETWORK #2

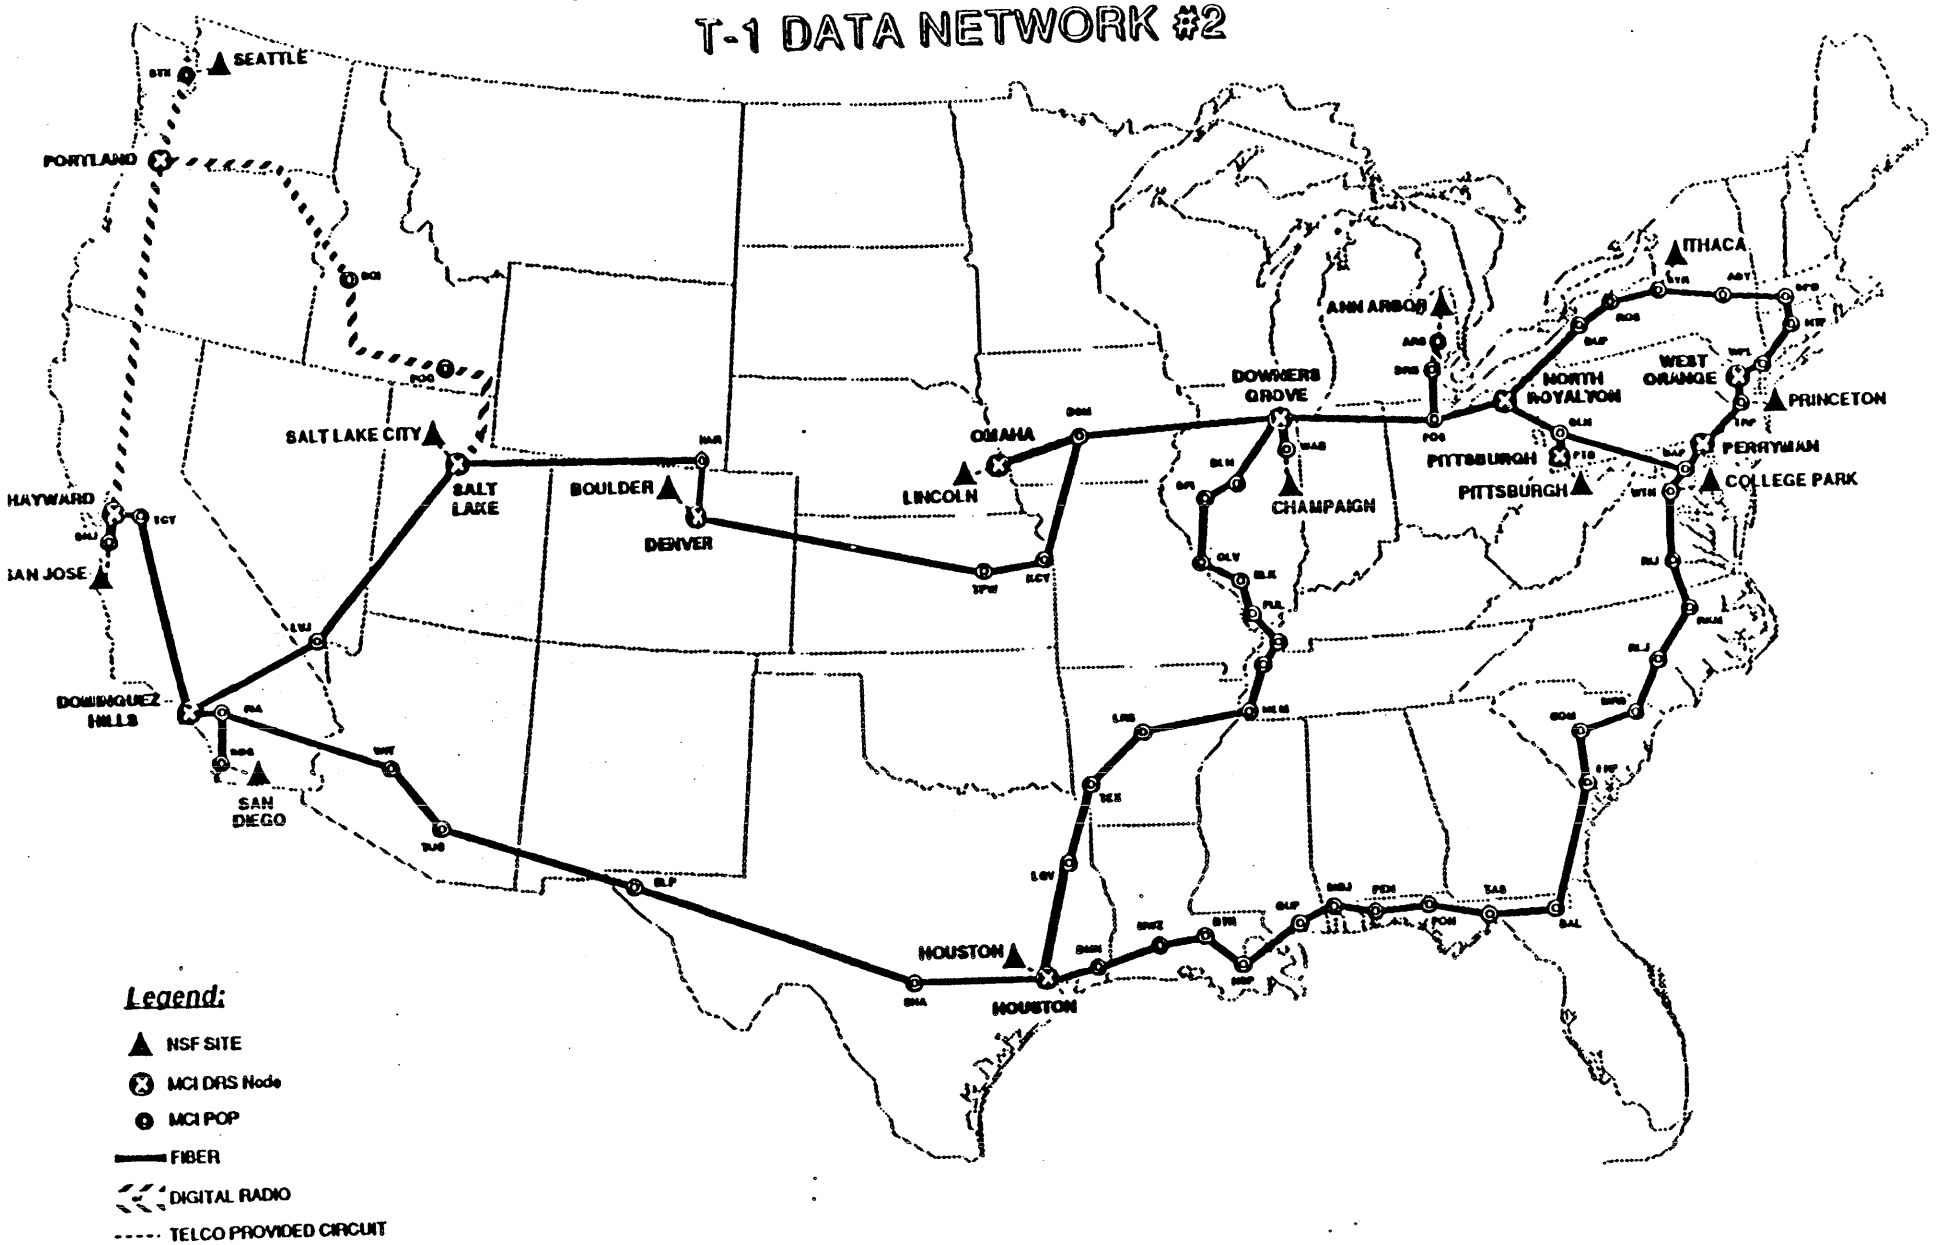

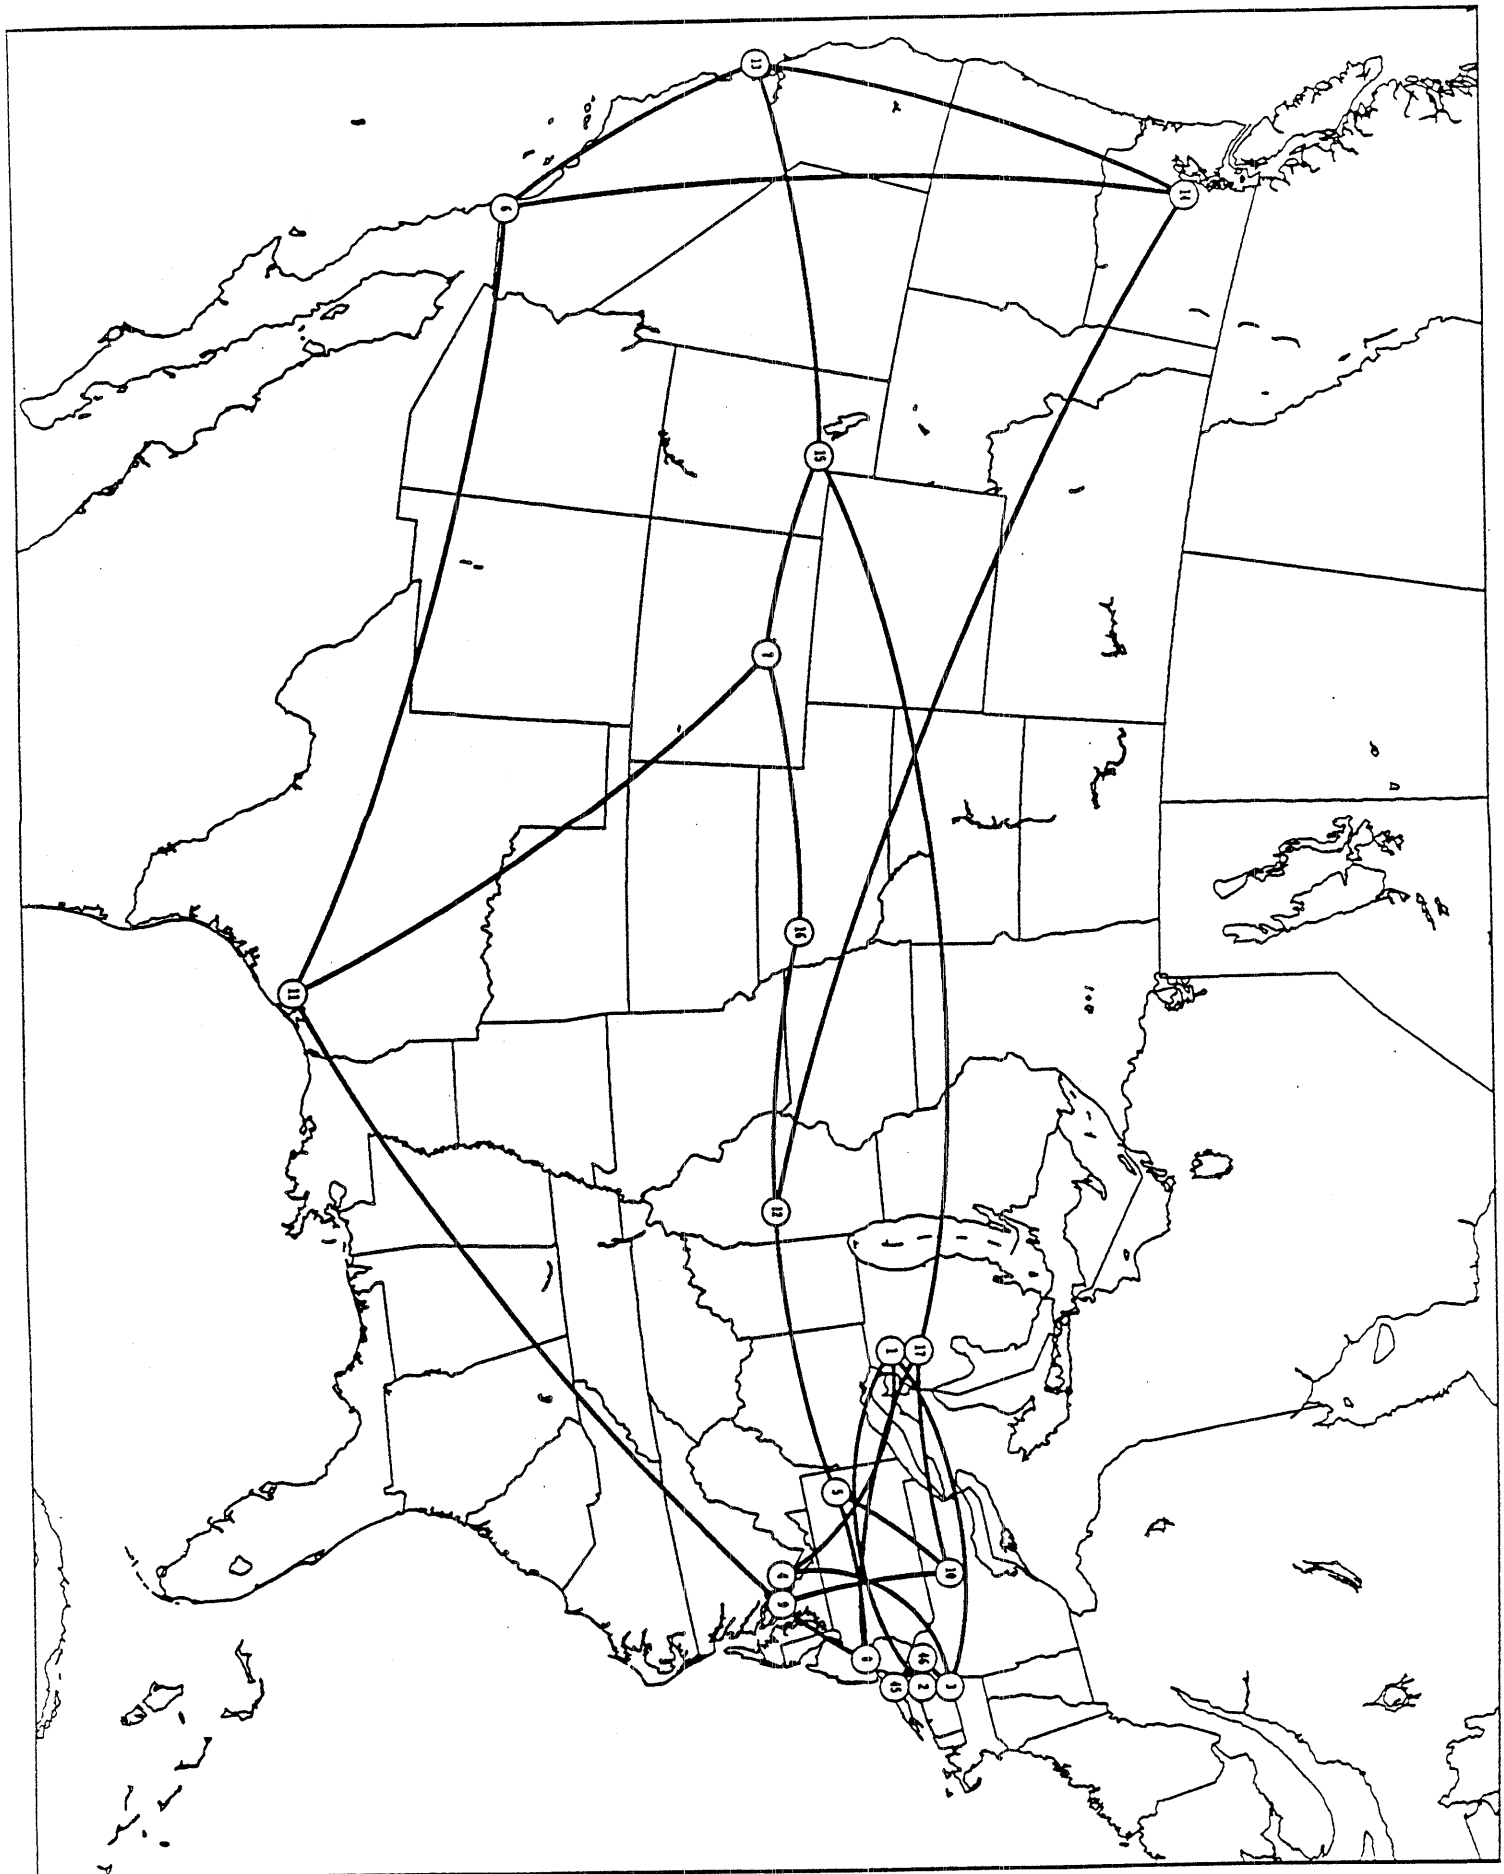

New NSFNET Backbone
Circles contain NSS number (1-4 and 45-46 are test nodes)

Prepared by NSFNET-Info@merit.edu at Thu Apr 6 16:57:09 1989
netmap-1.5 program by Brian Reid, map data from World Data Bank II
Lambert Conformal Projection [44°N,33°N], Map center: [40°N, 96° 30'W]
Image resolution 300/in., stroke limit 1 pixels

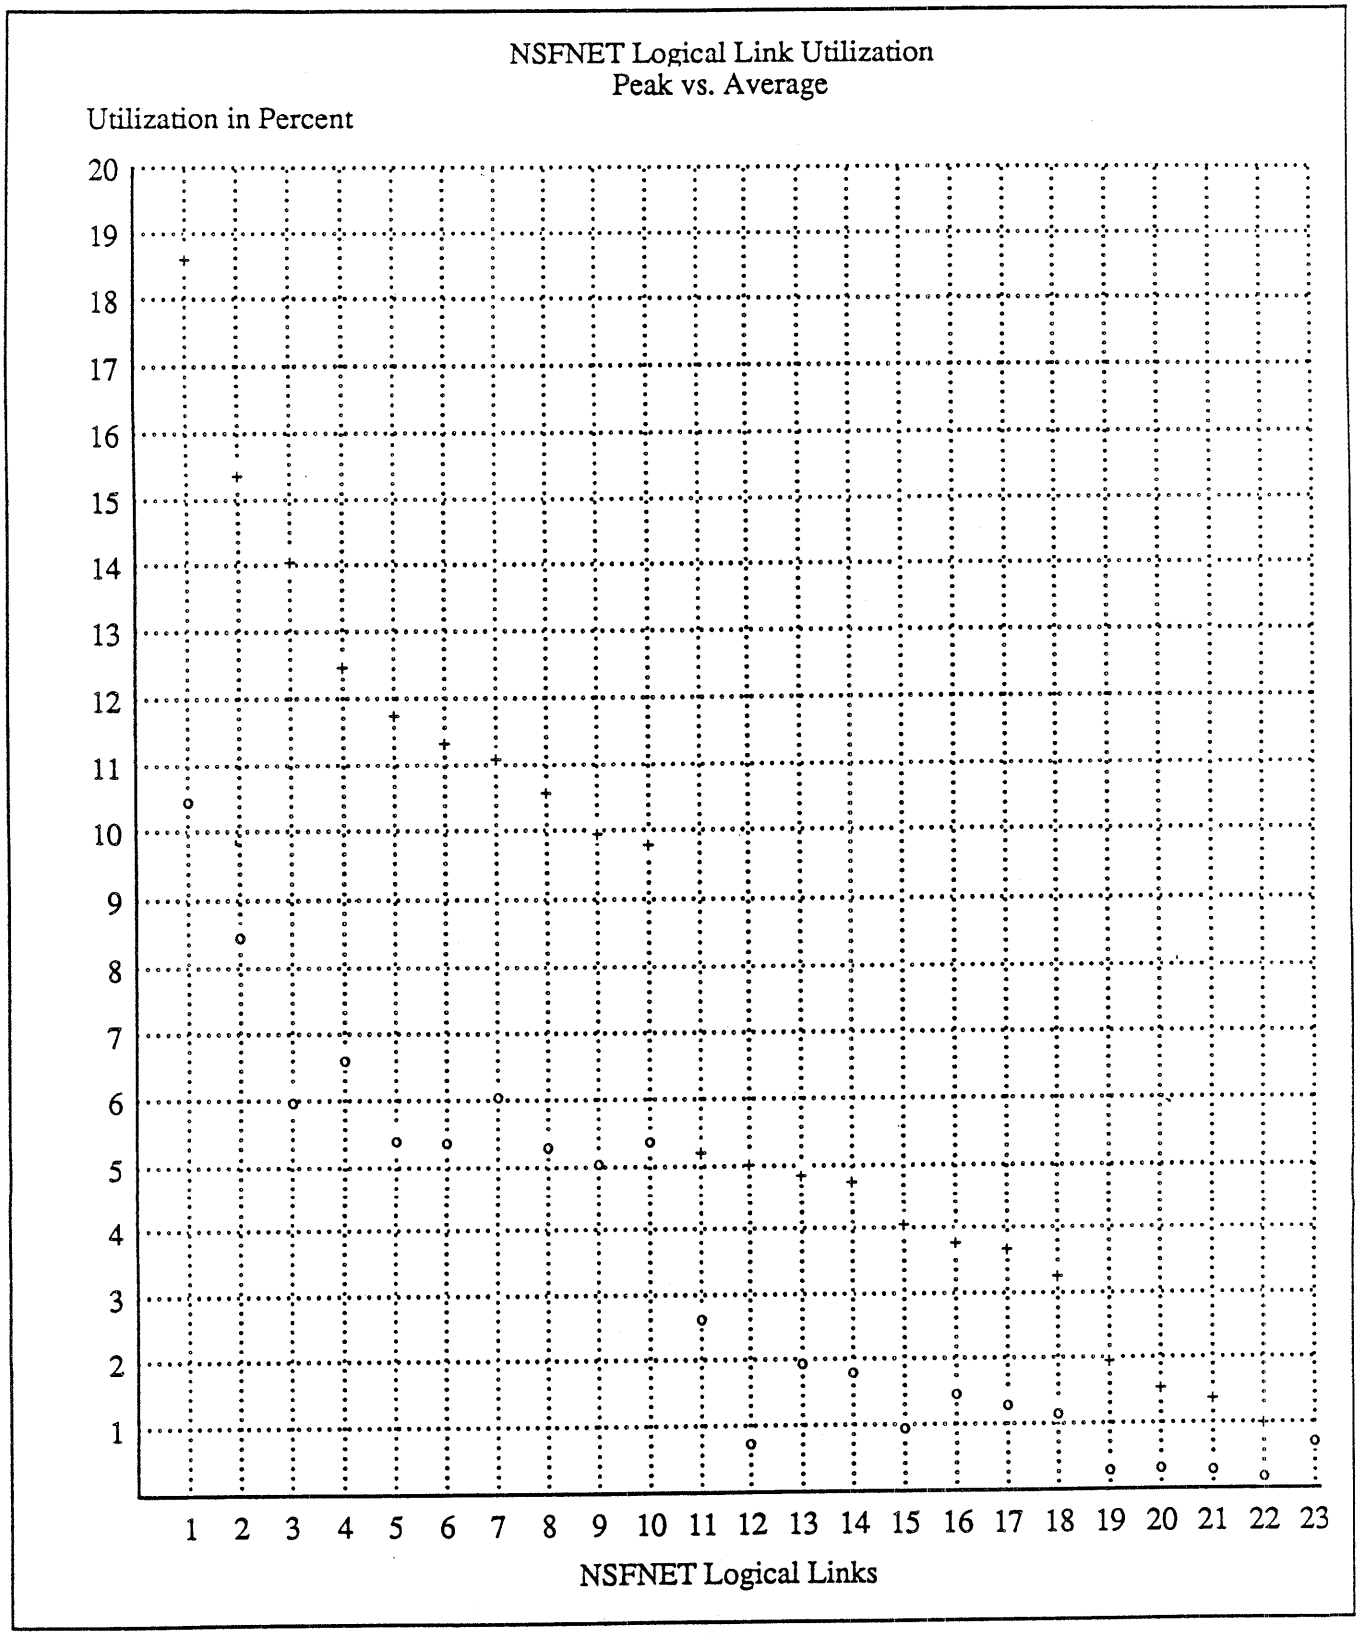
NSFNET Logical Link Utilization
Peak vs. Average
Utilization in Percent
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
NSFNET Logical Links

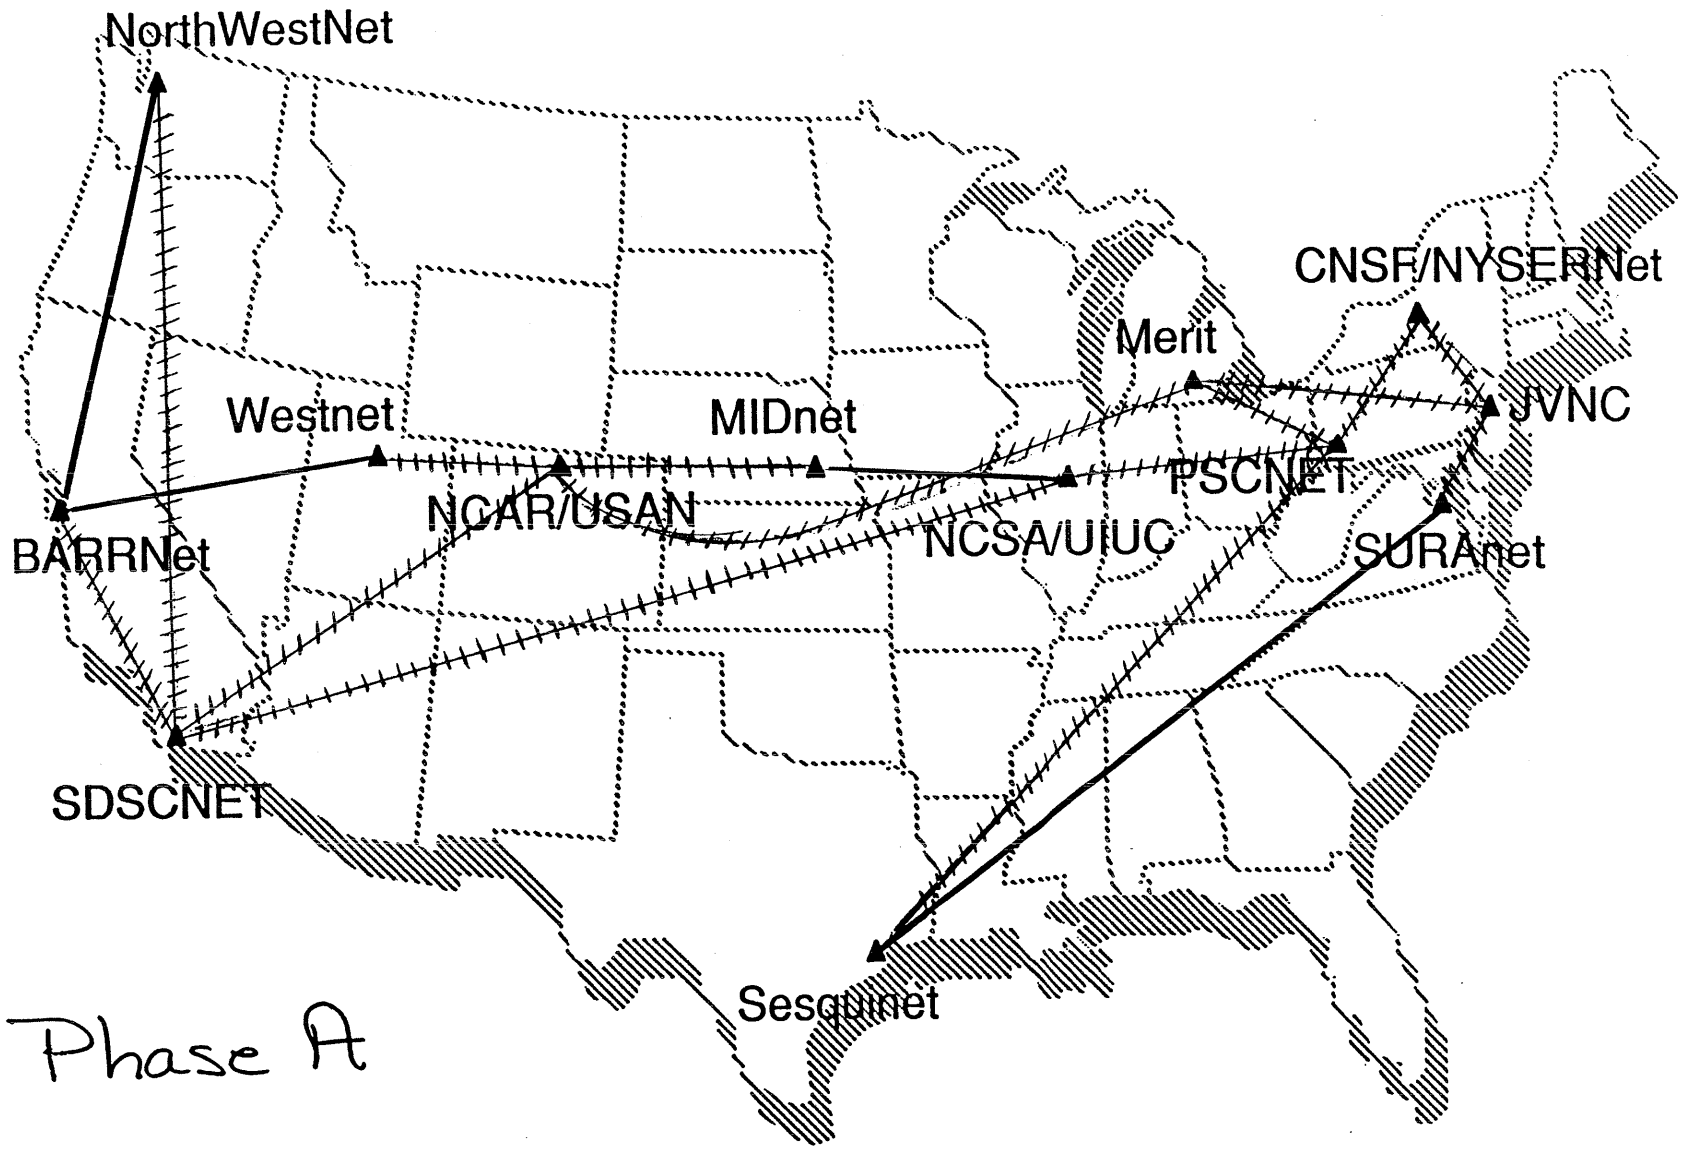
NorthWestNet
CNSF/NYSERNet
Merit
JVNC
Westnet
MIDnet
PSCNET
NCAR/USAN
NCSA/UIUC
SURAnet
BARRNet
SDSCNET
Sesquinet
Phase A

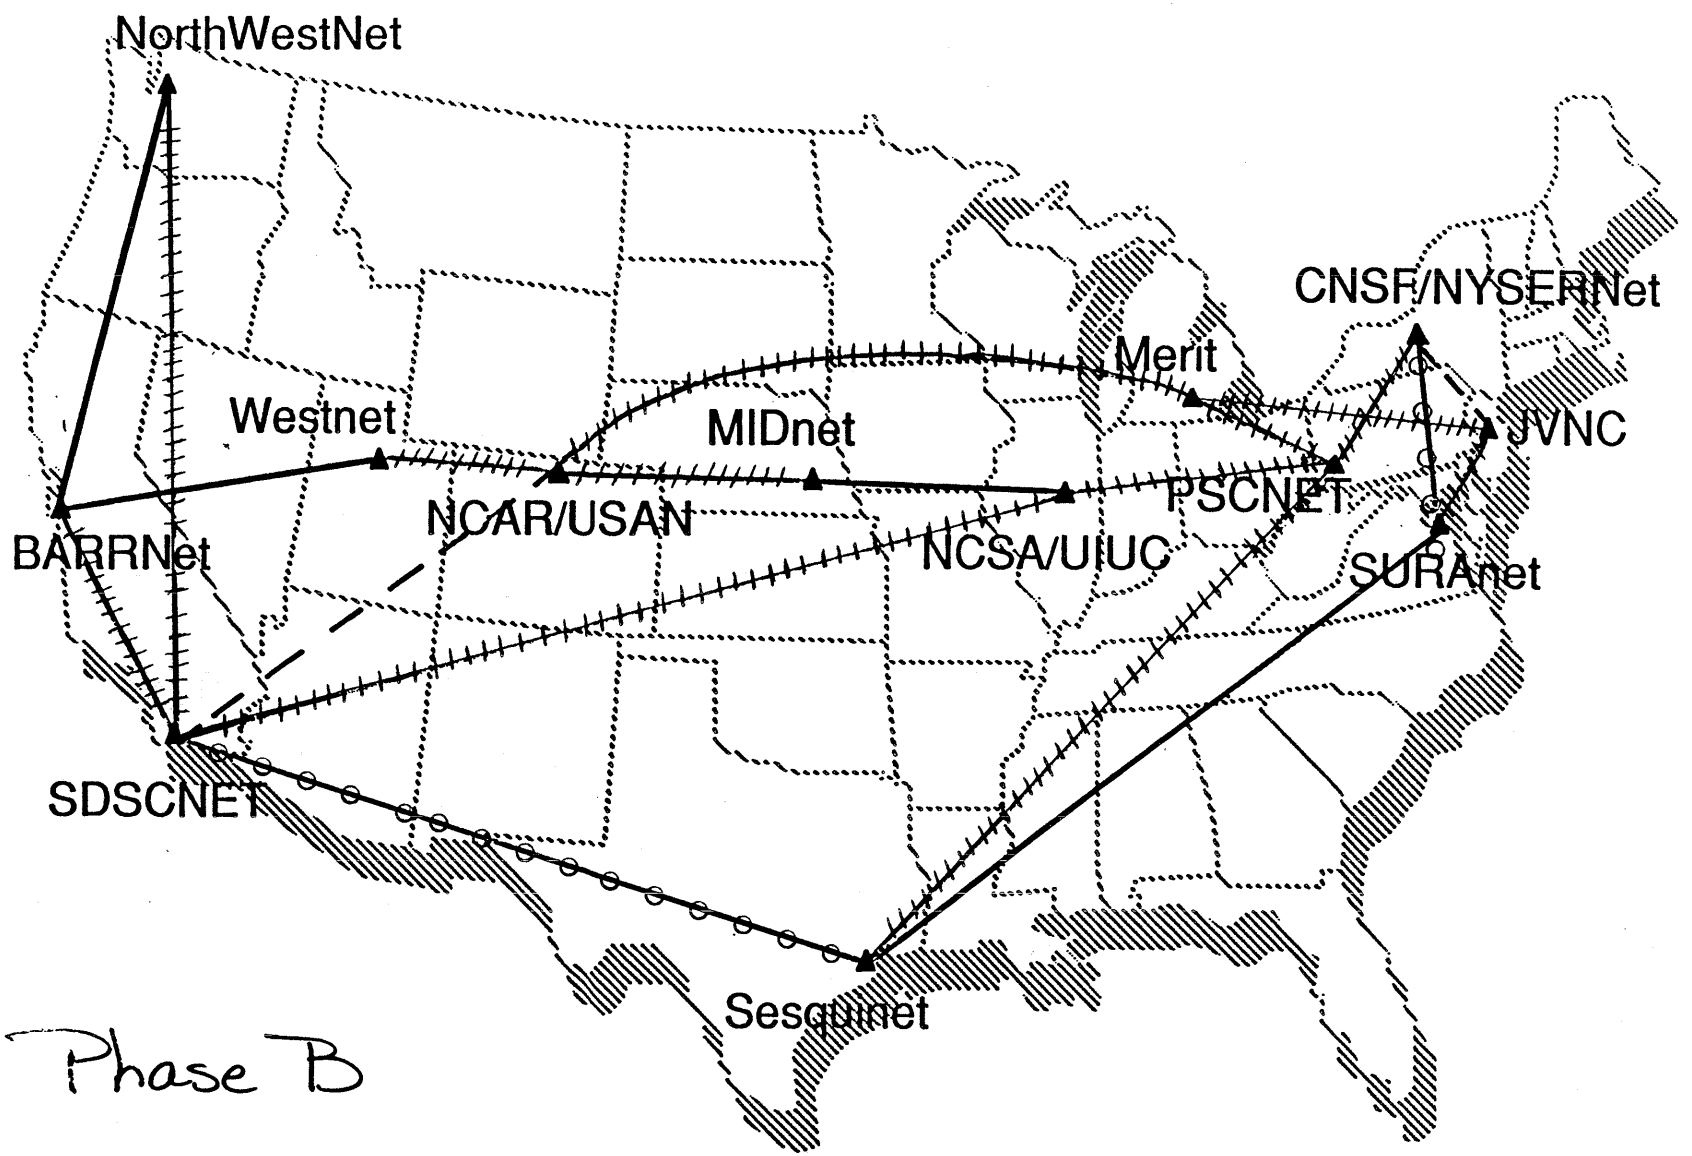
NorthWestNet
CNSF/NYSERNet
Merit
Westnet
MIDnet
JVNC
PSCNET
NCAR/USAN
BARRNet
NCSA/UIUC
SURAnet
SDSCNET
Sesquinet
Phase B

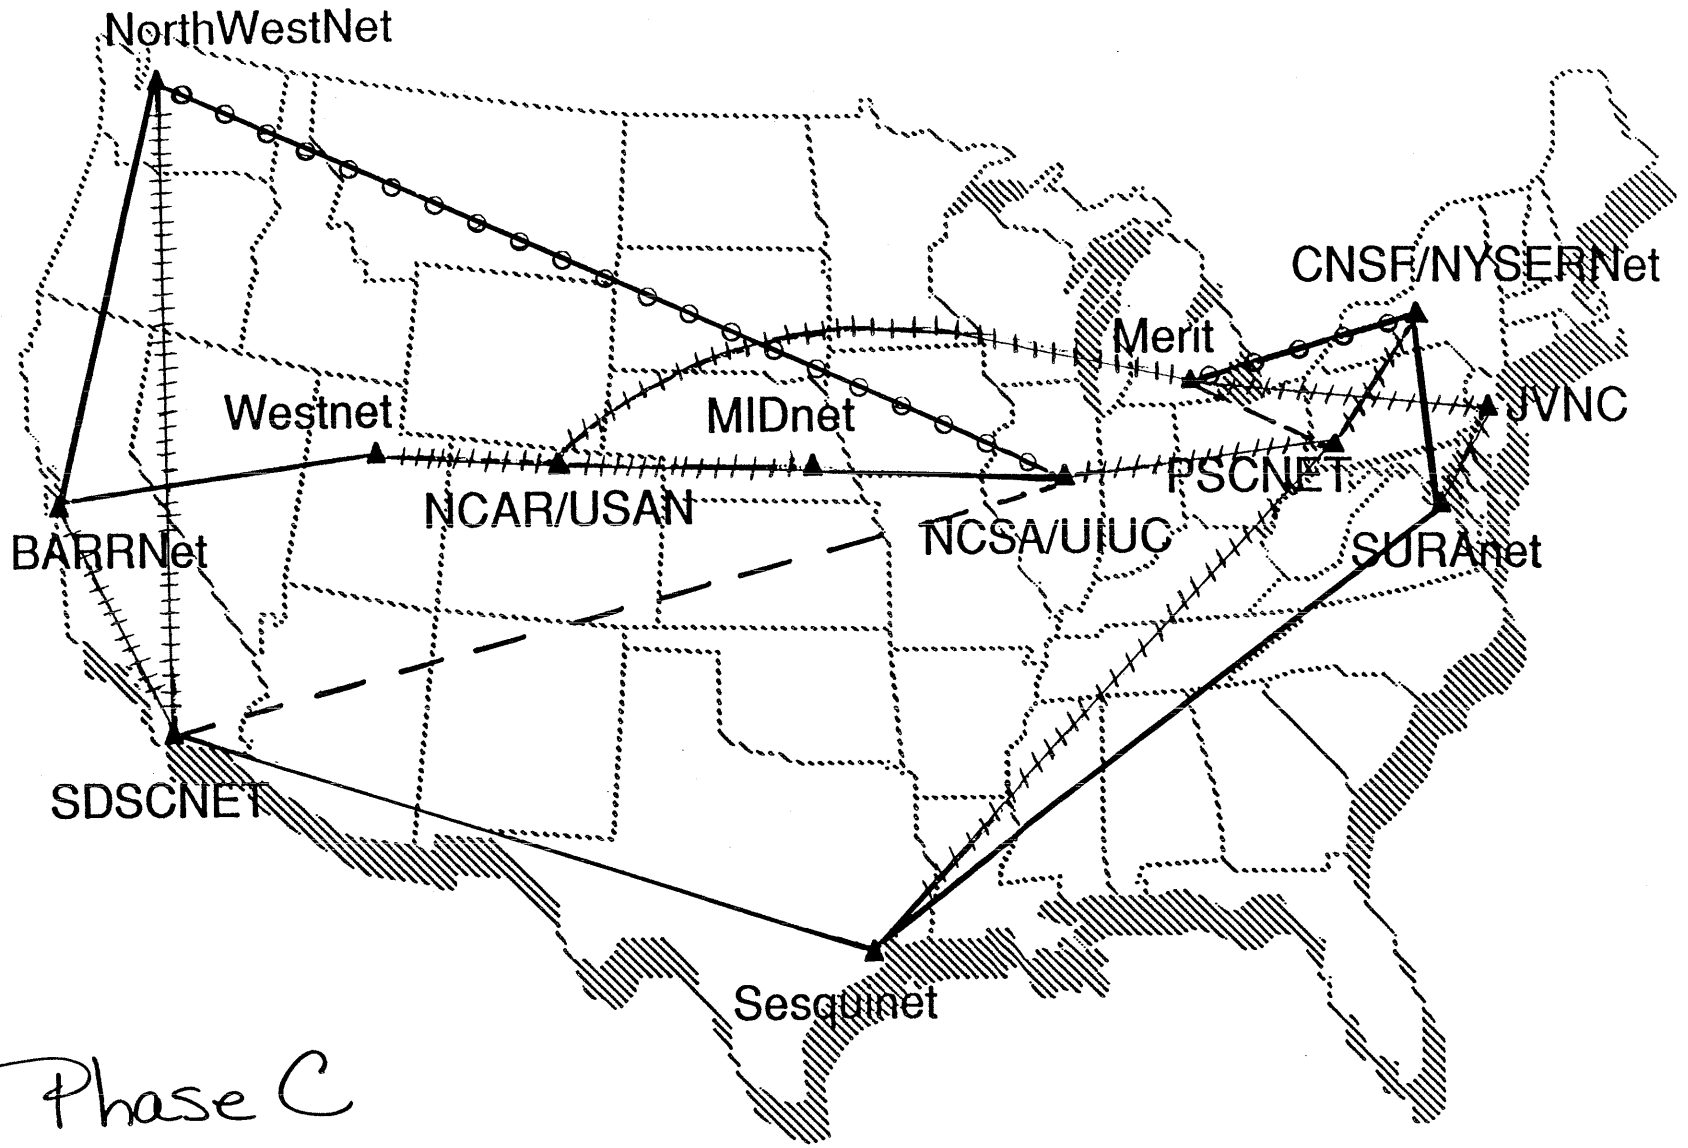
NorthWestNet
CNSF/NYSERNet
Merit
JVNC
Westnet
MIDnet
PSCNET
NCAR/USAN
NCSA/UIUC
SURAnet
BARRNet
SDSCNET
Sesquinet
Phase C

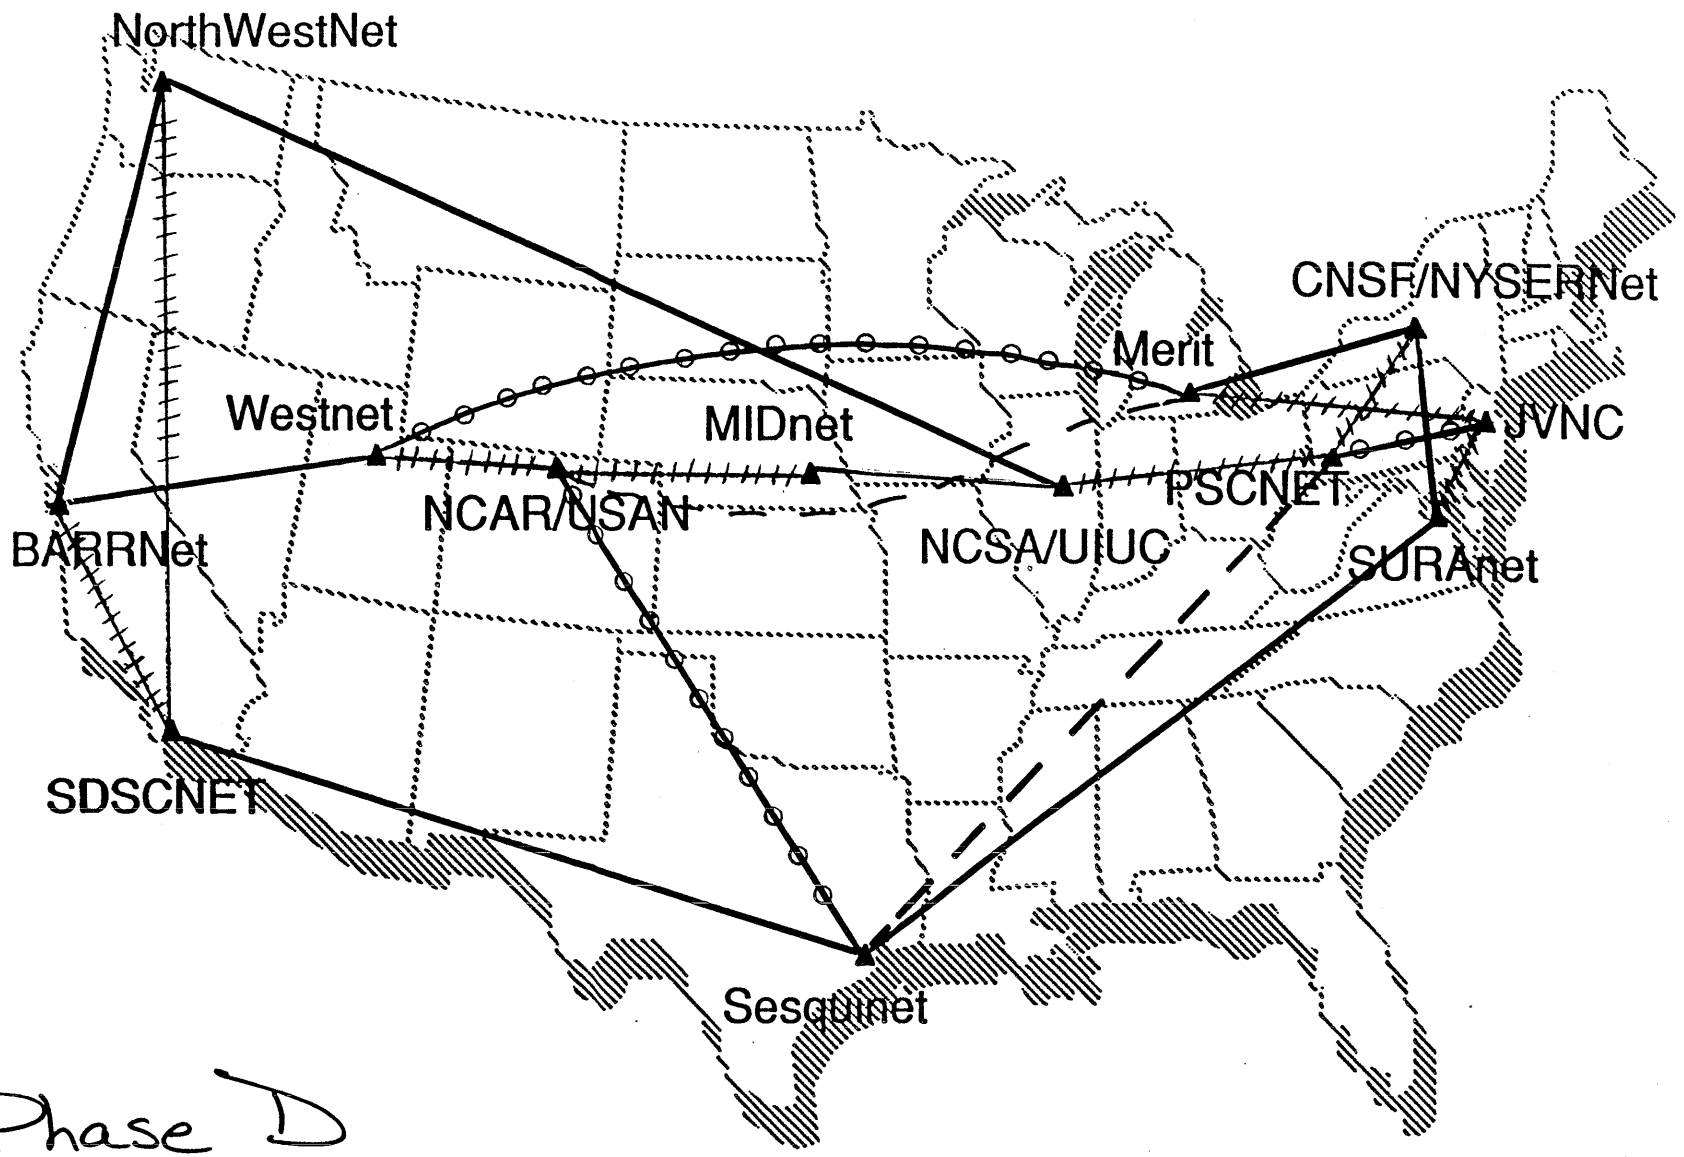
NorthWestNet
CNSF/NYSERNet
Merit
Westnet
MIDnet
JVNC
PSCNET
NCAR/USAN
NCSA/UIUC
SURAnet
BARRNet
SDSCNET
Sesquinet

Phase D

NEW TOPOLOGY

- " Fatter " pipes for packet switching

- No multiplexing and demultiplexing at sub T1 rates

  Clear channel T1

- Greater redundancy

  - No single tail circuit sites
  - No articulation points

- Optimized for MCI infrastructure

  - MCI redundancy
  - Optimum MCI routes

- Greater degree of connectivity

  3.07 v/s 2.15

Planned NSFNET/NSN/DDN connection at NASA Ames

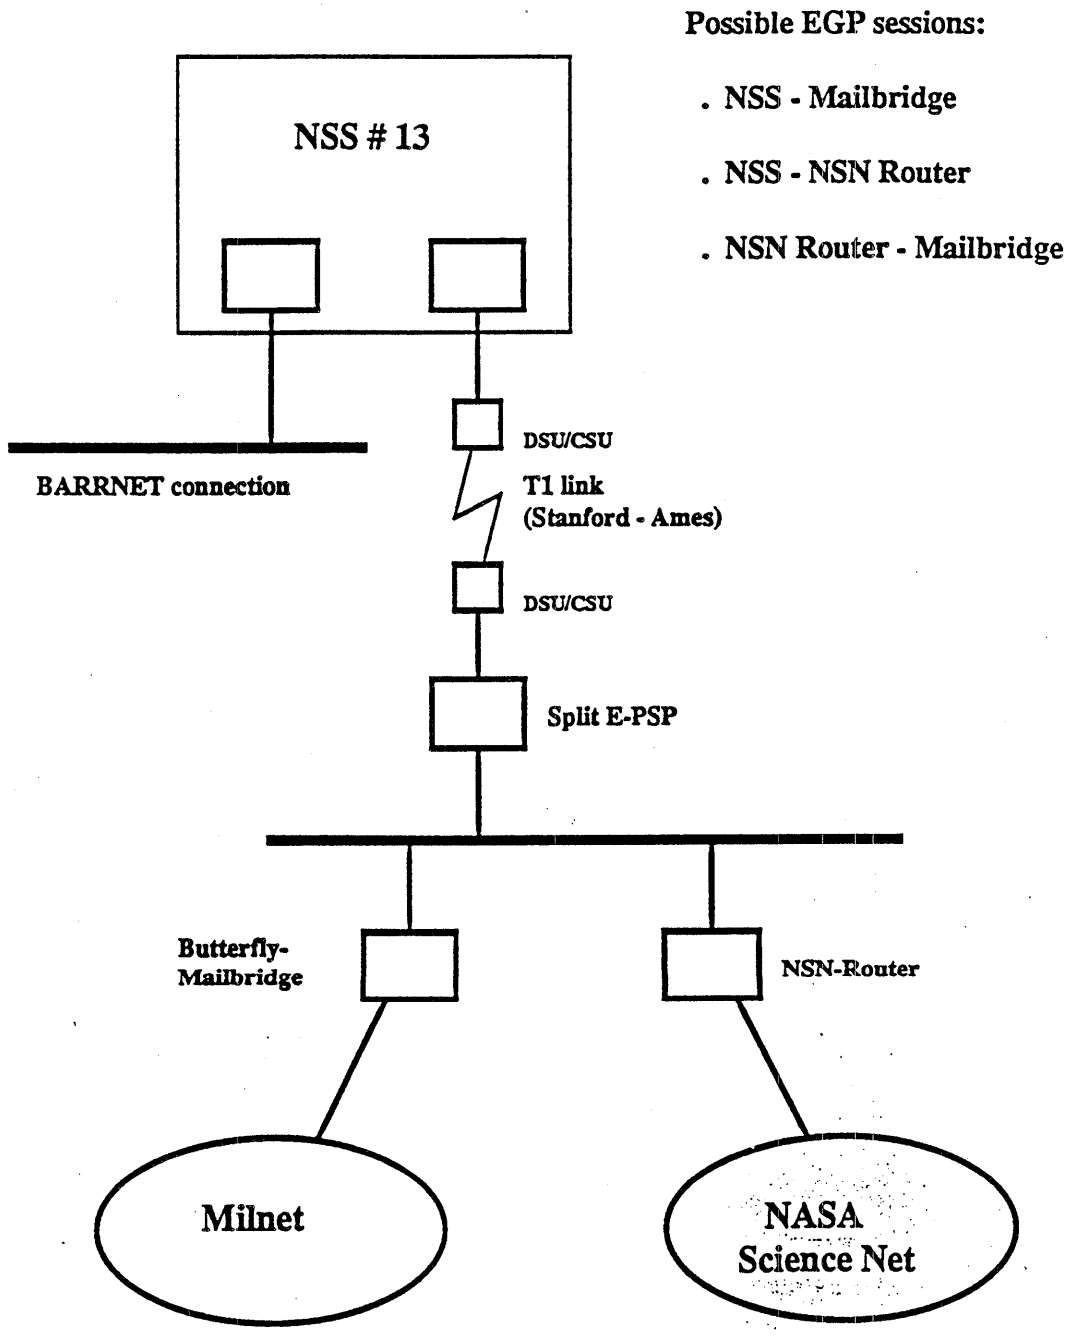

7 March 1989, HWB

Planned NSFNET/NSN/DDN connection at the University of Maryland

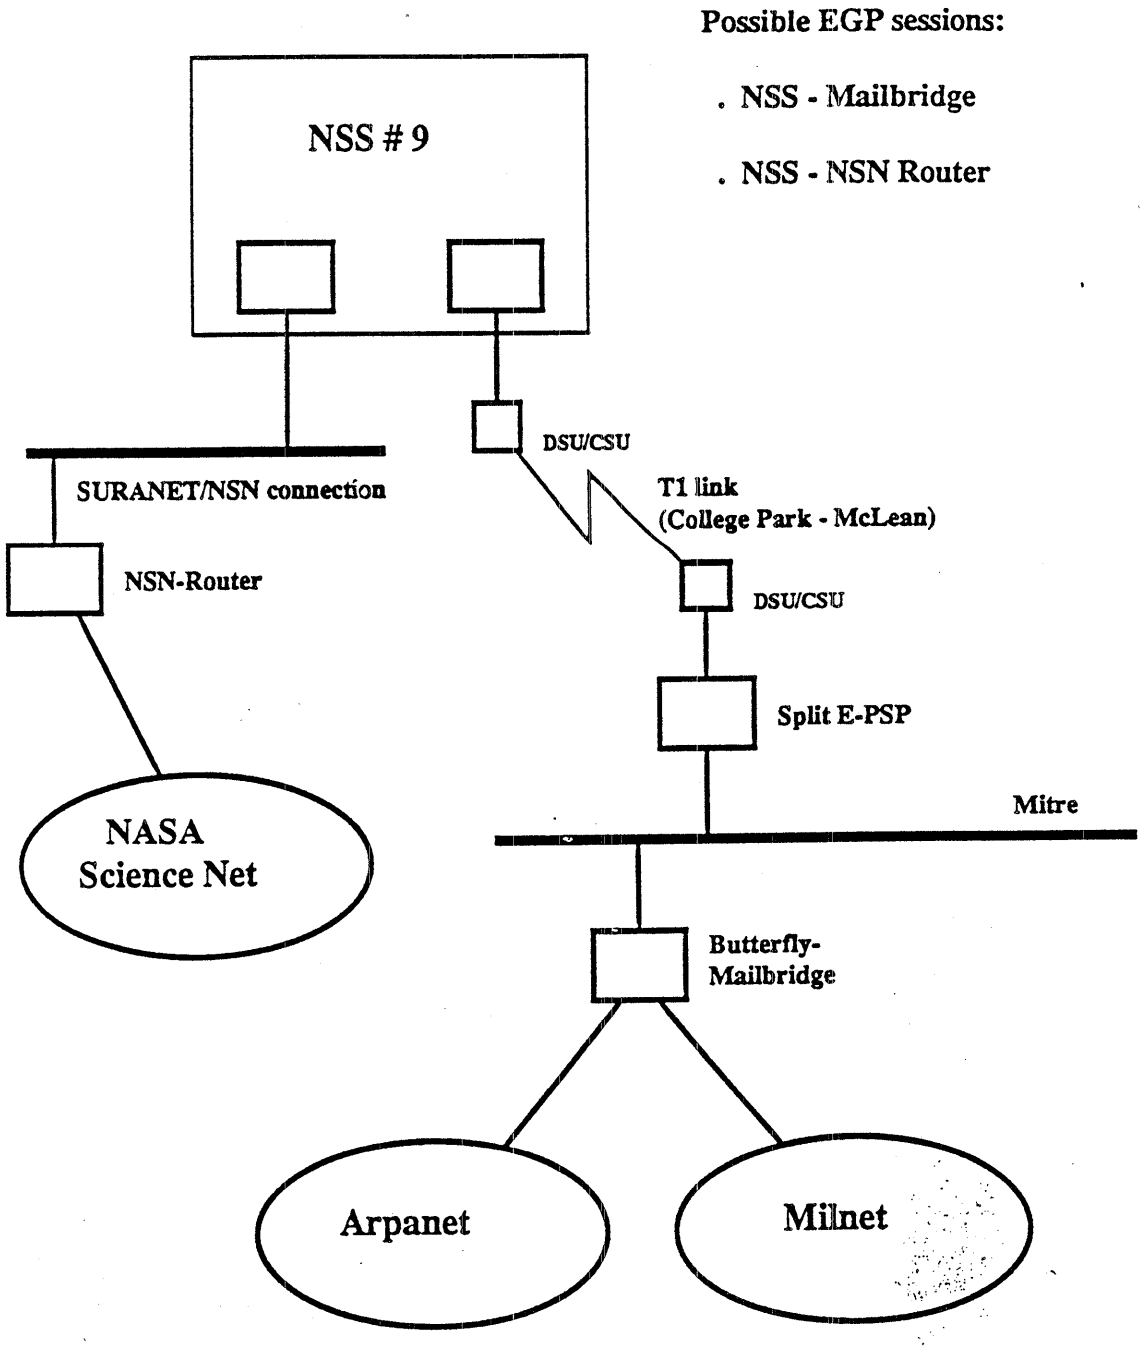

7 March 1989, HWB

Recently deployed:

. SNMP

Planned:

. ISO CLNP support

Proposed:

. T3 upgrade in 1990

Being considered:

. X.25 support

10 April 1989, HWB

Nifty NSFNET Stats, using NNStat
Presented by Elise Gerich
Reported by Dave Katz

When the National Science Foundation sent out the solicitation for the NSFNET in 1987, one of the charges was to provide continuous information gathering relative to the activity on the network and distribute the resulting data in electronic and hardcopy form.

This is being accomplished with the use of NNStat (NSF Net STATistics), a package written with NSF funding by Bob Braden and Annette DeSchon at the Information Sciences Institute at USC.

This package was originally coded to run in a Sun workstation attached to an Ethernet; changes were made to it and to the BSD 4.3 kernel so that it would run on an IBM RT PC attached to a token ring network.

The NNStat set of programs and utilities provide a flexible method for gathering traffic statistics, querying them interactively, and remotely collecting them for later analysis.

There are three major components to the package:

Monitoring--"Statspy" listens to the traffic on a network and builds statistical objects based on what it hears

Query--"Rspy" allows remote interactive queries to Statspy, providing the capability of dynamically altering its configuration and displaying gathered data

Collection--"Collect" remotely retrieves statistics from Statspy and logs it to disk for later processing

Installation

A cooperative effort was undertaken to install NNStat in the backbone. First of all, the NNStat application itself was ported to the IBM RT by Merit Internet Engineering staff. This entailed minor modifications to run in the RT environment, as well as the addition of a few features, the most notable of which is a security enhancement that restricts access to statistics data to a small number of machines.

The Berkeley 4.3 "Packet Filter" was installed in the NSS kernel by Merit IE staff to provide an application interface to the raw packet stream. This package was optimized and enhanced to meet the needs of NNStat.

The IBM Token Ring adaptor driver software was modified by Merit IE staff to support the Packet Filter interface and the "promiscuous" reception of packets.

A set of Read Only Memory (ROM) chips were obtained through the courtesy of Texas Instruments that allow the Token Ring adaptor to receive all token ring traffic. These chips were duplicated en masse by members of the Network Operations Center staff, and were installed by IBM's service representatives at each of the backbone sites.

Staff in the Merit Information Services group continue to work on a database into which the raw statistics information is placed. This database will provide a flexible way of retrieving subsets of sample data as well as aggregations of data. It will be possible to convert the data to graphical representations as well.

Configuration

A dedicated IBM RT in each NSS is used to run Statspy. The RT receives all data on the NSS token ring, which has a number of desirable characteristics:

- all user data traverses the token ring
- both transit traffic and traffic to/from the local NSS are present
- intra-backbone traffic (such as routing protocol traffic) is present
- no local site (intra-regional) traffic is visible

An IBM RT at the NSFNET NOC in Ann Arbor, running Collect, periodically calls out to each of the backbone statistics gatherers, logs the values of their statistical objects to disk, and clears the objects. The logs are transferred to the NSFNET IS mainframe. Once Berkeley Socket support is available on the IS mainframe, the Collect application will be ported to run there, thus eliminating staging the data on an RT.

On the mainframe, the raw statistics data is processed and loaded into a SPIRES database. Ancillary data, such as the mappings between network numbers and Autonomous System numbers, are added to the database as well. The database will allow flexible data retrieval by location, time, and other parameters. The raw data are archived to tape.

The gathering of statistics data has a negligible impact on network performance. The CPU-intensive parts of the process are performed on dedicated machines that are not involved in the switching of packets, and the data collected from the backbone is modest in volume compared to the traffic levels.

Data Collected

Approximately 10 million bytes of raw statistics data are collected daily. The bulk of the data are source/destination

network number pairs. Other data collected include the token ring packet switching rate, the distribution of well-known TCP and UDP ports, and the distribution of packet sizes. Data is collected with a granularity of fifteen minutes.

The accompanying slides contain the following:

A sample configuration for Statspy. This configuration defines the statistical objects to be built, as well as the access restrictions in effect.

Sample output from Rspy. This output corresponds to the configuration shown previously. The verbose logs from Collect are nearly identical. Terse logs from Collect contain the same information but in considerably less disk space.

Sample graphs. These graphs were put together quickly using relatively small amounts of data for illustrative purposes, although the data is real.

NNStat in the NSFNET

Dave Katz (NSFNET)

NNStat in the NSFnet Backbone
Dave Katz, Merit

NNStat: NsfNetSTATistics

Written with NSF funding by Bob Braden and Annette DeSchon at USC ISI

Collection of programs and scripts with three major components:

- Statspy: Listen to network and build statistical objects
- Rspy: Remote interactive query of Statspy
- Collect: Remote data collection

Originally written for Sun Workstation/Ethernet

Ported to IBM RT/Token Ring by Merit

Statspy

Listens "promiscuously" to network medium

Object types
- Simple frequency count
- Enumerated frequency count
- Pairwise frequency count
- Histogram (linear, exponential)
- Temporal clustering

Filters
- Range of values
- Set of values

Fields
- Ethernet, TCP, IP, UDP, ICMP headers
- Packets per second
- Interarrival time
- Raw packet

Flexible configuration language

Rspy

Provides remote query facility to any Statspy anywhere

Display, create, alter objects

Collect

Creates disk logs of object values

Programmable poll, checkpoint, clear times

Terse/verbose logs

NSFnet implementation

Dedicated RT in each NSS

Data collected on token ring, which carries
- All user traffic
- Transit traffic
- Intra-backbone traffic (routing, etc.)
- No local site traffic

Added access restrictions

Collect polls all NSSs every 5 minutes, checkpoints/clears every 15 minutes, data collected on local RT, staged to IS mainframe for database insertion

Eventually port Collect to CMS

Merit version available via anonymous FTP from merit.edu (NNStat/NNStat.Merit.tar.Z)


```
Rspy>host 129.140.17.16
Rspy>show ?
Connecting to Statspy on Host: 129.140.17.16
OK.
Acquired 177490771 packets in 1517644 secs => 116(avg) 997(max) 1100(inst) /sec
Fields are:
packet interval.ms per.sec Ether.src Ether.dst Ether.type IP.version IP.length IP.option IP.TOS
IP.offset IP.protocol IP.srcnet IP.dstnet IP.srchost IP.dsthost TCP.srcport TCP.dstport
UDP.srcport UDP.dstport ICMP.type
Rspy>read *
Connecting to Statspy on Host: 129.140.17.16
OK.

OBJECT: N17.all-persec Class= hist [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 308 (+0 orphans)
  Avg= 161 Min= 97 Max= 346
 [0-19]= 0
 [20-39]= 0
 [40-59]= 0
 [60-79]= 0
 [80-99]= 1
 [100-119]= 16
 [120-139]= 64
 [140-159]= 90
 [160-179]= 60
 [180-199]= 40
 [200-219]= 21
 [220-239]= 11
 [240-259]= 2
 [260-279]= 1
 [280-299]= 0
 [300-319]= 1
 [320-339]= 0
 [340-359]= 1
 [360-379]= 0

OBJECT: N17.10-pkts-in Class= freq-only [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 8503 (+0 orphans)
 #bins= 1 Other= 8503 (100.0%)
[0]= 0 (0%)

OBJECT: N17.10-ports Class= freq-all [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 5120 (+0 orphans)
 #bins = 10
[23 "Telnet"]= 1251 (24.4%) @ -0sec [119 "NetNews"]= 1162 (22.7%) @ -0sec
[20 "FTP data"]= 966 (18.9%) @ -0sec [25 "SMTP"]= 678 (13.2%) @ -0sec
[153 "SGMP"]= 588 (11.5%) @ -31sec [53 "Domains"]= 318 ( 6.2%) @ -10sec
[513 "rwho | rlogin"]= 83 ( 1.6%) @ -2sec [21 "FTP"]= 71 ( 1.4%) @ -21sec
[123 "NTP"]= 2 ( <.1%) @ -57sec [79 "Finger"]= 1 ( <.1%) @ -77sec
```

```
restrict readwrite 129.140.17.1 255.255.255.255
restrict readonly 129.140.0.0 255.255.0.0
restrict readonly 35.1.1.0 255.255.255.0
enum {
 *port* (20 "FTP data", 21 FTP, 23 Telnet, 25 SMTP,
     37 Time, 42 Name, 43 Whois, 53 Domains,
     69 TFTP, 79 Finger, 103 X.400, 104 "X.400-SND", 109 POP2,
     111 sunrpc, 115 SFTP, 119 NetNews, 123 NTP, 153 SGMP,
     512 exec, 513 "rwho|rlogin", 514 shell, 515 printer, 520 RIP)
}
attach {
   if Ether.dst is eqf (50:0:5a:1a:f:dc) {
      record IP.srcnet, IP.dstnet in N17.10-matrix matrix-all;
      record IP.length in N17.10-iplen hist(40);
      record IP.protocol in N17.10-IP.proto freq-all;
      if TCP.dstport is wellknownport setf(
         "FTP data", "FTP", "Telnet", "SMTP",
         "Time", "Name", "Whois", "Domains",
         "TFTP", "Finger", "X.400", "X.400-SND",
         "POP2", "sunrpc", "SFTP", "NetNews",
         "NTP", "SGMP", "exec", "rwho|rlogin",
         "shell", "printer", "RIP") {
         record TCP.dstport in N17.10-ports freq-all;
      } else if TCP.srcport is wellknownport {
         record TCP.srcport in N17.10-ports;
      }
      if UDP.dstport is wellknownport {
         record UDP.dstport in N17.10-ports;
      } else if UDP.srcport is wellknownport {
         record UDP.srcport in N17.10-ports;
      }
   }
   if Ether.src is eqf (50:0:5a:1a:f:dc) {
      if IP.dstnet is eqf(129.140.0.0) {
         record IP.srcnet, IP.dstnet in N17.10-bkbn-in matrix-all;
      }
      record IP.protocol in N17.10-pkts-in freq-only(0);
   }
   record per.sec in N17.all-persec hist(20);
}
```



```
OBJECT: N17.10-IP.proto  Class= freq-all  [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 5974 (+0 orphans)
 #bins = 3
[6 "TCP"]= 4738  (79.3%) @ -0sec    [17 "UDP"]= 956  (16.0%) @ -10sec
[1 "ICMP"]= 280   ( 4.7%) @ -0sec

OBJECT: N17.10-iplen  Class= hist  [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 5974 (+0 orphans)
  Avg= 78  Min= 28  Max= 1400
 [0-39]= 5
 [40-79]= 5135
 [80-119]= 395
 [120-159]= 122
 [160-199]= 10
 [200-239]= 3
 [240-279]= 7
 [280-319]= 7
 [320-359]= 5
 [360-399]= 6
 [400-439]= 3
 [440-479]= 5
 [480-519]= 7
 [520-559]= 239
 [560-599]= 3
 [600-639]= 0
 [640-679]= 0
 [680-719]= 1
 [720-759]= 0
 [760-799]= 0
 [800-839]= 0
 [840-879]= 1
 [880-919]= 7
 [920-959]= 0
 [960-999]= 0
 [1000-1039]= 0
 [1040-1079]= 2
 [1080-1119]= 0
 [1120-1159]= 0
 [1160-1199]= 0
 [1200-1239]= 0
 [1240-1279]= 0
 [1280-1319]= 0
 [1320-1359]= 0
 [1360-1399]= 1
 [1400-1439]= 10

OBJECT: N17.10-matrix  Class= matrix-all  [CreationTime: 22:45:47 02-13-89]
 ReadTime: 12:20:44 03-03-89, ClearTime: 12:15:36 03-03-89 (@ -308sec)
 Total Count= 5974 (+0 orphans)
 #bins = 82
[XXX.XXX.0.0 : 35.0.0.0]= 1124 (18.8%) @ -0sec
[XXX.XXX.0.0 : XXX.XXX.0.0]= 976  (16.3%) @ -0sec
[XXX.XXX.0.0 : XXX.XXX.0.0]= 425  ( 7.1%) @ -0sec
```

Packets/Sec vs. Time of Day

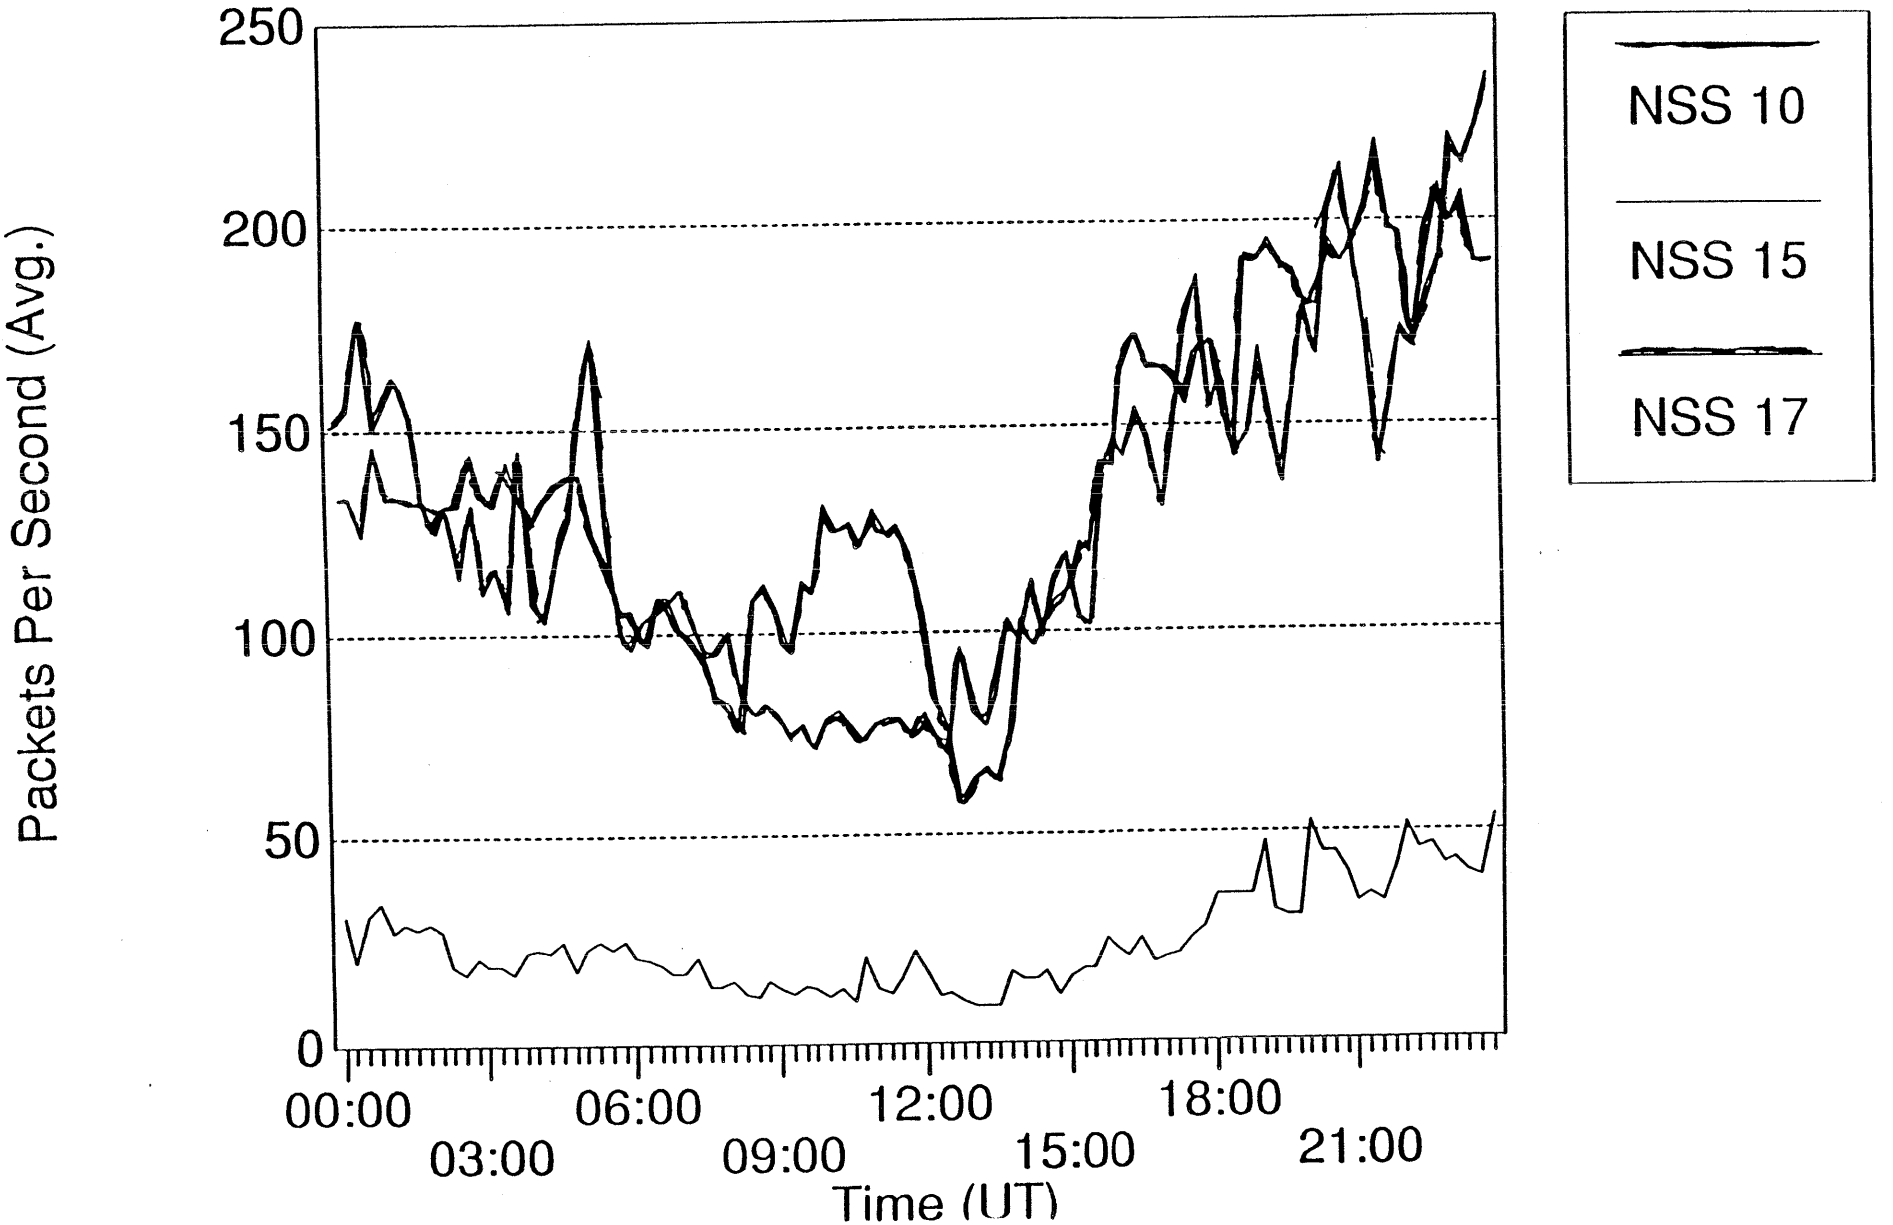

Switching Rate Distribution

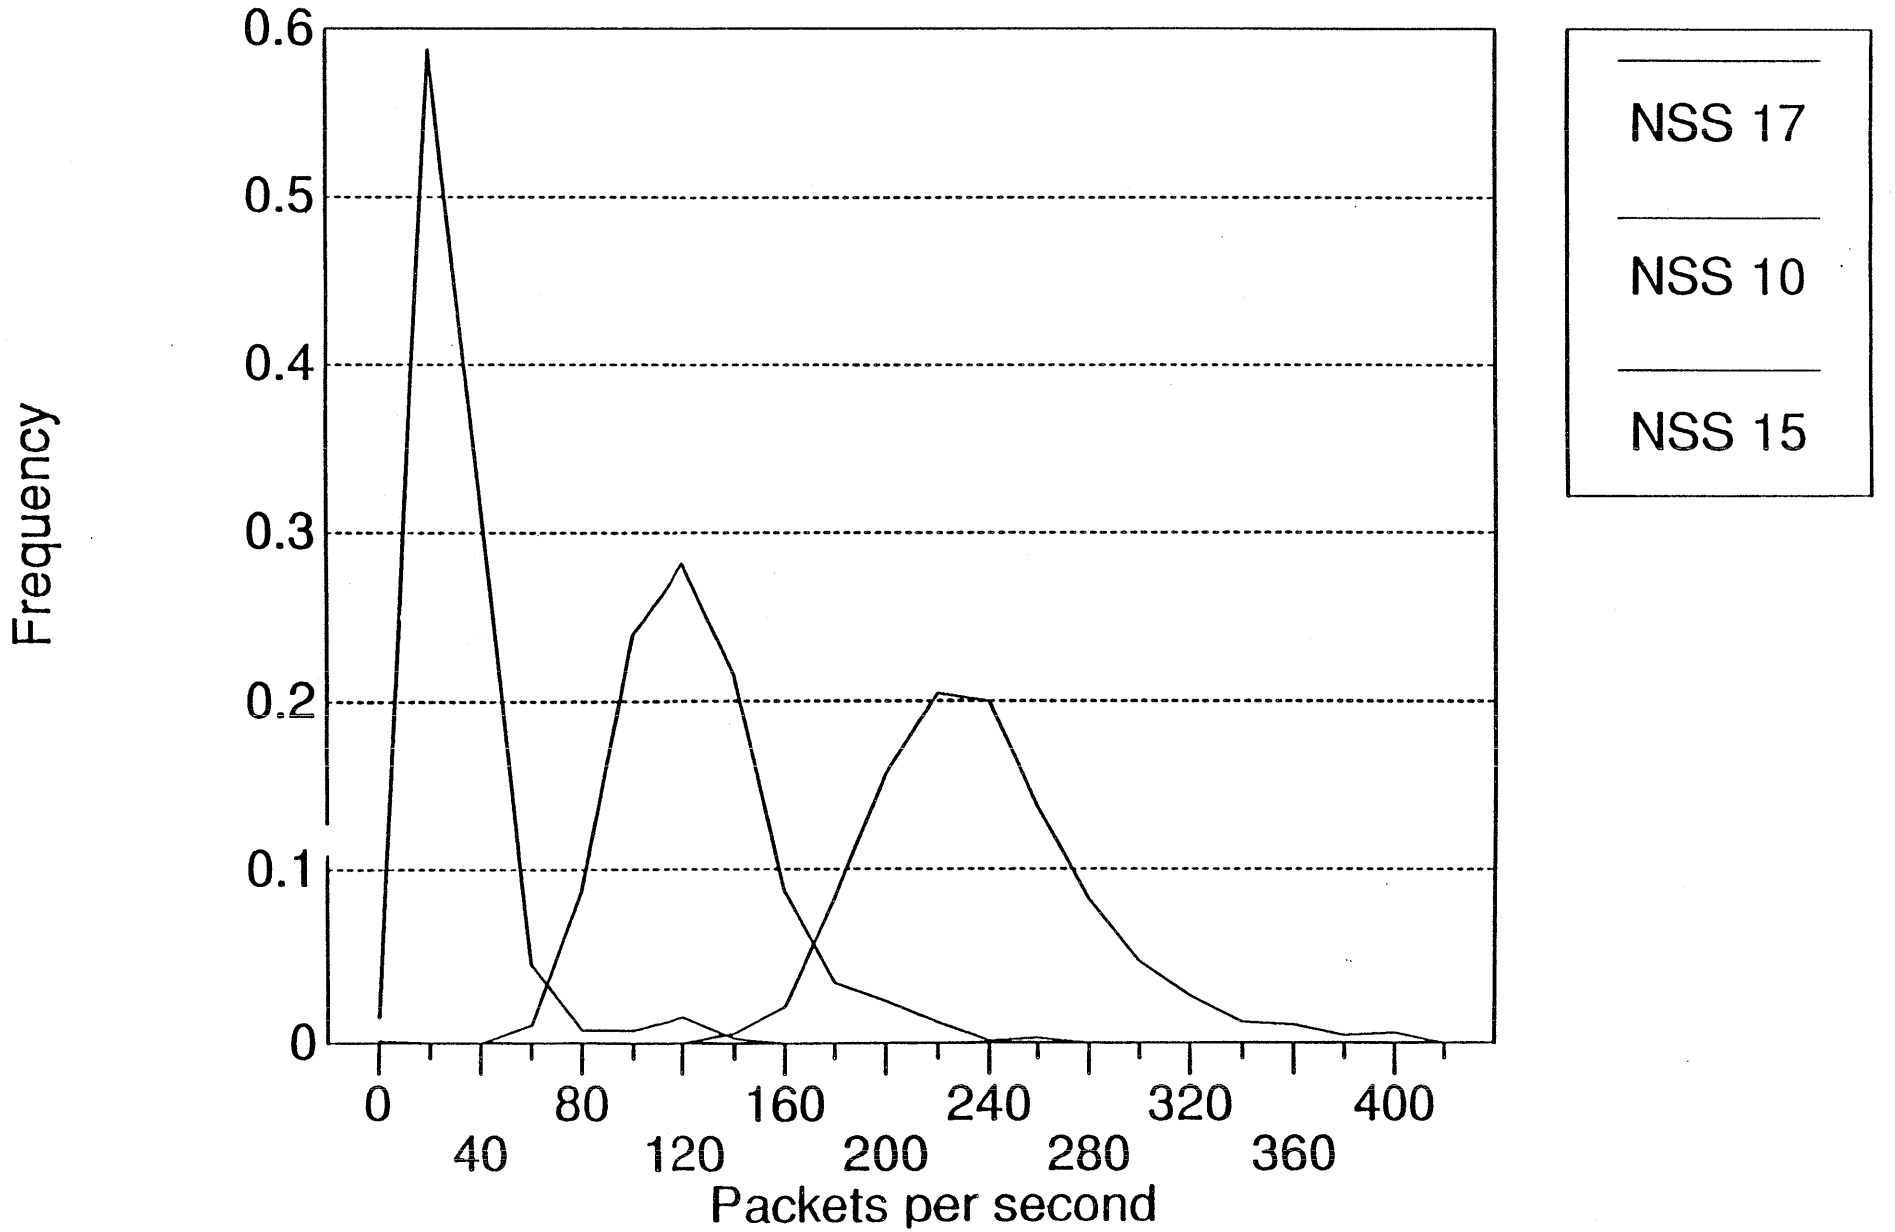

Packet Size Distribution

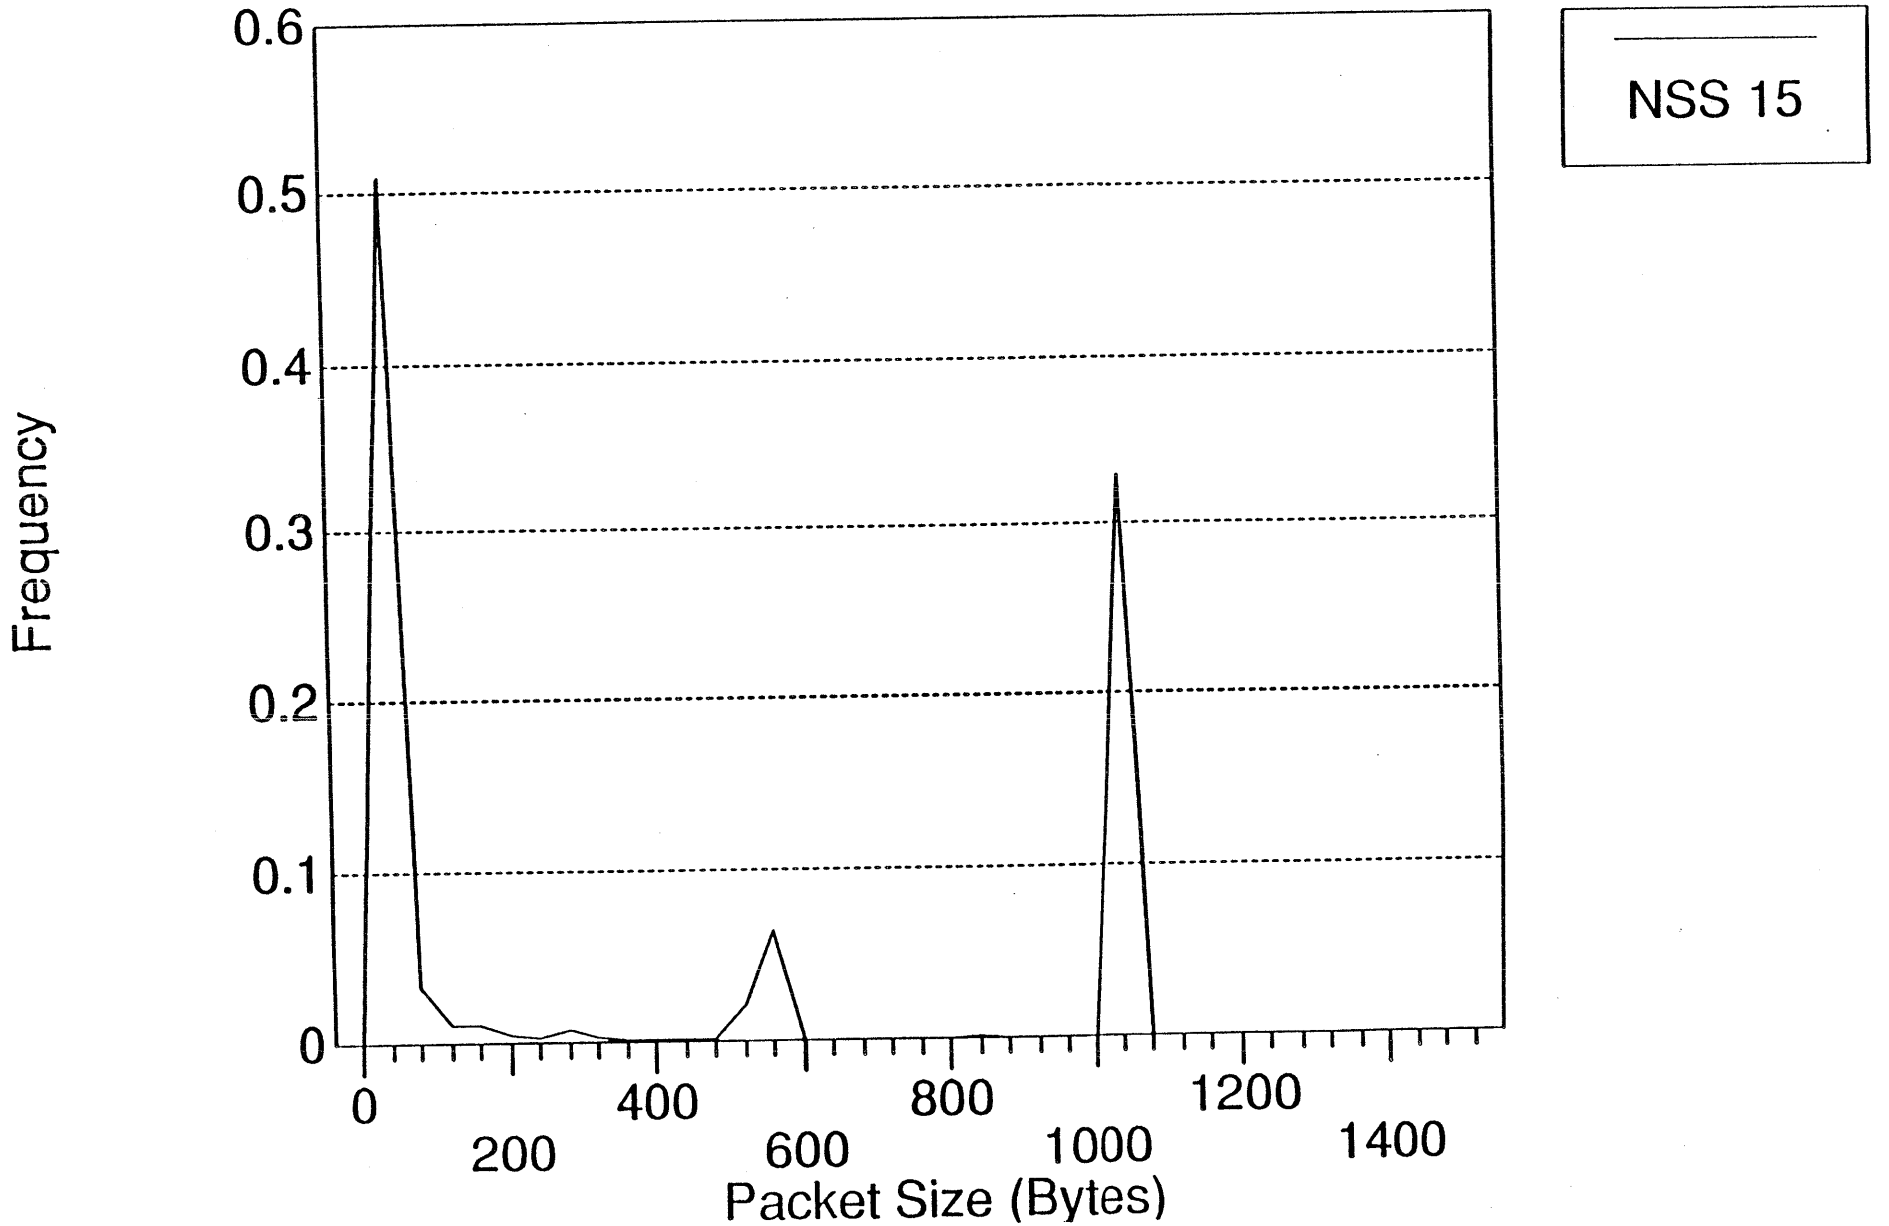

TCP/UDP Port Distribution

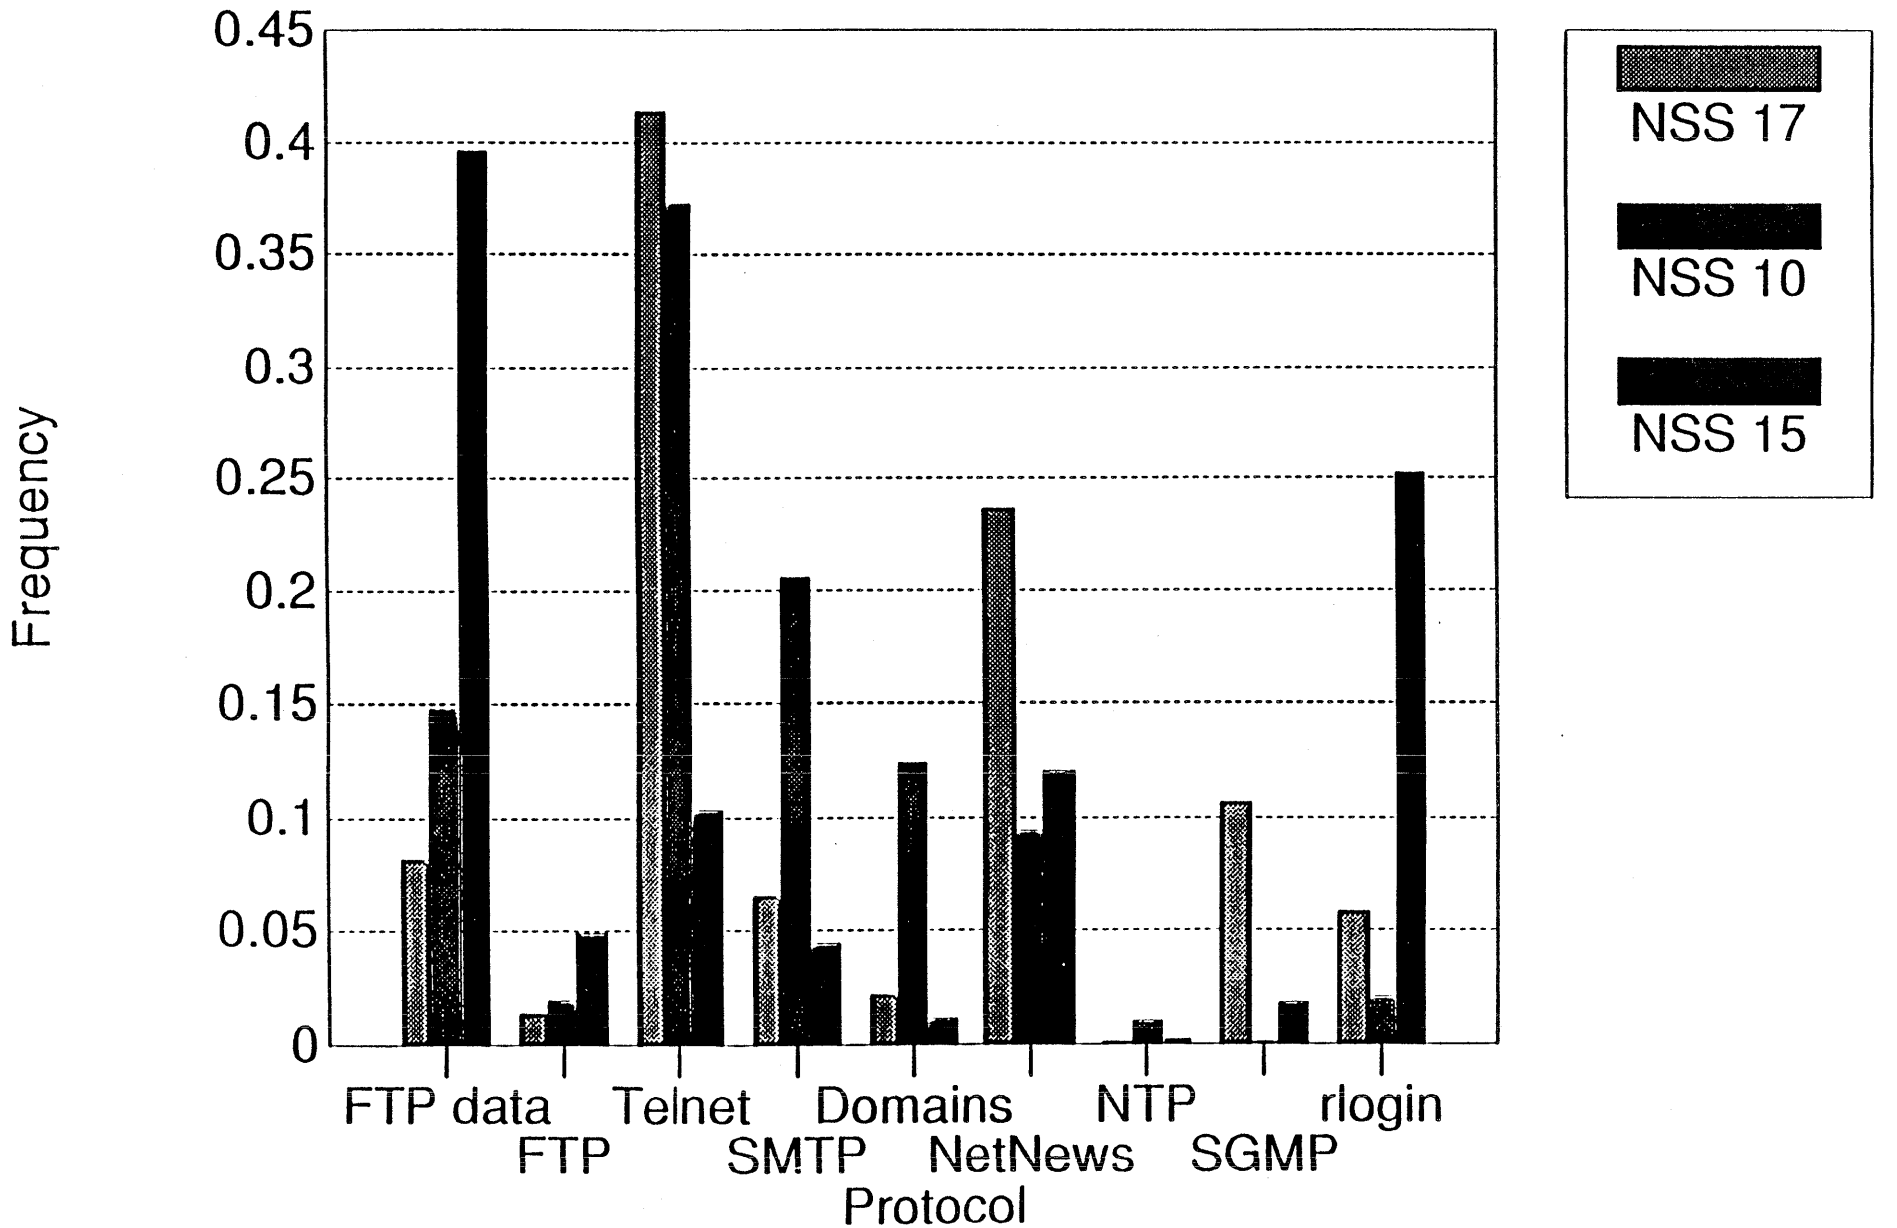

Inter Autonomous System Routing Alternatives
Presented by Yakov Rekhter/IBM

1) Since Inter-Domain Routing is not going to be available in the near future, and since using EGP as an inter-Autonomous system protocol becomes more and more unfeasible, work should be done on the EGP replacement. This is largely due to an increase in complexity for connectivity between Administrative Domains.

2) The IETF IWG group is concentrating on providing this intermediate solution for EGP replacement. Most of the members of IWG are deeply involved with NSFNET as well. Initially the basis for the next Inter-AD routing architecture seemed to be satisfyable by means of the EGP3 protocol. However, during the discussions in the working group it became obvious that the EGP3, as proposed, would only satisfy reachability, but not routing requirements.

3) Using some ideas from the IWG, IBM and CISCO came up with a draft proposal for the EGP replacement - The Border Gateway Protocol (BGP). BGP does not impose any requirements on the IGP within an Administraive Domain or Routing Domain.

4) IBM and CISCO already have an initial implementation of this protocol. Testing is done between NSS-1 (at Merit) and CISCO router (either at Merit or at CISCO). This is done by Jacob Rekhter (IBM, implementing the NSS code), Kirk Lougheed (Cisco, implementing the Cisco version) and Jessica Yu (Merit, testing and verifying interoperability).

5) Jeff Honig at the Cornell University Theory Center is working on a public domain BGP implementation (as part of GATED).

6) A draft paper on BGP is available. Send mail to hwb@merit.edu.

7) We have to make very rapid progress with BGP. It is currently in an experimental stage, but should be moved to an operational stage once the architecture proves feasible.

Optimistic view

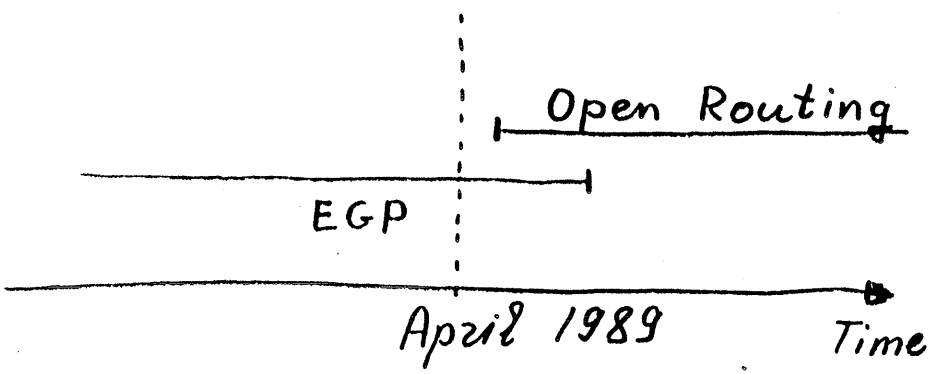

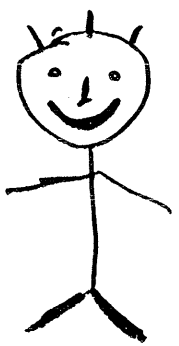

Pessimistic/Realistic View

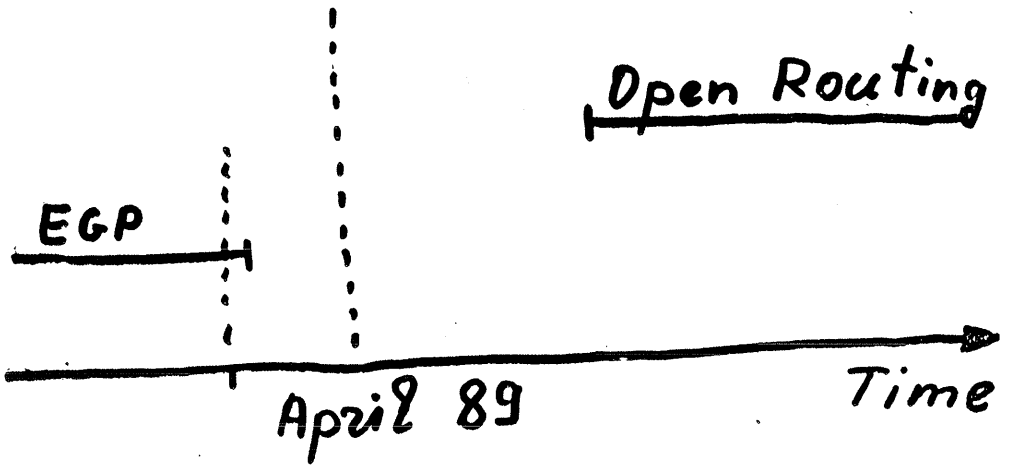

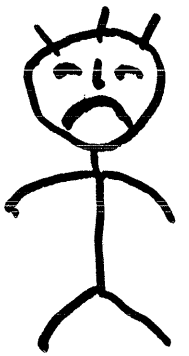

What to do now ?

Solution 1. Sit and wait

Comment:

You are probably waiting for disaster

Solution 2. Quickly come up with interim solution

- Active participation in IWG
- Mid-term Inter-AS Routing Architecture (MIRA)
- Protocol to support MIRA

"Border Gateway Protocol" - BGP

alias

"Route Server Protocol" - RSP

alias

BGP (alias RSP) Status

- First prototype implementations
  - IBM
    - GATED
    - NSS Routing Daemon
  - CISCO
- NSS-1 ⟷ CISCO
  BGP

April 10, 1989


```
From jyy@merit.edu Mon Apr 10 15:05:32 1989
Received: Mon, 10 Apr 89 15:05:27 PDT from merit.edu by kiddo.merit.edu (5.51/1.
6)
Received: Mon, 10 Apr 89 14:03:42 EST from kiddo.merit.edu by merit.edu (5.59/1.
5)
Received: Mon, 10 Apr 89 15:04:05 PDT by kiddo.merit.edu (5.51/1.6)
Date: Mon, 10 Apr 89 15:04:05 PDT
From: Jessica Yu <jyy@merit.edu>
Message-Id: <8904102204.AA04552@kiddo.merit.edu>
To: YAKOV%YKTVMX.BITNET@CUNYVM.CUNY.EDU
Subject: BGP testing
Cc: jyy@merit.edu
Status: RO


Jacob,

Good news!!

Five minutes after you left, I got some routes on the cisco.  And now, the
rcp-1-1 is learning these routes via bgp.

rcp-1-1(37): !n
netstat -Nn
Network             Region ID
192.35.163          178
192.35.162          178
192.35.165          182
192.35.164          182
192.35.169          179
128.104             231
130.126             231

(The two routes from AS 231 are from the cisco).

Anyway, I thought you'd like to know before going to the IETF meeting.  Kirk
has not called me back yet.

                                                        --Jessica
```

Future plans

- Public Domain BGP
  (as part of GATED)

- More experience

Requiem for the Arpanet

Vinton G. Cerf

Requiem for the ARPANET

Vint Cerf

Like distant islands sundered by the sea,
We had no sense of one community.
We lived and worked apart and rarely knew
that others searched with us for knowledge, too.

Distant ARPA spurred us in our quest
and for our part we worked and put to test
new thoughts and theories of computing art;
we deemed it science not, but made a start.

Each time a new machine was built and sold,
we'd add it to our list of needs and told
our source of funds "Alas! Our knowledge loom
will halt 'til it's in our computer room."

Even ARPA with its vast resources
could not buy us all new teams of horses
every year with which to run the race.
Not even ARPA could keep up that pace!

But, could these new resources not be shared?
Let links be built; machines and men be paired!
Let distance be no barrier! They set
that goal: design and build the ARPANET!

As so it was in nineteen sixty-nine,
a net arose of BBN design.
No circuit switches these, nor net complete
but something new: a packet switching fleet.

The first node occupied UCLA
where protocols and measurement would play
a major role in shaping how the net
would rise to meet the challenges unmet.

The second node, the NIC, was soon installed.
The Network Info Center, it was called.
Hosts and users, services were touted:
to the NIC was network knowledge routed.

Nodes three and four soon joined the other two:
UCSB and UTAH come on cue.
To monitor it all around the clock
at BBN, they built and ran the NOC.

A protocol was built for host-to-host
communication. Running coast-to-coast,
below the TELNET and the FTP,
we called this protocol the NCP.

The big surprise for most of us, although
some said they guessed, was another proto-
col used more than all the rest to shuttle
mail in content flaming or most subtle.

When we convened the first I Triple C,
the ARPANET was shown for all to see.
A watershed in packet switching art,
this demo played an overwhelming part.

Within three years the net had grown so large
we had to ask that DCA take charge
to operate a system guaranteed
for R&D and military need.

Exploring other packet switching modes,
we built the first spread spectrum mobile nodes.
The Packet Radio, the mobile net,
worked on the ground and even in a jet.

Deployed at SAC and Eighteenth Airborne Corps,
the Packet Radio unlocked the door
to what we now know as the Internet.
The driver for it all was PRNET.

The Packet Satellite, another new
technique, was added to the net milieu.
And then to shed more light upon the dark,
there came the Ethernet from Xerox PARC.

To these we added yet another thing
from MIT: a local token ring.
We saw the local net techniques compound
until the list could easily confound.

The Internet foundation thus was laid.
Its protocols from many sources made.
And through it all the ARPANET grew more;
It was, for Internet, the central core.

The hardware of the net was changing, too.
The Honeywell was first, and then the SUE,
which forms the heart of Pluribus today
though where this platform sits one cannot say.

The next big change was called the MBB.
It emulated Honeywell, you see,
so one by one they modified each node,
by means of closely written microcode.

Now known as 30 prefixed with a C,
these nodes are everywhere from A to Z.
The European MINET too was full
of nodes like these from Mons to Instanbul.

The second Autodin was long desired
but once accepted instantly expired.
Then to the rescue rode the ARPANET!
And soon the MILNET by its side was set.

By Nineteen-Eighty DoD opined
its data networks soon must be aligned
with Internetwork protocols, to wit:
by Eighty-Three the TCP was IT!

Soon every host that sat on ARPANET
became a gateway to a local net.
By Eighty-Six new long haul nets appeared
as ARPANET its second decade neared.

The NSFNET and its entourage
began a stately national dressage
and soon was galloping at T1 speed
outdistancing its aging peer indeed.

And so, at last, we knew its course had run,
our faithful servant, ARPANET, was done.
It was the first, and being first, was best,
but now we lay it down to ever rest.

Now pause with me a moment, shed some tears.
For auld lang syne, for love, for years and years
of faithful service, duty done, I weep.
Lay down thy packet, now, O friend, and sleep.

ARPANET Geographic Map After MAY 2, 1989

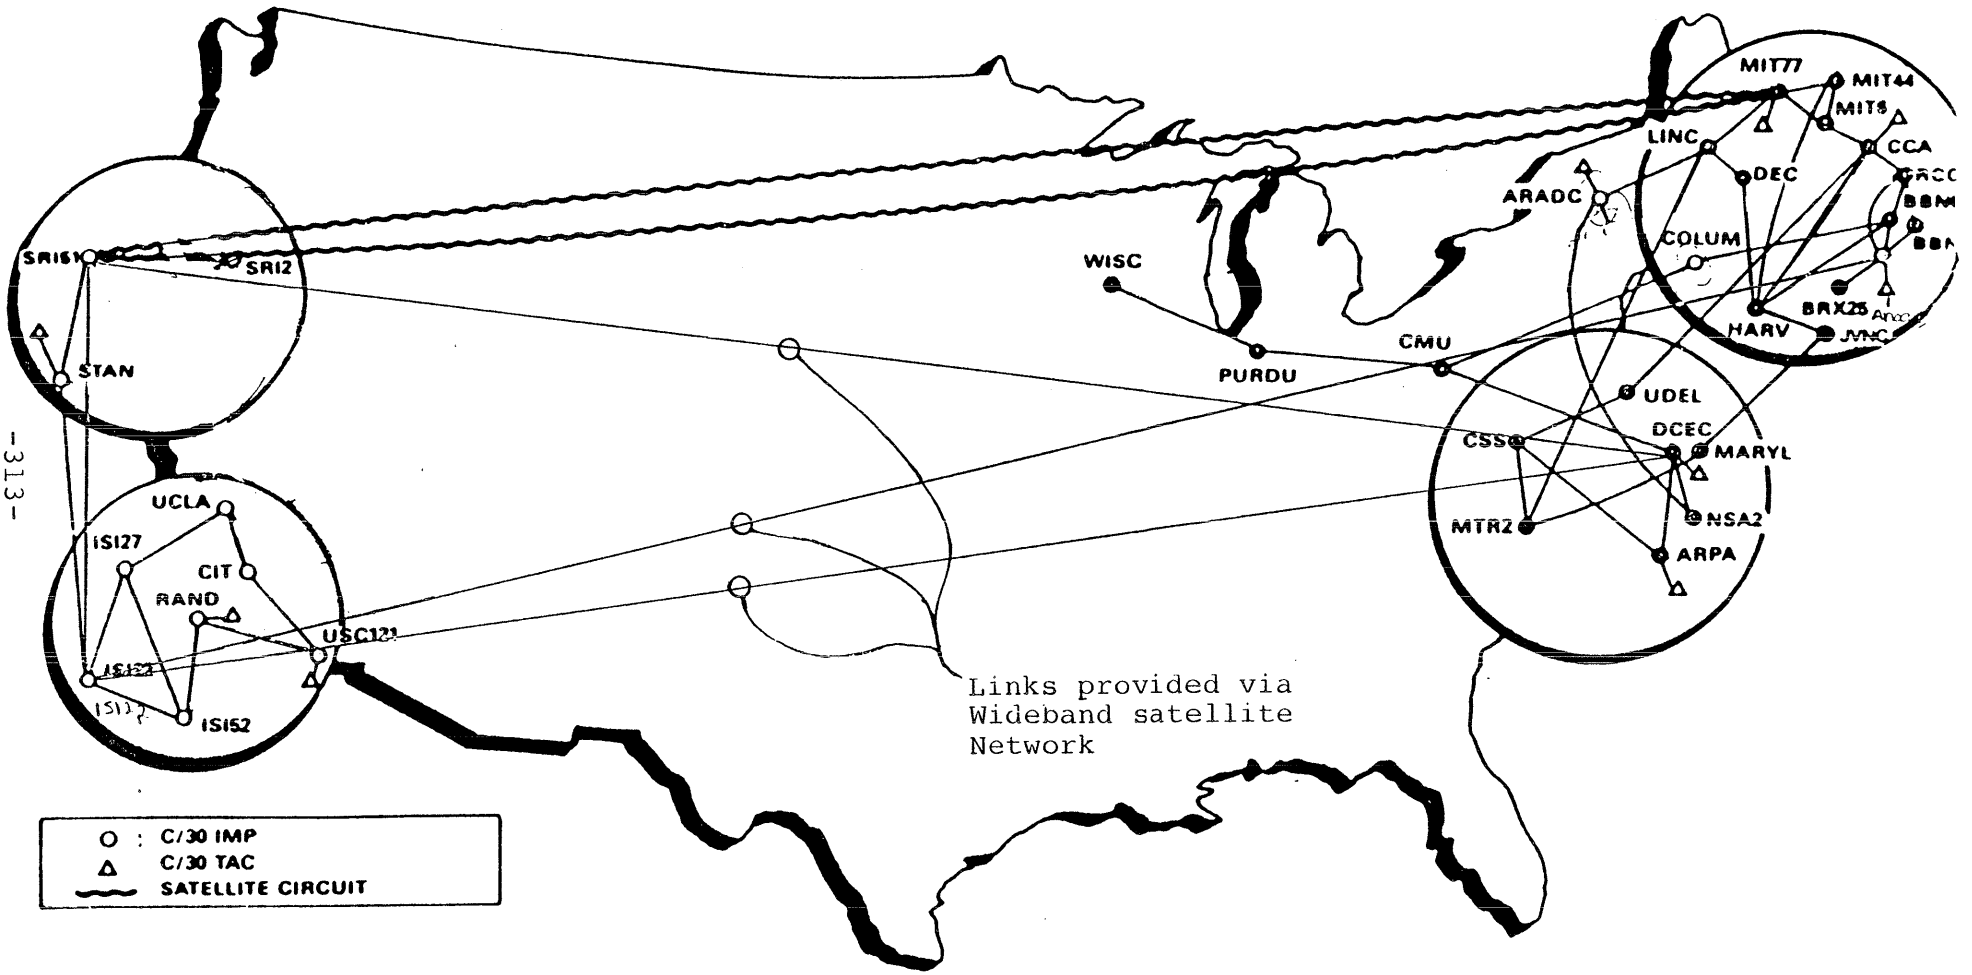

ARPANET Geographic Map After JUNE 2, 1989

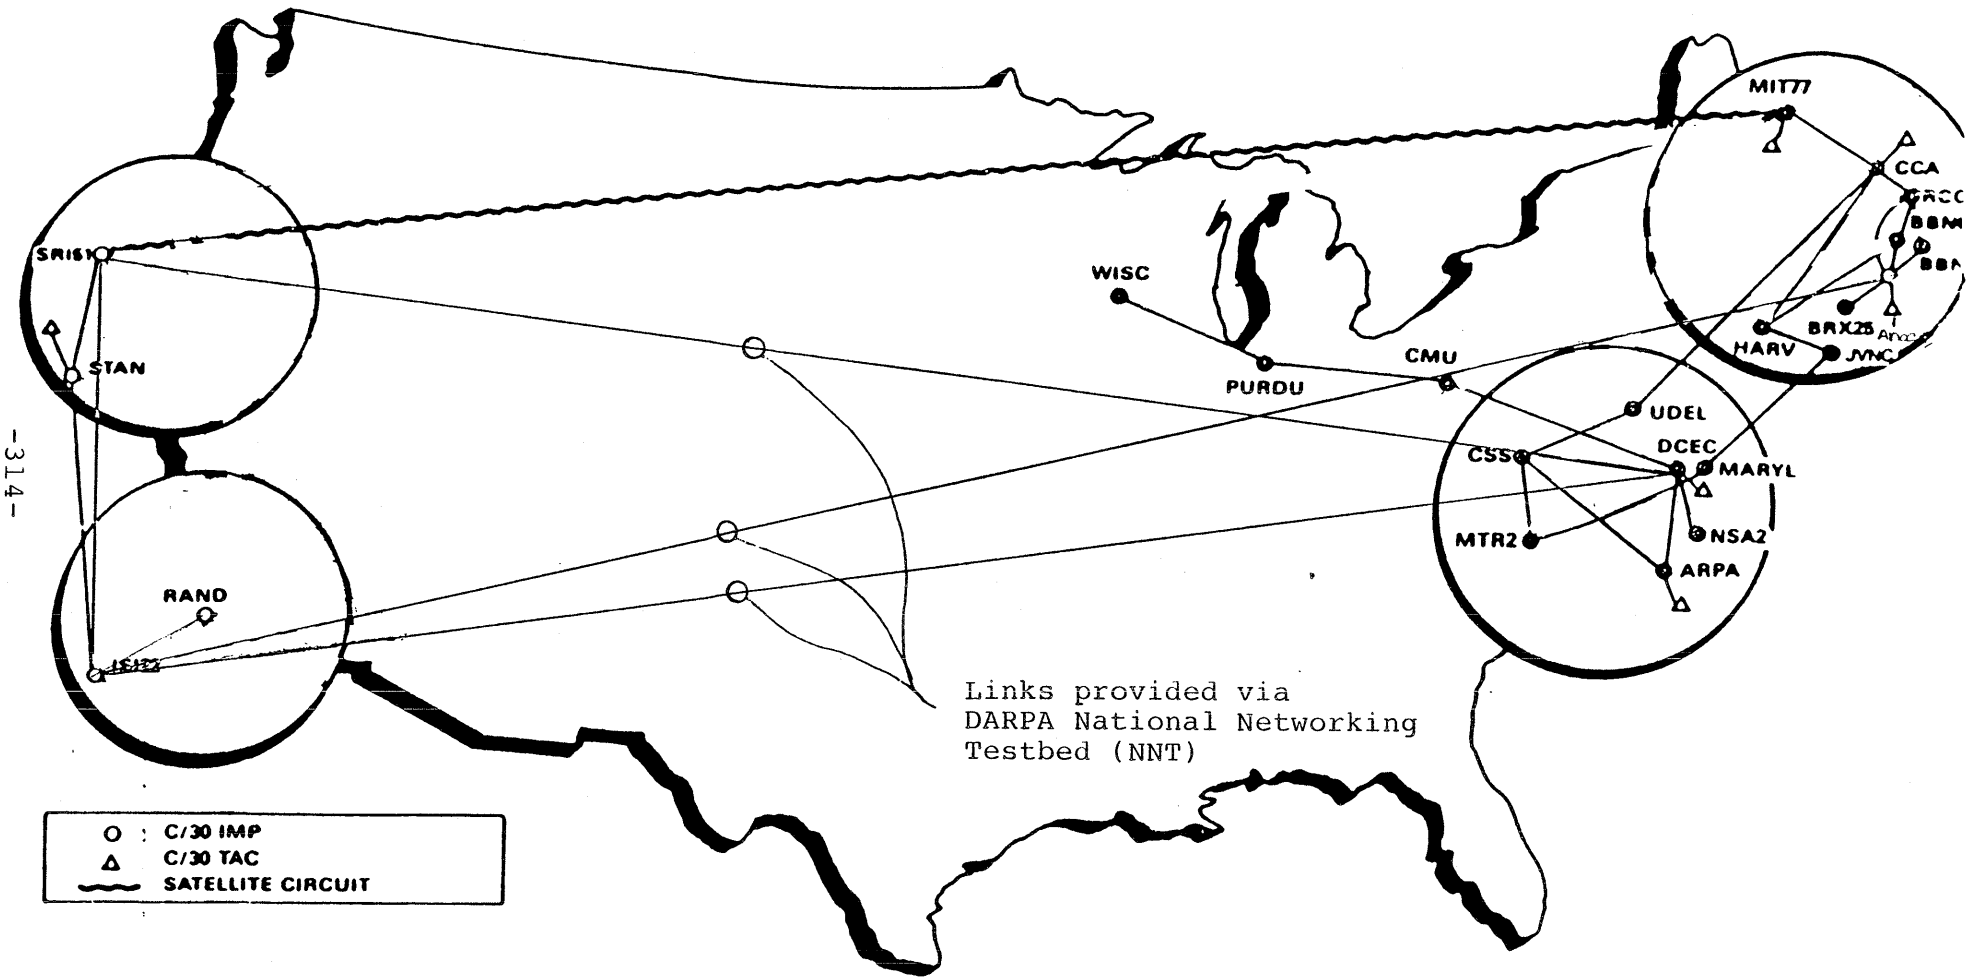

ARPANET Geographic Map After AUGUST 2, 1989

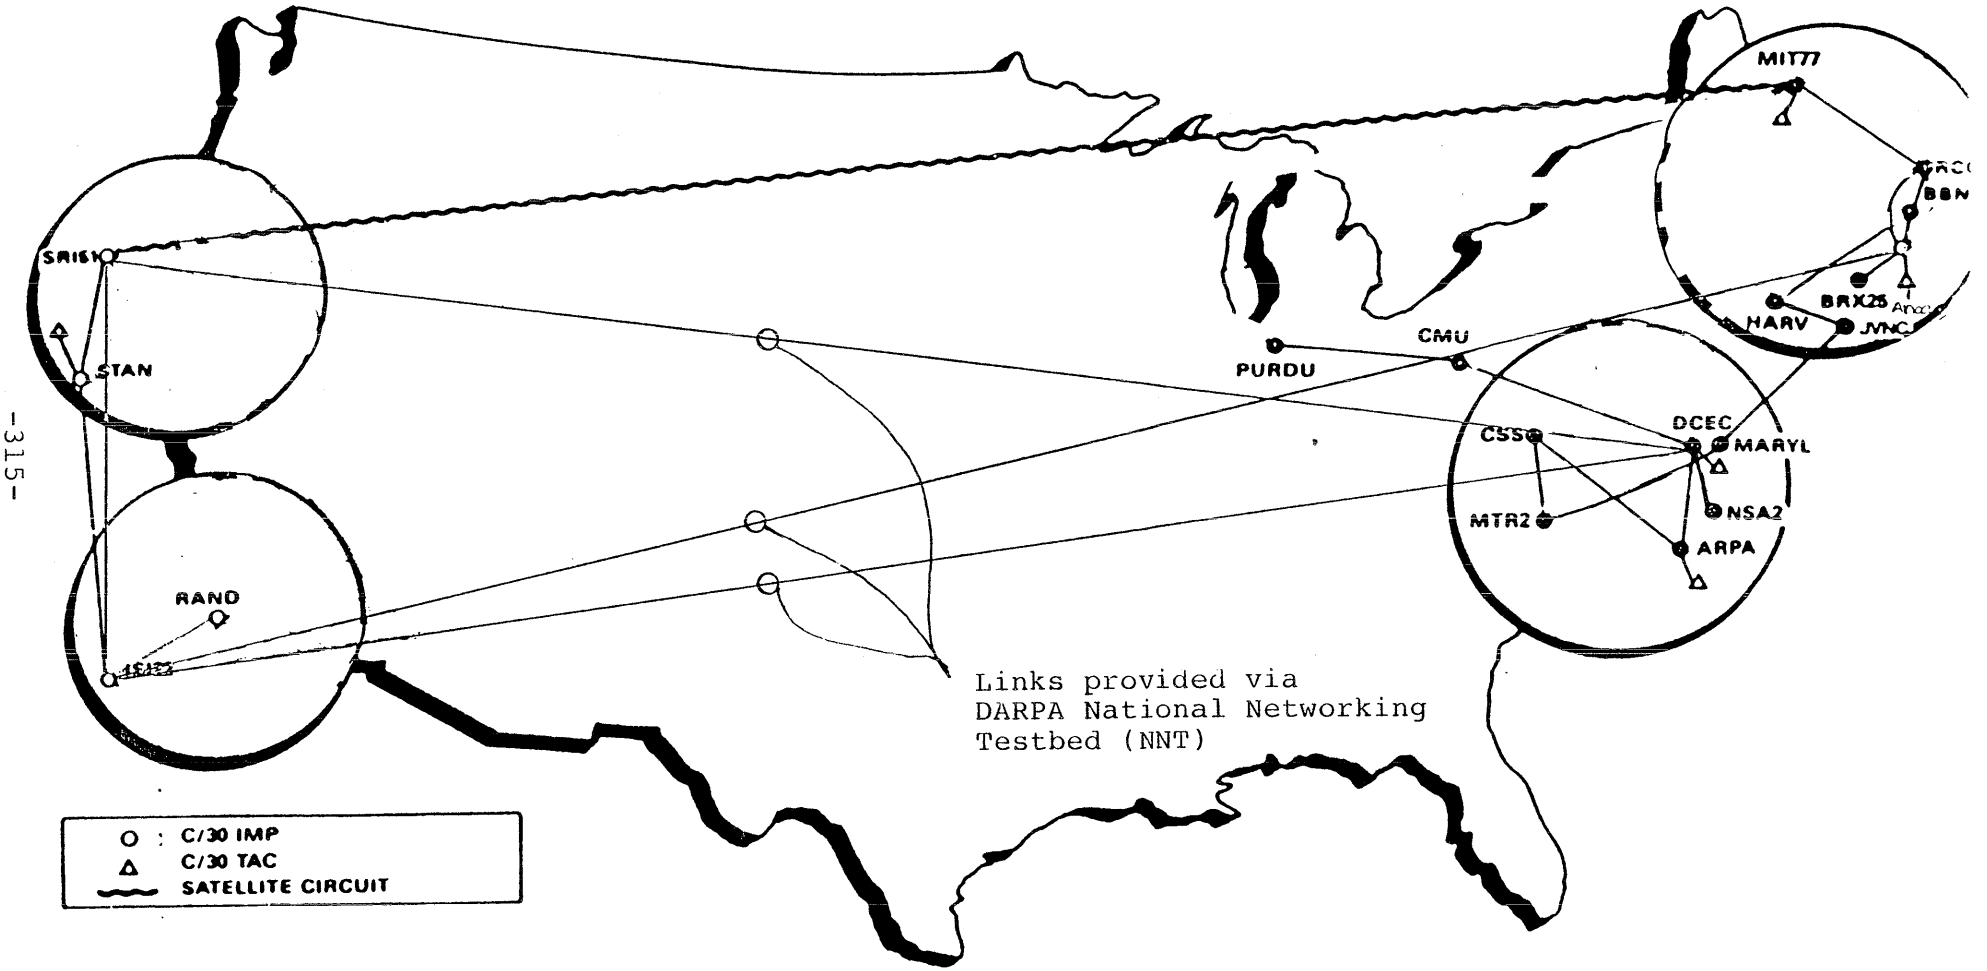

ARPANET Geographic Map After OCTOBER 2, 1989

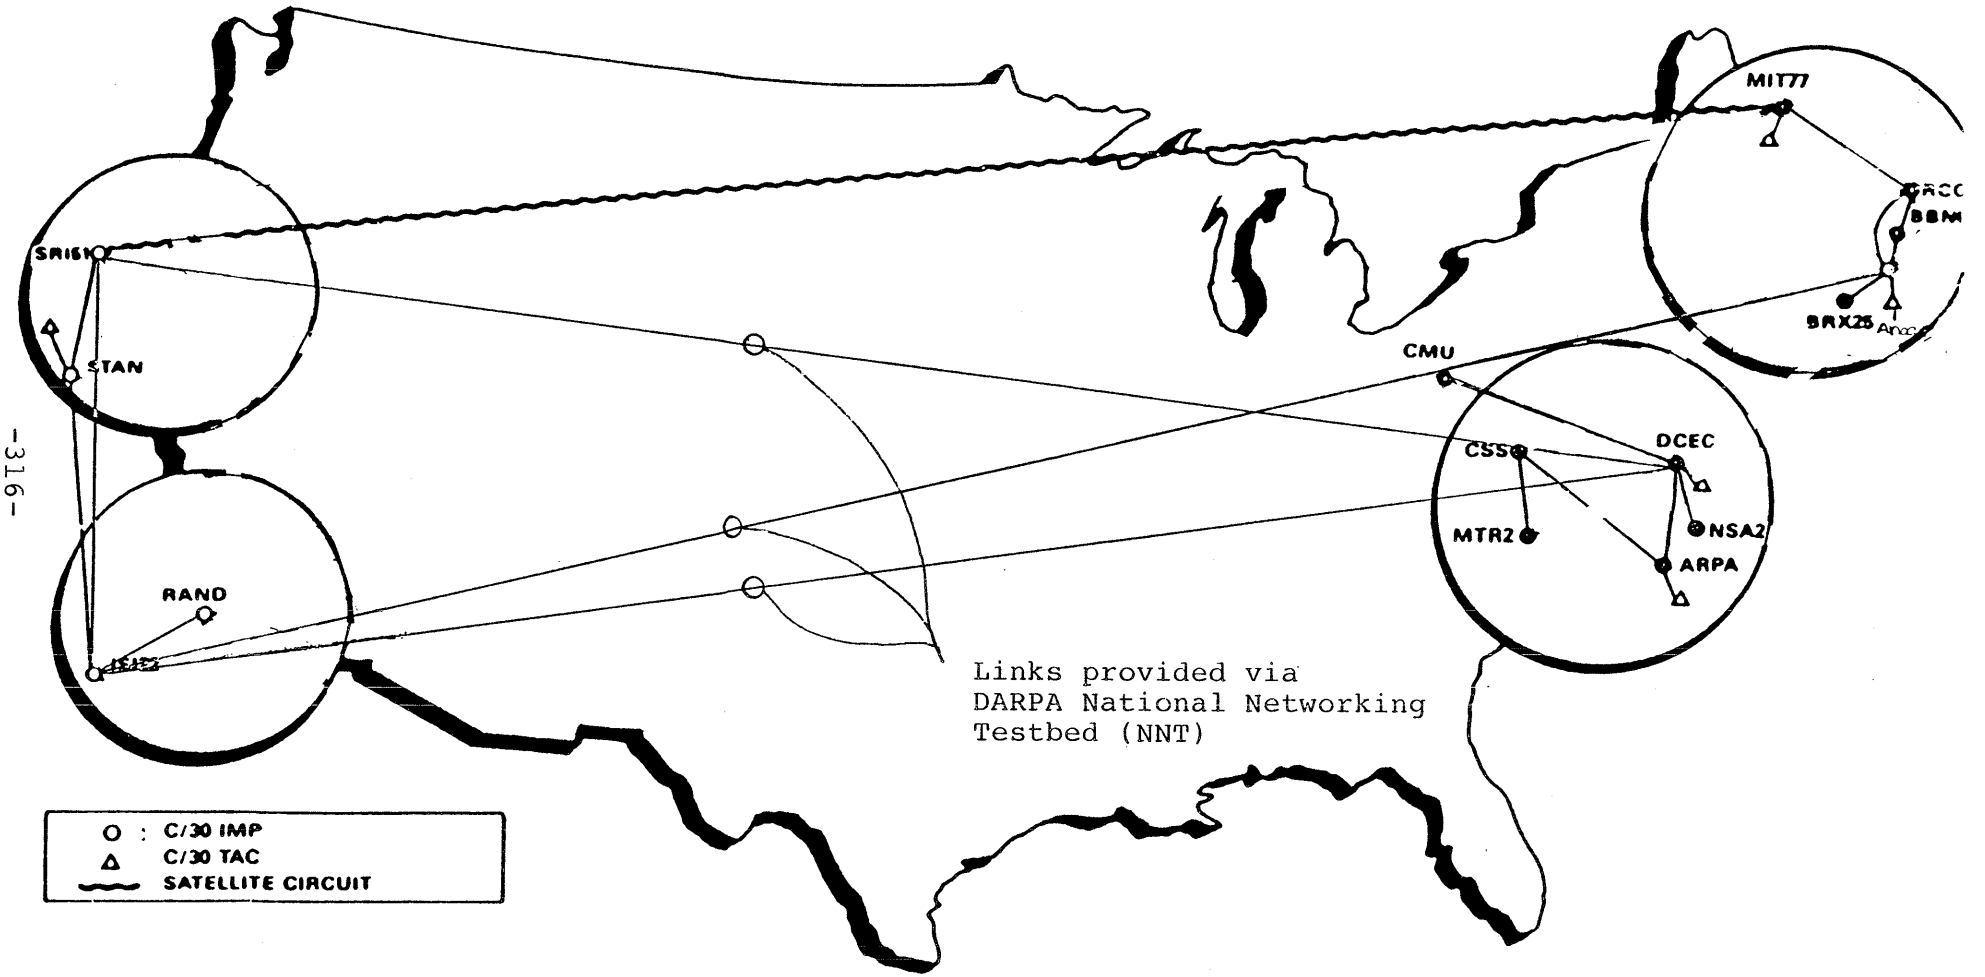

ARPANET Geographic Map After JANUARY 2, 1990

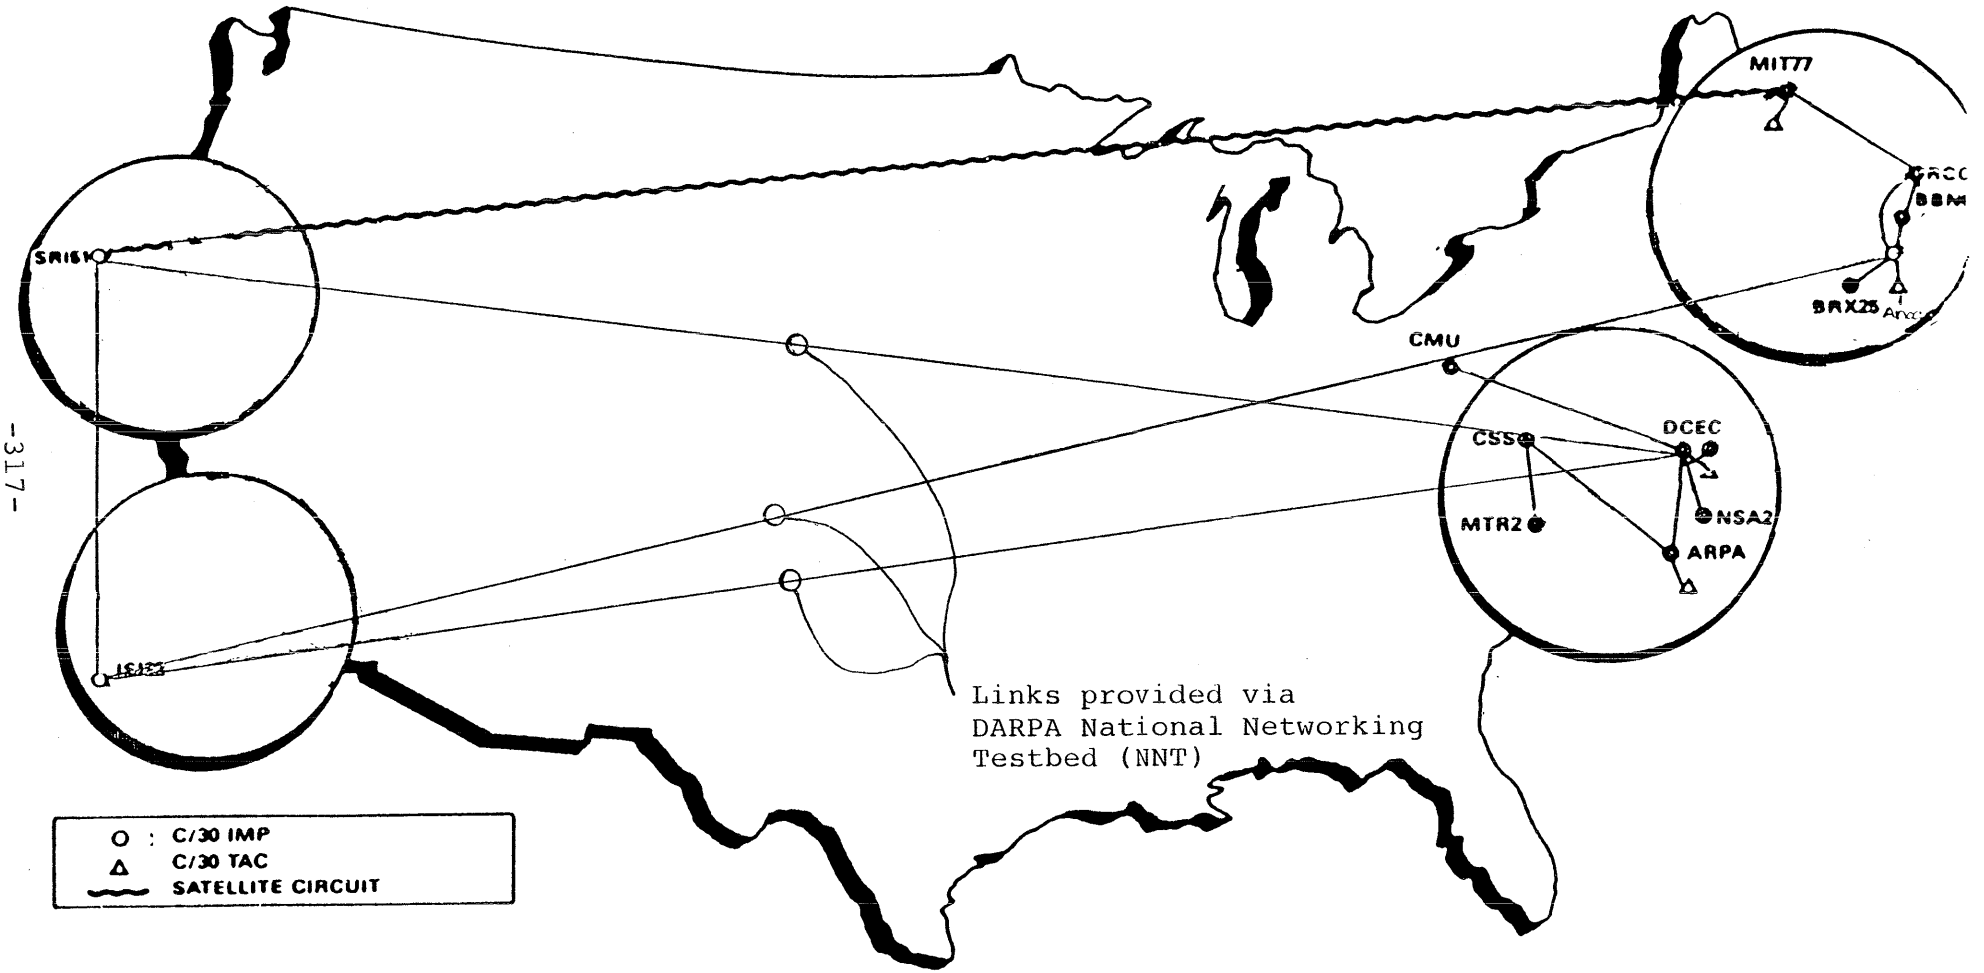

ARPANET Geographic Map After MARCH 2, 1990

Mailbridge Access Control

Marianne Lepp

WHY MAILBRIDGE ACCESS CONTROL

- Control Access to MILNET Resources
- Ability to Turn off "Undesired" Traffic in the Event of a "Break In"

WHAT IS A MAILBRIDGE?

- Full Function IP Gateway (Router)
- Plus
  - Access Control
  - Load Sharing
- Services
  - Gateway Between MILNET and ARPANET
  - Provides EGP Routing Services
  - Controls Access to MILNET

DDN MAILBRIDGE ACCESS CONTROL

Marianne Lepp

BBN Communications Corporation

MAILBRIDGE HISTORY

- 1983 - ARPANET Split Planned
  Mailbridge to be Electronic Mail Forwarding Host
- 1984 - Mailbridge Gateway (Release 1007)
  Original Access Control
- 1988 - New Mailbridge Gateway (Release 1008)
  Improved Access Control

ORIGINAL MAILBRIDGE ACCESS CONTROL SCHEME

- Enable / Disable
- If Enabled
  - All Electronic Mail (SMTP) Traffic
  - All Protocols if in Community of Interest (COI)
- Status
  - Installed in Original Mailbridges
  - Currently Disabled

IMPROVED MAILBRIDGE ACCESS CONTROL

- Arbitrary Decision Tree
- Questions Include
  - Source Address
  - Source Network
  - Destination Network
  - Previous Hop Network
  - Protocols
  - User or Server
  - Community of Interest
- Uses Cache for Performance

ACCESS CONTROL EXAMPLE 1

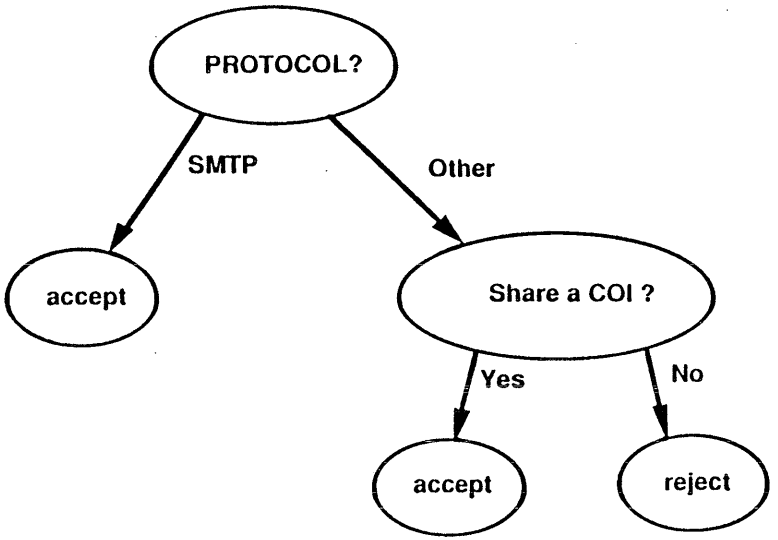

ACCESS CONTROL EXAMPLE 2

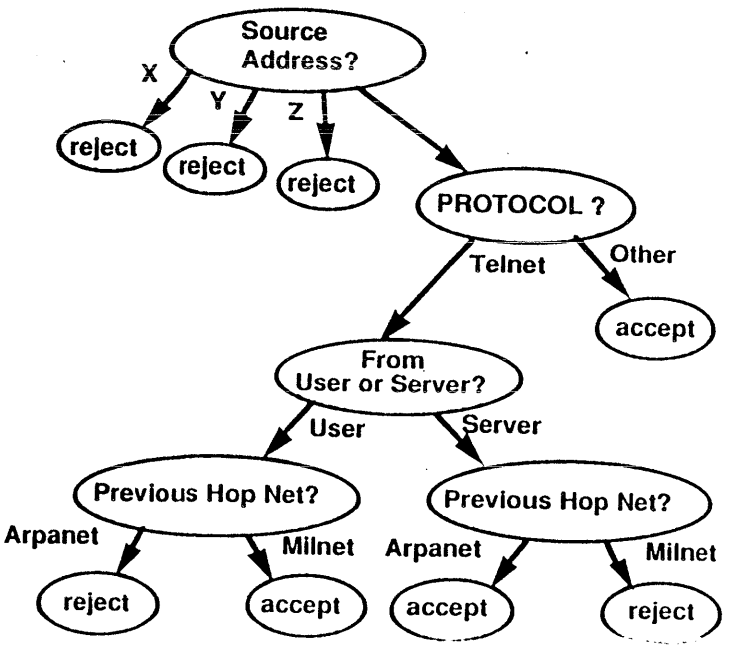

LESSONS LEARNED

- Hard to Anticipate Access Control Requirements
- Access Control Requirements Change Over Time
- Access Control Mechanisms Should be Flexible

The DCA TCP/IP Certification Program
Presented by Martin Gross/DCA/DCEC

In an effort to ensure compliance with its Military Standard High Level Data Communications Protocols (IP,TCP,FTP,SMTP,TELNET) and increase the probability of interoperability in its diverse multi-vendor environment, the Department of Defense (DoD) has initiated a program to certify vendor implementations of these products. As Executive Agent for the DoD Data Communications Protocols, the Defense Communications Agency (DCA) has been tasked with implementing this program.

The policy for high level protocol conformance testing was established in a memorandum from the Assistant Secretary of Defense for Command, Control, Communications and Intelligence dated 26 August 1988. The memorandum mandates conformance testing on all new contracts executed after 1 June 1989. Products procured under contract must be tested by a National Institute of Standards and Technology (NIST) accredited laboratory prior to operational use on any DoD network. The memorandum also establishes a Qualified Products List which will be maintained by DCA. For a product to be placed on the Qualified Products List, acceptable tests results must be presented to DCA from an accredited laboratory which is independent of the vendor.

DCA started an in-house testing program for DDN X.25 in 1983. Due to limited resources however, this program could not be continued in-house nor could a high level protocol test program be developed in the same manner. For this reason DCA turned to the National Voluntary Laboratory Accreditation Program (NVLAP) run by NIST. Under NVLAP, laboratories are recognized and accredited to perform specific testing services aimed at evaluating products to determine if they meet applicable standards. As the program's name specifies, this is a voluntary program and NVLAP accreditation does not imply the certification of products or test data. In July 1988, DCA requested that NIST establish a NVLAP for the DoD Protocols (DDN X.25 and the five High Level Protocols). Formal establishment of the program was announced in the Federal Register on 21 July 1988. The X.25 Program was established first and there are currently three accredited laboratories. The NVLAP for the High Level Protocols is now being developed.

The NVLAP Handbook for the High Level Protocols which presents the operational and technical requirements for an accredited laboratory was published in draft form on 22 March 1989. The document was mailed to all those who replied to the Federal Register Announcement and was open to public comment until the 14 April. Laboratory applications are now being accepted by NIST and the first laboratory will be accredited by the middle of June. It is expected that there will be a minimum of three accredited laboratories.

The laboratories will be required to use the DCA Upper Level Protocol Test System to perform the testing service. The system was developed under DCA contract to provide a standard testing capability for the DoD High Level Protocols. The system tests protocol functionality including; upper layer interfaces, validity of outputs, and input error handling. (See Connexions Volume 2, Number 8 for further details) The test system operates on a VAX* cpu running Ultrix* 1.1 and is publicly available from the National Technical Information Service. The system has been in use since December of 1987 and is currently in use by ten organizations for in-house testing.

To clarify testing policies and provide guidance to vendors, DCA will publish a DoD High Level Protocols Testing Circular. The circular will establish specific testing policies relating to the testing of products across hardware lines and the retesting of modified products. The circular will also establish the procedure for placement of products on the Qualified Products List. The DoD Protocol Conformance Testing Profile will also be included in the circular. This Profile establishes the set of mandatory features that must be implemented in each protocol. It also indicates those features which are optional. This Profile has been approved by the DoD's Protocol Standards Steering Group but is still available from DCA for public comment. The testing circular will be published in draft form by 1 June.

To help vendors prepare their products for laboratory certification, DCA has installed a test system that can be used by vendors. This system will be available for the next nine months on a first come first served basis. The system is accesible through the Internet or a dial-up link and is available for self testing with no on-line support. Information is provided below on how to obtain further information on this program.

Comments or questions relating to protocol testing can be addressed to:

Martin Gross
DCA Code R640
1860 Wiehle Ave.
Reston, VA 22090-5500 or
by email: martin@edn-unix.dca.mil or martin@protolaba.dca.mil.

For NVLAP information or documents contact: Jeff Horlick, NVLAP, NIST, Bldg 411, Gaithersburg, MD 20899, (301)975-4016.

For the DCA Upper Level Protocol Test System and Documents contact:

National Technical Information Service
Springfield, VA 22161
(703)487-4807. Product #AD-A204-558.

*VAX and Ultrix are Trademarks of the Digital Equipment Corporation.

DoD CONFORMANCE TESTING

PURPOSE

- Conformance to Military Standards
- Benefit Interoperability in Multi-vendor environment
- Aid in functioning of dual DoD/OSI components to be used in DoD

DCA UPPER LEVEL PROTOCOL TEST SYSTEM

- Developed by UNISYS under DCA contract
- Tests Protocol functionality
    - Upper Layer Interface
    - Validity of Output - Sending
    - Input Error Handling - Receiving
- System in use since Dec. 1987

DCA UPPER LEVEL PROTOCOL TEST SYSTEM

- Available from National Technical Information Service
  - Initial Release June 1988
  - Second Release 10 January 1989
- Ultrix 1.1 Available from DCA
- Ten sites using system for in-house testing

NATIONAL VOLUNTARY LABORATORY ACCREDITATION PROGRAM (NVLAP)

- Run by National Institute for Standards and Technology (NIST)
- Accredit laboratories to perform testing services
- Program established for DoD protocols in conjunction with DCA
  - Announced in Federal Register 21 July 1988
    - X.25
    - High Level Protocols (IP,TCP,FTP,SMTP,TELNET)

NVLAP for DoD High Level Protocols

- Operations and Technical Requirements Document
  - Published by NIST 22 March 1989
  - Mailed to those who responded to Federal Register Announcement
  - Initial comment period ends 14 April
- Accept Applicants starting 17 April
- Accredit first Lab Mid May
- Minimum 3 Labs

DoD Conformace Testing Profile

- Establishes Minimun Set of Mandatory Features
- Indicates Options That Can Be Implemented
- Approved by DoD Protocol Standards Group
- Still Available for Public Comment til June 89

OPEN LAB FOR VENDORS

- Help Vendors Prepare Products for Laboratory Certification
- Available for 9 months
- Test System at DCA
  - Dial-up or Network Access
  - Self Testing
  - No On-Line Support
  - First Come First Serve Scheduling

DoD Testing Policy

- Established by ASD Memo dated 26 August 1988
- Mandates conformance testing on contracts executed after 1 June 1989
- Tests must be performed by an independent NIST accreditated laboratory
- Product modifications require retesting
- Qualified Products List maintained by DCA

DoD TESTING CIRCULAR

- Specific Testing Policies and Requirements
- Provide guidance to services and agencies
- DoD Protocol Profile
- Published in Draft form 1 May 1989

COMMENTS or QUESTIONS

- Martin Gross
  DCA Code R640
  1860 Wiehle Ave.
  Reston, VA 22090-5500

- martin@edn-unix.dca.mil

  or

  martin@protolaba.dca.mil

DCA UPPER LEVEL PROTOCOL TEST SYSTEM SOFTWARE

Available from: U.S. Department of Commerce
National Technical Information Service (NTIS)
Springfield, VA 22161
(703)487-4807

NTIS Accession # AD-A204-558

Price 625.00

Media 1/2 inch tape; 1600 bpi; Unix tar format
Includes all documentation

DCA UPPER LEVEL PROTOCOL TEST SYSTEM DOCUMENTS

	NTIS Accession #	Price
1. Functional Description	AD-A195-129	12.95
2. Installation and Operations Manual	AD-A195-130	12.95
3. Test Operators Manual	AD-A195-131	12.95
4. Internet Protocol Mil-Std 1777 Remote Driver Specification	AD-A195-133	12.95
5. Transmission Control Protocol Mil-Std 1778 Remote Driver Specification	AD-A195-135	14.95
6. File Transfer Protocol Mil-Std 1780 Remote Driver Specification	AD-A195-137	12.95
7. Simple Mail Transfer Protocol Mil-Std 1781 Remote Driver Specification	AD-A195-142	12.95
8. TELNET Protocol Mil-Std 1782 Remote Driver Specification	AD-A195-138	12.95
9. Internet Protocol Mil-Std 1777 Test Traceability Index	AD-A195-132	19.95
10. Transmission Control Protocol Mil-Std 1778 Test Traceability Index	AD-A195-134	12.95
11. File Transfer Protocol Mil-Std 1780 Test Traceability Index	AD-A195-136	12.95
12. Simple Mail Transfer Protocol Mil-Std 1781 Test Traceability Index	AD-A195-140	12.95
13. TELNET Protocol Mil-Std 1782 Test Traceability Index	AD-A195-139	12.95
14. Transmission Control Protocol/ Internet Protocol (Tightly Coupled) Test Traceability Index	AD-A195-143	14.95
15. Internet Protocol Security Option Test Traceability Index	AD-A195-141	9.95

Note: All prices are for paper copy. Prices are 6.95 each for documents on microfiche.

National Bureau of Standards

[Docket No. 80741-8141]

National Voluntary Laboratory Accreditation Program

AGENCY: National Bureau of Standards, Commerce.

ACTION: Notice of formal establishment of a laboratory accreditation program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense.

SUMMARY: Under the National Voluntary Laboratory Accreditation Program (NVLAP), the National Bureau of Standards (NBS) announces the establishment of a laboratory accreditation program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense (Protocols Program). Laboratories that are interested in becoming accredited under the Protocols Program may indicate their interest in the program by informing the Manager, Laboratory Accreditation, National Bureau of Standards of their specific interests.

FOR FURTHER INFORMATION CONTACT: John L. Donaldson, Manager, Laboratory Accreditation, National Bureau of Standards, Admin A527, Gaithersburg, MD 20899, (301) 975-4016.

SUPPLEMENTARY INFORMATION:

Background

This notice is issued in accordance with § 7.17 of the NVLAP Procedures (15 CFR Part 7). Establishment of this program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense follows a request by the Defense Communications Agency. A **Federal Register** notice announcing the request for the Protocols LAP was published on December 3, 1987 (50 FR 45986-45988). Comments received in response to the announcement were reviewed by the Defense Communications Engineering Center whose director, Warren P. Hawryiko, has concluded that there were no valid reasons presented in the comment letters to prevent establishment of the Protocols LAP and therefore requested the National Bureau of Standards to proceed to establish the requested program.

The purpose of the LAP is to accredit and provide national recognition to laboratories capable of performing tests in accordance with the designated test methods. The scope of the LAP includes testing services for: (1) Defense Data Network (DDN) X.25 Link and Network Layer Protocols as specified in the DCA DDN X.25 Host Interface Specification; (2) the five DoD packet switching High Level Protocols (HLPs): (I) Internet Protocol (IP) MIL-STD 1777; (II) Transmission Control Protocol (TCP), MIL-STD 1778; (III) File Transfer Protocol (FTP), MIL-STD 1780; (IV) Simple Mail Transfer Protocol, MIL STD 1781; and (V) TELNET, MIL-STD 1782; and (3) the AUTODIN Mode I Protocol. Accreditation will be offered first for the X.25 protocol. Accreditation will be offered next, at least 45 days later, for the DoD HLPs (I)-(V). Accreditation for AUTODIN Mode I protocol will be offered last, after the initial X.25 protocol accreditations have been completed.

Procedure Prior to Application

Any testing laboratory interested in becoming accredited under this LAP should contact the Manager, Laboratory Accreditation, at the address shown above, specifying the protocols of interest. The laboratory will be sent the proposed technical documents for the requested protocol accreditation as they become available and will be invited to submit comments for their revision within 45 days of the publication date of this notice in the case of the X.25 protocol and of the mailing dates of the documents for the HLP and AUTODIN protocols. A meeting of all interested parties will be scheduled after each of the closing dates to resolve conflicting comments. If none arise, no meeting will be scheduled. The completed technical documents, instructions, fee schedules, and applications will be sent separately for each protocol as they become available to all laboratories that have previously requested them.

Earnest Ambler,

Director.

Dated: July 15, 1988.

[FR Doc. 88-16430 Filed 7-20-88; 8:45 am]

BILLING CODE 3510-13-M

ASSISTANT SECRETARY OF DEFENSE

WASHINGTON, D.C. 20301-3040

COMMAND, CONTROL, COMMUNICATIONS AND INTELLIGENCE

AUG 26 1988

MEMORANDUM FOR SECRETARIES OF THE MILITARY DEPARTMENTS
CHAIRMAN, JOINT CHIEFS OF STAFF
DIRECTORS OF THE DEFENSE AGENCIES

SUBJECT: Conformance Testing of Military Standard Data Communications Protocol Implementations

This memorandum establishes a conformance testing policy to be used for the acquisition of future implementations of military standard data communications protocols by the Military Services and Defense Agencies. This policy mandates that specific tests be conducted for each software version and/or hardware type. Testing provides a means of determining conformance to standards and is intended to significantly increase the probability of interoperation in a multi-vendor environment. This policy specifically applies to MIL-STDs 1777, 1778, 1780, 1781 and 1782. Implementations of these protocols must be tested by a National Bureau of Standards (NBS) accredited laboratory prior to first operational use on any DoD network.

The conformance testing requirement is mandatory for all contracts executed after 1 June 1989. A Qualified Products list will be established and maintained by the Defense Communications Agency. For a product to be placed on the Qualified Products List, acceptable test results must be presented to DCA from an accredited laboratory which is independent of the manufacturer. The modification of a Qualified Product will require retesting by an accredited laboratory prior to operational use.

A waiver of this requirement may be granted in instances where it can be clearly demonstrated that there are significant performance or cost advantages to be gained, or when the overall interests of the Department of Defense are best served by granting the waiver. Consideration will be given to waiving this testing requirement in cases of modifications to existing contracts or add-ons to existing programs. In instances where testing of products interferes with established DoD security requirements, consideration will be given to an alternative testing environment. Requests for waiver from this testing requirement should be forwarded to Defense Communications Agency, Code B100, Washington, D.C. 20305-2000.

Although conformance testing is mandatory only for future acquisitions, it is recommended that existing implementations of the MIL-STD protocols be tested to determine instances of nonconformance to the standards. In the event that such testing identifies nonconformance, the appropriate procurement office can determine where modifications may be prudently pursued.

With the adoption of the Open Systems Interconnection (OSI) protocols as co-standards to the above cited MIL-STD protocols, conformance of vendor implementations to protocol standards becomes paramount if proposed solutions such as dual DoD/OSI protocol hosts and dual DoD/OSI applications layer gateways are to effectively allow interchange between the two protocol suites. The establishment of DoD policy requiring conformance testing for DoD protocol implementations is a positive step towards a successful transition to the OSI protocol suite. OSI protocol suite implementations will also be required to successfully complete conformance testing once OSI conformance test capabilities are fully developed and implemented. NBS is addressing this issue and will be establishing a conformance testing policy as a Federal Information Processing Standard (FIPS). The NBS policy will apply to the OSI protocols covered by the U.S. Government Open Systems Interconnection Profile (GOSIP). The DoD and OSI protocol suite conformance test capabilities will provide the Military Services and Defense Agencies with important assurances of product conformance to specifications.

Gordon A. Smith

Header Compression for TCP/IP Datagrams

Van Jacobson
Lawrence Berkeley Laboratory

IETF
Cocoa Beach, FL
April 11–14, 1989

Executive Summary:

With modest amounts of storage ($<$1KB) and computation ($<100\mu$sec.), almost all TCP/IP datagram headers can be compressed down to one byte (plus the two byte TCP checksum). This makes telnet usable over even the slowest SLIP lines (e.g., 300 bps) and offers quite reasonable telnet, ftp, smtp, etc., service over 2400 bps and faster lines.

The problem:

People malfunction when **interactive response time** for things like typing echo is more than 100–200 ms.

There are three ways to lose over low-speed SLIP lines:

- Headers gobble all the bandwidth (character plus echo in TCP/IP packets = 82 bytes = 820 bits ÷ .2 sec ⇒ 4100 bps).

- Background file transfers grab all the bandwidth (440 byte packets = 4400 bits ÷ .2 sec ⇒ 22,000 bps).

- A modem trying to cheat the Shannon limit (≈22,000 bps) may thrash.

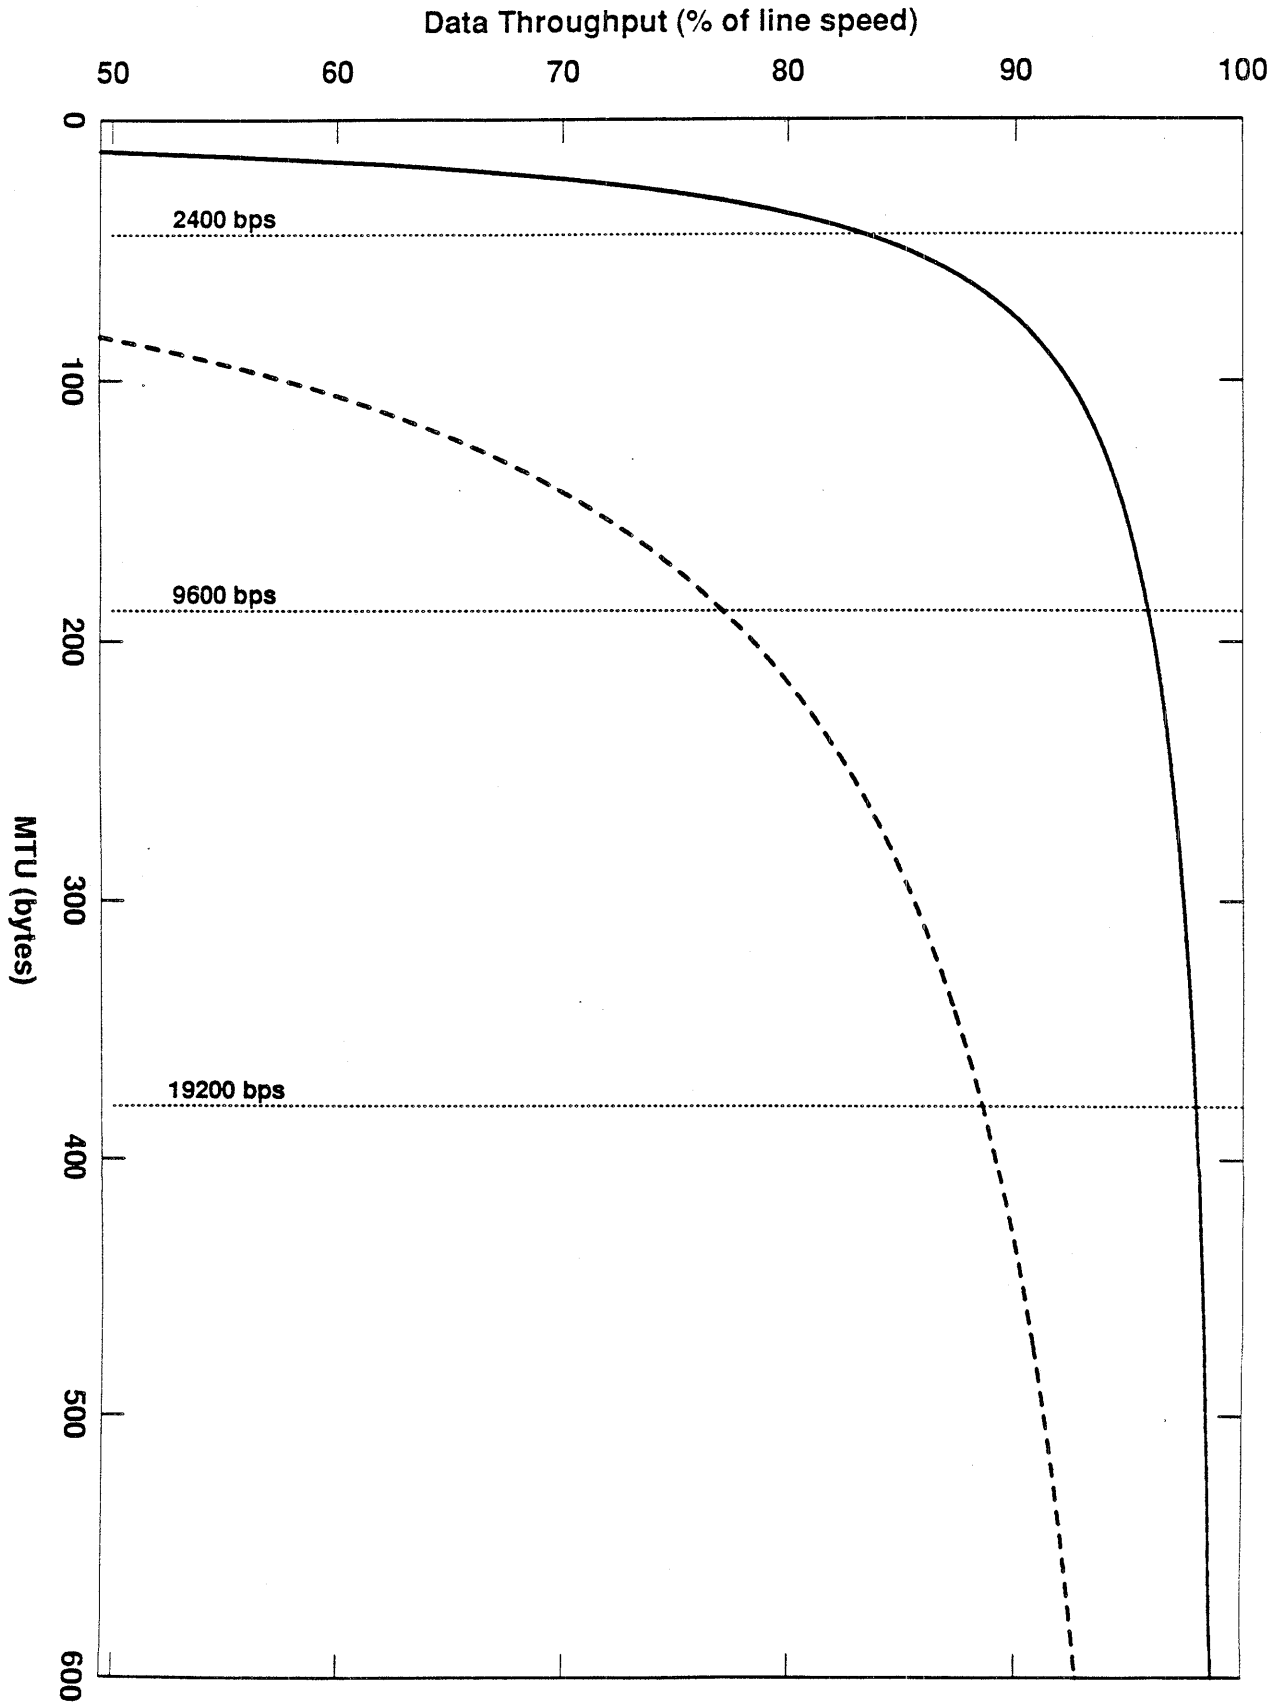
Data Throughput (% of line speed)
50
60
70
80
90
100
0
100
200
300
400
500
600
MTU (bytes)
2400 bps
9600 bps
19200 bps

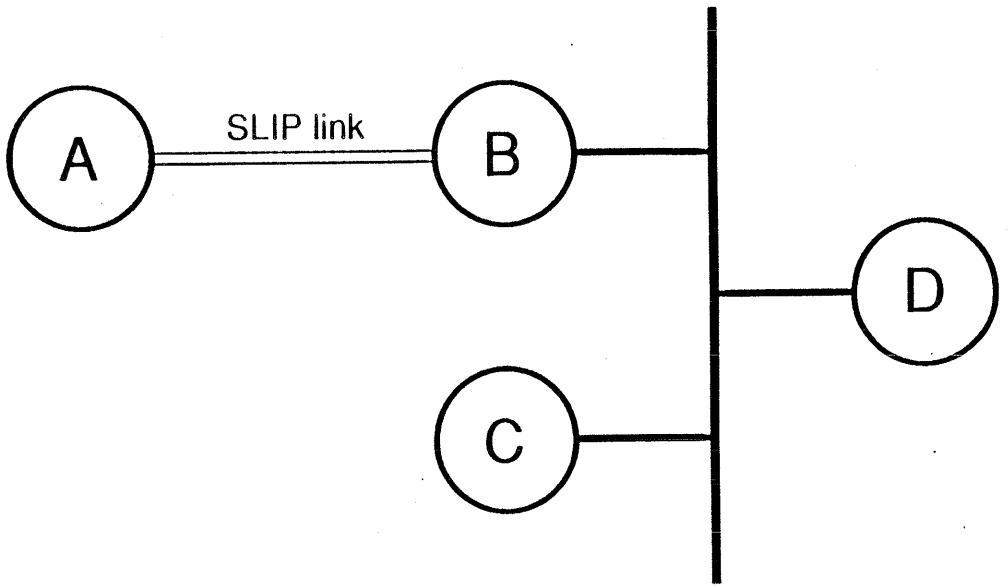
A
SLIP link
B
D
C

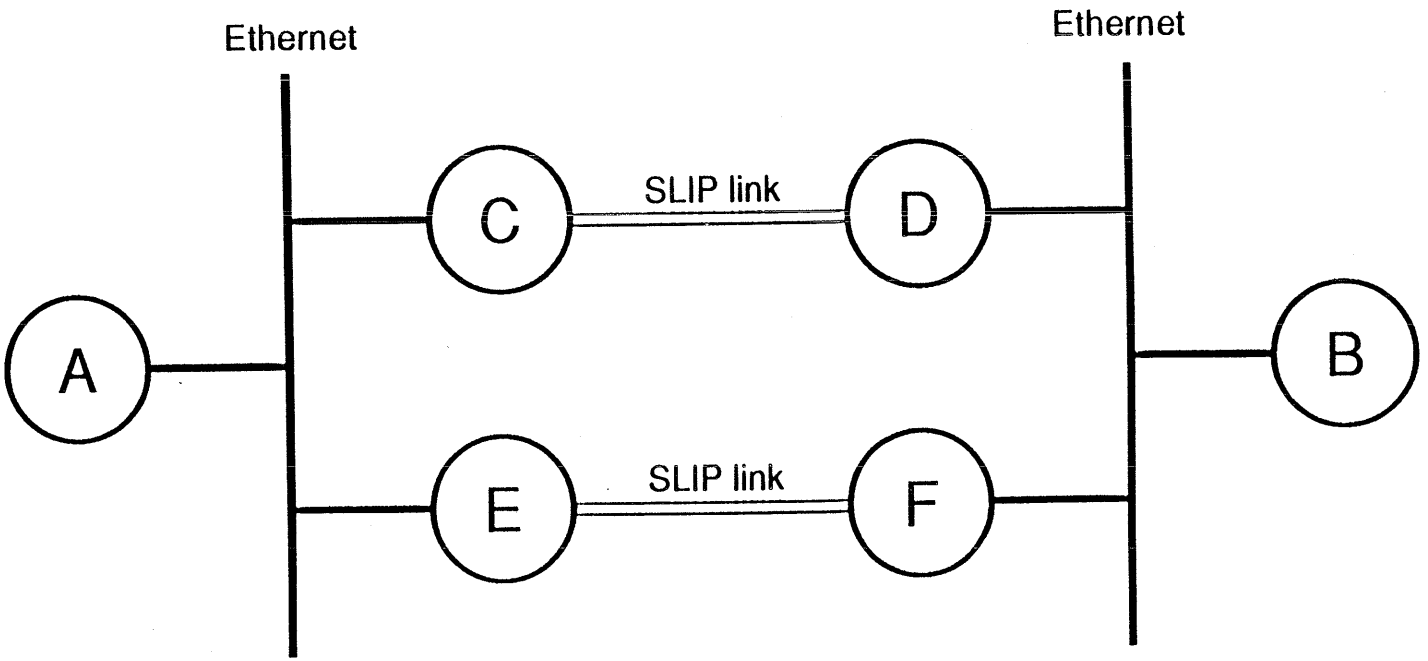
Ethernet
Ethernet
C
SLIP link
D
A
B
E
SLIP link
F

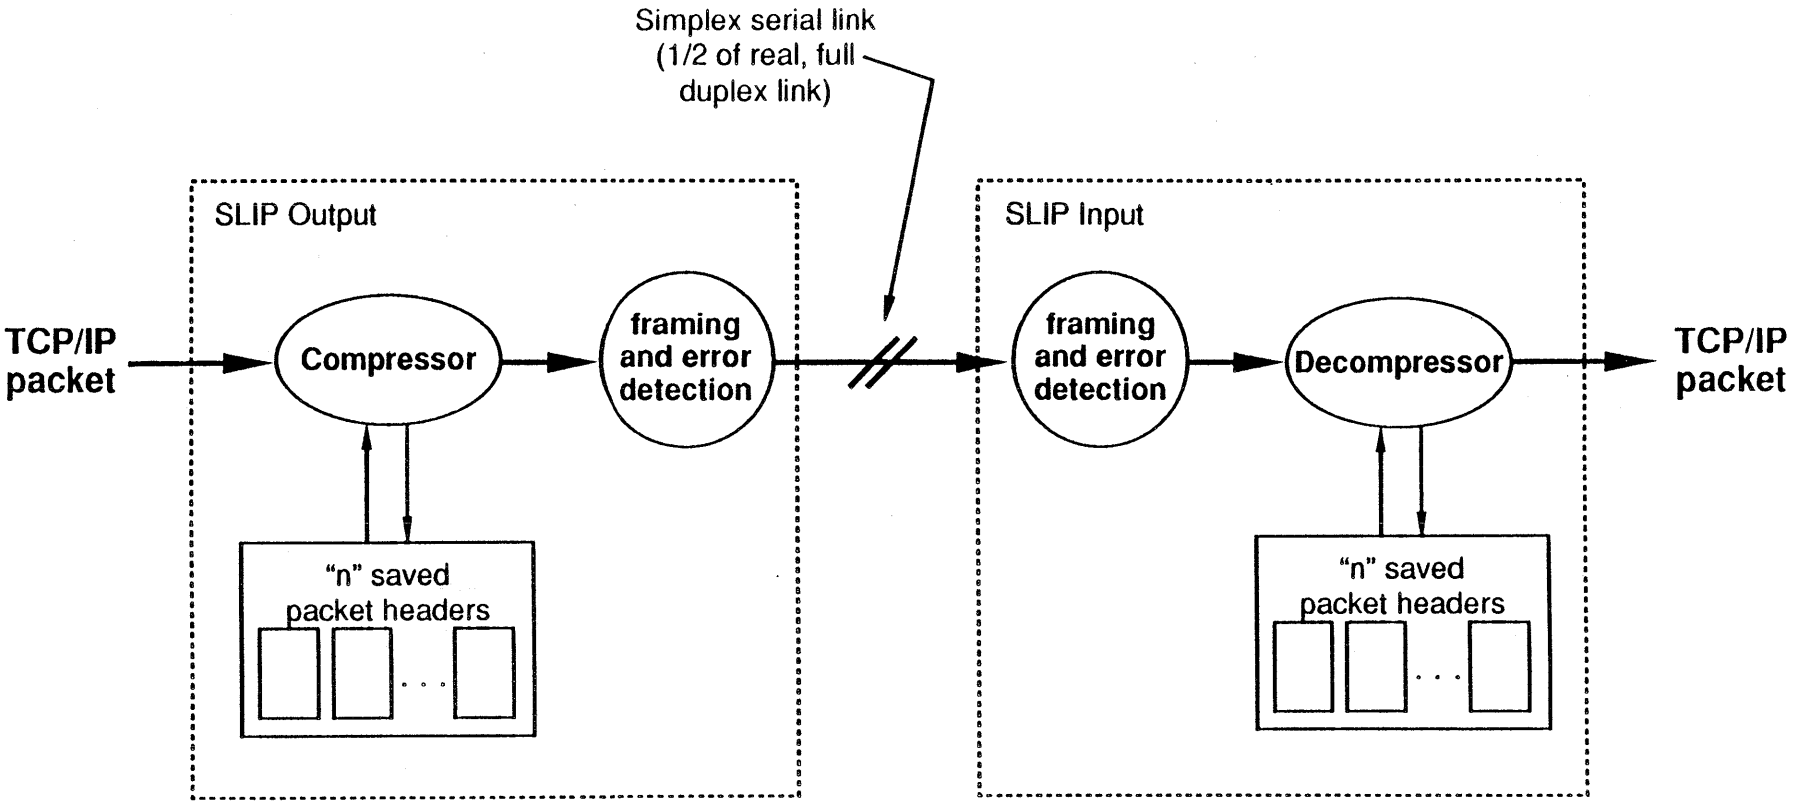
Simplex serial link
(1/2 of real, full
duplex link)
SLIP Output
SLIP Input
TCP/IP
packet
Compressor
framing
and error
detection
framing
and error
detection
Decompressor
TCP/IP
packet
"n" saved
packet headers
"n" saved
packet headers

IP / TCP
Datagram Format

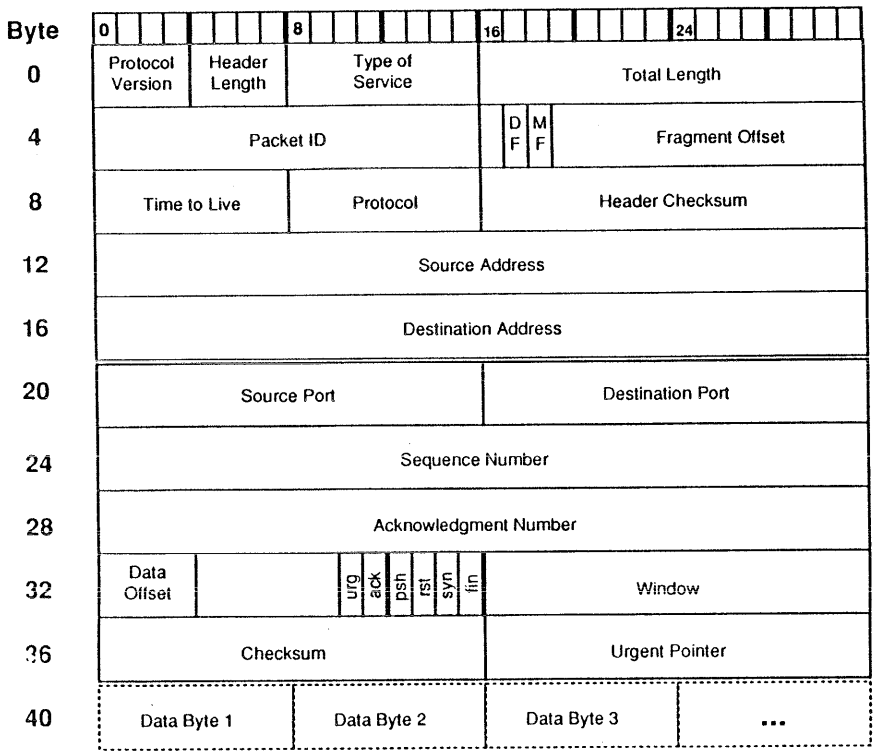

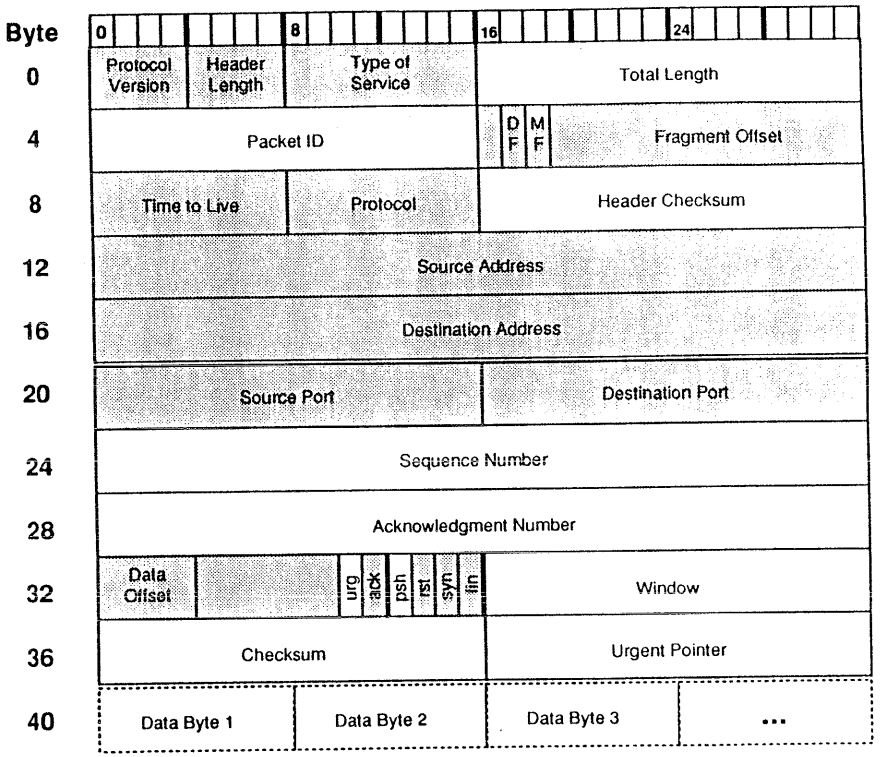

Telnet session between A and B

A.131 > B.telnet . 1(1) ack 100 id 41
B.telnet > A.131 . 100(1) ack 2 id 666
A.131 > B.telnet . 2(1) ack 101 id 42
B.telnet > A.131 . 101(1) ack 3 id 667
A.131 > B.telnet . 3(2) ack 102 id 43
B.telnet > A.131 . 102(2) ack 5 id 668
A.131 > B.telnet . 5(1) ack 104 id 44
B.telnet > A.131 . 104(1) ack 6 id 669

FTP data transfer from A to B

A.ftp > B.61 . 1(216)
B.61 > A.ftp . ack 217
A.ftp > B.61 . 217(216)
A.ftp > B.61 . 433(216)
B.61 > A.ftp . ack 649
A.ftp > B.61 . 649(216)
A.ftp > B.61 . 865(216)
B.61 > A.ftp . ack 1081

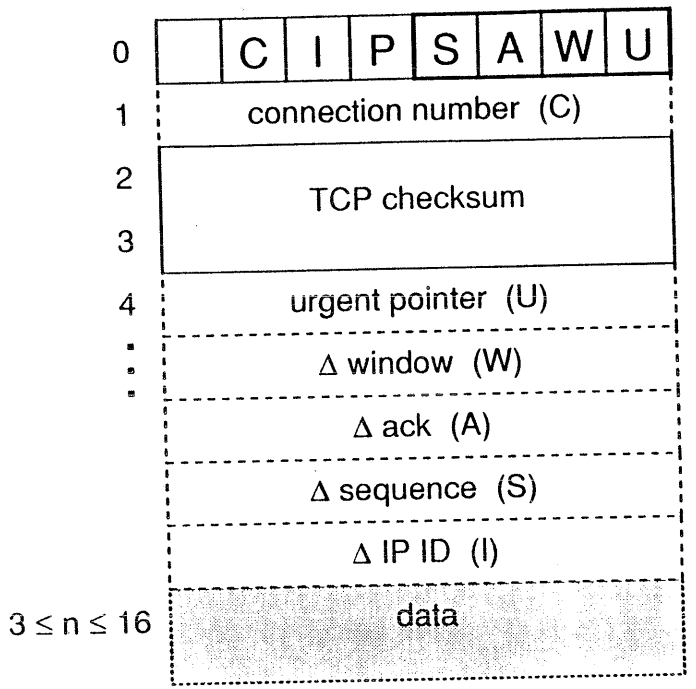

FTP data transfer from A to B

```
A.ftp > B.61 . 1(216)
```

ΔS = 216

```
A.ftp > B.61 . 217(216)
```

ΔS = 216

```
A.ftp > B.61 . 433(216)
```

ΔS = 216

```
A.ftp > B.61 . 649(216)
```

ΔS = 216

```
A.ftp > B.61 . 865(216)
```

Telnet session between A and B

```
A.131 > B.telnet . 1(1) ack 100  id 41
                ΔS = 1       ΔA = 1
A.131 > B.telnet . 2(1) ack 101  id 42
                ΔS = 1       ΔA = 1
A.131 > B.telnet . 3(2) ack 102  id 43
                ΔS = 2       ΔA = 2
A.131 > B.telnet . 5(1) ack 104  id 44
```


While any and all of the *sequence number, ack, window* and *urgent data* fields can change between two packets, two special cases describe almost all TCP traffic:

- Sequence number and ack change by the amount of data in the last packet seen; no change in window and no urgent data. (This is echoed, interactive traffic.)

- Sequence number changes by the amount of data in the last packet seen; no change in ack or window and no urgent data. (This is bulk data transfer or non-echoed interactive traffic.)

FTP data transfer from A to B

A.ftp > B.61 . 1(216)	U0
ΔS = 216	
A.ftp > B.61 . 217(216)	C- 0f xx xx dd...
ΔS = 216	
A.ftp > B.61 . 433(216)	C- 0f xx xx dd...
ΔS = 216	
A.ftp > B.61 . 649(216)	C- 0f xx xx dd...
ΔS = 216	
A.ftp > B.61 . 865(216)	C- 0f xx xx dd...

The Framer must:

- Provide a packet type field to distinguish three types (*IP, uncompressed TCP* and *compressed TCP*).
- Provide the length of the received packet.
- Provide good error detection (*not* correction — packets are discarded on error).

Error handling has to deal with two different problems:

Compression squeezes out almost all the header redundency $\Longrightarrow$ Decompressor will turn random line noise into a perfectly valid TCP packet.

Delta-encoding loses "DC" for sequence number, ack and window $\Longrightarrow$ dropped packet can invisibly shift down the rest of sequence space.

Compression rules: A type *IP* packet is sent if

- IPPROTO isn't "TCP".
- packet is fragment.
- SYN, FIN or RST is set or ACK is clear.

Compression rules (cont.): An *uncompressed TCP* packet is sent if

- there is no saved state for connection.
- header length, TOS, TTL, data offset, IP options or TCP options change.
- Sequence number or ack change is < 0 or $\geq 2^{16}$.
- nothing changed and data length zero or same as last packet.
- changes match special-case encodings.
- URG clear but urgent pointer changes.

A type *compressed TCP* packet is sent otherwise.

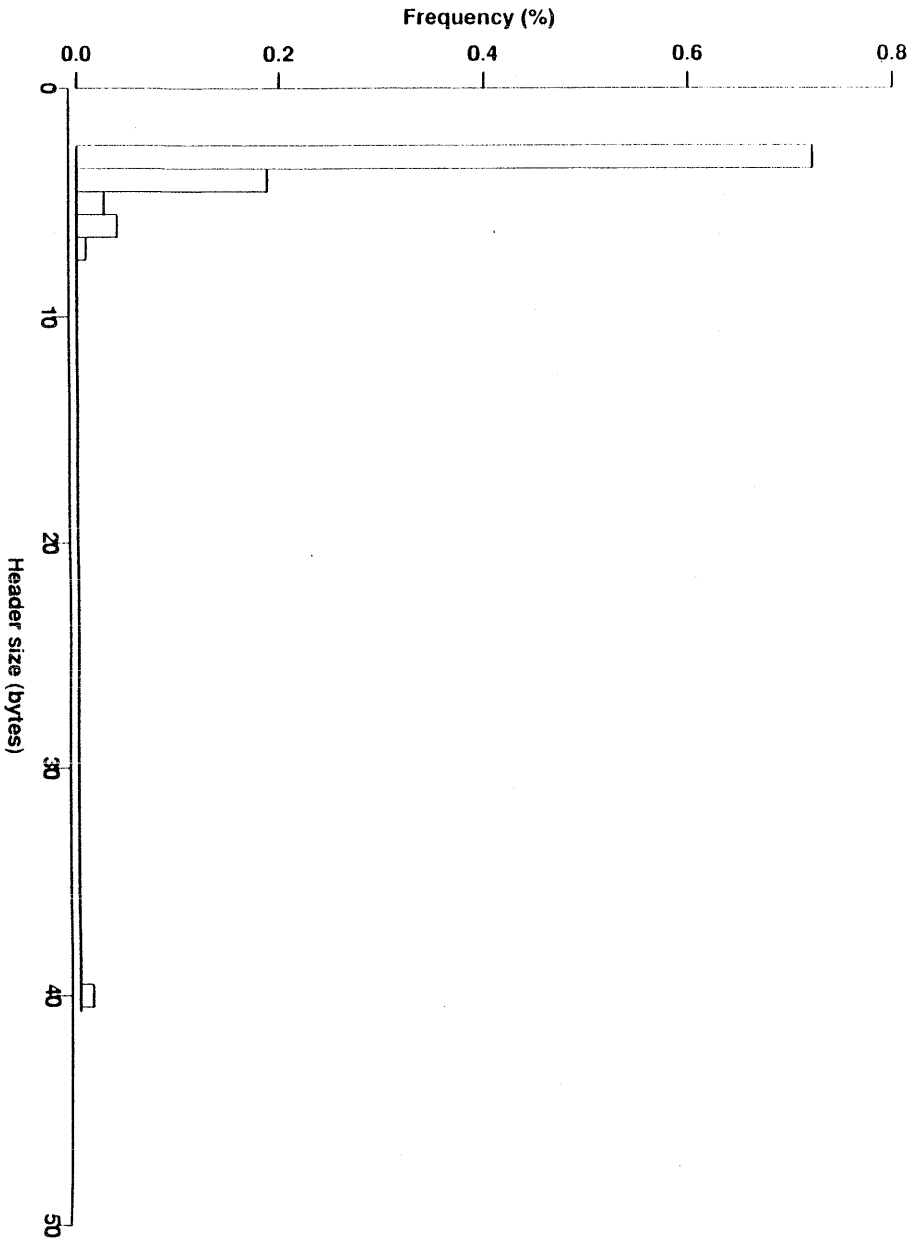

Compression effectiveness - client side of telnet

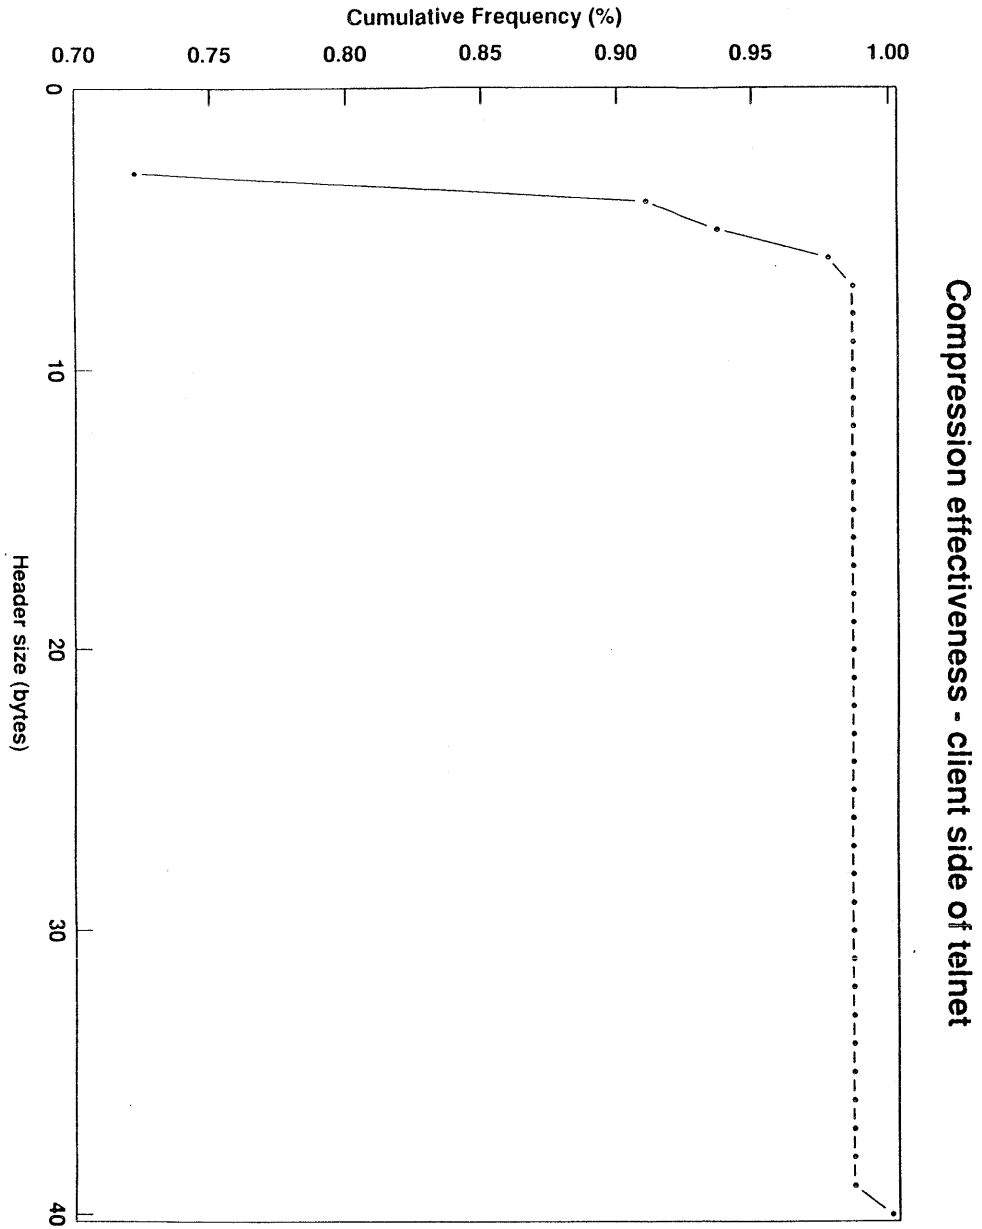

Compression effectiveness - client side of telnet

Code Timings

Machine	*Average per-packet processing time (μ sec.)*	
	Compression	*Decompression*
Sun 3/60	98	91
Sun 3/50	130	150
Sun 4/260	46	20
CCI	110	140
Vax 750	800	500
Vax 780	430	300

(Times measured using modified Berkeley Unix SLIP driver and mixed telnet/ftp traffic trace.)

VII. Papers Distributed at IETF

Five documents were made available at the April IETF meeting. Four are enclosed. The remaining document entitled "Inter- domain Intermediate Systems Routing", January 1989 Draft Technical Report, ECMA TR/ISK, Report # ECMA/TC32-TG10/89/... is not enclosed. This document was authored by the European Computer Manufacturers Association "to state the ECMA position with regard to Inter-Domain Routing; to serve as a vehicle for influencing decisions in other standard arenas; and to formalize the work carried out by ECMA". Inquiries concerning this particular document should be directed to Doug Montgomery/301-975-3630 or Tassos Nakassis/301-975-3632.

INTERNET CLUSTER ADDRESSING SCHEME AND ITS APPLICATION TO PUBLIC DATA NETWORKS

Carl-Herbert Rokitansky

Fern University of Hagen,
D-5860 Iserlohn, FRG
roki@A.ISI.EDU

The DARPA Internet protocol suite (TCP/IP, FTP, SMTP, TELNET, etc.) has developed into de facto industry standard for heterogeneous packet-switching computer networks. However, it seems that the Internet community has neglected the highly important public data network community so far. Thus, the current Internet gateway architecture does not provide dynamic algorithms to route Internet datagrams through public data networks.

In this paper a new concept of an addressing scheme is presented, in which a set of Internet networks is associated to an Internet cluster. Since this "Cluster Addressing Scheme" is of interest especially for wide-area networks (whose structure should be visible to the outside world for routing decisions), the application of the cluster addressing scheme to the system of X.25 public data networks is proposed. In addition the use of an address-mask (called "Cluster-Mask") for routing decisions within the cluster is discussed. Finally, due to the fact that VAN-gateways interconnect a datagram-oriented Internet world and a connection-oriented X.25 world, a new use of the IP Source Route option is proposed. The presented concept of the cluster addressing scheme provides a basis for the routing of Internet datagrams through X.25 public data networks and would therefore allow worldwide interoperation between the many local-area networks in various countries now using DARPA Internet TCP/IP protocols.

1. INTRODUCTION

The DARPA Internet system presently includes several thousand hosts connected to over 450 networks using over 250 gateways. It provides packet transport by means of a datagram service for hosts subscribing to the DARPA Internet protocol suite. A host may be connected to more than one network. Each datagram contains a 32-bit source and destination address and travels independently through the Internet. Datagrams can be sent over different routes and they can be delivered out of order; if they get lost or contain errors, duplicate datagrams are retransmitted.

The basic datagram protocol is the Internet Protocol (IP) [1]. Error reporting, flow control, first-hop gateway redirection and other control functions are provided by the Internet Control Message Protocol (ICMP) [2]. The Transmission Control Protocol (TCP) provides reliable end-to-end data stream service. A much simpler transport protocol than TCP is the User Datagram Protocol (UDP). All user level protocols above use either TCP/IP (e.g. FTP, TELNET, SMTP) or UDP/IP (e.g. NAMESERVER) as the basic packet transport mechanism. Due to their widespread implementation under various operating systems these protocols have developed into de facto industry standards for heterogeneous packet-switching computer networks.

The Internet model includes constituent networks, called local networks, to distinguish them from the Internet system as a whole. These local networks are connected together by means of Internet gateways. Each gateway is connected to two or more networks by a physical interface and has an address on each of the local nets between which it provides datagram transport service. Gateways belonging to different gateway systems ("autonomous systems") might use different intra-system routing mechanisms. In order to maintain the routing tables, an

---

This work was carried out at the German Aerospace Research Establishment (DFVLR) as part of the Internet research project and is continued at the Fern University of Hagen.

interior gateway protocol like the Gateway-Gateway Protocol (GGP) [3] can be used to exchange routing information between gateways of the same autonomous system, while the Exterior Gateway Protocol (EGP) [4] is used to exchange network reachability information with gateways belonging to a neighboring system.

Internet network numbers are assigned [5] to networks that are connected to the DARPA-Internet and DDN-Internet, and to independent networks that subscribe to the DARPA Internet protocol suite. Currently more than 10.000 networks numbers are officially assigned.

In this paper a new concept of a cluster addressing scheme is presented, which, applied to the international system of X.25 Public Data Networks (hereafter referred to as PDN), allows the implementation of improved algorithms to connect Internet networks to the DARPA Internet as well as to interconnect independent networks together by routing Internet datagrams via "VAN-gateways" through X.25 public data networks and to Internet hosts which are directly attached to a PDN ("Internet/PDN-host").

In Section 2 an outline of the current Internet routing model is given. In Section 3, we present the cluster addressing scheme and the use of a cluster-mask for routing decisions within a cluster. In Section 4, we discuss the application of the cluster addressing scheme to the international system of X.25 public data networks and a new use of the IP source route option. Finally, in Section 5, we summarize the changes to the existing Internet gateway system which would be necessary to support the cluster adressing scheme and the advantages and disadvantages of this clustering scheme and its application to PDN.

2. CURRENT INTERNET ROUTING MODEL

The current Internet routing model assumes that the route to a host can be computed by an algorithmic transformation on the destination address. The routing algorithms conforming to this model compute only a single (shortest) path from a given gateway to a given destination network, based on some metric such as hopcount or delay, without dependence on costs, type of service or anything else. Also, it is important to understand that the computed route does not depend on any network parameter (e.g. delay, costs, hopcount, link quality) within the destination network.

According to this "network-centric" routing model all hosts in the Internet must make the following routing decision when sending a datagram: Is the datagram addressed to a host on a directly connected network ("local" network) and can therefore be sent directly, or is it addressed to a host on a different Internet network ("foreign" network) which is reachable only via a local gateway. Note, that if subnets are in use (see [6] and [7]) only hosts/gateways on the same subnet appear to be local, while all other destinations are assumed to be reachable only indirectly via a (sub)gateway.

3. CLUSTER ADDRESSING SCHEME

The DARPA Internet is a complex of heterogeneous networks. Local-area networks (LANs), metropolitan-area networks (MANs) and wide-area networks (WANs) are connected together by means of Internet gateways. Usually, a distinct Internet network number is assigned to each of these networks, except that, according to the "Internet Standard Subnetting Procedure" [7], a single Internet network number might be assigned to a complex of LANs (each of it treated as a subnet). The fact that a network is subnetted (and therefore its internal structure) is invisible outside the network and is, in the case of a complex of LANs, of little interest to the outside world.

However the internal structure of a wide-area network (e.g. PDN, satellite network) might be of interest even outside such a network, especially for routing decisions.

Consider a (wide-area) network W, to which the Internet networks A and B are connected by means of gateways GAW and GBW as shown in Figure 3-1. To each network a different Internet network number is assigned. Host HA is attached to network A and hosts HW1 and HW2 are connected to network W.

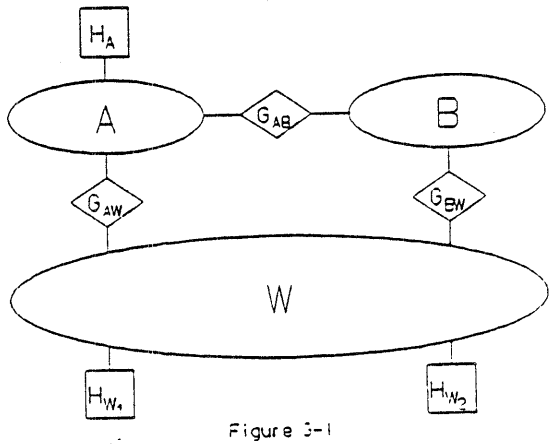

Figure 3-1

Assume that direct connections can be established between the gateways and hosts attached to network W.

According to the Internet routing model, packets from host HA to both HW1 and HW2 will be routed via gateway GAW (minimum hop count).

If W is a homogeneous network (i.e. delays, costs, quality of links, etc. between hosts/gateways on W do not differ very much) then routing through network B will be probably worse in any case.

Now assume that network W is inhomogeneous. For example the costs for a connection between GAW and HW2 are three times higher than between GBW and HW2. In this case it might be reasonable to route packets from HA to HW2 through network B via gateway GBW instead via GAW. Now, for routing decisions, the internal structure of network W would be of interest even outside it. However the current Internet gateway architecture does not provide any algorithms for this situation (due to "network-centric" routing), except that a user on HA could specify GBW explicitly in an IP source route option.

Note, that dividing network W into several subnets according to the internal structure of W would have no external effects, because it is invisible to the outside world as mentioned above. Therefore packets would still be routed from HA to HW2 via GAW and not through network B.

Another idea would be to assign different Internet network numbers to subdivisions (e.g. X and Y, see Figure 3-2) of W, instead of assigning a single Internet network number to network W. Thus, network W would become a complex of several Internet networks, and the structure of W would be visible even outside of it.

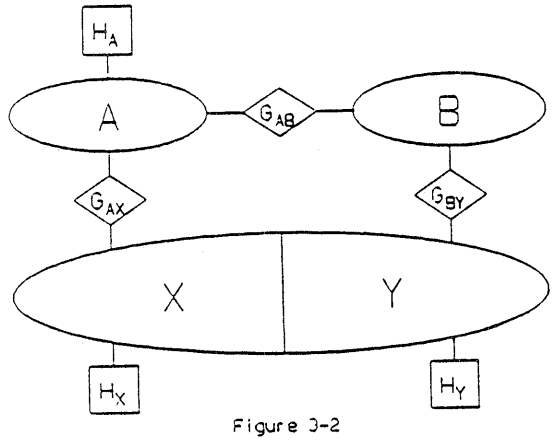

Figure 3-2

However this assignment would be inconsistent with the current routing model, because there are packets to "foreign" networks which need not be routed via a local gateway but can be sent directly: Consider packets to be sent from HX (former HW1) to HY (former HW2). Note that HX and HY are now hosts attached to different Internet networks, but direct connections can still be established between HX and HY without transiting an Internet gateway! According to the current routing algorithms, HX would determine that the destination HY is on a different ("foreign") Internet network and would therefore decide that the packets must be sent to a local gateway on the common network X. There is only one gateway GAX connected to network X, but transmitting packets to GAX would be unreasonable, since they can be sent directly to HY. But HY is neither a gateway nor is it a host on the local net. Similarly gateway GAX would encounter the same problems in its routing decision as HX when forwarding packets to HY. In addition gateway GAX cannot send an ICMP Redirect message to HX specifying HY as a better first hop on the route towards the destination, because HY is not a host on the same network.

Therefore, the following model of a clustering scheme is proposed, which adds an additional level to the interpretation of Internet addresses and is called "Cluster Addressing Scheme":

3.1. Cluster Addressing Scheme Model

Specific Internet network numbers are assigned to a set of nets between which direct connections can be established without transiting a gateway. These networks are associated to an "Internet Cluster". For all routing decisions within the cluster, and for the specification that different Internet networks are associated to a cluster, the use of an address-mask, called "Cluster-Mask", is proposed. By means of this cluster-mask, all hosts within the same cluster, even if attached to different Internet networks appear to be local. ICMP Redirect messages can be sent directly between gateways and hosts belonging to the same cluster. The fact that several Internet networks are aggregated to an Internet cluster is invisible to the outside world. However the internal structure of the cluster, which is a complex of Internet networks ("clusternets"), is visible outside the cluster.

The 32-bit INTERNET address consists of:

```
<INTERNET address> ::=
            <network-number><rest-field>
```

Now the cluster addressing scheme proposes that the <network-number> field is interpreted as

<network-number> ::=
    <cluster-number><cluster-net-number>

Although this subdivision of the <network-number> field is used for routing decisions within the cluster, it is invisible to the outside world.

Thus, if this clustering scheme is in use, the INTERNET address can be interpreted as

<INTERNET address> ::=
  <cluster-#><cluster-net-#><rest-field>

Consider the example of a (wide-area) network W as discussed above (Fig. 3-1). Assume that direct connections can be established (pairwise) between all hosts and gateways attached to network W. Further assume that network W is inhomogeneous; for example the costs for a connection between GAW1 (or HW1) and HW2 (or GBW2) are three times higher than between GAW1 and HW1 (or GBW2 and HW2). Now, according to the cluster addressing scheme specific Internet network numbers [W1] and [W2] are assigned to homogeneous subdivisions of W as shown in Figure 3.1-1. These Internet networks W1 and W2 are associated to an Internet cluster (W-cluster). If packets are to be sent from HW1 to HW2, host HW1 can determine by means of a cluster-mask that HW2 is a "local" host since it belongs to the same cluster. Therefore, the packets can be routed directly ("locally") by establishing a direct connection to HW2. If for some reason the packets are sent to gateway GAW1, then, according to the clustering scheme, GAW1 can send an ICMP Redirect message to HW1 specifying HW2 as a better first hop in this message. These routing mechanisms are described in detail in the following sections.

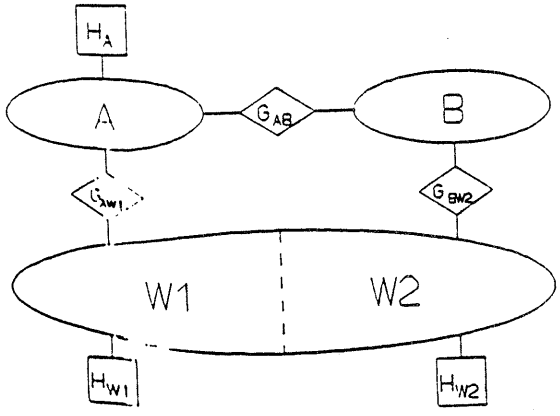

Figure 3.1-1

3.2. Cluster-Mask

A host uses a 32-bit mask, called "cluster-mask", to determine which bits of the <network-number> field are used to specify the cluster. In this cluster-mask all bits corresponding to the <cluster-number> field are set to one, while the remaining bits are set to zero. If the width of the <cluster-net-number> field is zero (e.g., if the cluster-mask contains all ones in the <network-number> field and zeros in the <rest-field>) the net does not belong to any Internet cluster.

For example, if a set of class B networks with Internet addresses whose 8 high-order bits are identical is associated with an Internet cluster, the cluster-mask would have the following value:

```
<network-number>< rest - field >
<cluste><cl.net>< rest - field >

11111111000000000000000000000000   binary
     255.      0.      0.      0  decimal
                                   cluster-mask
```

By means of this cluster-mask a host can determine if it is connected to a cluster-net. The host uses this mask for the routing decision if the destination IP-address specified in a datagram is either "local" or "foreign" depending whether the destination is in the same cluster or not. All datagrams to local destinations (even on different Internet networks (cluster-nets)) can be sent directly to the destination without transiting an Internet gateway.

If the bitwise AND of this cluster-mask with the destination IP address ("dg.ip_dest") matches the bitwise AND of the mask with the host's own IP address ("my_ip_addr"), the destination is assumed to be in the same cluster, and therefore the datagram can be sent directly ("locally"); if not, the destination is assumed on a network outside the cluster, reachable only via a gateway.

If an IP implementation supports subnets [7], (normally) no changes to the code are necessary to support the clustering scheme, except that "my_ip_mask" must be assigned the value of the cluster-mask:

```
IF bitwise_and(dg.ip_dest,my_ip_mask) =
   bitwise_and(my_ip_addr,my_ip_mask)

   THEN send_dg_locally(dg,dg.ip_dest)
   ELSE send_dg_locally(dg,gateway_to
    (bitwise_and(dg.ip_dest,my_ip_mask)))
```

3.3. ICMP Address Mask Request - Reply

To determine which cluster-mask (address mask) is in use, the two ICMP messages (specified in RFC-950 [7])

- ICMP Address Mask Request (AM1)
- ICMP Address Mask Reply (AM2)

can be used without changes. The address mask field of a Reply message contains the value of the 32-bit cluster-mask.

3.4. ICMP Redirect Messages

Due to the fact that all hosts and gateways within the same cluster appear to be reachable "locally", the cluster addressing scheme allows to send ICMP Redirect messages between gateways and hosts within the same CLUSTER and not only within a directly connected Internet NETWORK. This is a significant extension of the usage of the ICMP Redirect message.

3.5. Advantages and Disadvantages

The concept of associating a set of Internet networks to an Internet cluster and the specification of a cluster-mask has the following advantages:

- The internal structure of a cluster, consisting of a set of Internet networks, is visible to the outside world. This can be important for routing decisions outside the cluster.
- The fact that an Internet cluster has been formed is invisible outside the cluster. Therefore, no changes to the existing Internet gateway system are necessary to support the cluster addressing scheme.
- All hosts (gateways) within the same cluster appear to be reachable directly ("locally"). This is important for routing decisions within the cluster.
- No changes, or minor ones only, to hosts supporting subnets
- ICMP Address Mask Request and Address Mask Reply messages can be exchanged to determine which cluster-mask is in use.
- ICMP Redirect messages can be used between gateways and hosts on different Internet networks, but within the same cluster.

Disadvantages are:

- Specific Internet network numbers must be reserved for each cluster.

  For the implementation of the clustering scheme and for the reservation of Internet network numbers for specific clusters, it is proposed to assign (reserve) a number (depending on the number of bits (width) of the <cluster-net-number> field of the <INTERNET-address>) of the highest, not yet assigned network numbers of each class of networks.

  Remark: The disadvatage of reserving Internet network numbers for the clustering scheme seems to be acceptable due to the fact (see RFC-1020 [5]) that with regard to a maximum number of

126	Class A networks,	27 (21.4%)
16382	Class B networks,	301 (1.8%)
2097150	Class C networks,	7494 (0.4%)

  network numbers are assigned, and especially those network classes (B and C) with the higher percentage (98.2 % and 99.6 %) of available numbers are those which are of more interest for the clustering scheme.

4. APPLICATION OF THE CLUSTERING SCHEME TO X.25 PUBLIC DATA NETWORKS

The international system of X.25 Public Data Networks (PDN) is a typical wide-area network (WAN). In this system, the national packet-switched data networks in various countries are connected via gateways (CCITT, Rec. X.75 [8]) to allow international interworking between hosts on different national public data networks over international virtual circuits.

4.1. Costs

The costs for international virtual circuits differ from those for national virtual circuits and depend on:

- The charge for the call request
- The length of time the virtual circuit is open
- The data volume transmitted over that circuit in units of "segments"

Some facilities ("closed user group", "reverse charging", etc.) are usually not available for international calls.

4.2. X.121 Addressing

An X.121 address (CCITT, Rec. X.121 [8]) is assigned to each PDN host.

This international data number consists of:

<Internat. data number> ::= <DNIC><NTN>
or
<Internat. data number> ::= <DCC><NN>

where

DNIC....Data Network Identification Code
(fixed at 4 digits)
NTN.....Network Terminal Number
(up to 10 digits)

DCC.....Data Country Code (fixed at 3 digits, first 3 digits of DNIC)
NN......National Number (up to 11 digits)

DNIC	z denotes any digit from 2 thru 7
zxxn	(8..telex, 9..telephon)
DCC	x denotes any digit from 0 thru 9
	n denotes any digit from 0 thru 9 (network digit)

Thus, the system of data network identification codes (DNICs) as specified in Rec. X.121 [8] provides a theoretical maximum of 6000 (resp. 8000) DNICs. However, only about 100 national public data networks are in operation so far.

Currently an Internet class A network number [14.rrr.rrr.rrr] is assigned to the system of public data networks (PDN). For the time being the assignment of Internet addresses to hosts (gateways) on PDN is done successively in (chronological) order of request, regardless to which national public data network a host (gateway) belongs.

Connectivity with the Internet is provided by so called "VAN gateways", which are attached to the national public data networks.

4.3. Characteristics

The PDN can be characterized as follows:

- Wide-area network
- Complex of national public data networks
- Internat. virtual circuits are provided
- Different costs for international and national virtual circuits
- Costs depend on length of time and data volume transfered
- No broadcasting and no multicasting

4.4. Routing Thru PDN - Current Situation

Due to the characterization above, the following requirements seem to be reasonable for the routing of Internet datagrams through PDN:

- Routing decisions should be done with regard to the structure of the PDN (complex of national public data networks)
- Packets between PDN hosts should be sent directly through the PDN over virtual circuits

The current routing situation in the DARPA Internet with regard to PDN is very poor:

Due to the assignment of a class A network number to the system of public data networks (PDN), no internal structure of this PDN system is visible to the outside world. For this reason a division of the PDN into subnets would have no external effects.

Currently, the PDN is declared reachable via the "BBN-VAN-GATEWAY". However, since the "BBN-VAN-GATEWAY" does not provide international calls, PDN hosts on other national public data networks are unreachable from the Internet through PDN. (NOTE: A connection from (!) a PDN host to any Internet host via the "BBN-VAN-GATEWAY" is possible).

Packets sent from a PDN host via a VAN gateway on the same national public data network to any Internet host will reach the destination. However, due to the current routing in the Internet, which is "network-centric" not "gateway-centric", reply packets will be routed to the "BBN-VAN-GATEWAY". Assuming that an international virtual circuit between the "BBN-VAN-GATEWAY" and the PDN host does not exist, these reply packets will not reach the PDN host, since the "BBN-VAN-GATEWAY" does not make international calls. Therefore, connections from PDN hosts via VAN gateways other than the "BBN-VAN-GATEWAY" cannot be established.

4.5. PDN-Cluster Addressing Scheme

To allow an improved routing of Internet datagrams through PDN according to the Internet routing model, we propose to apply the "Cluster Addressing Scheme" (presented in Section 3 of this paper) to the system of X.25 public data networks:

a) Internet class B network numbers (with identical bits in the first (high-order) 8-bit field of the Internet address) are assigned to national public data networks.

b) The national public data networks are associated to an Internet cluster ("PDN-Cluster")

c) For the specification of this cluster and for routing decisions within the cluster, a cluster-mask is used (value <255.0.0.0>), thus all hosts within the PDN-cluster appear to be reachable "locally".

d) ICMP Redirect messages can be sent to any PDN host to manage the routing within the PDN-cluster.

NOTE: No changes to the existing Internet gateway system are necessary to support the cluster addressing scheme other than reserving a set of class B network numbers for the PDN-cluster and implementing this scheme on VAN-gateways and PDN-hosts. On hosts supporting subnets [7] this can be done very easily by simply setting the address-mask to the value of the cluster-mask. Following is a detailed description of the PDN-cluster addressing scheme.

4.6. PDN-Cluster

Due to reasons of homogenity (cost structure, available network options, etc.) a mapping between Internet network numbers and Data Network Identification Codes (DNICs, Recomm. X.121 [8]) is proposed.

Therefore Internet class B network numbers with identical bits in the first (high-order) 8-bit field of the Internet address (see below) are assigned to the different national public data networks. This allows a maximum number of 65.536 PDN hosts on each network.

The national public data networks are associated to an Internet cluster (PDN-cluster).

For this reason the 16-bit <network-number> field (class B network) is divided into an n-bit <cluster-number> field and a (16 minus n)-bit <cluster-net-number> field:

```
<network-number> ::=
     <cluster-number><cluster-net-number>
         (n bits)     +    (16-n bits)
```

In deciding how many bits of the <network-number> field should be used for the <cluster-number> field it seems to be reasonable to distinguish between:

- Theoretical number of addressable DNICs (see Recomm. X.121 [8])
- Number of reachable DNICs currently (and in the near future)

The system of Data Network Identification Codes (DNICs) as specified in Rec. X.121 [8] provides a theoretical maximum of 6000 (resp. 8000) DNICs.

However, only about 100 different national public data networks are reachable currently.

Even assuming an increase of additional public data networks, a maximum number of 256 (254) addressable DNICs seems to be sufficient for the near future.

Therefore, it is proposed to use an 8-bit <cluster-number>-field for the PDN-cluster and an 8-bit <cluster-net-number>-field:

```
<network-number> ::=
     <cluster-number><cluster-net-number>
         (8 bits)     +     (8 bits)
```

This allows to address 256 (254) different national public data networks with 65.536 PDN hosts on each.

According to the cluster addressing scheme, the reservation of Internet network numbers for an Internet cluster should start with the highest, not yet assigned network numbers of each class.

Therefore, the assignment of the cluster-number [191.nnn.rrr.rrr] to the PDN-cluster is proposed, thus to reserve the Internet network-numbers [191.001.rrr.rrr] up to [191.254.rrr.rrr] for the different national public data networks.

The assignment of Internet network numbers to the national public data networks (different DNIC's) could be done in the order of request (by grouping requests for zones 2 to 7).

The following Internet network numbers could be assigned so far:

DNIC	Public Data Network	INTERNET #
2041	DATANET (Netherlands)	191.001
2342	IPSS (U.K.)	191.002
2405	TELEPAK (Sweden)	191.003
2624	DATEX-P (West Germany)	191.004
.	.	.
3110	TELENET (USA)	191.096
.	.	.
.	last	191.254
reserved		191.255

4.7. PDN-Cluster-Mask

For the specification of the PDN-cluster and for internal routing decisions within the cluster (as described in detail in 3.2 above), corresponding to the width of the <cluster-number> field, a cluster-mask is used in which the first (high-order) 8 bits are set to "one", while the remaining bits are set to "zero" (value <255.0.0.0>).

```
<network-number>< rest - field >
<cluste><cl.net>< rest - field >

11111111000000000000000000000000   binary
     255.      0.      0.      0  decimal
                                cluster-mask
```

If the bitwise AND of this cluster-mask with the destination IP address ("dg.ip_dest") matches the bitwise AND of the mask with the host's own IP address ("my_ip_addr"), the destination is assumed to be in the same cluster and therefore the datagram can be sent directly ("locally"); if not, the destination is assumed on a network outside the cluster, reachable only via a gateway.

```
IF   bitwise_and(dg.ip_dest,my_ip_mask) =
     bitwise_and(my_ip_addr,my_ip_mask)
THEN send_dg_locally(dg,dg.ip_dest)
ELSE send_dg_locally(dg,gateway_to
    (bitwise_and(dg.ip_dest,my_ip_mask)))
```

If an IP implementation supports subnets, (normally) no changes to the code are

necessary to specify the PDN-cluster, except that "my_ip_mask" must be assigned the value of the "PDN-cluster-mask" <255.0.0.0>.

Thus all PDN hosts within the PDN-cluster appear to be reachable "locally". In fact, direct national and international virtual circuits can be established between PDN hosts.

4.8. ICMP Address Mask Request - Reply

To determine which cluster-mask (address mask) is in use, the two ICMP messages (specified in RFC-950 [7]):

- ICMP Address Mask Request (AM1)
- ICMP Address Mask Reply (AM2)

can be used without changes. The address mask field of a reply message contains the value of the PDN-cluster mask.

4.9. ICMP Redirect Messages

Due to the concept of the cluster addressing scheme all hosts and gateways within the same cluster (even on different Internet networks) appear to be reachable "locally". Therefore, to manage the routing within the PDN-cluster between PDN hosts and VAN gateways, ICMP Redirect messages can be sent to any host in the PDN-cluster, specifying any PDN-gateway/host as a better first-hop towards the destination in the redirect message.

4.10. IP Source Route Option Included by VAN-Gateways

While the PDN is based on connection oriented protocols, the Internet uses a a datagram oriented packet-switching technology. Therefore it might happen that packets between two Internet hosts are not routed along the same path.

Let us assume that there are two VAN gateways V1 and V2 and a PDN host HP1, which are attached to a national public data network P1 as shown in Figure 4.10-1. To send packets to a host HA on network A, HP1 establishes a switched virtual circuit to V1. The packets forwarded by V1, are received at host HA; however, reply packets from HA to HP1 are routed (for some reason) to V2. Now, V2 would have to open a virtual circuit to HP1 to forward the packets to the destination HP1.

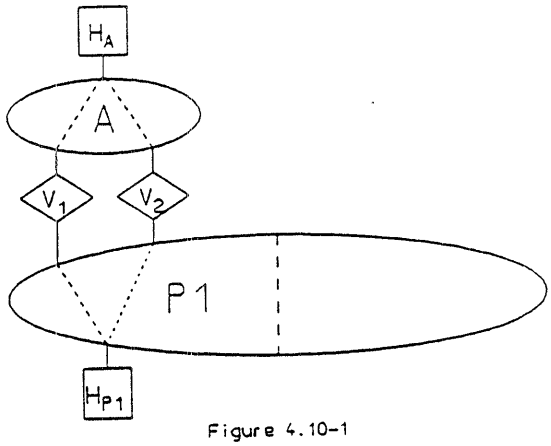

Figure 4.10-1

In most applications it seems to be reasonable to avoid such a situation for the following constraints:

- Technical: HP1 might be configurated with 1 logical channel only: Thus, an incoming call from V2 cannot be accepted, if the virtual circuit to V1 should remain established.

- Costs: Two virtual circuits would be in use. V1 and V2 might accept incoming calls only, but do not open a connection to HP1 since the calling site has to pay for the connection. (No reverse charging on international calls)

Since many Internet hosts use the IP source route option, when it is specified in received packets, even in their reply packets - a policy which is recommanded to be implemented on each host - a new use of the IP source route option is proposed to avoid the situation described above:

VAN gateways include their own Internet address (corresponding to the directly connected network interface through which a packet is transmitted) as an IP source route option in those datagrams, which are received from the "PDN-cluster-net" interface and forwarded to another network interface, by extending the Internet header by 8 (resp. 4) octets according to the following algorithm:

- If no source route option is specified in the original datagram, the Internet header is extended by 8 octets and the VAN gateway address is included as a Loose Source Route (LSR) option:

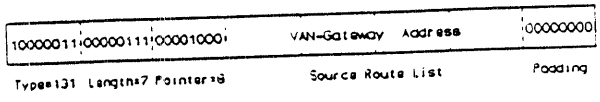

- If a Loose Source Route (LSR) or a Strict Source Route (SSR) option is already specified and the source route list does not contain the VAN gateways address, the source route list is extended by the VAN gateways address (4 octets, inserted as indicated by the pointer); the option length and the pointer are incremented by 4:

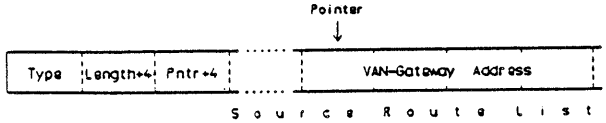

NOTE: Other options in the Internet header are uneffected; the source address and the destination address are left unchanged; if necessary, the Internet header is padded to a 32 bit boundary; the Internet Header Length (IHL), the Total Length and the Header Checksum are recomputed.

Thus, packets from HP1 via V1 to HA (see example above) are routed back (along the same path) from HA via V1 to HP1, if V1 includes its own Internet address as an IP source route option in packets from HP1 to HA, and HA uses this option in its reply packets to HP1.

5. SUMMARY

In this paper, a new concept of a cluster addressing scheme, in which a set of Internet networks is aggregated to an Internet cluster has been presented. Its application is of interest especially for WANs, since the structure of the WAN becomes visible even outside of it (important for Internet routing), while the fact that a cluster has been formed is invisible outside the cluster. Therefore no changes to the existing Internet gateway system are necessary.

For routing decisions the use of a cluster-mask is proposed; therefore all hosts within the cluster (even hosts on different Internet networks) appear to be reachable locally. If "my_ip_mask" is assigned the value of the cluster-mask, no changes, or minor ones, are necessary to support the cluster addressing scheme in hosts whose software allows the implementation of subnets. Even the two ICMP messages "Address Mask Request" and "Address Mask Reply" can be used without changes to determine which cluster-mask is in use. As a significant extension, ICMP Redirect messages can be used not only between gateways and hosts on the same Internet network, but also within the same cluster.

In addition, the application of the proposed cluster addressing scheme to the system of X.25 Public Data Networks (PDN) has been discussed. The PDN is a typical wide-area network, which consists of a number of national public data networks. Currently, due to the fact that only one Internet class A address is assigned to the PDN, it appears to be unstructured and therefore the routing of Internet datagrams through the PDN is very poor. To provide an improved routing the application of the cluster addressing scheme is proposed by assigning Internet class B network numbers to the national public data networks and associating these networks to the PDN-cluster. The proposal involves no changes to the existing Internet gateway system other than a policy reserving a set of class B network numbers for the PDN-cluster and implementing the cluster addressing scheme on VAN gateways and PDN hosts. For routing decisions within the cluster, the use of a cluster-mask has been discussed. PDN hosts whose software supports subnets can be equipped easily with the cluster addressing scheme. Finally, with regard to the connection-oriented characteristics of the PDN, a new use of the IP source route option (included by VAN gateways) has been discussed. The implementation of the proposed cluster addressing scheme would allow worldwide interoperation beween the many local area networks in various countries now using DARPA-Internet TCP/IP protocols.

ACKNOWLEDGMENT

It has been the poor routing of Internet datagrams through PDN, especially from the European point of view, which caused the development of the cluster addressing scheme and its application to PDN as described here. J. Noel Chiappa, Horst D. Clausen, J.J. Garcia-Luna, Dave Mills and Jon Postel provided helpful discussions and important suggestions.

REFERENCES

[1] Postel, J., ed., "Internet Protocol - DARPA Internet Program Protocol Specification", DARPA Network Working Group Request For Comments RFC-791, USC/Information Sciences Institute, Sept. 1981.

[2] Postel, J., ed., "Internet Control Message Protocol - DARPA Internet Program Protocol Specification", RFC-792, USC/Information Sciences Institute, Sept. 1981.

[3] Hinden, R., Sheltzer, A., "The DARPA Internet Gateway", RFC-823, Bolt, Beranek and Newman Inc., Sept. 1982.

[4] Mills, D.L., "Exterior Gateway Protocol Formal Specification", RFC-904, M/A-COM Linkabit, Apr. 1984.

[5] Romano, S., Stahl, M., "Internet Numbers", RFC-1020, Stanford Research International, Nov. 1987.

[6] GADS, "Toward an Internet Standard Scheme for Subnetting", RFC-940, DARPA Internet Gateway Algorithms and Data Structures Task Force, Apr. 1985

[7] Mogul, J., Postel, J., "Internet Standard Subnetting Procedure", RFC-950, Stanford University and USC/Information Sciences Inst., Aug. 1985.

[8] CCITT, Data Communication Networks. Transmission, Signalling and Switching, Network Aspects, Maintainance, Administrative Arrangements, Recommendations X.40 - X.180, Yellow Book, Vol. VIII - Fascicle VIII.3, Geneva, 1981

Carl-H. Rokitansky is with the Fern University of Hagen where he is engaged in research on computer networks and their protocols. He is particularly interested in the design, modeling and analysis of routing algorithms for large mobile networks, and in developing gateway algorithms for the interconnection of computer networks. Until recently he joined the Telecommunications Institute of the German Aerospace Research Establishment (DFVLR). He received his JD from the University of Salzburg and his MS degree in economics (special field mathematical statistics and computer science) from the University of Vienna. In 1986/1987 Dr. Rokitansky was a faculty member in the Electrical and Computer Engineering department of the University of Kansas in Lawrence. He is a member of the Internet Engineering Task Force and is chairing the Public Data Network Routing working group of the IETF.

HIERARCHICAL VAN-GATEWAY ALGORITHMS AND PDN-CLUSTER ADDRESSING SCHEME FOR WORLDWIDE INTEROPERATION BETWEEN LOCAL TCP/IP NETWORKS VIA X.25 NETWORKS

Carl-Herbert Rokitansky

Fern University of Hagen
Data Processing Techniques
D-5860 Iserlohn, FRG
roki@DHAFEU52.BITNET
roki@A.ISI.EDU

Within the last few years, the DARPA Internet protocol suite (TCP/IP, FTP, SMTP, TELNET, etc.) has developed into a de facto industry standard for heterogeneous packet-switching computer networks. However, the current Internet gateway architecture does not provide dynamic algorithms to route Internet datagrams between local TCP/IP networks via the system of X.25 public data networks (PDN).

In this paper, the application of the Internet cluster addressing scheme to the system of X.25 public data networks is described and hierarchical VAN-gateway algorithms for worldwide network reachability information exchange between local TCP/IP networks are presented. In addition, the mapping between Internet network numbers and X.121 Data Network Identification Codes (DNICs) is discussed and PDN routing algorithms are described. The presented concept of the PDN-cluster addressing scheme and the described VAN-gateway algorithms provide a basis for the routing of Internet datagrams through X.25 public data networks, and would therefore allow worldwide interoperation between the many local-area networks in various countries, now using DARPA Internet TCP/IP protocols.

1. INTRODUCTION

The Internet system provides packet transport by means of a datagram service for hosts subscribing to the DARPA Internet protocol suite. Currently, in the DARPA Internet several thousand hosts are connected to over 550 networks, using over 300 gateways. In addition, there are several thousand registered networks [RFC1020], and an unknown number of unregistered local area networks, using DARPA Internet TCP/IP protocols, which are not interconnected so far.

This work was granted by the German BMFT and the German automotive industry under subcontract No TV8815-7.

The basic datagram protocol is the Internet Protocol (IP) [RFC791]. Error reporting, flow control, first-hop gateway redirection and other control functions are provided by the Internet Control Message Protocol (ICMP) [RFC792]. Internet transport layer protocols are the Transmission Control Protocol (TCP), which provides reliable end-to-end data stream service, and is equivalent to ISO/OSI Transport Protocol class 4 (TP4), while the much simpler User Datagram Protocol (UDP) is equivalent to the ISO/OSI Transport Protocol class 0 (TP0). All Internet user level protocols use either TCP/IP (e.g. FTP, TELNET, SMTP) or UDP/IP (e.g. NAMESERVER) as the basic packet transport mechanism. Due to the widespread implementation under various operating systems (currently about 180), these protocols have developed into a de facto industry standard for heterogeneous packet-switching computer networks.

The Internet model includes constituent networks, called local networks, to distinguish them from the Internet system as a whole. These local networks are connected together by means of Internet gateways. Each gateway is connected to two or more networks by a physical interface and has an address on each of the local nets between which it provides a datagram transport service. Gateways belonging to different gateway systems ("autonomous systems") might use different intra-system routing mechanisms. In order to maintain the routing tables, an interior gateway protocol like the Gateway-Gateway Protocol (GGP) [RFC823], can be used to exchange routing information between gateways of the same autonomous system, while the Exterior Gateway Protocol (EGP) [RFC904] is used to exchange network reachability information with gateways belonging to a neighbor system.

Internet network numbers are assigned [RFC1020] to networks that are connected to the DARPA-Internet and DDN-Internet, and to independent networks that subscribe to the DARPA Internet protocol suite. Currently more than 10.000 networks numbers are officially assigned.

In the USA, the ARPANET, the NSF-Net, the WIDEBAND, etc. serve as the backbone of the DARPA Internet system, currently consisting of around 550 networks. In Europe, the situation is completely different: Hundreds of local networks are using TCP/IP protocols, but most of them are not connected to any other Internet network. The only considerable network, by which these stand-alone TCP/IP networks could be interconnected worldwide, would be the system of X.25 public data networks (PDN), but no dynamic PDN-routing and VAN-gateway algorithms have been developed so far.

In this paper the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks (hereafter referred to as PDN), is discussed and the new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is presented. This concept allows the implementation of improved algorithms to connect Internet networks to the DARPA Internet as well as to interconnect independent networks together by routing Internet datagrams via "VAN-gateways" through X.25 public data networks and to Internet hosts which are directly attached to a PDN ("Internet/PDN-host").

In Section 2 an outline of the current Internet routing model is given. In Section 3, we discuss the application of the Internet cluster addressing scheme to the international system of X.25 public data networks and the use of a cluster-mask for routing decisions within the PDN-cluster. In addition, the mapping between Internet PDN-cluster network numbers and DNICs is described. In Section 4, we present the new concept of hierarchical VAN-gateway algorithms, which are used for worldwide Internet network reachabilty information exchange between VAN-gateways. Also, the use of a route server for routing Internet datagrams through the PDN is considered. Finally, in Section 5, we summarize the advantages of the PDN-cluster addressing scheme and we specify the requirements for VAN-gateways, which are necessary to allow worldwide intercperation between local TCP/IP networks, by routing Internet datagrams through the international system of X.25 public data networks.

2. CURRENT INTERNET ROUTING MODEL

The current Internet routing model assumes that the route to a host can be computed by an algorithmic transformation on the destination address. The routing algorithms conforming to this model compute only a single (shortest) path from a given gateway to a given destination network, based on some metric such as hopcount or delay, without dependence on costs, type of service or anything else. Also, it is important to understand that the computed route does not depend on any network parameter (e.g. delay, costs, hopcount, link quality) within the destination network.

According to this "network-centric" routing model all hosts in the Internet must make the following routing decision when sending a datagram: Is the datagram addressed to a host on a directly connected network ("local" network) and can therefore be sent directly, or is it addressed to a host on a different Internet network ("foreign" network) which is reachable only via a local gateway. Note, that if subnets are in use (see [RFC940] and [RFC950]) only hosts/gateways on the same subnet appear to be local, while all other destinations are assumed to be reachable only indirectly via a (sub)gateway.

3. INTERNET PDN-CLUSTER ADDRESSING SCHEME FOR X.25 PUBLIC DATA NETWORKS

The international system of X.25 public data networks (PDN) is a typical wide-area network (WAN). In this system, the national packet-switched data networks in various countries are connected via gateways (CCITT, Rec. X.75 [CCITT]) to allow international interworking between hosts on different national public data networks over international virtual circuits.

3.1 X.121 Addressing

An X.121 address (CCITT, Rec. X.121 [CCITT]) is assigned to each PDN host. This international data number consists of the Data Network Identification Code (DNIC, fixed at 4 digits) and the Network Terminal Number (NTN, up to 10 digits).

The system of data network identification codes (DNICs) as specified in Rec. X.121 [CCITT] provides a theoretical maximum of 6000 (resp. 8000) DNICs. However, only about 200 national public data networks are in operation so far.

Currently an Internet class A network number [14.rrr.rrr.rrr] is assigned to the system of public data networks (PDN). For the time being the assignment of Internet addresses to hosts (gateways) on PDN is done successively in (chronological) order of request, regardless to which national public data network a host (gateway) belongs.

Connectivity with the Internet is provided by so called "VAN gateways", which are attached to the national public data networks.

3.2 Characteristics

The PDN can be characterized as follows:
- Wide-area network
- Complex of national public data networks
- International virtual circuits are provided
- Different charges for international and national virtual circuits
- Charges depend on length of time and data volume transferred
- No broadcasting and no multicasting

3.3 Routing Through PDN - Current Situation

Due to the characterization above, the following requirements seem to be reasonable for the routing of Internet datagrams through PDN:

- Routing decisions should be made with regard to the structure of the PDN (complex of national public data networks)
- Packets between PDN hosts should be sent directly through the PDN over virtual circuits

The current routing situation in the DARPA Internet with regard to PDN is very poor:

Due to the assignment of a class A network number to the system of X.25 public data networks (PDN), no internal structure of this PDN system is visible to the outside world.

No parameters (e.g. charges) within the destination network are taken into account in Internet routing decisions.

The nearest VAN-gateway, which declares the PDN system as reachable through it, will receive all data traffic for the PDN. If this VAN-gateway does not make (international) X.25 calls actively, all packets for hosts reachable only via the PDN will be discarded, unless an (international) virtual circuit, which was established previously by the destination host (or another VAN-gateay), is still open.

3.4 Motivation for a Clustering Scheme

Consider the system of X.25 Public Data Networks (network P), to which the Internet networks A and B are connected by means of gateways GAP and GBP as shown in Figure 3.4-1. Host HA is attached to network A and hosts HPT and HPD are connected to network P.

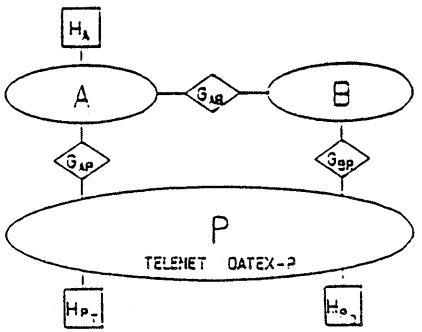

Figure 3.4-1

Direct connections (switched virtual circuits) can be established between the gateways and hosts attached to network P.

According to the Internet routing model, packets from host HA to both HPT and HPD will be routed via gateway GAP (minimum hopcount).

Assume that HPT is a host on the X.25 national public data network TELENET (USA), and HPD is a host on DATEX-P (Germany). Then the charges for a connection between GAP and HPD will be around three times higher than between GBP and HPD. In this case it might be reasonable to route packets from HA to HPD through network B via gateway GBP instead via GAP. Therefore, for routing decisions, the internal structure of the PDN system (network P) would be of interest even outside it. However the current Internet gateway architecture does not provide any algorithms for this situation (due to "network-centric" routing), except that a user on HA could specify GBP explicitly in an IP source route option.

Note, that dividing network P into several subnets according to the internal structure of P would have no external effects, because it is invisible to the outside world as mentioned above. Therefore packets would still be routed from HA to HPD via GAP and not through network B.

Another idea would be to assign different Internet network numbers to subdivisions (e.g. T and D, see Figure 3.4-2) of P, instead of assigning a single Internet network number to network P. Thus, the PDN would become a complex of several Internet networks, and the structure of the PDN would be visible outside of it.

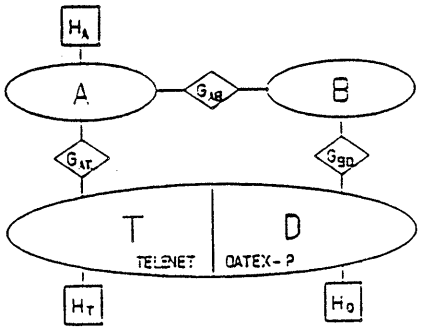

Figure 3.4-2

However this assignment would be inconsistent with the current Internet routing model, because there are packets to "foreign" networks which need not be routed via a local gateway but can be sent directly: Consider packets to be sent from HT (former HPT) to HD (former HPD). Note that HT and HD are now hosts attached to different Internet networks, but direct connections can still be established between HT and HD without transiting an Internet gateway! According to the current routing algorithms, HT would determine that the destination HD is on a different ("foreign") Internet network and would therefore decide that the packets must be sent to a local gateway on the common network T. There is only one gateway GAT connected to network T, but transmitting packets to GAT would be unreasonable, since they can be sent directly to HD. But HD is neither a gateway nor is it a host on the local net. Similarly, gateway GAT would encounter the same problems in its routing decision as HT when forwarding packets to HD. In addition, gateway GAT cannot send an ICMP Redirect message to HT specifying HD as a better first hop on the route towards the destination, because HD is not a host on the same network.

Therefore, the following model of a clustering scheme is proposed [ROKI88], which adds an additional level to the interpretation of Internet addresses and is called "Cluster Addressing Scheme":

3.5 Cluster Addressing Scheme Model

Specific Internet network numbers are assigned to a set of nets between which direct connections can be established without transiting a gateway. These networks are associated to an "Internet Cluster". For all routing decisions within the cluster, and for the specification that different Internet networks are associated to a cluster, the use of an address-mask, called "Cluster-Mask", is proposed. By means of this cluster-mask, all hosts within the same cluster, even if attached to different Internet networks appear to be local. ICMP Redirect messages can be sent directly between gateways and hosts belonging to the same cluster. The fact that several Internet networks are associated to an Internet cluster is invisible to the outside world. However the internal structure of the cluster, which is a complex of Internet networks ("cluster-nets"), is visible outside the cluster.

The 32-bit INTERNET address consists of:

<INTERNET address> ::= <network-number><rest-field>

Now the cluster addressing scheme proposes that the <network-number> field is interpreted as

<network-number> ::= <cluster-number><cluster-net-number>

Although this subdivision of the <net-work-number> field is used for routing decisions within the cluster, it is invisible to the outside world.

Thus, if this clustering scheme is in use, the INTERNET address can be interpreted as

<INTERNET address> ::= <cluster-#><cluster-net-#><rest-field>

3.6 PDN-Cluster Addressing Scheme

To allow an improved routing of Internet datagrams through PDN according to the Internet routing model, the application of the cluster addressing scheme to the system of X.25 public data networks is proposed:

- Internet class B network numbers (with identical bits in the first (high-order) 6-bit field of the Internet address) are assigned to national public data networks.
- The national public data networks are associated to an Internet cluster ("PDN-Cluster")
- For the specification of this cluster and for routing decisions within the cluster, a cluster-mask is used (value <252.0.0.0>), thus all hosts within the PDN-cluster appear to be reachable "locally".
- ICMP Redirect messages can be sent to any PDN host to manage the routing within the PDN-cluster.

NOTE: No changes to the existing Internet gateway system are necessary to support the cluster addressing scheme other than reserving a set of class B network numbers for the PDN-cluster and implementing this scheme on VAN-gateways and PDN-hosts. On hosts supporting subnets [RFC950] this can be done very easily by simply setting the address-mask to the value of the cluster-mask.

Consider the example of network P (system of X.25 public data networks) as discussed above (Fig. 3.4-1). Now, according to the cluster addressing scheme specific Internet network numbers [P1] and [P2] are assigned to national public data networks as shown in Figure 3.6-1. These Internet networks P1 and P2 are associated to an Internet cluster (PDN-cluster). If packets are to be sent from HP1 to HP2, host HP1 can determine by means of a cluster-mask that HP2 is a "local" host since it belongs to the same cluster. Therefore, the packets can be routed directly ("locally") by establishing a direct connection (switched virtual circuit, SVC) to HP2. If the packets are sent for some reason to gateway GAP1, then, according to the clustering scheme, GAP1 can send an ICMP Redirect message to HP1 specifying HP2 as a better first hop in this message.

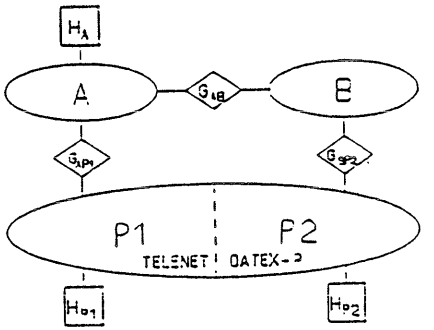

Figure 3.6-1

3.7 Mapping between Internet Network Numbers and DNICs

Due to reasons of homogenity (cost structure, available network options, etc.) a mapping between Internet network numbers and Data Network Identification Codes (DNICs, Recomm. X.121 [CCITT]) is proposed.

Therefore Internet class B network numbers with identical bits in the first (high-order) 6-bit field of the Internet address (see below) are assigned to the different national public data networks. This allows a maximum number of 65.536 PDN-hosts/ VAN-gateways on each network.

The national public data networks are associated to an Internet cluster (PDN-cluster).

For this reason the 16-bit <network-number> field (class B network) is divided into an n-bit <cluster-number> field and a (16 minus n)-bit <cluster-net-number> field:

```
<network-number> ::= <cluster-number><cluster-net-number>
                        (n bits)    +   (16-n bits)
```

In deciding how many bits of the <network-number> field should be used for the <cluster-number> field it seems to be reasonable to distinguish between:

- Theoretical number of addressable DNICs (see Recomm. X.121 [CCITT])
- Number of reachable DNICs currently (and in the near future)

The system of Data Network Identification Codes (DNICs) as specified in Rec. X.121 [CCITT] provides a theoretical maximum of 6000 (resp. 8000) DNICs.

However, only about 200 different national public data networks are reachable currently.

According to the "International Numbering Plan for Public Data Networks", Rec. X.121 [CCITT], the world is divided into six zones. The following table shows how many Data Network Identification Codes (DNIC) are already assigned to national public data networks, how many Internet PDN-cluster networks should be reserved for each zone, how many spares are available, and which Internet network numbers should be reserved for the PDN-cluster.

Zone	Area	assigned	reserve	spare	Internet/PDN-Cluster Networks
2	Europe	49	256	207	[188.001.r.r] - [188.254.r.r]
3	North America	39	256	217	[189.001.r.r] - [189.254.r.r]
4	Asia	40	192	152	[190.001.r.r] - [190.191.r.r]
5	Pacific	20	64	44	[190.192.r.r] - [190.254.r.r]
6	Africa	8	64	56	[191.192.r.r] - [191.254.r.r]
7	South America	44	192	148	[191.001.r.r] - [191.191.r.r]
Total		200	1024	824	

Therefore, it is proposed to use a 6-bit <cluster-number>-field for the PDN-cluster and a 10-bit <cluster-net-number>-field:

```
<network-number> ::= <cluster-number><cluster-net-number>
                        (6 bits)    +     (10 bits)
```

This allows to address 1024 (1016) different national public data networks with 65.536 PDN hosts on each.

According to the cluster addressing scheme, the reservation of Internet network numbers for an Internet cluster should start with the highest, not yet assigned network numbers of each class.

Therefore, the assignment of the cluster-numbers [188.n.r.r]-[191.n.r.r] to the PDN-cluster is proposed, thus to reserve the Internet network-numbers [188.001.rrr.rrr] up to [191.254.rrr.rrr] for the different national public data networks.

3.8 PDN-Cluster-Mask

For the specification of the PDN-cluster and for internal routing decisions within the cluster, corresponding to the width of the <cluster-number> field [ROKI88], a cluster-mask, called "PDN-cluster mask" is used, in which the first (high-order) 6 bits are set to "one", while the remaining bits are set to "zero" (value <252.0.0.0>).

```
<network-number>< rest - field >
<clus><clus.net>< rest - field >

11111100000000000000000000000000   binary
    252.      0.      0.      0  decimal
                    PDN-cluster mask
```

By means of this cluster mask a host can determine if it is connected to a cluster-net. Each PDN-host/VAN-gateway uses this mask for the routing decision if the destination IP-address specified in a datagram is either "local" or "foreign" depending whether the destination is in the same cluster or not. All datagrams to local destinations (even on different Internet networks (cluster-nets)) can be sent directly to the destination without transiting an Internet gateway.

Thus all PDN-hosts/VAN-gateways within the PDN-cluster appear to be reachable "locally". In fact, direct national or international virtual circuits can be established between PDN hosts.

If an IP implementation supports subnets, (normally) no changes to the code are necessary to specify the PDN-cluster, except that "my_ip_mask" must be assigned the value of the "PDN-cluster mask" <252.0.0.0>.

3.9 ICMP Redirect Messages

Due to the fact that all hosts and gateways within the same cluster appear to be reachable "locally", the cluster addressing scheme allows to send ICMP Redirect messages between gateways and hosts within the same cluster and not only within a

directly connected Internet network. This is a significant extension of the usage of the ICMP Redirect message. Therefore, to manage the routing within the PDN-cluster between PDN-hosts and VAN-gateways, ICMP Redirect messages can be sent to any host/gateway in the PDN-cluster, specifying any VAN-gateway/PDN-host as a better first-hop towards the destination in this redirect message.

3.10 Advantages and Disadvantages of the PDN-Cluster Addressing Scheme

The concept of associating a set of Internet networks to an Internet/PDN-cluster and the specification of a PDN-cluster mask has the following advantages:

- The internal structure of the system of X.25 PDNs, consisting of national public data networks, becomes visible to the outside world. This is important for routing decisions outside the cluster.
- The fact that an Internet/PDN-cluster has been formed is invisible outside the cluster. Therefore, no changes to the existing Internet gateway system are necessary to support the PDN-cluster addressing scheme.
- All hosts (gateways) within the PDN-cluster appear to be reachable directly ("locally"). This is important for routing decisions within the PDN-cluster.
- No (or minor) changes to hosts supporting subnets are needed only
- ICMP Address Mask Request and Address Mask Reply messages can be exchanged to determine which cluster-mask is in use.
- ICMP Redirect messages can be used within the PDN-cluster between VAN-gateways and PDN-hosts on different Internet networks

The requirement to reserve specific Internet network numbers for the PDN-cluster might be regarded as a disadvantage.

4. HIERARCHICAL VAN-GATEWAY ALGORITHMS

In order to route Internet datagrams to the destination network, Internet gateways must exchange routing and network reachability information ([RFC823] and [RFC904]). Due to the fact, that the system of X.25 public data networks (PDN) is a wide-area network, with no broadcasting feature, and with a complex tariff structure for national and international calls, new concepts of hierarchical Internet/VAN-gateway algorithms must be developed to provide worldwide interoperability between TCP/IP networks by routing Internet datagrams through PDNs. These hierarchical VAN-gateway algorithms are presented in this section.

4.1 Modified EGP Between VAN-Gateways

To advertise the reachability of PDN-cluster-networks and of networks beyond the PDN, EGP messages [RFC904] can be exchanged between VAN-gateways and other Internet gateways.

However, if a particular network does not have any external connectivity (outside the PDN) to the rest of the Internet and is reachable only via a VAN-gateway, it will be necessary to exchange network reachability information between VAN-gateways through the PDN. For this purpose (due to charges) a modified version of EGP on an event driven basis could be used, in which neither "Hello" packets nor "Network Reachability Updates" are sent periodically. After the "Neighbor Acquisition" phase, network reachability updates are exchanged only on an event driven basis, i.e. only if changes in network reachability occur. (Specification of such a protocol is outside the scope of this paper and will be discussed elsewhere)

4.2 Worldwide Internet Network Reachability Exchange Between VAN-Gateways

For a dynamic routing of Internet datagrams through the system of X.25 public data networks, a worldwide exchange of network reachability information between VAN-gateways is required. It is expected that, according to the number of local TCP/IP networks, which will be connected through the PDN, hundreds of VAN-gateways on the national public data networks in various countries will take part in this process. Due to the expected number of participating VAN-gateways and the requirement for worldwide communication, in order to limit the number of direct neighbor gateways of each VAN-gateway, and to minimize the amount of network reachability information, which has to be exchanged between VAN-gateways, hierarchical gateway algorithms are used, to maintain and update the routing tables of VAN-gateways.

First, consider a 2-level hierarchy for VAN-gateways as shown in Figure 4.2-1. Each level-1 VAN-gateway (V1D, V1I, V1T) on a national public data network reports on an event driven basis those local networks, which are reachable through it to a level-2 VAN-gateway on the same national public data network (V2D, V2I, V2T). Level-1 gateways may or may not receive worldwide Internet network reachability information from a level-2 VAN-gateway. If not, they use the level-2 gateway as the default gateway; if it is necessary, the level-2 gateway replies with a ICMP Redirect message to specify a better first hop VAN-gateway within the PDN-cluster. Normally, level-1 gateways do not exchange network reachability information among each other and with level-2 gateways on other national public data networks.

To each national public data network (at least) one Internet level-2 VAN-gateway is attached. Each level-2 gateway receives network reachability information from level-1 gateways on the same national public data network and exchanges this information with level-2 gateways on different national public data networks. In this scheme each

VAN-gateway has only one or two "active" neighbor gateways, to which switched virtual circuits (SVC) have to be opened actively by an X.25 "Call Setup Request", but Internet network reachability information is distributed worldwide. Note, that each VAN-gateway might have several "passive" neighbors (SVC opened by the neighbor gateway). In the example in Figure 4.2-1, the level-2 VAN-gateway V2D receives network reachability information from V1D about networks M and N, and reports the network reachability for networks D,L,M and N to level-2 gateway V2I, from which it receives worldwide network reachability information about networks I,F,G,H,K and (due to information exchange between V2I and V2T) also for T,A,B,C,E.

Now assume, that packets are to be sent from a host on network N (HN) to a host on network A (HA). These packets are routed to V1D. If V1D uses V2D as the default VAN-gateway for foreign networks (networks outside the PDN-cluster, D,I,T), then, according to the PDN-cluster addressing scheme, an ICMP Redirect message will be sent by V2D to V1D, specifying V2T (being in the same cluster) as a better first hop towards the destination. After opening a SVC to V2T, the packets will be routed to HA via V2T and network B.

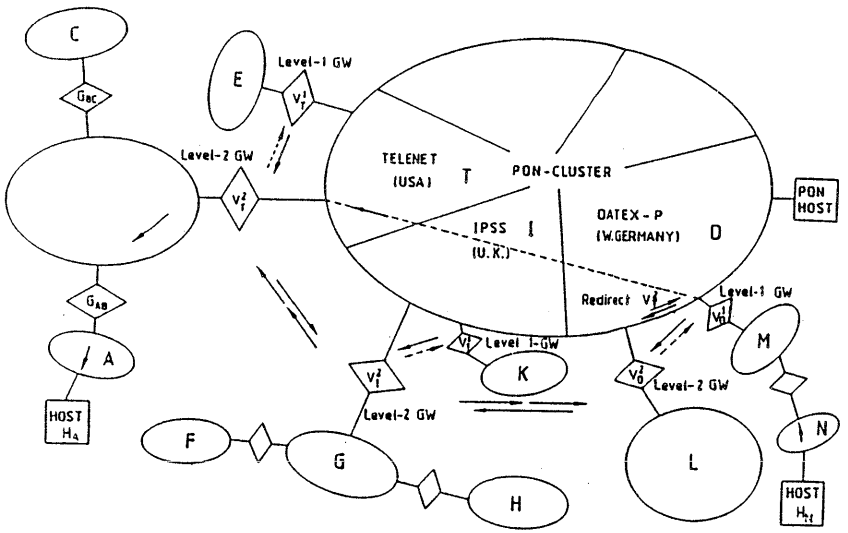

Figure 4.2-1

4.3 Route Server for Network Reachability Information

To avoid that network reachability for thousands of Internet networks must be distributed worldwide to all Internet gateways, although connections to a host on a network in another continent are established relatively seldom, the implementation of route servers for countries or zones might be reasonable.

Consider Figure 4.3-1 and assume that instead of only three networks A,B and C, there are 500 US-Internet networks behind level-2 VAN-gateway V2T. In this case the network reachability information about networks I,F,G,H,K and D,L,M,N is to be distributed to hundreds of Internet gateways in the USA, although connections from most US-networks to Europe are established relatively seldom. If a route server is implemented on a host on network C, then network reachability information about the European networks must be sent only to this route server. If the route server is implemented on level-2 VAN-gateway V2T, then even this information exchange could be avoided.

Assume that packets from a host on network A (HA) should be sent to a host on network N (HN). The packets will be routed to the default gateway GAB. Having no information about network N, this gateway would send a "Route Request" message to the route server, specifying network N as the destination network. The route server would respond with a "Route Reply" message, specifying V2T as the VAN-gateway in the USA through which the destination network N in Europe is reachable. Routed via V2T, V1D, and network M, the packets would finally reach the destination HN.

Now we assume that there are more intermediate gateways between GAB and V2T. In this case, each gateway would have to send a "Route Request" message to the route server. This could be avoided if gateway GAB uses the Internet address of V2T as a Source Route option (modified use, similar to the procedure as described in [ROKI88], 4.10). Then the packets could be routed to V2T without the requirement of requesting an additional route information from the route server by an intermediate gateway.

The described algorithms would have to be implemented in all Internet gateways. The "Route Request" and "Route Reply" messages could probably be defined as new ICMP messages.

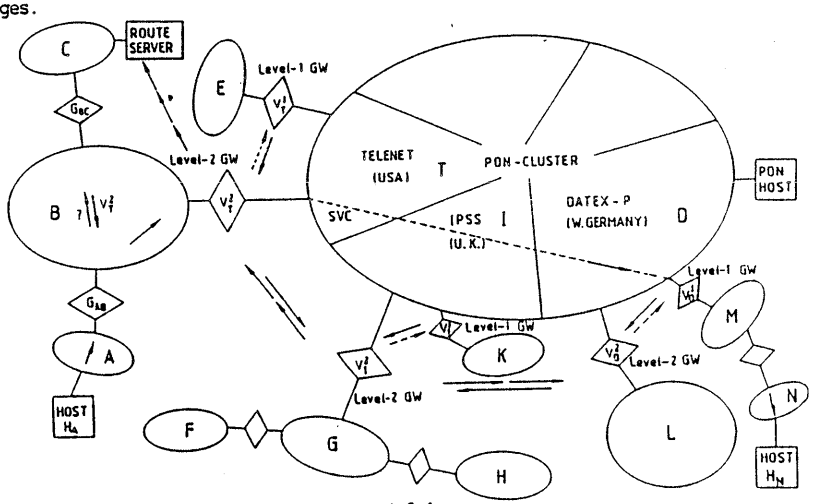

Figure 4.3-1

4.4 VAN-Gateway Level Classes

In the examples above, for reasons of simplicity, only a 2-level hierarchy of VAN-gateways has been considered. However, due to the fact, that currently about 200 national public data networks are in operation, in order to minimize the hopcount required to distribute network reachability information worldwide, additional VAN-gateway levels must be specified. The following four classes of VAN-gateways, according to a 4-level hierarchy, are defined:

- Level-1 LOCAL-VAN-Gateway:
  This gateway connects one or more local TCP/IP networks to the system of X.25 public data networks. It reports information about the local networks to a level-2 gateway, from which it may or may not receive worldwide Internet network reachability information.

- Level-2 DATA-NETWORK-Gateway:
  To each national public data network (at least) one level-2 Data-Network-Gateway is attached. Each level-2 gateway receives network reachability information from level-1 Local-VAN-gateways and reports this information to a level-3 gateway, from which it receives worldwide network reachability information.

- Level-3 COUNTRY-Gateway:
  This gateway reports the network reachability information it receives from all the level-2 Data-Network-gateways on the various national public data networks within the same country, to a level-4 gateway, from which it receives worldwide network reachability information.

- Level-4 ZONE-Gateway:
  This gateway collects the network reachability information from all the level-3 Country-gateways within the same zone, and exchanges worldwide network reachability information with adjacent level-4 Zone-gateways.

According to this hierarchical VAN-gateway scheme, Internet network reachability information is distributed worldwide as shown in Figure 4.4-1.

Note, that the described hierarchical VAN-gateway scheme is only a logical one, and that the same physical VAN-gateway, for example the GMD-VAN in Figure 4.4-1, could act as the Local-VAN-gateway for the GMD Net, as the Data-Network-gateway for DATEX-P, as the Country-gateway for Germany, and finally as the Zone-gateway for Europe.

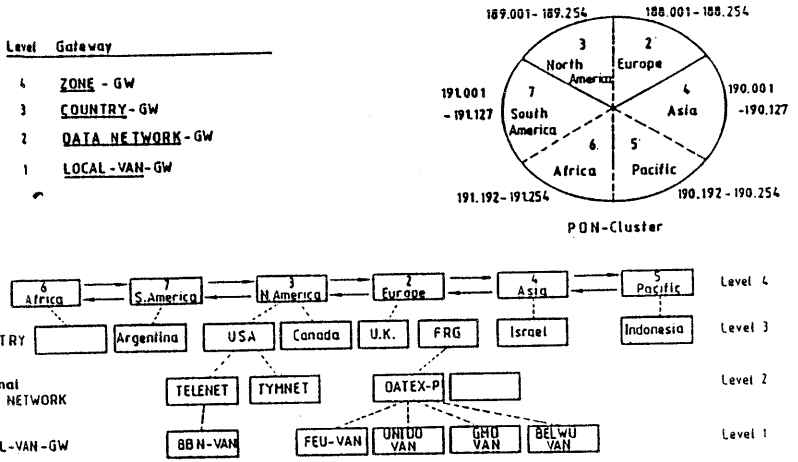

Figure 4.4-1

A modification of the proposed hierarchical VAN-gateway scheme might be reasonable (for example an additional level-5 World-gateway, or an additional Region-gateway between the Country-gateway and the Zone-gateway) and will be discussed in the Public Data Network Routing working group of the Internet Engineering Task Force (IETF), and with PDN-test partners in Europe, the United States and Australia, and possibly in Argentinia, Canada, Indonesia and Japan.

5. SUMMARY

In this paper, the application of the cluster addressing scheme to the system of X.25 public data networks (PDN) has been discussed. Different Internet network numbers are assigned to the national public data networks (about 200 worldwide) and associated to the PDN-cluster. Thus, the internal structure of the PDN, which currently appears to be unstructured, becomes visible to the outside world, which is important for Internet routing decisions. However, the fact that a PDN-cluster has been formed is invisible outside the cluster. Therefore, no changes to the existing Internet gateway system are necessary.

By means of the PDN-cluster mask, which is used for routing decisions, all hosts/gateways within the PDN-cluster (even on different Internet networks) appear to be reachable locally, and in fact, direct virtual circuits between PDN-hosts and VAN-gateways can be established. As a significant extension, ICMP Redirect messages can be sent not only between hosts and gateways on the same Internet network, but within the whole PDN-cluster. In addition, a hierarchical VAN-gateway scheme has

been presented, by which Internet network reachability information can be distributed worldwide very effectively with a few number of hops. The PDN-cluster addressing scheme and the presented hierarchical VAN-gateway algorithms involve no changes to the existing Internet gateway system, and must be implemented only on VAN-gateways and hosts, which are directly attached to the PDN. PDN-hosts whose software supports subnets can be equipped easily with the PDN-cluster addressing scheme. The implementation of the proposed PDN-cluster addressing scheme and the hierarchical VAN-gateway alogorithms would allow worldwide interoperation beween the many local networks in various countries now using DARPA-Internet TCP/IP protocols, in a very short time, since no changes to the existing Internet gateway system are necessary.

ACKNOWLEDGMENT

J. Noel Chiappa, Horst D. Clausen, Dave Mills, Jon Postel, Bernhard Walke and the members of the Public Data Network Routing working group of the Internet Engineering Task Force (IETF) have provided helpful discussions and important suggestions.

REFERENCES

[CCITT] CCITT, Data Communication Networks. Transmission, Signalling and Switching, Network Aspects, Maintainance, Administrative Arrangements, Recommendations X.40 - X.180, Yellow Book, Vol. VIII - Fascicle VIII.3, Geneva, 1981

[RFC791] Postel, J., ed., "Internet Protocol - DARPA Internet Program Protocol Specification", DARPA Network Working Group Request For Comments RFC-791, SC/Information Sciences Institute, Sept. 1981.

[RFC792] Postel, J., ed., "Internet Control Message Protocol - DARPA Internet Program Protocol Specification", RFC-792, USC/Information Sciences Institute, Sept. 1981.

[RFC823] Hinden, R., Sheltzer, A., "The DARPA Internet Gateway", RFC-823, Bolt, Beranek and Newman Inc., Sept. 1982.

[RFC904] Mills, D.L., "Exterior Gateway Protocol Formal Specification", RFC-904, M/A-COM Linkabit, Apr. 1984.

[RFC940] GADS, "Toward an Internet Standard Scheme for Subnetting", RFC-940, DARPA Internet Gateway Algorithms and Data Structures Task Force, Apr. 1985

[RFC950] Mogul, J., Postel, J., "Internet Standard Subnetting Procedure", RFC-950, Stanford University and USC/Information Sciences Institute, Aug. 1985.

[RFC1020] Romano, S., Stahl, M., "Internet Numbers", RFC-1020, Stanford Research International, Nov. 1987.

[ROKI88] Rokitansky, C.-H., "Internet Cluster Addressing Scheme and Its Application to Public Data Networks", in Proceedings of the 9th International Conference on Computer Communication (ICCC'88), pp. 482-491, Tel Aviv, Israel, Oct./Nov. 1988.

National Institute of Standards and Technology

NVLAP

National Voluntary Laboratory Accreditation Program

PROGRAM HANDBOOK

COMPUTER NETWORK INTERFACE PROTOCOLS DoD HIGH LEVEL PROTOCOLS

DRAFT

MARCH 1989

NVLAP PROGRAM HANDBOOK COMPUTER NETWORK INTERFACE PROTOCOLS DoD HIGH LEVEL PROTOCOLS

REQUIREMENTS FOR ACCREDITATION

Jeffrey Horlick

National Voluntary Laboratory Accreditation Program

DRAFT March 1989

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
Gaithersburg, Maryland 20899

Table of Contents

continued....

I. PROGRAM SUMMARY

This document presents the operational and technical requirements of the National Voluntary Laboratory Accreditation Program (NVLAP) related to Department of Defense specifications for communications network interfacing. Explanations of NVLAP's technical requirements indicate how criteria apply to the laboratory accreditation process.

The Defense Communications Agency (DCA) is the provider of the common-user packet-switched communications network, the Defense Data Network (DDN), for the Department of Defense. Systems using any of the five DoD High Level Protocols (Internet Protocol (IP), Transmission Control Protocol (TCP), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and the TELNET Protocol) on the DDN or any DoD Network must have their implementation conformance tested as specified in a memorandum from the Assistant Secretary of Defense for Command, Control, Communications and Intelligence dated 26 August 1988.

In July 1988, DCA requested that the National Bureau of Standards, now the National Institute of Standards and Technology (NIST), establish a NVLAP program to recognize and accredit laboratories that produce reliable test data of vendor products to assure that they conform to the DoD High Level Protocol Standards.

Any laboratory (including commercial, manufacturers', university, Federal, State, or local government) that uses test methods listed in this document may apply for NVLAP accreditation. Accreditation will be granted to a laboratory that complies with conditions for accreditation defined in the NVLAP Procedures: Title 15, Part 7 of the Code of Federal Regulations. The names of NVLAP-accredited laboratories are published in the Federal Register, NVLAP Annual Directories, and other media to which information is regularly provided.

Testing services covered: DoD High Level Protocols: Internet Protocol (IP), Transmission Control Protocol (TCP), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and the TELNET Protocol

Period of accreditation: One year, renewable annually

On-site assessment: Visit by a technical expert to determine compliance with the NVLAP criteria before initial accreditation and every two years thereafter. Monitoring visits as required.

Assessors: Selected from technical experts with experience in the appropriate field(s) of testing.

Proficiency testing: Each laboratory is required to test and analyze a known reference protocol implementation. The completed test report is sent to NVLAP for analysis. Proficiency testing is required for initial accreditation and is conducted annually. Advance notice and instructions are given before testing is scheduled.

II. INTRODUCTION

Background

The U.S. Department of Commerce, National Institute of Standards and Technology (NIST), formerly the National Bureau of Standards (NBS), administers the National Voluntary Laboratory Accreditation Program (NVLAP). NVLAP's function is to accredit public and private testing laboratories based on evaluation of their technical qualifications and competence for conducting specific test methods in specified fields of testing. Accreditation is granted on the basis of conformance with criteria published in the Code of Federal Regulations as part of the NVLAP procedures (15 CFR Part 7) (see Appendix A).

This document is intended for information and use by staff of accredited laboratories, those seeking accreditation, other laboratory accreditation systems, and others needing information on the requirements for accreditation under this NVLAP program. This document is generally included in the NVLAP Application Package along with General Application Forms, Test Method Selection Lists, and other materials needed to apply for or renew accreditation. It presents the administrative and operational procedures and technical requirements of the accreditation program and should be retained and be readily accessible to laboratory personnel.

NVLAP Accreditation

NVLAP accreditation is available to commercial laboratories, manufacturers' in-house laboratories, university laboratories, Federal, State, and local government laboratories. Foreign-based laboratories may be accredited by NIST if they meet the same requirements as domestic laboratories and pay any additional fees required (see Appendix D).

NVLAP is self-supporting and operates on a cost reimbursable basis by charging fees to those who pursue accreditation. The program receives no appropriated public funds.

Accreditation is granted only after thorough evaluation of the applicant has demonstrated that all NVLAP criteria have been met. Laboratories which successfully demonstrate compliance with the criteria are issued two documents to attest to that compliance: (1) a Certificate of Accreditation, and (2) a Scope of Accreditation which states the specific test methods and services for which the laboratories has been accredited (see Appendix I).

Why NVLAP Accreditation?

A laboratory may wish to be accredited for many reasons such as; legal requirements, regulations or codes, contract specifications, or the desire to be recognized as demonstrably competent to meet the needs of its clients.

NVLAP provides formal recognition of the competence of accredited laboratories to the user community. Information about accredited laboratories, including the name and scope of accreditation, is disseminated in various media.

For accreditation to be meaningful, it must be granted by a clearly credible organization. NVLAP provides an unbiased third party evaluation and recognition of performance as well as expert technical assistance to upgrade laboratory performance.

Testing Laboratory Defined

NVLAP defines "testing laboratory" as an organization that provides services to measure, examine, test, calibrate, or otherwise determine the characteristics or performance of materials, products or systems. See Appendix E for information about laboratory main and sub facilities.

Accreditation Defined

NVLAP accreditation signifies recognition of a testing laboratory's competence to perform specific test methods in specified fields of testing. It means that the laboratory's quality system, staff, facilities and equipment, calibration procedures, test methods and procedures, records, and test reports have all been evaluated and found to meet NVLAP criteria. NVLAP accreditation does not mean a guarantee (certification) of laboratory performance or of product test data; it is solely a finding of laboratory competence.

NVLAP Programs

Laboratories may participate in as many NVLAP programs as they wish, provided that they meet all NVLAP criteria for each program. Programs currently available are:

- Acoustical Testing Services
- Asbestos in Bulk Insulation and Air Carpet
- Commercial Products Testing
  - Paints and Coatings
  - Paper and Related Products
  - Plastics
  - Seals and Sealants
- Computer Network Interface Protocols
- Construction Testing Services
- Electromagnetic Compatibility and Telecommunications
- Personnel Radiation Dosimetry
- Thermal Insulation Materials
- Wood Stoves

For further information about NVLAP, or for assistance in understanding and meeting the NVLAP requirements and criteria, please write or call:

NVLAP
National Institute of Standards and Technology *
ADMIN A527
Gaithersburg, MD 20899
Phone: (301) 975-4016

* formerly the National Bureau of Standards (NBS)

III. OPERATIONAL INFORMATION AND REQUIREMENTS

The information and requirements presented in this section are generally applicable to all NVLAP programs. Technical and proficiency testing requirements presented in subsequent sections are specifically applicable to the field(s) of testing covered by this Handbook.

Laboratory Code Number (LAB CODE)

Each participating laboratory is assigned a four-digit laboratory code number. The code number is used by the NVLAP staff for identification, filing, recordkeeping, and database management. Participants are requested to put their Lab Code number on all correspondence with NVLAP. The Lab Code number is cross-referenced with the laboratory name and location in the NVLAP Directory of Accredited Laboratories.

Accreditation Period

Accreditation is granted for a period specified in the Accreditation Application Package (usually one year). The accreditation period begins on one of four dates: January 1, April 1, July 1, or October 1. Once a laboratory has been assigned an accreditation date, it retains that date as long as it remains in the program. Accreditation both expires and is renewed on that date.

Approved Signatory

The laboratory must designate one or more staff members as Approved Signatories. The name of at least one Approved Signatory must appear on all test reports endorsed with the NVLAP logo (see section Use of the NVLAP Logo elsewhere in this Handbook). This person is responsible for the technical contents of the report and is the one to be contacted by NVLAP, laboratory clients, or others in case of questions or problems with the report.

There is no formal requirement for nomination or approval of persons designated as Approved Signatories. The laboratory must inform NVLAP of its appointments by completing the appropriate sections in the General Application for accreditation. Approved Signatories should be persons with adequate responsibility or authority within the organization, with adequate and appropriate technical capabilities, and without conflict of interest.

Laboratory test reports carrying the NVLAP logo need not be signed individually by the Approved Signatory. Test report forms may be preprinted with the required information. Forms that are electronically or computer generated may have the information printed along with the test results.

Authorized Representative

The laboratory must designate an Authorized Representative to sign the application form and commit the laboratory to fulfill the NVLAP requirements. The Authorized Representative is the only one who can authorize a change in the scope or nature of the laboratory's application. The Authorized Representative may also be an Approved Signatory.

Renewal

Each participating laboratory will be sent a renewal Application Package, well in advance of the expiration date of its accreditation, to allow sufficient time to complete the renewal process. The technical requirements and fees for renewal are generally the same as for initial accreditation.

The application and fees must be received by NIST prior to expiration of the laboratory's current accreditation to avoid a lapse in accreditation. If an on-site assessment is required, the application and fees must be received to allow sufficient time for the visit to be completed and deficiencies corrected prior to expiration of accreditation. In addition, any current proficiency testing requirements must be met.

Keeping NVLAP Informed

During the accreditation period, a laboratory must inform NVLAP:

of any major changes involving the location, ownership, management structure, authorized representative, technical director, approved signatories, or facilities;

if it wishes to delete a test method; or

if it is no longer capable of performing test methods or services for which it is accredited.

If a laboratory elects not to renew or wishes to voluntarily terminate its accreditation at any time, the notification of such intention should be forwarded to NVLAP in writing.

Additions to Scope of Accreditation

During the accreditation period, a laboratory may request the addition of test methods or services to its Scope of Accreditation. The laboratory must meet all NVLAP criteria for the additional test methods or services such as fees, proficiency testing, technical requirements, etc. The need for an additional on-site assessment will be determined on a case-by-case basis.

NVLAP Directory

NVLAP publishes an annual Directory of Accredited Laboratories. The Directory contains the name and address, scope of accreditation, contact person, and the accreditation renewal date for each accredited laboratory. Supplements to the Directory are published quarterly to cover interim accreditation actions including initial accreditations, renewals, suspensions, terminations, and revocations. The Directory is distributed nationally and internationally to manufacturers, suppliers, retailers, professional and trade associations, code groups, and government agencies.

Referencing Your Accredited Status and Use of the NVLAP Logo

Accredited laboratories are encouraged, within specified limits, to announce their accredited status. The NVLAP logo may be used in such announcements. Photographic copies of the logo are available from the NVLAP office.

A laboratory must limit the representation of the scope of its accreditation to only those tests or services for which accreditation has been granted. The following statement is recommended: "Accredited by the National Institute of Standards and Technology, National Voluntary Laboratory Accreditation Program for selected test methods or service."

In Advertising

Laboratory advertising of accredited status must be limited to professional, technical, trade, or other laboratory services publications. Letterhead referencing NVLAP accreditation may be used in direct solicitation for business from potential customers. It is recommended that a copy of the NVLAP Certificate and Scope of Accreditation be appended to such a solicitation.

News stories and advertising by laboratories of their accredited status in the trade press is permissible and encouraged. The use of advertisements in the trade press is consistent with NVLAP procedures.

Laboratories may not reference their accredited status in consumer media, in product advertising, or on product labels, containers and packaging. The nature or type of product advertising prohibited by NVLAP procedures includes any advertising that is intended to encourage a consumer to purchase a product because it was tested by an accredited laboratory, whether that advertising appears in consumer media, the business media, or at a point of sale to consumers. Advertising must not imply product certification by NVLAP, NIST, or the U.S. Government.

On Laboratory Documents

As long as a laboratory is NVLAP accredited, it may use the NVLAP logo on letterhead and brochures, preferably with the qualifying quote given above. The logo may be used on test reports that are within the scope of accreditation. These reports must bear the name of an Approved Signatory in accordance with the guidelines given in the Approved Signatory section of this Handbook.

Compliance With Existing Laws

Accreditation does not relieve the laboratory of the need to observe and comply with existing Federal, State, and local statutes, ordinances, or regulations that may be applicable to its operations, including consumer protection and antitrust laws.

IV. TECHNICAL EXPERTS

NVLAP uses Technical Experts (TEs) as assessors and evaluators. They may be engineers or scientists currently active in the field, consultants, college professors or retired persons. They are selected on the basis of their professional and academic achievements, experience in the field of testing, management experience, and tact in dealing with people. Their services are generally contracted as required; they are not NVLAP staff members.

Assessors are TEs selected to conduct an on-site assessment of a particular laboratory on the basis of how well their individual experience matches the type of testing to be assessed, as well as absence of conflicts of interest. The laboratory has the right to appeal the assignment of an assessor and may request an alternate.

Evaluators are TEs selected to review the record of the laboratory as a whole, including the application, assessment report, deficiencies, corrections to deficiencies, and proficiency test results and, based on this record, to recommend whether or not a laboratory should be accredited. The evaluators are matched to the type of testing being evaluated and are selected to avoid conflicts of interest.

V. ACCREDITATION PROCESS

Accreditation is granted following successful completion of a process which includes submission of an application and payment of fees by the laboratory, an on-site assessment, resolution of deficiencies identified during the on-site assessment, participation in proficiency testing, technical evaluation, and administrative review. The process is described in the following sections.

Application and Fees

An **Application Package** is sent to a laboratory on request. It includes: **General Application Forms, Fee Calculation** forms, and the program Handbook. The **General Application Form** must be completed and signed by the authorized representative of the laboratory. Before completing and signing the application, the authorized representative should review all documents and become familiar with NVLAP requirements.

In general, the accreditation fee is composed of several parts, some of which are fixed while others depend on the scope of accreditation desired and the specifics of the program. The total accreditation fee must be paid before accreditation can be granted. The individual parts of the accreditation fee include, as appropriate: an Administrative and Technical Support fee, a Test Method fee, a Proficiency Testing fee, the cost of reference materials and quality assurance samples, and an On-Site Assessment fee. The fees for this accreditation program are shown in the **Fee Calculation Sheet** included in the **Application Package.**

The laboratory will be contacted to schedule a mutually acceptable date for the on-site assessment after payment of all required fees and will be notified of any additional information which must be supplied, and of any applicable proficiency testing requirements which must be completed, for the technical evaluation.

On-site Assessment

Before initial accreditation and periodically thereafter, an on-site assessment of each laboratory is conducted to determine compliance with the NVLAP criteria. The assessment is conducted by one or more NVLAP assessors selected on the basis of their expertise in the field of testing to be reviewed. Assessors use checklists developed by NVLAP so that each laboratory receives an assessment comparable to that received by others. However, assessors have some latitude to make judgments about a laboratory's compliance with the NVLAP criteria.

Each laboratory will be contacted to arrange a mutually agreeable date for an assessment. An assessment normally takes one to three days depending on the extent of the laboratory's application. Every effort is made to conduct an assessment with as little disruption as possible to the normal operations of the laboratory. During the assessment the assessor will:

- meet with management and supervisory personnel responsible for the laboratory's activities (for which accreditation is being sought) to review the assessment process with the individuals involved and to set the assessment agenda.

- examine the quality assurance system employed by the laboratory. The assessor may select and trace the history of one or more samples from receipt to final issuance of test reports. The assessor will conduct a thorough review of the laboratory's quality manual or equivalent, evaluate the training program, examine notebooks or records pertaining to the samples, check sample identification and tracking procedures, determine whether the appropriate environmental conditions are maintained, and examine copies of completed test reports.

- review records of periodic internal audits, use of check-samples or participation in round robin testing or other similar programs.

- review personnel records including resumes and job descriptions of key personnel, competency evaluations for all staff members who routinely perform the testing for which accreditation is sought, calibration or verification records for apparatus used, test reports, and sample control records.

- observe demonstrations of testing techniques and discuss them with the technical personnel to assure their understanding of the procedures.

- examine major equipment, apparatus, and facilities for appropriateness, capability, adherence to specifications, etc.

At the conclusion of the assessment, the assessor will conduct an exit briefing to discuss observations with responsible laboratory staff. A written assessment report will be left with the laboratory. The assessor will forward the assessment forms and a copy of the report to NVLAP.

Monitoring Visits

In addition to regularly scheduled assessments, monitoring visits may be conducted by assessors or by NIST staff at any time during the accreditation period. The scope of a monitoring visit may range from checking a few designated items to a complete review. Monitoring visits may occur for cause or on a random selection basis. These visits serve to verify reported changes in the laboratory's personnel, facilities, and operations or to explore possible reasons for poor performance in proficiency testing.

Proficiency Testing

Proficiency testing is an integral part of the NVLAP accreditation process. Demonstration of appropriate facilities, equipment, personnel, etc., is essential, but may not be sufficient for a complete evaluation of laboratory competence. The actual performance of tests and reporting of results using special proficiency testing samples provides NVLAP with a way to determine the overall effectiveness of the laboratory (see Appendix C).

Proficiency testing is a process for checking actual laboratory testing performance, usually by means of inter-laboratory comparisons. Each accreditation program has unique proficiency testing requirements. The data are analyzed by NVLAP and summary reports of the results are sent to the participants.

Information obtained from proficiency testing helps to identify problems in a laboratory. When problems are found, NVLAP staff members work with the laboratory staff to solve them. If problems with the test method are suspected, NVLAP provides information to the appropriate standards writing bodies.

The specific proficiency testing requirements for this Program are included elsewhere in this document.

Deficiency Notification and Resolution

A deficiency is the failure of a laboratory to meet a NVLAP criterion. Deficiencies may be determined during on-site assessments, monitoring visits, proficiency testing, NVLAP staff review, and Technical Evaluation. Laboratories are informed of deficiencies during the on-site assessment and through other correspondence.

When a laboratory is notified by NVLAP of deficiencies, the laboratory must respond in writing to NVLAP within 30 days of the notification. The response must provide documentation, signed by the authorized representative, that the specified deficiencies have either been corrected or that specific actions are being taken to make corrections. A timetable for completion of corrections should be included.

A laboratory which is currently accredited must correct all deficiencies noted within 30 days of notification or face possible revocation, suspension, or expiration without renewal of its accreditation.

Test equipment that is identified as deficient should not be used until corrective action has been completed. Evidence of correction must be sent to NVLAP.

If substantial deficiencies have been cited, NVLAP may conduct an additional on-site assessment prior to granting accreditation. All deficiencies and resolutions will be subject to thorough review and corrective actions verified during subsequent assessments and technical evaluations.

Technical Evaluation

When a laboratory is ready for an accreditation action, a final technical evaluation is conducted by experts chosen for their experience and knowledge of the pertinent test methods. They review records on each applicant laboratory and base their evaluation on:

- information provided on the application;
- on-site assessment reports;
- actions taken by the laboratory to correct deficiencies;
- results of proficiency testing; and
- information from any monitoring visits of the laboratory.

If the technical evaluation reveals additional deficiencies, written notification describing them will be made to the laboratory. The laboratory must respond within 30 days of such notification and provide documentation, signed by the authorized representative, that the specified deficiencies have been corrected. Clarification of some issues may be requested by telephone. All deficiencies must be corrected before accreditation can be granted or renewed.

Administrative Review

After the technical evaluation has been completed, the NVLAP staff prepares an administrative recommendation that the laboratory either be granted or denied accreditation. This recommendation is based on a review of the technical evaluation and other records to ensure that all NVLAP technical, financial and administrative requirements have been satisfied.

Accreditation Actions

The following accreditation actions may be taken by NIST:

Accreditation — If accreditation is recommended, the recommendation forms the basis for granting accreditation. A **Certificate of Accreditation** and a **Scope of Accreditation** will be issued to the laboratory.

Denial — If denial is recommended, the laboratory is notified of a proposal to deny accreditation and the reason(s) therefor.

Suspension — If a laboratory is found to have violated the terms of its accreditation, the accreditation can be suspended. The laboratory will be notified of the reasons for and conditions of the suspension and the action(s) that the laboratory must take to have accreditation reinstated.

Revocation — If a laboratory is found to have violated the terms of its accreditation, the laboratory is notified of a proposal to revoke accreditation and the reasons therefor. The laboratory may be given the option of voluntarily terminating accreditation. If accreditation is revoked, the laboratory must return its Certificate of Accreditation and cease use of the NVLAP logo on any of its reports, correspondence, or advertising.

If denial or revocation has been proposed, the laboratory may request, in writing, a hearing, under United States Code 5 U.S.C. 556, within 30 days of the date of receipt of the notification. If a hearing is not requested, the action becomes final upon the expiration of that 30-day period.

When accreditation has been terminated, whether voluntarily or through adverse action, the accreditation certificate must be returned to NVLAP.

VI. TECHNICAL REQUIREMENTS

Section 7.33 of the NVLAP Procedures, found in Appendix A, contain the Criteria for accreditation expressed in general terms. The following interpretive comments and additional requirements make the criteria specifically applicable to the Computer Network Interface Protocols Accreditation Program for the DoD High Level Protocols. The requirements listed in Section 7.33 and those specified in this section must be met in order to gain accreditation.

Scope of the Program

The NVLAP Program for the DoD High Level Protocols offers accreditation for testing to Mil-Std 1777 Internet Protocol, Mil-Std 1778 Transmission Control Protocol, Mil-Std 1780 File Transfer Protocol, Mil-Std 1781 Simple Mail Transfer Protocol, and Mil-Std 1782 TELNET Protocol using the DCA Upper Level Protocol Test System. For information about DoD specifications, contact DCA, Code B672, Washington, DC 20305-2000, telephone (703) 285-5337.

QUALITY SYSTEM (see Procedures Sec. 7.33a)

The Quality System requirements are designed to promote laboratory practices which ensure technical integrity of the analyses and adherence to quality assurance practices. The laboratory must maintain a Quality Manual which documents the laboratory's practices and the specific steps taken to ensure quality testing. The Quality Manual must contain or refer to documentation which describes and details the laboratory's implementation of procedures covering all of the technical requirements in this section. This information will be reviewed by NVLAP assessors during on-site assessments.

The Quality System must provide for routine checks of the competence of technicians and others involved in the conduct and evaluation of tests. The Quality Manual must contain a detailed test plan for the conduct of DoD High Level Protocol conformance testing and describe how the laboratory assures the accuracy and consistency of its results. A description of the testing required by DoD is given in the first document listed below.

Records must be kept of all quality assurance activities. Test data from quality assurance checks performed in the laboratory (or with other laboratories) must be summarized and retained for use by the laboratory in monitoring its performance.

The most recent release of the following documents should be available in the laboratory as reference in developing and maintaining the Quality System (see Appendix H for sources of documents).

Conformance Testing Profile for Department of Defense Military Standard Data Communications High Level Protocol Implementations, DRAFT 1989

DCA Upper Level Protocol Test System Functional Description, May 1988.

DCA Upper Level Protocol Test System Installation and Operations Manual, May 1988.

DCA Upper Level Protocol Test System Test Operators Manual, May 1988.

DCA Upper Level Protocol Test System Internet Protocol Mil-Std 1777 Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System Transmission Control Protocol Mil-Std 1778 Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System File Transfer Protocol Mil-Std 1780 Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System Simple Mail Transfer Protocol Mil-Std 1781 Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System TELNET Protocol Mil-Std 1782 Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System Transmission Control Protocol/ Internet Protocol (Tightly Coupled) Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System Internet Protocol Security Option Test Traceability Index, May 1988.

DCA Upper Level Protocol Test System Internet Protocol Mil-Std 1777 Remote Driver Specification, May 1988.

DCA Upper Level Protocol Test System Transmission Control Protocol Mil-Std 1778 Remote Driver Specification, May 1988.

DCA Upper Level Protocol Test System File Transfer Protocol Mil-Std 1780 Remote Driver Specification, May 1988.

DCA Upper Level Protocol Test System Simple Mail Transfer Protocol Mil-Std 1781 Remote Driver Specification, May 1988.

DCA Upper Level Protocol Test System TELNET Protocol Mil-Std 1782 Remote Driver Specification, May 1988.

Military Standard 1777 Internet Protocol, August 1983.

Military Standard 1778 Transmission Control Protocol, August 1983.

Military Standard 1780 File Transfer Protocol, May 1984.

Military Standard 1781 Simple Mail Transfer Protocol, May 1984.

Military Standard 1782 TELNET Protocol, May 1984.

RFC 791 Internet Protocol, September 1981.

RFC 792 Internet Control Message Protocol, September 1981.

RFC 793 Transmission Control Protocol, September 1981.

RFC 854 TELNET Protocol Specification, May 1983.

RFC 959 File Transfer Protocol, October 1985.

RFC 821 Simple Mail Transfer Protocol, August 1982.

STAFF (See Procedures Sec. 7.33b)

The laboratory shall maintain a complete listing of position descriptions and the staff members assigned to those positions. For each staff member, the laboratory must maintain a personnel folder which includes: a resume of qualifications, training, laboratory procedures to which assigned, and the results of periodic testing performance reviews. Performance reviews of staff members may include intra-operator tests, inter-operator tests and between-laboratory tests. The laboratory shall have a description of its training program for ensuring that staff are able to perform tests properly.

Training:

The laboratory must assure that test system operators, analysts, and administrators have adequate qualifications and training to conduct DoD High Level Protocol testing. At a minimum, the laboratory shall be staffed with personnel experienced/knowledgeable in the following areas:

Unix Operating System Installation, Maintenance, and Administration
DCA Upper Level Protocol Test System Operation
DoD Upper Level Protocols (IP,TCP,FTP,SMTP,TELNET)
Data Communications
Packet Switching Networks
Ability to successfully conduct tests, obtain and analyze results, and develop final test reports.

A laboratory must ensure that each new staff member is trained for the testing duties assigned and that staff members are retrained when they are assigned new responsibilities or when test methods are updated. Each staff member must receive (or have had) training for assigned testing duties either through on-the-job training or formal classroom sessions.

Competency:

In addition to training, the laboratory must evaluate the competency of each staff member either through an observation of performance, or an oral or written examination for each test method the staff member is authorized to conduct. The performance must be observed annually by the immediate supervisor or a designee appointed by the laboratory director, and must be adequately documented. A record of the annual evaluation of each staff member must be dated and signed by the supervisor and the employee, and placed in the personnel file.

FACILITIES AND EQUIPMENT (see Procedures Sec. 7.33c)

The laboratory shall have at least the following hardware configuration:

- VAX Series CPU, 4MB Main Memory, .8-MIPS/Whetstones rating
- Magnetic Tape I/O, standard 1/2 inch tape, 1600 bpi
- 200 MB disk storage space
- DDN approved X.25 and/or HDH and/or ethernet network interfaces (See Note 1)
- 2 CRT terminals
- Printer
- 2 9600 bps V.32 Trellis Modulation Modems

Note 1 - An X.25 or HDH interface is required for proficiency testing. Any hardware configuration that contains only an ethernet interface is not acceptable. Also note that bus connected X.25 interfaces do not work under the required operating system (Ultrix 1.1).

The laboratory shall have the following software installed on their test system:

- DEC Ultrix 1.1 Operating System (see Note 2)
- Latest release of the DCA Upper Level Protocol Test System software available from the National Technical Information Service (see Appendix H for ordering information)

Note 2 - DEC no longer distributes Ultrix 1.1. DCA has permission from DEC to distribute Ultrix 1.1 to laboratories which have an Ultrix license. DCA must receive written confirmation from DEC that the laboratory has a license before the software is released. Contact DCA, Code R640, 1860 Wiehle Avenue, Reston, VA 22090, telephone (703) 437-2165.

The laboratory test system shall be connected to the DoD internet either directly or through a gateway. Internet Protocol testing requires that the vendors implementation under test reside on the same network as the lab's test system, therefore a dial-up modem link to the laboratory test system's network is also required. Figures 1, 2 and 3 in Appendix J show several possible network configurations.

NVLAP policy on facilities is given in Appendix E.

CALIBRATION (see Procedures Sec. 7.33d)

Calibration requirements do not apply to the DoD High Level Protocols Program.

TEST METHODS AND PROCEDURES (see Procedures Sec. 7.33e)

Any errors or problems experienced with the test system or procedures should be reported in writing to DCA, Code R640, 1860 Wiehle Avenue, Reston, VA 22090, telephone (703) 437-2165.

The laboratory should not attempt to interpret DCA policy for vendors. If vendors have questions on policy or interpretation, the laboratory should direct the vendor to submit comments on issues in question to DCA.

DCA requires that implementations which have been tested and approved be retested whenever the product is modified. In these cases, allowances may be made by DCA for any test criteria that have been changed since the original testing. Requests should be addresses to DCA.

When an implementation has successfully passed the test, copies of the official report will be furnished to the vendor and, if requested by the vendor, to DCA for consideration for inclusion on the Department of Defense High Level Protocols Qualified Products list.

RECORDS (see Procedures Sec. 7.33f)

The laboratory shall maintain a functional record keeping system. Records must be easily accessible and contain complete information on the subject. Magnetic media must be logged and properly marked. Records covering the following are required:

Staff training dates and results
Staff competency review dates and results
Test equipment name and description
Manufacturer, model, and serial number
Test system hardware and maintenance logs
Test system software

Software and documentation updates
Comprehensive logs of test activities
Problems with test system and documentation
Test data and official reports

Vendor test results and official reports shall be kept by the laboratory for a period of three years following the completion of testing. This includes hard copies of the official results and the test results files. This requirement supercedes the one-year criterion stated in the NVLAP Procedures Section 7.33(f)(2).

TEST REPORTS (see Procedures Sec. 7.33g)

There are two basic types of test reports; one intended for use by the vendor only, and one that is submitted to DCA for acceptance of the protocol implementation on the approved products list.

Reports intended for use only by the vendor shall meet vendor/laboratory contract obligations and be complete, but need not necessarily meet all DCA requirements. The test report must contain sufficient information for the exact test conditions to be reproduced at a later time if a retest is necessary.

Tests reports to be submitted to DCA shall contain the official results from all tests conducted by the laboratory. DCA requires that implementations meet the guidelines for protocol conformance in the DoD High Level Protocol Conformance Profile (to be published). All mandatory requirements and implemented options must pass (see Note) the tests associated with them to be certified conformant by DCA.

The report shall be on letterhead with the NVLAP logo, shall be approved and signed by the laboratory manager, and contain, at a minimum the following information:

Name and address of the laboratory;
Pertinent dates and identifying numbers;
Name of vendor;
Description and name, product no. of implementation;
Vendor information sheet;
Test results and comments related to passing criteria;
Test audit trails and result logs;
Signature of NVLAP approval signatory and
All other items required by the test methods.

Note - It is possible to pass a test without meeting the functional requirements intended to be tested. DCA requires that implementations possess the functionality tested. "Not implemented" responses by the IUT will be considered a failed test.

VII. PROFICIENCY TESTING

Proficiency testing is an integral requirement of the NVLAP process. Applicant laboratories will be required to participate satisfactorily in proficiency testing prior to initial accreditation and annually thereafter. Laboratories renewing accreditation must have satisfactorily participated in all required proficiency testing during their previous accreditation period (see Appendix C).

To evaluate the effective and proper operation of a laboratory, proficiency testing may consist of several parts. The proficiency testing concept is designed to allow the evaluation of the laboratory's ability to produce repeatable and reproducible test data. Portions of the testing process may be "highlighted" in proficiency testing, eg., software, hardware, data analysis, etc. Proficiency testing may include:

1) Testing a known reference protocol implementation to demonstrate the laboratory's ability to properly perform test procedures and operations.

2) Analyzing and developing final test reports from audit trails and results logs, provided by NVLAP, which contain characteristics that are unknown to the laboratory. These audit trails and result logs will be specially generated for proficiency testing purposes.

3) Laboratory access via telephone lines to a remote protocol implementation.

The results of the proficiency testing program will be reported to the participants in appropriate documents and reports. The identities and performance of individual laboratories will remain confidential.

The results of proficiency testing will be made available to on-site assessors for use during laboratory assessment visits. If problems are indicated by proficiency testing, they will be discussed with appropriate laboratory personnel, who will then be responsible for developing and implementing plans for resolving the problems.

Deficiencies identified by proficiency testing, whether during an on-site or not, must be resolved in a manner similar to the process for on-site deficiency resolution.

VIII. ON-SITE ASSESSMENT

A laboratory must undergo a successful on-site assessment and resolve any deficiencies (departures from the NVLAP criteria) noted during the assessment before accreditation can be initially granted for the high level protocols. Deficiencies noted during subsequent on-sites must be resolved in order to maintain accreditation.

The laboratory should be in good working order and prepared to demonstrate testing using the DCA Upper Level Protocol Test System. All observations made by the NVLAP assessor are held in strictest confidence.

The assessor will use NVLAP checklists containing specific questions about all aspects of the visit. The checklists, based on the NVLAP criteria for accreditation and the Critical Elements for High Level Protocols (see Appendix G), serve to ensure a complete assessment and that all assessors cover the same items at each laboratory. The assessor will need to take breaks during the day to fill in the NVLAP checklists and to prepare the Assessment Report.

The laboratory will be responsible for demonstrating its competence to conduct the tests, analyze the test data, and prepare a test report.

The agenda for a typical one-day on-site visit is given below.

1. Assessor conducts an entry briefing with laboratory manager to explain the purpose of the on-site and to discuss the schedule for the day. At the discretion of the laboratory manager, other staff may attend the briefing.

2. Assessor reviews equipment and maintenance records, software versions, record keeping procedures, quality system manuals, laboratory test reports, and personnel competency records. Although there must be a laboratory staff member available to answer questions, the assessor may wish to review the documents alone. The assessor does not usually ask to take any laboratory documents with him.

3. Assessor physically examines equipment and facilities. Assessor observes the demonstration of selected procedures and interviews the personnel. The demonstrations must include test system set-up and the use of all major equipment. The assessor may request a specific demonstration for use as a Proficiency Test.

4. An exit briefing is held with the laboratory manager and staff to discuss the assessor's findings. Deficiencies are discussed and resolutions are mapped out. Items that must be addressed before accreditation can be granted are emphasized. Items that have been corrected during the on-site and any recommendations are specially noted.

5. As a part of the exit briefing, the assessor completes an Assessment Report detailing his findings. The report is signed by the assessor and the laboratory representative and a copy is given to the representative.

APPENDICES

A - NVLAP Procedures, Subpart D,
"Conditions and Criteria for Accreditation "

B - (Reserved)

C - NVLAP Lab Bulletin Nó. 19
"Satisfactory Proficiency Testing Is a Requirement for Accreditation"

D - NVLAP Policy Bulletin #4
"Accreditation of Foreign Laboratories"

E - NVLAP Policy Guide 10
Main laboratory facilities and subfacilities

F - Federal Register Notice
Formal establishment of the Protocols Program
July 21, 1988

G - Critical Elements

H - Sources of Documents

I - Certificate of Accreditation
Scope of Accreditation

J - Acceptable Network Configurations

APPENDIX A

NVLAP PROCEDURES - TITLE 15, PART 7, CODE OF FEDERAL REGULATIONS

SUBPART D - CONDITIONS AND CRITERIA FOR ACCREDITATION

Sec. 7.31 Application of accreditation conditions and criteria.

(a) To become accredited and maintain accreditation, a laboratory must meet the conditions for accreditation set out in Section 7.32 and the criteria set out in Section 7.33 as tailored for specific LAPs.

(b) The conditions leading to accreditation include acceptance of the responsibilities of an accredited laboratory and requirements for information disclosure.

(c) The criteria are tailored and interpreted for the test methods, types of test methods, products, services or standards of the relevant LAP. These tailored criteria are the technical requirements for accreditation developed through the procedures of Section 7.15.

(d) In applying the conditions, criteria, and technical requirements for accreditation, the Director of OPSP shall not:

(1) Prohibit accreditation solely on the basis of a laboratory's affiliation or nonaffiliation with manufacturing, distributing, or vending organizations, or because the laboratory is a foreign firm; or

(2) Develop, modify, or promulgate test methods, standards, or comparable administrative rules.

Sec. 7.32 Conditions for accreditation.

(a) To become accredited and maintain accreditation, a laboratory shall agree in writing to:

(1) Be assessed and evaluated initially and on a periodic basis;

(2) Demonstrate, on request, that it is able to perform the tests representative of those for which it is seeking accreditation;

(3) Pay all relevant fees;

(4) Participate in proficiency testing as required.

(5) Be capable of performing the tests for which it is accredited according to the latest version of the test method within one year after its publication or within another time limit specified by the Director of OPSP;

(6) Limit the representation of the scope of its accreditation to only those tests or services for which accreditation is granted;

(7) Limit all its test work or services for clients to those areas where competence and capacity are available;

(8) Limit advertising of its accredited status to letterheads, brochures, test reports, and professional, technical, trade, or other laboratory services publications, and use the NVLAP logo under guidance provided by the Director of OPSP;

(9) Inform its clients that the laboratory's accreditation or any of its test reports in no way constitutes or implies product certification, approval, or endorsement by NBS;

(10) Maintain records of all actions taken in response to testing complaints for a minimum of one year;

(11) Maintain an independent decisional relationship between itself and its clients, affiliates, or other organizations so that the laboratory's capacity to render test reports objectively and without bias is not adversely affected;

(12) Report to the Director of OPSP within 30 days any major changes involving the location, ownership, management structure, authorized representative, approved signatories, or facilities of the laboratory; and

(13) Return to the Director of OPSP the certificate of accreditation for possible revision or other action should it:

(i) be requested to do so by the Director of OPSP;

(ii) voluntarily terminate its accredited status; or

(iii) become unable to conform to any of these conditions or the applicable criteria of Section 7.33 and related technical requirements.

(b) To become accredited and maintain accreditation, a laboratory shall supply, upon request, the following information:

(1) Legal name and full address;

(2) Ownership of the laboratory;

(3) Organization chart defining relationships that are relevant to performing testing covered in the accreditation request;

(4) General description of the laboratory, including its facilities and scope of operation;

(5) Name and telephone number of the authorized representative of the laboratory;

(6) Names or titles and qualifications of laboratory staff nominated to serve as approved signatories of test reports that reference NVLAP accreditation; and

(7) Other information as may be needed for the specific LAP(s) in which accreditation is sought.

Sec. 7.33 Criteria for accreditation.

(a) Quality System.

(1) The laboratory shall operate under an internal quality assurance program appropriate to the type, range, and volume of work performed. The quality assurance program must be designed to ensure the required degree of accuracy and precision of the laboratory's work and should include key elements of document control, sample control, data validation, and corrective action. The quality assurance program must be documented in a quality manual or equivalent (e.g., operations notebook) which is available for use by laboratory staff. A person(s) must be identified as having responsibility for maintaining the quality manual.

(2) The quality manual must include as appropriate:
  (i) The laboratory's quality assurance policies including procedures for corrective action for detected test discrepancies;
  (ii) Quality assurance responsibilities for each function of the laboratory;
  (iii) Specific quality assurance practices and procedures for each test, type of test, or other specifically delineated function performed;
  (iv) Specific procedures for testing, control charts, reference materials, and interlaboratory tests; and
  (v) Procedures for dealing with testing complaints.
(3) The laboratory shall periodically review its quality assurance system by or on behalf of management to ensure it's continued effectiveness. These reviews must be recorded with details of any corrective action taken.

(b) Staff.
(1) The laboratory shall:
  (i) Be staffed by individuals having the necessary education, training, technical knowledge, and experience for their assigned functions; and
  (ii) Have a job description for each professional, scientific, supervisory and technical position, including the necessary education, training, technical knowledge, and experience.
(2) The laboratory shall document the test methods each staff member has been assigned to perform.
(3) The laboratory shall have a description of its training program for ensuring that new or untrained staff are able to perform tests properly and uniformly to the requisite degree of precision and accuracy.
(4) The laboratory shall be organized:
  (i) So that staff members are not subjected to undue pressure or inducement that might influence their judgment or results of their work; and
  (ii) In such a way that staff members are aware of both the extent and the limitation of their area of responsibility.
(5) The laboratory shall have a technical manager (or similar title) who has overall responsibility for the technical operations of the laboratory.
(6) The laboratory shall have one or more signatories approved by the Director of OPSP to sign test reports that reference NVLAP accreditation. Approved signatories shall:
  (i) Be competent to make a critical evaluation of test results; and
  (ii) Occupy positions within the laboratory's organization which makes them responsible for the adequacy of test results.

(c) Facilities and Equipment.
(1) The laboratory shall be furnished with all items of equipment and facilities for the correct performance of the tests and measurements for which accreditation is granted and shall have adequate space, lighting, and environmental control, and monitoring to ensure compliance with prescribed testing conditions.
(2) All equipment must be properly maintained to ensure protection from corrosion and other causes of deterioration. Instructions for a proper maintenance procedure for those items of equipment which require periodic maintenance must be available. Any item of equipment or component thereof

which has been subjected to overloading or mishandling, gives suspect results, or has been shown by calibration or otherwise to be defective, must be taken out of service and clearly labelled until it has been repaired. When placed back in service, this equipment must be shown by test or calibration to be performing its function satisfactorily.
(3) Records of each major item of equipment must be maintained. Each record must include:
  (i) The name of the item of equipment;
  (ii) The manufacturer's name and type, identification and serial number;
  (iii) Date received and date placed in service;
  (iv) Current location, where appropriate;
  (v) Details of maintenance; and
  (vi) Date of last calibration, next calibration due date, and calibration report references.

(d) Calibration. The laboratory shall:
(1) Calibrate new testing equipment before putting it into service;
(2) Recalibrate, at regular intervals, in-service testing equipment with the calibration status readily available to the operator;
(3) Perform checks of in-service testing equipment between the regular calibration intervals, where relevant;
(4) Maintain adequate records of all calibrations and recalibrations; and
(5) Provide traceability of all calibrations and reference standards of measurement where these standards exist. Where traceability of measurements to primary (national or international) standards is not applicable, the laboratory shall provide satisfactory evidence of the accuracy or reliability of test results (e.g., by participation in a suitable program of interlaboratory comparison).

(e) Test Methods and Procedures. The laboratory shall:
(1) Conform in all respects with the test methods and procedures required by the specifications against which the test item is to be tested, except that whenever a departure becomes necessary for technical reasons the departure must be acceptable to the client and recorded in the test report;
(2) Have data to prove that any departures from standard methods and/or procedures due to apparatus design or for other reasons do not detract from the expected or required precision of the measurement;
(3) Maintain a test plan for implementing testing standards and procedures including adequate instructions on the use and operation of all relevant equipment, on the handling and preparation of test items (where applicable), and on standard testing techniques where the absence of such instructions could compromise the test. All instructions, testing standards, specifications, manuals, and reference data relevant to the work of the laboratory must be kept up-to-date and made readily available to the staff;
(4) Maintain measures for the detection and resolution of in-process testing discrepancies for manual and automatic test equipment and electronic data processing equipment, where applicable;

(5) Maintain a system for identifying samples or items to be tested, which remains in force from the date of receipt of the item to the date of its disposal, either through documents or through marking to ensure that there is no confusion regarding the identity of the samples or test items and the results of the measurements made; and

(6) Maintain rules for the receipt, retention, and disposal of test items, including procedures for storage and handling precautions to prevent damage to test items which could invalidate the test results. Any relevant instructions provided with the tested item must be observed.

(f) Records. The laboratory shall:

(1) Maintain a record system which contains sufficient information to permit verification of any issued report;

(2) Retain all original observations, calculations and derived data, and calibration records for one year unless a longer period is specified; and

(3) Hold records secure and in confidence, as required.

(g) Test Reports.

(1) The laboratory shall issue test reports of its work which accurately, clearly, and unambiguously present the specified test results and all required information. Each test report must include the following information as applicable:

(i) Name and address of the laboratory;

(ii) Identification of the test report by serial number, date, or other appropriate means;

(iii) Name and address of client;

(iv) Description and identification of the test specimen, sample, or lot of material represented;

(v) Identification of the test specification, method, or procedure used;

(vi) Description of sampling procedure, if appropriate;

(vii) Any deviations, additions to, or exclusions from the test specifications;

(viii) Measurements, examinations, and derived results supported by tables, graphs, sketches, and photographs, as appropriate, and any failures identified;

(ix) A statement of measurement uncertainty, where relevant;

(x) Identification of the organization and the person accepting technical responsibility for the test report and date of issue;

(xi) A statement that the report must not be reproduced except in full with the approval of the laboratory; and

(xii) A statement to the effect that the test report relates only to the items tested.

(2) The laboratory shall issue corrections or additions to a test report only by a further document suitably marked, e.g. "Supplement to test report serial number ....." which meets the relevant requirements of Section 7.33(g)(1).

(3) The laboratory shall retain a copy of each test report issued for one year unless a longer period is specified by the Director of OPSP.

(4) The laboratory shall ensure that all test reports endorsed with the NVLAP logo are signed by an approved signatory.

National Voluntary Laboratory Accreditation Program

NVLAP

U.S. Department of Commerce in cooperation with the National Bureau of Standards

Bulletin

Policy Bulletin No. 19
August 1986

SATISFACTORY PROFICIENCY TESTING IS A REQUIREMENT FOR ACCREDITATION

Accreditation by the National Bureau of Standards, under the National Voluntary Laboratory Accreditation Program (NVLAP), requires that a laboratory meet all performance requirements and criteria as determined by on-site assessments and proficiency testing.

If, as the result of on-site assessments, deficiencies are found, the laboratory must satisfactorily resolve those deficiencies, in order to obtain initial accreditation or maintain accreditation.

Unsatisfactory participation in any NVLAP proficiency testing program is a technical deficiency which must be resolved in order to obtain initial accreditation or maintain accreditation.

Unsatisfactory participation in NVLAP proficiency testing programs is defined as, but not limited to, one or more of the following:

1. Failure to meet specified proficiency testing performance requirements prescribed by a standard or test method for which the laboratory is seeking accreditation. (Example: ANSI Standard N13.11 for the Dosimetry LAP.)

2. Failure to participate in a regularly scheduled "round" of proficiency testing for which the laboratory has received instructions and/or materials.

3. Failure to submit laboratory control data as required. (Example: Within laboratory control data to be submitted twice annually for the Concrete LAP.)

4. Performance as a statistically outlying laboratory in two successive rounds of proficiency testing or showing a general pattern of outlying test results over three or more rounds.

5. Failure to produce test data within acceptable limits of error when testing NBS Standard Reference Materials or special artifacts whose properties are well characterized and known to NBS/NVLAP.

NVLAP will notify the laboratory of proficiency testing deficiency(s) and actions to be taken to resolve the deficiency(s). Denial or suspension of accreditation will result from failure to resolve deficiencies.

APPENDIX D

NVLAP Policy Bulletin #4
(revised April 1987)

ACCREDITATION OF FOREIGN LABORATORIES

Foreign laboratories, located outside of the continental United States, may be accredited by NVLAP on the same basis as U.S. domestic laboratories. Foreign laboratories must meet the same requirements and criteria as domestic laboratories. The criteria are defined in the NVLAP Procedures and technical Handbooks provided to all applicants. Accreditation is granted based on compliance with all NVLAP criteria as determined by on-site assessments and the results of proficiency testing programs.

Since NVLAP is a cost-reimbursable program, the fees charged foreign laboratories must cover all costs in excess of those associated with the accreditation of domestic laboratories. Additional fees will be charged to foreign laboratories for travel by assess outside of the United States and for shipment of proficiency testing materials to the laboratories.

Upon application, a foreign laboratory must forward payment of NVLAP fees (as calculated on the Fee Calculation Sheet) in U.S. currency. The laboratory will be notified of additional travel and proficiency testing costs which must be paid to NVLAP before an assessor leaves to perform the on-site assessment.

In cases where laboratory documents are not in English, or laboratory personnel do not speak English, it is the responsibility of the laboratory to provide a translator to assist the NVLAP assessor during the inspection. The translator will assist the assessor to converse directly with laboratory management and technical staff and to review laboratory documentation. Documents such as quality control manuals, protocols, standards, and test reports need not be translated into English solely for NVLAP purposes.

SPECIAL REQUIREMENT FOR THE PROTOCOLS PROGRAM

An export license, issued by the U.S. Department of Commerce, may be required for certain equipment to be sold outside the United States. A foreign laboratory applying for NVLAP accreditation must own the required equipment and must have a valid export license for it.

For export license information call (202) 377-4811 or write to:

U.S. Department of Commerce
Export Administration
Exporter Assistance
P.O. Box 273
Washington, DC 20230

NVLAP POLICY GUIDE 10

This Policy Bulletin presents NVLAP definitions of the types of laboratory facilities which may be granted NVLAP accreditation, the requirements and conditions that must be satisfied in order to achieve accreditation, and procedures that NVLAP will follow in evaluating various types of facilities for their conformance to accreditation criteria.

Definitions:

a. Main (laboratory) facility:

(1) permanently (at all times) maintains staff, equipment, procedures, documentation, and facilities necessary to perform the tests, for which it seeks accreditation;
(2) implements all quality assurance procedures;
(3) maintains and retains all records, and issues test reports; and
(4) may be a permanently fixed site or a permanent mobile facility.

b. Sub-facility is physically separate from, but considered an extension of, its main facility. Although it may have all staff, equipment, procedures, and documentation necessary to perform the requisite tests, it receives technical direction and quality assurance management from the main facility.

1. A permanent sub-facility maintains staff, equipment, procedures, documentation, and facilities necessary to perform the tests, for which it seeks accreditation, at all times. It may be a permanently fixed site or a permanent mobile facility and is expected to remain in operation for at least one year.

2. A temporary sub-facility is provided with staff, equipment, procedures, documentation, and facilities necessary to perform the tests, for which it seeks accreditation, on an interim basis, to meet the needs of the main facility. A temporary sub-facility may be established at a fixed site or in a mobile facility and is expected to remain in operation less than one year.

Conditions for Accreditation:

NVLAP accreditation of a laboratory main facility does not extend to accreditation of sub-facilities unless the sub-facilities have been separately evaluated. These facilities are uniquely identified in the NVLAP accreditation documents. A NVLAP-accredited laboratory must not present or report test data, produced at any non-accredited, sub-facility as having been produced under the status of NVLAP accreditation.

NVLAP offers accreditation to laboratories that are found competent to perform specific test methods or types of tests in specified fields of testing. Competence is defined as the ability to meet specific technical criteria

relating to quality assurance, staff, equipment, facilities, procedures, records, and reports. Technical criteria may or may not be equally applicable to main facilities and sub-facilities. Accreditation of sub-facilities may require NVLAP criteria that address the use and maintenance of equipment and facilities, and the implementation of procedures, that are particularly applicable to the performance of specific test methods in sub-facilities. NVLAP must develop specific technical criteria upon which to base an objective evaluation of staff, facilities, equipment, and procedures employed in applicable sub-facilities.

NVLAP will accredit a main facility if the facility complies with all applicable NVLAP criteria.

NVLAP will accredit a sub-facility (in addition to the main facility) if:

a. the laboratory main facility meets all NVLAP accreditation criteria;

b. the laboratory main facility satisfactorily documents and maintains quality assurance procedures addressing the applicable sub-facility; and,

c. the sub-facility complies with all applicable NVLAP criteria.

Procedures:

In principle, NVLAP will require that sub-facilities, to be included in a laboratory's accreditation, undergo on-site assessments and participate in proficiency testing. NVLAP staff, with the guidance of NVLAP technical experts, will determine the need for and extent of such evaluations based on the number and location of similar sub-facilities managed by the laboratory, the nature of the quality assurance system, and any special technical considerations. Decisions on the need for and extent of the evaluations may not be made until after the accreditation of the main facility. The conditions and requirements for evaluation of sub-facilities providing specific testing services are described in NVLAP documents pertaining to the relevant accreditation program.

Laboratories seeking NVLAP accreditation should clearly state, on the NVLAP Application Form, what type(s) of sub-facilities are to be included in the accreditation. NVLAP fees for on-site assessments and proficiency testing will be based on the number of facilities seeking accreditation that are required to undergo on-sites and participate in proficiency testing. A single administrative/technical support fee is charged to the laboratory (main facility).

APPENDIX F

National Bureau of Standards

[Docket No. 80741-8141]

National Voluntary Laboratory Accreditation Program

AGENCY: National Bureau of Standards, Commerce.

ACTION: Notice of formal establishment of a laboratory accreditation program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense.

SUMMARY: Under the National Voluntary Laboratory Accreditation Program (NVLAP), the National Bureau of Standards (NBS) announces the establishment of a laboratory accreditation program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense (Protocols Program). Laboratories that are interested in becoming accredited under the Protocols Program may indicate their interest in the program by informing the Manager, Laboratory Accreditation, National Bureau of Standards of their specific interests.

FOR FURTHER INFORMATION CONTACT: John L. Donaldson, Manager, Laboratory Accreditation, National Bureau of Standards, Admin A527, Gaithersburg, MD 20899, (301) 975-4016.

SUPPLEMENTARY INFORMATION:

Background

This notice is issued in accordance with § 7.17 of the NVLAP Procedures (15 CFR Part 7). Establishment of this program for laboratories that test the computer industry's implementation of communications protocols used by the Department of Defense follows a request by the Defense Communications Agency. A Federal Register notice announcing the request for the Protocols LAP was published on December 3, 1987 (50 FR 45986-45988). Comments received in response to the announcement were reviewed by the Defense Communications Engineering Center whose director, Warren P. Hawrylko, has concluded that there were no valid reasons presented in the comment letters to prevent establishment of the Protocols LAP and therefore requested the National Bureau of Standards to proceed to establish the requested program.

The purpose of the LAP is to accredit and provide national recognition to laboratories capable of performing tests in accordance with the designated test methods. The scope of the LAP includes testing services for: (1) Defense Data Network (DDN) X.25 Link and Network Layer Protocols as specified in the DCA DDN X.25 Host Interface Specification; (2) the five DoD packet switching High Level Protocols (HLPs): (I) Internet Protocol (IP) MIL-STD 1777; (II) Transmission Control Protocol (TCP), MIL-STD 1778; (III) File Transfer Protocol (FTP), MIL-STD 1780; (IV) Simple Mail Transfer Protocol, MIL STD 1781; and (V) TELNET, MIL-STD 1782; and (3) the AUTODIN Mode I Protocol. Accreditation will be offered first for the X.25 protocol. Accreditation will be offered next, at least 45 days later, for the DoD HLPs (I)-(V). Accreditation for AUTODIN Mode I protocol will be offered last, after the initial X.25 protocol accreditations have been completed.

Procedure Prior to Application

Any testing laboratory interested in becoming accredited under this LAP should contact the Manager, Laboratory Accreditation, at the address shown above, specifying the protocols of interest. The laboratory will be sent the proposed technical documents for the requested protocol accreditation as they become available and will be invited to submit comments for their revision within 45 days of the publication date of this notice in the case of the X.25 protocol and of the mailing dates of the documents for the HLP and AUTODIN protocols. A meeting of all interested parties will be scheduled after each of the closing dates to resolve conflicting comments. If none arise, no meeting will be scheduled. The completed technical documents, instructions, fee schedules, and applications will be sent separately for each protocol as they become available to all laboratories that have previously requested them.

Earnest Ambler,
Director.

Dated: July 15, 1988.

[FR Doc. 88-16439 Filed 7-20-88; 8:45 am]
BILLING CODE 3510-13-M

APPENDIX G

National Voluntary Laboratory Accreditation Program (NVLAP)
Computer Network Interface - DoD High Level Protocols

CRITICAL ELEMENTS

NVLAP TEST METHOD DESIGNATION:

TEST METHOD: DoD Military Standard Data Communications High Level Protocols: Mil-Std 1777 Internet Protocol, Mil-Std 1778 Transmission Control Protocol, Mil-Std 1780 File Transfer Protocol, Mil-Std 1781 Simple Mail Transfer Protocol, and Mil-Std 1782 TELNET Protocol

TEST METHOD SCOPE:To ensure that the DoD high level protocol of the Implementation Under Test (IUT) conforms to the corresponding military standard, per the Defense Communications Agency (DCA) Upper Level Protocol Test System.

ENVIRONMENTAL/TEST SAMPLE CONDITIONING: N/A

TEST EQUIPMENT AND APPARATUS:

1. VAX Series CPU, 4MB Main Memory (min.), .8-MIPS/Whetstones rating (min.).
2. Magnetic Tape I/O, standard 1/2 inch tape, 1600 bpi.
3. 200 MB disk storage space.
4. DDN approved X.25 and/or HDH and/or ethernet interface.
5. 2 CRT Terminals.
6. Printer.
7. 2 9600 bps V.32 Trellis Modulation Modems.
8. DEC Ultrix 1.1 Operating System
9. DCA Upper Level Protocol Test System Software.

TESTING PROCEDURES: (Reference: DCA Upper Level Protocol Test System Test Operators Manual)

1. Laboratory has appropriate information from vendor to initialize tests.
2. Network connection is available between IUT and test system.
3. Test System is properly initialized.
4. Appropriate tests are executed.
5. Audit Trails and Result Logs are printed and examined.

TEST REPORTS:

1. Contain all information required by client or by DCA.

SPECIAL CONSIDERATIONS:

1. Proficiency testing is required.
2. Tests are properly initialized.
3. Appropriate tests are executed.
4. Hard copies of test data are obtained.
5. Test data and reports are stored for future reference.

SOURCES OF DOCUMENTS AND SOFTWARE

The following document(s) are available from: DCA, Code R640, Reston, VA 22090, telephone (703) 437-2261.

- Conformance Testing Profile for Department of Defense Military Standard Data Communications High Level Protocol Implementations, DRAFT 1989

The DCA Upper Level Protocol Test System Software and the following documents are available from: National Technical Information Service, 5285 Port Royal Road, Springfield, VA 22161, telephone (703) 487-4650

- DCA Upper Level Protocol Test System Functional Description
- DCA Upper Level Protocol Test System Installation and Operations Manual
- DCA Upper Level Protocol Test System Test Operators Manual
- DCA Upper Level Protocol Test System Internet Protocol Test Traceability Index
- DCA Upper Level Protocol Test System Transmission Control Protocol Test Traceability Index
- DCA Upper Level Protocol Test System File Transfer Protocol Test Traceability Index
- DCA Upper Level Protocol Test System Simple Mail Transfer Protocol Test Traceability Index
- DCA Upper Level Protocol Test System TELNET Protocol Test Traceability Index
- DCA Upper Level Protocol Test System Transmission Control Protocol/ Internet Protocol (Tightly Coupled) Test Traceability Index
- DCA Upper Level Protocol Test System Internet Protocol Security Option Test Traceability Index
- DCA Upper Level Protocol Test System Internet Protocol Remote Driver Specification
- DCA Upper Level Protocol Test System Transmission Control Protocol Remote Driver Specification
- DCA Upper Level Protocol Test System File Transfer Protocol Remote Driver Specification
- DCA Upper Level Protocol Test System Simple Mail Transfer Protocol Remote Driver Specification
- DCA Upper Level Protocol Test System TELNET Protocol Remote Driver Specification

The following documents are available from: Naval Publications and Forms Center, Code 3015, 5801 Tabor Drive, Philadelphia, PA 19120

- Military Standard: 1777, 1778, 1780, 1781, and 1782

The following documents are available from: DDN Network Information Center, SRI International, 333 Ravenswood Avenue, Room EJ291, Menlo Park, CA 94025

- RFC: 791, 792, 793, 854, 959, and 821

National Institute of Standards and Technology

NVLAP

National Voluntary Laboratory Accreditation Program

SCOPE OF ACCREDITATION

COMPUTER NETWORK INTERFACE PROTOCOLS

NVLAP LAB CODE 0000

LABORATORY, INC.
1 Main Street, Anytown, MD 00000
John Doe Phone: 301-555-1212

Accreditation Renewal Date: January 1, 19--

NVLAP Test Method Code	Test Method Designation
17/H01	Department of Defense Military Standard Data Communications High Level Protocols, per Defense Communications Agency Upper Level Protocol Test System: Mil-Std 1777 Internet Protocol Mil-Std 1778 Transmission Control Protocol Mil-Std 1780 File Transfer Protocol Mil-Std 1781 Simple Mail Transfer Protocol Mil-Std 1782 TELNET Protocol
17/X25	Defense Data Network X.25 Host Interface Qualification Tests per Defense Communications Agency Circular 370-P195-(5)

DEPARTMENT OF COMMERCE • UNITED STATES OF AMERICA

For the National Institute of Standards and Technology

APPENDIX I

United States Department of Commerce
National Institute of Standards and Technology

NVLAP

Certificate of Accreditation

LABORATORY, INC.
ANYTOWN, MD

is recognized under the National Voluntary Laboratory Accreditation Program for satisfactory compliance with criteria established in Title 15, Part 7 Code of Federal Regulations. Accreditation is awarded for specific services, listed on the Scope of Accreditation, for:

COMPUTER NETWORK INTERFACE PROTOCOLS

January 1, 19--
Effective until

DEPARTMENT OF COMMERCE • UNITED STATES OF AMERICA

For the National Institute of Standards and Technology

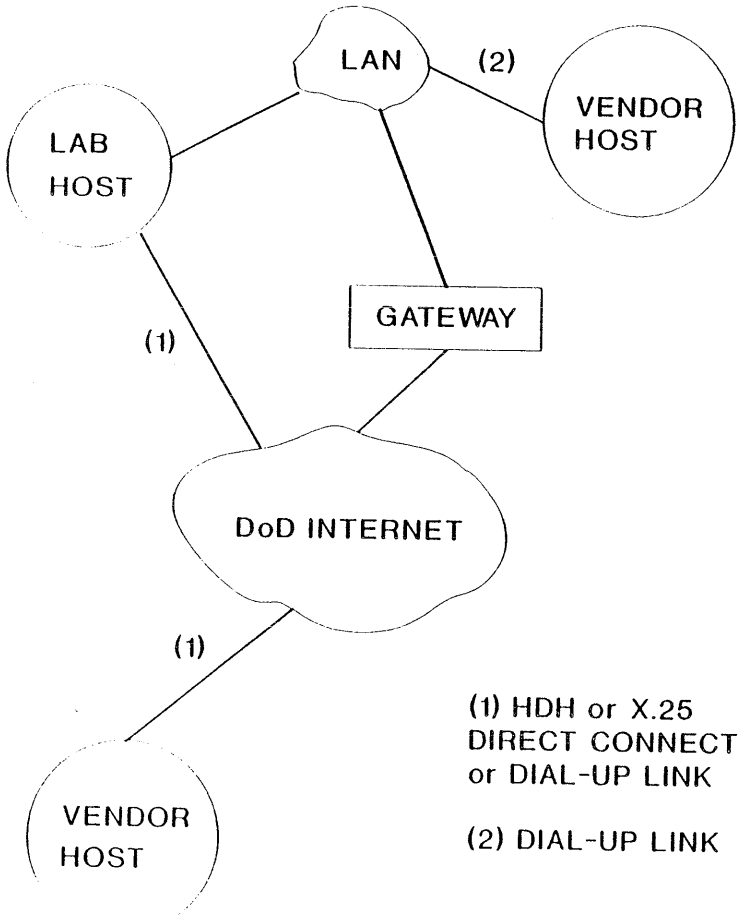

FIGURE 1. ACCEPTABLE NETWORK CONFIGURATION

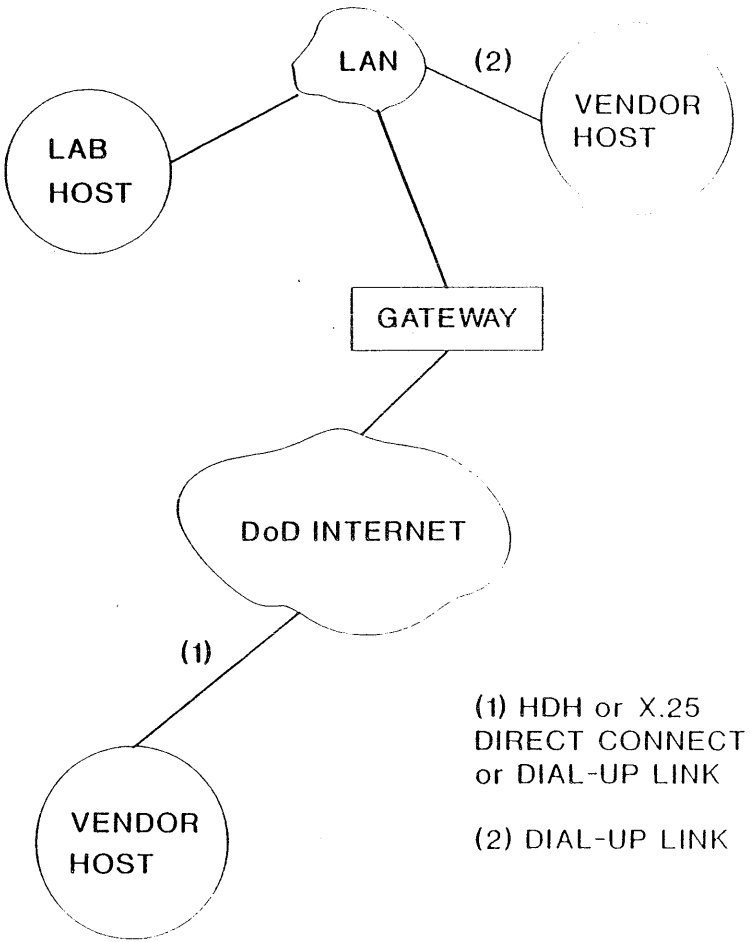

FIGURE 2. ACCEPTABLE NETWORK CONFIGURATIONS

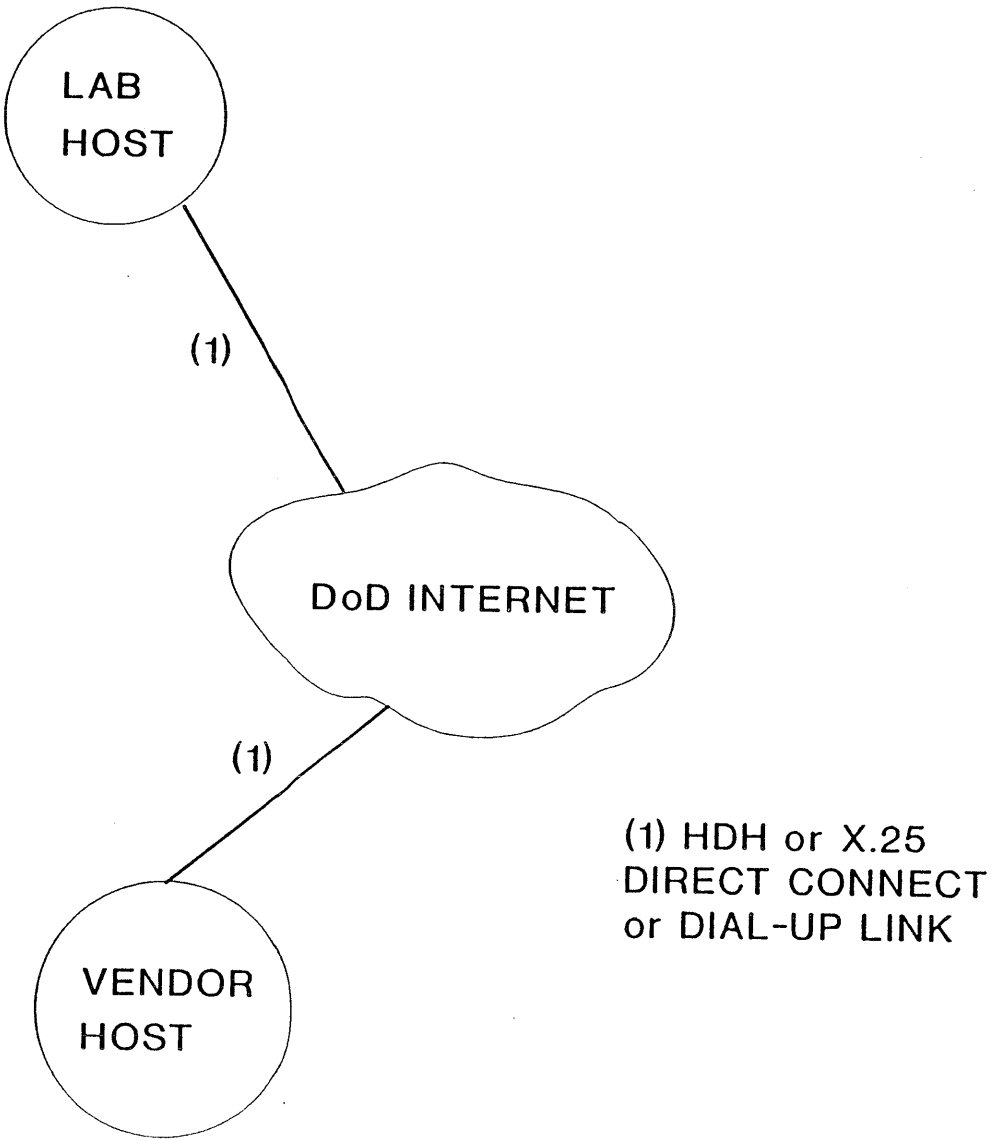

FIGURE 3. ACCEPTABLE NETWORK CONFIGURATION

THE COMPUTER WORM

A Report to the Provost of Cornell University on an Investigation Conducted by The Commission of Preliminary Enquiry:

Ted Eisenberg, Law
David Gries, Computer Science
Juris Hartmanis, Computer Science
Don Holcomb, Physics
M. Stuart Lynn, Office of Information Technologies (Chair)
Thomas Santoro, Office of the University Counsel

February 6, 1989

Cornell University
308 Day Hall
Ithaca, NY 14853-2801
(607) 255-3324

THE COMPUTER WORM

A Report to the Provost on an Investigation Conducted by
The Commission of Preliminary Enquiry:

Ted Eisenberg, Law
David Gries, Computer Science
Juris Hartmanis, Computer Science
Don Holcomb, Physics
M. Stuart Lynn, Office of Information Technologies (Chair)
Thomas Santoro, Office of University Counsel

---

1. INTRODUCTION

This is a report of the Commission of Preliminary Enquiry appointed in response to Provost Barker's letter of November 7, 1988, to Vice President for Information Technologies, M. Stuart Lynn. Provost Barker's letter requested an investigation of the apparent use of Cornell computers to construct and launch the "worm"[1] that disrupted computer networks and systems nationwide beginning November 2, 1988. Provost Barker's letter was prompted by widespread press reports alleging that a Cornell first-year computer science graduate student, Robert Tappan Morris, had created the worm and had unleashed it on the Internet, a collection of national computer networks linking research and instructional facilities in universities as well as government and industrial research establishments.

The worm reportedly disrupted the operations of over 6,000 computers nationwide[2] by exploiting certain security loopholes in applications closely associated with the operating system[3]. Computers affected were limited to those

[1] The press popularly referred to the worm as a "virus", which was the early "diagnosis" of some technical experts before the program had been fully analyzed. However, technically the program was a "worm" since it did not attach itself to a host program in order to propagate itself across the networks.

[2] This estimate by the press may not be accurate. See Section 6, "Impact of the Worm".

[3] The operating system of a computer is a complex piece of software that controls the operations of the computer, providing the environment in which applications software can function. Every computer requires an operating system.

running a version of the UNIX operating system[4] known as 4.3BSD, which was developed by the Computer Systems Research Group (CSRG) of the University of California, Berkeley, and distributed at no charge other than distribution costs to universities and research institutions around the country. It also affected versions of UNIX that were derived from the CSRG work, in particular versions of SUN, which ran on SUN Microsystems computers.

The Commission was charged to:

(1) Accumulate all evidence concerning the potential involvement of Mr. Robert Tappan Morris in the computer worm attack, and to assess such evidence to determine whether or not Morris was the likely perpetrator.

(2) Accumulate all evidence concerning the potential involvement of any other member of the Cornell community, and to assess such evidence to determine whether or not any other member of the Cornell community was involved in the worm attack or was aware of the potential worm attack.

(3) Evaluate relevant computer policies and procedures to determine which, if any, were violated and to make preliminary recommendations as to whether any of such policies and procedures should be modified to inhibit potential future security violations of this general type.

4 UNIX is a registered trademark of AT&T, the original developers of the system.

2. SUMMARY OF FINDINGS AND COMMENTS

Findings:

Based on the evidence presented to the Commission, the Commission finds[5] that:

Responsibility for the Acts:

- The worm attack occurred as described in Section 3.
- Robert Tappan Morris, a first year computer science graduate student at Cornell, created the worm and unleashed it on the Internet.
- In the process of creating and unleashing the worm, Morris violated Computer Science Department policy on the use of departmental research computing facilities.

Impact of the Worm:

- The performance of computers "infected" by the worm degraded substantially, unless remedial steps were taken. Eventually such infected computers would come to a halt. These symptoms were caused by uncontrollable replication of the worm clogging the computer's memory. The worm, however, did not modify or destroy any system or user files or data.
- Based on anecdotal and other information, several thousand computers were infected[6] by the worm. The Commission has not systematically attempted to estimate the exact number infected. Many thousands more were affected in the sense that they had to be tested for infection and preventive measures applied even if the computers were not infected. It appears that the operation of most infected and potentially affected computers and of the research done on those computers was brought to a halt in order to apply remedial or preventive measures, all of which required the diversion of

5 The Commission has chosen not to adopt an express standard of proof for its findings. The findings are only qualified where the Commission cannot reach a definitive conclusion.

6 We use the term "infect" to signify that at least one copy of the worm was left on the penetrated computer.

considerable staff time from more productive efforts.

Mitigation Attempts:

- Morris made only minimal efforts to halt the worm once it had propagated, and did not inform any person in a position of responsibility as to the existence and content of the worm.

Violation of Computer Abuse Policies:

- The Cornell Computer Science Department "Policy for the Use of the Research Computing Facility" prohibits "use of its computer facilities for browsing through private computer files, decrypting encrypted material, or obtaining unauthorized user privileges". All three aspects of this Policy were violated by Morris.

- Morris was apparently given a copy of this Policy but it is not known whether he read it. Probably he did not attend the lecture during orientation when this Policy was discussed, even though he was present on campus.

Intent:

- Most probably Morris did not intend for the worm to destroy data or other files or to interfere with the normal functioning of any computers that were penetrated.

- Most probably Morris intended for the worm to spread widely through host computers attached to the network in such a manner as to remain undiscovered. Morris took steps in designing the worm to hide it from potential discovery, and yet for it to continue to exist in the event it actually was discovered. It is not known whether he intended to announce the existence of the worm at some future date had it propagated according to this plan.

- There is no direct evidence to suggest that Morris intended for the worm to replicate uncontrollably. However, given Morris' evident knowledge of systems and networks, he knew or clearly should have known that such a consequence was certain, given the design of the worm. As such, it appears that Morris failed to consider the most probable consequences of his actions. At the very least, such failure constitutes reckless disregard of those probable consequences.

Security Attitudes and Knowledge:

- This appears to have been an uncharacteristic act for Morris to have committed, according to those who knew him well. In the past, particularly while an undergraduate at Harvard University, Morris appears to have been more concerned about protecting against abuse of computers rather than in violating computer security.

- Harvard's policy on misuse of computer systems contained in the Harvard Student Handbook clearly prohibited actions of the type inherent to the creation and propagation of the worm. For this and other reasons, the Commission believes that Morris knew that the acts he committed were regarded as wrongful acts by the professional community.

- At least one of the security flaws exploited by the worm was previously known by a number of individuals, as was the methodology exploited by other flaws. Morris may have discovered the flaws independently.

- Many members of the UNIX community are ambivalent about reporting security flaws in UNIX out of concern that knowledge of such flaws could be exploited before the flaws are fixed in all affected versions of UNIX. There is no clear security policy among UNIX developers, including in the commercial sector. Morris explored UNIX security issues in such an ambivalent atmosphere and received no clear guidance about reporting security flaws from his peers or mentors at Harvard or elsewhere.

Technical Sophistication:

- Although the worm was technically sophisticated, its creation required dedication and perseverance rather than technical brilliance. The worm could have been created by many students, graduate or undergraduate, at Cornell or at other institutions, particularly if forearmed with knowledge of the security flaws exploited or of similar flaws.

Cornell Involvement:

- There is no evidence that anyone from the Cornell community aided Morris or otherwise knew of the worm prior to its launch. Morris did inform one student earlier that he had discovered certain

security weaknesses in UNIX. The first that anyone at Cornell learned that any member of the Cornell community might have been involved came at approximately 9.30 p.m. on November 4 when the Cornell News Service was contacted by the Washington Post.

Ethical Considerations:

- Prevailing ethical beliefs of students towards acts of this kind vary considerably from admiration to tolerance to condemnation. The computer science profession as a whole seems far less tolerant, but the attitudes of the profession may not be well communicated to students.

Community Sentiment:

- Sentiment among the computer science professional community appears to favor strong disciplinary measures for perpetrators of acts of this kind. Such disciplinary measures, however, should not be so stern as to damage permanently the perpetrator's career.

University Policies on Computer Abuse:

- The policies and practices of the Cornell Computer Science Department regarding computer abuse and security are comparable with those of other computer science and many other academic departments around the nation.

- Cornell has policies on computer abuse and security that apply to its central facilities, but not to departmental facilities.

- In view of the pervasive use of computers throughout the campus, there is a need for university-wide policy on computer abuse. The Commission recommends that the Provost establish a committee to develop such policy, and that such policy appear in all legislative and policy manuals that govern conduct by members of the Cornell community.

- In view of the distributed nature of computing at Cornell, there is also a need for a university-wide committee to provide advice and appropriate standards on security matters to departmental computer and network facility managers. The Commission recommends that the Vice President for Information Technologies be asked to establish such a committee.

Comments:

The Commission believes that the acts committed in obtaining unauthorized passwords and in disseminating the worm on the national network were wrong and contrary to the standards of the computer science profession. They have little if any redeeming technical, social or other value. The act of propagating the worm was fundamentally a juvenile act that ignored the clear potential consequences. The act was selfish and inconsiderate of the obvious effect it would have on countless individuals who had to devote substantial time to cleaning up the effects of the worm, as well as on those whose research and other work was interrupted or delayed.

Contrary to the impression given in many media reports, the Commission does not regard this act as an heroic event that pointed up the weaknesses of operating systems. The fact that UNIX, in particular BSD UNIX, has many security flaws has been generally well-known, as indeed are the potential dangers of viruses and worms in general. Although such security flaws may not be known to the public at large, their existence is accepted by those who make use of UNIX. It is no act of genius or heroism to exploit such weaknesses.

A community of scholars should not have to build walls as high as the sky to protect a reasonable expectation of privacy, particularly when such walls will equally impede the free flow of information. Besides, attempting to build such walls is likely to be futile in a community of individuals possessed of all the knowledge and skills required to scale the highest barriers.

There is a reasonable trust between scholars in the pursuit of knowledge, a trust upon which the users of the Internet have relied for many years. This policy of trust has yielded significant benefits to the computer science community and, through the contributions of that community, to the world at large. Violations of such a trust cannot be condoned. Even if there are unintended side benefits, which is arguable, there is a greater loss to the community as a whole.

This was not a simple act of trespass analogous to wandering through someone's unlocked house without permission but with no intent to cause damage. A more apt analogy would be the driving of a golf-cart on a rainy day through most houses in a neighborhood. The driver may have navigated carefully and broken no china, but it should have been obvious to the driver that the mud on the tires would soil the carpets and that the owners would later have to clean up the mess.

Experiments of this kind should be carried out under controlled conditions in an isolated environment. Cornell Computer Science Department faculty would certainly have cooperated in properly establishing such an experiment had they been consulted beforehand.

The Commission suggests that media exaggerations of the value and technical sophistication of this kind of activity obscures the far more accomplished work of those students who complete their graduate studies without public fanfare; who make constructive contributions to computer science and the advancement of knowledge through their patiently constructed dissertations; and who subject their work to the close scrutiny and evaluation of their peers, and not to the interpretations of the popular press.