# Proceedings of the Twelfth

# Internet Engineering Task Force

## January 18-20, 1989 at University of Texas - Austin


Compiled and Edited
by

Karen Bowers
Phill Gross



March 1989

## Acknowledgements

I would like to extend a very special thanks to Bill Bard and Allison Thompson of the University of Texas-Austin for hosting the January 18-20, 1989 IETF meeting. We certainly appreciated the modern meeting rooms and support facilities made available to us at the new Balcones Research Center. Our meeting was <u>especially enhanced</u> by Allison's warmth and hospitality, and her timely response to an assortment of short notice meeting requirements.

I would also like to thank Sara Tietz and Susie Karlson of ACE for tackling all the meeting, travel and lodging logistics and adding that touch of class we all so much enjoyed. We are very happy to have you on our team!

Finally, a big thank you goes to Monica Hart (NRI) for the tireless and positive contribution she made to the compilation of these Proceedings.

Phill Gross

## Table of Contents

# 1. Chairman's Message

Phill Gross
NRI

## Chair's Message

In the last Proceedings, I mentioned that we were working to improve the services of the IETF. Perhaps the most visible example of that was the improved meeting logistics provided at the January meeting by Sara Tietz and Susan Karlson of ACE. In the future, ACE will be working closely with Karen Bowers to make our meetings run as smoothly as possible.

At the January meeting, I announced that IDEAS would undergo a facelift and become INTERNET-DRAFTS. The functionality will be the same, with the important difference that INTERNET-DRAFTS is NOT a document series. It is a repository for draft documents that will ultimately become RFCs. The new rules are simple and are summarized on the next two pages. With these simple new guidelines, we once again have a way to manage the many documents produced in the IETF.

At long last, we have scheduled the dates and locations for the next five meetings (beyond April 89). As always, the most difficult issue was making sure our dates did not conflict with other important events. Karen, with the help of the host sites and some others, deserves the credit for finally pulling off this bit of magic. So mark your calendars! These dates are being announced in ConneXions and ACM SIGCOMM CCR. The dates are:

25-28 Jul 89 (confirmed)
31 Oct - 3 Nov 89 (confirmed)
6-9 Feb 90 (tentative site)
1-4 May 90 (confirmed)
31 Jul - 3 Aug 90 (confirmed)

Karen also deserves the credit for producing this very detailed Proceedings. Karen actually went to the trouble of listening to all the tapes from the last meeting so that she could accurately capture questions and comments from the audience. Please feel free to comment to Karen or I about the format of the Proceedings. We are evolving toward a standard format and your comments are welcome.

A current listing of the IETF Working Groups is also given below. This listing gives a brief summary of the status of the groups. Several new groups are in formation and may hold their first meetings in Cocoa Beach. The CMOT WG is to be congratulated for submitting their implementor's agreement for release as an RFC.

I am looking forward to a productive meeting in Florida!

Phill Gross

IETF Working Group Status (March 1989)

| Working Groups | RFC/ Draft | Met 1/88 | Cur Rpt | Mtg 4/89 | Chair or POC (address) |
|---|---|---|---|---|---|
| Authentication | Yes | No | Yes | Yes | Jeff Schiller (MIT) jis@athena.mit.edu |
| CMIP-over-TCP (CMOT) | Yes | Yes | Yes | Yes | Lee LaBarre (MITRE) cel@mitre-bedford.arpa |
| Domain (new) | | | | Yes | Paul Mockapetris pvm@venera.isi.edu |
| Host Requirements | Yes | Yes | Yes | No | Bob Braden (ISI) braden@isi.edu |
| Interconnectivity | No | Yes | Yes | Yes | Guy Almes (Rice) almes@rice.edu |
| Internet MIB | Yes | Yes | Yes | No | Craig Partridge (BBN) craig@nnsc.nsf.net |
| LAN Manager (a subgroup of MIB) | | | | Yes | Amatzia Ben Artzi amatzia@amadeus. stanford.edu |
| NSFnet/Reg Monitoring | No | Yes | Yes | Yes | Susan Hares (Merit) skh@merit.edu |
| Open SPF-based IGP | Yes | No | Yes | Yes | Mike Petry (UMD) petry@trantor.umd.edu |
| Open Systems Routing | Yes | No | Yes | Yes | Marianne Lepp (BBN) mlepp@bbn.com |
| OSI Interoperability | Yes | Yes | Yes | Yes | Ross Callon (DEC) callon@erlang.dec.com |
| PDN Routing Group | No | No | Yes | Yes | C-H Rokitansky roki@isi.edu |
| Performance and CC | Yes | Yes | Yes | Yes | Allison Mankin (MITRE) mankin@gateway.mitre.org |
| Pt-Pt Protocol | Yes | Yes | Yes | Yes | Drew Perkins (CMU) ddp#@andrew.cmu.edu Russ Hobby (UC Davis) rhobby@ucdavis.edu |
| ST and CO-IP | Yes | Yes | Yes | Yes | Claudio Topolcic (BBN) topolcic@bbn.com |
| TELNET Linemode | Yes | Yes | Yes | No | Dave Borman (Cray) dab@cray.com |
| User Services (New) | No | Yes | Yes | Yes | Karen Bowers (NRI) bowers@sccgate.scc.com (Interim Chair) |

Future IETF Meeting Sites

| | |
|---|---|
| 11-14 April 89(confirmed) | Kennedy Space Center/Cocoa Beach Hilton |
| 25-28 July 89(confirmed) | Stanford University |
| 31 Oct- 3 Nov 89(confirmed) | University of Hawaii |
| 6-9 Feb 90(tentative site) | Florida State University |
| 1-4 May 90(confirmed) | Pittsburgh Supercomputer Center |
| 31 Jul- 3 Aug 90(confirmed) | University of Washington |

# What Ever Happened to IDEAS ??!!??

Current:

1 Directory -- IETF

Contains IDEAS and IETF information

Proposed:

2 Directories -- IETF and INTERNET-DRAFTS

IETF will contain information on IETF

INTERNET-DRAFTS will contain draft documents for the Internet community

Motivation: Provide a well known location for documents-in-progress.

The IETF directory will contain the following:

- Description of the IETF
- Description each Working Group
- Current WG status reports
- Information on past and future IETF meeting

  (dates, locations, agendas, meeting notes, attendance lists, etc)

E RULES FOR THE INTERNET-DRAFTS DIRECTORY:

1) ONLY DRAFT DOCUMENTS THAT WERE MEANT TO BE SUBMITTED ULTIMATELY AS RFCS WOULD BE PLACED IN THE INTERNET-DRAFT DIRECTORY.
2) ALL DOCUMENTS WILL BE IN RFC FORMAT, WITH (A) 'DRAFT' IN THE TITLE AND ON EVERY PAGE, AND (B) THE 'STATUS OF THIS MEMO' PARAGRAPH WILL BE LEFT BLANK. (THESE TWO EFFECTS WILL BE CHANGED ONLY WHEN SUBMITTED AS AN RFC).
3) DOCUMENTS IN THE DIRECTORY WOULD HAVE A FINITE LIFETIME. AFTER A CERTAIN PERIOD OF TIME (SUGGESTED: 6 MONTHS), THE DOCUMENTS WOULD EITHER BE ADVANCED TO RFCS OR REMOVED FROM THE DIRECTORY. THE AUTHOR (OR EDITOR FOR A GROUP) WOULD BE CONSULTED BEFORE SUCH A MOVE IS MADE. IF THE DOCUMENT IS RELEASED AS AN RFC, IT WILL BE REMOVED FROM THE DIRECTORY.
4) A STANDARD NAMING SCHEME WILL BE USED TO MAKE IT EASY FOR USERS TO LOCATE DOCUMENTS OF INTEREST. THE STANDARD NAMING SCHEME IS:

DRAFT-<TFNAME>-<WGNAME>-<ABBREVTITLE>-<REVNO>.TXT

WHERE 'TFNAME' AND 'WGNAME' ARE ABBREVIATIONS FOR THE TASK FORCE AND WORKING GROUP, RESPECTIVELY. 'ABBREVTITLE' IS A SHORTENED FORM OF THE TITLE, AND 'REVNO' IS THE REVISION NUMBER. IF THE DOCUMENT IS NOT BEING AUTHORED IN A TASK FORCE, THEN THE AUTHORS NAME WILL BE SUBSTITUTED FOR 'TFNAME' AND 'WGNAME'.

E CURRENT CROP OF IDEAS WILL HAVE ONE OF THREE DISPOSITIONS:

1) TO BE SUBMITTED AS RFCS (AFTER CONSULTATION WITH THE AUTHOR)
2) DELETE BECAUSE ALREADY RELEASED AS RFCS (OR BECAUSE OBSOLETE)
3) INSTALL IN INTERNET-DRAFTS DIRECTORY AS A DOCUMENT STILL-IN-PROGRESS.

- IDEAs would have 1. of 3 dispositions
  1) submit as RFC
  2) Delete (either already RFC or obsolete)
  3) Install in new dir. as work-still-in-progress

- Approximate breakdown of IDEAs
  1) 10 to be submitted as RFC
  2) 8 delete
     5 - already RFCs
     3 - obsolete
  3) 5 installed in new dir
  4) ~10 backlogged documents

# 2. IETF Attendees

---

Guy Almes

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

713-527-6038 almes@rice.edu

---

Phil Almquist

Stanford University
285 Clinton Park
San Francisco, CA 94103

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

---

William Anderson

MITRE Corporation
Burlington Rd.
MS.E066
Bedford, MA 01730

617-271-3388 wda@mitre-bedford.org

---

Bill Bard

University of Texas-System Office
COM 1
Austin, TX 78712

512-471-3241 bard@utadnx.cc.utexas.edu

---

Amatzia Ben-Artzi

3Com
3165 Kifer Rd.
Santa Clara, CA 95052

408-970-1177 amatzia@amadeus.stanford.edu

---

Larry Besaw

Hewlett-Packard
3404 East Harmony Road
Fort Collins, CO 80526

303-229-6022 lmb%hpcndaw@hplabs.hp.com

---

---

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@cray.com

---

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 22102

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

---

Tim Boyce

IBM
11400 Burnet Road
Austin, TX 78758

512-823-4100 none

---

Robert Braden

USC/ISI
4676 Admiralty Way
Marina del Rey, CA 90292

213-822-1511 braden@isi.edu

---

---

Ronald Broersma

Naval Ocean Systems Center
271 Catalina Blvd.
Bldg. A33
San Diego, CA 92152-5144

919-553-4470 ron@nosc.mil

---

Jeffrey Burgan

Wollongong Group
20350 Stevens Creek Blvd. #509
Cupertino, CA 95014

415-962-7100 jeff@twg.com

---

John Burruss

Wellfleet Communications, Inc.
12 DeAngelo Drive
Bedford, MA 01730

617-275-2400 burruss@bbn.com

---

Ross Callon

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

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

---

Stephen Casner

USC/ISI
4676 Admiralty Way
Marina del Rey, CA 90292

213-822-1511 casner@isi.edu

---

J. Noel Chiappa

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

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

---

---

Joseph Choy

NCAR
PO Box 3000
Boulder, CO 80307-3000

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

---

Tony Chung

Sytek
1225 Charleston Road
Mountain View, CA 94043

415-966-7430 sytek!anthony@hplabs.hp.com

---

Danny Cohen

USC/ISI
4676 Admiralty Way
Marina del Rey, CA 90272

213-822-1511 cohen@isi.edu

---

Michael Collins

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

415-422-4018 collins@nmfecc.arpa

---

James Davin

MIT LCS
Computer Science Lab, NE43-504
545 Technology Square
Cambridge, MA 02139

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

---

Clive Dawson

MCC
3500 W. Balcones Ctr.
Austin, TX 78759

512-338-3430 clive@mcc.com

---

---

Phil Draughon

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

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

---

Charles Eldridge

Sparta, Inc.
7926 Jones Branch Drive
McLean, VA 22102

703-448-0210 none

---

Robert Enger

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

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

---

Dino Farinacci

3Com Corporation
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

---

Jim Forster

cisco Systems
1350 Willow Road
Menlo Park, CA 94025

415-326-1941 forster@cisco.com

---

---

Craig Fox

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

612-424-4888 foxcj@network.com

---

Jose Joaquin Garcia-Luna

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

415-859-5647 garcia@sri.com

---

Farrell Gerbode

Rice University
Office of Network. & Planning
PO Box 1892
Houston, TX 77251-1892

713-527-4988 farrell@rice.edu

---

Chuck Gerlach

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

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

---

Phill Gross

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

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

---

Martin Gross

DCA/DCEC
1861 Wiehle Avenue
Reston, VA 22090

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

---

---

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 National Lab
PO Box 5509, L-561
Livermore, CA 94550

415-422-4200 hain@nmfecc.arpa

---

Nancy Hall

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

608-262-5945 nhall@cs.wisc.edu

---

Gary Haney

University of Tennessee
Y12 Plant, Bldg.9103
M/S 8141, PO Box 2009-8141
Oak Ridge, TN 37831

615-690-9639 haneyg@utkuxl.utk.edu

---

Susan Hares

MERIT/NSFNet
1075 Beal Ave., NDSB
Ann Arbor, MI 48109-2112

313-936-2095 skh@merit.edu

---

Bob Harris

BBN Communications
50 Moulton St.
Cambridge, MA 02238

617-873-1817 bharris@bbn.com

---

Gene Hastings

Pittsburgh Supercomputing Center
5000 Forbes Ave.
Pittsburgh, PA 15213

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

Sergio Heker

John von Neumann Center
665 College Rd. East
Princeton, NJ 08540

609-520-2000 heker@jvnca.csc.org

Russell Hobby

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

916-752-0236 rdhobby@ucdavis.edu

Steve Hunter

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

415-423-2219 hunter@nmfecc.arpa

Ole-Jørgen Jacobsen

Advanced Computing Environments
480 San Antonio Rd.
Suite 100
Mountain View, CA 94040

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

Van Jacobson

Lawrence Berkeley Laboratory
M/S 46A
Berkeley, CA 94720

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

---

Keith Jarett

Teknekron
2121 Allston Way
Berkeley, CA 94704

415-548-4100 none

---

Dan Jordt

University of Washington
Academic Computing Services
3737 Brooklyn Avenue NE
Seattle, WA 98105

206-543-7352 danj@blake.acs.washington.edu

---

Michael Karels

University of California at Berkeley
1107 Liberty Street
CSRG, CS Division
Berkeley, CA 94530

415-642-4948 karels@berkeley.edu

---

Dave Katz

Merit Computer Network/NSFNet
1075 Beal Avenue
Ann Arbor, MI 48103

313-763-4898 dave_katz@um.cc.umich.edu

---

Dave Kaufman

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

508-898-2800 dek@proteon.com

---

Kenneth Key

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

615-974-6721 key@utkuxl.utk.edu

---

---

Lee LaBarre

MITRE Corporation
Burlington Rd
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

---

Walter Lazear

MITRE Corporation
7525 Colshire Drive
M/S W422
McLean, VA 22102

703-883-6515 lazear@gateway.mitre.org

---

John Lekashman

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

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

---

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

---

Kirk Lougheed

cisco Systems
1360 Willow Rd. #201
Menlo Park, CA 94025

415-326-1941 lougheed@cisco.com

---

Dan Lynch

Advanced Computing Environments
21370 Vai Ave.
Cupertino, CA 95014

415-941-3399 lynch@isi.edu

---

Charles Lynn

BBN Communications
10 Moulton Street
Cambridge, MA 02238

617-873-3367 clynn@bbn.com

---

Allison Mankin

MITRE Corporation
7525 Colshire Dr.
McLean, VA 22102

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

---

Matt Mathis

Pittsburgh Supercomputer Center
4400 5th Ave.
Pittsburgh, PA 15213

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

---

---

Tony Mazraani

Washington University
CCRC
Campus Box 1115
St. Louis, MO 63130

314-726-4163 tonym@flora.wustl.edu

---

Keith McCloghrie

Wollongong Group
1129 San Antonio Rd
Palo Alto, CA 94303

415-962-7160 kzm@twg.com

---

Dan McKernan

Sterling Software
M/S 233-18
NASA Ames Research Center
Moffett Field, CA 94035

415-694-4795 mckernan@nsipo.nasa.gov

---

Leo McLaughlin

Wollongong Group
1129 San Antonio Road
Palo Alto, CA 94303

415-962-7100 ljm@twg.com

---

Milo Medin

Sterling Software
1121 San Antonio Road
Palo Alto, CA 943036

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

---

Bill Melohn

Sun Microsystems, Inc.
2550 Garcia Ave
Mountain View, CA 94043

415-336-2941 melohn@sun.com

---

---

Donald Merritt

US Army
Attn: AMXBR-SECAD
Aberdeen Proving Grounds, MD 21005-5066

301-278-6808 merritt@brl.mil

---

Paul Mockapetris

USC-ISI
4676 Admiralty Way
Marina del Rey, CA 90292

213-822-1511 pvm@isi.edu

---

Chris Moore

Wollongong Group
1129 San Antonio Road
Palo Alto, CA 94303

415-962-7160 cwm@twg.com

---

Ron Morita

Hewlett-Packard
19420 Homestead Rd
Cupertino, CA 95014

408-447-2197 ron_morita%y0%hp6600@hplabs.hp.com

---

Russ Mundy

DCA
Code B602
Washington, DC 20305-2000

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

---

Gerard Newman

GA Technologies Inc., SDSC.
POB 85608
San Diego, CA 92138

619-534-5076 gkn@sds.sdsc.edu

---

---

Carolyn Nguyen

AT&T Bell Laboratories
200 Laurel Ave.
Room 4C-222
Middletown, NJ 07748

201-957-5348 mhn@caelum.att.com

---

Rebecca Nitzan

Lawrence Livermore National Lab
L-561
PO Box 5901
Livermore, CA 94550

415-422-9775 nitzan@nmfecc.arpa

---

Zbigniew Opalka

BBN Communications
50 Moulton St.
Cambridge, MA 02238

617-873-2888 zopalka@bbn.com

---

Phillipe Park

BBN STC
10 Moulton Street
Cambridge, MA 02238

617-873-2892 ppark@bbn.com

---

Craig Partridge

BBN Systems & Technologies Corp.
10 Moulton Street
Cambridge, MA 02138

617-873-2459 craig@bbn.com

---

Francine Perillo

cisco Systems
1360 Willow Road
Menlo Park, CA 94025

415-326-1941 perillo@cisco.com

---

---

Drew Perkins

Carnegie Mellon University
4910 Forbes Ave
Pittsburgh, PA 15213

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

---

John Quarterman

TIC
701 Brazos
Suite 500
Austin, TX 78701-3243

512-320-9031 jsq@longway.tic.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

---

Gladys Reichlen

MITRE Corporation
7525 Colshire Dr.
McLean, VA 22102

703-883-7878 reichlen@gateway.mitre.org

---

Jacob Rekhter

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

914-945-3896 yakov@ibm.com

---

Robert Reschly

Ballistic Research Laboratory
ATTN: SLCBR-SF (Reschley)
APG, MD 21005-5066

301-278-6678 reschly@brl.mil

---

---

Jim Robertson

3Com/Bridge
2081 Stierlin Rd
Mountain View, CA 94043

415-940-7683 none

---

Larry Robinson

Hewlett-Packard
Colorado Networks Division,MS 45
3404 East Harmony Rd.
Fort Collins, CO 80525

303-229-2294 none

---

Milt Roselinsky

Communication Machinery Corporation
125 Cremona Drive
Santa Barbara, CA 93117

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

---

Karen Roubicek

BBN Communications-nnsc
10 Moulton St.
Cambridge, MA 02138

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

---

Greg Satz

cisco Systems
1360 Willow Road
Menlo Park, CA 94025

415-326-1941 satz@cisco.com

---

Martin Schoffstall

NYSERNet Inc
Rensselaer Technology Park
165 Jordan Road
Troy, NY 12180-3590

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

---

---

Bruce Schofield

DCEC
1860 Wiehle Avenue
Reston, VA 22090

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

---

Dallas Scott

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

202-883-6700 dscott@gateway.mitre.org

---

Rajeev Seth

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

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

---

Jim Sheridan

IBM
PO Box 334
Whitmore Lake, MI 48189

313-393-6537 jsherida@ibm.com

---

Keith Sklower

University of California at Berkeley
Evans Hall, Rm. 570
Berkeley, CA 94720

415-642-9587 sklower@okeeffe.berkeley.edu

---

Dave Solo

Sparta
7926 Jones Branch Drive
Suite 1070
McLean, VA 22102

703-448-0210 solo@dockmaster.arpa

---

---

Charles Spurgeon

University of Texas
Computation Center
Austin, TX 78712

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

---

Michael St. Johns

DCA
Code B612
Washington, DC 20305-2000

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

---

Louis Steinberg

IBM
472 Wheelers Farms Rd
MS 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 guest@edn-unix.dca.mil

---

Geof Stone

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

612-424-1693 geof@network.com

---

Zaw-Sing Su

SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025

415-859-4576 zsu@sri.com

---

---

James Sweeton

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

313-936-3000 jcs@merit.edu

---

Allison Thompson

Univ of Texas
10100 Burnet Road
Computation Center, COM1
Austin, TX 78712

512-471-2444 xxak163@chpc.brc.utexas.edu

---

Claudio Topolcic

BBN Communications
10 Moulton St.
Cambridge, MA 02238

617-873-3874 Topolcic@bbn.com

---

Lance Travis

Apollo Computer
330 Billerica Rd
Chelmsford, MA 01824

508-256-6600 cmt@apollo.com

---

James VanBokkelen

FTP Software
PO Box 150
Kendall Square Branch
Boston, MA 02142

617-868-4878 jbvb@ftp.com

---

Ross Veach

Univ. of Illinois at Urbana-Champaign
1304 West Springfield
Urbana, IL 61801

217-244-4274 rrv@uxc.cso.uiuc.edu

---

Steve Waldbusser

Carnegie-Mellon University
4910 Forbes Avenue
Pittsburgh, PA 15213

412-268-6628 waldbusser@andrew.cmu.edu

---

Asher Waldfogel

Wellfleet Communications, Inc.
12 DeAngelo Drive
Bedford, MA 01730

617-275-2400 ...linus!wellflt!awaldfog

---

Unni Warrier

Unisys
Defense Systems
2400 Colorado Avenue
Santa Monica, CA 90406

213-453-5196 unni@sm.unisys.com unni@cs.ucla.edu

---

Richard Wilder

MITRE Corporation
7525 Colshire Dr
M/S W425
McLean, VA 22102

703-883-6174 rick@gateway.mitre.org

---

Wayne F. Wilson

USAF Communications Command
1842EEG/EEM
Scott AFB, IL 62225-6001

618-256-4585 none

---

Bradley Wright

USAF
AF DDN Program Office
SSC/AQFC
Gunter AFB, AL 36114

205-279-4075 afddn.wright@gunter-adam.arpa

---

Jeff Wu

Hewlett-Packard
3404 E Harmony Rd
Ft. Collins, CO 80525

303-229-2243 jeff%hpcndm@hplabs.hp.com

# 3. Final Agenda

# Final IETF Agenda, 18-20 January 1989

This was the final agenda for the January 18-20 IETF meeting at the Balcones Research Center, University of Texas - Austin. The meeting was hosted by Bill Bard and Allison Thompson of the University of Texas System, Office of Telecommunication Services.

TUESDAY, JANUARY 17th, 9:00 am - 5:00 pm
Pre-IETF Working Group Meetings

- Internet Management Information Base (MIB) WG
(Craig Partridge, BBN)

- ST and Connection-Oriented IP WG
(Claudio Topolcic, BBN)

WEDNESDAY, JANUARY 18TH

9:00 AM Opening Plenary and local arrangements
(Gross, NRI / Thompson, U of Texas)

Working Group Sessions

9:15 am - 12:00 pm

- CMIP-over-TCP Net Management (Lee LaBarre, MITRE)

9:15 am - 5:00 pm

- Performance and Congestion Control (Mankin, MITRE)
- Point-to-Point Protocol (Perkins, CMU / Hobby, UCDavis)
- OSI/Internet Interoperability
(Callon, DEC / Hagens, UWisc)
- User Services (Bowers, NRI)
- St and Connection-Oriented IP WG
(Claudio Topolcic, BBN)

1:00 pm - 5:00 pm

- Interconnectivity and EGP3 (Almes, Rice)

5:00 pm - 7:00 pm

- Domain Name System WG (Mockapetris, ISI)

THURSDAY, JANUARY 19th

9:00 am Opening Plenary

9:15 am - 12:00 pm Working Group Sessions

- o Host Requirements (Braden, ISI)
- o TELNET Linemode (Borman, Cray)
- o Working Group for Joint Monitoring Access for Adjacent Networks focusing on the NSFnet Community (Hares, Merit)
- o Interconnectivity and EGP3 (Almes, Rice) Members only

1:15 pm - 5:00 pm Technical Sessions and Network Status Reports

- o Merit NSFnet Report (Hares, UMich)
- o Internet Report (Opalka, BBN)
- o DOE ESNET Report (Hain, LLL)
- o CSNET Report (Partridge, BBN)
- o Domain System Statistics (Lotter, NIC)
- o Report on the Interconnectivity and EGP3 WG (Almes, Rice)
- o Support for OSI Protocols in 4.4 BSD (Hagens, UWisc)
- o Report and Discussion on the Internet Worm (Karels, UCB)

5:00 pm - 6:15 pm FILM and tour of University of Texas Supercomputer Facility

FRIDAY, JANUARY 20TH

9:00 am Working Group Reports and Discussion

- o CMIP-over-TCP (CMOT) (LaBarre, MITRE)
- o Domain Name System (Mockapetris, ISI)
- o Host Requirements (Braden, ISI)
- o Internet MIB (Partridge, BBN)
- o Joint NSFnet/Regional Monitoring (Hares, Merit)
- o OSI Internet Interoperability (Callon, DEC / Hagens, UWisc)
- o Performance and CC (Mankin, MITRE)
- o Point-to-Point Protocol (Perkins, CMU)
- o ST and CO-IP (Topolcic, BBN)
- o TELNET Linemode (Borman, Cray)
- o User Services (Bowers, NRI)

12:00 pm - 12:30 pm Concluding Plenary Remarks and Group Discussion (Gross, NRI)

# 4. Working Group Reports/Slides
University of Texas-Austin
January 18-20, 1989

# CMIP-over-TCP (CMOT)

Lee LaBarre
Mitre Corporation

January 1989 IETF Meeting
CMIP-over-TCP (CMOT) Working Group Summary
Reported by Lee LaBarre (MITRE)

The CMIP-over-TCP (CMOT) Working Group resolved the issues remaining on an RFC describing the use of ISO network management architecture and protocols for management of TCP/IP networks. Coordination with the IETF MIB Working Group resulted in registration of ASN.1 identifiers in the Internet MIB which were necessary for operation of the CMIP protocol. Agreement was also reached on the ASN.1 format of a managed system identifier necessary for proxy management and inter domain management. All identifiers, and the single variable defined, are specific to CMOT, and do not effect the core MIB implemented in agents.

The CMOT Working Group voted to forward the document for registration as an RFC, with the recommendation that it be a Draft Standard with the status of Recommended. The title of the document is "The Common Management Information Services and Protocol over TCP/IP".

Subsequent to the IETF meeting, an updated document was sent to Phill Gross in accordance with the new procedures initiated in the IETF. The document is currently being reviewed by experts outside the CMOT Working Group, and will be forwarded by Phill Gross for registration as an RFC by February 10, 1989. Phill will also make the document available in the recently established IETF directory at SRI-NIC.

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: CMIP-Over TCP Net Management
Chairperson: Lee LaBarre

| Attendees: | E-Mail Address: |
|---|---|
| 1. Amatzia Ben-Artzi | amatzia@amadeus.stanford.edu |
| 2. Larry Besaw | lmb%hpcndaw@hplabs.hp.com |
| 3. Tim Boyce | 512-823-4100 |
| 4. Charles Eldridge | 703-448-0210 |
| 5. Gary Haney | haneyg@utkuxl.utk.edu |
| 6. Sue Hares | skh@merit.edu |
| 7. Steven Hunter | hunter@nmfecc.arpa |
| 8. Keith Jarett | 415-548-4100 |
| 9. Ken Key | key@utkuxl.utk.edu |
| 10. Lee LaBarre | cel@mitre.org |
| 11. Charles Lynn | clynn@bbn.com |
| 12. Keith McCloghrie | kzm@twg.com |
| 13. Dan McKernan | mckernan@nsipo.nasa.gov |
| 14. Carolyn Nguyen | mhn@caelum.att.com |
| 15. Jim Robertson | 415-940-7683 |
| 16. Milt Roselinsky | cmcvax!milt@hub.ucsb.edu |
| 17. Rajeev Seth | rajs%hpindbu@hplabs.hp.com |
| 18. Jim Sheridan | jsherida@ibm.com |
| 19. Lou Steinberg | louiss@ibm.com |
| 20. Unni Warrier | unni@sm.unisys.com |
| 21. Bradley P. Wright | afddn.wright@gunter-adam.arpa |
| 22. Jeff Wu | jeff%hpcndm@hplabs.hp.com |

- Completed RFC XXXX

  Common Management Information Services and Protocol over TCP/IP
  (Editors: L. Besaw & U. Warrier)

- Content
  - Part I : Tutorial
  - Part II ; Protocol Agreements
  - Appendix I ; Use with Internet MIB (RFC 1066)
  - Appendix II : Examples

The CMOT Protocol Architecture

(Based on ISO & Internet Stds)

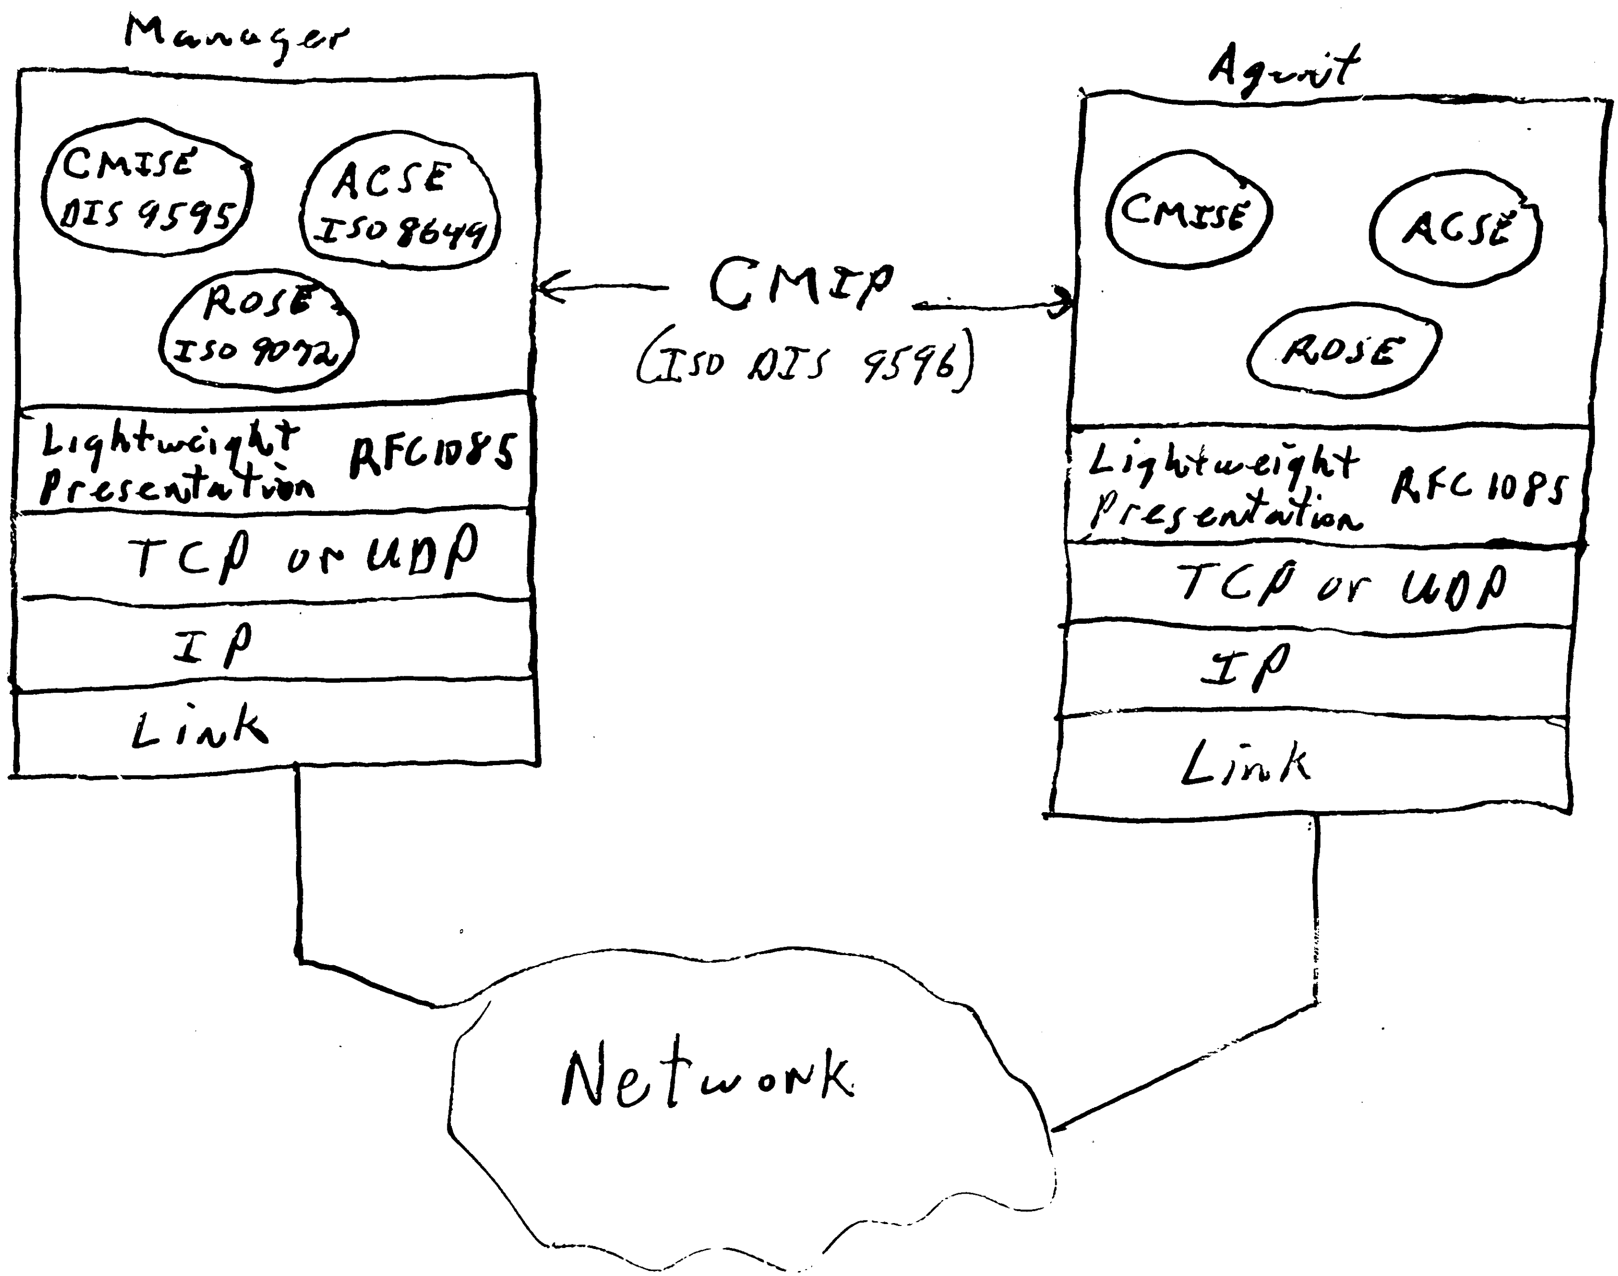

# Domain Name System

Paul Mockapetris
USC/ISI

January 1989 IETF Meeting
Domain Working Group Summary
Reported by Paul Mockapetris (USC-ISI)

The Domain WG met at the Austin IETF, made a list of areas for work and discussion, and came to an agreement about a method for naming networks and subnets.

The list was:

1. How should we map netnames to netnumbers (e.g. 26<=>MILNET) in the DNS?

2. What should the WG suggest to the Host Requirements WG?

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

4. A test/validation suite for the DNS.

5. How do DNS processes appear in the MIB.

6. Implementation catalog for DNS software.

7. Enhancements to the DNS in general.

8. Policy on load balancing.

9. Should MX be required in the Host Requirements.

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

Most of these received little, if any, discussion because of time limits on the WG meeting.

The general discussion regarding Host Requirements was that there were a lot of "NEVER DO THIS" items, but most were really covered in RFC 1034 and 1035, and hence didn't merit inclusion in HR. A majority of the WG felt that MX entries should be created for all hosts which appear in the RHS of mail addresses, but everyone recognized that there was no way to enforce this.

PVM is building a DNS test suite and implementation catalog. Submissions or suggestions to pvm@isi.edu are welcome. Writing to appear.

Drew Perkins and others brought up the issue of load balancing in the DNS. For example, suppose a given domain had multiple mail exchanges. Could its domain name server dynamically alter the precedence of MX RRs? Might a set of hosts behind one name hand out addresses based on load average? The weakest approach is to allow sorting/randomization of RRs. This is OK and even recommended so long as it doesn't create a situation where some order is required.

A name server can be set up to generate answering RRs on request, an use load average, etc. as criteria. This is OK so long as it also updates the zone serial and doesn't output long TTLs; that is, its OK so long as it simulates a behavior that could have happened due to zone updates. A caution here is that BIND has a severe bug related to short TTLs, and this can make compliance difficult. It is not clear what the exact policy here should be, talk to Drew if you feel strongly about it.

The remainder of the discussion was taken up by the issue of adding information which would allow translation between network names and addresses. The main issues were whether this feature should be added in a manner which would serve as a precedent for future additions, or be grafted onto the current IN-ADDR.ARPA tree. Another concern was whether the control and responsibility for the mapping should be with the NIC (which allocates network numbers) or the delegated authority for the network. There was also the question of whether network names would become domain names with lots of components and dots.

The results:

Network names become domain names.

The mapping is entered and maintained by the local authority.

Subnet naming is possible.

The mapping uses the host zero address in nets and subnets.

A general method should be specified to create a precedent, perhaps for UDP/TCP sockets, but this case should be handled specially.

Paul Mockapetris to draft RFC describing all this.

## IETF Working Group Roster
University of Texas at Austin
18–20 January 89

Working Group Session: Domain Working Group
Chairperson: Paul Mockapetris

| Attendees: | E-Mail Address: |
|---|---|
| 1. David Borman | dab@cray.com |
| 2. Jeffrey Burgan | jeff@twg.com |
| 3. James Davin | jrd@ptt.lcs.mit.edu |
| 4. Robert Enger | enger@gburg.scc.com |
| 5. Phill Gross | gross@sccgate.scc.com |
| 6. Ken Key | key@utkuxl.utk.edu |
| 7. Walt Lazear | lazear@gateway.mitre.org |
| 8. John Lekashman | lekash@orville.nas.nasa.gov |
| 9. Charles Lynn | clynn@bbn.com |
| 10. Dan McKernan | mckernan@nsipo.nasa.gov |
| 11. Bill Melohn | melohn@sun.com |
| 12. Don Merritt | merritt@brl.mil |
| 13. Russ Mundy | mundy@beast.ddn.mil |
| 14. Zbigniew Opalka | zopalka@bbn.com |
| 15. Drew Perkins | ddp@andrew.cmu.edu |
| 16. Robert Reschly Jr. | reschly@brl.mil |
| 17. Martin Schoffstoll | schoff@nisc.nyser.net |
| 18. Mike St. Johns | stjohns@beast.ddn.mil |
| 19. James VanBokkelen | jbvb@ftp.com |

/

# DOMAIN W.G.

ISSUES:

1) NETNAME <-> NUMBER MAPPING
   e.g. 26 <-> MILNET
2) HOST REQUIREMENTS VS DOMAINS
3) SEMANTICS FOR NEW TYPES, CLASSES, *
4) DNS TEST SUITE
5) MIB
6) Implementation CATALOG
7) ENHANCEMENTS
8) LOAD BALANCING
9) MX REQUIRED (HOST REQ)
10) DYNAMIC ADD & DELETE
~~11) BIND BASHING~~

DOMAIN W.G.

NETNAME MAPPING
GENERAL VS TAILORED APPROACH
2 APPROACHES:

I
- SETUP & MAINTAINED (INITIALLY) BY NIC
- COULD BE FLAT NAME SPACE OR COULD EVENTUALLY BE DELEGATED

II
- NETNAMES BECOME DOMAIN NAMES
- INSTANTLY DISTRIBUTE CONTROL ACCORDING TO ADD RETS CONTROL (IN-ADDR. ARPA)
- SUBNET (FEATURE?)

⇒ WRITE UP SCHEME I FOR USE FOR TCP-PORTS or SIMILAR

⇒ USE SCHEME II FOR NETNAMES

---

DOMAIN W.G.

[ DNS TEST SUITE
IMPLEMENTATION CATALOG
ask PUM @ ISI.EDU / NIC ]

[ LOAD BALANCING?
ask PUM or D Perkins ]

[ MIB
ask them ]

# Host Requirements

Robert Braden
USC/ISI

January 1989 IETF Meeting
Host Requirements Working Group Summary
Reported by Bob Braden (ISI)

The Host Requirements Working Group was scheduled to meet for one half-day session. We scheduled only a half day for this WG because there was every reason to believe its work was nearly complete. This WG had worked intensively on the Host Requirements draft all through 1988, with 6 meetings of the group and every extensive use of email for discussion (nearly 2Mbytes of messages so far). The document had grown to 180 pages, and the rate of change had noticeably diminished. Some of the detailed requirements were starting to oscillate from one meeting to the next, depending upon which subset of passionate advocates were present. Half a dozen 20-40 page critiques contributed by different people had been discussed and incorporated.

We hoped to miss by only a little our goal of publishing the RFC by December 31, 1988. The meeting in Austin was meant to be a final wrap-up.

It did not work out that way. A number of people came with extensive lists of complaints; Mike Karels in particular contributed hundreds of proposed changes. As a result, and in an effort to reach closure and publication, the chair scheduled an evening meeting; this lasted until midnight. In one area in particular, multihoming, the group was unable to reach a concensus, and decided to appeal to higher authority.

Minutes of the regular meeting were taken by Paul Mockapetris. These have been distributed to the mailing list. No minutes were taken at the evening meeting.

ATTENDEES:

Almquist, Anderson, Borman, Braden, Burgan, Chiappa, Davin, Eldridge, Gross, Jacobson, Karels, Key, Lazear, Lekashman, Lynn, Mathis, Maxraani, McCloghrie, Melohn, Merritt, Mockapetris, Partridge, Perkins, Reschly, Robertson, Seth, VanBokkelen, and Wilson.

FUTURE:

Another meeting of the Working Group will be required before the document can be published.

UNFINISHED BUSINESS:

The WG has endeavored to nucleate efforts to solve outstanding problems in the host architecture. The resulting efforts are generally incomplete, but ought to go forward. The most important are as follows:

A. Dead-Gateway Detection

This topic was discussed quite extensively in the Dave Clark Five, but the lessons learned in the DARPA-funded Internet research program were somehow lost after 1984. The Host Requirements RFC contains a general approach that is an elaboration of Dave Clark's ideas. A specific algorithm has been suggested, but it needs to be described in an RFC and tested experimentally.

B. MTU Discovery Option

RFC-1063 was published, proposing a pair of IP options to discover the mininum MTU on a path, using the ideas from the Kent Mogul SIGCOMM paper. A host implementation was made but not tested by Partridge at BBN. It was noted that this RFC involves a reliable delivery mechanism at the IP layer and is moderately complex to implement for a host, but trivial for a gateway.

To make progress in this area, we need to have gateways implement the MTU option. In the future, IETF ought to consider adding minimum-MTU computation to all routing protocols.

C. ICMP Gateway Discovery

An RFC describing a mechanism to inform hosts of the available gateways went through several different drafts but was not finished. However, the HRWG feels that this is only a small part of an important problem, the whole process of network initialization of a host, especially a diskless host. The HRWG recommends that a new WG be convened to develop a coherent approach to host initialization.

D. Type-of-Service

The break out of the chicken-and-egg stalemate between hosts and gateways, the Host Requirements RFC requires that hosts be able to configure and set TOS values in datagrams. The WG furthermore tried to frame "reasonable" values for the major application protocols, but did not reach closure on this. If values can be agreed upon, they can be published in a future Assigned Numbers RFC.

NOTABLE CONCLUSION:

The most constructive outcome of the meeting was the contribution of a new aphorism by Dan Lynch:

> "The difference between a host and a gateway is...
> one is paranoid, the other is schizophrenic."

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: Host Requirements
Chairperson: Bob Braden

| Attendees: | E-Mail Address: |
|---|---|
| 1. Philip Almquist | almquist@jessica.stanford.edu |
| 2. Bill Anderson | wda@mitre-bedford.org |
| 3. David Borman | dab@cray.com |
| 4. Bob Braden | braden@isi.edu |
| 5. Jeffrey Burgan | jeff@twg.com |
| 6. Noel Chiappa | jnc@lcs.mit.edu |
| 7. James Davin | jrd@ptt.lcs.mit.edu |
| 8. Charles Eldridge | 703-448-0210 |
| 9. Martin Gross | martin@protolaba.dca.mil |
| 10. Van Jacobson | van@helios.ee.lbl.gov |
| 11. Mike Karels | karels@berkeley.edu |
| 12. Ken Key | key@utkuxl.utk.edu |
| 13. Walt Lazear | lazear@gateway.mitre.org |
| 14. John Lekashman | lekash@orville.nas.nasa.gov |
| 15. Charles Lynn | clynn@bbn.com |
| 16. Matt Mathis | mathis@faraday.ece.cmu.edu |
| 17. Tony Mazraani | tonym@flora.wustl.edu |
| 18. Keith McCloghrie | kzm@twg.com |
| 19. Bill Melohn | melohn@sun.com |
| 20. Don Merritt | merritt@brl.mil |
| 21. Paul Mockapetris | pvm@isi.edu |
| 22. Craig Partridge | craig.bbn.com |
| 23. Drew Perkins | ddp@andrew.cmu.edu |
| 24. Robert Reschly Jr. | reschly@brl.mil |
| 25. Jim Robertson | 415-940-7683 |
| 26. Rajeev Seth | rajs%hpindbu@hplabs.hp.com |
| 27. James VanBokkelen | jbvb@ftp.com |
| 28. Wayne Wilson | 618-256-4585 |

"THE DIFFERENCE BETWEEN
A HOST AND A GATEWAY
IS...

ONE IS PARANOID AND
THE OTHER IS
SCHIZOPHRENIC"

DAN LYNCH

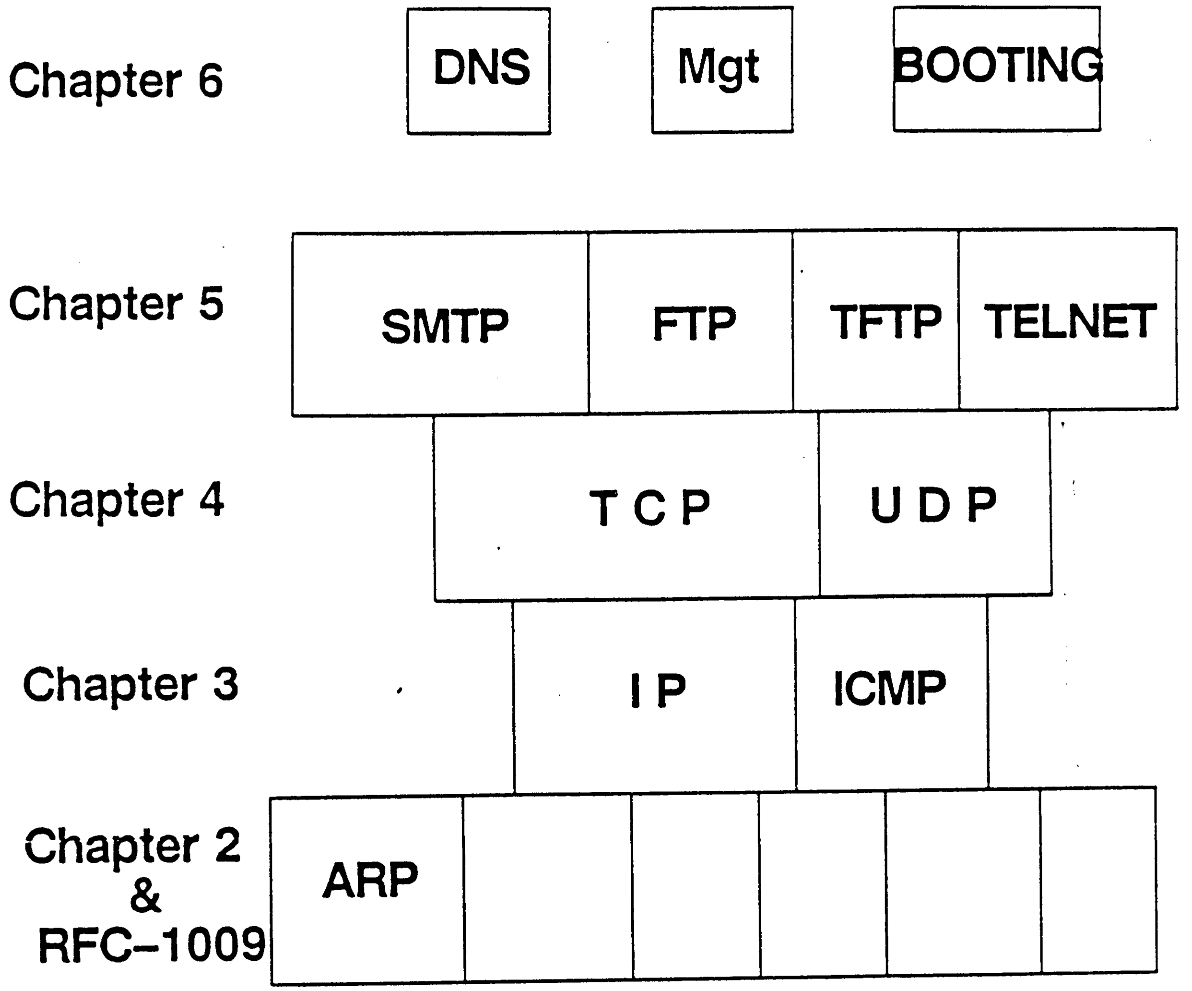


# HOST REQUIREMENTS RFC

- Requirements ...
  - "MUST" , "REQUIRED"
- Recommendations ...
  - "SHOULD" , "RECOMMENDED"
- Options --
  - "MAY" , "OPTIONAL"

## 7.4 TRANSPORT LAYER -- TCP

| FEATURE | SECTION | MUST | SHOULD | MAY | SHOULD NOT | MUST NOT | Footnote |
|---|---|---|---|---|---|---|---|
| **Ports** | | | | | | | |
| Well-Known Ports (1 - 255) | 4.2.2.1 | x | | | | | |
| Reserved ports (256 - X) | 4.2.2.1 | | | x | | | |
| **Push flag** | | | | | | | |
| Aggregate un-pushed data | 4.2.2.2 | | | x | | | |
| Send max-size seg even if not PSH | 4.2.2.2 | | x | | | | |
| Collapse successive PSH flags | 4.2.2.2 | | | x | | | |
| Application can specify PSH\|not | 4.2.2.2 | | | x | | | |
| Notify receiving user of PSH | 4.2.2.2 | | | x | | | |
| **Window** | | | | | | | |
| Treat as unsigned number | 4.2.2.3 | x | | | | | |
| Slow Start & Congestion Avoidance | 4.2.2.13 | x | | | | | |
| Probe zero window | 4.2.2.14 | x | | | | | |
| Allow rcvr window to stay closed | 4.2.2.14 | x | | | | | |
| **Urgent Data** | | | | | | | |
| Multiple bytes of urgent data | 4.2.2.4 | x | | | | | |
| Pointer points to last octet | 4.2.2.4 | x | | | | | |
| **TCP Options** | | | | | | | |
| Receive TCP option in any segment | 4.2.2.5 | x | | | | | |
| TCP ignore unsupported options | 4.2.2.5 | x | | | | | |

# PLANS...

- MEETING IN FEB '89
- SUBMIT IN MAR '89

ietf-hosts-request@
NNSC.NSF.NET

Host: venera.isi.edu
Path: pub/ietf-hosts.rfc.txt

---

## UNFINISHED BUSINESS

- DEAD GATEWAY DETECTION
  - General approach in RFC
  - Specific algorithm designed
  - Need: RFC
  - Need: Experiments
- MTU DISCOVERY OPTION
  - RFC published
  - Partial implementation
  - Recommendations (?)
    - Gateways implement option
    - Future routing protocols compute min{MTU}

# MORE U. B. ...

- ICMP GATEWAY DISCOVERY
  - RFC incomplete
  - Recommendation:
    - ▶ WG on Internet System Startup

- Type-of-Service
  - Required configurable TOS values
  - Tried to frame "reasonable" values

- RFC's Needed
  - ARP
  -

# Interconnectivity and EGP3

Guy Almes
Rice University

January 1989 IETF Meeting
Interconnectivity Working Group Summary
Reported by Guy Almes (Rice University)

This WG has been active since its creation in September 1988. We are focussing on a gap we perceive between (1) the time when the current EGP structure is no longer adequate and (2) the time when the Open Routing Working Group fields a really modern external routing mechanism. We hope to be more long-term than a mere fire-fighting effort, yet less long-term than the ORWG.

The work at the Austin meeting focussed on the routing information to be passed from one Autonomous System (AS) to its neighboring ASes. In EGP, this consists of a list of <network number, metric, next hop> triples. We had an extensive discussion of the benefits and complications of including a different set of information, namely a "full-AS-path" and a list of <network number> entries that are reachable through that full-AS-path. The full-AS-path consists of * a sequence of ASes, the first being the source AS, the last being the destination AS, and each AS being "directly connected" to the preceding AS in the path, * a metric associated with that path, and * possibly some other information associated with that path. Thus, at the cost of including slightly more information about each path, we can get rid of any per-network-number information. The anticipation is that the number of distinct full-AS-paths will be a few percent more than the number of distinct destination ASes, which in turn will be much fewer than the number of destination network numbers. Thus, the total amount of information exchanged will be smaller. Also, since almost all computation at machines involved in external routing will be on a per-AS rather than a per-network-number basis, there may be some savings here.

Given this inter-AS information exchange, the inter-AS routing database present at each AS will be somewhat richer, and we discussed at some length cases where this database could be used to make external routing decisions to reflect * performance preferences: for example, it might be known that one AS is implemented with full T1 circuits while another is satellite-based 9600b/s. Although information on bandwidth and delay is not present in the information exchanged, it might be known via an external source of information and coupled with the dynamic reachability information obtained via the protocol. * policy constraints: for example, it might be known that a given AS corresponds to a transit network with restricted usage. Paths using that restricted AS might be used for some source or destination networks but not for others. We do not imply that exchange of full-AS-paths makes such performance or policy choices easy, but it does provide fuller AS-level reachability which can be combined with static knowledge about performance or policy preferences.

As a collateral benefit, this database will allow ASes to detect anomolous situations with more complete information on their nature. This should make it much easier to debug misconfigurations.

Finally, we hope this will allow removal of the classical "stub model" constraint on inter-AS topology.

On the down side, we are very keen on minimizing the impact of such a scheme on vendors or other implementors of gateways. Several vendors present at the meeting expressed distinct misgivings about the possibility of many vendors having to implement a new complex protocol in synchrony.

Finally, some members of the ORWG voiced an interest in whatever scheme we came up with making some use of ideas accepted as important concepts within the ORWG. If we put those concepts to use and test them, then that might have a useful impact on the work of the ORWG.

Work within the IWG since the Austin meeting has focussed both on the details of the database and structuring the protocols so as to minimize the impact on gateway implementors while maximizing technical quality.

## IETF Working Group Roster

University of Texas at Austin
18–20 January 89

Working Group Session: Interconnectivity and EGP3
Chairperson: Guy Almes

| Attendees: | E-Mail Address: |
|---|---|
| 1. Guy Almes | almes@rice.edu |
| 2. Bob Braden | braden@isi.edu |
| 3. Jeffrey Burgan | jeff@twg.com |
| 4. Joe Choy | choy@ncar.ucar.edu |
| 5. Dino Farinacci | dino@bridge2.3com.com |
| 6. Mark Fedor | fedor@nisc.nyser.net |
| 7. Craig Fox | foxcj@network.com |
| 8. Jose Garcia-Luna | garcia@sri.com |
| 9. Susan Hares | skh@merit.edu |
| 10. David Kaufman | dek@proteon.com |
| 11. Mike Little | little@macom.arpa |
| 12. Kirk Lougheed | lougheed@cisco.com |
| 13. Milo Medin | medin@nsipo.nasa.gov |
| 14. Russ Mundy | mundy@beast.ddn.mil |
| 15. Becca Nitzan | nitzan@nmfecc.arpa |
| 16. Yakov Rekhter | yakov@ibm.com |
| 17. Greg Satz | satz@cisco.com |
| 18. Mike St. Johns | stjohns@beast.ddn.mil |
| 19. Geof Stone | geof@network.com |
| 20. Zaw-Sing Su | zsu@sri.com |
| 21. Ross Veach | rrv@uxc.cso.uiuc.edu |
| 22. Asher Waldfogel | ...linus!wellflt!awaldfog |

Membership:

Guy Almes, Rice University chair
Hans-Werner Braun, NSFnet
Mike Brescia, BBN
Scott Brim, Cornell
Joe Choy, NCAR
Phill Gross, NRI ex officio
Milo Medin, NASA
Russ Mundy, DDN
Jacob Rekhter, IBM
Mike St. Johns, DDN
Marianne Lepp, BBN
Rebecca Nitzan, ESnet
Ross Veach, UIUC

The Problem:

Address Inter-AS Routing during a period before results of ORWG are deployed.

Brush fires IWG ORWG

History:

[ (1987 – mid-1988)   Short-term routing ]

[ (July 1988)   NSFnet Backbone ]

| | | | |
|---|---|---|---|
| September | 1988 | NASA Ames | EGP3 |
| October | 1988 | IETF @ Ann Arbor | |
| January | 1989 | pre-IAB Workshop | |

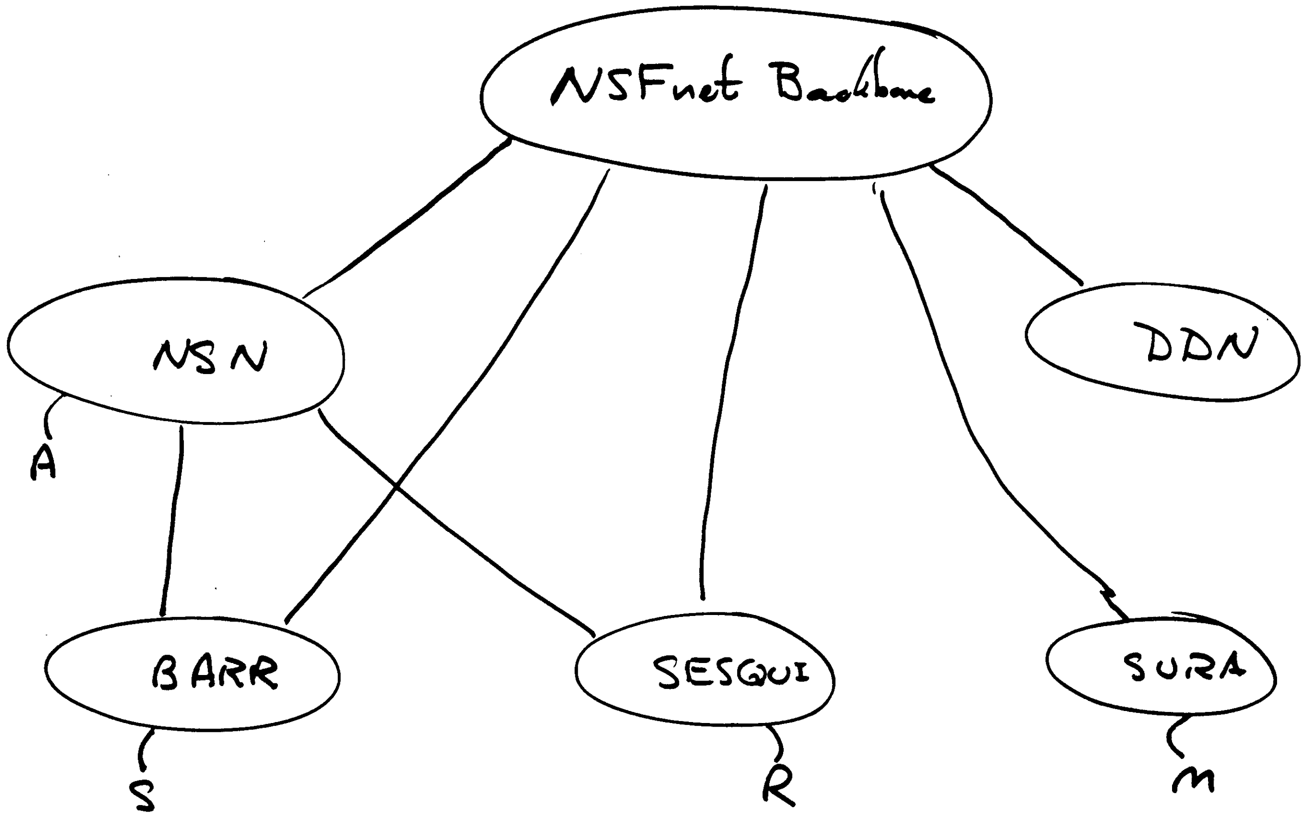


Essentially a DAG

Vertices are Autonomous Systems

Arcs are Inter-AS connections
Reachability exchanges

Two-phase rule: { Up } { Down }

# Interconnectivity Working Group ? EGP3

## Previous AS Field

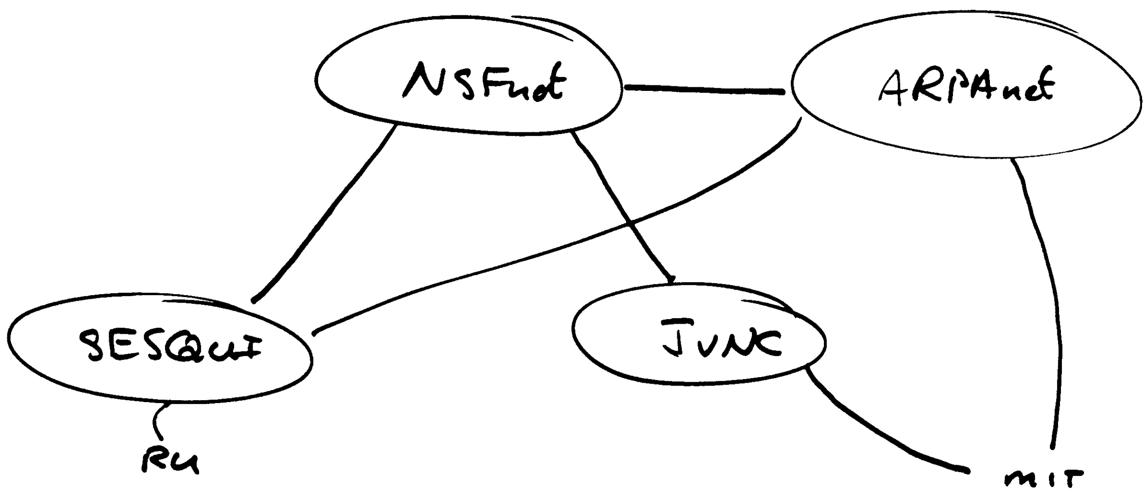


**Normal Case:**

| | | my AS | prev AS |
|---|---|---|---|
| JVNC | : | JVNC | — |
| NSF | : | NSF | JVNC |
| ARPA | : | ARPA | - |

∴ prefer NSFnet

**JVNC Fails:**

| | | | |
|---|---|---|---|
| NSF: | | NSF | ARPA |

∴ prefer ARPAnet

# General EGP3 Gains

- Avoidance of need for fragmentation.

- Incremental updates =>

  Less Transmission Cost

  Less CPU Cost

# DAG Enforcement

- Assign each Autonomous System a Level Number
- Constrain propagation of advertisements

  Enforce Two-Phase Rule

  Prevent Looping of Advertisements
- Possibilities Considered:

  Prefer Lowest Level #

  Inter-AS Hop Count

  Full Paths instead of Prev AS

  Remove Lateral Arcs

# Summary of Recommendations

1. Affirm EGP3 as in IDEA paper

2. Simplify metric

3. Add PreviousAS field for each net

4. Do not attempt generalized topology at this time.

5. Due to:
   - Current needs and status
   - Improved ORWG forecast

   call for implementation of EGP3

6. Continue steps to strengthen current Inter-AS routing.

# Internet MIB

Craig Partridge

BBN Systems & Technologies Corp.

January 1989 IETF Meeting
MIB Working Group Summary
Reported by Craig Partridge (BBN)

The MIB Working group met and discussed guidelines for the revision of the core MIB. At the meeting we decided that the key goals for the revision were:

1. To bring the core MIB into compliance with the Host Requirements specification (which at the time we expected to see issued at the end of January. The document is now delayed which makes compliance a little harder).

2. To enhance the generic error information stored in the generic interface section of the MIB.

3. To find a way to allow the generic interface structures to reference device-specific structures. The idea was to allow more detailed reporting on particular devices while retaining the advantages of having a generic structure that one can query without knowing exactly to what networks the device is connected.

We also decided on general criteria for changes to the core MIB in this revision:

- the MIB must continue to interface with both SNMP and CMOT

- we will allow "strong" writeable variables, but implementing the write function is optional. (We remain hamstrung by the lack of a standard authentication mechanism in the management protocols.)

- some experience with any proposed variable is required

- redundant variables are to be avoided

- no implementation specific variables are permitted in the core MIB

- heavy instrumentation in critical code sections is to be avoided

The group also decided that the core MIB should not expand to try to include information for more specialized systems such as bridges or TACs nor should it expand to include applications. Instead the group recommended that sub-groups be established to generate suggested MIB variables for particular applications (e.g. TELNET and FTP) and devices (e.g. TACs). The method by which sub-group sections of the MIB will become part of the Internet standard MIB was debated but not decided upon. Some groups favored an annual standards meeting at which the MIB working group decided which proposals were stable enough to become part of the Internet standard (at the risk of delaying upgrading a proposal to a standard for several months). Others preferred allowing each proposal to become a standard at its own pace (with the danger that the standard MIB would change very frequently).

Since the meeting, a list of guidelines for such groups has been issued. Interested parties should contact Craig Partridge (craig@bbn.com) for a copy before setting out to establish a sub-working group.

Finally, there was a debate about how to handle version numbers in the MIB. Some members believe the first MIB's practice of putting the version number into the name space was a mistake and want to change it. The proponents of both sides have promised position papers on the subject.

# MIB WG

- Decided we are working on just CORE MIB

- other groups to appear to do specialized vars for TACs, Bridges, etc.

---

- Criteria for MIB 2 CORE
  - Host Req mandated
  - must interface w/ CMOT + SNMP
  - Experience req
  - writeable vars (strong) OK but writeability is optional
  - avoid redundant vars
  - no implementation specific vars
  - avoid heavy instrumentation of critical code
  - core excludes applications etc

---

## MIB WG cont.

- Decisions made
  - Host Req to be reviewed + vars incorporated
  - IF layer to be enhanced
    - more generic error info
    - mechanism (pointer?) to point/refer to device/media specific info
  - CMOT gets section {mib 9} in the tree for CMOT vars
    SNMP gets {mib 10}

- Issues
  - How non-core subtrees get added?
    snowball vs incremental
  - can we find a better version # mechanism? (and what about versions for subtrees)

# Joint NSFNET/Regional Monitoring

Susan Hares

MERIT/NSFNET

January 1989 IETF Meeting
Jo-MANN Working Group Summary
Reported by Susan Hares (Merit)

The NSFnet Jo-MANN working group met on January 19th at the Houston IETF meeting. The meeting lasted from 9:30 - 11:30 am, the following people attended the meeting:

| | |
|---|---|
| Gene Hastings/PSCNET | (412)268-4960<br>hastings@morgul.psc.edu |
| Ross Veach/UIUC | (217)244-4274<br>rrv@uxc.cso.uiuc.edu |
| Russ Hobby/UC Davis | (916)752-0236<br>rdhobby@ucdavis.edu |
| Sergio Hecker/JVNC | (609)520-2000<br>heker@jvnca.cs.org |
| Gerard K Newman/SDSC | (619)534-5076<br>gkn@sds.sdsc.edu |
| Steve Waldbusser/CMU | (412)268-6628<br>waldbusser@andrew.cmu.edu |
| Dan McKernan/NASA | (415)694-4795<br>mckernan@nsipo.nasa.gov |
| Martin Schoffstall/NYSERNET | (518)283-8860<br>schoff@nisc.nyser.net |
| Geof Stone/Networks Systems Corporation | (612)424-1693<br>geof@network.com |
| Farrell Gerbode, Rice/Sesquinet | (613)527-4988<br>farrell@rice.edu |
| Robert Stone/SPARTA | (703)448-0210<br>quest@edn-unix.dca.mil |
| Bob Harris/BBN | (617)873-1817<br>bharris@bbn.com |
| Tim Boyce/IBM | (512)823-4100 |
| Rick Boivie/IBM | (203)783-7284<br>rboivie@ibm.com |
| Jim Sheridan/IBM | (313)393-5637<br>JSheridan@ibm.com |
| Zaw-Sing Su/SRI | (415)859-4576<br>ZSu@sri.com |
| Dan Jordt/NWNet, Seattle | (206)543-7352<br>danj@blake.acs.washington.edu |
| Sue Hares/NSFnet | (313)936-2095<br>skh@merit.edu |

Next Meeting:

March 6th at Ann Arbor, an evening session during an Advanced NSFnet Topic Seminar.

Agenda for the Meeting was:

1. Introduction

   NSFnet Jo-MANN is an IETF Working Group
   Duration: 6-9 months

Please note the increase in the time frame to 9 months. Map collection and getting tools in place is going slower than expected.

memo group: njm@merit.edu

2. Aids to Problem Solving

a) Maps - on-line are available via anonymous ftp on nis.nsf.net, the "cd maps" subdirectory.

Only a few on-line maps are on is.nsf.net. More Maps were collected at the IETF (and in the weeks after IETF. Sue Hares will try to get these on-line ASAP. Please report problems to Sue Hares skh@merit.edu.

Please send additional maps to NSFnet-ADMIN@merit.edu. Sue will post on the nis.nsf.net machine in the maps directory.

(note: Craig Partridge suggested at the IETF general session that we define an ASCII text format to describe maps. This ASCI format could be read into a program which would generate a Postscript map form.)

b) Maps in paper form will be mailed out in 2-3 weeks. Hopefully more maps will appear by then.

c) Contacts

Contacts for the NSS sites are available on the nis.nsf.net machine for anonymous ftp under the "cd NSFnet" subdirectory contact information is also available via email. Send to nis-info@nis.nsf.net a mail message with the following in the body: help. Instructions on how to use the email system will return to you. If you have any problems, either send a mail message to NSFnet-info@merit.edu or call (800)66-MERIT.

Please send in contacts for your Administrative Domain/AS. Please send two administrative contacts, two technical contacts, and a 24 hour security contact. The 24 hour security contact would be used, in case of "virus" like attack, to reach people.

Contacts for networks and Administrative domains will be available via email in the near future. Please ask the networks within your region to keep the SRI-NIC's who's database up to date. Sue Hares will send to Administrative Domains/Regional networks a list of network contacts for the networks within the regional networks. Please check these contacts to see if they are valid.

d) Configuration Information

Reports are available via anonymous ftp from nis.nsf.net. "cd nsfsites" will get you to the right subdirectory. (PLEASE note this is a change from the location reported at the IETF Jo-Mann meeting. (Sue Hares))

Configurations files for NSS are available via anonymous tp from nis.nsf.net, then cd nsfconfg, and "quote acct pwd" (pwd is the password sent to all NSS site people for this account)

(please note this subdirectory is different than the one announced at the IETF Jo-Mann meeting. (Sue Hares)) It is a goal of the NSFnet to make the configuration reports available via email. Watch for an announcement of this service.

e) SGMP/SNMP tool writeups available via anonymous ftp from nis.nsf.net, "cd tools". However, not all manufacturers are represented.

All sites should have an SGMP/SNMP session set-up with the community name "monitor" on all their gateways.

Proteon and CISCO both have SNMP support for their routers. Other vendors also have SNMP support. Please send any information about SNMP or SGMP support to Sue Hares. Any major time invested in tool development should be spent with SNMP tools. The NSFnet will migrate within the near future to SNMP.

f) Traceroute and unix kernel modification available from Van Jacobson for Suns, and unix systems. Russ Hobby has made modifications to KA9Q tcp-ip code for PC DOS. You can get a copy of the traceroute modifications for the KA9Q code via anonymous ftp from ucdavis.edu under the dist subdirector. Both archive and tar forms are available.

g) NSS outage reports

NSS outage reports are being sent to NSFNET-SITE-PEOPLE@merit.edu. The mail group NSFNET-SITE-PEOPLE@merit.edu will only report status information. The change has been made in response to the Jo-MANN request that a status only information group get created for the outage reports. Configuration information and technical information will be sent to NWG@merit.edu.

Farrell Gerbode noted that during the virus it took 8 hours to get mail from inside MERIT to the Internet. Sue Hares took it as an Action Item to find out what could be done about it.

[Note: Outage Mail is no longer being sent from the previously used machine. Mail should not be delayed within MERIT any more. Please contact the NSFnet NOC if you see large delays in your mail.]

3. 2 EGP PEERs

Currently MERIT, BARRNET, and Sesquinet have 2 EGP peers per Autonomous System. Sue Hares will add a 2nd EGP peer for any requesting the change. However, care must be taken with the local EGP metrics. The NSS will pay attention to the EGP metrics in the UPDATE messages when routing things from the NSS to the gateways within the region. As always, the local EGP metrics do not affect the primary/secondary pathway for a network within the NSS. The primary or secondary pathway to a networks is set by the NSFnet Routing Policy Database.

Contact Sue Hares to add your 2nd EGP PEER.

[UPDATE; JVNC has added a second EGP PEER. UIUC will add a second EGP within the week.]

4. Outage Reports

Currently NSFNET-SITE-PEOPLE receives outage reports for NSSs. Comments are requested if the outage reports would be useful for Administrative Domains and for networks.

What form should these outage reports take? On the NISC.NYSER.NET machine under the stats directory, you may find 7 different types of reports. The NSFnet monthly report also gives performance and statistics. Please send feedback to njm@merit.edu.

Marty Schoffstall asked which NOCs support 24 hour service. NYSERNET is currently expanding toward 18 hour service. However, Marty Stated Cornell does not staff their site on weekends or off hours. Marty stated this caused problems for NYSERNET. Sue Hares took it as an Action Item to follow up on this.

Are the network count statistic reports useful? These were shown as part of the NSFnet talk. Should these be posted?

5. NSFnet Security Issues

(The Security portion of the discussion was limited to members of the working group.)

Sue Hares suggested some ideas on security:

a) Each network should have responsible party you can contact if a host on a network has:

- "inappropriate" behavior - such as spreading a virus or trying to break into other computers.
- problems such as the SUN machine that send continuous name resolver requests to the MIDNET machines.

b) The responsible party should provide a 24 hour security number to be called only for these problems.

Networks are responsible that hosts behave properly. Regional Networks are responsible for networks.

Much discussion resulted from this topic. Two types of difficulties were discussed: political and technical. The political/legal problems were:

1. Suggesting some guarantee of security might imply some responsibility for problems if they occur. The X>25 Telenet contract is an example of a legal document for a network which has the customer sign away any idea that the network is responsible for any damage due to the network.

2. Gerard ask some questions of Steve Wolff regarding security and was going to forward this mail to the njm group.

Technical issues were:

1. How to set-up a phone tree for calling people.

   Normally we use email. However, during the virus mail got delayed due to the network shut downs. Farrell Gerbode suggested a phone tree.

2. How to get tools to detect virus or usage which is not normal sooner. Can SNMP monitoring tools find this problem? Does anyone have a script for this?

6. Debugging through National Networks

What problem areas are there in debugging through multiple networks? Please send your problems to njm@merit.edu. During the week of March 6-10th, NSFnet will be holding a workshop on Routing. One portion of the workshop will be on debugging through multiple networks.

Action Items:

Holdover's from the first meeting:

a) Get your Postscript Maps to Sue Hares
b) Add a SGMP or SNMP session called "monitor" to your regional gateways. Not all regions have added this. Sue will report on who has these sessions.
c) Send more SGMP/SNMP tool descriptions to Sue Hares
d) Sue Hares will post tool descriptions to "tools" subdirectory on nis.nsf.net
e) Sue Hares will post all MAPs to "maps" subdirectory on nis.nsf.net. Also, Sue will mail out a paper copy to NSFnet sites, and peer network.

New Action Items:

1. Sue will contact Cornell NSS site about weekend availability.

2. Gerard Newman will post his correspondence about Security to the jnm@merit.edu list.

3. Members of Jo-MANN will send comments about outage reports for Administrative Domains or networks to njm@merit.edu.

4. Send problems you have tracking problems through multiple peer networks. An example of this is a path between two networks which goes through both the NSFnet and the ARPANET.

5. Sue Hares will follow up on "Status only" mail group to check if sites want to change the mail IDs listed under NSFNET-SITE-PEOPLE@merit.edu.

Short Summary of the Jo-Mann Working Group January 19th Meeting

Annotated Agenda:

1. Aids gathered to solve problems:

a) Maps
on-line: nis.nsf.net - maps subdirectory paper copies in 3 weeks,

b) Contacts for
- o NSS sites
- o Regional Networks
- o Networks
- o International Contacts

NSS sites on nis.nsfnet - NSFnet directory, also available via email. Send to nis-info@nis.nsf.net a mail message with "help" in body.

Please send contacts for your Administrative Domain to nsfnet-info@merit.edu. Two administrative contacts, two technical contacts, and a 24 hour security contact are needed for each AS. Please make sure SRI-NIC database is correct for networks within your regional network.

c) Configuration Information

Reports in nis.nsf.net - nsfsites.
Configurations files in nis.nsf.net - nsfconfg.
Configuration information will be sent to NWG@merit.edu.
NSFNET-SITE-PEOPLE@merit.edu will be status only mail group.

d) SGMP/SNMP Tools

Please add "monitor" session to SGMP/SNMP session on gateways within your network. Send any information on SNMP/SGMP support in gateways or SNMP/SGMP tools to njm@merit.edu.

e) Traceroute

Available for unix, PC traceroute changes for KA9Q code are available from Russ Hobby via anonymous ftp from ucdavis.edu under dist subdirectory.

f) NSS outage reports

Send comments about the NSS outage reports or what you would like to see in outage reports to njm@merit.edu.

2. Outage Reports for networks within NSFnet community

Phone tree needed for worm or virus outbreaks.

3. NSFnet Security Issues

Lost of discussion on political and technical methods. Please refer to longer notes.

4. Debugging problems through National Peer Networks

Topic not discussed. Please send examples of problems to njm@merit.edu.

Next meeting:

March 6th - at MERIT during Advanced NSFnet Topics Seminar (very important meeting to attend)

April IETF Meeting - in Florida

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: Working Group for Joint Monitoring Access for Adjacent Networks Focusing on the NSFNet Community

Chairperson: Susan Hares

| Attendees: | E-Mail Address: |
|---|---|
| 1. Rick Boivie | rboivie@ibm.com |
| 2. Tim Boyce | 512-823-4100 |
| 3. Farrell Gerbode | farrell@rice.edu |
| 4. Bob Harris | bharris@bbn.com |
| 5. Gene Hastings | hastings@morgul.psc.edu |
| 6. Sergio Heker | heker@jvnco.csc.org |
| 7. Russ Hobby | rdhobby@ucdavis.edu |
| 8. Dan Jordt | danj@blake.acs.washington.edu |
| 9. Dan McKernan | mckernan@nsipo.nasa.gov |
| 10. Gerard Newman | gkn@sds.sdsc.edu |
| 11. Martin Schoffstoll | schoff@nisc.nyser.net |
| 12. Jim Sheridan | jsherida@ibm.com |
| 13. Robert Stine | guest@edn-unix.dca.mil |
| 14. Geof Stone | geof@network.com |
| 15. Zaw-Sing Su | zsu@sri.com |
| 16. Ross Veach | rrv@uxc.cso.uiuc.edu |
| 17. Steve Waldbusser | waldbusser@andrew.cmu.edu |

# Introduction

**NSFNET Jo-MAAN**

- IETF Working Group
- 6-9 month duration
- njm@merit.edu mail group

Sue Hares
NSFNET

# Agenda

1.) Introduction

2.) Aids to Problem Solving

3.) Multiple EGP PEERS
per AS

4.) Outage Reports

5.) Debugging through
National Networks

# Aids to Problem Solving

- Maps - online and hardcopy
- Contact information
- Configuration Reports
- SNMP/SGMP tools

SEND a MAP to

skh@merit.edu/

Correct a MAP

&

WIN a MAP

Packet

# Tools

- SGMP/SNMP tool writeup via anonymous ftp from nis.nsf.net, "cd tools"

- traceroute

# OSI Internet Interoperability

Ross Callon

Digital Equipment Corporation

January 1989 IETF Meeting
OSI Interoperability Working Group Summary
Reported by Ross Callon (DEC)

This was the first meeting of the working group in nearly a year, and the first to be co-chaired by Rob Hagens (Univ. of Wisconsin) and Ross Callon (DEC). We therefore spent some time discussing what the working group should be doing, particularly with respect to the "short term tasks" that were listed in the draft scope that we had sent out prior to the meeting. The short term part of the scope was extended, and each of the items on the list was discussed in some detail.

1. Working Group Name and Structure

The OSI working group has alternatively been called the "OSI Transition working group" and the "OSI Interoperability working group", as well as other names. We very quickly agreed that we were talking about interoperability, rather than transition, and therefore that the proper name for the group was "IETF OSI Interoperability working group".

Since we are involved in the entire issue of getting OSI to work in the Internet, our charter ranges from the application layer down to internetwork protocols (ISO 8473, ES-IS, and IS-IS). This is a huge scope. It is unlikely that any one person will be an expert on all relevant protocols. This suggests that we may want to split up into sub-groups at some time. One possible split would be into three groups: (i) Network and Transport layers; (ii) Session, Presentation, and Application layers; and (iii) Network Management. However, there are some issues which almost everyone will be interested in (such as to what extent, if at all, mixed protocol stacks are a good idea), and some issues for which the solution is so obvious that there is no need for a group to do any work. Also, we are not yet clear on precisely which issues will need work. We therefore thought that it was premature to decide whether or not to split into sub-groups, and if so which sub-groups to split into.

2. Berkeley Release 4.4

The next release of the BSD operating system (BSD 4.4) will include OSI protocols (in addition to the TCP/IP suite). Rob Hagens and Keith Sklower (Berkeley) gave an overview of the features to be included. A brief summary is:

| | |
|---|---|
| Data Link Layer: | LLC type 1 802.3 passive support for XID and test uncertain: X.25 link layer |
| Network Layer: | arbitrary OSI NSAP addresses<br>ISO 8473 (including source routing, record route, QOS and "DEC bit")<br>ISO 9542 (ES-IS) supports additional |

| | |
|---|---|
| | ECHO packet may be used as gateway<br>DOES NOT INCLUDE IS-IS partial support for ISO 8878 X.25 device driver |
| Transport Layer: | class 0 (over X.25) and 4 (over 8473)<br>all functions for classes listed<br>(including response to congestion bit)<br>also allows TP4 over DoD IP uses sockets |
| Upper Layers: | Session, Presentation ...<br>reliable transfer and remote operations<br>ASN.1 Directory Service (which not decided) MHS (X.400) FTAM VTP X.400/822 gateway FTAM/FTP gateway |

There are a few outstanding issues. Most notably, the Network Layer code does not contain any IS-IS routing code. They would like to include IS-IS, and would be "pleased as punch" if someone would send them a C implementation which is compatible with BSD. They also have not decided what directory services to use.

Someone asked if there were any plans to have the BSD code COS-certified, but nobody knew the answer. They hope to release BSD 4.4 in about a year. This Summer for Beta-Test sites.

3. Internet OSI NSAP Addressing Format

We discussed the proposed RFC written by Ross Callon and Hans-Werner Braun (from Merit/NSFNET/Univ. of Michigan), which has been waiting to be published for many months. Some frustration was expressed that it took so long to get documents through the RFC process. After some discussion, we agreed unanimously to urge that the RFC be published quickly in its current form (i.e., no changes are required).

Concern was raised that the RFC addressing format was different from the GOSIP format for the non-DoD Internet. We agreed that it would be best if we had only one format.

4. EON (Experimental OSI-based Network)

The University of Wisconsin is involved in an effort to experiment with the use of ISO Network layer protocols over the DoD Internet, using the DoD Internet as if it were a single subnetwork. Several months ago they produced an RFC (not yet published) giving an overview of this experiment.

In order to experiment with multicast subnets, given a DoD Internet which doesn't have multicast (or at least didn't back when they wrote the RFC), they make use of a sort of "send n times" multicast layer between the DoD IP and the ISO IP. There was considerable concern expressed that this would not scale at all to any reasonable size network. The counter-arguement was that the RFC was intended only for experimental use on a small-scale. This led to a discussion of the problem that anything put into an RFC is likely to be taken seriously, and find its way into government RFPs (requests for proposals) and other operational specifications or products. After considerable discussion, it was agreed that a section should be added at the beginning of the RFC as follows:

WARNING

The methods proposed in this RFC are suitable ONLY for experimental use on a limited scale. These methods are not suitable for use in an operational environment.

It was felt that anyone who ignored such a blunt warning deserved what he will get. With this addition, it was unanimously agreed to publish the RFC (possibly with an additional change to clarify a paragraph near the end.

5. Use of Mixed Protocol Stacks

Some people expressed interest in "mixed stacks", in which either (i) OSI application, presentation, and session layers are run over TCP/IP; or (ii) OSI application, ... are run over TP4, which is run over DoD IP. They seemed to be interested in an architecture in which a variety of different mixed and pure protocol stacks would be used in different end systems, implying a need for protocol converters which would convert the middle layers (Transport and/or Internet) while leaving the OSI applications alone.

Quite a few other people expressed opinions that this could lead to a logistical nightmare, and that it was a whole lot easier to transition the gateways to deal with both IPs in a dual internet (using some combination of dual gateways, and of separate DoD IP and OSI IP gateways on the same subnets). In addition, one might also want to have special gateways which make use of one Internet as if it were a single subnet of the other Internet. This should be able to interact smoothly with "real" gateways. In any case mixed stacks in the end system are particularly messy since: (i) it means you have to transition end sytems twice (assuming you don't want the end system to use a mixed stack permanently); (ii) it makes interoperability more complicated.

Concern was raised about the practice of using application layer gateways as a way to gain connectivity between like end systems (thus, for example, two pure OSI end systems might communicate through a DoD Internet using two application layer gateways). We agreed that application layer gateways were only for Interoperability between end systems with different applications.

6. GOSIP

Several non-DoD government representatives expressed concern that the DOD part of the GOSIP spec may in some cases (including NSAP addressing) propose different solutions relative to the non-DoD part of GOSIP. However, the Internet spans both DoD and non-DoD parts of the government (as well as universities and ...). It is clearly desireable for the Internet to make use of consistent protocol features (such as consistent address formats).

There were also questions about what to do if we find things in Gosip that we think need to be fixed. There really was no resolution of this, except that there will be future versions of Gosip and correcting errors is at least possible.

7. Network Management

We agreed that CMIP should be used to manage OSI protocols in the Internet, and that the MIBs developed in ANSI X3S3.3 are the best available (with the understanding that they may be changed in ANSI and ISO). It was not clear what MIBs should be used for higher layers.

8. Routing (Short Term)

It was suggested that we might want to use fixed tables for short term use. This was greeted by widespread dis-favour, at least with respect to intra-domain routing. It was then agreed that the full "compliant" (to the extent that this means anything at this time) implementations of the most recent ANSI routing spec was the correct thing to use for intra-domain routing (i.e., no subsetting, no enhancements, no attempt to convert other routing protocols for short term use).

There were two suggestions for short term inter-domain routing: (i) use fixed tables; (ii) use a version of EGP (basically update EGP 3 to include OSI-style address prefixes, call it EGP 4). We did not have enough time to discuss the relative merits of both, although there was some concern that fixed tables would not be sufficient, even for short term use. There was also concern expressed that it might take a long time for ISO to standardize a long term Inter-Domain routing protocol.

9. Application Gateways

As mentioned above, there was agreement that application gateways are useful for interoperability between dissimilar end systems, and not for connectivity between similar end systems via dissimilar intermediate systems. There was concern expressed about the limitations of application gateways (mostly functional, but also performance). There was a complaint that the group did not spend enough time on application gateway issues. We generally agreed, but were lacking a combination of time and expertise on these issues. We probably need to produce a sub-group to look at these issues.

10. Schedule, etc..

We didn't feel that we had enough of a handle on what the working group needs to do to try to produce a schedule for the short term tasks.

We discussed the possibility of having an intermediate meeting before the next IETF. If the meeting were held in "some single location that required any air travel" then only 9 (out of about 30) people thought that they would be able to attend. However, if the meeting were held by multi-media conference involving the four experimental multimedia sites (at BBN, Darpa, SRI, and ISI) then 19 people thought that they could attend. It was my impression that the problem was primarily that people are burned out w.r.t. travel, rather than the cost. I was therefore requested to look into the possibility of setting up a multimedia conference.

Subsequent to the meeting I discovered that the experimental multimedia conference facilities will be down for approximately the month of March, and will only be reliable for two-site conferences until after the work that will be done in March. I am therefore trying to get a two site (west coast and east coast, probably ISI and BBN) conference arranged for the end of February.

At the end of the meeting we went though each of the areas mentioned above and tried to list the documents which one ought to be familiar with.

## IETF Working Group Roster
University of Texas at Austin
18–20 January 89

Working Group Session: OSI/Internet Interoperability
Chairpersons: Ross Callon/Rob Hagens

| Attendees: | E-Mail Address: |
|---|---|
| 1. Rick Boivie | rboivie@ibm.com |
| 2. John Burruss | burruss@bbn |
| 3. Ross Callon | callon@erlang.dec.com |
| 4. Noel Chiappa | jnc@lcs.mit.edu |
| 5. Anthony Chung | sytek!anthony@hplabs.hp.com |
| 6. Jim Forster | forster@cisco.com |
| 7. Chuck Gerlach | cag@iwlcs.att.com |
| 8. Martin Gross | martin@protolaba.dca.mil |
| 9. Rob Hagens | hagens@ls.wisc.edu |
| 10. Tony Hain | hain@nmfecc.arpa |
| 11. Bob Harris | bharris@bbn.com |
| 12. Dan Jordt | danj@blake.acs.washington.edu |
| 13. Dave Katz | dave_katz@um.cc.umich.edu |
| 14. Norbert Leser | nl@osf.org |
| 15. Kirk Lougheed | lougheed@cisco.com |
| 16. Leo McLaughlin | ljm@twg.com |
| 17. Chris Moore | cwm@twg.com |
| 18. Ron Morita | ron_morita%y0%hp6600@hplabs.hp.com |
| 19. Russ Mundy | mundy@beast.ddn.mil |
| 20. Gerard Newman | glen@sds.sdsc.edu |
| 21. Becca Nitzan | nitzan@nmfecc.arpa |
| 22. Zbigniew Opalka | zopalka@bbn.com |
| 23. Gladys Reichlen | reichlen@gateway.mitre.org |
| 24. Yakon Rekhter | yakov@ibm.com |
| 25. Larry Robinson | 303-229-2294 |
| 26. Dallas Scott | dscott@gateway.mitre.org |
| 27. Keith Sklower | sklower@okeeffe.berkeley.edu |
| 28. Dave Solo | solo@dockmaster.arpa |
| 29. Lance Travis | cmx@apollo.com |
| 30. Ross Veach | rrv@uxc.cso.uiuc.edu |
| 31. Rick Wilder | rick@gateway.mitre.org |

# OSI INTEROPERABILITY WORKING GROUP

ROSS CALLON (DEC)
ROB HAGENS (U.Wisc.)

---

## MAIN GOALS

- FACILITATE INCORPORATION OF O.S.I. IN THE INTERNET
- OSI SUITE IN PARALLEL WITH TCP/IP
- FACILITATE COEXISTENCE

## SHORT TERM SCOPE

- W.G. STRUCTURE & SCHEDULE
- INTERNET NSAP ADDR. FORMAT
- EON
- PURE VS MIXED STACKS
- BSD 4.4
- GOSIP (DoD vs NON-DOD INTERNET)
- GWY ISSUES, SWITCHING OSI-Grams in the INTERNET

## SHORT TERM SCOPE-CONT'D

- NETWORK MANAGEMENT
  - -PROTOCOL (CMIP)
  - -MIB (ANSI)
- ROUTING
  - -INTRA-(ANSI)
  - -INTER-(??)
- APPL'N GATEWAYS

# Performance and CC

Allison Mankin

MITRE Corporation

January 1989 IETF Meeting
Performance and Congestion Control Working Group Summary
Reported by Allison Mankin

The Performance and Congestion Control Working Group met in Austin on January 18. The meeting had two parts. In the morning, the draft of the working group paper ''Gateway Congestion Control Policies'' was the subject. About a third of the draft received detailed review. In the afternoon, IETF colleagues who implement gateways were invited to give an early reaction to the paper and to talk with the group about the feasibility of implementing various congestion control policies in their gateways. The speakers were Yakov Rekhter (IBM), Bob Harris and Zbigniew Opalka (BBN), Noel Chiappa (MIT/Proteon), Greg Satz (cisco), and Geof Stone (NSC). We thank them all for taking part. Thanks also to John Lekashman and Walt Lazear for their useful notes on the morning and afternoon respectively.

Morning Session

The reason for writing a paper on gateway congestion control is that within the timeframe of six months to two years, we expect significant Internet congestion will be a problem again. Infinitely persistent name resolvers are just one example of unpredictable large contributors to traffic. Various remote file systems are being modified to work well over the Internet; their tuning allows them to consume large amounts of bandwidth. Gateways will need to have strong congestion control for these within six months (Van).

Since the need is in the short-term, the recommended policy must be one that is easy to add to existing gateway implementations. Group consensus is that we can and should now recommend random drop in the paper. The draft needs to state this up front, and to correspondingly emphasize the section detailing random drop over the sections on other policies. [Note: the new draft incorporates this and other changes suggested by the group.]

Some assumptions of our paper were clarified during the morning. On fairness, we were aware that not all gateways should allocate their resources equally across users; some users, such as mailers on relay hosts, properly have much higher demand than others. The term 'user' here means (as in the draft) whatever process or entity makes use of an Internet transport service. As we have noted in the draft, the means of allocating and administering the unequal shares is outside the scope of congestion control. But during this meeting we also agreed that gateways in network backbones should not prefer or discriminate in this way. Our recommendations directly address this type of gateway.

The inclusion of the ''congestion experienced bit''in the paper was debated, with objectors arguing that it cannot be relevant to the DoD Internet. One argument for its relevance is the not so far off arrival of dual DoD-ISO routers on the Internet scene. The bit will probably be part of a heterogeneous congestion control environment; what can be done about this is anticipate the interactions among policies.

Concerning the sections of the draft on metrics and measurements, Bob Braden asked if there was experimental evidence that power was a good metric, since it is used very centrally in the draft; power is the defined as (throughput ^ alpha) / delay. The only pitfall is in trying to obtain a global power through the measurement of the individual gateways (K.K.). Computed and optimized for each gateway separately, it is useful and has the property of getting worse in both directions from the optimum. It is very important to state all the metrics are time-varying and that algorithms need to use averages over appropriate time-scales; transient effects of performance policies must not be given too much weight.

End-system cooperation (such as provided by Slow-start TCP) is needed for gateway congestion control to be effective. Non-cooperating users are a great concern. The group agreed only on a minimum game theory approach to them: guarantee that they cannot get better service by not cooperating. It may not be possible to prevent noncooperative users from causing other users of the gateway to have worse service (Van). We need to put more emphasis on the second of the purposes of congestion control in the gateway:

1. Protecting oneself from end-systems and the upstream gateways, defending one's own resources.

2. Protecting the downstream gateways from oneself.

A policy that both protects the 'downstream' and carries out the game theory goal would be for the gateway to reduce all of its output in the direction of the noncooperating user's data stream (Van). This startled the group, though it was appreciated that such behavior could induce the administrators of the gateway gateway to ''dynamically buy more bandwidth.'' A question was raised as to whether we needed to emphasize non-cooperating users quite this much... [Yes, because significant expected problems in the short-term will be due to uncontrolled traffic such as NFS.] An obvious question is why not administer selective punishment to the badly behaved? With random drop, non-cooperating users are quickly and accurately distinguished (their packets get dropped repeatedly). In Van's implementation, a ''hot list'' of the offenders is generated. Their traffic can be selectively blocked by putting it on a minimum service queue, attended to only when no other traffic is present. An additional suggestion would be to give selective feedback as in DEC TR-510 based on a list like Van's (K.K.).

Finally, concerns surfaced about random drop's fairness. Matt Mathis (CMU) described scenarios of end-system user types which appear cooperative, but still use substantially more than their fair share of resources. Theoretically a transport protocol could enable the process using it to appear to all gateway tests as if it was two or more statistically independent users. A transport protocol could also do Slow-start congestion avoidance but ''use a higher gain'' and, depending on the eventual implementation of the gateway congestion control, not have its excess demand curtailed. At this point, we have to admit that random drop will be vulnerable to such counter-measures. This is one of the reasons that we see random drop as a starting point for gateway congestion control in the Internet, to be supplemented by unequal service and type-of-service algorithms when these are better understood.

Afternoon Session

A copy of the draft paper and an agenda were provided to each of the speakers (either the day before or on the morning of the 18th). The intention was to have the same range of topics covered by each. The agenda was as follows:

> The WG is working on a draft paper on Gateway Congestion Control Policies. We have invited gateway implementation experts to join us to review the assumptions about feasible and useful policies that are reasonable for this paper. The following is a list of some of the points that would be helpful to know about. At no point is the intention to ask for proprietary information. All questions on implementation are oriented toward the review of the paper.
>
> What are the performance goals set for your router? That is, what kind of yardstick is used when performance enhancement is undertaken? If possible, what is a high level overview of the implementation? For example, how many levels of buffering are there? Would it be possible to implement Random Drop? What would happen if TTL was required to be a time? Would a gateway system delay bound, suggested as an an alternative to TTL as a time, be possible and effective? How accurate a clock resolution could be assumed?
>
> Is Fair Queueing implemented in the implementation and what is the aggregation? What is the Source Quench generation policy?
>
> If the implementation includes an X.25 interface, what is done to handle fair provision of virtual circuits? To what extent could built-in instrumentation allow the router's

fairness to be assessed? Has operational measurement been done much? Is there any cooperative operational measurement among the router's users?

All the speakers were generous with details of their implementations. Many of their details naturally concerned the optimization of packet processing and bus performance, and need not be reported here. The performance of their gateways as part of the Internet system was somewhat less of a natural focus. As a general conclusion from the afternoon, however, it appeared that gateway implementors would be hospitable towards a policy for gateway congestion control recommended by the IETF.

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: Performance Congestion Control
Chairperson: Allison Mankin

| Attendees: | E-Mail Address: |
|---|---|
| 1. Philip Almquist | almquist@jessica.stanford.edu |
| 2. David Borman | dab@cray.com |
| 3. Bob Braden | braden@isi.edu |
| 4. Jeffrey Burgan | jeff@twg.com |
| 5. Noel Chiappa | jnc@ks.mit.edu |
| 6. Michael Collins | collins@nmfecc.arpa |
| 7. James Davin | jrd@ptt.lcs.mit.edu |
| 8. Mark Fedor | fedor@nisc.nyser.net |
| 9. Nancy Hall | nhall@cs.wisc.edu |
| 10. Bob Harris | bharris@bbn.com |
| 11. Van Jacobson | van@helios.ee.lbl.gov |
| 12. Walt Lazear | lazear@gateway.mitre.org |
| 13. John Lekashman | lekash@orville.nas.nasa.gov |
| 14. Charles Lynn | clynn@bbn.com |
| 15. Matt Mathis | mathis@faraday.ece.cmv.edu |
| 16. Keith McCloghrie | kzm@twg.com |
| 17. Dan McKernan | mckernan@nsipo.nasa.gov |
| 18. Leo McLaughlin | ljm@twg.com |
| 19. DonMerritt | merritt@brl.mil |
| 20. Zbigniew Opalka | zopalka@bbn.com |
| 21. K. K. Ramakrishnan | rama%erlang.dec@decwrl.dec.com |
| 22. Robert J. Reschly Jr. | reschly@brl.mil |
| 23. Bruce Schofield | schofield@edn-unix.arpa |
| 24. Robert Stine | guest@edn-unix.dca.mil |
| 25. Geof Stone | geof@network.com |
| 26. Zaw-Sing Su | zsu@sri.com |
| 27. James VanBokkelen | jbvb@ftp.com |
| 28. Steve Waldbusser | waldbusser@andrew.cmu.edu |

PERFORMANCE + CONGESTION CONTROL WG MEETING - JANUARY 18, 1989

MORNING - REVIEWED DRAFT OF GATEWAY CONGESTION CONTROL POLICIES PAPER

OBJECTIVE - RECOMMEND (WITHIN COHERENT FRAMEWORK) RANDOM DROPPING AS OUR LOWER BOUND CC POLICY FOR GATEWAYS

- FRAMEWORK NEEDED INCLUDES

  RESOURCE POWER METRIC - SHARED MEANING OF INTERNET OPERATING POINT

  FAIRNESS

  MONITORABLE, CONTROLLABLE METRICS

- PAPER IS AGREEING ON REQUIREMENTS OF THOSE ALGORITHMS THAT WILL BE IMPLEMENTED

SCHEDULE:

COMPLETE DRAFT WAS PROVIDED TO GROUP ON MAIL LIST BEFORE THIS MEETING - OUR WG DISCUSSION GOT THRU ~1/2.

CONTINUE BY E-MAIL

PERF WG
JAN 1989

AFTERNOON - INFORMAL PRESENTATIONS TO A SET OF QUESTIONS BY 5 GATEWAY IMPLEMENTATION EXPERTS

MOST OF THE INFO - WHAT IS DONE TO GET THE LEAST PROCESSING LATENCY IN THE GW IMPLEMENTATIONS

VALUE: VERIFIED WE COULD RECOMMEND R.D. AND IT WOULD BE FEASIBLE

SEVERAL CONTRIBUTED INFO DIRECTLY IN LINE WITH PAPER'S GOALS - WHAT ARE THE DESIGN/IMPL. POINTS THEY HAVE FOR INTERACTION AMONG GATEWAYS

# Point-to-Point Protocol

Drew Perkins
Carnegie Mellon University

and

Russ Hobby
UC-Davis

January 1989 IETF Meeting
Point-to-point Protocol Working Group Summary
Reported by Russ Hobby (UCDavis)

## I. Introduction

The PPP WG met for the day on January 18. There was a review from the previous meeting covering the physical layer and link framing of the protocol. There was also a discussion of the draft document written by Drew Perkins. The remainder of the meeting was spent working on the PPP Packet format.

The work group plans to have the current work written up and have a video conference half way towards the next IETF meeting.

## II. Physical Layer and Link Framing

The physical layer for PPP can be any of the standards currently in use, such as RS232, RS422, RS423, RS449, V.35. Use of control signals on the physical interface is encouraged. There was a request from Telebit for raising and lowering control signals on packet boundaries.
Van Jacobson feels that this is unnecessary and will discuss it with Telebit.

The link framing will be considered for only synchronous and asynchronous circuits for the time being. For both synch and asynch the PPP standard will be ISO 3309 (HDLC). HDLC has been extended to asynch by Addendum 1 to ISO 3309

## III. PPP Packet Formatting

The PPP packet is encapsulated in the data field of the HDLC frame and consists of a protocol field and a data field. The protocol field defines what type of network packet the data field contains, thus PPP will allow multiple network protocols to operate on the same circuit.

The protocol field can be either 8 or 16 bits in length with the length being determined by the standard ISO extension method. Protocol numbers for the protocol field will be odd numbers starting at 33. This avoids sending control characters for asynch circuits.

## IV. PPP Control Packet Format

PPP Control Packets (PCP) are used for line control and maintenance. This includes functions such as establishing the initial configuration, determining loopback, up/down control, circuit disconnect and other line oriented functions. The protocol number for the PPP Control Packets is 33. The PCP packet format is as follows:

```
 0                   1                   2                   3 3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      protocol (value = 33)    |    version    |     type      |
+---------------------------------------------------------------+
|                         magic number                          +
+---------------------------------------------------------------+
|                            data                               +
+---------------------------------------------------------------+
```

a. Version - The version of PCP supported.

b. Type - Type of PCP packet. Defined types are:

```
 1 - Configuration           2 - Config Ack
 3 - Config Huh?             4 - Config Nak
 5 - Echo Request            6 - Echo Reply
 7 - Version Reject          8 - Type Reject
 9 - Terminate Request      10 - Terminate Ack
11 - NOP
```

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 PCP packets sent out must contain the senders magic number (see discussion on loopback detection).

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

PCP Packet Types

1. Configuration - This packet type is sent out the line to indicate pertinent configuration information and is used to establish a connection. Configuration Items (CI) are placed in the data field of the PCP. Multiple CI's may be included in each packet. The format of a CI is:

```
 0                   1                   2                   3 3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            length             |              type             |
```

```
+-------------------------------------------------------------+
|                            data                             +
+-------------------------------------------------------------+
```

a. Length - inclusive length of the CI.

b. Type - Type number of the CI.

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

CI's will provide information on MTU, asynch character mapping and, perhaps, address exchange, compression, icryption, authentication. There is a current discussion on if these last items belong in configuration information for the line or if they should be done independently for each protocol within that protocol's packets.

2. Config Ack - This packet type is sent in response to a configuration packet and indicates acceptance of the other ends CI's. Any information in the PCP data field may be ignored.

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 PCP 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 PC data field will contain the CI entries of the unacceptable CI(s).

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

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

7. Version Reject - This packet type is sent in response to a PCP packet of an unacceptable version. Any information in the PCP data field may be ignored.

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

9. 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 PCP data field may be ignored.

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

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

V. Loopback Detection

Since symmetrical protocols are often used on point-to-point circuits, it is often difficult to determine if you are truly talking to the other end or a loopback of your own data. The Magic Number field of the PCP is used at various times for loopback detection. Magic Numbers are selected so as to maximize the probability the number will be unique to the device. Possible sources for the Magic Number are the device serial number, ethernet number, low order time-of-day bits, or any other random number source.

Initial loopback detection is done when sending out the first Configuration packet. If a Configuration packet is returned that contains the Magic Number that was selected, the line is probably looped back. To be more certain of the loopback condition a new Magic number can be selected and tried. Once a magic number has been selected for a connection, that same Magic Number must be used for the whole connection.

During a connection, loopback may be detected by sending an Echo Request packet. If an Echo Request with the devices own magic is received, it is very highly probable that the line is looped back.

VI. Connection Sequence

The connection will be in the down state until all necessary conditions for the physical circuit are met (control signals up, etc.), at which time it will go into the connecting state. While waiting for a connection, a device will send Configuration packets at regular intervals and wait to receive a Config Ack or a Configuration packet from the other side. A connection has been established when both sides have sent Configuration packets and have received Config Acks for those packets.

At this point both sides have also established their Magic Numbers for the connection. If, while connected, a device receives a PCP that contains a Magic Number that does not belong to the device at the other end of the connection, the connection is broken. Other reasons for breaking the connection may be: receipt of a Termination Request, poor line quality as determined by Up/Down Control, or loss of interface control signals. When possible, send a Terminate Request packet and wait for a Terminate Ack before breaking the connection. (State diagram needs to be done)

VII. Up/Down Control

Facilities for Up/Down Control are provided through the Echo Request and Echo Reply packets. With echos the line quality can be determined (packet loss). Once the line quality has degraded to a certain point the connection can be broken or simply not used as a route. For poor quality switched circuits, disconnection and establishing a new connection may solve the problem. For poor quality fixed circuits, the best option may be to stop using the circuit as a route, but keep the connection open and continue testing it. It is advisable to have some hysteresis in the Up/Down control to prevent rapid route flapping.

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: Point–Point Protocol
Chairperson: Drew Perkins

| Attendees: | E-Mail Address: |
|---|---|
| 1. Ron Broersma | ron@nosc.mil |
| 2. Jeffrey Burgan | jeff@twg.com |
| 3. Dino Farinacci | dino@bridge2.3com.com |
| 4. Craig Fox | foxcj@network.com |
| 5. Russ Hobby | rdhobby@ucdavis.edu |
| 6. David Kaufman | dek@proteon.com |
| 7. Mark Lottor | mkl@sri-nic.arpa |
| 8. Milo Medin | medin@nsipo.nasa.gov |
| 9. Bill Melohn | melohn@sun.com |
| 10. Drew Perkins | ddp@andrew.cmu.edu |
| 11. Greg Satz | satz@cisco.com |
| 12. Steve Waldbusser | waldbusser@andrew.cmu.edu |
| 13. Asher Waldfogel | ...linus!wellflt!awaldfog |

# PROTOCOL 33 - PPP LINK CONTROL

| VERSION | MAGIC # | TYPE | OPTIONS |
|---|---|---|---|
| 8 BITS | 32 BITS | 8 BITS | * |

VERSION - PPP LINK CONTROL VERSION

MAGIC # - UNIQUE INTERFACE NUMBER

TYPE

| | |
|---|---|
| 1 - CONFIGURATION | 7 - VERSION REJECT |
| 2 - CONFIG ACK | 8 - TYPE REJECT |
| 3 - CONFIG HUH? | 9 - TERMINATE REQUEST |
| 4 - CONFIG NAK | 10 - TERMINATE ACK |
| 5 - ECHO REQUEST | 11 - NOP |
| 6 - ECHO REPLY | |

OPTIONS - MULTIPLE ENTRIES ALLOWED

| LENGTH | TYPE | DATA |
|---|---|---|
| 16 BITS | 16 BITS | * |

| TYPE | LENGTH |
|---|---|
| 1 - MAXIMUM RECEIVE UNIT | 16 BITS |
| 2 - ASYNC CONTROL CHARACTER MAP | 32 BITS |

YET TO BE CONSIDERED

ADDRESS EXCHANGE; COMPRESSION; ENCRIPTION; AUTHENTICATION

# POINT-TO-POINT PROTOCOL (PPP)

## PHYSICAL LAYER

MOST COMMON INTERFACES

RS 232, RS 422, RS423, V.35, ...

USE OF CONTROL SIGNALS ADDS FUNCTIONALITY

## DATA LINK LAYER

SYNC & ASYNC FRAME STRUCTURE

ISO 3309 [HDLC] (ADDENDUM 1 DEFINES ASYNC)

| FLAG | ADDRESS | CONTROL | INFORMATION | FCS | FLAG |
|---|---|---|---|---|---|
| 01111110 | 8 BITS | 8 BITS | * | 16 BITS | 01111110 |

| PROTOCOL | DATA |
|---|---|
| 8 OR 16 BITS (ISO EXTENSION) | * |

SYNC ALWAYS 16 BITS

PROTOCOL NUMBERS START AT 33 TO AVOID CONTROL CHARACTERS

33 - PPP LINK CONTROL

35 - DOD IP

37 - ISO CLNS

⋮

# ST and CO-IP

Claudio Topolcic
BBN Communications

January 1989 IETF Meeting
ST and COIP Working Group Summary
Reported by Claudio Topolic (BBN)

The working group held two meetings, which correspond to the two tracks we are pursuing. The meeting held on Tuesday 17 January covered the high level and long term issues of connection oriented internet protocols. The meeting held on Wednesday 18 Jan covered the short term need to finalize the ST specification.

17 January 1989

In discussing connection oriented internet protocols, we identified a number of issues. We decided to write them up in a paper. After discussion, we felt this paper would be more useful than the "requirements document" we had planned to write when we met at the last meeting. This is because this paper is more focused and is intended to to describe the problem we are addressing. The issues we discussed are the following:

1. Our bias. We don't pretend to solve all the problems in the Internet. We feel that there are applications that are not supported ideally by IP as it now stands. We further feel that simple extensions to IP may not support those applications well enough. We feel a new approach is worth investigating, and we want to do so without our thinking being constrained by the current IP. Since we are looking at connection oriented protocols, we tried to determine what we really mean by "connection". A definition of was proposed for discussion. It states that in a connection,

   [gateways have per-"flow" state, and can recognize the data packets in a "flow" (using Dave Clark's definition of "flow"), and b) the end systems send a "connection setup" message before sending data packets.]

   It was not agreed to. The need for an explicit connection setup message was specifically disagreed with by some. Nevertheless, it serves as a good starting point. It was agreed, however, that a connection does not imply the guaranteed delivery and ordering of a circuit, and does not imply the reserved bandwidth of today's ST implementation.

2. The applications. Ultimately, this protocol attempts to support certain applications better than datagram oriented protocols. For the purposes of this discussion, we chose to ignore the distinction among the transport layer and above, and simply called the aggregate an application. We are very interested in the requirements of applications, and how those

requirements are expressed and transferred to the internet layer. We assume that the requirements apply to a connection, not individual packets. These requirements are simply Quality Of Service (QOS) parameters. Some QOS parameters that may be relevant are:

- Some measure of the distribution of the offered load, such as the peak and average bandwidth and a burst factor.

- The required delay distribution, that may include average and some form of worst case.

- The reliability requirements, possibly expressed in terms of m out of n packets that can be lost.

This transfer of information cannot be only one-way. The internet layer must give the application layer some information as well. Minimally, the internet layer must say if it can or cannot deliver the service specified by the QOS. Further, it can supply hints to the application. We identified two forms that these hints can take. If the internet layer cannot meet the requested QOS, it may nevertheless suggest modifications to the QOS that it can meet. Alternately, the application may query the internet layer to determine what QOS is likely to be feasible before attempting to establish a connection. In either case, the application may modify its behavior to use the service that is available. This information is only a hint since the application may not be able to use it, and the internet layer may not be able to support a subsequent connection since the state of the Internet may have changed.

Finally, the control information between the application and the internet layer must also support asynchronous changes in state. For example, a network or a gateway may fail, or a connection may be pre-empted by a higher priority connection. As a result, an existing application's QOS may no longer be supported. The application and the internet may have to renegotiate that connection with a different QOS.

3. The underlying networks. To do its job, the internet layer must map the QOS requested by the applications it supports into the services it can obtain from the networks on which it resides. Ideally, the protocol should operate across any network. So we must answer two questions: "What services will such a protocol need from the underlying networks?" and "What services can it expect to get?" It is hard to answer the first

question at this point because we have to look at the QOS parameters first. The second question is hard to answer since we don't know what networks we may want to operate over. Nevertheless, we feel the protocol should operate over many different kinds of networks. We decided that the internet layer should operate it terms of the QOS parameters, and that there may need to be an intermediate layer between the internet and network to supply the services that the internet layer needs but the network does not provide.

4. Gateway functions. We made a first pass at identifying the functions that should be performed by the gateways in support of this internet protocol. Some of these functions deal with managing the per-connection state information that the gateways are expected to maintain. This is the list:

- Resource management. The resources that must be managed are:

  - network bandwidth
  - gateway processing power
  - gateway buffers

This function differs from datagram based protocols in that in this protocol we expect unequal sharing of the resources. An established connection will preferentially receive resources before one that is not yet established. Resource management does not imply that specific resources are assigned to specific connections. That is merely one resource management technique, and we intend to explore others.

- Routing. We feel that the gateways will be involved in making routing decisions. Furthermore, the routing decisions may have to take into account the currently available resources, not only within the gateway or on the locally attached networks, but in remote networks and gateways as well.

- Connection setup. However a connection is established, the gateways must be involved so that they can obtain the appropriate per-connection state. The connection setup will eliminate the need for later sending redirects or fragmenting packets.

- Reacting to asynchronous state changes. State changes such as the failure of a gateway, network or host, or the pre-emption of an established connection by a higher priority connection, cause the gateway's state information to become obsolete. The gateways will have to be involved in correcting the obsolete state and also in reconstructing the connection.

Our short term plan is to discuss these ideas by exchanging electronic mail, and then to assign writing tasks and to write a short paper. We also identified that we should keep in contact with other IETF working groups and IAB task forces. Specifically the Open Routing Working Group and the Autonomous Systems Task Force.

18 January 1989

We made a lot of progress in agreeing to changes to the ST specification. We discussed the goals we are trying to achieve our proposed changes. We began this discussion with video conferences and agreed to the goals at this meeting. Specific proposals and alternatives will be measured against how they support these goals. The goals are as follows:

- ST must support current and planned activities that use ST. For example multimedia conferencing.

- ST must be practical to implement. It must be reasonably simple. The specification must be complete. Areas that have not been completely thought out should specify some well defined alternative.

- ST must be usable in different environments. Specifically, it should be possible to use ST across IP-only systems, at least in the case when resources are plentiful and don't need to be managed.

- ST must support research into connection oriented communication. It should be flexible and extensible. It must support performing experiments suggested by the CIP working group or other sources.

- ST should be able to support applications different from those that currently use it. For example, it may be interesting to see if there is any advantage to running TCP above ST.

We discussed the changes that had been proposed in our draft ST specification.

Once we identified the underlying changes, and we all understood them we generally achieved consensus. We did not discuss the mechanisms that would be required to support these changes, though a number had been discussed previously. The specific changes we agreed to are the following:

1. The Access Controller (AC) will be eliminated. The reason is primarily that it does not support real access control very well, it ties ST and the next higher protocol layer too closely together, and it adds unnecessary complexity into the protocol. However, since the AC provided a number of useful functions, those functions will have to be supported by some alternate mechanisms. Specifically, those functions include:

   - Access control. Access control can only be performed properly by the end systems.

   - Source of conference information. Information that is specific to a conference or is conference-wide will no longer be available from the AC. ST will depend on the next higher protocol layer for some of this information. Other such information will be passed from one ST agent to another within the connection control packets.

   - Unique connection IDs. The AC will no longer be the source of globally unique connection IDs. ST will require some form of unique IDs. ST will use only hop-by-hop IDs or possibly network-locally unique IDs.

2. The Forwarding Bit Map will be eliminated. This feature did not perform any function that we felt was necessary or useful, and added unnecessary complexity into the protocol. ST forwarding will be based on the ID.

3. The structure of an "ST stream" will not be "omniplex", but will be a simplex tree. The omniplex connections are too omplex to manage. Simplex trees are much simpler. Each participant in a conference will create a tree that is routed at itself, and that sends packets to all the receivers.

4. The Envelope header will be eliminated. As implemented, the envelope does not support aggregation as it was intended, the type of aggregation it supports is of questionably utility and we can think of different mechanisms to support aggregation better. However, we have not yet thought those mechanisms through in detail, so this subject needs further work.

5. The connection setup message will contain a "conference name" as well as a "connection name". This is necessary since certain resource management strategies, such as "multi-destination half-duplex" need to identify different ST connections that correspond to the same conference.

   At this point we need to identify the mechanisms that will support all these changes. We will do this by exchanging electronic mail and by using multi media conferences. As a result, we will write a new ST specification.

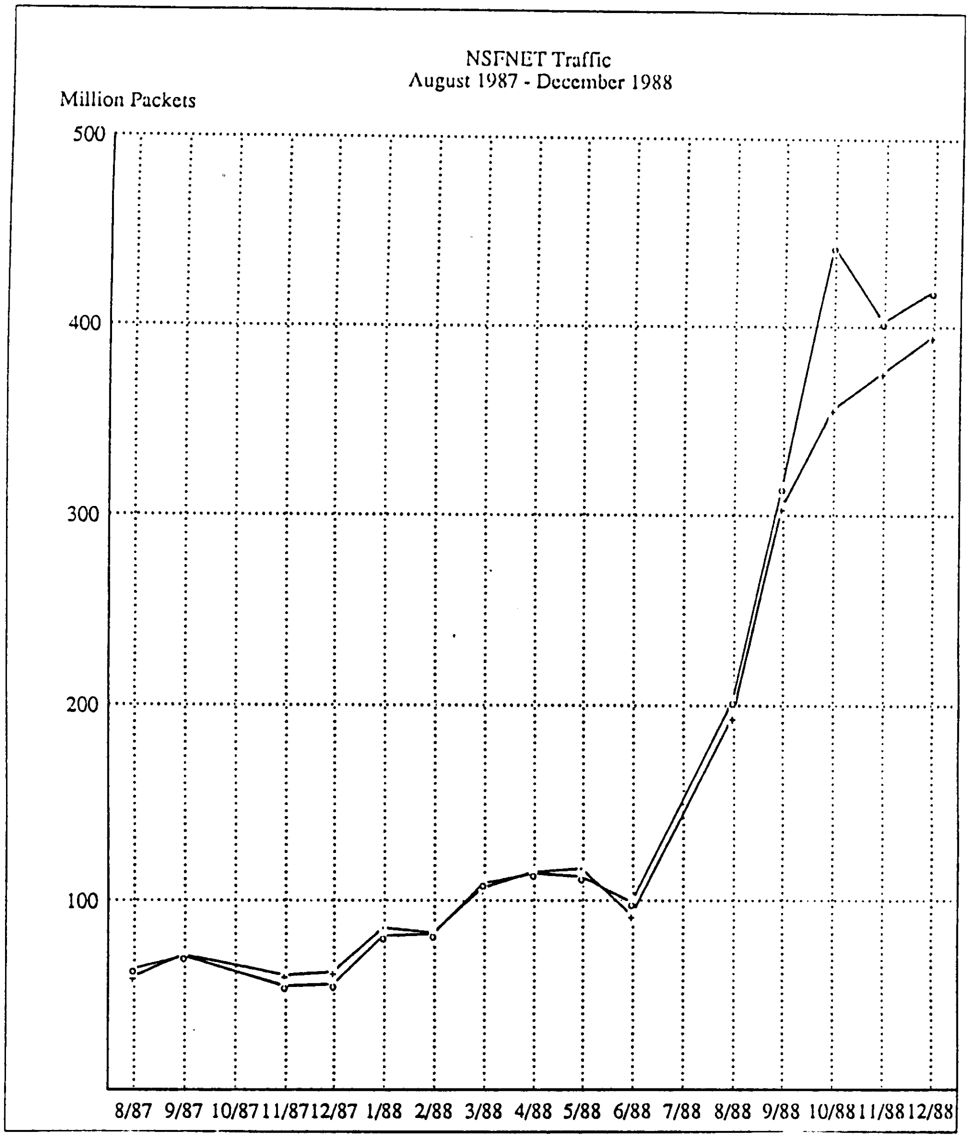


HIGHLIGHTS:

NSFnet Traffic Statistics

Dark line - Packets in
Light line - Packets out

Oct 87 - NSFnet Phase I Traffic

Aug 88 - Introduction of NSFnet Phase II (400 million packets total for the month)

Sept 88 - Dec 88

Traffic gradually climbing. Dip in Oct statistic a result of problem with aberrant name server in Hawaii.

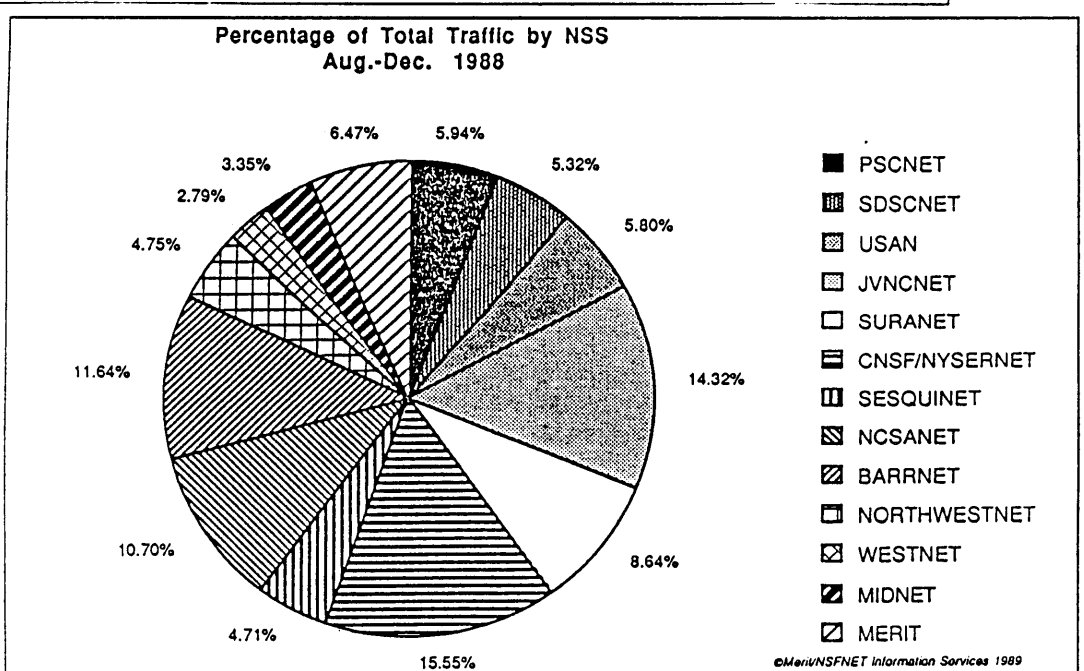


HIGHLIGHTS: Multi-colored pie chart reflects percentage of the total traffic during August through December by NSS site. This is essentially the marking period since the beginning of the NSFnet.

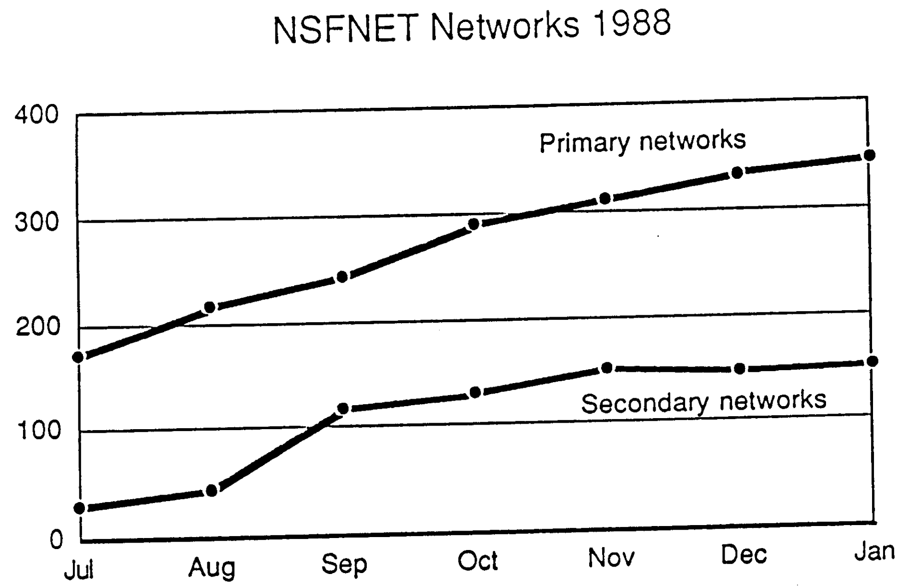


HIGHLIGHTS: The NSFnet is a policy-based system. It will only let authorized traffic pass through the backbone. This chart represents "configured" networks in the NSFnet. At the beginning of July about 180-190 nets were configured, with growth to 360-370 networks allowed by the first of January 89.

The top line represents networks that have a "primary" pathway to the NSFnet. The lower line represents networks with a "secondary" or backup pathway to the NSFnet. As you can see, about half, 150 or so, have a secondary path.

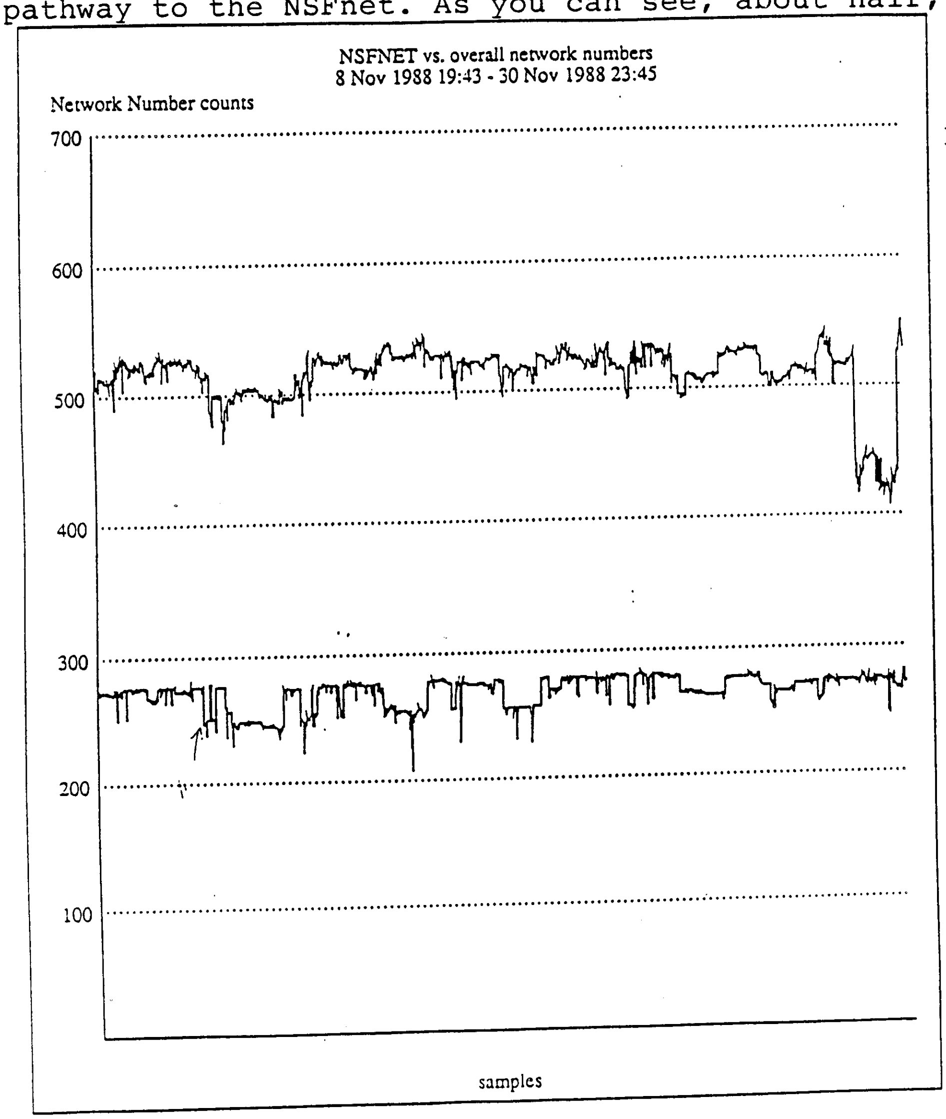


HIGHLIGHTS:

Samples - 15 minute snap shots of routing tables.

Network Number Count - Counts of networks which are announced (not hosts).

Bottom graph display - Networks directly connected to the NSFnet and networks NSFnet "hears" from the Arpanet.

Dips are related to different problems: power outage, loss of circuits, etc.

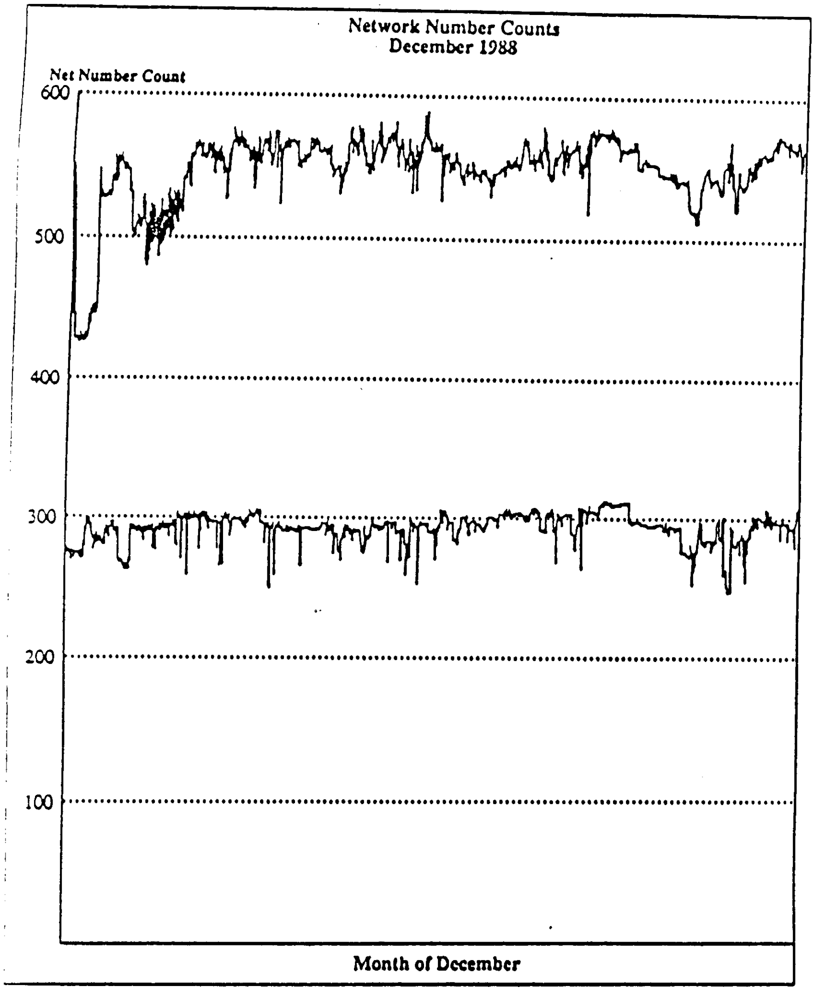


HIGHLIGHTS: Throughout December the total count tends to grow quite significantly. At about the 570 Net Number Count was the time at which deployment of the butterflies was needed. That level was exceeded before the butterflies were deployed but they are now in place and operating.

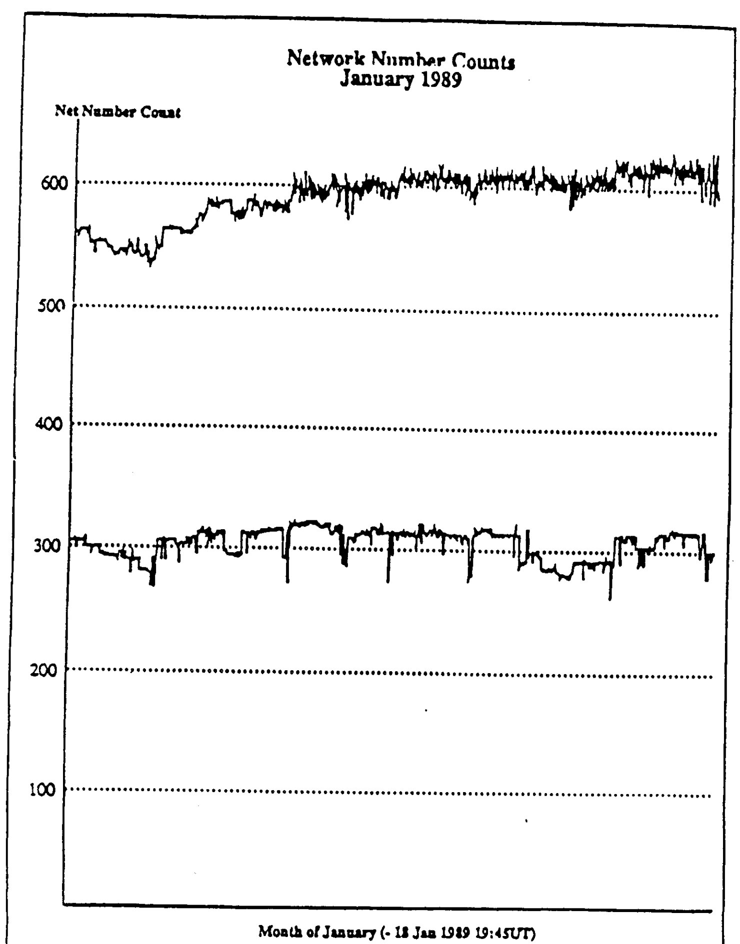

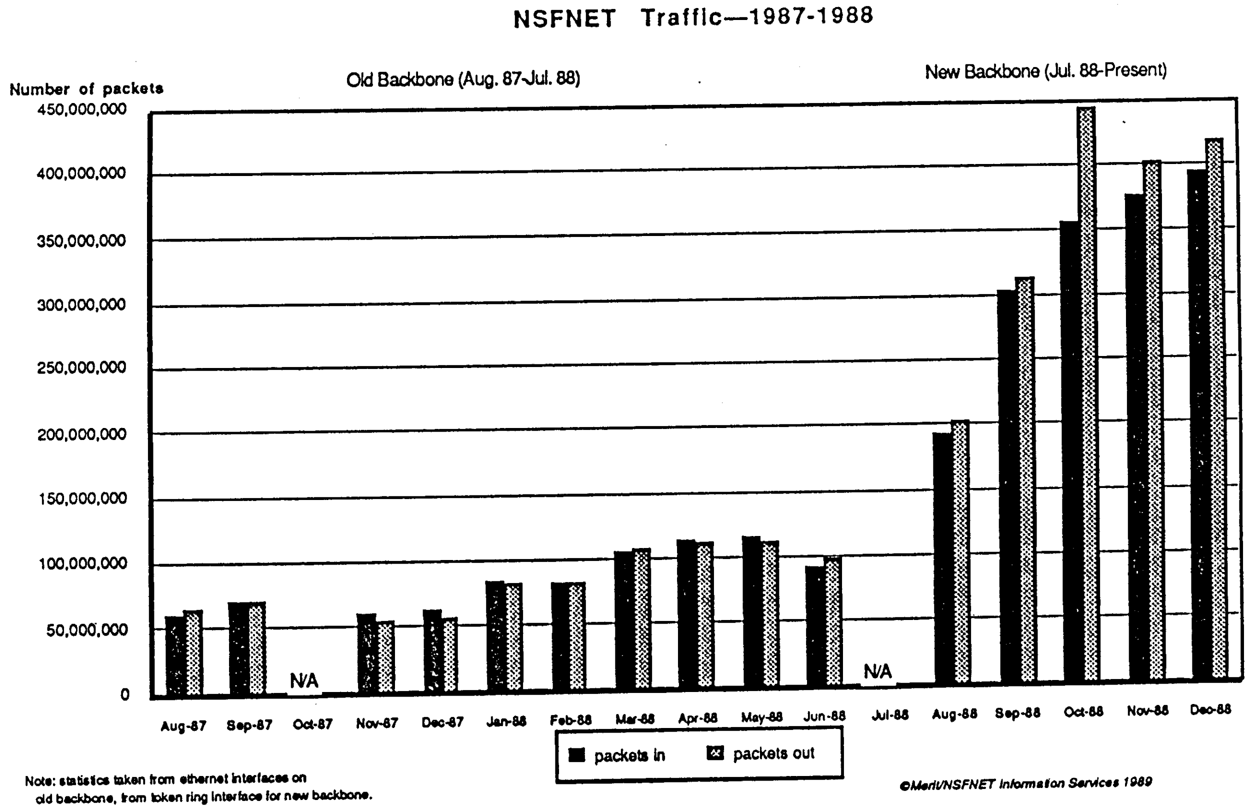
NSFNET Traffic—1987-1988
Number of packets
Old Backbone (Aug. 87-Jul. 88)
New Backbone (Jul. 88-Present)
450,000,000
400,000,000
350,000,000
300,000,000
250,000,000
200,000,000
150,000,000
100,000,000
50,000,000
0
N/A
N/A
Aug-87
Sep-87
Oct-87
Nov-87
Dec-87
Jan-88
Feb-88
Mar-88
Apr-88
May-88
Jun-88
Jul-88
Aug-88
Sep-88
Oct-88
Nov-88
Dec-88
packets in
packets out
Note: statistics taken from ethernet interfaces on old backbone, from token ring interface for new backbone.
©Merit/NSFNET Information Services 1989

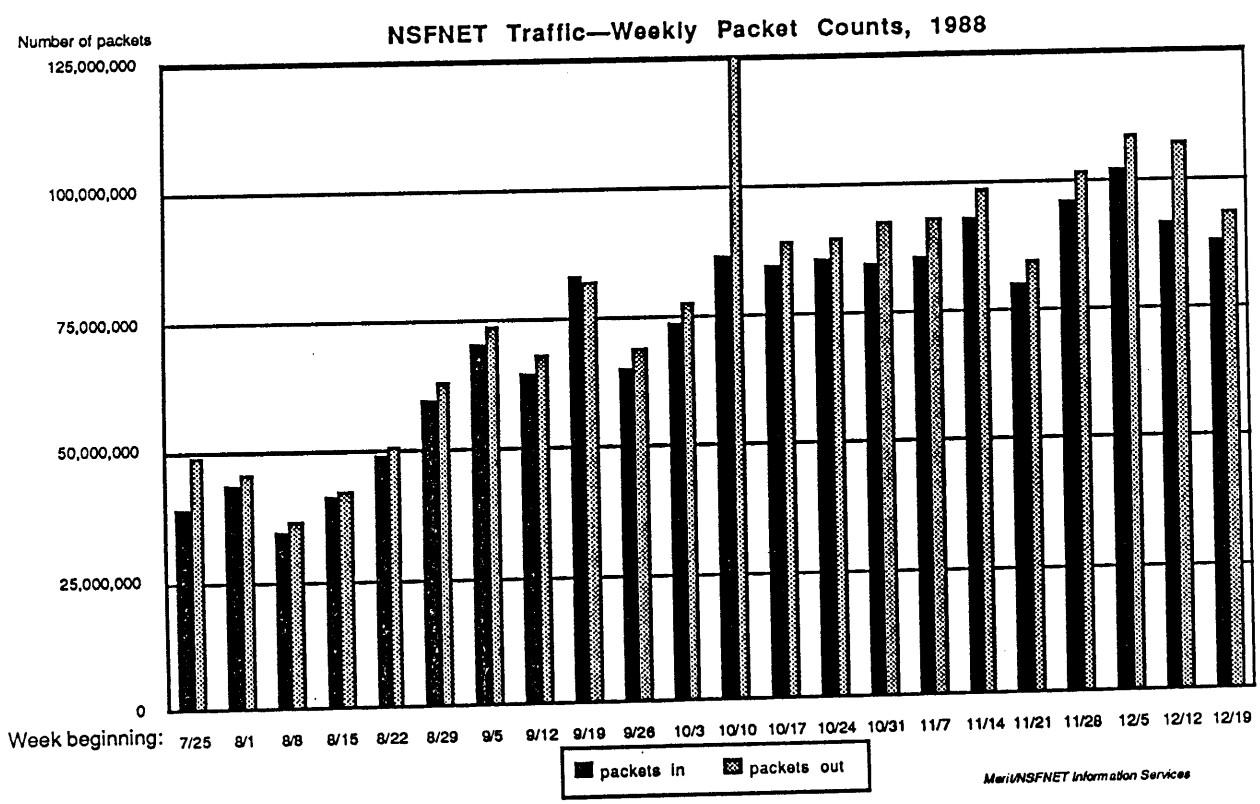
NSFNET Traffic—Weekly Packet Counts, 1988
Number of packets
125,000,000
100,000,000
75,000,000
50,000,000
25,000,000
0
Week beginning:
7/25
8/1
8/8
8/15
8/22
8/29
9/5
9/12
9/19
9/26
10/3
10/10
10/17
10/24
10/31
11/7
11/14
11/21
11/28
12/5
12/12
12/19
packets in
packets out
Merit/NSFNET Information Services

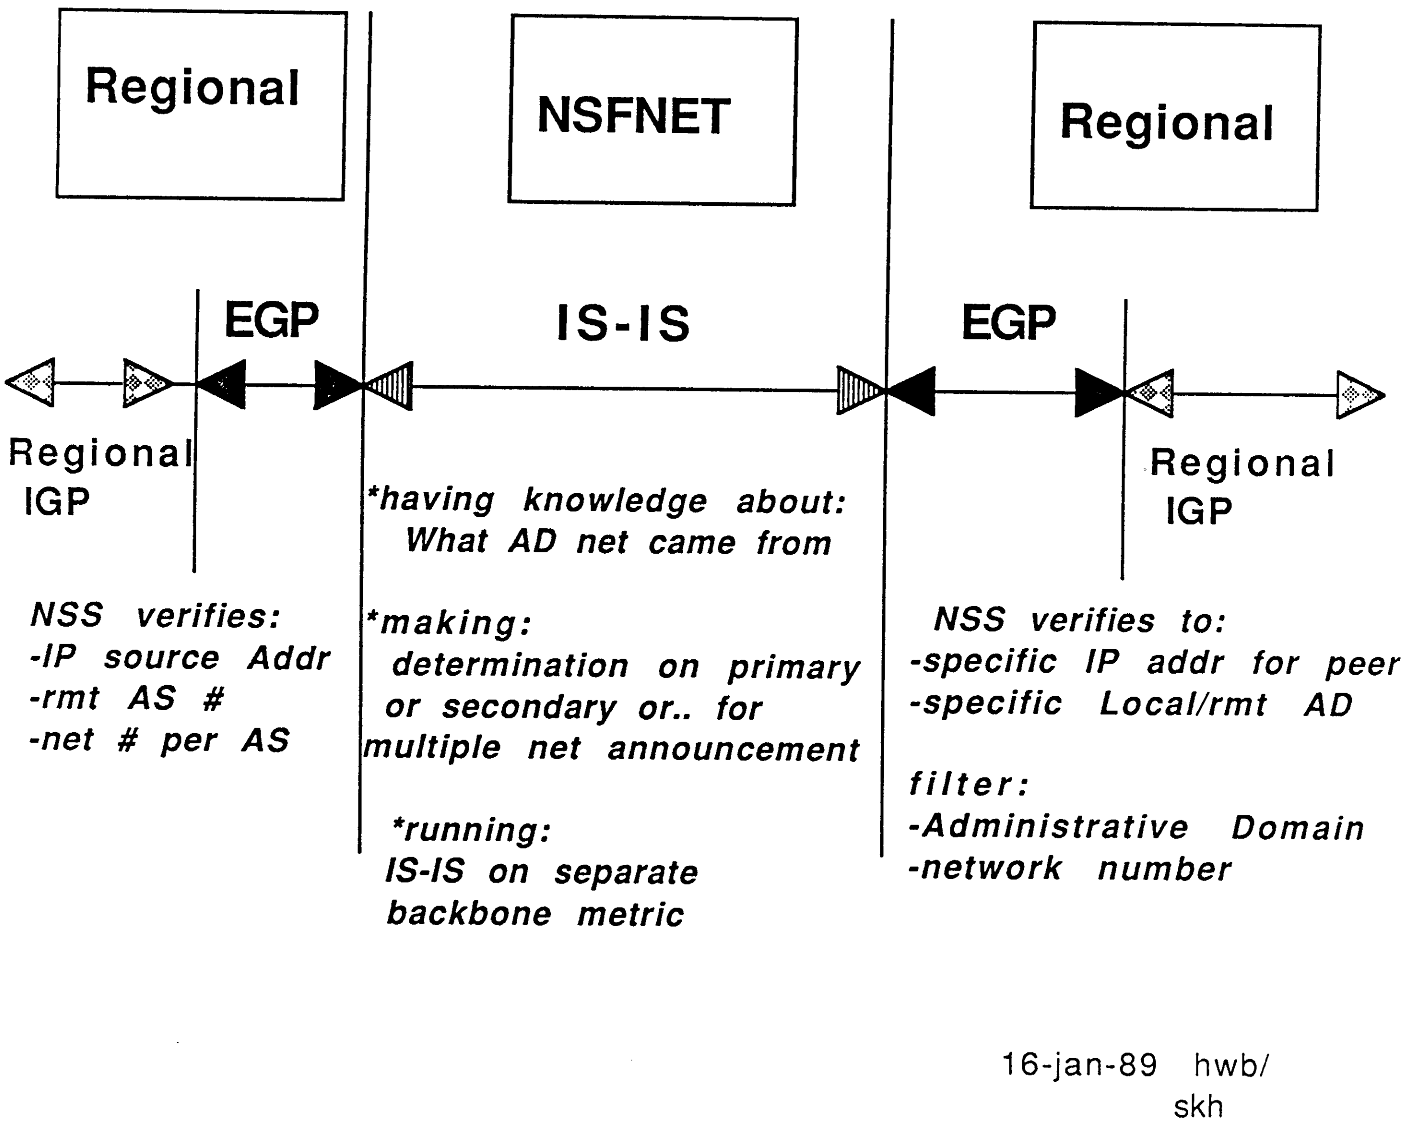


HIGHLIGHTS:

Q: A detailed question about the IS to IS and its delay metrics; are they dynamically computed or are they static numbers?

A: Static numbers, right now, based on statistics taken every 15 minutes. Static but configurable; tuned every few weeks based on statistics collected on a 15 minute basis.

Q: Given the filtering you have mentioned is actually on the EGP route and not on a packet-by-packet basis on the routers/gateways, what mechanism is there to prevent packets from being source routed through some network that is allowed to send to some destination?

A: There is a certain amount of packet filtering that can be done on the NSSs on a packet-by-packet basis. Most of that capability is not enabled at this time.

# Internet Report

Zbigniew Opalka
BBN Communications

January 1989 IETF Meeting
Internet Report Briefing
Presented by Zbigniew Opalka

This discusses the significant events in the Internet in the time frame of late 1988 and early 1989. The most significant event was the introduction of the Butterfly as the replacement for the aging LSI-11 as the base of the Mailbridge.

In the LSI-11 Core system there were 7 Mailbridges used to connect the Milnet to the ARPANET and 6 EGP servers, 3 on the ARPANET and 3 on the MILNET, making a total of 13 LSI-11 gateways. In new Butterfly-based Core system, there is a total of 6 Butterfly Mailbridges serving both as the "Mailbridges" and as the EGP servers.

One of the major reasons for the change was the significant increase of the number of nets in the Internet. The LSI-11's were no longer able to accommodate, because of memory limitations, the large number of networks. The LSI-11 EGP servers were eventually limited to only 620 networks in their Routing Table, while the LSI-11 Mailbridges were able to handle only 570 networks. This inconsistency was adversely impacting the operation of the Internet.

The Butterfly Mailbridges, with significantly more memory, provided the solution. In the current configuration they have enough room for approximately 1,000 networks and can be easily reconfigured for more when the time arises. The Butterfly has other advantages over the LSI-11, it is composed of 3 MC68000 processors with a total of 3 Megabytes of memory. The Butterfly also adds improved versions of "load sharing" and "access control" to its functionality.

The Butterfly Mailbridge can maintain several "access control" tables in its memory. The access control parameters based in the tables can be switched on demand to provide the required amount of access control. Also hosts and gateways are "homed" to appropriate Mailbridges to provide fairness in spreading the traffic across the six Mailbridges. Gateways are "told" to use the appropriate Mailbridge through EGP exchanges with the Mailbridges, hosts are directed to their Mailbridge through the use of the ICMP Redirect message. Presumably the hosts were configured properly based on information supplied in the DDN Management Bulletin.

The six Mailbridges are:

| GW # | Site Name | Addresses | |
|---|---|---|---|
| 1 | BMILAMES | 10.2.0.8<br>26.20.0.16 | west coast |
| 2 | BMILBBN | 10.3.0.5<br>26.1.0.49 | east coast |
| 3 | BMILDCEC | 10.6.0.20<br>26.21.0.104 | east coast |
| 4 | BMILISI | 10.6.0.22<br>26.6.0.103 | west coast |
| 5 | BMILLBL | 10.2.0.68<br>26.20.0.34 | west coast |
| 6 | BMILMTR | 10.3.0.111<br>26.20.0.17 | east coast |

The Mailbridges were Beta tested in December. In late December and early January two DDN Management Bulletins (49 and 53) were sent. These Bulletins specified the default Mailbridges for hosts and new EGP neighbors for gateways to acquire, respectively.

During a "grace" period both systems were up and running, though the LSI-11 Mailbridges were no longer serving as conduits for the ARPANET and MILNET. Rather, they were receiving traffic and forwarding it to the LSI-11 EGP servers who would forward it to the Butterflies, while at the same time sending ICMP Redirects when appropriate.

The last part of this transition is to decommission both the LSI-11 Mailbridges and EGP servers. This was done on 6 March 1989 as announced by DDN Management Bulletin 55. There were several problems encountered during this transition. The change of Autonomous System number for the Core, changed from 1 to 164, caused gateways in some AS which were higher in the LSI case to be lower in the Butterfly. This caused problems as they used the "less than" relationship to determine "passive-active" EGP peer relationship.

The most prevalent problem encountered was the growth of the routing updates. The size of the updates increased to over 2

Kbytes causing many UNIX systems to improperly handle EGP updates. Once the maximum IP packet size was increased to 4kbytes, the systems were able to reassemble the EGP updates properly.

The Butterfly Mailbridges also made "reasonableness" checks that were determined to be not proper. This was based partly on EGP specifications that were not firmly defined.

The Mailbridges are currently passing over 7 million packets per day and dropping on the average less than .01% of the packets received.

17 January 1989

# STATE OF THE INTERNET

Zbigniew Opalka

BBN Communications Corporation

17 January 1989

## CURRENT INTERNET CORE

- 7 LSI-11 Mailbridges
- 6 LSI-11 Servers
  - 3 on the ARPANET
  - 3 on the MILNET
- More than 630 Nets Connected to the Core

BBN Communications Corporation

HIGHLIGHTS: The old internet core consists of 7 LSI-11 Mail Bridges and 6 LSI-11 EGP servers, with 630 networks connected, as seen through EGP. Some instability is going away with the new mailbridges because the butterflies are configured to handle in the neighborhood of 1,000 networks, whereas the EGP server is configured to handle 620 networks and the LSI-11 mailbridges handle only 570 networks.

# LSI LIMITATIONS

- Memory
  - Routing Table Holds only 620 Nets
- Single Processor
- Extra Hop Problem
- GGP Overhead Consumes Excessive Internet Bandwidth





# NEW INTERNET CORE

- 6 Butterfly Mailbridges
  - Mailbridge Function
  - EGP Server Function



HIGHLIGHTS: The new butterfly mailbridges serve both as mailbridges and EGP servers, which eliminates the extra hop problem.

# MAILBRIDGE TRANSITION SCHEDULE

| | |
|---|---|
| * December 8, 1988 | Start of Beta Test for Butterfly Mailbridges |
| December 23, 1988 | DDN Management Bulletin #49 Announcing Host Rehomings (New Default Gateways for Hosts) |
| January 12, 1989 | DDN Management Bulletin #53 Announcing Gateway Rehomings (New EGP Neighbors for Gateways to Acquire) |
| January 19, 1989 | LSI Mailbridges to Issues Redirect-Host Messages Instead of Forwarding all Traffic |
| February 1, 1989 | LSI Mailbridges to be Decommissioned |
| Mid February, 1989 | LSI GEGPSerer to be Decommissioned |





# BUTTERFLY MAILBRIDGE ADDRESSES

| GW# | SITE NAME | ADDRESSES |
|---|---|---|
| 1 | BMILAMES | 10.2.0.8 |
| | | 26.20.0.16 |
| 2 | BMILBBN | 10.3.0.5 |
| | | 26.1.0.49 |
| 3 | BMILDCEC | 10.6.0.20 |
| | | 26.21.0.104 |
| 4 | BMILISI | 10.6.0.22 |
| | | 26.6.0.103 |
| 5 | BMILLBL | 10.2.0.68 |
| | | 26.20.0.34 |
| 6 | BMILMTR | 10.3.0.111 |
| | | 26.20.0.17 |



*HIGHLIGHTS: In early December BBN started soliciting beta testing for EGP and, among others, Dave Mills and Bob Enger started pinging on EGP code. As more people became involved, DDN Management Bulletins announced host rehomings, then gateway rehomings.

# BUTTERFLY MAILBRIDGE FEATURES

- Parallel Processing
  - 3 Processors (MC68000)
- 3 Megabytes of Memory Total
- SPF Replaces GGP
- Improved Load Sharing Functionality
- Improved Access Control Capability



HIGHLIGHTS: The installed butterfly mailbridges are not the newest 68030 or 40 but are definitely better than the LSI-11s, and have three megs of memory. With respect to the improved load sharing functionality, there is a static load sharing table that identifies PSNs. For hosts, hosts are sent redirects to appropriate mailbridges. For gateways, load sharing information is sent to them via EGP. So, a gateway talking to a mailbridge will be redirected (using EGP) to the appropriate mailbridge. With the improved access control capability, the access control structure is a tree structure that can be walked down, permits access control to be turned on and off, and permits several different trees and several different policies to reside in a mailbridge simuntaneously. Access control is based on source-destination pairs, community of interest, protocols, user-server pairs, and input interface; any of this information can exist in a tree structure and determines whether or not a packet goes through. BBN does not determine the access control policy; DCA does. The access control methods are documented in a BBN report, which will be made into a public release document by DCA and distributed to this community through Bob Braden. An RFC(s) will probably be pursued and will address both load sharing and access control.

17 January 1989

# OPERATIONAL NETWORKS

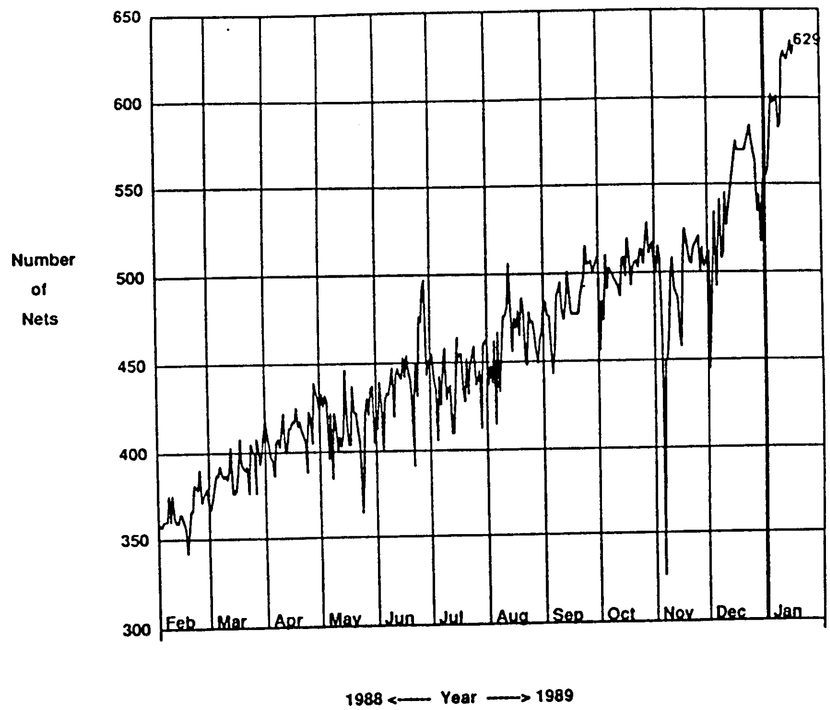


BBN Communications Corporation

17 January 1989

# OPERATIONAL NETWORKS LONG TERM2

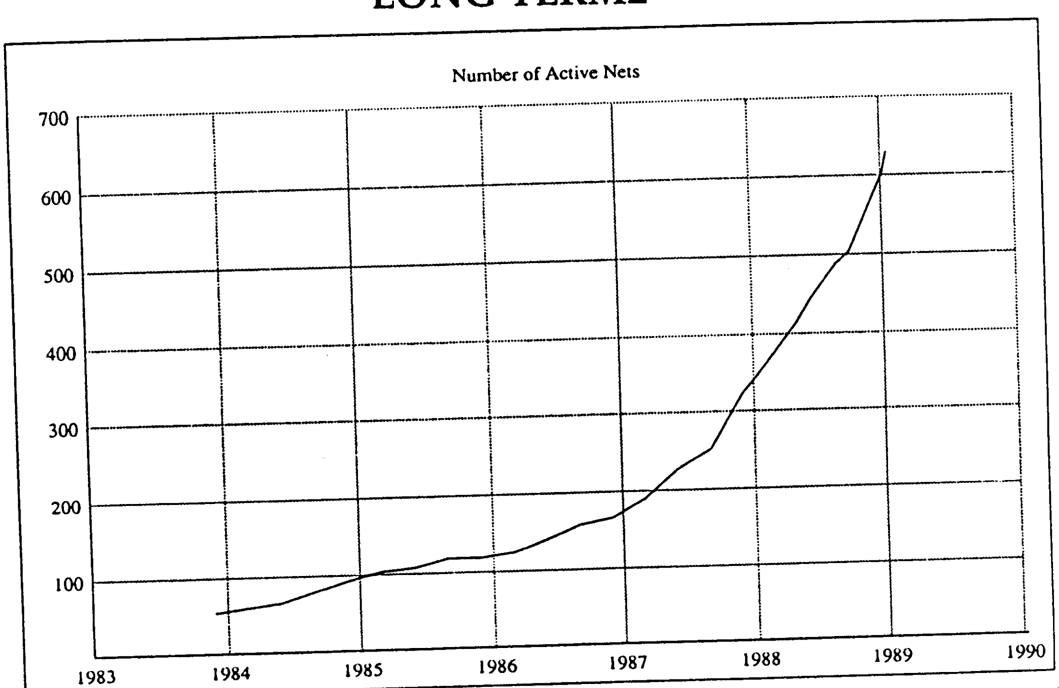


BBN Communications Corporation

HIGHLIGHTS: The butterflies are currently configured for just under 1,000 networks. It is purely a matter of a configuration parameter; they can be configured for numbers higher than that.

# PROBLEMS ENCOUNTERED AND RESOLVED

- Max IP Packet Size should be Increased to 4K, otherwise EGP NR Updates may be discarded by IP Reassembly Code
- EGP must Support the same Large Packet Size for NR Routing Udates (Gated, Kirton Code, Etc.)
- Butterfly Check for Reasonable EGP Parameters to be Changed



HIGHLIGHTS:

One of the major problems encountered was when Max IP packet sizes went over 2K; due to fragmentation, people had to reassemble IP packets bigger than 2K. The solution is to make the buffers capable of handling at least 4K worth of data.

The other problem was that the butterfly had an "unhealthy" sanity check for some EGP parameters. BBN had an instance where a gateway was talking to a butterfly, telling the butterfly it was willing to accept EGP "hellos" every five seconds. The butterfly considered that unreasonable and said "we are not going to be neighbors". What happened was the code checked the sanity check first and didn't even get to the point of determining "passive" or "active" mode. BBN will release new disks that will bring the sanity check down to one. This should solve the problem.

If you encounter EGP problems, you can send inquiries to NB-Transition@BBN.COM.

# DOE ESNET Report

Tony Hain

Lawrence Livermore National Lab

January 1989 IETF Meeting
ESnet Briefing
Presented by Tony Hain (LLNL)

ESnet is a communications network for use by researchers funded through the DoE Office of Scientific Computing. A limited release configuration, 8 Micro-Vax 2's as IP gateways, has been in operation since August '88. These gateways are connected by 56 Kbps lines between LLNL, GA, ANL, LBL, ANL, PPL, FSU, and UTA. A limited number of hosts at each site have been provided with the NSP/IP suite for access to the NMFECC Crays. ESnet employs a proprietary network core routing protocol to allow concurrent and equal service to multiple OSI model level 3 protocols. Current plans are to provide service for Internet IP, X.25, and OSI CLNS. The addition of concurrent X.25 routing is projected for June 89.

The project expansion is on hold pending a DoE IG audit report which was due out in November '88. Presuming a favorable outcome of the audit, plans are progressing for T1 service and 17 sites have been identified as candidates for CY '89. Initial T1 service will provide subchannels for the existing MFEnet and private X.25 backbones as well as a 448 Kbps cross country ring.

ESnet policy restricts the use of its resources to that authorized by the DoE/OSC. Much effort has gone into finding a mechanism to implement this policy while maintaining our informal policy of being "good citizens" within the Internet. This combination of issues restricts us from advertising routes to networks where policy says we are unable to carry all of the traffic. Because there are no current solutions to these problems, our site coordinators have agreed to implement static host routes network wide. This step will allow us to move forward for some time without creating "black holes" within the Internet. Some of the routing options explored include multi-homing, host redirects, and address translation. Multi-homing (2 IP addresses for the node) was rejected as unavailable for the single interface BSD workstation community. We rejected host redirects because we could not reliably determine the legitimate gateway to redirect to. Address translation (replacing the source and destination IP addresses) held promise that we could use traditional routing techniques and advertise a virtual network. It fell apart because several applications were reported to pass the IP source as data above the TCP header during session establishment.

The ESnet NOCC is currently under construction at LLNL, and several components are running at the proof of concept stage.

# ESnet Status Jan. '89

- 8 MV II's RUNNING as IP GATEWAYS since Aug. '88

  NSP/IP SUITE IN LIMITED RELEASE

- PROJECT HOLD PENDING DOE-IG AUDIT REPORT

- X.25 ROUTING SUPPORT PROJECTED JUNE '89

- T1 UPGRADE PLANS PROGRESSING

  17 SITES IDENTIFIED FOR T1 IN CY '89

- POLICY & ACCOUNTING DEVELOPMENT CONTINUES

  ESnet POLICY RESTRICTS RESOURCE BY USAGE

- CURRENT ROUTING PLAN:

  STATIC HOST ROUTES EVERYWHERE

- NETWORK OPERATIONS UNDER CONSTRUCTION

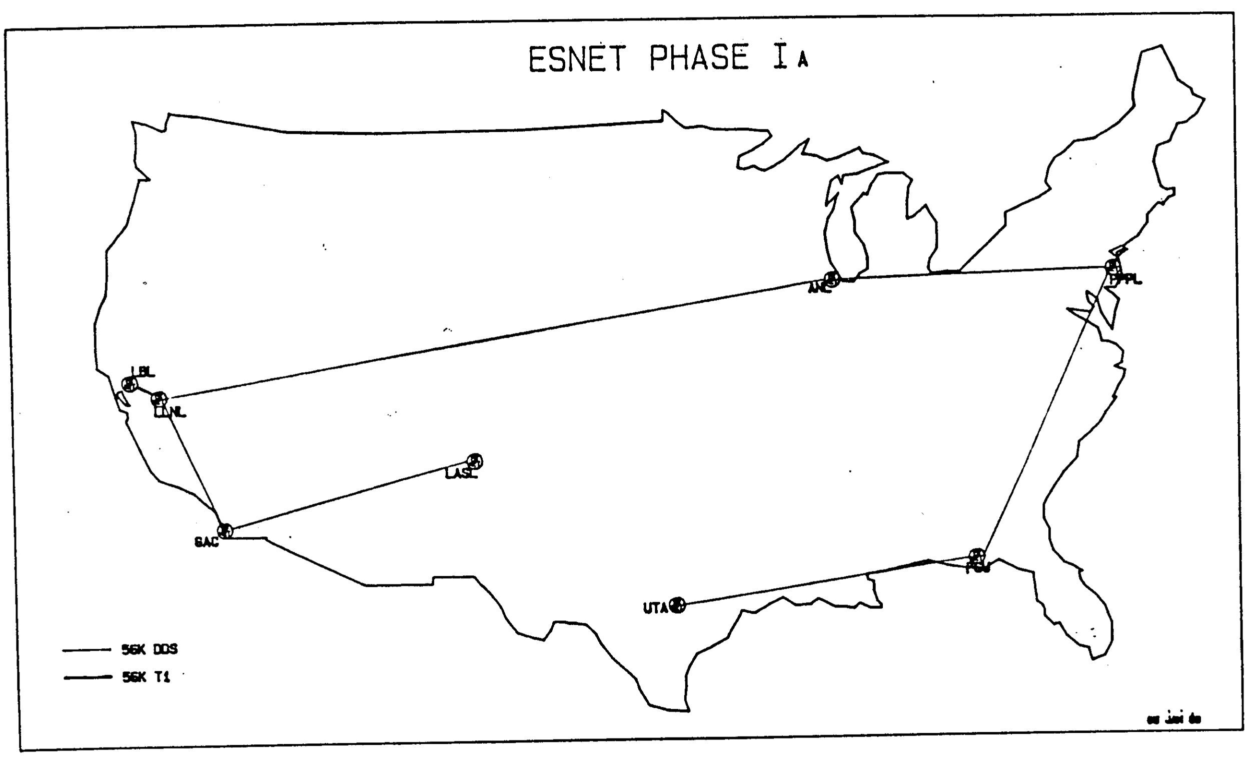

# CSNET Report

Craig Partridge

BBN Systems & Technologies Corporation

January 1989 IETF Meeting
CSNET Briefing
Presented by Craig Partridge

## History of CSNET

CSNET was started with an NSF grant in 1981 and was completely self-supporting by 1985. It originally was run from three sites but was consolidated into a single operation and information center at BBN in 1983.

## CSNET Now

CSNET current has 183 members of which 78 run PhoneNet (our mail-only software package that runs over dial-up phone links) and 105 use one of our four IP services (X.25Net, Cypress, Dial-up IP, Lease Lines). CSNET also provides a wide variety of user services including a 24 hour hotline, domain backup, two newsletters (one on-line), Postmaster services and the Info Server, a program for retrieving documents via e-mail.

The network is currently in the process of merging with BITNET.

## CSNET

- History of CSNET
- CSNET Now
- CSNET Future

## The Computer + Science Network (CSNET)

CSNET Coordination and Information Center (CIC)
at BBN Systems and Technologies Corporation
10 Moulton St
Cambridge MA 02138

cic@sh.cs.net

## History of CSNET

- Original Proposal 1979
- CSNET Funded and Starts Operation 1981
- CSNET Completely Self Supporting 1985

## History of CSNET (continued)

- CIC Originally Three Sites (Rand,UWisc,UDel)
- Consolidated at BBN in 1983

## CSNET Now

- CSNET Statistics
- CSNET Communications Services
- CSNET User Services

## CSNET Statistics

- 183 Members
- 78 PhoneNet, 105 IP

## CSNET Communication Services

- PhoneNet
- Cypress
- X.25Net
- Dial-up IP
- Leased Lines

## CSNET Users Services

- 24 Hour Hotline
- Domain Backup
- Newsletters (CSNET News and CSNET Forum)
- Postmaster Services
- Nameserver
- Info Server

## CSNET Future

- Financially Healthy Despite NSFNET
- Merger With Bitnet (Pending Vote of BITNET Members)

HIGHLIGHTS: CSNET users can send IP datagrams out on the connected Internet and establish telnet connections.

# Domain System Statistics

Mark Lottor
SRI International

January 1989 IETF Meeting
Domain System Statistics Briefing
Presented by Mark Lotter (SRI)

Slides:

1 - networks and gateways
2 - host address counts
3 - random string searches

The first slide is from NIC database statistics. It shows the number of networks and gateways registered with the NIC. As you can see not too many people register gateways. The number of registered and CONNECTED networks is about double that of what is actually in use (see BBN slides for ACTIVE networks).

The second and third slide are from statistics gathered by the domain survey program. The second slide shows the number of hosts that have the corresponding number of host addresses. Zero addresses was probably for an MX-only domain entry. When you get up past 10 or so addresses, it is usually people doing strange non-standard things and not necessarily a single host with that many addresses.

The third slide shows some counts for 'interesting' string searches done on the resulting host list from the survey program. A few interesting things to note are that COM and EDU hosts make up most of the hosts on the Internet. Also note that hosts within HP and SUN aren't actually directly connected to the Internet.

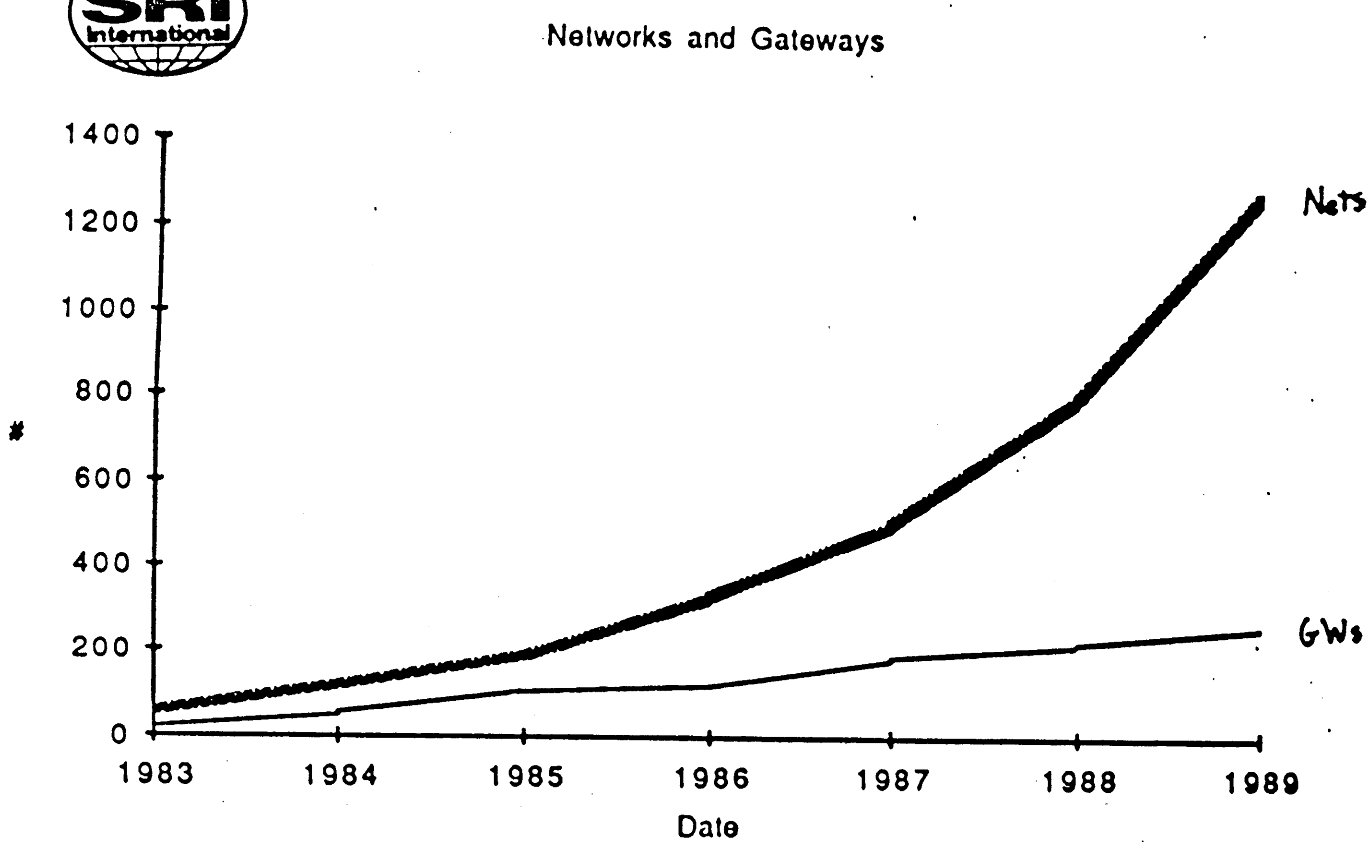

## string searches

| string | Jan 89 | Oct 88 |
|---|---|---|
| total | 85000 | 60000 |
| unix | 16883 | 14200 |
| vax | 6402 | 5500 |
| sun | 21489 | 17000 |
| ibm | 5848 | 4700 |
| gw | 2091 | 2165 |
| edu | 37966 | |
| com | 34034 | |
| mil | 3121 | |
| hp.com | 14049 | |
| sun.com | 8412 | 7500 |
| * | 713 | 550 |

## Host address counts

| addresses | hosts |
|---|---|
| 0 | 4448 |
| 1 | 78299 |
| 2 | 1728 |
| 3 | 248 |
| 4 | 83 |
| 5 | 30 |
| 6 | 15 |
| 7 | 1 |
| 8 | 7 |
| 9 | 3 |
| 10 | 1 |
| 11 | 2 |
| 12 | 1 |
| 15 | 1 |
| 22 | 1 |
| 25 | 1 |
| 30 | 2 |

# Support for OSI Protocols in 4.4 BSD

Rob Hagens
University of Wisconsin

January 1989 IETF Meeting
OSI Protocol Support in BSD 4.4 Briefing
Presented by Robert Hagens (UWis)

In 1988, the National Institute for Science and Technology (NIST) funded the OSI-POSIX project. The goal of this project is to provide an openly available, OSI stack conforming to GOSIP in a POSIX compliant operating system. The motivation for this project was based on the observation that the first openly available implementation of TCP/IP promoted the acceptance and use of TCP/IP. Likewise, it is hoped that the dissemination of a reference implementation of the OSI protocol stack would promote interest and acceptance of the OSI protocol suite.

Due to its widespread distribution, the BSD 4.4 release was selected as the mechanism to distribute this reference implementation. The remainder of this note is a terse description, layer by layer, of the OSI protocols to be distributed.

The data link layer will support IEEE 802.2 LLC type 1 (connection-less) on top of an IEEE 802.3 media. The LLC will provide passive support for XID and TEST; that is, the LLC will respond to XID and TEST commands but will not generate them. Support for an X.25 public data network interface will be provided.

Both the Connection-Less Network Service (CLNS) and the Connection Oriented Network Service (CONS) will be implemented. The CLNS will operate over the Connection-Less Network Protocol (CLNP). The CLNP will support three options: source route, record route, and the globally unique quality of service. The latter option is required in order to provide access to the congestion experienced bit. A mechanism to encapsulate CLNP packets as described by RFC 1070 will also be provided.

The CONS will be provided by utilizing the X.25 interface. The X.25 connection management support does not strictly follow the rules set forth in ISO 8878.

The CLNP will support any arbitrary NSAP address format. The kernel routing table has been redesigned to allow arbitrary OSI NSAP addresses up to 20 bytes in length. The End System to Intermediate System routing protocol (ES-IS) has been implemented. This implementation will support the report, record, and query configuration functions; the configuration response function; and the redirect generation and receipt functions. Values for the Multicast addresses "All End Systems" and "All Intermediate Systems" are taken from the December 1987 NBS agreements.

The Transport Protocol (TP) implementation will support both class 0 and 4. All TP functions for each class will be supported.

In addition to operating TP class 4 over CLNP and TP class 0 over CONS, TP class 4 will operate over DoD IP.

The interface to TP will utilize BSD sockets. The semantics of the accept() system call have been changed slightly for the TP case. Specifically,the CC TPDU will not be transmitted when the user returns from accept(). Rather, the CC TPDU will only be transmitted when the user tries to send or receive data. The user data portion of the CR and CC TPDU will be accessed by the new ancillary data field of the msghdr structure (used in conjunction with sendmsg() and recvmsg()). The end of TSDU indication will be supplied to the user through the flags parameter of the recvmsg() system call.

The Session, Presentation, Reliable Transfer, Remote Operations, and ACSE protocols as well as ASN.1 support will be provided by the ISODE package. An X.500 directory service will be provided as well.

The 5 applications supported in this stack will be the Message Handling System (MHS); File Transfer, Access, and Management (FTAM); Virtual Terminal Protocol (VTP); MHS/822 gateway; and an FTAM/FTP gateway.

The principle players in the OSI-POSIX project are UC Berkeley - responsible for the subnet support, CONS, X.25 device driver, integration and testing; UW Madison - responsible for TP0 and TP4, CONS, CLNP and ES-IS; The Wollongong Group - responsible for ISODE; The University College London - responsible for MHS; MITRE - responsible for VTP; and NIST - responsible for the MHS/822 gateway, FTAM/FTP gateway, X.500, conformance testing, and overall coordination.

# OSI PROTOCOL SUPPORT

# IN

# 4.4 BSD

Robert Hagens
University of Wisconsin - Madison

HIGHLIGHTS: Some of the protocols that are going to be in the 4.4 BSD release were originally developed at Wisconsin on PC RTs under an IBM project.

# THE OSI-POSIX PROJECT

# NIST

- The Mission:

Provide an openly available GOSIP conformant OSI stack in a POSIX compliant operating system

HIGHLIGHTS: Berkeley and others were funded by NIST under the OSI-POSIX Project, the purpose being to get OSI out to the user community. When Berkeley sends out the next release, it will contain a complete OSI stack, which will help immensely with the "furtherment" of OSI.

# BSD 4.4 OSI PROTOCOL SUPPORT

## NETWORK LAYER (Connectionless)

- Supported options
    - source route
    - record route
    - QOS

- Supports echo packet type

- Supports congestion experienced bit

# BSD 4.4 OSI PROTOCOL SUPPORT *

## NETWORK LAYER (Connection Oriented)

- Partial support of ISO 8878

- Connection manager glue

- X.25 device driver

# BSD 4.4 OSI PROTOCOL SUPPORT

# NETWORK LAYER (ES-IS)

- System type statically determined

- Supported functions
  - report configuration
  - record configuration
  - query configuration
  - configuration response
  - redirect generation
  - redirection receipt

- Multicast addresses follow December 1987 NBS agreements

# Kernel R.I.B management in 4.4BSD

HIGHLIGHTS: Very briefly, the way the Routing Information Base (RIB) is updated and the way routes are looked up will be completely rewritten in the 4.4 release. Basically, the Kernel RIB will be updated through a socket interface rather than ioctl. Route is now represented as a pair of a "destination" and a "mask". The two are used together to denote an equivalence class of addresses that are reachable via that route. Routes are stored right now in a radix tree. (For more detailed information you may review the following five slides and/or direct your questions to Keith Sklower and Van Jacobsen.)

Keith Sklower

sklower@Berkeley.EDU

# Routing

- Kernel RIB updated through socket interface

- Route is now a pair (DST, Mask)

- Routes are stored in radix tree

---

- ## User Interface

  - Pass messages to/from Kernel via special socket:

    S = socket(SOCK_RAW, PF_ROUTE, proto);

    where the values 0 mean all protocols, PF_INET mean inet only related info, etc.

  - Manually alter R.I.B. by message types

    RTM_ADD, RTM_DELETE, RTM_CHANGE,

    success/failure indicated in returned message

  - Learn of spontaneous events:

    RTM_REDIRECT, RTM_MISS, RTM_LOSING

  - Someday, maybe do remote management by splicing into stream.

- New definition of route:

  Route represented by pair ($D$ , $M$).

  $D$ is a *struct sockaddr*.

  $M$ is a bitmask stored as a *sockaddr*.

  Route is collection of trial destinations $T$ ,
  so that $T$ & $M$ == $D$ & $M$ .
  (Here, ''&'' represents bit–wise ''and'').

- If $M$ is all 1's, route contains only D.
  (Shorthand: $M$ == 0 means this case.)

- If $M$ is all 0's, route contains all hosts.
  (Represents default!)

- New Routing Table:

  - Radix tree based on destinations $D$ .
    (We eliminate one-way branching).

  - Represented by a binary tree with separate
    internal nodes and leaves.

    Each node indicates bit position to test,
    and lists masks which might apply to in subtree.

- Searching Algorithm:

  If bit at position is on
  go right;
  else
  go left;
  until you reach leaf.

  At leaf, perform masking & compare operation.
  Possibly backtrack to find hierarchical defaults.

# BSD 4.4 OSI PROTOCOL SUPPORT

## NETWORK LAYER (General)

- Supports arbitrary OSI NSAP address formats
- Supports CLNP
- Supports ES-IS
- Supports CONS*
- Supports EON

HIGHLIGHTS: The Network Layer will support the Connectionless Network Layer Protocol (which is ISO 8473 or simply OSI IP), the End System to Intermediate System routing protocol, and the Connection Oriented Network Service (which is basically using X.25 as a connection-oriented network service). There are currently no plans to support IS to IS routing protocol but Berkeley would be more than happy to include that in the release if someone would kindly contribute one.

# BSD 4.4 OSI PROTOCOL SUPPORT

## DATALINK LAYER

- Supports 802.2 LLC type 1
- Supports 802.3
- Passive support for XID and TEST
- Uses OSI NL SAP

HIGHLIGHTS: In describing the stack present in the 4.4 release, let's begin with the Datalink Layer and go bottom up. The Datalink Layer is going to support 802.2, LLC type 1 (connectionless) on top of 802.3 networks. The 802.2 and 802.3 support will be only for the OSI Network Layer. There will be passive support for the 802.2 passive support for the 802.2 XID and test commands. Also, X.25 can be supported at the Datalink Layer, as well as the Network Layer, assuming an appropriate X.25 interface can be acquired.

- Idea stolen from Van Jacobsen,
  who has somewhat different algorithm.

  - Tree shape independent of insertion order.
    Random input generates mostly balanced tree!

  - Searching time:
    (log (# routes)) * (# of germane masks)

  - Insertion + Deletion time:
    (log (# routes)) + (# of germane masks)

# BSD 4.4 OSI PROTOCOL SUPPORT

## TRANSPORT LAYER

- Supports class 0* and 4
- Supports all TP functions for classes listed
- Supports TP4/DoD IP path
- Uses BSD sockets
  - CC TPDU sent after read/write
  - CR & CC data accessed as ancillary data
  - EOTSDU signaled through msghdr structure

HIGHLIGHTS: The Transport Layer will support ISO transport class Ø and 4. There will be an additional interoperability path that will allow TP4 to operate on top of DOD IP, in parallel with operating on top of ISO IP. The interface to the Transport Layer will be essentially Berkeley sockets with some changes to support some of the different features of TP.

# BSD 4.4 OSI PROTOCOL SUPPORT

## UPPER LAYERS

- Session
- Presentation
- Reliable Transfer, Remote Operations
- ASN.1
- Directory Service

HIGHLIGHTS: The upper layers are composed primarily of the ISODE OSI offering: IS session and IS presentation, reliable transfer, remote operations and ACSE, ASN.1, and a directory service. It is not clear yet whether the directory service will be the ISODE directory service or one that is being worked on at NBS.

# BSD 4.4 OSI PROTOCOL SUPPORT

# APPLICATIONS

- MHS
- FTAM
- VTP
- MHS/822 Gateway
- FTAM/FTP Gateway

HIGHLIGHTS: Applications will include MHS (an implementation done by UCL), FTAM (the ISODE FTAM), Virtual Terminal Protocol (done by Mitre) and two gateways. Who will supply the gateways is unknown at this time. NIST has worked on the MHS/822 gateway and, I believe, have come up with a specification for the FTAM/FTP gateway. Mitre has also done an FTAM/FTP gateway. And, the UCL MHS implementation has an MHS/822 gateway.

# BSD 4.4 OSI PROTOCOL SUPPORT

# THE PLAYERS (LOWER LAYERS)

- UC Berkeley
  - Subnet support
  - CONS*
  - X.25
  - Integration
  - Testing

- UW Madison
  - Transport class 0, 4
  - CONS*
  - CLNP
  - ES-IS

HIGHLIGHTS: This is an overview of who has done what at the lower layers. Berkeley is doing the subnet support, integration of the CONS to a new X.25 driver, integration of the whole thing, and testing. Wisconsin provided transport class Ø and 4, a connection-oriented network service implementation, CLNP and the End System to Intermediate System routing protocol.

- Efficency Hacks:

  1. Fill in missing bits in "Mask with holes"

     (+): no lookaside, match is guaranteed

     (−): more ESIS activity.

  2. Re-order bits.

     (−): How?

  3. Real PATRICIA

- Other Concerns:

  1. How to store both

     128.32 (ffff0000)

     + 128.32.0 (ffffff00) ?

     (for now, dup. leaves)

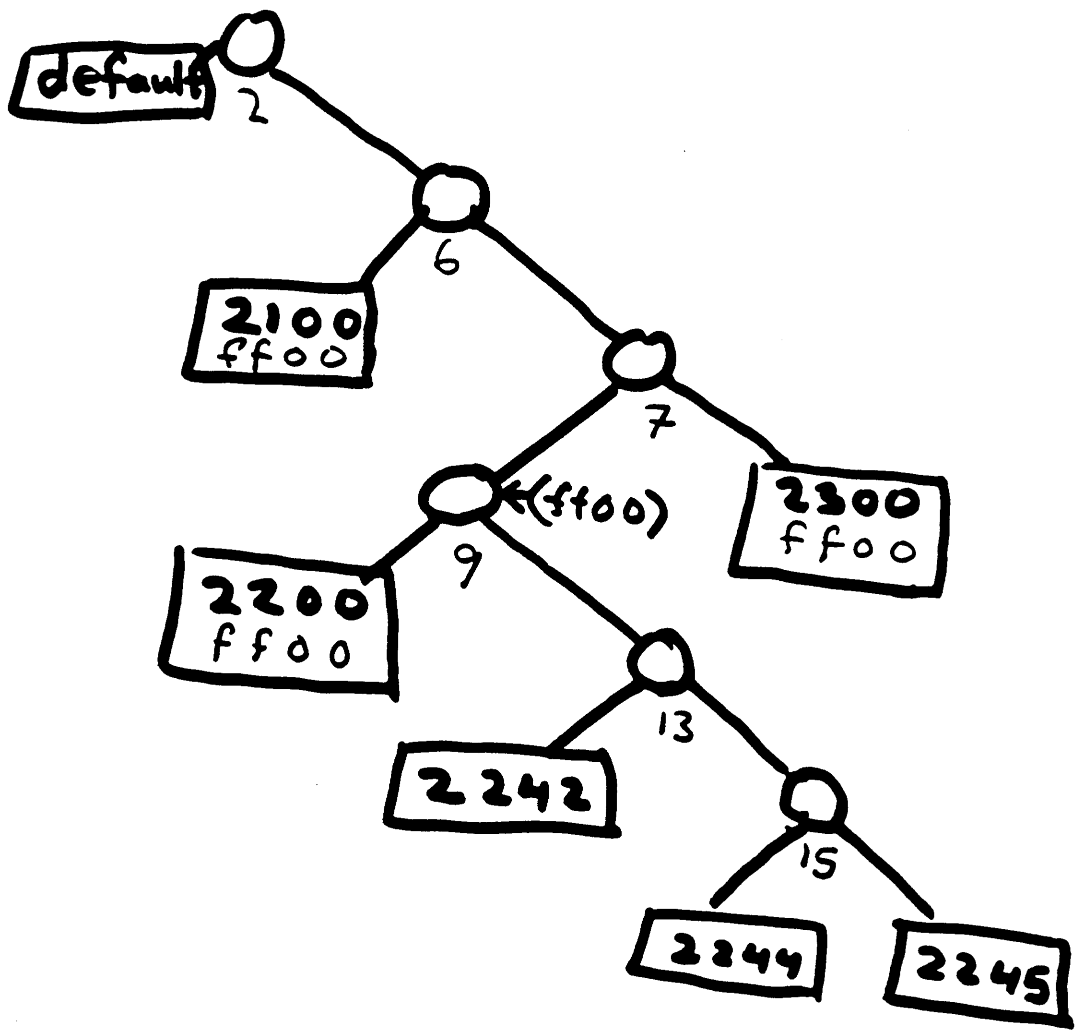


Reordered Bits

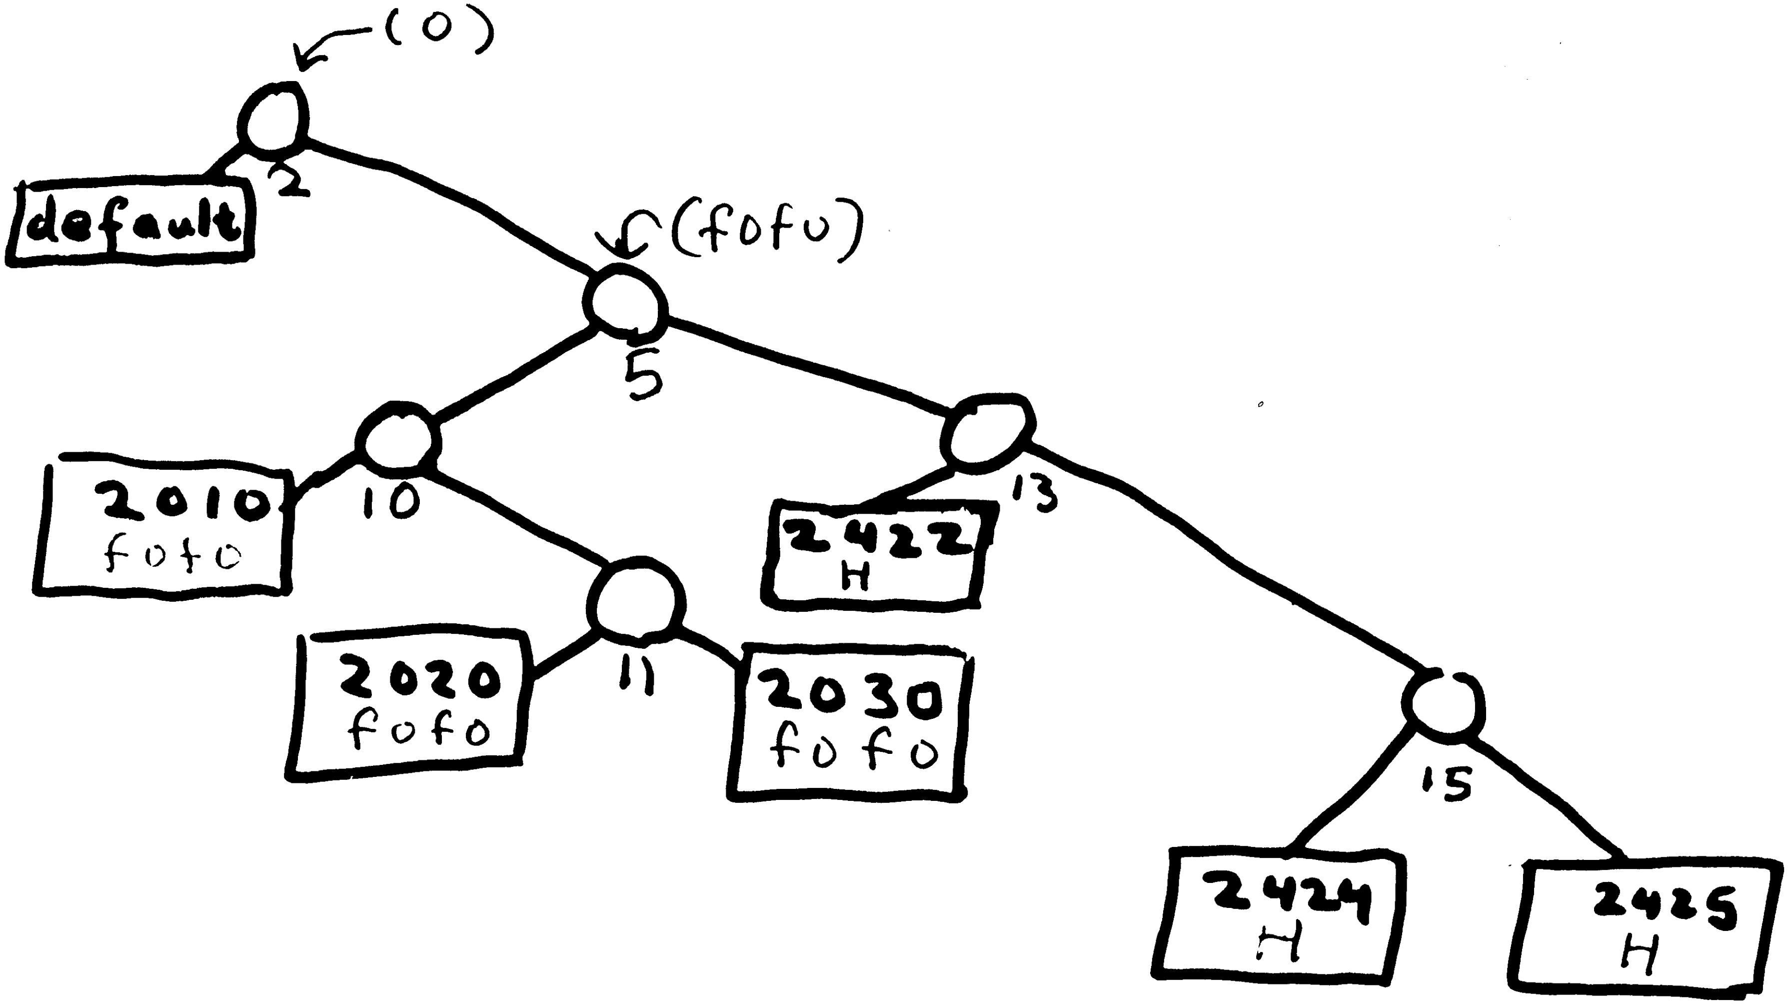


Semi-pathological

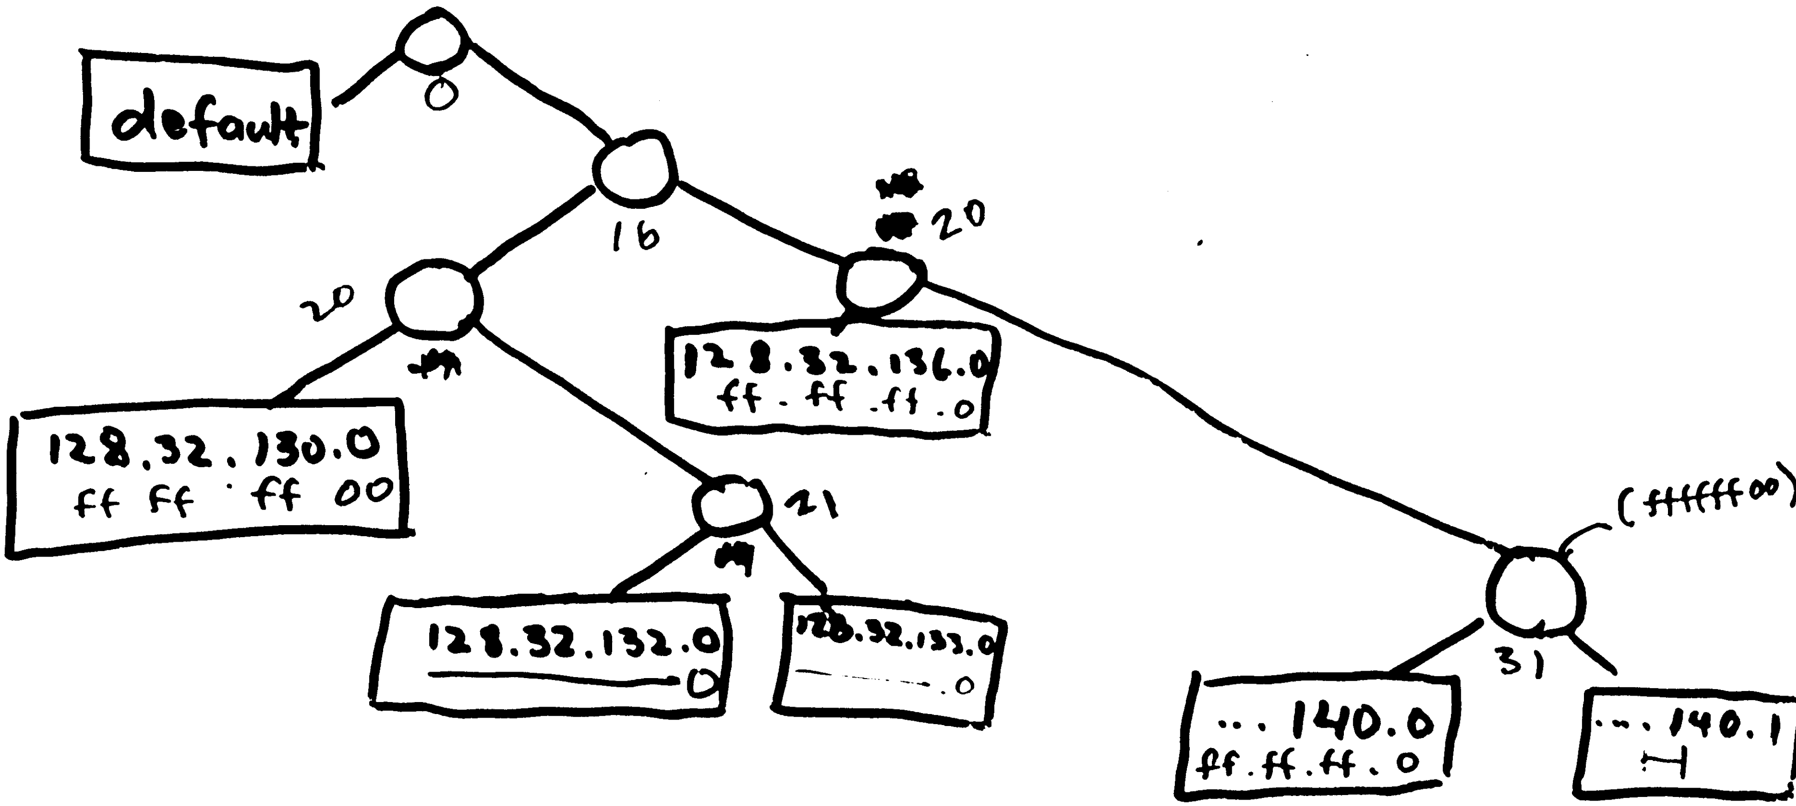


6 Berkeley Routes.

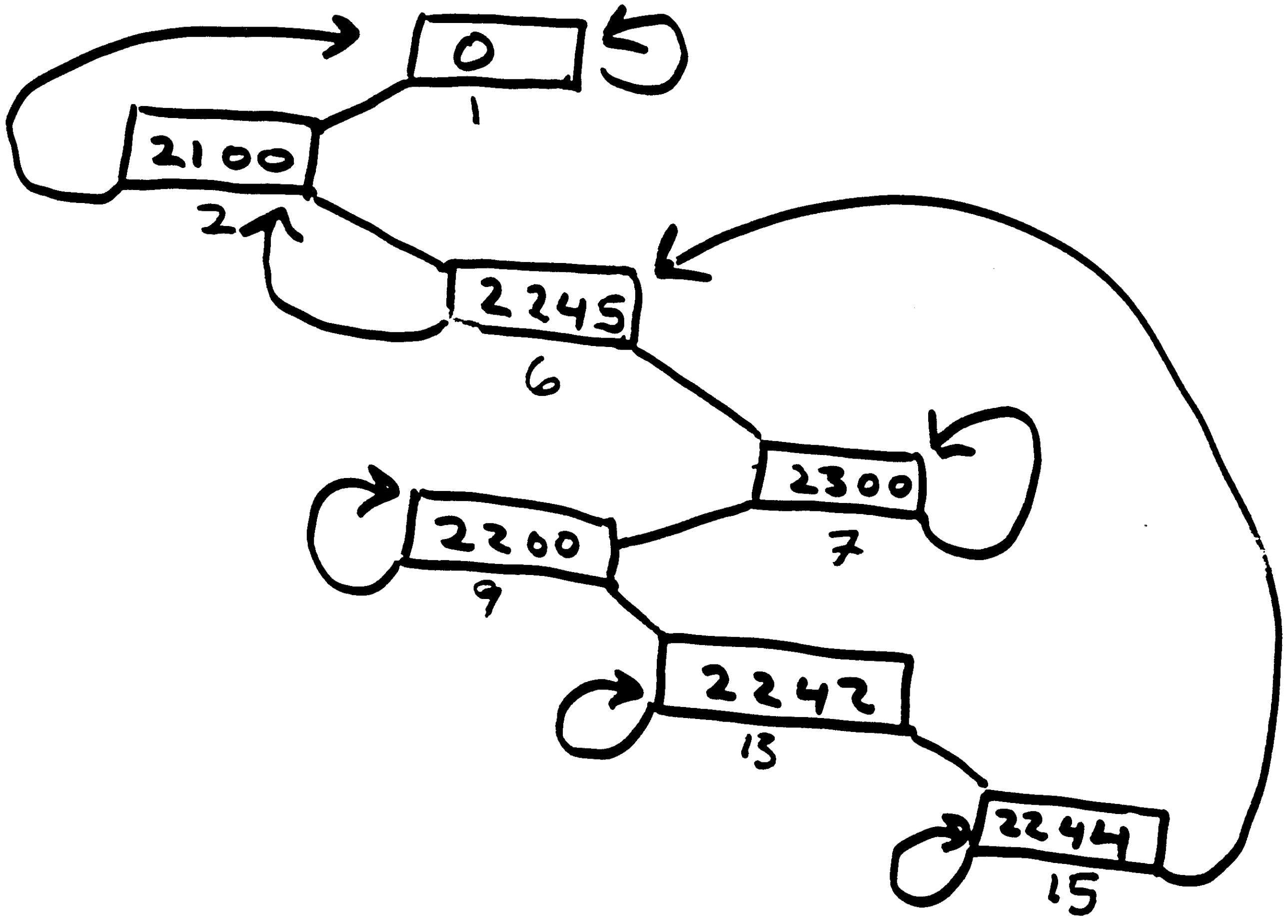


Patricia, reordered.

# BSD 4.4 OSI PROTOCOL SUPPORT

# THE PLAYERS (UPPER LAYERS)

- The Wollongong Group (ISODE)
  - Session
  - Presentation
  - Reliable Transfer, Remote Operations
  - FTAM
  - ASN.1
- UCL
  - MHS
- MITRE
  - VTP
- NIST
  - MHS/822 Gateway
  - FTAM/FTP Gateway
  - Directory Service
  - Conformance Tests

HIGHLIGHTS: At the upper layers, Wollongong is responsible for session, presentation, generally what you see here; UCL for MHS; Mitre for VTP and, perhaps, the FTAM/FTP gateway; and NIST for the remainder and general organization.

Q. Has there been any thought of using something from one of the two Express Project people (Michigan or CMU) to provide more capability for mail, i.e., the multi-media editor?

A. Not that I know of. Last time I spoke with Marshall Rose, Wollongong was going to write some type of X-based interface to the X.400 User Agent. I do not know anything about the current status of that.

Q. What sorts of problems do you expect people to have when they get this release? We all remember the 4.1C days when TCP wasn't tuned, or 4.2 days when there were interesting problems. So now we are going to see 4.4 with OSI. What sorts of problems do you expect users to come crashing into in the OSI suite that we are going to find ourselves fixing in 4.5?

A. Routing... There is a large effort in the OSI Working Group to try and identify those problems and provide feedback to Berkeley, if possible, before the final release. Also, unlike 4.1C, this will have another networking protocol implementation included and that one will have been fairly well tuned and debugged.

Q. What is the "target" release date?

A. There isn't a definite plan for the release, or the contents of 4.4, yet. What is planned is the order in which we are going to be doing different projects. The current plan is that there will be a release to the people involved in this project, basically the ones shown on the slides and various other interested parties, as we do the final integration in the next few months. Probably three or four months after that, sometime mid-summer, one of two things will happen: a beta test release of 4.4 or a "4.1C style" release which is an intermediate release of "here's the stuff done so far, we are still working on other things". In either case we would expect that some time in roughly that time frame this code would be available to basically anyone who wants it, as long as they don't expect to be able to call us with problems, unless they have fixes!

# Internet Worm

Michael Karels
University of California - Berkeley

HIGHLIGHTS: There were three main things the program (worm) did in order to gain access to machines and move on from there. It attacked by means of sendmail, finger server, and password guessing.

There were two specific software flaws exploited, and certain software features exploited after that. The first software flaw was a "trojan horse" in the sendmail SMTP server that was there pretty much intentionally from a long time ago and had never been removed. The "trojan horse" was a "debug" command added to the SMTP command set in this implementation. The debug command for the most part just printed all kinds of debugging information on the SMTP socket, quite contrary to SMTP protocol specifications, but it also overrode a couple of checks in the program for certain non-standard mail destinations, in particular the ability to sendmail to programs. The result was that it sent a message to a command interpreter using a command that stripped off the header and than interpreted the rest as commands. The body of the message contained a command to create a small .C file (C source file) which was called the "grappling hook" or "downloader", and then sufficient commands to compile that and run it, along with an address and magic number to connect back to the previous machine and "suck over the rest of itself". The "rest of itself" was a 4.3 VAX binary >0 file, not loaded but compiled, and a SUN 3 object file, the source for the loader program itself. [We would run commands using that command interpreter to compile, using each of those in turn (or as we loaded the thing), and then try to run it and it would go on and try to crack other machines and perform these other attacks.]

# WORM ATTACKS

- sendmail

  SMTP "debug" command overrides checks for mail to programs

  - send message to command interpreter, stripping header
  - body creates C source file for loader, compiles and runs
  - loader connects to "server" on infecting host, fetches VAX (4.3) and Sun objects, loader source, then compiles and runs

HIGHLIGHTS: The other method used was by way of the finger server, which is a trivial program, but it managed to have three different bugs, fixing anyone of which would have prevented this from working. The main problem was that the finger server program failed to check the length of the input string for the user name so it would just overrun its buffer. The buffer happened to be on the stack and so the worm sent a very carefully constructed string that would overrun that buffer and overwrite the return frame on the stack that would replace the return PC. The finger server did a return from the main program as opposed to an exit (which is another thing that could have prevented this from working).

To exploit these bugs the worm included some code in this string as well (VAX code only; there wasn't a SUN version) that the return PC pointed to. This code was three or four instructions that set up a call to exact command interpreter which of course would mean that you would get a prompt and could run commands, like sending across the grappling hook, compiling it and running it.

These were the main two methods of entry that the worm exploited: sendmail and finger server.

>1 and connection oriented internet protocol working group
Claudio Topolcic

- Two parallel tracks
  - longer term connection oriented protocol issues
  - short term ST specification

1. Connection oriented protocol issues
   - our bias
     - there are applications that aren't supported ideally by IP
     - simple extensions to IP will not support them well enough
     - our thinking need not be constrained by current IP
     - proposed (not agreed) definition of "connection"
       - Gateways have per-"flow" state and can recognize the data packets of a "flow"
       - The end systems send a "connection setup" message before sending data packets

       * NOT a circuit

   - Underlying Networks
     - What services will such a protocol need from the underlying nets?
     - what can it expect?
     - How to reconcile the two?

   - Gateway functions
     - Resource management
       - unequal sharing of resources
       - don't necessarily reserve specific resources for a specific connection
       - network bandwidth
       - gateway processing
       - gateway buffering
     - Routing
     - Connection setup
       - eliminates need for later redirects and fragmentation
     - Reacting to asynchronous state changes, such as breaks.

   - We are writing this up

2. ST specification

- Goals
  - support current and planned activities
    Ex: conferencing
  - must be practical to implement
  - usable in different environments
    Ex: across IP-only systems where resource management is not a problem
  - support research in connection oriented communication
    Ex: from the longer term effort
  - support other applications
    Ex: TCP

- Specific changes to current ST spec
  - Remove Access Controller
    - must shift access control into "Application"
    - need source of non-globally unique connection ID
      - hop-by-hop ID
      - local (to net) unique ID
    - eliminate "Forwarding Bit Map"
      - do forwarding with ID
  - Connections result in simplex tree rather than omniplex streams
  - Eliminate Envelope header
    - aggregation as implemented didn't do what ~~not~~ was intended so we must re-examine this

- We will have video conferences and mail to discuss mechanisms and write them up

- The Applications
  → what information must flow between the "Application" and the internet layer - QOS
    - distribution of offered bandwidth
      Ex: peak, average, burst factor
    - required delay distribution
    - ~~red~~ reliability requirements
    - can/cannot be supported
    - hints
    - changes in state

# Telnet Linemode

Dave Borman
Cray Research

January 1989 IETF Meeting
TELNET Linemode Working Group Summary
Reported by Dave Borman

The agenda was rather short, we met for about an hour.

1. RFC 1080, Telnet Flow Control Option, which overlapped our spec.
2. SLC_DSUSP: Do we need it?
3. SLC_AYT: Should we add it?
4. Timing-mark / TELNET Synch
5. SLC_SET: new SLC_NOSUPPORT bit
6. Output character processing

Item 1:

RFC 1080 showed up a couple of weeks before the meeting. There was annoyance over the fact that it overlapped things that were happening in the our document, and we had not been consulted before it was released. The damage was done, though. It was decided that we would bow to 1080, and use it to negotiate flow control.

Item 2:

It was decided that SLC_DSUSP was not needed, and so it was eliminated form the spec.

Item 3:

SLC_AYT was added, it was an oversight in previous versions.

Item 4:

It was agreed that more wording needed to be added about how to do flushing using TIMING-MARK and the Telnet "synch".

Item 5:

The SLC_NOSUPPORT bit had been added since the previous meeting, and there was a brief discussion on what it did, and a general approval that it was a good thing to add.

Item 6:

It was decided to make it explicit in the spec that output character processing continues to be done by the remote system. Going into linemode does not change this. We do not want to change the NVT. So, for output, a "newline" is mapped to CR LF, an explicit carriage return is CR NULL, and an explicit line feed is just LF.

Also, the name of the option in previous drafts had been LOCALEDIT. It was changed to LINEMODE.

Attendees:

| | |
|---|---|
| Dave Borman | dab@cray.com |
| Bruce J. Schofield | schofield@edn-unix.arpa |
| Drew D. Perkins | ddp@andrew.cmu.edu |
| Jeffrey Burgan | jeff@twg.com |

Developments since the Austin Meeting:

The SLC section has been re-organized to be split into two parts: those special characters that get translated into TELNET commands, and those special characters that are interrpreted locally.

Option values have been assigned: LINEMODE is 34, ABORT is 236, SUSP is 237, and ABORT is 238.

On the mailing list, discussion is going on about how to decide when output/input should be flushed when various Telnet commands are sent (IP and ABORT, most notably). Once this is resolved, a new copy of the draft document (draft 5) will be made available for general discussion. There are also efforts underway to implement the spec, so that there is at least one working implementation before the spec is submitted as an RFC.

The working group will not be meeting at the next IETF meeting. We are in a hold state right now, once the RFC has been submitted we will officially disband the working group.

**IETF Working Group Roster**
University of Texas at Austin
18–20 January 89

Working Group Session: TELNET Linemode
Chairperson: Dave Borman

| Attendees: | E-Mail Address: |
|---|---|
| 1. David Borman | dab@cray.com |
| 2. Jeffrey Burgan | jeff@twg.com |
| 3. Drew Perkins | ddp@andrew.cmu.edu |
| 4. Bruce Schofield | schofield@edn-unix.arpa |

# Telnet Linemode

- RFC 1080 - FLOWCONTROL - obsolete it, or defer to it? Probably the latter

- Eliminate SLC_DSUSP, add SLC_AYT, SLC_SYNCH
- Add discussion for TIMING-MARK/synch
- Output processing happens on server, not client

Futures:

1) Implement it, to flush out any bugs
2) Submit as an RFC
3) disband the working group

linemode@uc.msc.umn.edu
linemode-request@uc.msc.umn.edu
draft on uc.msc.umn.edu through anonymous ftp, in "linemode"

# User Services

Karen Bowers

Corporation for National Research Initiatives

January 1989 IETF Meeting
User Services Working Group Summary
Reported by Karen L. Bowers (NRI)

The newly formed User Services Working Group convened the first time during the IETF meeting held at University of Texas, Balcones Research Center in Austin, Texas, 18-20 Jan 89. The agenda was fairly ambitious and included:

- Introduction of the new User Services Working Group and its relationship to such IETF Working Groups as Joint Monitoring, OINOCS and Interoperability, and to INTERNICS;

- A brief introduction of the chairperson, Karen L. Bowers, who is currently on the technical staff at Corporation for National Research Initiatives;

- The review and rewrite of the draft charter and objectives;

- Establishment of "selection criteria" for actions/projects to be undertaken by the USWG; and

- A first step in selecting and prioritizing actions/projects to be addressed.

The morning of the first day was spent defining what the goals and objectives of the USWG should be. This led to the development of the DRAFT CHARTER: to identify service requirements for "people who help end-users" (e.g. local net managers); to develop tools and materials to aid in productivity of end users; and to coordinate these actions with existing/newly forming organizations such as IETF working groups, NICS, NOCs, and user services organizations.

To ensure the efforts of the USWG are "production-oriented", that is will produce results, SELECTION CRITERIA were presented by the Working Group Chairperson and further expanded by the USWG members. These SELECTION CRITERIA will aid in determining what projects/requirements will be undertaken by the USWG:

1. Project/selected action must lend itself to accomplishment within a reasonable timeframe (say 1-3 years);

2. Must culminate in a measurable/quantifiable end result (RFC, network users directory, etc.); and,

3. Must address user assistance needs and not technology specific requirements (e.g. routing) (must be user-oriented).

4. Products/tools resulting from these efforts must not only address user information requirements but must be designed to be both easily maintained and easily updated. Accountability must be built in to ensure these products/tools are in fact maintained and updated.

Several members voiced the concern that it is very important the USWG does not duplicate efforts, but rather surveys existing resources and determines the most appropriate approach. Three alternative approaches were defined: to produce a totally new product, to enhance/improve/influence an existing resource(s), or to simply table the action for later consideration. With this in mind, the decision was made to survey the internet environment for existing tools (documentation/techniques) currently supporting user service requirements. A list of AREAS FOR CONSIDERATION was drawn up during the afternoon.

Individual members volunteered to research what tools are currently available and provide this information to the USWG in a mini-briefing at the next USWG meeting. At that time specific issues/actions to be addressed by the USWG will be identified and prioritized. One or two projects derived from the list will be selected for immediate action.

The AREAS FOR CONSIDERATION and respective "responsible persons" are:

- Internet Connection Checklist (Craig Partridge and Karen Bowers) - such as, a standardized format describing how to connect: permission requirements, network specific procedures, guidance on physical (circuit/equipment) interface requirements and software (protocol) requirements, and Internet specific procedures (initial configuration requirements: net# assigned, name server, subnets, hand configure routing tables, etc...)

- Network Resources Directory (Karen Roubicek, Tracy LaQuey and Mary Stahl) - could include: short/concise description of each network, net #, and net maps; and POCs for various actions such as permission to connect, network engineering, network ops, 800#s, support services (assistance with routing/performance problems), etc.

- Good Neighbor Standards (Nethics) (Wayne Wilson and Craig Partridge)

- Network To Network Mail (Tabled: John Quarterman to revitalize later)

- How To Setup A Campus NIC, NOC (Tracy LaQuey)

- Bibliography Of Documents Every Nic Should Have (Marty Schoffstall and Francine Perillo): 1) end users, 2)local net managers

- Mailing List Management (Jim Sweeton)

- Simple Configuration Control (Bill Anderson)

- Simple-Minded Debugging Tools (Bob Enger)

- Consolidated Common End User Questions w/ Answers (Sergio Heker et al)

- "How To" Manuals for local net managers covering key/common systems and configurations (LANS, network management, NIC management, NOC management, etc.) (Tabled: to be further defined next session.)

- How To "Get Fixed" (Tabled)

- User Email Address Book (by discipline?); define standards/ guidelines for Postmaster and Net Manager (Need to review Dave Clark's White Pages proposal/ tabled until the next meeting.)

The remainder of the first day was spent identifying possible invitees for future USWG meetings. This list follows:

| | |
|---|---|
| Mary Stahl (NIC) | Pat Sisson (SPAN) |
| Mel Pleasnat (Rutgers) | Scott Early (BITNET) |
| Steve Miller (U of MD) | Rick Adams (UUNET) |
| Mark Horton (previous UUCP) | Laura Breeden (CSN) |
| Philip Prindeville (Montreal) | Juha Heinanen (Finn Network) |
| Marilyn Martin (UDCNET) or | John Vemco (UDCNET) |
| Henry Nussbaker (ILAN) | Daniel Karrenberg (EUNET) |
| Dennis Jennings (EARN) or | Eric Thomas (EARN) |
| Gene Spafford (USENET) | Betty Schermerhorn (NPAC) |
| Jill Foster (JANET) | |

The second day's session was very brief. Members were given the opportunity to voice any concerns/issues that had occurred to them during the night. A few points from the previous day's activites were rehashed and the results further reinforced the USWG plan of attack.

Attendees of the 18-19 Jan 89 USWG meeting follow. The email addresses of these members will be placed on a USWG mailing list, as well as others who have requested to be included in future correspondence. Karen Roubicek kindly volunteered to provide this service.

ATTENDEE ROSTER

| NAME | POSTAL ADDRESS/ PHONE NUMBER | EMAIL ADDRESS |
|---|---|---|
| Karen L. Bowers | Corp for National Research Initiatives | bowers@sccgate.scc.com (temporary) |

| | | |
|---|---|---|
| | 1895 Preston White Dr<br>Suite 100<br>Reston, VA<br>(703) 620-8990 | |
| Jose Joaquin,<br>Garcia-Luna | SRI International<br>333 Ravenswood Ave<br>Menlo Park, CA 94025<br>(415) 859-5647 | garcia@sri.com |
| Tracy LaQuey | Computation Center<br>The University of Texas<br>Austin, TX 78712<br>(512) 471-5046<br>(512) 471-3241 | tracy@emx.utexas.edu |
| Sergio Heker | John von Neumann Ctr<br>665 College Road<br>Princeton, NJ 08540<br>(609) 520-2000 | heker@jvnca.csc.org |
| Martin<br>Schoffstall | NYSERNet<br>165 Jordan Road<br>Troy, NY 12180<br>(518) 283-8860 | schoff@nisc.nyser.net |
| Robert Enger | CONTEL Federal Systems<br>1300 Quince Orchard Blvd<br>Gaithersburg, MD 20878-4199<br>(301) 840-4040 | enger@gburg.scc.com |
| Paul Mockapetris | USC/ISI<br>4676 Admirality Way<br>Marina del Rey, CA 90292<br>(213) 822-1511 | pvm@isi.edu |
| Karen Roubicek | NNSC<br>BBN Systems Technologies Corp<br>10 Moulton Street<br>Cambridge, MA 02138<br>(617) 873-3361 | roubicek@nnsc.nsf.net |
| Jim Sweeton | Merit<br>1075 Beal Avenue<br>Ann Arbor, MI 48109-2112<br>(313) 936-3000 | jcs@merit.edu |
| Wayne F. Wilson | AFCC<br>1842 EEG/EEM<br>Scott AFB, IL 62225-6001<br>(618) 256-4585 | none |

| | | |
|---|---|---|
| Craig Partridge | NNSC<br>BBN Systems and Tech Corp<br>10 Moulton Street<br>Cambridge, MA 02138<br>(617) 873-2459 | craig@bbn.com |
| Joe Choy | NCAR<br>P.O. Box 3000<br>Bouldar, CO 80307<br>(303) 497-1222 | choy@ncar.ucar.edu |
| Mike StJohns | DCA<br>Code B612<br>Washington, DC 20305-2000<br>(703) 285-5133 | stjohns@beast.ddn.mil |
| Francine Perillo | cisco Systems<br>1360 Willow Road<br>Menlo Park, CA 94025<br>(415) 326-1941 | perillo@cisco.com |
| Ole Jacobsen | ConneXions/ACE<br>480 San Antonio Rd, Suite 100<br>Mountain View, CA 94040<br>(415) 941-3399 | ole@csli.stanford.edu |
| Bill Anderson | MITRE Corp<br>Burlington Road MS E066<br>Bedford, MA 01730<br>(617) 271-3388 | wda@mitre-bedford.org |
| Stewart Vance | University of Texas<br>Balcones Research Center<br>10,100 Burnet Road<br>Austin, TX 78758 | none listed |
| John Quarterman | TIC<br>701 Brazos Suite 500<br>Austin, TX 78701-3243<br>(512) 320-9031 | jsq@longway.tic.com |
| Mark Fedor | NYSRNet, Inc.<br>Rensselear Tech Park<br>165 Jordan Road<br>Troy, NY 12180<br>(518) 283-8860 | fedor@nisc.nyser.net |

Agenda for Jan Meeting

- Introduction of new User Services Working Group and its relationship to such IETF WGs as Joint Monitoring, OINOCS and Interoperability, and to INTERNICS (Phill Gross)
- Brief Introduction of Chairperson
- Review/Modify/Finalize Draft Charter, Objectives and Selection Criteria ( Karen Bowers)
- Bráinstorm Action Items/Issues to be addressed
- (If time permits) Select and Priortize Actions/Items to be addressed and, as necessary, assign a Responsible Person.

Draft

CHARTER: to identify service requirements for "people who help end users" (e.g. local net mgrs) and to develop tools & materials to aid in productivity of end users. This group will coordinate these actions with existing/newly forming organizations such as IETF WGs, NICs, NOCs, & user services organizations.

NAT'L NET

↑

Network Mgr. NSF, DCA

↑

NICs/NOCs

↑

Regional Net Mgr

↑

Local Net Mgr

↑

End User

W.G

SELECTION CRITERIA (for projects/requirements to be addressed by the User Services Working Group):

1. Project/selected action must lend itself to accomplishment within a reasonable timeframe (say 1-3 years).

2. Must culminate in a measurable/quantifiable end result (production oriented; e.g., RFC, network users directory, etc.)

3. Must address user assistance needs and not technology specific requirements (e.g. routing) = user oriented

4. Products/tools resulting from these efforts must not only address user information requirements but must be designed to be both maintainable and easily updateable.

5. Not duplicate efforts / survey existing resources

6. Build in Accountability (carrots vs the stick)

1. Survey environment for existing tools (documentation/techniques) supporting user service requirements.

2. Determine most appropriate approach:
- produce totally new product
- enhance/improve/influence existing resource
- table action for later consideration

3. Prioritize issues to be addressed; select 1 or 2 for immediate action.

# Areas for Consideration

- Internet Connection Checklist (Craig Partridge, Karen Bowers)
- Network Resources Directory (Karen Roubicek, Tracy LaQuey, Mary Stahl)
- Good Neighbor Standards - Nethics (Wayne Wilson, Craig Partridge)
- Network to Network Mail {Tabled: John Q. to revitalize later}
- How to setup Campus NIC, NIC (Tracy LaQuey)
- Bibliography of documents every NIC should have:
  1) end users, 2) local net mgrs (Marty Schoffstall; Francine Perillo)
- mailing List Mgt (Jim Sweeton)
- Simple config. control (Bill Anderson)

- Simple-minded debugging tools (Bob Eager)
- Consolidated Common End User Questions w/ Answers (Sergio Heker et al)
- "How To" manuals (Tabled: to be further defined next session)
    - How to get "fixed"
    - 
    - 

Do Not Forget
White pages!!

ISSUES/ACTIONS FOR CONSIDERATION:(to be further expanded)

- A national directory (or directories) of existing networks and associated points of contact to include:

Resource Handbook

1. short/concise description of each network, net#, and net maps,

2. POCs for various actions: permission to connect, network engineering, network ops, 800#s, support services (such as assistance with routing/performance problems), etc.

Internet Connect. Checklist ?

3. a standarized format describing how to connect: permission requirements, network specific procedures, guidance on physical (circuit/equipment) interface requirements and software (protocol) requirements, and Internet specific procedures (initial configuration requirements: net # assigned, name server, subnets, hand configure routing tables..)

- How to set up and establish national procedures for net connections = RFC.

- How to best answer new user problems

- A national (international) network users directory (as a phone book)

- A guide to user training resources

- What is the Internet? Both
- Am I on the Internet? Are you on the Internet? User
- How to hook your campus to the Internet; technology to plug in. (Connection check list) NetMgr
- How to get fixed. Net Mgr
- User email address book (telephone book); define standards/ guidelines for Postmaster by discipline (Assist Dave Clark?) User+Netmgr
  NetMgr
- Network to Network mail
- "Good Neighbor" standards (Netiquette) - Users through Local NetMgr
- Network Resources Directory - Both
- Simple minded debugging tool ← Both
- Consolidated Common User Questions/Answers - Both
- Prepackaged: "How To" Manuals for key/Common Systems! Configurations (LANs, Network Mgt, NIC Mgt, NOC Mgt. etc.) → Local/NetMgrs
- Videos - everyone!

- Simple configuration controls (NOC)
- Non-documentation: ex: Listserv (Eric thomas)

(duplicate interface not implement.)

$\boxed{\text{MAILING LIST MGT}}$

# 5. Network Status Briefings and Technical Presentations

# Merit NSFnet Report

Susan Hares

MERIT/NSFNET

# WORM ATTACKS

- finger server
  failed to check length of user name string
- send string that overflows buffer on stack, overwriting return PC in frame
- string also includes VAX code to invoke command interpreter; return PC set to address of this code.
- commands to transfer, compile and run loader are run.

HIGHLIGHTS: Once it was on a system the other thing that the worm did was to try to attack user accounts by guessing their passwords and then moving around to other machines that they had either the same password on or permission to connect to without supplying a password. So it had a couple of different algorithms for trying to find people's passwords. The first thing it would do is to try very trivial passwords like the null password, people's first and last names, their first and last names written backwards, etc. It also had a word list of about 400 favorite passwords and would try each of those. If all of those failed on all the various accounts, it would then go on to brute force searching the dictionary. Somebody computed how long it would take to go through the average password file with all of these things and it was a couple of months, I think. (I am not sure.) In order to speed up the process, the worm had its own DES encryption package which was approximately nine times faster than the standard library version. Once it did guess a user's password it would encrypt various guesses (the way it could do this was to read the Unix password file which has the encrypted password in it) and then compare the two strings to see if it had guessed correctly. Once it did guess correctly it had that user's account, at least on that machine. It made the assumption that many users would have the same password on many other machines and it would go through its list of known machines and would try logging in as that user using the rexec (rather obscure) network utility, which takes user name and password and then runs a command (if the name and password work). It would try to connect to various known hosts with that user name and password. If that failed, it would then use rexec on the local machine which was guaranteed to succeed if the server was there. Using that to run a command called rsh, which requests a remote shell on various other machines, it would again go through its list of machines, looking for machines that would let this user log in without a password. Therefore, it was able to attack machines that didn't have either of the two security holds if they trusted some other machine that had one of them.

Q: Have they determined whose DES code was used?

A: Rumors are that it came from Bell Labs implementation some time back but I do not know if that is true or not. There are a couple of papers about this that various people have written. One of them was by Don Sealy and he had a little bit of information on the nature of the changes. There are lots of loop unrollings and things- it is not amazingly faster but it is definitely faster than Unix code, which was designed to be slow originally.

# WORM ATTACKS

- Password guessing
  - guess trivial passwords (null, first/last name, etc)
  - guess ~400 favorite words
  - try dictionary words

uses fast DES encryption

when password guessed:

- use rexec to enter known hosts with user name and password
- use rexec to run rsh locally, connecting to known hosts.

HIGHLIGHTS: One thing there was a great deal said about in the press, which is mostly wrong, is the attempts that it made "to limit the infections to one per machine" so it might go undetected. This was generally unsuccessful for various reasons and this is the supposed "minor flaw" that caused the problem to go totally out of control. Having looked at the source code (at least one version of the source code) I don't see any "minor flaws". I see some rather major strategy problems if that is what was actually intended.

The way it <u>tried</u> to limit itself was to create a listening TCP socket when it started up. (Since it is creating a server on a well known port of its choice, it can only do that in <u>one</u> process and so any other instances would fail.) When another program go started it would try connecting to that port and go through secret handshakes exchanging magic numbers. If that succeeded (if the other program was there already and it knew the magic number) then it would do some minimal amount of work, like asking a few passwords, and then go away. In part because of timing problems on this (loads started going up on machines) and in part because of <u>the next thing</u>, this was not very successful.

The <u>next thing</u> is that one time in seven, rather than going through <u>this</u> procedure, it would just skip it and that meant a program would become immortal. It would never check to see if there was another worm on the same machine. These of course accumulated and the load started going way up on a lot of machines. That is why it was noticed so quickly in a lot of places.

# Infection Limitation
## ("population control")

- Worm creates listening TCP socket
- new worms connect to existing worm (if present); after short run, new worm exits if secret handshake succeeds
- 1 in 7 worms skip check, become immortal.

HIGHLIGHTS: There are a number of things the worm did to make it hard to figure out what it was doing (some of them are Unix specific) and how to defend against it. One is that it named the program that it produced. Rather than "worm" or "virus", it called itself "sh". So if you did a command list on Unix, you would see this funny program running called "sh". "Sh" is the name of the Unix command interpreter, so that didn't look so very strange. Once it got started one of the very first things it did was to erase the command-line arguments (like L1.c, Sun 3.0, and VAX.0, and various other things), so when you looked at the list of running processes you wouldn't see much information. It would also do a fork and periodically change its process ID. It would have been very nasty to kill if it had done this more often but it only did it about every five minutes. (If it did it often enough that you couldn't look at what the process ID was and then kill it, that would have been "annoying".) Also, once it got started it removed its own binary file and other files (SUN 3, VAX.0 and downloader program) from the file system so they would not appear there; it just loaded them into its own memory image and encoded them (by XORing with a random pattern, which was the code for the XOR routine). Also, it stored its internal strings in a similarly encoded fashion by just XORing with 80 or 81 hex, depending on which set of things it was. It pretty much had itself contained in its memory image. Also, if it were to fail catastrophically, normally that would write a core image in the file system and somebody could go along and look at it with a debugger. To conteract this, the worm turned that process off by setting the limit on the core dump size to zero. It was trying fairly hard to make sure that people couldn't look at it. Fortunately there <u>is</u> a facility that will generate a core dump of a running program, and <u>that</u> it couldn't turn off. Last thing is that these .0 files were pre-compiled and loaded with everything <u>but</u> the standard system library. Most of the extraneous symbols were removed so it was hard to find the names of sub-routines and that sort of thing (so there were a few global symbols but not very many which somewhat complicated figuring out what the thing was doing).

# WORM DEFENSES

- named program "sh" (command interpeter)
- erased command-line arguments
- changed process ID with fork occassionally
- removed program and files, stored internally (encoded)
- turned off core dumps
- encoded all strings in binary image (XOR)
- eliminated resolved symbols from symbol table

HIGHLIGHTS: There are a couple of other peculiarities which are sort of amusing, One is the way the worm found hosts to try to attack. First thing it did curiously was to run the command that reports the kernel routing tables and it would try to attack the gateways that supported telnet. Therefore, there were very strange things in places where people had gateways that supported telnet (but none of the other things that were used to try to attack them). It would connect to the telnet port. If it got a connection it would immediately close it, usually ringing alarm bells and things. then it would try to do various other things to these hosts, which usually didn't work unless they actually happened to be BBSD systems. It also found hosts to attack by looking through the system list of "equivalent" hosts, hosts whose users "would be allowed to log into this host without supplying a password". It would look at both the root and individual users' accounts to find the lists of accounts that <u>they</u> would allow to log in, assuming there might be reciprocal privileges. Then it would also look at things like mail forwarding information; look at users' forwarding files. (If <u>that</u> is where users forward their mail then maybe they have the same account and password there.) Finally it would look at the list of directly links and all kinds of good things) and would look at the network numbers and then choose various low numbered hosts on that network. (It also understood about Class A networks and IMPs so it carefully tried about 20 low numbered IMPs and various random, low numbers on each of those IMPs.) It was fairly clever in terms of that!

# Choice of hosts

- gateways(!) from netstat -r that support telnet.
- "equivalent" hosts
- root-trusted hosts
- individual users' trusted hosts, mail forwarding
- random low-numbered hosts on local (sub-)nets

# Timeline

- 6pm, Wed:
  first known infection, University of Pittsburgh.
- 7pm, Wed:
  Berkeley infected.
- 1am, Th.
  sendmail bug fixed, Berkeley systems mostly cleaned up.
- 3am, Th.
  sendmail bug posted to USENET, tcp-ip mailing list.
- 5pm, Th.
  decompilation in progress, fingerd bug proven
- 9pm, Th.
  full-scale analysis, fingerd bug posted
- 6am, Fri.
  decompilation finished
- 3pm, Fri.
  last bug reports posted

HIGHLIGHTS: CERT (the Computer Emergency Response Team) was formed by DARPA and several other government agencies to do preliminary screening of problem reports and to call in experts as required. CERT which is located at SEI(CMU), is staffed with 4 personnel and provides 24 hour service via an 800 #.

# 6. Papers Distributed at IETF

Eight Documents were distributed at the January IETF meeting. Two are enclosed: Conformance Testing Profile for Department of Defense Military Standard Data Communications Upper Level Protocol Implementations, and Center for High Performance Computing. The remaining six are listed below. If you are interested in obtaining a copy (copies) of one or more of these, please contact the respective individual.

| TITLE | AUTHOR | DISTRIBUTED BY |
| --- | --- | --- |
| Proposal for Possible Restructuring of Datalink Objects within the MIB | Wellfleet Communications, Inc | John Burress and Terry Bradley (617)275-2400 |
| Proposed Changes to the Experimental Internet Stream Protocol (ST) (Draft) | BBN | Claudio Topolcic and Philippe Park (617)873-3874/ (617)873-2892 |
| Use of the DARPA/NSF Internet as a Subnetwork for Experimentation with the OSI Network Layer (Network WG Request for Comments:XXXX) | U of Wisconsin-Madison and Wollongong | R.A. Hagens (608)262-1017 N.E. Hall and M.T. Rose |
| Guidelines for the Use of Internet IP Addresses in the ISO Connectionless Mode Network Protocol (Network WG Request for Comments: 986 revised) | BBN and U of Michigan | Ross Callon (508)486-5009 |
| The Point-to-Point Protocol (PPP): A Proposed Standard for the Transmission of IP Datagrams Over Point-to-Point Links | CMU | Drew D. Perkins (412)268-8576 |
| Xerox Synchronous Point to Point Protocol (Xerox System Integration Standard) | Xerox Corporation | Drew D. Perkins (412)268-8576 |

Center for High Performance Computing

# The University of Texas System
# Center for High Performance Computing

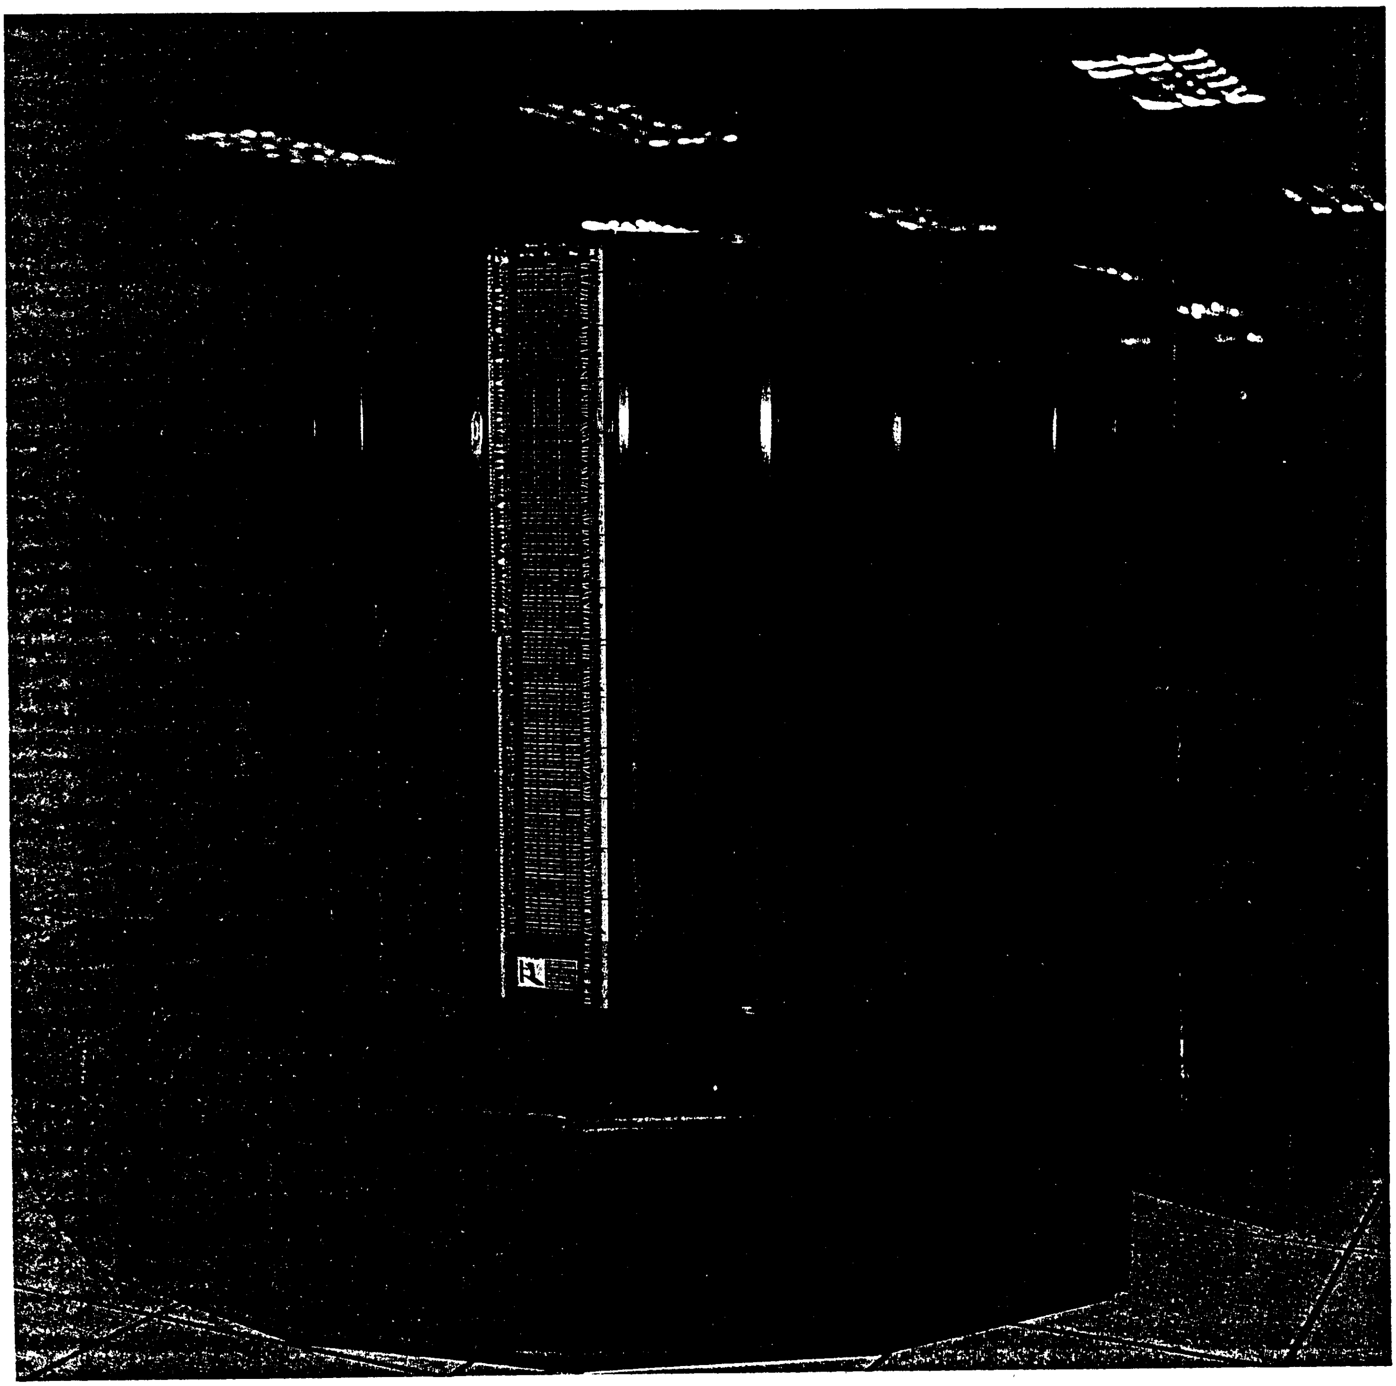

# Center for High Performance Computing

Few developments have so changed our world as the invention of the computer. But the capacity of early computers is trivial when compared with the power available today. A calculation that would have taken years can now be done in a few seconds. The challenge in the past was to find the means to answer complex questions. Now our calculating capacity enables us to answer those questions and to search out answers to questions previously believed to be beyond our understanding.

This colossal power is provided by supercomputers.

Through the Center for High Performance Computing, The University of Texas System supercomputing facility in Austin, brilliant minds and powerful machines are working together to answer questions on which advances in science, engineering, medicine, social sciences, and other fields depend.

The supercomputer and the intellectual environment supported by the Center are designed to bring out the most creative approaches to discovery through computing. Through collaborative efforts, researchers and students in the thirteen University of Texas System academic and health institutions share their disciplinary knowledge and computer skills to foster breakthrough research and enhance teaching.

At the Center scientists and computer professionals work with the best computing tools known at this time. Obstacles that often impede problem solving have been minimized. Academic users in all The University of Texas System institutions have convenient access to the supercomputing facility directly from their campuses. In such an environment, even the most gifted researchers can accomplish far more than they once thought possible.

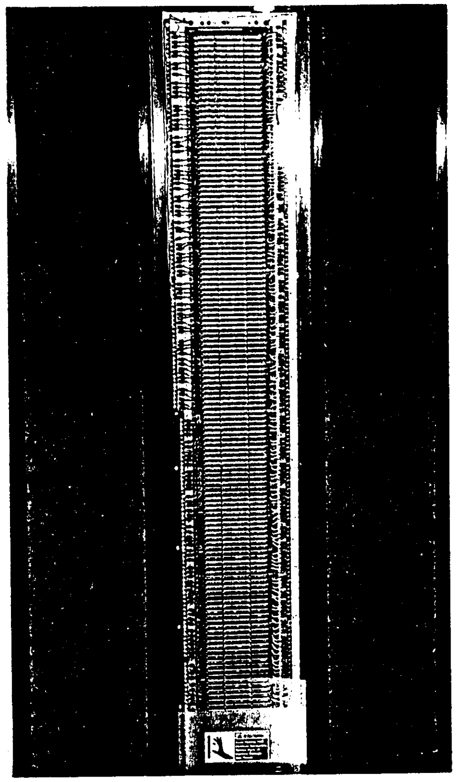

## The Supercomputer

The Center for High Performance Computing offers state-of-the-art performance in computing. The centerpiece of The University of Texas System facility is the Cray X-MP/24. This supercomputer is a two-processor machine with vector processors and a large amount of high-speed memory. It is ideal for complex calculations which involve large quantities of data.

The Cray X-MP/24 supercomputer is capable of sustaining 250 million computations per second, and each processor has a peak performance of 200 million floating point operations per second. The two processors can operate independently on separate jobs or jointly on a single job.

The supercomputer has a multi-layer memory system with a high capacity and fast retrieval. It contains a four million word main memory coupled with a thiry-two million word external fast memory through a high-speed (a billion bytes per second) channel. Lower-speed disk storage units can hold ten billion characters of information. A large capacity (trillions of bits) hierarchical mass storage facility holds user permanent files.

## Supercomputer System

The integrated supercomputer system, with the Cray X-MP as its centerpiece, includes two powerful front-end computers, a DEC VAX 8600 system and an IBM 4381 system. At each University of Texas System institution a DEC VAXstation II functions as a network interface processor for direct access to the supercomputer.

The VAX 8600 serves as a front-end communications computer for the supercomputer and handles communications between the Center for High Performance Computing and The University of Texas System campuses. It provides general timesharing services for preparing, submitting, and monitoring the progress of supercomputer jobs and for receiving program output from the supercomputer. The IBM 4381 provides the hierarchical file storage facility with over twenty billion bytes of disk storage and magnetic tape drives for tape cartridges and reels. The magnetic tape drives are also accessible to the supercomputer itself.

**Cray X-MP/24 Computer System Configuration**

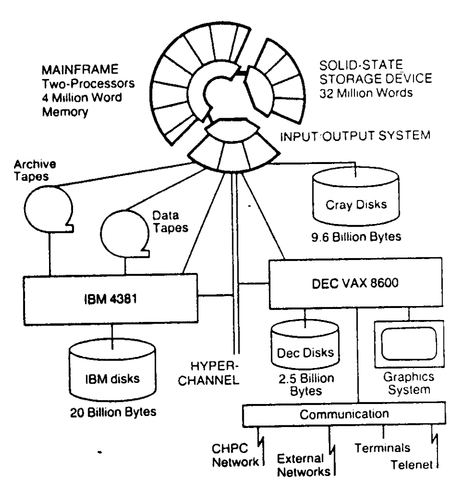


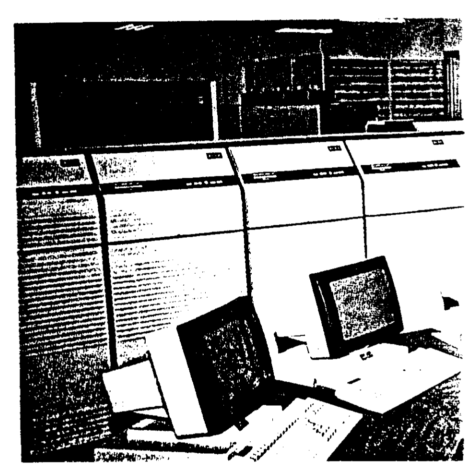

The University of Texas System
Center for High Performance Computing

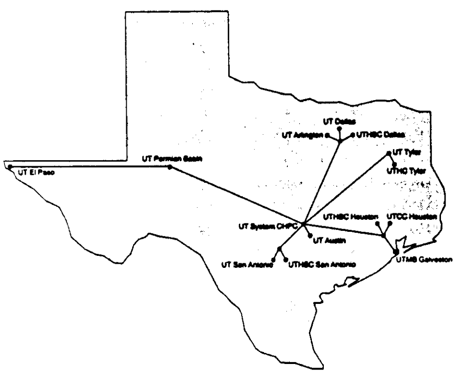


### The University of Texas System Computer Network

Faculty and students on The University of Texas System campuses communicate with the supercomputer and with each other across the state through an integrated and highly efficient network. Existing mainframe computers, minicomputers, and individual workstations on each campus tie into a local DEC VAXstation II, which, in turn, connects through the DECnet network to the Center's VAX 8600 computer to permit high-volume transfer of information and interactive communication with the Cray supercomputer. The Center's VAX 8600 computer may also be accessed from the national ARPANET research network and the international BITNET network.

### Graphics Capability

The supercomputer's speed and capacity to generate data make graphic representation of numerical output a necessity. The University of Texas System users can view the data graphically at their campus workstations. Equipment at the Center can handle large-volume user needs for graphics and includes laser printers, plotters, microfiche, 35mm slide and 16mm film production, and high-resolution video display and recording.

### User Support

Users receive expert advice on appropriate software tools and on the coding of research problems for optimal use of the supercomputer. Experienced computer professionals are on hand at the Center to provide short courses for users on such topics as the supercomputer system, vectorization, debugging, mathematical software, and using the supercomputer in a workstation environment. The Center staff also provides training for consultants at The University of Texas System institutions so they may, in turn, support users on their own campuses.

Consultations at the Center are available by telephone and electronic mail Monday through Friday to aid researchers in using the Center's facilities. Researchers are encouraged to visit the Center for training. Network and consulting staff are also available on a limited basis to visit The University of Texas System campuses.

### Executive and Advisory Committees

An Executive Committee oversees the administration of the Center for High Performance Computing for The University of Texas System. This five-member committee, appointed by the Office of the Chancellor, consists of three members from the academic institutions and two members from the health institutions.

A User Advisory Committee advises the Center on matters of policy and issues of importance to Center users, and an Operations Advisory Committee advises on matters relating to the operation of the Center and its interaction with academic computing facilities on The University of Texas System campuses. Each committee includes representatives from The University of Texas System academic and health institutions.

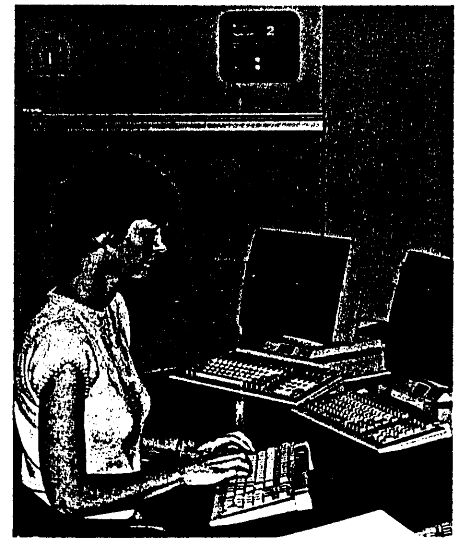

### Balcones Research Center

The Center for High Performance Computing is located in the Commons Building at the Balcones Research Center in Austin. The Center's computer room contains 3,500 square feet of floor space and houses all of the supercomputer system. Staff offices, workstations, reference materials, consultation, and input/output services are adjacent and are available to Center users. The Commons Building is particularly well-designed to house the Center, with an auditorium which seats 330 people, a conference facility which seats 140 people, a spacious cafeteria, and other support facilities.

**Center for High Performance Computing**
Commons Building
Balcones Research Center
Austin, Texas 78758-4497
(Area code 512) 471-2472 (or, 471-CHPC)

## Supercomputer Research

Supercomputers speed research on a wide spectrum of complex problems. The spectrum encompasses genetics, artificial intelligence, molecular biology, and neuropsychology, as well as structural engineering, materials science, geophysics, and computer engineering.

The powerful calculation abilities of the Center's supercomputer are expected to foster new research and development programs on The University of Texas System campuses in such areas as:

*Chemistry*
Molecular modeling, quantum chemistry, crystallography, chemical kinetics.

*Physics and Astronomy*
Astrophysics, plasma physics, magnetohydrodynamics, nuclear dynamics, galaxy formation, supernova structure.

*Health Sciences*
Biomathematics, magnetic resonance imaging, cancer cell modeling, DNA physical and biological properties, drug modeling, neural circuit analysis.

*Engineering*
Image and signal analysis, finite element analysis of structures, semiconductor materials, ocean mapping, orbital mechanics, VLSI design, turbulent fluid flow, robotics, biomedical engineering.

*Geophysics*
Enhanced oil and gas recovery, plate tectonics, physical oceanography, seismic analysis.

*Mathematics and Computer Sciences*
Algorithm development, parallel processing, artificial intelligence, high-resolution graphics, highly parallel computer architectures, network communications.

*Economics and Business*
Micro- and macro-economic modeling, operations research, financial market analysis.

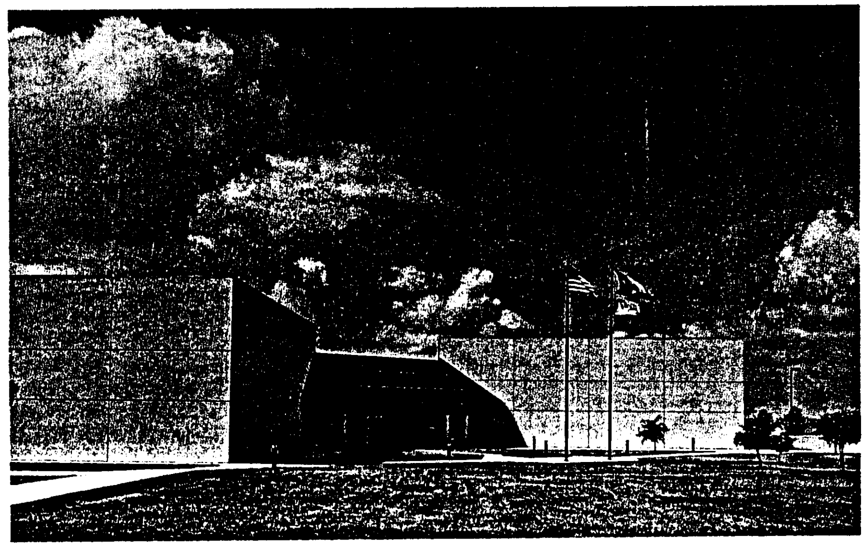

# The University of Texas System

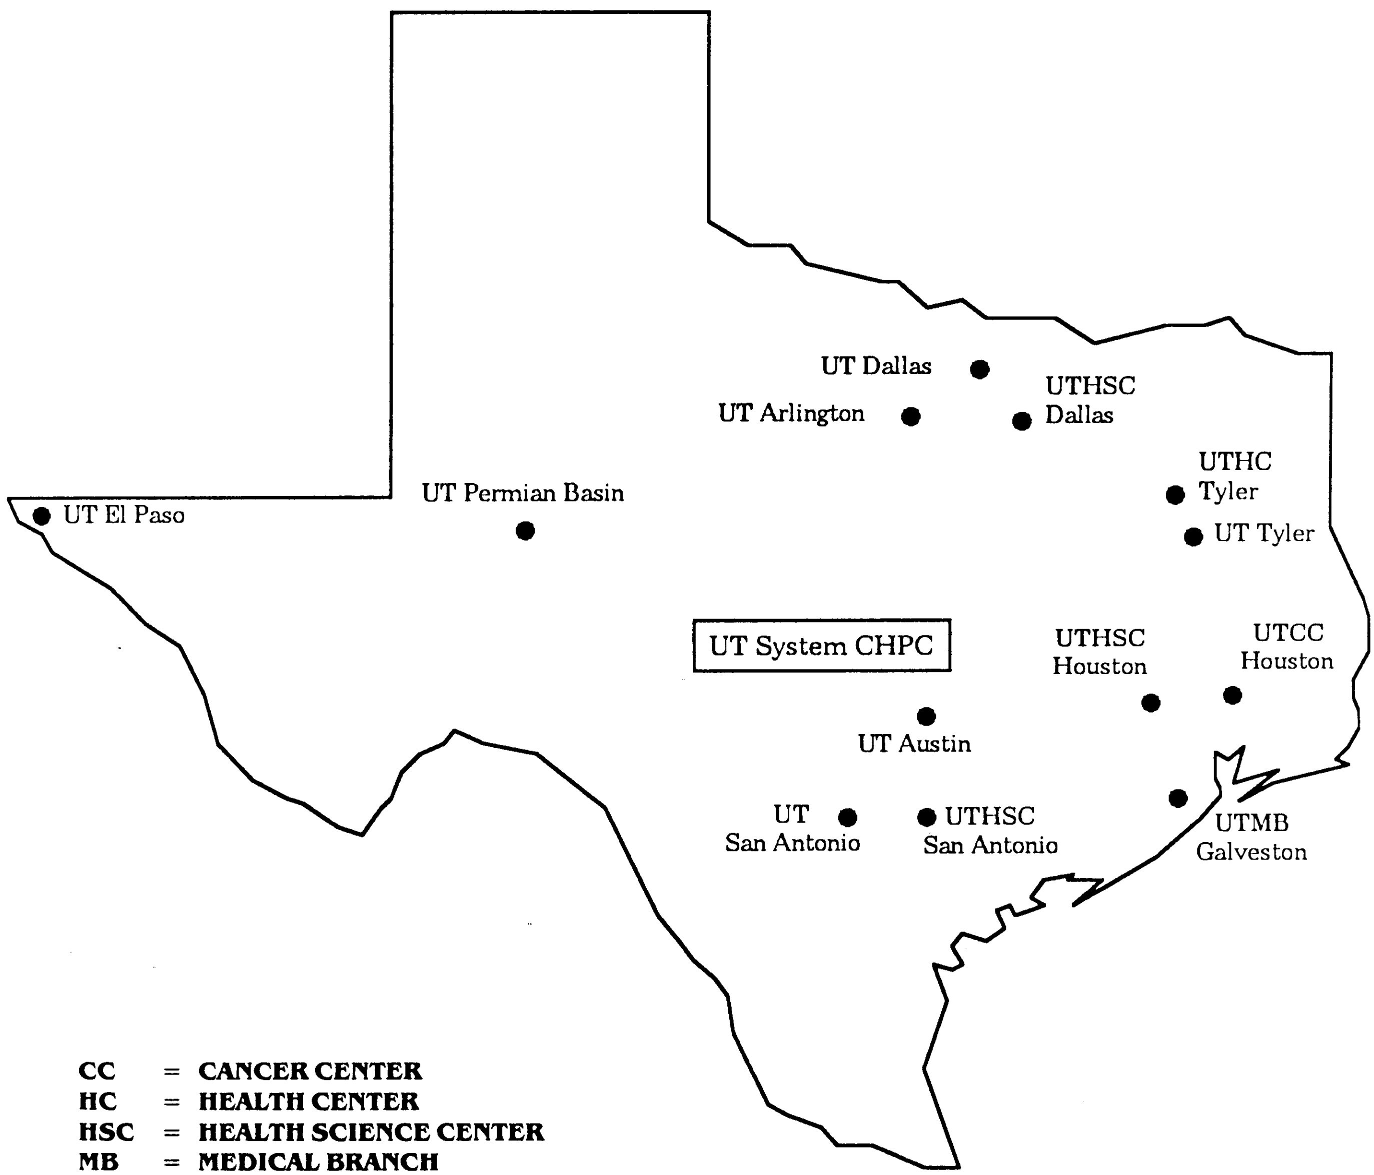


**CC** = **CANCER CENTER**
**HC** = **HEALTH CENTER**
**HSC** = **HEALTH SCIENCE CENTER**
**MB** = **MEDICAL BRANCH**
**CHPC** = **CENTER FOR HIGH PERFORMANCE COMPUTING**

CHPC Newsletter

The University of Texas System

Center for High Performance Computing

Vol. 4, No. 3 28 November 1988

# COS to UNICOS Migration Tools

R. Harkness

The new Cray X-MP/14se arrived on 21 November 1988 as scheduled and is undergoing acceptance testing and software installation in preparation for its role as a platform for migration of the CHPC facility from COS to UNICOS. Although the new machine will not be immediately available to the general user community, you may want to begin taking a closer look at the software available under COS that will help in the migration to UNICOS.

Cray Research has provided a set of software tools to help users move applications from COS to UNICOS systems. This note describes features available now on the COS system. Some of the tools run only under COS, or only under UNICOS, and some under both operating systems. A description of the UNICOS-specific tools will follow in a subsequent article.

Those users planning to move applications to the new system will need to convert COS JCL to UNICOS shell scripts, convert COS datasets to UNICOS file format, and in some cases, convert programs, since not all COS library software is available under UNICOS. Some aspects of the migration process are addressed below.

## Program Migration

Numerous issues are associated with the differences between the COS and UNICOS Fortran environments. A detailed description of these issues will be ready before the Cray X-MP/14se system is available for general use. Both the CFT and CFT77 compilers will be available under UNICOS, but the original COS loader, LDR, will not be available under UNICOS. Switching from LDR to SEGLDR is likely to be a trivial exercise for most users, except for those with codes requiring overlays. Users are advised to convert to SEGLDR on both operating systems.

[*continued on page 46*]

## Contents

**LD2** Current LDR users may find the new LD2 command useful. LD2 has the same arguments as LDR, but instead calls SEGLDR to generate the directives necessary to perform the same function as LDR. The LD2 command is fully described in the *COS Reference Manual* (SR-0011, Revision O).

Consider the following (non-overlay) LD2 statement:

```
LD2,LLD,LIB=OMF,MAP=PART,C=OFF,NX,AB=EXEC.
```

The extra argument LLD is required to save the SEGLDR command and directives (and any CAL routines generated for applications using OVERLAYs). In this example, the relevant local datasets are $ILDR and $DLDR. The $ILDR dataset contains the new SEGLDR command line, which in this case is simply:

```
SEGLDR(I=$DLDR)
```

and the $DLDR directive file contains:

```
TITLE=SEGLDR substituted for LDR
*LD2 version 0.00  (09/27/88 14:15:33)
ABS=EXEC
MAP=PART
LIB=OMF
ALIGN=MODULES
```

Thus, LD2 comes as close as possible to the original intent of the equivalent LDR statement. (Note that LLD outputs the SEGLDR command line using parentheses instead of the comma and period convention).

**DMYLIB** Programs that run under COS with CFT77 and SEGLDR are unlikely to need modifications to run correctly under UNICOS, but there are exceptions: some COS library functions have no UNICOS equivalents. The DMYLIB library can be used to spot these COS dependencies.

Use DMYLIB as follows:

```
JOB,JN=TEST,US=ABCD123,RT=99,P=1.
ACCOUNT,UPW=TESTING.
FETCH,DN=PROGRAM,TEXT='.....'.
CFT,I=PROGRAM,L=0.            ! Compile
LIB,DN=NAG.                   ! Access any non-Cray libraries
LIB,DN=DMYLIB.                ! Access the DMYLIB package
SEGLDR,CMD='TRIAL;NODEFLIB;LIB=DMYLIB,NAG'.
                              ! Load excluding COS libraries
```

The SEGLDR output will list unsatisfied external references if any required COS library subroutines are not available under UNICOS. The NODEFLIB directive instructs SEGLDR not to search the usual COS system libraries such as $UTLIB, $SCILIB, etc. Note that execution following the load is not possible; thus, the TRIAL directive is used for efficiency. A typical output is shown below:

```
TEST      SEGMENT LOADER      VERSION 3.1

     UNSATISFIED EXTERNAL REFERENCE SUMMARY

     ENTRY NAME      MODULES REFERENCING ENTRY

       ASSIGN          TEST
       GETPARAM        TEST
       MEMORY          TEST
       SAVE            TEST
```

Clearly, any codes that manipulate COS datasets directly will require modification. Also, any codes that use dynamic common blocks will need to be adapted to use the heap manager, as explained below.

GETPARAM is a very useful Fortran subroutine that evaluates parameters passed in a COS control statement. Since this routine is COS-specific, it will not be available under UNICOS, but a GETPARAM migration routine will be part of the UNICOS migration tools package installed on the Cray X-MP/14se.

**Heap Manager**

User programs that currently use overindexing of dynamic common blocks should be modified to make use of the heap manager, as this approach is common to all Cray systems. As shown in the previous example, the MEMORY subroutine is not available under UNICOS. However, the heap manager routines are identical on both COS and UNICOS systems. For example:

```
              COS                  |          COS OR UNICOS
      EXPANDING BLANK COMMON       |        EXPANDING THE HEAP
   ----------------------------------------------------------------
                                   |
      SEGLDR DIRECTIVE:            |    NO SEGLDR DIRECTIVES
         DYNAMIC = WORK            |          NECESSARY
   ----------------------------------------------------------------
                                   |
       PROGRAM TEST                |     PROGRAM TEST
                                   |
       COMMON /WORK/ X(1)          |     COMMON /WORK/ IPTR
                                   |     POINTER ( IPTR , X(1) )
                                   |
       CALL MEMORY('UC',100000)    |     CALL HPALLOC(IPTR,100000,
                                   |    -             ERR,ABORT)
                                   |
       DO 100 I=1,100000           |     DO 100 I=1,100000
       X(I) = RANF()               |     X(I) = RANF()
   100 CONTINUE                    | 100 CONTINUE
       END                         |     END
```

The pointer declaration and the common block WORK are required in all subroutines that reference the dynamic array.

**COS JCL Conversion**

The COS JCL to UNICOS shell command conversion package will obviously become more important when the new system is available. However, users may want to get familiar now with the format of UNICOS job control. In addition, the conversion package can point out the commands that have no UNICOS equivalent. This can help you plan strategies to avoid COS dependence.

Under UNICOS, batch jobs are managed by NQS (Network Queueing System). The command conversion package normally produces output suitable for submission to NQS. A detailed article on NQS is being prepared.

There are four utilities in the conversion package:

- *cjcl* Reads a COS JCL file and converts it to UNICOS shell commands.
- *cjproc* Preprocesses COS PROCs for input to *cjcl*.
- *cjplib* Preprocesses COS PROC libraries for input to *cjcl*.
- *cjprnt* Reformats a COS JCL file for easy reading.

All of these procedures require a "Unix-style" set of arguments to be specified as a string in COS JCL. Examples are shown below. Note that these utilities cannot convert all occurrences of JCL, and they are of little use in converting very simple JCL files. To produce usable UNICOS job files may require some hand coding to reflect local CHPC requirements and to optimize the output. Also, at the present time, not all of the options can be exercised under COS. These commands will also be available under UNICOS.

**cjcl**

```
cjcl  [-d pdir] [-f flags] [-1 pass1] [-2 pass2]
      [-i input] [-o output]

   -d pdir     Name of the directory in which to store
               converted PROCs.   Default PROCLIB.

   -f flags    Control flags:  lowercase to enable,
                               uppercase to disable.

          b    Generate NQS (batch) directives
          c    Include original COS statements as comments
          e    Include error-checking
          j    Copy JOB and ACCOUNT statements
          r    Defines COS variables
          s    Include shell control flags
          v    Identify cjcl and cjproc version

   -1 pass1    Name of template file for pass 1:  default PASS1
   -2 pass2    Name of template file for pass 2:  default PASS2
   -i input    Input file: default $IN (redirect using <)
   -o output   Output file: default $OUT (redirect using >)
```

If the COS JCL file contains any PROCs, SDATA, or data (e.g., Fortran, data, or SEGLDR directives) following /EOF, the file must be preprocessed using *cjproc* before calling *cjcl.*

**cjproc**

```
cjproc  [-s] [-v] [-p newlib] [proclib1] [proclib2]

     -s           Suppress statistics
     -p newlib    Create a cjproc library
     -v           Output detailed information
     proclib1     Additional procedure library: default none
     proclib2     Additional procedure library: default none
```

*cjproc* reformats PROC definitions, in-line data, etc.

**cjplib**

*cjplib* reads a COS JCL file from *stdin* and writes the processed file to *stdout* ($IN and $OUT, respectively), unless redirected. *cjplib* modifies a COS JCL file to produce acceptable PROC definitions for the *cjproc* utility.

**cjproc / cjcl Example**

The following example illustrates the kind of output you can expect for a simple COS batch job. The first section is the job to run *cjcl:*

```
JOB,JN=UNICOS,US=ABCD123,RT=99,P=1.
ACCOUNT,UPW=PASSWORD.
FETCH,DN=COSJCL,TEXT='cosjclfile'.
LIB,DN=PASS1.               ! Template file for PASS1
LIB,DN=PASS2.               ! Template file for PASS2
LIB,DN=CJCL.                ! Access cjcl command
LIB,DN=CJPROC.              ! Access cjproc command
CJPROC,'<COSJCL >TEMP'.     ! Run COS JCL through pre-processor
REWIND,DN=TEMP.             !
CJCL,'<TEMP >UNIX'.         ! Run CJCL, output to file UNIX
```

```
REWIND,DN=UNIX.             !
COPYF,I=UNIX,S=2.           ! Copy to $OUT and shift 2 cols
```

The following is the simple COS JCL file to be translated to UNICOS:

```
JOB,JN=LINP,US=XXAL604,RT=99,MFL=3600000.
ACCOUNT,UPW=PWD.
DISPOSE,DN=$OUT,DEFER,TEXT='TEMP:[XXAL604]'.
FETCH,DN=CODE,TEXT='TEST.CFT'.
CFT77,I=CODE,OFF=P,ON=F.
ASSIGN,DN=CHECK,A=FT08.
SEGLDR,GO,CMD='MAP=ADDRESS'.
DISPOSE,DN=CHECK,DC=ST,TEXT='TEMP:[XXAL604]CHECK.OUT'.
EXIT.
```

The output of *cjcl* is quite complex. The first lines beginning with # contain NQS directives for use if the job is to be run in batch mode. If the job is executed interactively, these lines appear as comments. Also, for UNICOS 4.0 and above, the NQS directive "@$" will become "QSUB".

```
cjproc: version 1.1  13 November, 1987
cjproc: 9 input lines
cjproc: 80 characters allocated for COS line
cjcl: version 1.0  13 November, 1987
cjcl: flags = <bceJRsV>

 # user=xxal604 pw=pwd
 # @$-s /bin/sh        # select Bourne shell
 # @$-r linp           # Job name
 # @$-eo               # combine stderr & stdout
 # @$-lM 3600000       # MFL
 #
 #   Set shell control flags
 #
 set -xu
 #
 #   The COS $OUT file is translated to stdout.  This assumes
 #   the job is to be run under NQS to capture the stdout file.
 #
 #   COS simulated file system is as follows for permanent files.
 #       $HOME/COSPFS/ID/PDN/ed              (default OWN)
 #       $HOME/(OWN)/ID/PDN/ed               (explicit owner)
 #   COS simulated local (temporary) files, this directory
 #       is removed.
 #          $HOME/JOBTMP{process number}/dn
 #
 #
 #   Path for files with implicit owner - redefine if needed:
 #
 COSPFS=${HOME}/COSPFS export COSPFS
 #
 #   Path for PROC library - redefine if needed:
 #
 PROCLIB=$HOME/PROCLIB export PROCLIB
 #
 #
 #   Path for simulated COS local files:
 #   The current process number is used as part of the directory
```

```
#   name and the directory is created if it does not exist.
#
JOBTMP=${HOME}/JOBTMP$$ export JOBTMP
mkdir ${JOBTMP};cd ${JOBTMP}
while test 0 -eq 0
# 1      JOB,JN=LINP,US=XXAL604,RT=99,MFL=3600000.
COSLN=1
# 2      ACCOUNT,UPW=PWD.
COSLN=2
# 3      DISPOSE,DN=$OUT,DEFER,TEXT='TEMP:[XXAL604]'.
COSLN=3
dispose ${JOBTMP}/$OUT -dPR -tTEMP:[XXAL604] || break
# 4      FETCH,DN=CODE,TEXT='TEST.CFT'.
COSLN=4
fetch ${JOBTMP}/CODE -tTEST.CFT || break
# 5      CFT77,I=CODE,OFF=P,ON=F.
COSLN=5
ln ${JOBTMP}/CODE ${JOBTMP}/CODE.f || break
echo 'CFT77: default list is OFF'
cft77  -b ${JOBTMP}/$BLD.o -ef -dp ${JOBTMP}/CODE.f || break
rm ${JOBTMP}/CODE.f
# 6      ASSIGN,DN=CHECK,A=FT08.
COSLN=6
echo 'ASSIGN: partial support only'
env $FILENV assign -a ${JOBTMP}/CHECK fort.08 || break
# 7      SEGLDR,GO,CMD='MAP=ADDRESS'.
COSLN=7
echo 'SEGLDR: default support only'
echo 'map=ADDRESS'>>dir
segldr -o ${JOBTMP}/$ABD ${JOBTMP}/$BLD.o dir || break
${JOBTMP}/$ABD || break
# 8      DISPOSE,DN=CHECK,DC=ST,TEXT='TEMP:[XXAL604]CHECK.OUT'.
COSLN=8
dispose ${JOBTMP}/CHECK -dST -tTEMP:[XXAL604]CHECK.OUT || break
# 9      EXIT.
cd $HOME;rm -fr ${JOBTMP}
exit
echo "EXIT processing entered at COS line $COSLN" >&2
echo 'THIS IS THE END OF THE JOB'
cd $HOME;rm -fr ${JOBTMP}
exit
echo "Break at COS line number $COSLN"
cd $HOME;rm -fr ${JOBTMP}
exit
```

*cjcl* performs a literal interpretation of the original COS job and includes error checking. Obviously, the output is much more cumbersome than it needs to be for such a simple task, but *cjcl* could be of some use to UNICOS novices with large and complex COS JCL files.

**File Migration**

The file format conversions run only under UNICOS and will be described as soon as the new system comes on line.

## Holiday Schedule

D. Nobles

The University of Texas Christmas Holiday schedule is as follows:

5 P.M. Friday, 23 December 1988 to 8 A.M. Tuesday, 3 January 1989

The CHPC consulting, systems, and administrative staff will be generally unavailable during this period, except for periodic perusal of the REMARK bulletin board and personal mail. The CHPC will have full operator coverage throughout the holidays, and the operators can be reached at (512) 471-2472 or (800) 262-2472 (toll-free within Texas).

---

## UNICOS Presentation at UTMB, Galveston

J. Werst

Bill Eue from Cray Research will present a two-hour "Introduction to UNICOS" from 10 A.M. to noon on Wednesday, 14 December 1988 in room 449 of the Administration Building at the UT Medical Branch, Galveston.

For more information, call Cecil Denney at (409) 761-1813, Office of Academic Computing, UTMB, Galveston.

This presentation is available to other UT System component institutions if prior arrangements are made through the User Advisory Committee representative listed on the directory page in this *Newsletter*.

---

## Cray Software: CFT77 3.0 Upgrade

J. Respess

On Monday, 5 December 1988, CFT77, Release 3.0, will replace Release 2.0 as the default CFT77 compiler on the Cray X-MP/24. CFT77 3.0 has many new features, including:

- Debugging support for optimized code,
- In-line subroutine expansion to improve performance,
- New intrinsic functions that provide additional capabilities for bit manipulation,
- Enhanced RANSET and RANGET intrinsic functions to improve usability and compatibility with the CFT compiler.

Until 5 December, users can experiment with the new version of CFT77 by using SELECT to specify this version. Use the command:

```
SELECT,PROD=CFT77,VERS=NEW.
```

Documentation for CFT77, Release 3.0 is contained in the Cray Research publication, *CFT77 Reference Manual*, (SR-0018, Revision C). For more information, type the DCL commands:

```
HELP @COS CFT77
HELP @COS SELECT
HELP CHPC MANUAL_ORDERS
```

---

## Cray Software: Pascal 4.0 Upgrade

A. Kochis

Effective 5 December 1988, Pascal 4.0 will become the default compiler on the Cray X-MP/24. This new release fixes known significant problems in earlier versions and contains the following enhancements:

- New predefined functions MIN, MAX, LDZERO, FIRST, and LAST.
- Variable-length string processing package. (Strings may be assigned, passed as parameters, read from and written to TEXT files, and manipulated with several new predefined functions: LENGTH, MAXLENGTH, SUBSTR, DELETE, INDEX, TRIM, LTRIM and COMPRESS.)

- Extended function result types.
- Improved interlanguage support.
- Larger sets. (The limit on the number of elements in a set has been increased from 128 to 4,096. Vector code is generated to handle sets with more than 128 elements.)
- Improved vectorization.
- INCLUDE facility.
- Support for large data structures.

These new features are described in the Cray Research publication, *Pascal Reference Manual* (SR-0060, Revision E).

---

## Cray Software: GRADSCF Installed

K. Milfeld

GRADSCF has been installed on the Cray X-MP/24 and will be available for general use on 5 December 1988.

GRADSCF is an *ab initio* gradient program from Polyatomics Research Institute in Mountain View, California. The program calculates SCF (Self-Consistent Field) and MC-SCF (Multi-Configuration) wave functions and energies for atoms or molecules. Post-SCF refinements to the energy can be determined by several methods: MP2, second-order Moller-Plesset perturbation; GVB (generalized valence bond); and pairwise MC-SCF.

External electric fields can be included in the SCF calculations and some one-electron properties are determined from the wave function. These include Mulliken populations, dipole and quadrupole moments, and polarizability tensors. Location of minimum energy structures (often referred to as geometric optimization) is available for automatically determining the geometric configuration of ground states, excited states, and transition states. Either numerical or analytic second-derivative methods can be used to determine frequencies and molecular force constants. With analytic derivatives, analytic IR (infrared) intensities can be obtained through atomic polar tensors.

The wave function forms are closed shell, restricted and unrestricted open shell, GVB (generalized valence bond) perfect pairing, and pairwise MC-SCF. Gradients of the energy are available for all wave functions, and second derivatives are limited to closed-shell SCF wave functions. Both derivative algorithms are highly vectorized. Geometry optimization is allowed in either internal or Cartesian coordinates. Symmetry group techniques for reducing computational effort have not been incorporated into GRADSCF.

Documentation for GRADSCF has been distributed to the Operations Advisory Committee member for each UT System component institution. For information on execution, notable comments, and references, type the DCL command:

HELP @CRAY_SOFTWARE GRADSCF

---

## VAX Software: VMS Cray Station 4.01

S. Kneuper

On Monday, 12 December 1988, the new release 4.01 of the VMS Cray Station software will be available on the VAX 8600 front-end system. This software enables job submission and interactive access to the Cray X-MP/24.

Release 4.01 is only slightly different from the previous version, 3.06. The new VMS Cray Station offers interactive access using the CINT program only, whereas version 3.06 offered interactive access using both CINT and the Cray "INT" command. Batch access remains the same. Some displays are different and convey more information; for example, the TAPE command displays more than 9999 blocks correctly.

The CINT program has a different look and feel compared to the old "CRAY INTER" command. For example, a <control-c> will ask the Cray to abort the current command, and <control-d> will send an end-of-file indication to the Cray.

For more information about the commands available from CINT, you can use the internal CINT help facility or type "HELP CINT" at the VMS prompt.

---

## VAX Software: New DI-3000 Metafile Translator

B. Chauvin

On 5 December 1988, a new DI-3000 Metafile Translator will be installed on the VAX 8600 front-end system. This translator contains the ability to interpret and create CGM (Computer Graphics Metafile) or DI-3000 metafiles.

Several new device drivers have also been installed with the new metafile translator. Below is a list of the device drivers now available:

| | |
|---|---|
| CRT | "dumb" terminal |
| ERG | Micro-Term Ergo |
| GIG | DEC Gigi |
| LPR | line printer |
| MVP | Matrix MVP Rasterizer for Matrix film recorder |
| PTX | Printronix plotter |
| PST | PostScript driver |
| T10 | Tektronix 4010 |
| 405 | Tektronix 4105 |
| 407 | Tektronix 4107/4109 |
| 415 | Tektronix 4115 |
| TL8 | Talaris 800/1200/2400 |
| 125 | VT 125 |
| VTK | Versatec V-80 |
| ZTA | Nicolet Zeta plotters |
| RJZ | UT Austin Computation Center Zeta plotters |

For more information, type the DCL command:

HELP @CRAY_SOFTWARE DI-3000 Metafile_Translator

---

## Cray Disk Storage Rate Reduced

L. Keeler

Due to the increased volume of archived dataset storage, the rate charged for permanent dataset storage on the Cray X-MP/24 has been reduced from \$.08 to \$.04 per megabyte per day. This change actually went into effect on the first day of the billing period beginning 26 October 1988. The change should significantly reduce the costs to users who have large storage requirements. The disk charges for the VAX 8600 will remain at \$.30 per megabyte per day.

For more information about charges for CHPC services, type the DCL command:

HELP RATES

Commands for transferring datasets to Cray archival storage are also described in HELP files, which may be viewed by typing the following DCL commands:

| | |
|---|---|
| HELP CHPC POLICY DISK | HELP CRAUDIT |
| HELP @COS SAVE | HELP CRPUT |
| HELP @COS ACCESS | HELP CRGET |

---

## Updated THEnet Directory

S. Barron

The *Users' Directory of Computer Networks* accessible to the Texas Higher Education Network (THEnet) member institutions was updated in July, 1988 by the Office of Telecommunication Services and is now available on-line on the CHPC VAX 8600. You can copy, view, or print individual sections from this on-line manual by calling the DOCUMENT utility at the VMS prompt. DOCUMENT is menu-driven and will prompt you appropriately, or you can request items directly. For more information about DOCUMENT, type the DCL command:

```
HELP DOCUMENT
```

As previously, the THEnet manual is broken into 6 sections. For example, the first section which contains the cover, title page, introduction, and table of contents can be viewed by typing:

```
DOCUMENT OTSNETDIR00
```

A hard copy of this document may be obtained from The UT System Office of Telecommunication Services, Balcones Research Center, 10100 Burnet Road, Austin, TX 78758-4497. The price is $15.00.

## Using Pascal Modules from Fortran

A. Kochis

Every so often it becomes necessary to incorporate a program module written in a different language into your code. The Pascal, CFT, CFT77, and C manuals all allude to being able to accomplish this task, with certain restrictions. The problem with the manuals is that they do not present any examples. To correct this oversight, I have developed two examples of the following task:

A Fortran program (FTEST) calls a Pascal module (ptest) to initialize a two-dimensional array by calling a Fortran function (SECOND) which is defined in the standard Fortran Library.

In the first example, the array is passed in a named common block; in the second example, the array is passed as a parameter. The numbers in the right margin refer to notes explaining the inter-language connection.

**Example 1:** **Using a common block to pass the data.**

Pascal Code Notes

```
MODULE ptest;                                                1

FUNCTION second : REAL ; FORTRAN;                            2

PROCEDURE init;  EXPORTED;                                   3
COMMON  x : ARRAY[1..20,1..10] OF INTEGER;                   4
VAR i,j : INTEGER;
BEGIN
   FOR i := 1 TO 20 DO
     FOR j := 1 TO 10 DO
         x[i,j] := ROUND(second);                            5
END;
```

| Fortran Code | Notes |
|---|---|
| ```      PROGRAM FTEST``` | |
| ```      INTEGER Y(10,20)``` | 4 |
| ```      COMMON /X/ Y``` | 6 |
| ```      CALL INIT``` | 7 |

```
      PROGRAM FTEST
      INTEGER Y(10,20)
      COMMON /X/ Y
      CALL INIT
      DO 40 I = 1,10
        DO 20 J = 1,20
           PRINT 10,Y(I,J)
10         FORMAT(' ',I3)
20      CONTINUE
        PRINT 30
30      FORMAT('-----')
40    CONTINUE
      END
```

**Example 2:** **Using parameters to pass data.**

| Pascal Code | Notes |
|---|---|
| `PROCEDURE init(VAR x : mat);  EXPORTED;` | 8 |

```
MODULE ptest;
TYPE mat = ARRAY[1..20,1..10] OF INTEGER;
FUNCTION second : REAL ; FORTRAN;
PROCEDURE init(VAR x : mat);  EXPORTED;
VAR i,j : INTEGER;
BEGIN
   FOR i := 1 TO 20 DO
     FOR j := 1 TO 10 DO
         x[i,j] := ROUND(second);
END;
.
```

| Fortran Code | Notes |
|---|---|

```
      PROGRAM FTEST
      INTEGER Y(10,20)
      CALL INIT(Y)
      DO 40 I = 1,10
        DO 20 J = 1,20
           PRINT 10,Y(I,J)
10         FORMAT(' ',I3)
20      CONTINUE
        PRINT 30
30      FORMAT('-----')
40    CONTINUE
      END
```

**Example JCL**

| JCL to bind Pascal and Fortran together. | Notes |
|---|---|
| `JOB,JN=FTOP,US=ABCD123,RT=9.` | |
| `ACCOUNT,UPW=zotz.` | |
| `PASCAL,B=PC.` | 9 |
| `CFT77,B=FC,ON=S.` | 10 |
| `SEGLDR,GO,CMD='MAP=full;BIN=PC,FC'.` | 11 |
| `/EOF` | |
| `Pascal code` | |
| `/EOF` | |
| `Fortran code` | |
| `/EOF` | |

**Notes**

1. Define a Pascal module without defining a program entry point for SEGLDR. This is described in Chapter 12 of the *Pascal Reference Manual.*
2. Refer to the function SECOND in the *Programmer's Library Reference Manual.* The FORTRAN directive is defined in Chapter 9 of the *Pascal Reference Manual.*
3. The procedure must use the EXPORTED directive, defined in Chapter 9 of the *Pascal Reference Manual*, to create a load point for SEGLDR.
4. The two-dimensional data structure. Note that Pascal and Fortran define arrays in the opposite order. To Fortran, this is a 10 x 20 array, while Pascal sees it as a 20 x 10 array.
5. After definition, the Fortran function can be used the same way a Pascal function is used.
6. Data to be passed must be in a *named* common block. The name of the block is the same as the data name, but uppercase, in Pascal.
7. The Pascal procedure is then called from Fortran.
8. All parameters to be passed must be called by name (denoted by the VAR in the parameter definition).
9. Compile Pascal module, writing binary load module to temporary dataset, PC.
10. Compile Fortran program, writing binary load module to temporary dataset, FC.
11. Load and execute the Pascal and Fortran binary modules using SEGLDR.

**Reference Manuals**

*Pascal Reference Manual* (SR-0060)
Chapter 12 describes Modules.
Chapter 9 defines the Procedure directives EXPORTED and FORTRAN.

*CFT77 Reference Manual* (SR-0018)
Appendix F explains how to utilize non-Fortran routines.

*CFT Reference Manual* (SR-0009)
Appendix F explains how to utilize non-Fortran routines.

*Programmer's Library Reference Manual* (SR-0113)
Chapter 15 on the timing routines describes the Fortran routine SECOND.

*Segment Loader (SEGLDR) Reference Manual* (SR-0066)

## Using Commercial Software Packages at CHPC

L. Keeler

*"Respect for intellectual labor and creativity is vital to academic discourse and enterprise. This principle applies to works of all authors and publishers in all media. It encompasses respect for the right to acknowledgment, right to privacy, and right to determine the form, manner, and terms of publication and distribution.*

*"Because electronic information is volatile and easily reproduced, respect for the work and personal expressions of others is especially critical in computer environments. Violation of authorial integrity, including plagiarism, invasion of privacy, unauthorized access, and trade secret and copyright violations, may be grounds for sanctions against members of the academic community."*

— from *Using Software, A Guide to the Ethical and Legal Use of Software for Members of the Academic Community,* EDUCOM & ADAPSO, 1987.

Many of the application software packages available on CHPC computer systems are proprietary products licensed from commercial software vendors. All of our commercial vendors recognize the academic nature of the computational research conducted with CHPC resources and provide their products to us at substantially discounted rates. In turn, The University of Texas agrees that the CHPC user community will comply with the terms of the agreements in the academic license.

The requirements of these terms are normally published with the installation announcement of the software in this *Newsletter*. For the convenience of our users, we have summarized the requirements of our current commercial software agreements below.

### COMMERCIAL VENDORS, PRODUCTS, AND SPECIAL REQUIREMENTS

**BCSLIB**
**BCS VectorPak**

Boeing Computer Services
Software and Education
Products Group
P.O. Box 24346
Mail Stop 7K-10
Seattle, WA 98124-0346

No special requirements, but see "Acknowledgments and CHPC Requirements" below.

**Gaussian 86**
**Gaussian 82**

Gaussian 86
Department of Chemistry
Carnegie-Mellon University
Attn: J. A. Pople
4400 Fifth Avenue
Pittsburg, PA 15213

If you intend to publish a result obtained through the use of this software, the license requires an appropriate citation that includes the name of the product ("Gaussian 82, Release A" or "Gaussian 86, Release C"), the source ("Carnegie-Mellon University"), and the authorship ("Michael Frisch, Stephen Binkley, H. Bernard Schlegel, Krishnan Raghavachari, Richard Martin, James J. P. Stewart, Frank Bobrowicz, Douglas DeFrees, Rolf Seeger, Robert Whiteside, Douglas Fox, Eugene Fluder and J. A. Pople").

**ABAQUS**

Hibbitt, Karlsson, & Sorensen, Inc.
100 Medway Street
Providence, RI 02906

This proprietary package is licensed to the UT System for research and academic work only. Users external to the UT System will be denied access to ABAQUS.

You are required to provide suitable acknowledgment of HKS's provision of ABAQUS in all publications and reports that use results generated with the program, and to forward one copy of any such report or publication to HKS, provided that such a report is not private.

"Commercial usage" of the program requires the user to notify CHPC and arrange to pay HKS a monthly surcharge based on payments received by the user or his institution on account of the work. "Commercial usage" is any use of the program for a commercial purpose or to produce calculations having commercial value, including any project or calculation of a type performed in the normal course of a business or practice.

**IMSL Library PDE/PROTRAN**

IMSL, Inc.
2500 ParkWest Tower One
2500 CityWest Boulevard
Houston, TX 77042-3020

Government laboratories and government research facilities located at educational institutions do not qualify to use IMSL under our educational license agreement.

Users are encouraged to contribute articles and reports to the IMSL quarterly newsletter, *DIRECTIONS*.

**MSC/NASTRAN**

MacNeal-Schwendler Corporation
815 Colorado Boulevard
Los Angeles, CA 90041-1777

Users engaged in sponsored research are expected to obtain funding for MSC/NASTRAN CPU time from the sponsoring agencies. CHPC users with sponsored research funds, or users external to the UT System, are surcharged and billed for the use of this software as follows:

| | |
|---|---|
| UT System users with sponsored research funds | $10.78 / CPU minute |
| External users | 35.93 / CPU minute |
| Other users | No surcharge |

The use of MSC/NASTRAN in consulting activities or commercial purposes is specifically forbidden. All requests for assistance with NASTRAN must be directed to the CHPC.

**NAG Library**

Numerical Algorithms Group, Inc.
1101 31st Street, Suite 100
Downers Grove, IL 60515-1263

When publishing results of research, users should refer to the software with acknowledgment that it is licensed from NAG.

**PATRAN**

PDA Engineering
PATRAN Division
2975 Redhill Avenue
Costa Mesa, CA 92626

The number of simultaneous PATRAN users is limited to 40.

In order to renew our license annually, the vendor requires CHPC to submit two PATRAN models on magnetic media or 35 mm slides to be included in various vendor publications. Please contact CHPC if you wish to contribute. Users are also encouraged to submit photographs and articles about unique uses of PATRAN for publication in the PATRAN newsletter.

**GRADSCF**

Polyatomics Research Institute
c/o Andrew Komornicki, Director
1101 San Antonio Road, Suite 420
Mountain View, CA 94043

Use of this software under our license is specifically limited to a non-profit or university environment. Under no circumstances shall this software be used by a commercial company. If such use is desired, a separate agreement must be assigned.

**CHARMm**

Polygen Corporation
200 Fifth Avenue
Waltham, MA 02254
Mountain View, CA 94043

The UT System CHPC will provide a central technical support resource for all interaction between Polygen and CHPC users.

**Contouring, DI-3000**
**PicSure, DI-TEXTPRO**
**Grafmaker, GK-2000**

Precision Visuals, Inc.
6260 Lookout Road
Boulder, CO 80301

No special requirements, but see "Acknowledgments and CHPC Requirements" below.

**SLAM II**

Pritsker & Associates, Inc.
1305 Cumberland Avenue
P.O. Box 2413
West Lafayette, IN 47906-0413

SLAM II, executable derivative code, and all related materials, documentation, and written information are of a confidential and trade secret nature and are restricted from disclosure to ineligible users. Users external to the UT System will be denied access to SLAM II.

Users are also encouraged to share interesting applications or suggestions on SLAM II in P&A's monthly *SOFTLETTER*.

**ARI/RANDOM**

Wyle Laboratories
Applied Research Division
6151 W. Century Boulevard, Suite 1000
Los Angeles, CA 90045-5399

No special requirements, but see "Acknowledgments and CHPC Requirements" below.

## ACKNOWLEDGMENTS AND CHPC REQUIREMENTS

**Acknowledgment**

It is common professional practice to acknowledge all contributors when publishing the results of your research. All software suppliers expect this courtesy and appreciate receiving copies of published work involving their products.

**CHPC Requirements**

No personal or commercial use of UT System CHPC computers is permitted.

CHPC cannot make proprietary products available in any form on a computer system other than the one designated in the license — usually the Cray X-MP/24.

**Acknowledgment of CHPC**

In order to encourage the continued development of the UT System CHPC, it is essential that our users help us make its usefulness visible to The University of Texas System and the research community. You can help as follows:

- Give appropriate acknowledgment to the Center in each publication presenting results obtained through the use of the CHPC facility. A suggested citation is:

  **"Computing resources for this work were provided by The University of Texas System Center for High Performance Computing."**

  Please forward a copy of each publication to the Director, UT System CHPC (see the directory page of this *Newsletter*).

- In proposals for research in which the use of the CHPC facility is foreseen, the amount of CHPC computing resources to be used should be stated in the proposal, and funding for those resources should be requested. If an award is made based on such a proposal, please notify the Director of the CHPC as to the amount of the award and explain the award's significance to the UT System and to the research community (the Computer Project Registration (CPR) form can be used for this purpose). If the proposal is rejected due to any deficiency of the facilities offered by the CHPC, please make these facts known so that we may explore ways to remedy the deficiency.

---

# Information and Documentation Services at CHPC

J. Werst

The UT System Network is the gateway to the supercomputing resources of CHPC. Because easy access to information is essential for effective use of system resources, most CHPC documentation and information services are available through this network. You can use the network to gain access to on-line documentation, to order copies of published manuals, and to get consulting assistance.

**On-Line Help**

On-line help is available on the VAX 8600 for both VAX and Cray topics by typing HELP after the VMS prompt. If you are unfamiliar with how help works, type HELP, then INSTRUCTION, HINTS, or HELP/PROMPT after the TOPIC? prompt.

Useful HELP topics for new CHPC users are:

| | |
|---|---|
| CHPC | for information on CHPC services, policies, and accounting |
| @COS | for Cray Operating System (COS) commands |
| @CRAY_SOFTWARE | for descriptions of applications software on the Cray X-MP |
| @NETWORKING | for descriptions of various methods of access to CHPC |
| BULLETIN | for help on how to read bulletins on the VAX 8600 |
| NOTES | for an explanation of the VAX Notes conferencing system |
| MAIL | for an explanation of the VMS MAIL system |
| DOCUMENT | for help on using the on-line DOCUMENT utility |

**On-Line Documents**

On-line documents are larger and more detailed than HELP files. They are accessible through DOCUMENT, a locally written utility on the VAX 8600. To call it, type DOCUMENT at the VMS prompt. It will respond with a numbered list of on-line documents. Enter the number corresponding to the document you want. DOCUMENT then prompts you for the action to be performed (view, copy, or print).

DOCUMENT can also be run in command mode, in which you can enter the name or number of the document in the DOCUMENT command line together with the action to be performed. For example, to view the document CRAYCFT01, which is chapter 1 of the *Cray Fortran (CFT) Reference Manual*, type the command:

```
DOCUMENT "CRAYCFT01" /VIEW
```

**Published Documentation**

Cray user's manuals may be purchased through CHPC with the costs automatically billed to your computer supplies account. For instructions on how to order manuals on-line, log on to the VAX 8600 and type:

HELP CHPC MANUAL_ORDERS

DEC VAX/VMS manuals may be ordered directly from DEC by calling 1-800-258-1710, or by writing:

Digital Equipment Corporation
P.O. Box CS2008
Nashua, New Hampshire 03061

Convex manuals may be ordered by calling (214) 952-0200, or by writing:

Convex Computer Corporation
P.O. Box 833851, 701 Plano Road
Richardson, Texas 75083-3851

**Recommended Manuals:**

*COS Version 1 Reference Manual* (SR-0011, Revision O) Cray Research, Inc.

*COS Version 1 Ready Reference* (SQ-0023, Revision G) Cray Research, Inc.

*CFT77 Reference Manual* (SR-0018, Revision C) Cray Research, Inc.

*COS Message Manual* (SR-0039, Revision E) Cray Research, Inc.

*Introduction to VAX/VMS* (AA-Y500A-TE) Digital Equipment Corp.

*UNICOS Primer* (SG-2010, Revision D) Cray Research, Inc.

*UNICOS User Commands Reference Manual* (SR-2011, 5.0) Cray Research, Inc.

Please note that updates and revisions to manuals are not automatically sent to individual users. Notices will be posted on the VAX 8600 in VAX Notes under CHPCNews when updates become available.

**Usage Notes**

CHPC User Services has several free publications for beginning VAX 8600 and Cray users, called Usage Notes. They are available through the DOCUMENT utility explained above, or by sending a request via electronic mail to REMARK. Please include your name and mailing address in your request.

**Reference Materials**

Each UT System campus is sent copies of the most up-to-date reference manuals. These materials are mailed to the Operations Advisory Committee members (see directory page) who place the materials in a convenient campus location.

See the *CHPC Newsletter*, Vol. 3, No. 7, for a list of current reference materials and their location at each institution. This information is also available on-line through VAX Notes under the conference CHPCNews.

The CHPC site has a reference copy of all manuals available to users.

**On-Site Consulting**

Your UT System component institution may have on-site consulting. Contact your User Advisory Committee member for further details. See the directory page for the User Advisory Committee member on your campus.

**CHPC Site Consulting**

Consulting at the CHPC site in Austin is provided Monday through Friday, 9 A.M. to 5 P.M., except on Wednesdays, when the schedule is 10 A.M. to 3 P.M.. Consulting is also closed every day from noon to 1 P.M..

You will get the best response to your questions and problem reports by sending electronic mail to REMARK on the CHPC VAX 8600 (see HELP MAIL), since the User Services staff monitors REMARK many times each day. Please include the names of your .CPR, JOB, and JCL files if your question relates to a job which failed to run to completion.

For telephone consulting, call 1-800-262-2472 (471-2472 in Austin) during the hours listed above.

# Directory of the UT System CHPC

## UT System Center for High Performance Computing

Dr. James C. Almond, Director, (512) 471-2442 STS# 821-2442
Balcones Research Center, Commons Building
10100 Burnet Road, Austin, Texas 78758-4497
(512) 471-2472; (800) 262-2472 [toll-free, from Texas phones only]; STS# 821-2472

**Network IDs**

ARPAnet: CHPC.BRC.UTEXAS.EDU
BITNET: UTCHPC
GTE/TELENET: ;512131
UT System Network (DECnet): UTCHPC
MICOM dialup: (512) 471-9400 (1200 bps)
(512) 471-9420 (2400 bps)
MICOM class: CHPC

**Computer Systems**

Cray X-MP/24 – 2-CPU, 4-megaword shared central memory, 32-megaword SSD, 9.6 GB disk storage, COS 1.16 operating system.
VAX 8600 front-end – 32 MB central memory, 4 GB disk storage, VMS 4.6 operating system.
IBM 4381 file server – 16 MB central memory, 25 GB disk storage, MVS/XA 2.2.0 operating system.

**Accounting Information:** Beverly Millard, Accountant
(512) 471-3241, ext. 210/211 STS# 821-3241

**Administration:** Janet McCord, Senior Administrative Associate (512) 471-2418 STS# 821-2418

**Network Communications Services:** William C. Bard, Director of UT System Office of Telecommunication Services, (512) 471-2444 STS# 821-2444

**Newsletter Information:** John J. Werst, Editor, (512) 471-2463 STS# 821-2463

**Operations/Scheduling/Tape Librarian:** Dean Nobles, Operations Supervisor
(512) 471-2401 STS# 821-2401

**Operating Systems and Technical Support:** Gary Smith, Associate Director
(512) 471-2411 STS# 821-2411

**User Services:** Luther Keeler, Assistant Director, (512) 471-2425 STS# 821-2425
Consulting Office, (512) 471-2404 STS# 821-2404

**UT System CHPC Executive Committee**

Dr. Gerhard J. Fonken, Executive Vice President and Provost, UT Austin, Chairman
Dr. Robert H. Rutford, President, UT Dallas
Dr. Frederick F. Becker, Vice President for Research, UT System Cancer Center
Dr. K. W. Tolo, Associate VP for Academic Affairs and Research, UT Austin (*ex officio*)
Dr. Terry Mikiten, Associate Dean, Graduate School of Biomedical Sciences, UTHSC San Antonio
Dr. John McElroy, Dean, College of Engineering, UT Arlington
Dr. Hans Mark, Chancellor, UT System (*ex officio*)
Dr. James P. Duncan, Executive Vice Chancellor for Academic Affairs, UT System (*ex officio*)
Dr. Charles B. Mullins, Executive Vice Chancellor for Health Affairs, UT System (*ex officio*)
Dr. Charles H. Warlick, Exec. Director, Academic Information Systems, UT System (*technical liaison*)

**CHPC Advisory Committees**

| Component Institutions | Operations Advisory Committee | User Advisory Committee |
|---|---|---|
| UT Arlington | Mr. Melvin Pierce | Dr. Dennis Marynick |
| UT Austin | Dr. Charles H. Warlick | Dr. Bob E. Schutz |
| UT Dallas | Dr. Ronald Briggs | Dr. Cyrus D. Cantrell III |
| UT El Paso | Dr. John Starner | Dr. John Starner |
| UT Permian Basin | Dr. Douglas F. Hale | Dr. Douglas F. Hale |
| UT San Antonio | Mr. John W. Massey | Dr. Dennis Kern |
| UT Tyler | Mr. Steve Wilson | Dr. George Whitson |
| UTSMC Dallas | Mr. William A. Sholar | Dr. Steve Sprang |
| UTHSC Houston | Mr. Rick Miller | Dr. Nizar Mullani |
| UTHSC San Antonio | Dr. Clair W. Goldsmith | Dr. Paul Horowitz |
| UTHC Tyler | Mr. Johnnie Collier | Dr. Jerry McLarty |
| UTMB Galveston | Mr. Cecil Denney | Dr. Gilbert Hillman |
| UT MDA Cancer Center | Mr. William W. Hargrove, Jr. | Dr. Stuart O. Zimmerman |

# PUBLICATION ORDER AND ADDRESS CHANGE FORM

Mail this form to:

Publication Distribution
UT System CHPC
Balcones Research Center
10100 Burnet Rd.
Austin, TX 78758-4497

Please Supply the Following Publications:
(Use this Form for Free Publications Only)

| Number of Copies | Series Number | Title |
|---|---|---|

( ) Use mailing label on reverse side to send publications.

( ) Mail to the address shown below.

## ADDRESS CORRECTION

( ) Delete the address shown on label on the reverse side.

( ) Add or correct mailing label as shown below.

## ADDRESS

Name ______________________________

Department ______________________________

Institution ______________________________

Address ______________________________

______________________________

______________________________ Zip Code ____________

Requested by: ______________________________ Date: ____________

Center for High Performance Computing
Balcones Research Center
10100 Burnet Road
Commons Building
Austin, Texas 78758-4497

# THE UNIVERSITY OF TEXAS SYSTEM

# CENTER FOR HIGH PERFORMANCE COMPUTING

Balcones Research Center, Commons Building, 10100 Burnet Road, Austin, Texas 78758-4497

## INTRODUCTION

The University of Texas System Center for High Performance Computing (UT System CHPC) has been established by the UT System Board of Regents to serve the research and instructional supercomputing needs of the seven academic and six health component institutions of the UT System. The UT System CHPC, which became operational on May 15, 1986, is located on the UT Balcones Research Center campus in North Austin. Access to the services of the UT System CHPC from the component institutions is provided via a dedicated telecommunication network and by academic and common carrier data communication networks. The UT System CHPC reports to the Office of Academic Information Systems in the Office of the Chancellor, The University of Texas System. Policy guidance for the UT System CHPC is provided by a five-member Executive Committee appointed by the Office of the Chancellor. The computational resources of the UT System CHPC are allocated to the component institutions for support of their diverse research and instructional programs.

## COMPUTER EQUIPMENT, COMMUNICATION EQUIPMENT AND OPERATING SYSTEMS

The computing resources in the UT System CHPC include the Cray Research Inc. Model X-MP/24 "supercomputer"; two front-end computers, an IBM 4381-13 and a Digital Equipment (DEC) VAX 8600; a Network Systems Corporation (NSC) "HYPERchannel"; and appropriate input/output equipment and graphics workstations. The features of these systems are as follows:

### Cray X-MP/24 Computer System

- Two processors, giving a peak computing speed of 420 million floating point operations per second (MFLOPS)
- Four million 64-bit words (32 million bytes) of central memory
- Thirty-two million 64-bit words (256 million bytes) of "solid state disk" memory, capable of transferring data at the rate of one billion bytes per second
- I/O equipment directly connects to the Cray I/O channels:
    - 9.8 billion bytes of disk storage (Eight Cray Model DD49 disk drives)
    - Four 6250 bpi (bit per inch) nine-track magnetic tape drives (IBM Model 3420)
    - Six 37,500 bpi magnetic tape cartridge drives (IBM Model 3480)
- Cray Operating System (COS)

### IBM 4381-13 Front-end Computer

- Single processor system with 16 million bytes central memory
- Twenty billion bytes of disk storage (Four IBM Model 3380, dual density)
- Magnetic tape drives, same as above with Cray system. All tape drives are connected to both the Cray X-MP/24 and the IBM 4381-13.
- Line printer, 2000 lines per minute (IBM 4245)
- Communication control unit (IBM 3725)
- IBM MVS operating system, with hierarchical file management system and magnetic tape management system

### DEC VAX 8600 Front-end Computer

- Single processor system with 32 million bytes central memory
- Four billion bytes of disk storage (Six DEC RA81)
- Two 6250 bpi magnetic tape drives (DEC TA78)
- Communication Controllers:
    - DECnet router interface (for UT System-wide data communication)

- X.25 interface (for Telenet common carrier data network communication)
- Ethernet interface (for local CHPC workstation communication)
- Asynchronous line interface (for simultaneous connection of up to 24 asynchronous terminals via the dial telephone network)

- DEC VMS operating system

### NSC HYPERchannel System

- Three NSC HYPERchannel adapters connected to the Cray X-MP/24, IBM 4381-13 and DEC VAX 8600 permit serial data transmission between the computer systems at data rates up to fifty million bits (50 megabits) per second.

### High Performance Graphics Workstation Laboratory

- A separate room in the UT System CHPC will house a variety of high performance image graphics workstations. Initially, four DEC VAXstation II graphics workstations have been installed in this facility. Additonal equipment also will be acquired.

### Computer Output Services

- Separate output service facility with two 600 line per minute (lpm) printers, and an 8-pen plotter
- Printer service directly from IBM 4381-13 and DEC VAX 8600
- Laser printer for printed or plotter output on 8.5x11 in. paper
- Film recorder for graphics output, to include slide, sheet film and movie formats

## CRAY X-MP/24 COMPUTER SYSTEM CONFIGURATION

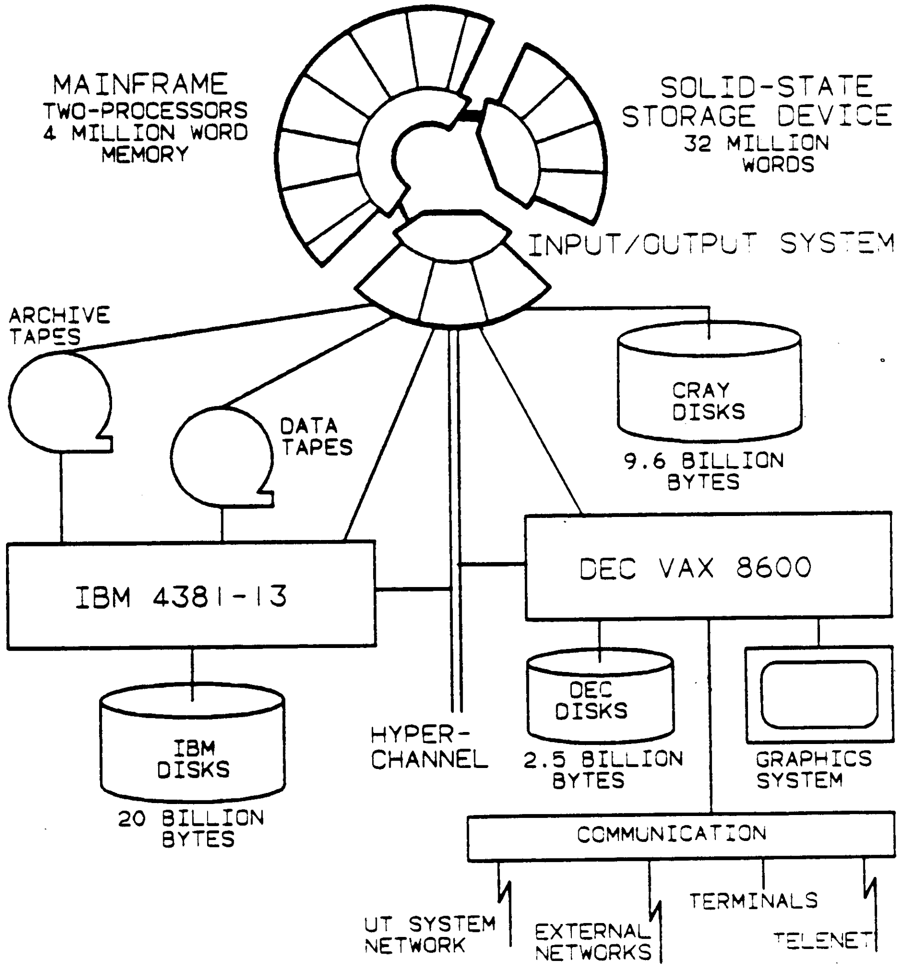

## AVAILABILITY OF COMPUTING SERVICES

The Cray X-MP/24 and the front-end computers are operated twenty-four hours per day, seven days per week. The computers are available for use at all times except for short periods of engineering maintenance, usually outside of "office hours".

## NETWORKS AND COMMUNICATION

### The UT System Computer Network

The University of Texas System computer network is operated by the UT System Office of Telecommunication Services for the component institutions of the UT System. The UT System computer network provides dedicated 56 kbps (thousand bits per second) circuits between the UT System CHPC and the component institutions in the Austin, Dallas, Houston and San Antonio metropolitan areas, and 9.6 kbps circuits to the institutions in east and west Texas. UT Austin also operates a 45 megahertz (million bits per second) microwave circuit between UT Austin and the UT System CHPC.

The network protocol for the UT System computer network is DECnet. The Office of Telecommunication Services has placed a DEC VAXstation II on each of the thirteen component campuses to serve as "Network Interface Processors" (NIPs). Mainframe computers, minicomputers and individual terminal workstations at the component campuses can connect to the NIPs to provide direct connection to the CHPC DEC VAX 8600 communication front-end computer and then to the Cray X-MP/24. Individual students, teachers and researchers at the component institutions can thereby connect the terminals and workstations in their offices and laboratories to the Cray supercomputer through the UT System computer network. Special provision has been made to permit connection of IBM as well as DEC computers to the UT System CHPC through the UT System computer network. The UT System CHPC provides remote login and interactive access through the virtual terminal capability, file transport, remote output queuing and graphics access.

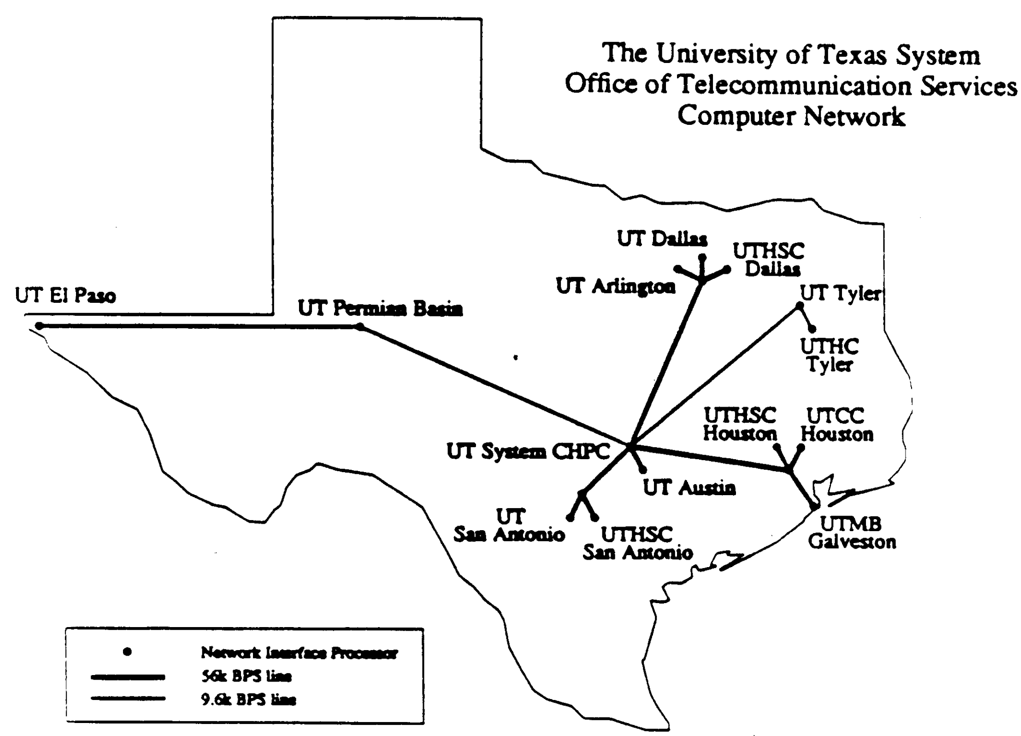

## Other Computer Networks and Access Methods

- Telenet. Access is available from all component campuses via the Telenet X.25 commercial data communication network. Terminals may be connected to the DEC VAX 8600 front-end computer by a local telephone call to Telenet throughout the United States.
- Bitnet. Through the Bitnet interface on the DEC VAX 8600, users of the UT System CHPC are able to communicate with the Cray X-MP/24 from over 525 universities connected to the Bitnet in the USA, Europe and Japan.
- ARPAnet. The Defense Advanced Research Projects Agency (DARPA) operates the ARPAnet which provides 56 kbps access to more than 100 universities and research facilities. Access to the UT System CHPC is provided through the ARPAnet inter-message processor (IMP) located on the UT Austin campus.
- CSnet. The computer science departments of colleges and universities are connected together through the CSnet. UT Austin is a member of CSnet via the ARPAnet network. UT System researchers visiting the schools on the CSnet are able to access the UT System CHPC using the CSnet.
- DECnet. In addition to the UT System computer network described earlier, there is a state-wide DECnet to which are connected many other Texas colleges and universities. The UT System computer network is a sub-network within the larger statewide DECnet network.
- Dial-up access. The DEC VAX 8600 provides 2400 bps dial access to 24 simultaneous users.

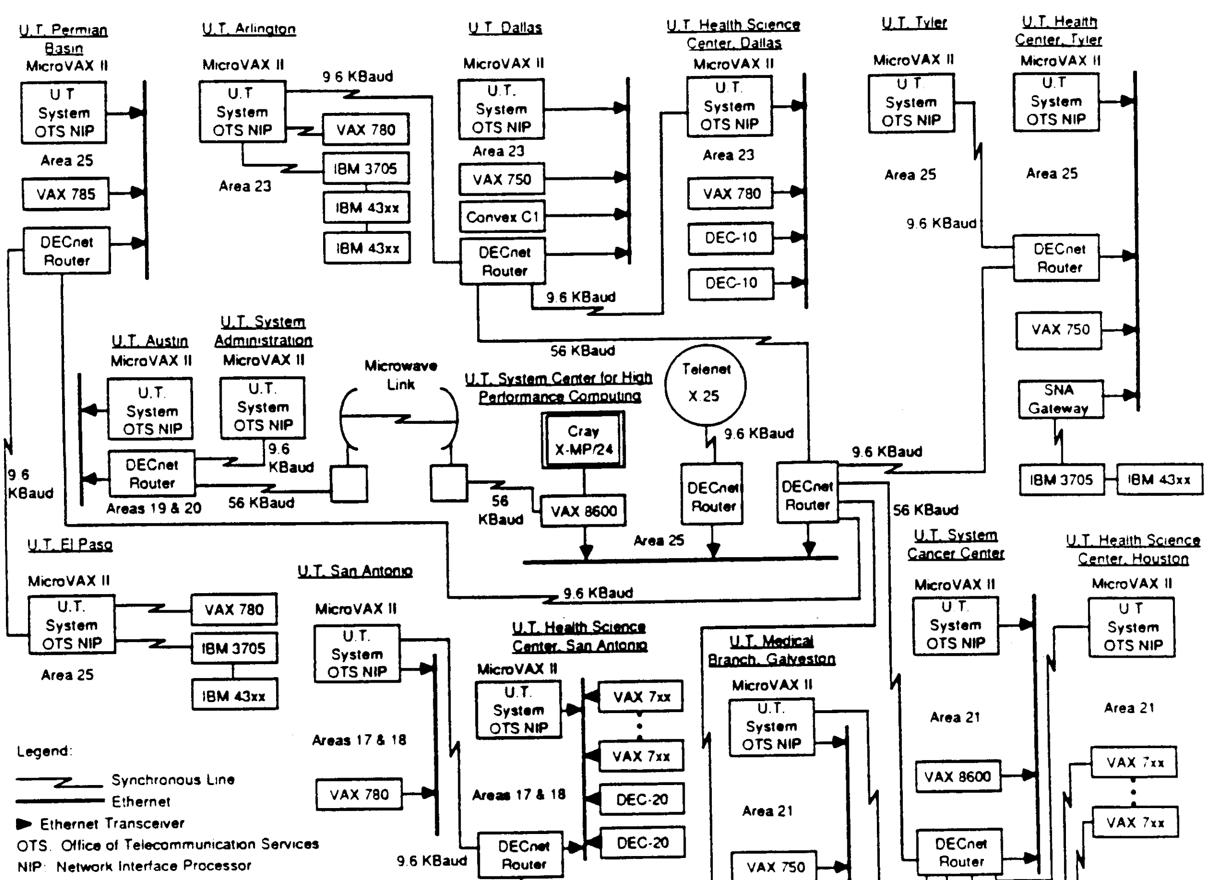

U.T. SYSTEM COMPUTER NETWORK 1 Feb 1987

## LANGUAGES AND SOFTWARE PACKAGES ON THE CRAY X-MP/24

### Computer Languages

| | |
|---|---|
| CFT | Cray Fortran |
| CAL | Cray Assembly Language |
| PROLOG | Non-procedural programming language |
| LISP | Symbolic data processing (Portable standard LISP) |
| PASCAL | (ISO Level 1 Pascal) |
| C | Based on portable C compiler |

### Mathematical and Statistical Software

| | |
|---|---|
| BCSLIB | Standard mathematical and statiscal subroutine library |
| BCS/Vectorpak | High optimized SCILIB routines, sparse-vector operations, and multi-dimensional FFTs |
| IMSL | Fortran subroutine library for standard math |
| LSODE | FORTRAN subroutine library of Gear-type ODE solvers |
| MINPACK | Numerical solutions to linear equations and non-linear least-squares problems |
| NAG | Fortran subroutine library for standard math |
| PDE/PROTRAN | Finite element PDE solver |
| SCILIB | Contains LINPACK, EISPACK. Fortran subroutine library for solving linear systems |
| SLATEC | Mathematical subroutine libraries developed by national laboratories |
| ITPACKV | Iterative solutions for sparse linear systems |

### Utilities

| | |
|---|---|
| BENCHLIB | FORTRAN code conversion and optimization aids |
| FITPACK | Curve and surface fitting by splines under tension |
| NAGGRF | NAG/DI-3000 graphics interface |
| PLOT-10 | Techtronix terminal graphics generator |
| ZETALIB | ZETA ink plotter graphics generator |

### Application Software

| | |
|---|---|
| ABAQUS | Engineering analysis and modeling system |
| ARI/RANDOM | MSC/NASTRAN post-processor |
| DI-3000 | Device independent Fortran graphics systems |
| GAUSSIAN-82 | Programs for performing molecular orbital calculations |
| MOVIE.BYU | Programs for displaying and manipulating graphical renderings of solid model MSC/NASTRAN. |
| MSC/NASTRAN | Programs for performing static and dynamic structural analysis |
| PATRAN | Programs for constructing, displaying, and editing 3-D goemetric and finite element models |
| SLAM II | FORTRAN-based simulation and modeling language |

## ALLOCATION OF UT SYSTEM CHPC COMPUTING RESOURCES

The computing resources of the UT System CHPC are allocated to the component institutions of The University of Texas System by the CHPC Executive Committee. The individual institutions are then responsible for the further allocation and management of the computing resources within the institutions for their instructional and research programs. Allocations to the UT System component institutions and thence to departments and individual research projects are made on a dollar basis rather than on a basis of hours or "service units".

Sponsored research projects are expected to pay for UT System CHPC computer resources used in support of the research projects. Those projects with insufficient funds to support their total computing requirements may request additional funds from the allocation to the researchers' home component institution. Research projects with computing requirements that exceed the level supportable within an individual component's allocation may apply to the UT System CHPC Executive Committee for additional resource allocation.

Computer time provided by the UT System CHPC is charged on either a prime or deferred basis. Programs run on a prime rate basis are run as soon as possible and are charged at the full computer rate. Deferred jobs are those designated by the user to run during late night and weekend periods, and such deferred jobs are charged at a rate less than the full rate. The availability of such differential pricing permits users to stretch their allocations of UT System CHPC computer time funds or their grant dollars.

## LOCATION OF THE UNIVERSITY OF TEXAS SYSTEM CENTER FOR HIGH PERFORMANCE COMPUTING

The UT System CHPC is located in the Commons Building of the Balcones Research Center of UT Austin at 10100 Burnet Road, Austin, Texas 78758-4497 (see map on page 10). The UT System Office of Telecommunication Services is also located in the Commons Building. The Commons is particularly well designed to house the UT System CHPC, with an auditorium which seats 330 people, a conference center which seats 140 people, a large cafeteria, a branch of the UT Austin library and facilities for recreational sports. The computer room contains 3,500 square feet of floor space and houses all of the equipment described earlier. In addition to the offices for the UT System CHPC staff, there are ample facilities for users who choose to go to the CHPC site to use the computers. Workstations, reference materials, consultation and input/output services are available to the users.

## ORGANIZATION OF THE UNIVERSITY OF TEXAS SYSTEM CENTER FOR HIGH PERFORMANCE COMPUTING

The UT System Center for High Performance Computing has been established to serve the supercomputing needs of the component institutions of the UT System. The Director of the UT System CHPC reports to the Executive Director, Office of Academic Information Systems in the Office of the Chancellor of the University of Texas System. Dr. Charles H. Warlick is the UT System Executive Director, Academic Information Systems and his office is located in Ashbel Smith Hall, 201 West 7th Street, Austin, TX 78701; (512) 499-4240.

The Office of the Chancellor has assigned policy oversight for the Center for High Performance Computing to the UT System CHPC Executive Committee. The five-member CHPC Executive Committee is appointed by the Office of the Chancellor and consists of three members from the seven academic components and two members from the six health components.

The members of the UT System Executive Committee are:

*Academic Component Representatives*

Dr. Gerhard J. Fonken
Executive Vice President and Provost
The University of Texas at Austin

Dr. John W. Rouse
Dean of Engineering
The University of Texas at Arlington

Dr. Robert H. Rutford
President
The University of Texas at Dallas

*Health Component Representatives*

Dr. Frederick F. Becker
Vice President for Research
The University of Texas System Cancer Center

Dr. Terry Mikiten
Associate Dean, The University of Texas Graduate School of Biomedical Sciences
The University of Texas Health Science Center at San Antonio

The University of Texas System
Center for High Performance Computing

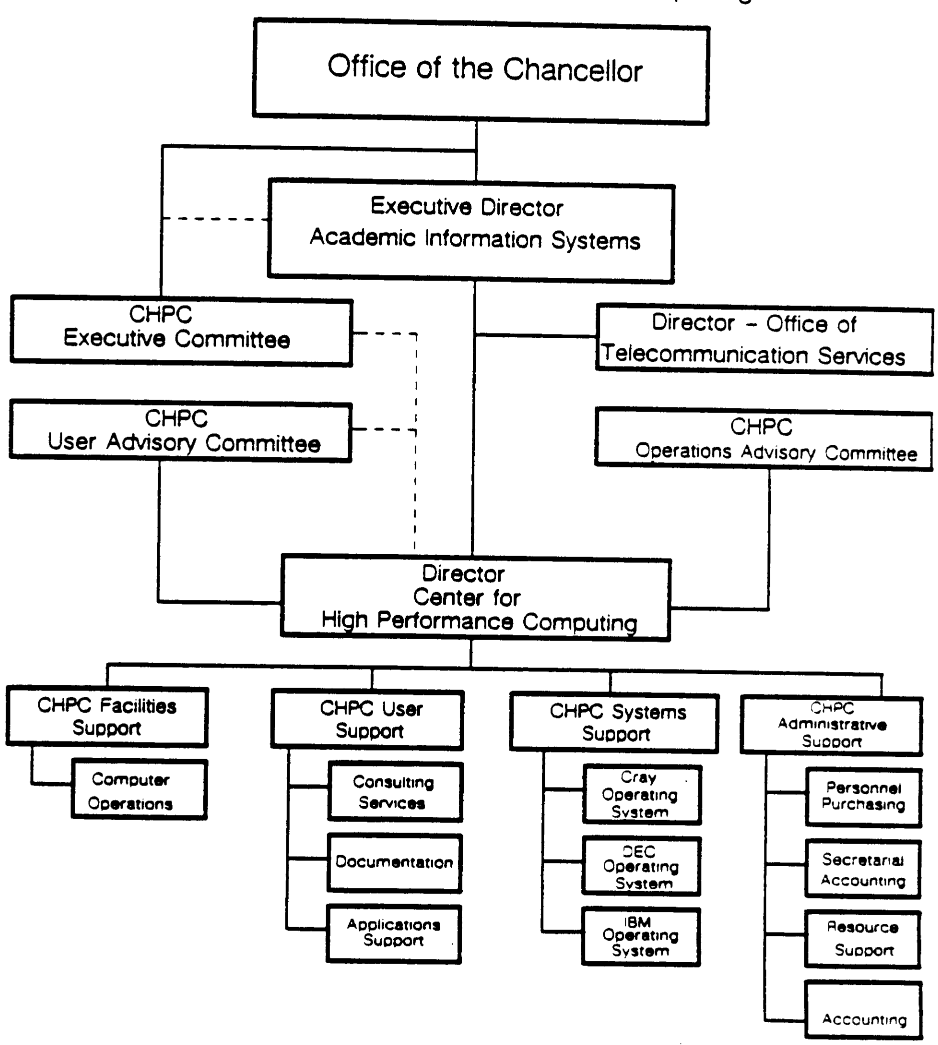


10 Feburary 1987

Also serving with the UT System CHPC Executive Committee:

*ex-officio Members*

Dr. Hans Mark
Chancellor
The University of Texas System

Dr. James P. Duncan
Executive Vice Chancellor for Academic Affairs
The University of Texas System

Dr. Charles B. Mullins
Executive Vice Chancellor for Health Affairs
The University of Texas System

*Technical Liaison*

Dr. Charles H. Warlick
Executive Director,
Academic Information Systems
The University of Texas System

**UT System CHPC User Advisory Committee**

The head of each UT System academic and health component institution has appointed a member of the faculty to serve on a committee to advise the Executive Committee and the Director, UT System CHPC, on matters of policy and on issues of importance to the users of the UT System CHPC. The thirteen members of the UT System CHPC User Advisory Committee are:

Dr. Dennis Marynick
Associate Professor, Chemistry
The University of Texas at Arlington

Dr. Stuart O. Zimmerman
Professor and Chairman, Biomathematics
The University of Texas System Cancer Center

Dr. Julian Peterson
Professor, Biochemistry
The University of Texas Health
Science Center at Dallas

Dr. Gilbert Hillman
Professor, Pharmacology and Toxicology
The University of Texas Medical
Branch at Galveston

Dr. Douglas F. Hale
Professor, Math./Computer Science
The University of Texas of the Permian Basin

Dr. Paul M. Horowitz
Professor, Biochemistry
The University of Texas Health
Science Center at San Antonio

Dr. George Whitson
Associate Professor, Computer Science
The University of Texas at Tyler

Dr. Bob E. Schutz
Professor, Aero. Eng./Eng. Mech.
The University of Texas at Austin

Dr. Cyrus D. Cantrell, III
Professor, Physics
The University of Texas at Dallas

Dr. John Starner
Assistant Professor, Computer Science
The University of Texas at El Paso

Dr. Nizar Mullani
Associate Professor, Medicine
The University of Texas Health
Science Center at Houston

Dr. Dennis Kern
Assistant Professor, Math./Computer Science
The University of Texas at San Antonio

Dr. Jerry McLarty
Chief, Epidemiology/Biomathematics
The University of Texas Health
Center at Tyler

## UT System CHPC Operations Advisory Committee

The head of each academic and health component institution has appointed the director of academic computing for the component campus to serve on a committee to advise the Director of the UT System CHPC on matters relating to the operation of the UT System CHPC and the interaction of the Center with the academic computing facilities at the component institutions. The thirteen members of the UT System CHPC Operations Advisory Committee are:

Mr. Melvin Pierce
Director, Academic Computing
The University of Texas at Arlington

Mr. William W. Hargrove, Jr.
Director, Computing Resources
The University of Texas System Cancer Center

Dr. Julian Peterson
Director, Academic Computing Services
The University of Texas Health Science Center at Dallas

Mr. Cecil Denney
Director, Academic Computing
The University of Texas Medical Branch at Galveston

Dr. Douglas F. Hale
Director, Computer Services
The University of Texas of the Permian Basin

Dr. Clair W. Goldsmith
Director, Computing Resources
The University of Texas Health Science Center at San Antonio

Dr. Charles H. Warlick
Director, Computation Center
The University of Texas at Austin

Dr. Ronald Briggs
Director, Academic Computing
The University of Texas at Dallas

Dr. John Starner
Director, Computer Center
The University of Texas at El Paso

Mr. Stephen M. Sokol
Assistant Vice President, Information Management
The University of Texas Health Science Center at Houston

Dr. Johnnie Collier
Systems Analyst
The University of Texas Health Center at Tyler

Mr. Steve Wilson
Director, Computer Services
The University of Texas at Tyler

## Management of the UT System CHPC

The Office of the Chancellor of the UT System and the President of UT Austin have implemented an agreement under which UT Austin is providing the necessary facility management services for the establishment and operation of the UT System CHPC. Dr. Charles H. Warlick, Director of the UT Austin Computation Center, is serving as Interim Director of the UT System CHPC. The staff of the UT System CHPC is budgeted for 31 people, including seven system programmers, eight user services personnel, nine computer operators plus management and administrative personnel.

February 1,1987

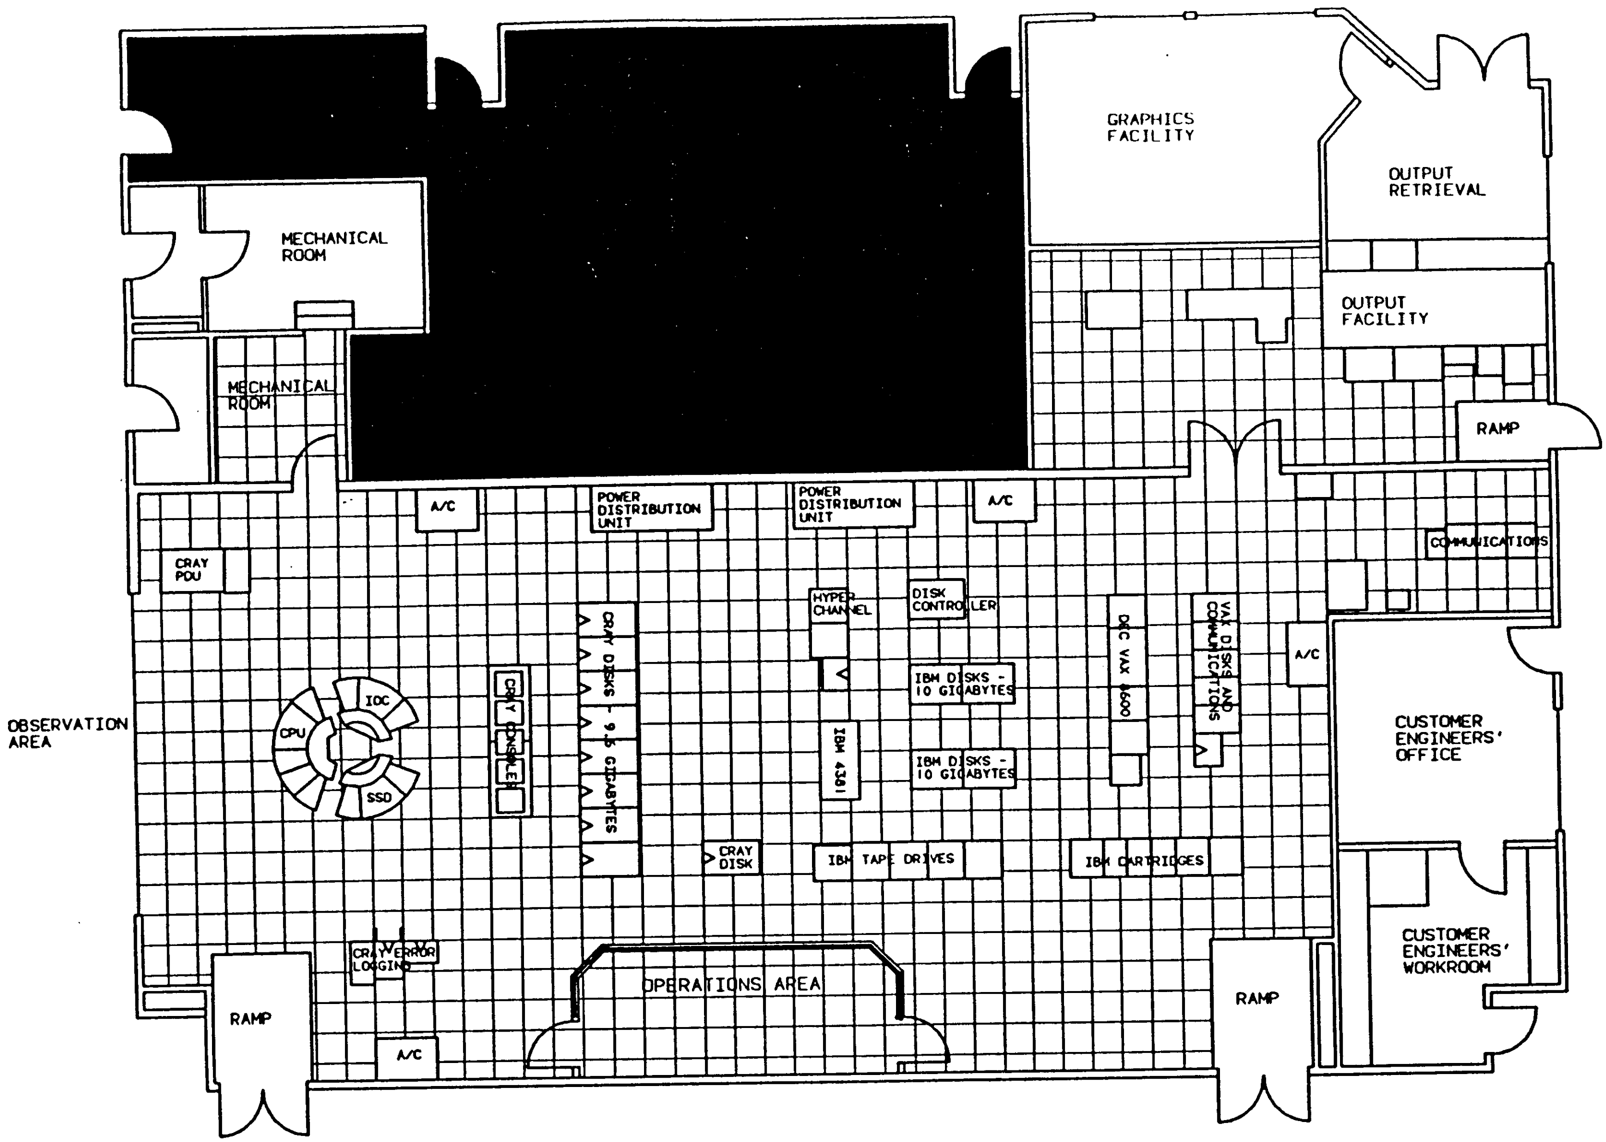


**THE UNIVERSITY OF TEXAS SYSTEM CENTER FOR HIGH PERFORMANCE COMPUTING**
**COMPUTER FACILITY AREA**
**AUSTIN, TEXAS**

SCALE: 1/8" = 1'
JULY 3, 1986

## COMPUTER EQUIPMENT, COMMUNICATION EQUIPMENT AND OPERATING SYSTEMS

The computing resources in The University of Texas System Center for High Performance Computing (UT System CHPC) include:

- o the Cray Research Inc. Model X-MP/24 "Supercomputer"
- o an IBM 4381-13
- o a Digital Equipment (DEC) VAX 8600
- o a Network Systems Corporation (NSC) "HYPERchannel"

Features of the Cray X-MP/24 Computer System are:

- o 2 processors (420 million floating point operations per second)
- o 4 million 64-bit words (32 million bytes) of memory
- o 32 million 64-bit words (256 million bytes) as random-access memory ("solid state disk" memory)
- o I/O equipment to Cray I/O Channels:
    - 9.8 billion bytes of disk storage (8 Cray Model DD49 disk drives)
    - four 6250 bpi (bits per inch) 9-track magnetic tape drives (IBM Model 3420)
    - six 37,500 bpi magnetic tape cartridge drives (IBM Model 3480)
    - Cray Operating System (COS)

## NETWORKS AND COMMUNICATION

The University of Texas System Board of Regents has established The University of Texas System Center for High Performance Computing (UT System CHPC) to serve the research and instructional supercomputing needs of the seven academic and six health component institutions of the University of Texas System. Component institutions have access to the services of the UT System CHPC through a dedicated telecommunications network and through academic and common carrier data communication networks.

University of Texas System
Center for High Performance Computing
Communications Network

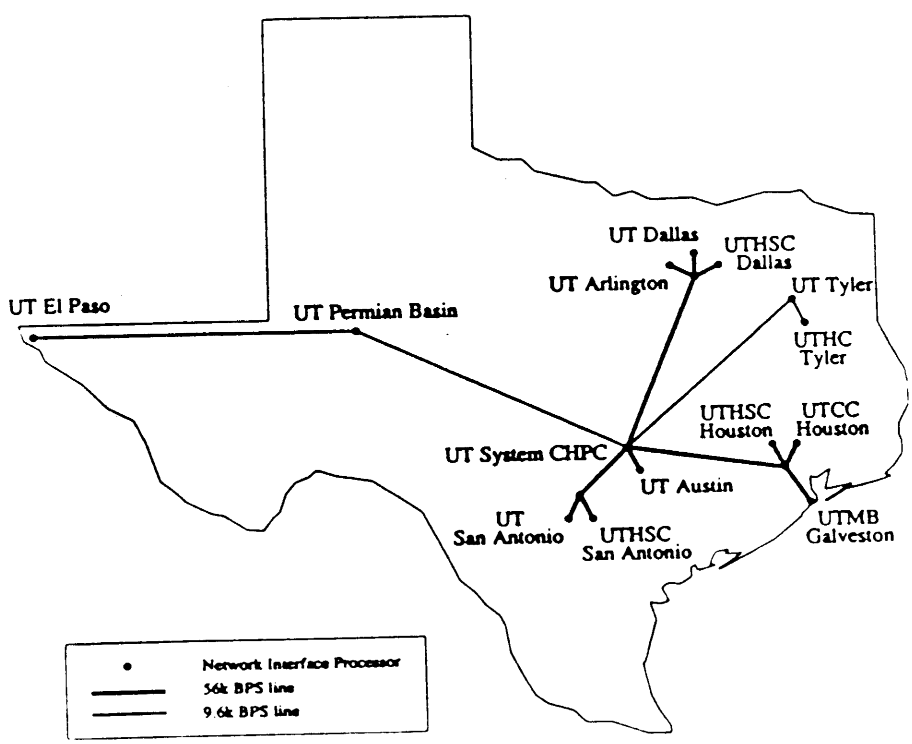

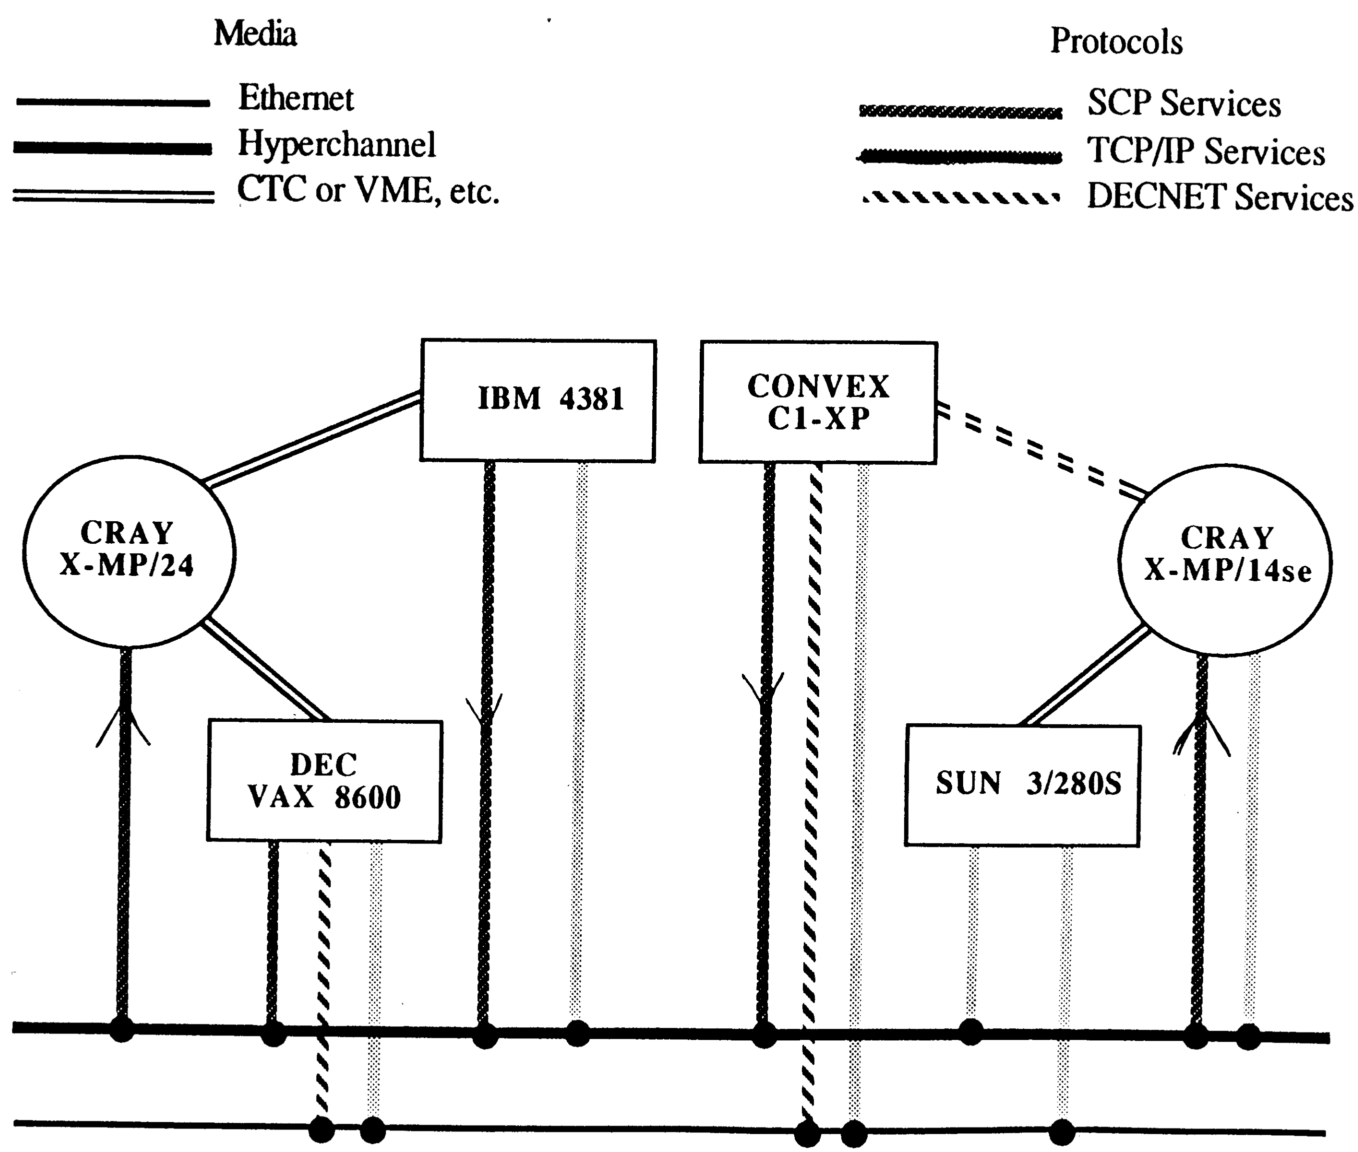
Media
Ethernet
Hyperchannel
CTC or VME, etc.
Protocols
SCP Services
TCP/IP Services
DECNET Services
CRAY X-MP/24
IBM 4381
CONVEX C1-XP
CRAY X-MP/14se
DEC VAX 8600
SUN 3/280S

# Conformance Testing Profile for Department of Defense Military Standard Data Communications High Level Protocol Implementations

Conformance Testing Administrator
DCA Code R640

# CONFORMANCE TESTING PROFILE FOR DEPARTMENT OF DEFENSE MILITARY STANDARD DATA COMMUNICATIONS HIGH LEVEL PROTOCOL IMPLEMENTATIONS

DRAFT

1/26/89

Comments can be sent to:

DCA Code R640
1860 Wiehle Avenue
Reston, VA 22090
ATTN: Conformance Testing Administrator.

Conformance Testing Profile For Department of Defense
Military Standard Data Communications High Level
Protocol Implementations

## Introduction

This document presents the conformance testing requirements for the five Department of Defense (DoD) Military Standard Data Communications High Level Protocols (IP,TCP,FTP,SMTP,TELNET). The conformance testing requirements are grouped in two classes: MANDATORY (M) and OPTIONAL (O). All protocol implementations must sucessfully pass the Mandatory conformance tests to be certified for use on DoD networks. If the protocol implementation contains options it must successfully pass those conformance tests associated with that option to be certified for use on DoD networks. When the conformance tests are run, using the Defense Communications Agency Upper (DCA) Level Protocol Test System, each tests results in one of four messages: OK, PROBLEM, INCONCLUSIVE, and OBSERVATION. A test is considered successful, when its results are OK or OBSERVATION and the required functionality is available in the Implementation Under Test (IUT). A PROBLEM message indicates non-conformance to the specification. INCONCLUSIVE means that the test ended before there was sufficient data to evaluate conformance to the tested functional requirement; therefore the test should be re-run until the pass/fail criteria is met. Test results which contain OBSERVATION are useful in determining specific parameters associated with the IUT. An additional response, NOT IMPLEMENTED, can be sent by the IUT. This indicates that the IUT has not implemented the tested functional requirement. If the functional requirement is mandatory or optionally implemented by the IUT, then the response is interpreted as a PROBLEM.

The test numbers in this document refer to the same numbers in the DCA Upper Level Protocol Test System. For further information on individual tests refer to the following documents: 1) DCA Upper Level Protocol Test System Internet Protocol Mil-Std 1777 Test Traceability Index 2) DCA Upper Level Protocol Test System Transmission Control Protocol Mil-Std 1778 Test Traceability Index 3) DCA Upper Level Protocol Test System Transmission Control Protocol / Internet Protocol Tightly Coupled Test Traceability Index 4) DCA Upper Level Protocol Test System File Transfer Protocol Mil-Std 1780 Test Traceability Index 5) DCA Upper Level Protocol Test System Simple Mail Transfer Protocol Mil-Std 1781 Test Traceability Index 6) DCA Upper Level Protocol Test System TELNET Protocol Mil-Std 1782 Test Traceability Index.

# INTERNET PROTOCOL MIL-STD 1777 PROFILE
# (IP TEST SUITE)

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 207<br>210 - 213<br>238 - 253<br>300 - 307<br>310 - 312<br>316 - 317 |
| OPTIONS | | |
| | STREAM | 214 - 218 |
| | TIMESTAMP | 219 - 229 |
| | RECORD ROUTE | 230 - 237 |
| | ICMP TIMESTAMP | 308 - 309 |
| | ICMP INFO | 313 - 315 |

# INTERNET PROTOCOL MIL-STD 1777 PROFILE
# (IP TEST SUITE)

| Test No. | Class | Purpose |
|---|---|---|
| IP-Tests | | |
| 1 | M | Deliver and Send Datagram |
| 2 | M | Precedence Values - Acceptance |
| 3 | M | Low Delay - Acceptance |
| 4 | M | High Reliability - Acceptance |
| 5 | M | High Throughput - Acceptance |
| 6 | M | Type of Service Combinations - Acceptance |
| 7 | M | Illegal Time to Live - Rejection |
| 8 | M | Too Small Time to Live - Rejection |
| 9 | M | Range of Valid Time to Live Values - Acceptance |
| 10 | M | Invalid Version Number - Rejection |
| 11 | M | Invalid Checksum - Rejection |
| 12 | M | Illegally Small Header Length - Rejection |
| 13 | M | Inconsistent Header and Total Length - Rejection |
| 14 | M | Illegally Small Total Length - Rejection |
| 15 | M | Total Length Greater Than Actual Length - Rejection |
| 16 | M | Total Length Smaller Than Actual Length - Rejection |
| 17 | M | More Fragments Field - Recognition |
| 18 | M | Reassembly of 2-Fragment Datagram |
| 19 | M | Reassembly of 3-Fragment Datagram |
| 20 | M | Reassembly of 576-Octet Datagram |
| 21 | M | Reassembly of Out-of-Order Fragments - Mixed |
| 22 | M | Reassembly of Fragments Received in Reverse Order |
| 23 | M | Expired Time to Live in Arriving Fragment - Rejection |
| 24 | M | Duplicate Fragment in Reassembly |
| 25 | M | Inconsistent Protocol Fields in Fragment Reassembly - Rejection |
| 26 | M | Inconsistent Precedence Fields in Fragment Reassembly - Rejection |
| 27 | M | Inconsistent Source Fields in Fragment Reassembly - Rejection |
| 28 | M | Expiration of Time to Live during Reassembly - Rejection |
| 29 | M | Setting and Restarting Reassembly Timer |
| 30 | M | Reassembly of Two Intermixed Datagrams |
| 31 | M | Reassembly of Many Intermixed Datagrams |
| 32 | M | Datagram with NOP and End-of-Options List |
| 33 | M | Datagram with 2 NOP, 1 End-of-Options List Options - Acceptance |
| 34 | M | Datagram with 3 NOP, 1 End-of-Options List Options - Acceptance |
| 35 | M | Datagram with 4 NOP Options - Acceptance |
| 36 | M | Datagram with Invalid Options - Rejection |

| Test No. | Class | Purpose |
|---|---|---|
| **IP-Tests** | | |
| 37 | M | Stream Option - Acceptance |
| 38 | M | Stream Option with Invalidly Large Option Length - Determination of Acceptance |
| 39 | M | Stream Option with Invalidly Small Option Length - Determination of Acceptance |
| 40 | M | Duplicate Stream Option - Rejection |
| 41 | M | Stream Option without Copy Flag - Determination - Acceptance |
| 42 | M | Stream Option on Fragmented Datagram - Reassembly |
| 43 | M | Stream Option Not Present on all Fragments - Rejection |
| 44 | M | Full Range of Stream Id Values - Acceptance |
| 46 | M | Timestamp Option with Format 0 - Acceptance |
| 47 | M | Timestamp Option with Format 1 - Acceptance |
| 48 | M | Timestamp Option with Format 3 - Acceptance |
| 49 | M | Timestamp Option with Overflow - Acceptance |
| 50 | M | Timestamp Option with Empty value fields - Acceptance |
| 51 | M | Timestamp Option with Non-standard Timestamp - Acceptance |
| 52 | M | Timestamp Option with Standard and Non-standard Timestamp fields - Acceptance |
| 53 | M | Timestamp Option with Illegally Small Option Pointer - Determination of Acceptance |
| 54 | M | Timestamp Option with Invalid Option Ptr - Determination of Acceptance |
| 55 | M | Timestamp Option with Illegally Small Option Length - Determination of Acceptance |
| 56 | M | Timestamp Option with Illegally Large Option Length - Determination of Acceptance |
| 57 | M | Timestamp Option with Invalid Option Length - Determination of Acceptance |
| 58 | M | Timestamp Option with Copy Flag - Determination of - Acceptance |
| 59 | M | Timestamp Option with Invalid Format - Determination of - Acceptance |
| 60 | M | Duplicate Timestamp Option - Rejection |
| 61 | M | Timestamp Option on Datagram Fragments - Reassembly |
| 62 | M | Timestamp Option Duplicated on Datagram Fragments - Determination of Acceptance |
| 63 | M | Timestamp Option with Incomplete Timestamp Field - Determination of Acceptance |
| 64 | M | Record Route Option - Acceptance |
| 65 | M | Record Route Option with all Field Filled - Acceptance |
| 66 | M | Record Route Option with all Route Fields Not Filled - Acceptance |
| 67 | M | Record Route Option with Invalidly Small Option Length - Determination of Acceptance |

| Test No. | Class | Purpose |
|---|---|---|
| **IP-Tests** | | |
| 68 | M | Record Route Option with Invalidly Large Option Length - Determination of Acceptance |
| 69 | M | Record Route Option with Illegally Small Option Pointer - Determination of Acceptance |
| 70 | M | Record Route Option with Invalid Option Ptr - Determination of Acceptance |
| 71 | M | Record Route Option with Invalid Copy Flag - Determination of Acceptance |
| 72 | M | Duplicate Record Route Option - Rejection |
| 73 | M | Record Route Option with Invalid Relationship Between Length and Pointer - Determination of Acceptance |
| 74 | M | Record Route Option on Datagram Fragments - Reassembly |
| 75 | M | Record Route Option Duplicated on Datagram Fragments - Rejection |
| 76 | M | Strict Source Option - Acceptance |
| 77 | M | Strict Source Option with Multiple Record Fields - Acceptance |
| 78 | M | Strict Source Option with Not All Gateways Traversed - Rejection |
| 79 | M | Strict Source Option with Invalidly Small Option Length - Determination of Acceptance |
| 80 | M | Strict Source Option with Invalidly Large Option Length - Determination of Acceptance |
| 81 | M | Strict Source Option With Illegally Small Option Pointer - Determination of Acceptance |
| 82 | M | Strict Source Option With Invalid Option Ptr - Determination of Acceptance |
| 83 | M | Strict Source Option Without Copy Flag - Rejection |
| 84 | M | Duplicate Strict Source Option - Rejection |
| 85 | M | Strict Source Option Not Duplicated on all Fragments - Rejection |
| 86 | M | Strict Source Option Duplicated on Fragments - Reassembly |
| 87 | M | Loose Source Option - Acceptance |
| 88 | M | Loose Source Option with Multiple Record Fields - Acceptance |
| 89 | M | Loose Source Option with Not All Gateways Traversed - Rejection |
| 90 | M | Loose Source Option with Invalidly Small Option Length - Determination of Acceptance |
| 91 | M | Loose Source Option with Invalidly Large Option Length - Determination of Acceptance |
| 92 | M | Loose Source Option With Illegally Small Option Pointer - Determination of Acceptance |

| Test No. | Class | Purpose |
|---|---|---|
| **IP-Tests** | | |
| 93 | M | Loose Source Option With Invalid Option Ptr - Determination of Acceptance |
| 94 | M | Loose Source Option Without Copy Flag - Rejection |
| 95 | M | Duplicate Loose Source Option - Rejection |
| 96 | M | Loose Source Option Not Duplicated on all Fragments - Rejection |
| 97 | M | Loose Source Option Duplicated on Fragments - Reassembly |
| 98 | M | 2 Options on Datagram, Record Route and Strict Source - Acceptance |
| 99 | M | 3 Options on Datagram, Record Route, Loose Source and Timestamp, - Acceptance |
| 100 | M | 4 Options on Datagram, Record Route, Strict Source, Stream and Timestamp, - Acceptance |
| 200 | M | Precedence - Setting of Values 0 - 3 |
| 201 | M | Low Delay - Setting |
| 202 | M | High Reliability - Setting |
| 203 | M | High Throughput - Setting |
| 204 | M | Type of Service Combinations - Setting with Precedence Values 0 - 3 |
| 205 | M | Precedence - Setting of Values 4 - 7 |
| 206 | M | Type of Service Combinations - Setting with Precedence Values 4 - 7 |
| 207 | M | Time To Live Illegally Small - Refusal to Send Range of time to Live Values - Setting Don't Fragment Flag - Setting |
| 210 | M | 1 NOP and 3 EOL options - Sending |
| 211 | M | 2 NOP and 2 EOL options - Sending |
| 212 | M | 3 NOP and 2 EOL options - Sending |
| 213 | M | 4 NOP - Sending |
| 214 | O | Stream Option - Sending |
| 215 | O | Stream Option with Illegal Option Length - Refusal to Send |
| 216 | O | Stream Option with Incorrect Stream ID - Refusal to Send |
| 217 | O | Two Stream Options - Refusal to Send |
| 218 | O | Range of Stream Option Ids - Setting |
| 219 | O | Timestamp Option with Format 0 - Sending |
| 220 | O | Timestamp Option with Format 1 - Sending |
| 221 | O | Timestamp Option with Format 3 - Sending |
| 222 | O | Timestamp Option with Overflow Set - Refusal to Send |
| 223 | O | Timestamp Option with Illegally Small Pointer - Refusal to Send |
| 224 | O | Timestamp Option with Invalid Pointer - Refusal to Send |

| Test No. | Class | Purpose |
|---|---|---|
| IP-Tests | | |
| 225 | 0 | Timestamp Option with Illegally Small Option Length - Refusal to Send |
| 226 | 0 | Timestamp Option with Illegally Large Option Length - Refusal to Send |
| 227 | 0 | Timestamp Option with Copy Flag - Refusal to Send |
| 228 | 0 | Timestamp Option with Unassigned Format Code - Refusal to Send |
| 229 | 0 | Two Timestamp Options - Refusal to Send |
| 230 | 0 | Record Route Option - Sending |
| 231 | 0 | Record Route Option with Multiple Record Fields - Sending |
| 232 | 0 | Record Route Option with Illegally Small Length - Refusal to Send |
| 233 | 0 | Record Route Option with Illegally Large Length - Refusal to Send |
| 234 | 0 | Record Route Option with Illegally Small Pointer - Refusal to Send |
| 235 | 0 | Record Route Option with Invalid Pointer - Refusal to Send |
| 236 | 0 | Record Route Option With Invalid Copy Flag - Refusal to Send |
| 237 | 0 | Two Record Route Options - Refusal to Send |
| 238 | M | Strict Source Record Route Option - Sending |
| 239 | M | Strict Source Option with Multiple Fields - Sending |
| 240 | M | Strict Source Option with Illegally Small Option Length - Refusal to Send |
| 241 | M | Strict Source Option with Invalidly Large Option Length - Refusal to Send |
| 242 | M | Strict Source Option with Illegal Small Pointer Value - Refusal to Send |
| 243 | M | Strict Source Option with Invalid Strict Source Pointer Value - Refusal to Send |
| 244 | M | Strict Source Option with No Copy Flag - Refusal to Send |
| 245 | M | Two Strict Source Options - Refusal to Send |
| 246 | M | Loose Source Record Route Option - Sending |
| 247 | M | Loose Source Option with Multiple Fields - Sending |
| 248 | M | Loose Source Option with Illegally Small Option Length - Refusal to Send |
| 249 | M | Loose Source Option with Invalidly Large Option Length - Refusal to Send |
| 250 | M | Loose Source Option with Illegal Small Pointer Value Value - Refusal to Send |
| 251 | M | Loose Source Option with Invalid Strict Source Pointer Value - Refusal to Send |
| 252 | M | Loose Source Option with No Copy Flag - Refusal to Send |
| 253 | M | Two Loose Source Options - Refusal to Send |
| 254 | 0 | Two Different Options - Sending |
| 255 | 0 | Three Different Options - Sending |
| 256 | 0 | Multiple Options - Sending |

| Test No. | Class | Purpose |
|---|---|---|
| **ICMP Tests** | | |
| 300 | M | ICMP Destination Unreachable - Acceptance |
| 301 | M | ICMP Destination Unreachable - Sending |
| 302 | M | ICMP Time Exceeded - Acceptance |
| 303 | M | ICMP Time Exceeded - Sending |
| 304 | M | ICMP Time Exceeded in Reassembly - Sending |
| 305 | M | ICMP Parameter Problem - Acceptance |
| 306 | M | ICMP Parameter Problem - Sending |
| 307 | M | ICMP Redirect - Acceptance |
| 308 | O | ICMP Timestamp Reply - Acceptance |
| 309 | O | ICMP Timestamp/Timestamp Reply<br>- Acceptance of Request, Sending of Reply |
| 310 | M | ICMP Echo Reply - Acceptance |
| 311 | M | ICMP Echo/Echo Reply - Acceptance of Request,<br>- Refusal of Reply |
| 312 | M | ICMP Invalid Checksum - Refusal to Accept |
| 313 | O | ICMP Information Reply - Acceptance |
| 314 | O | ICMP Information Request/Reply<br>- Acceptance of Request, Sending of Reply |
| 315 | O | ICMP Information Request/Reply for Network Address<br>- Acceptance of Request, Sending of Reply |
| 316 | M | ICMP Source Quench - Acceptance |
| 317 | M | ICMP Source Quench - Sending |

INTERNET PROTOCOL MIL-STD 1777 PROFILE
(TCP/IP TIGHTLY COUPLED TEST SUITE)

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 101 - 137 |
| OPTIONS | | |
| | ICMP TIMESTAMP | 138 |
| | ICMP INFO | 139 |

## INTERNET PROTOCOL MIL-STD 1777 PROFILE
## (TCP/IP TIGHTLY COUPLED TEST SUITE)

| Test No. | Class | Purpose |
|---|---|---|
| **TCP/IP Tests** | | |
| 101 | M | Deliver and Send Datagram |
| 102 | M | Precedence Values - Acceptance |
| 103 | M | Low Delay - Acceptance |
| 104 | M | High Reliability - Acceptance |
| 105 | M | High Throughput - Acceptance |
| 106 | M | Type of Service Combinations - Acceptance |
| 107 | M | Illegal Time to Live - Rejection |
| 108 | M | Too Small Time to Live - Rejection |
| 109 | M | Range of Valid Time to Live Values - Acceptance |
| 110 | M | Invalid Version Number - Rejection |
| 111 | M | Invalid Checksum - Rejection |
| 112 | M | Illegally Small Header Length - Rejection |
| 113 | M | Inconsistent Header and Total Length - Rejection |
| 114 | M | Illegally Small Total Length - Rejection |
| 115 | M | Total Length Greater Than Actual Length - Rejection |
| 116 | M | Total Length Smaller Than Actual Length - Rejection |
| 117 | M | More Fragments Field - Recognition |
| 118 | M | Datagram with NOP and End-of-Options List Options - Acceptance |
| 119 | M | Datagram with 2 NOP, 1 End-of-Options List Options - Acceptance |
| 120 | M | Datagram with 3 NOP, 1 End-of-Options List Options - Acceptance |
| 121 | M | Datagram with 4 NOP Options - Acceptance |
| 122 | M | Datagram with Invalid Options - Acceptance |
| 123 | M | Reassembly of 2-Fragment Datagram |
| 124 | M | Reassembly of 3-Fragment Datagram |
| 125 | M | Reassembly of 576-Octet Datagram |
| 126 | M | Reassembly of Out-of-Order Fragments - Mixed |
| 127 | M | Reassembly of Fragments Received in Reverse Order |
| 128 | M | Expired Time to Live in Arriving Fragment - Rejection |
| 129 | M | Duplicate Fragment in Reassembly |
| 130 | M | Inconsistent Protocol Fields in Fragment Reassembly - Rejection |
| 131 | M | Inconsistent Precedence Fields in Fragment |
| 132 | M | Expiration of Time to Live during Reassembly - Rejection Reassembly - Rejection |
| 133 | M | Setting and Restarting Reassembly Timer |
| 134 | M | Reassembly of Two Intermixed Datagrams |
| 135 | M | Reassembly of Many Intermixed Datagrams |
| 136 | M | Precedence Values - Transmission |
| 137 | M | Echo and Echo Reply |
| 138 | O | Timestamp and Timestamp Reply |
| 139 | O | Information Request and Information Reply |

## TRANSMISSION CONTROL PROTOCOL MIL-STD 1778 PROFILE
## (TCP AND TCP/IP TIGHTLY COUPLED TEST SUITES)

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 8<br>12 - 23<br>25 - 37<br>52 - 54<br>59 - 73 |
| OPTIONS | FULLY SPECIFIED PASSIVE OPEN | 9 - 10 |
| | ACTIVE OPEN WITH DATA (SENDING) | 11 |
| | STATUS | 24 |
| | SECURITY | 38 - 50 |
| | ALLOC | 51 |
| | ULP TIMEOUTS | 55 - 58 |

## TRANSMISSION CONTROL PROTOCOL MIL-STD 1778 PROFILE
## (TCP AND TCP/IP TIGHTLY COUPLED TEST SUITES)

| Test No. | Class | Purpose |
|---|---|---|
| TCP-Tests | | |
| 1 | M | Unspecified Passive Open Request |
| 2 | M | Active Open Request |
| 3 | M | Basic Data Transfer |
| 4 | M | Remote Driver Interpretation of Command LCN |
| 5 | M | Determine IUT Standard Send Buffer |
| 6 | M | Closing Handshake - IUT initiates close |
| 7 | M | Closing Handshake - IUT peer initiates close |
| 8 | M | Ability to Reconnect Remote Driver Command |
| 9 | O | Fully Specified Passive Open Request |
| 10 | O | Illegal Fully Specified Passive Open Request |
| 11 | O | Active Open with Data (Sending) |
| 12 | M | Active Open with Data (Ack of Data) |
| 13 | M | Port Number Range |
| 14 | M | Graceful Closing - Completion of data transfer after ULP close |
| 15 | M | Graceful Closing - Data transfer after receipt of peer's FIN |
| 16 | M | Graceful Closing - Peer data transfer after IUT initiates close |
| 17 | M | ULP abort |
| 18 | M | Peer abort |
| 19 | M | ULP abort - Data queued for sending |
| 20 | M | Peer abort - Data queued for sending |
| 21 | M | Precedence - Mismatched |
| 22 | M | Precedence - Matched |
| 23 | M | Precedence Negotiation |
| 24 | O | Status |
| 25 | M | Out-of-Order Data |
| 26 | M | Overlapping Data |
| 27 | M | Lost Data |
| 28 | M | TCP Bad Checksum Detection |
| 29 | M | Sequence Number Wraparound |
| 30 | M | Multiplexing - Two connections with unique 4-tuple IUT opens passively |
| 31 | M | Multiplexing - Common destination port in 4-tuple common IUT port |
| 32 | M | Multiplexing - Common destination port in 4-tuple common REF port |
| 33 | M | Multiplexing - Two connections with unique 4-tuple IUT opens actively |
| 34 | M | Multiplexing - Three connections with common IUT port in 4-tuple |
| 35 | M | Duplicate Connection Attempt - IUT Passive |
| 36 | M | Multiplexing - Same sequence numbers on two connections |
| 37 | M | Duplicate Connection Attempt - IUT Active |

| Test No. | Class | Purpose |
|---|---|---|
| TCP-Tests | | |
| 38 | 0 | Setting Security in Active Open |
| 39 | 0 | Setting Security in Passive Open |
| 40 | 0 | Setting Security in Fully Specified Passive Open |
| 41 | 0 | Secure IUT rejecting connection to unsecured peer |
| 42 | 0 | Secure IUT rejecting connection from unsecured peer |
| 43 | 0 | Security option placement in sending data |
| 44 | 0 | Response to data with mismatched security class |
| 45 | 0 | Response to data with mismatched security protection authority |
| 46 | 0 | Response to data with extra protection authority |
| 47 | 0 | Use of security option for unclassified connections |
| 48 | 0 | Recognition of UNCLASS and GENSER as unsecured |
| 49 | 0 | Unsecured IUT response to connection attempt by secured host |
| 50 | 0 | Unsecured IUT response to data marked with classified security |
| 51 | 0 | Alloc |
| 52 | M | Maximum segment size option |
| 53 | M | Retransmission after acknowledgement of data |
| 54 | M | Retransmission after acknowledgement of SYN and FIN |
| 55 | 0 | ULP timeout service in Active Open |
| 56 | 0 | ULP timeout service in Send |
| 57 | 0 | ULP timeout service in Passive Open |
| 58 | 0 | ULP timeout notify action tested |
| 59 | M | TCP in window mechanism |
| 60 | M | Urgent service |
| 61 | M | Urgent service when peer has zero window |
| 62 | M | Urgent data delivery |
| 63 | M | Push service - Service not requested |
| 64 | M | Push service - Service requested |
| 65 | M | Reset - as response to connection refusal |
| 66 | M | Reset - partial rest prior to connection establishment |
| 67 | M | Reset - response to reset received while sending data |
| 68 | M | Reset segment format on receipt of Active Open with no listening port |
| 69 | M | Reset segment format on receipt of Active Open with data with no listening port |
| 70 | M | Reset segment format on receipt of invalid segment with ACK set |
| 71 | M | Reset segment format on receipt of invalid segment with SYN and ACK set |
| 72 | M | Reset - no reset sent on receipt of segment with bad acknowledgment number |
| 73 | M | Determine number of connections resources will allow |

FILE TRANSFER PROTOCOL MIL-STD 1780
SERVER PROFILE

| | | |
|---|---|---|
| | | TESTS |
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 7 |
| | | 10 - 16 |
| | | 19 - 21 |
| | | 23 - 43 |
| | | 44 - 45* |
| | | 47* |
| | | 48 - 53 |
| | | 55 - 69 |
| | | 75 |
| OPTIONS | | |
| | ALLO COMMAND | 8 - 9 |
| | REIN COMMAND | 17 |
| | ABOR COMMAND | 18 |
| | SITE COMMAND | 22 |
| | COMPRESSED MODE | 54 |
| | 3-WAY TRANSFER | 70 - 74 |

* Note: Tests 44,45, and 47 are mandatory only if the command is not implemented in the server. Test 46 is invalid since the implementation of the STAT command is mandatory.

FILE TRANSFER PROTOCOL MIL-STD 1780
SERVER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| FTP-Tests | | |
| 1 | M | Connection establishment |
| 2 | M | Response to USER command |
| 3 | M | Response to PASS command |
| 4 | M | Response to ACCT command |
| 5 | M | Response to QUIT command |
| 6 | M | Response to APPE command |
| 7 | M | Response to APPE nonexistent file |
| 8 | O | Response to ALLO command without record option |
| 9 | O | Response to ALLO command with record option |
| 10 | M | Response to RNFR command |
| 11 | M | Response to RNTO command |
| 12 | M | Response to LIST command with file parameter |
| 13 | M | Response to LIST command without parameter |
| 14 | M | Response to NLST command |
| 15 | M | Response to CWD command |
| 16 | M | Response to DELE command |
| 17 | O | Response to REIN command |
| 18 | O | Response to ABOR command |
| 19 | M | Response to STAT command without parameter |
| 20 | M | Response to STAT command with directory param |
| 21 | M | Response to STAT command with file parameter |
| 22 | O | Response to SITE command |
| 23 | M | Response to HELP command |
| 24 | M | Response to HELP command with parameter |
| 25 | M | File transfer (default transfer params: SFA) |
| 26 | M | File transfer (explicit SFA) |
| 27 | M | File transfer (explicit SFI) |
| 28 | M | File transfer (explicit SFE) |
| 29 | M | File transfer (explicit SRA) |
| 30 | M | File transfer (explicit SRI) |
| 32 | M | Response to NOOP command |
| 33 | M | Response to STOR command when not logged in |
| 34 | M | Response to unknown User |
| 35 | M | Response to out-of-sequence ACCT command |
| 36 | M | Response to out-of-sequence PASS command |
| 37 | M | Response to known User and incorrect Password |
| 38 | M | Response to incorrect APPE command |
| 39 | M | Response to RNFR of nonexistent file |
| 40 | M | Response to incorrect LIST command |
| 41 | M | Response to incorrect NLST command |
| 42 | M | Response to CWD to nonexistent directory |
| 43 | M | Response to DELE of nonexistent file |
| 44 | M | Response to unimplemented REIN command |
| 45 | M | Response to unimplemented ABOR command |
| 46 | M | Response to unimplemented STAT command |
| 47 | M | Response to unimplemented SITE command |

| Test No. | Class | Purpose |
|---|---|---|
| **FTP-Tests** | | |
| 48 | M | Response to HELP command with unknown parameter |
| 49 | M | Response to incorrect STOR command |
| 50 | M | Response to incorrect RETR command |
| 51 | M | Response to incorrect STRU command |
| 52 | M | Response to incorrect TYPE command |
| 53 | M | Response to incorrect MODE command (Block) |
| 54 | O | Response to incorrect MODE command (Compressed) |
| 55 | M | Response to command syntax error |
| 56 | M | Response to APPE command corruption |
| 57 | M | Response to DELE command corruption |
| 58 | M | Response to LIST command corruption |
| 59 | M | Response to NLST command corruption |
| 60 | M | Response to Rename command sequence corruption |
| 61 | M | Response to RETR command corruption |
| 62 | M | Response to RNFR command corruption |
| 63 | M | Response to RNTO corruption |
| 64 | M | Response to STOR command corruption |
| 65 | M | Response to ABOR command corruption |
| 66 | M | Verification of file transfer via checksum - STOR |
| 67 | M | Response to HELP command corruption |
| 68 | M | Response to SITE command corruption |
| 69 | M | Response to STAT command corruption |
| 70 | O | File transfer over multiple connections |
| 71 | O | Establishment of multiple connections |
| 72 | O | Three-way file transfer |
| 73 | O | Response to QUIT command in three-way context |
| 74 | O | Response to unimplemented PASV command |
| 75 | M | Response to PORT command |

FILE TRANSFER PROTOCOL MIL-STD 1780
USER PROFILE

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 3<br>5 - 10<br>13<br>15 - 22<br>35<br>40 |
| OPTIONS | ALLO COMMAND | 4 |
| | REIN COMMAND | 11 |
| | ABOR COMMAND | 12 |
| | SITE COMMAND | 14 |
| | BLOCK TRANSFER | 23 - 28 |
| | COMPRESSED TRANSFER | 29 - 34 |
| | MULTIPLE CONNECTIONS | 36 - 39 |

FILE TRANSFER PROTOCOL MIL-STD 1780
USER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| FTP-Tests | | |
| 1 | M | Connection establishment and login (USER, PASS, ACCT commands) |
| 2 | M | Connection closing (QUIT command) |
| 3 | M | Transmission of APPE command |
| 4 | O | Transmission of ALLO command |
| 5 | M | Transmission of RNFR command |
| 6 | M | Transmission of RNTO command |
| 7 | M | Transmission of LIST command |
| 8 | M | Transmission of NLST command |
| 9 | M | Transmission of CWD command |
| 10 | M | Transmission of DELE command |
| 11 | O | Transmission of REIN command |
| 12 | O | Transmission of ABOR command |
| 13 | M | Transmission of STAT command |
| 14 | O | Transmission of SITE command |
| 15 | M | Transmission of HELP command |
| 16 | M | File transfer (default transfer params: SFA) |
| 17 | M | File transfer (explicit SFA) |
| 18 | M | File transfer (explicit SFE) |
| 19 | M | File transfer (explicit SFI) |
| 20 | M | File transfer (explicit SRA) |
| 21 | M | File transfer (explicit SRE) |
| 22 | M | File transfer (explicit SRI) |
| 23 | O | File transfer (explicit BFA) |
| 24 | O | File transfer (explicit BFE) |
| 25 | O | File transfer (explicit BFI) |
| 26 | O | File transfer (explicit BRA) |
| 27 | O | File transfer (explicit BRE) |
| 28 | O | File transfer (explicit BRI) |
| 29 | O | File transfer (explicit CFA) |
| 30 | O | File transfer (explicit CFE) |
| 31 | O | File transfer (explicit CFI) |
| 32 | O | File transfer (explicit CRA) |
| 33 | O | File transfer (explicit CRE) |
| 34 | O | File transfer (explicit CRI) |
| 35 | M | Transmission of NOOP command |
| 36 | O | Second connection extablishment and login (USER, PASS, ACCT) |
| 37 | O | Switch connection, PASV command |
| 38 | O | Switch connection, PORT command |
| 39 | O | Ability to close connection to first reference (QUIT command) |
| 40 | M | Transmission of PORT command |

## SIMPLE MAIL TRANSFER PROTOCOL MIL-STD 1781
## SENDER PROFILE

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1<br>5 - 8 |
| OPTIONS | SOML COMMAND | 2 |
| | SAML COMMAND | 3 |
| | SEND COMMAND | 4 |

## SIMPLE MAIL TRANSFER PROTOCOL MIL-STD 1781
## SENDER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| SMTP-Tests | | |
| 1 | M | Basic mail transaction command sequence (HELO, MAIL, RCPT, DATA, QUIT) |
| 2 | O | Transmission of SOML command |
| 3 | O | Transmission of SAML command |
| 4 | O | Transmission of SEND command |
| 5 | M | Transmission of all printable ASCII characters |
| 6 | M | Response to incorrect command reply syntax |
| 7 | M | Response to incorrect command reply code |
| 8 | M | Response to excessively long command reply |

SIMPLE MAIL TRANSFER PROTOCOL MIL-STD 1781
RECEIVER PROFILE

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 11<br>13 - 18<br>23 - 24<br>33<br>34<br>37 - 38<br>40 - 42 |
| OPTIONS | MAIL RELAY | 20 - 22, 43 |
| | RSET following SAML | 25 |
| | RSET following SOML | 26 |
| | SEND COMMAND | 27 - 28 |
| | SOML COMMAND | 29 |
| | SAML COMMAND | 30 |
| | TURN COMMAND | 32 |
| | EXPN COMMAND | 35 - 36 |

## SIMPLE MAIL TRANSFER PROTOCOL MIL-STD 1781
## RECEIVER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| SMTP-Tests | | |
| 1 | M | Response to HELO-QUIT sequence |
| 2 | M | Response to multiple HELO commands |
| 3 | M | Basic mail transaction command sequence (HELO, MAIL, RCPT, DATA, QUIT) |
| 4 | M | Ability to prefix reverse path to message |
| 5 | M | Ability to prefix null reverse path to message |
| 6 | M | Notification of Undeliverable Mail |
| 7 | M | Ability to prepend timestamp to mail message |
| 8 | M | Response to data transparency sequence |
| 9 | M | Ability to process all printable ASCII characters |
| 10 | M | Response to NOOP command outside a transaction |
| 11 | M | Response to NOOP command during a transaction |
| 13 | M | Response to excessively long command line |
| 14 | M | Response to excessively long test line |
| 16 | M | Response to invalid command sequence |
| 17 | M | Response to nonexistent command |
| 18 | M | Response to incorrect command syntax |
| 20 | O | Ability to relay message |
| 21 | O | Ability to relay multiple messages to multiple recipients |
| 22 | O | Response to MAIL requiring relay to nonexistent recipients |
| 23 | M | Response to RSET command during a transaction |
| 24 | M | Response to RSET command following RCPT |
| 25 | O | Response to RSET command following SAML |
| 26 | O | Response to RSET command following SOML |
| 27 | O | Response to SEND command |
| 28 | O | Response to SEND to nonexistent recipient |
| 29 | O | Response to SOML command |
| 30 | O | Response to SAML command |
| 32 | O | Response to TURN command and ability to fulfill basic Sender SMTP role |
| 33 | M | Response to VRFY command outside a transaction |
| 34 | M | Response to VRFY command during a transaction |
| 35 | O | Response to EXPN command outside a transaction |
| 36 | O | Response to EXPN command during a transaction |
| 37 | M | Response to HELP command outside a transaction |
| 38 | M | Response to HELP command during a transaction |
| 40 | M | Response to various hostname formats |
| 41 | M | Response to multiple RCPT commands |
| 42 | M | Ability to process multiple, simultaneous mail transactions |
| 43 | O | Ability to relay message through multiple hosts |

TELNET PROTOCOL MIL-STD 1782
USER PROFILE

| | | TESTS |
|---|---|---|
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 8<br>11 - 12<br>14 - 22 |
| OPTIONS | NETWORK VIRTUAL TERMINAL GO AHEAD STATUS | 22 |
| | GENERATION OF EC | 9 |
| | GENERATION OF EL | 10 |
| | GENERATION OF BRK | 13 |
| | GENERATION OF GA | 19 |

TELNET PROTOCOL MIL-STD 1782
USER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| TELNET-Tests | | |
| 22 | M | Open Connection |
| 22 | O | Network Virtual Terminal Go Ahead Status |
| 22 | M | Network Virtual Terminal Echo Status |
| | | Response To DO Option Request |
| 1 | M | Remote echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response To DO Option Request for Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to DON'T Option Request for Disabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to WILL Option Request for Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| | | Response to WILL Option Request for Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |

| Test No. | Class | Purpose |
|---|---|---|
| TELNET-Tests | | |
| | | Response to WONT Option--Announcement of Disabling |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to Request to Enable Option Peer Has Enabled |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Correct Implementation of Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 6 | M | Status |
| | | Corrected implementation after Disabling Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Status |
| | | Correct Implementation of Option When Both Sides Enabled |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 6 | M | Status |
| 7 | M | Generation of Synch Embedding Data Mark (DM) |
| 8 | M | Generation of Are You There Command (AYT) |
| 9 | O | Generation of Erase Character Command(EC) |
| 10 | O | Generation of Erase Line Command (EL) |
| 11 | M | Generation of Abort Output Command (AO) |
| 12 | M | Generation of No Operation (NOP) |
| 13 | O | Generation of Break Command (BRK) |
| 14 | M | Transmission of ASCII Printable Characters |
| 15 | M | Transmission and Receipt of Newline (CRLF) |

| Test No. | Class | Purpose |
|---|---|---|
| **TELNET-Tests** | | |
| 16 | M | Transmission and Receipt of Carriage Return (CR NULL) |
| 17 | M | Generation of Interrupt Process Command (IP) |
| 18 | M | Non-transmission of request for previously refused option - Binary |
| 19 | O | Generation of GoAhead Command (GA) |
| 20 | M | Non-transmission of subnegotiation for Disabled option - Status |
| 21 | M | Close Connection |

TELNET PROTOCOL MIL-STD 1782
SERVER PROFILE

| | | |
|---|---|---|
| | | TESTS |
| MANDATORY IMPLEMENTATION REQUIREMENTS | | 1 - 8<br>11 - 18<br>19 |
| OPTIONS | NETWORK VIRTUAL TERMINAL GO AHEAD STATUS | 19 |
| | RESPONSE TO EC | 9 |
| | RESPONSE TO EL | 10 |

TELNET PROTOCOL MIL-STD 1782
SERVER PROFILE

| Test No. | Class | Purpose |
|---|---|---|
| **TELNET-Tests** | | |
| 19 | M | Open connection |
| 19 | O | Network Virtual Terminal Go Ahead Status |
| 19 | M | Network virtual Terminal Echo Status |
| | | Response to DO Option Request |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to DO Option Request for Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to DON'T Option Request for Disabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to WILL Option Request for Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoHead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Response to WONT Option--Announcement of Disabling |
| 1 | M | Remote echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |

| Test No. | Class | Purpose |
|---|---|---|
| **TELNET-Tests** | | |
| | | Response to Request to Enable Option Peer Has Enabled |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Timing Mark |
| 5 | M | Extended Options |
| 6 | M | Status |
| | | Correct Implementation of Enabled Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 6 | M | Status |
| | | Correct Implementation After Disabling Option |
| 1 | M | Remote Echo |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 4 | M | Status |
| | | Correct Implementation of Option When Both Sides Enabled |
| 2 | M | GoAhead |
| 3 | M | Binary |
| 6 | M | Status |
| 7 | M | Response to Synch |
| 8 | M | Response to Are You There (AYT) |
| 9 | O | Response to Erase Character (EC) |
| 10 | O | Response to Erase Line (EL) |
| 11 | M | Receipt of ASCII Printable Characters |
| 12 | M | Transmission and Receipt of Newline (CRLF) |
| 13 | M | Transmission and Receipt of Carriage Return (CR NULL) |
| 14 | M | Response to No Operation (NOP) |
| 15 | M | Response to Data Mark (DM) with No TCP Urgent |
| 16 | M | Response to Abort Output (AO) |
| 17 | M | Response to Interrupt Process (IP) |
| 18 | M | Close connection |