The INTEROP 88 Network: Design, Problems, and Lessons Learned*

Philip Almquist

* WARNING: do not try this at home. Professional stunt driver required.

Introduction

- Large scale demonstration of TCP/IP Interoperability
  - 49 vendors
  - Approximately 250 hosts and gateways
  - Almost 2 miles of cabling
  - High-speed connections to ARPANet, MILNet, NSFNet, ...
- Standalone network for CMOT (NETMAN) demonstration
- Very successful
- Purposes of this talk
  - Inform
  - Stimulate IETF action

Description of the network

- Designed by Peter DeVries and myself
- Subnetted class B net
- Multiple media
    - Ethernet
    - Thin Ethernet
    - Ethernet over twisted pair
    - Ethernet over fiber
    - PRONet-80
    - IBM/802.5 token ring
    - SLIP
    - Packet radio
    - (also Hyperchannel, PRONet-10, T-1, and Ethernet over broadband in individual booths)
- Tree topology - no alternate routes
- Small subnets
- All backbone routers in NOC
- Built in 5 1/2 days by Peter, myself, 3 part-time technicians, and a horde of volunteers

INTEROP 88

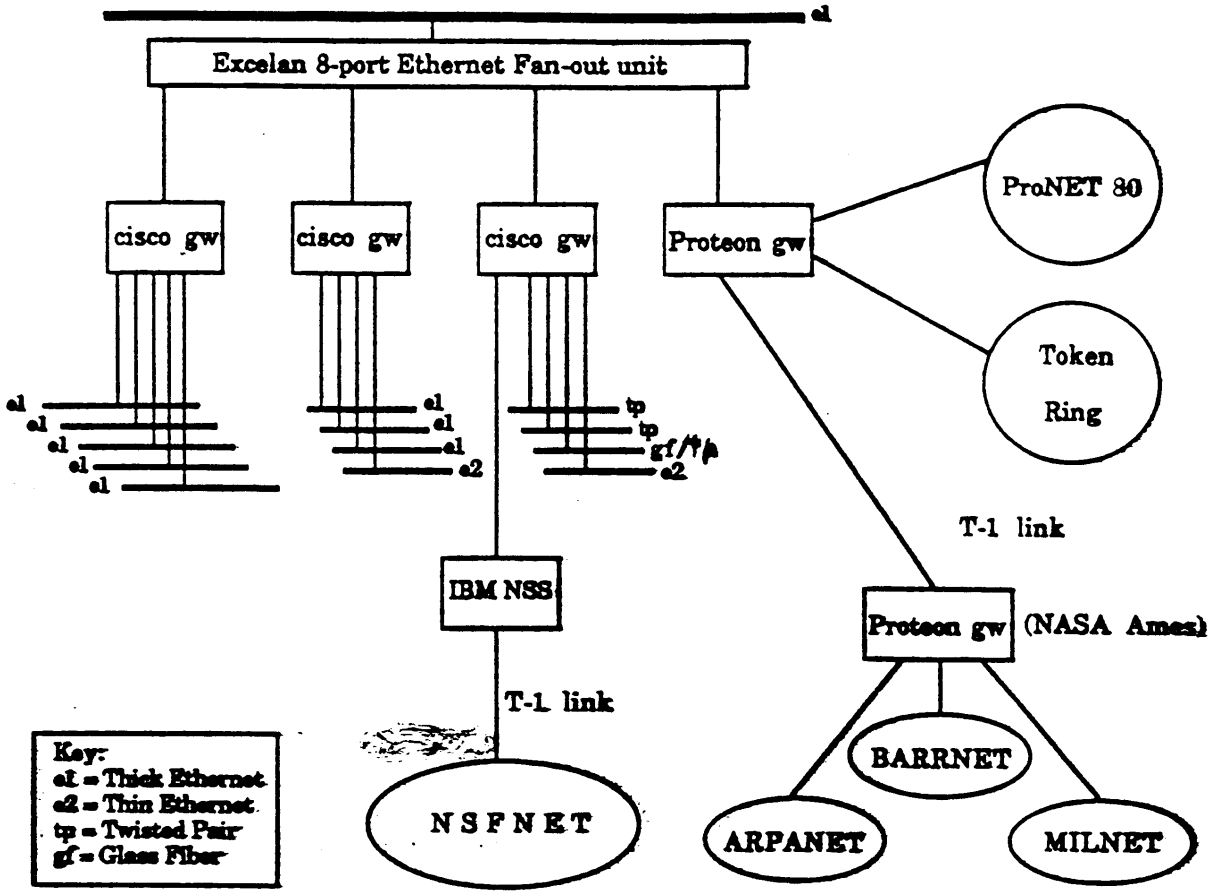

Show and Tel-Net Topology

Participating Vendors:

3Com
ACC
Apple Computer
Banyan Systems
BBN Communications
COMPUTERWORLD
CMC
Computer Network Technology
Concurrent Computer
Convergent Technologies
cisco Systems
DCA/SRI International
DEC
Encore
Eon Systems
Excelan/TGV/Kinetics
FTP Software
Hallay Systems
Hewlett-Packard
Highland Software
IBM/MCI/Merit/CMU
Interactive Systems
InterCon
Interphase
Lachman Associates
Mitre/Unisys (NetMan)
Network General
Network Research
Network Solutions
Network Systems
Prentice-Hall
Prime Computer
Process Software
Proteon
Sirius Systems
Spider Systems
Sun Microsystems
SynOptics Communications
Syntax Systems/10Net
Sytek
Tandem Computers
TCL
TRW
Ungermann-Bass
UNIX World
Vitalink Communications
VXM Technologies/MIPS
Wellfleet Communications
Western Digital
The Wollongong Group
Xyplex

INTEROP '88
Show and Tel-Net
September 28-30, 1988
Santa Clara Convention Center

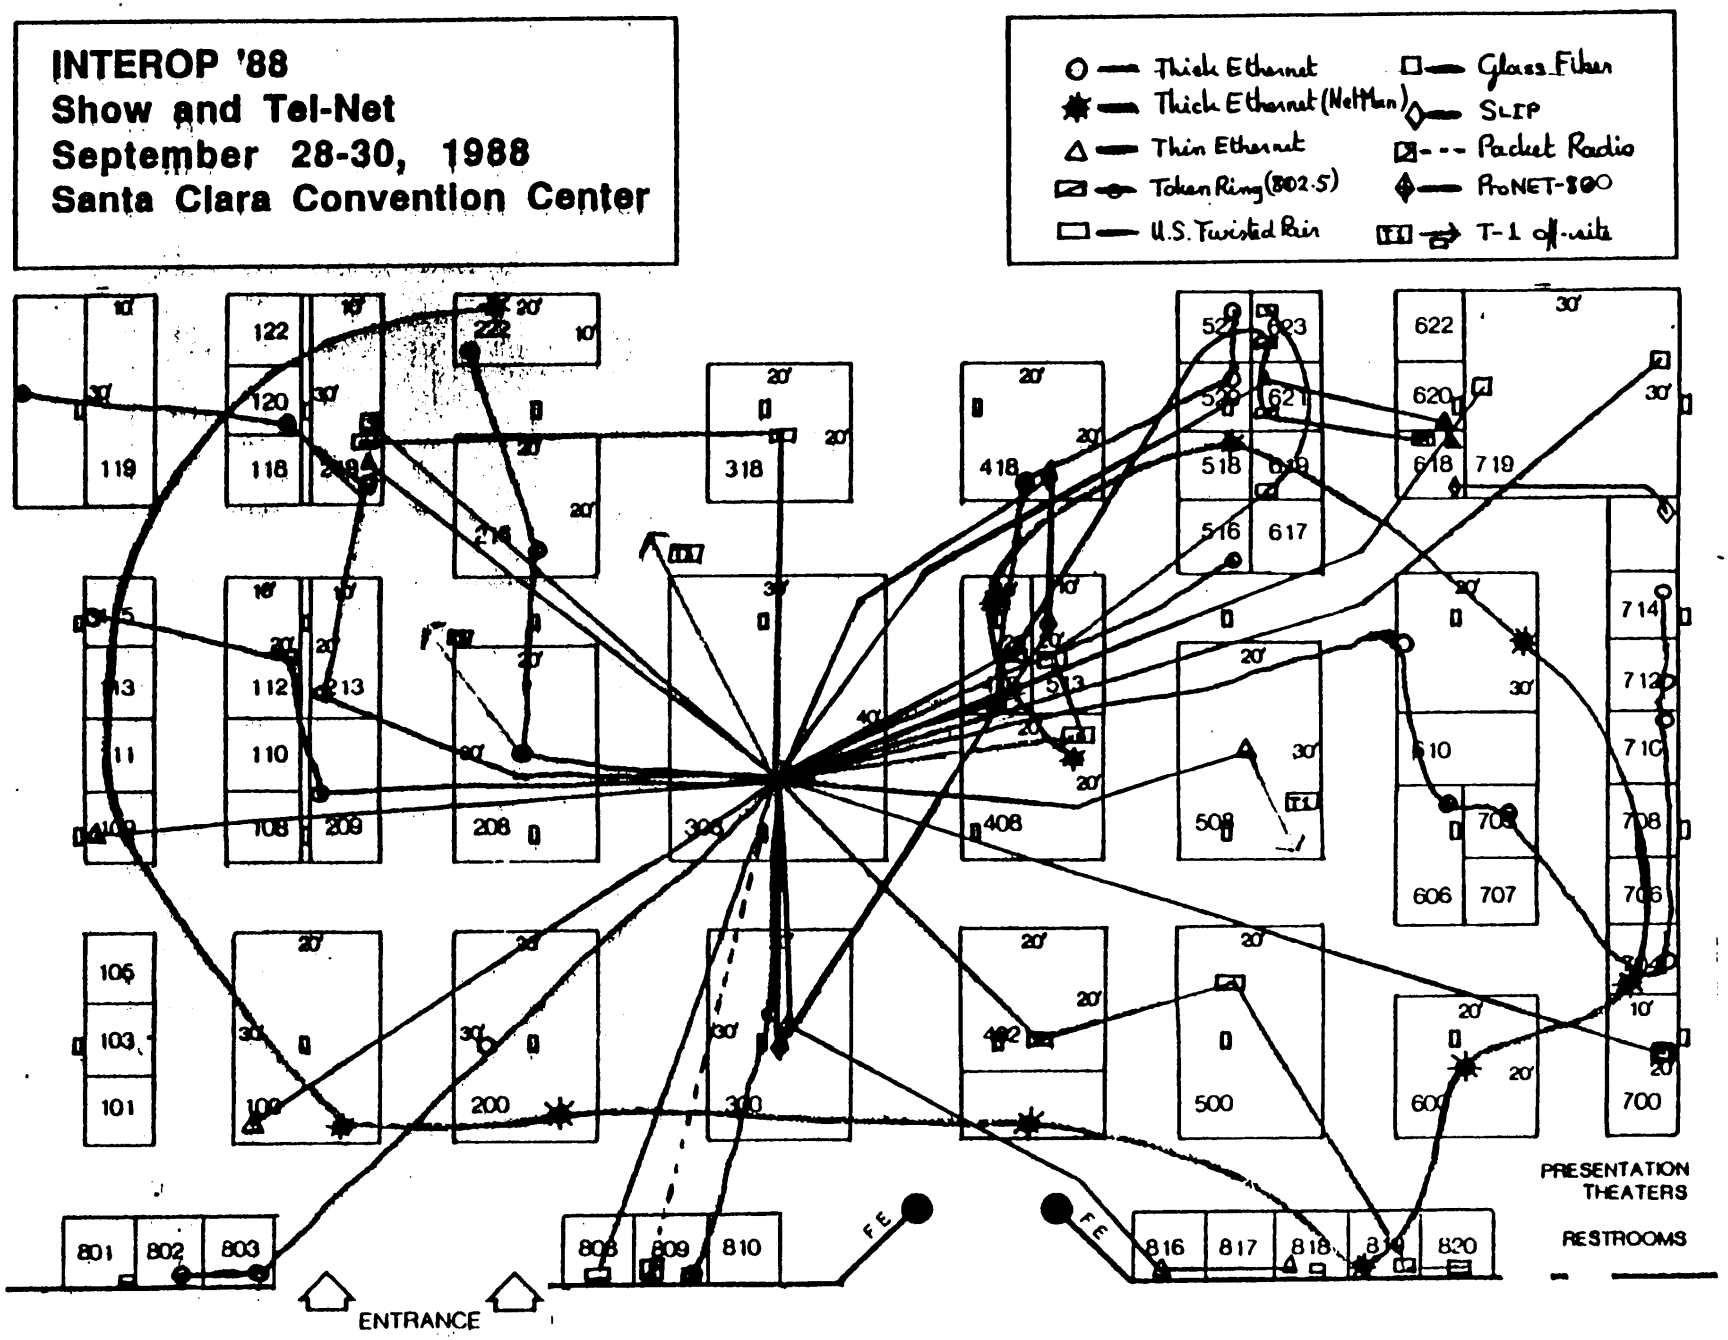

Sponsored by Advanced Computing Environments, Inc.

Designed, installed and managed by The Wollongong Group, Inc.

Cabling

- What we did
  - Cabling hung from ceiling
  - Intentionally very visible
  - Tranceivers reachable with a ladder
- Problems
  - Ran out of cable
  - T-1 didn't want to work (of course!)
  - Too many people inside the wiring center
  - One booth on wrong subnet because vendor rewired it!
  - Mysterious temporary failure of one Ethernet segment on second day of show
  - The usual minor problems...

IP address assignment/host table creation

- What we did
  - We obtained a domain: ShowNet.COM
  - Vendors filled out host questionnaires
  - We assigned IP addresses and created a zone file
  - A program read the zone file to generate the IN-ADDR.ARPA zone files and a HOSTS.TXT
- Problems
  - Questionnaires were returned late and filled out incorrectly
  - No host table czar
  - Zone file inaccessible until T-1 came up
  - Some vendors required /etc/hosts format

Domain service

- What we did
  - 3 authoritative servers (two off-site)
  - Off-site servers set up as secondaries
  - Small TTL's and refresh times
- Problems
  - Syntax errors in the master files
  - Little familiarity with domain software on primary
  - Miscommunication between the NIC and Wollongong
  - Root server update procedure failed
  - Primary not installed until the day before the show

Lessons

- Make sure domain requests get honored well before you need them
- Root server updates are probably not as robust as they should be
- Hand-typed zone files require a syntax checker program

Network Management

- What we did
  - SUN running Wollongong/NYSERNet SNMP tools
  - Protocol analyzer
  - Smart Ethernet terminator
- Problems
  - pre-SNMP code on cisco routers the first day
  - bug in Proteon SNMP
  - SUN had incomplete/incorrect SNMP configuration files
  - Most segments didn't have extra tranceivers for monitoring
  - NOC personnel unfamiliar with the particular management tools available
- Lessons
  - Network management tools are useless if they can't be used quickly and easily when problems occur

Internet Protocol Police

Notice of Protocol Violation

IP Address of Offender: _____

Domain Name of Offender: _____

Improper Configuration

- ☐ Wrong IP Address
- ☐ Wrong IP broadcast address
- ☐ Wrong Subnet Mask (or subnets not supported)
- ☐ Excessive Broadcasting
- ☐ ARPing for Broadcast Address
- ☐ Invalid Ethernet/Subnet address
- ☐ _____

Warnings

- ☐ Disabling UDP checksums
- ☐ Dropping packets while resolving addresses
- ☐ Tinygram generation
- ☐ Improper round-trip-timing
- ☐ Lack of congestion avoidance
- ☐ _____

Protocol Violations

- ☐ Forwarding broadcast packets
- ☐ TCP response to broadcast
- ☐ ICMP response to broadcast
- ☐ Ignoring ICMP redirects
- ☐ Ignoring ICMP source quench
- ☐ Broadcast TCP packets
- ☐ TCP Keepalives
- ☐ TCP aborts on ICMP message while connected
- ☐ Misc. protocol error
  - ☐ TCP ☐ IP
  - ☐ UDP ☐ ICMP
  - ☐ ARP ☐ _____
- ☐ _____

Inspector: _____ **Date:** ___/___/___

Internal routing

- What we planned
  - Use RIP throughout
  - Back doors were allowed only if not advertised
- What we actually did
  - Core routers sent all routes via RIP
  - Core routers believed RIP only from other core routers
  - Core routers had static routes to subnets behind non-core routers
  - Hosts and non-core routers to avoid RIP and use a static default route
  - Reasoning: possible bogus routes from misconfigured RIP-speakers
- Problems
  - Large and unnecessary RIP broadcasts (from NSFNet routes) caused prolems for PC's
- Lessons
  - Static routing is a b*tch

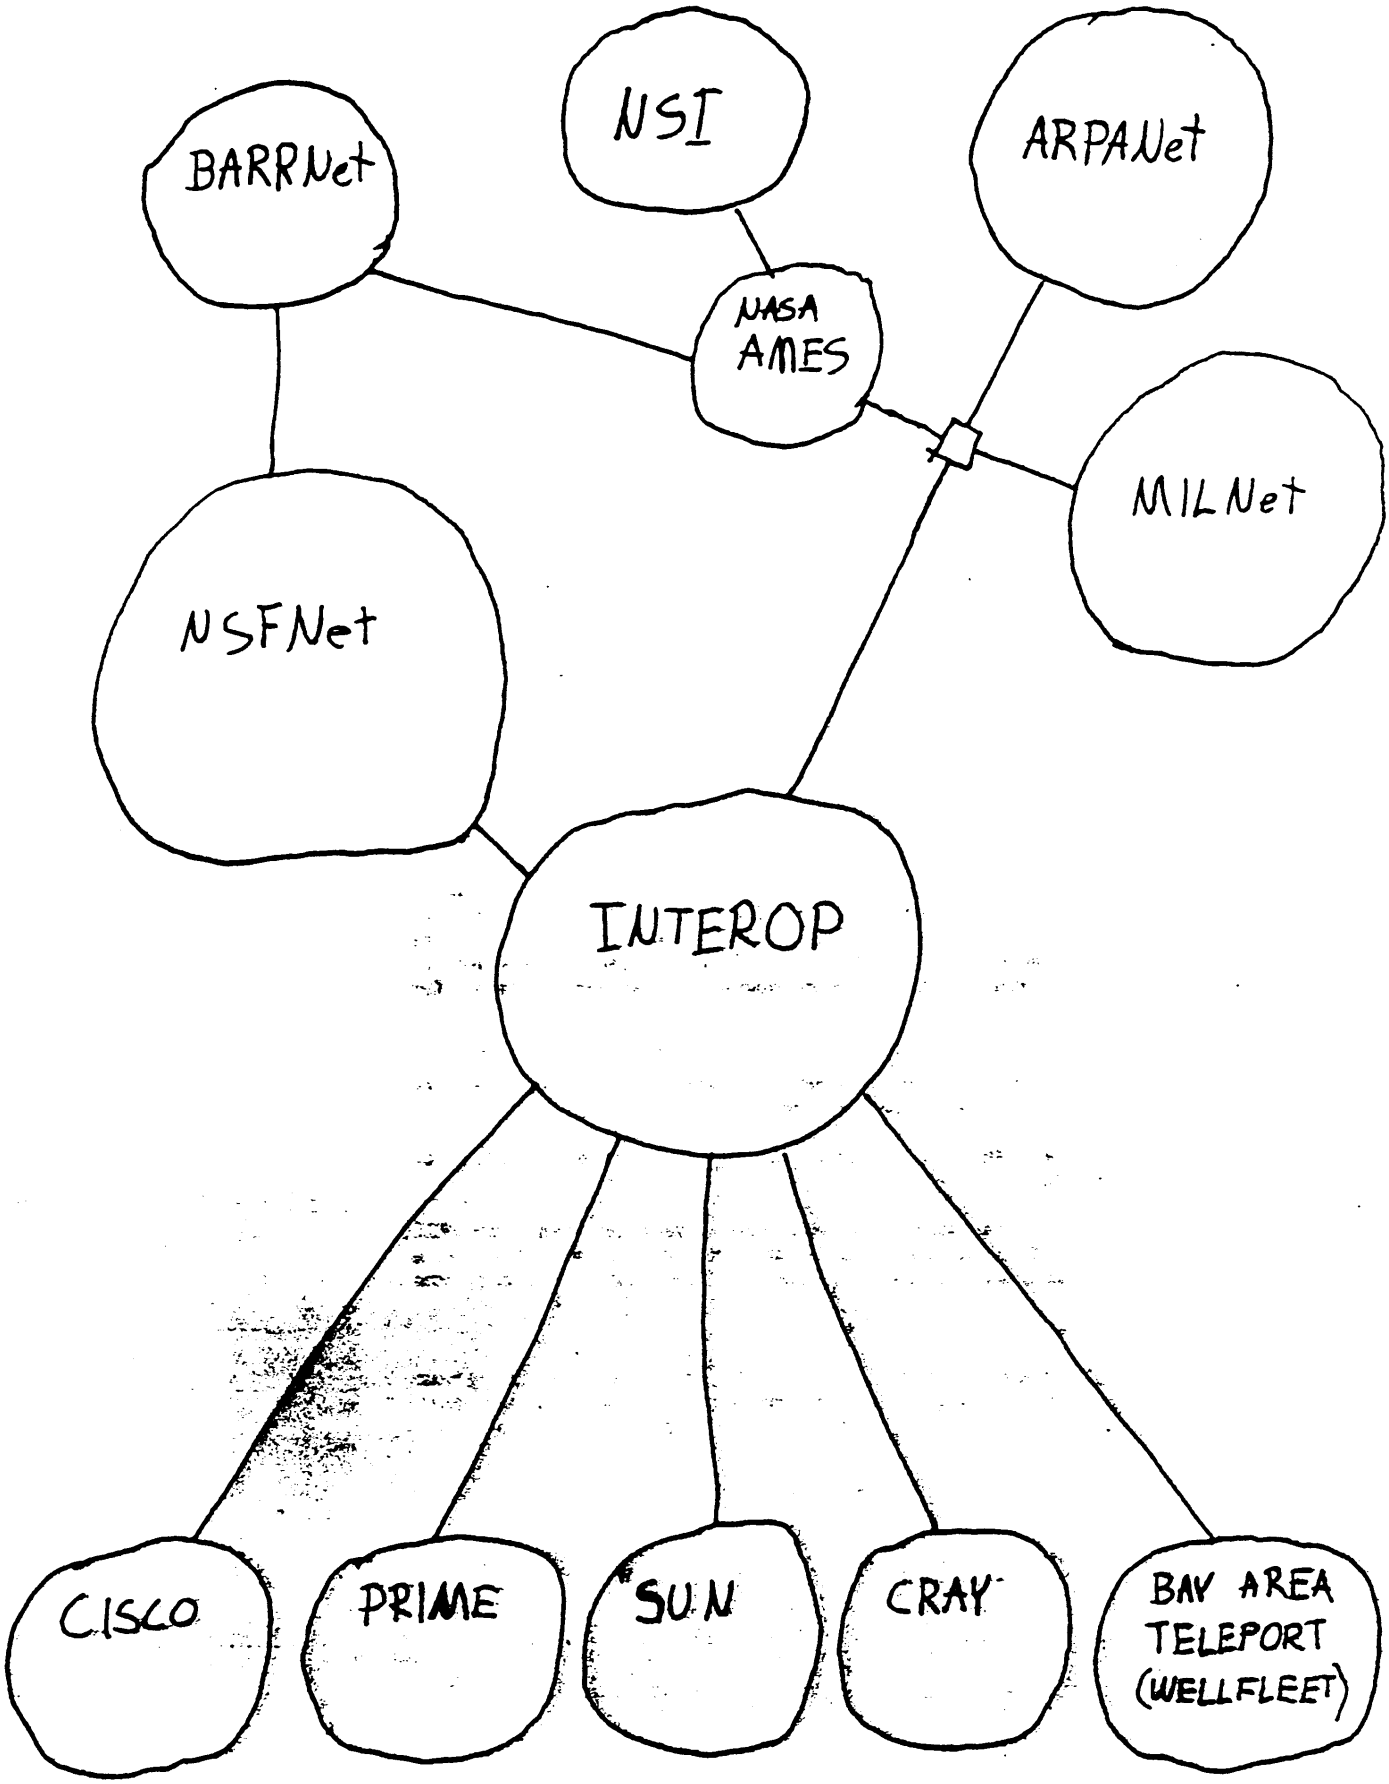
BARRNet
NSI
ARPANet
NASA AMES
MILNet
NSFNet
INTEROP
CISCO
PRIME
SUN
CRAY
BAY AREA TELEPORT (WELLFLEET)

External routing

- What we did
  - T-1 between core Proteon and AMES ARPANet/MILNet gateway
  - static routing over T-1
  - Proteon advertised RIP default
  - static routes to cisco, Prime, SUN Cray, Bay Area Teleport
  - Explicit RIP routes for NSFNet routes through IBM's NSS
- Lessons
  - cisco routers ignore RIP default

External routing - NSFNet

- What we did
  - NSFNet NSS in IBM booth
  - Secondary NSFNet path through BARRNet
  - IBM "subnet" was a class C net so EGP could treat it differently
  - PC/RT in IBM booth EGP peered with NSS and distributed RIP routes on the class C net
  - cisco core gateway also EGP peered with the NSS and distributed RIP routes on the class B net
  - Result: routing policy decisions by IBM and the NOC were independent of each other
  - NOC policy decision: always believe NSF routes (except for one afternoon when the NSFNet T-1 was flapping)

- Problems
  - We started out the show running old cisco code without NSFNet fixes to EGP
  - The NOC policy decision somewhat controversial...
  - Black holes occurred due to bad mixtures of static routes and firewalls in some of the regionals
- Lessons
  - Because of firewalls, it is dangerous to add a network to NSFNet without informing the regional networks.

Disappointments

- Network took one day too long to build
  - No time for interoperability testing
  - Network management not set up
  - No time for packet watching
  - Vendors pretty much left to sink or swim on their own
  - Network would have been more solid if it had run for a day before the show

Things I was particularly happy about

- It worked well enough...
- We got a tremendous amount of help from the Internet community

The reasons it all worked

Rick Boivie
Len Bosack
David Bridgham
Eric Brunner
Jeff Burgan
Myu Campbell
Mario Castro
Shelly DeVries
Steve Knowles
Susan Hares
Alex Latzko
Sandy Lerner
Milo Medin
Robert Michaels
Paul Mockapetris
Mike Moesler
Vince Raya
Sue Romano
Greg Satz
Mick Scully
Jim Shimoto
Mike St. Johns
James VanBokken
John Veizades

People who contributed to this talk

Peter DeVries
Milo Medin