Network Management

Who's the boss?

You or the network?

Parts of Network Management

- network management is not just keeping bits moving
- OSI network management components
  - fault management
  - performance management
  - security management
  - configuration and name management
  - accounting management
- add policy-based management
- reporting

Fault Management

- detect network problems
  - transient/persistent
  - failure/overload
- detect server problems
- isolating problems
- reporting mechanism
  - link to help desk
  - notify on-call personnel
- setup & control alarm procedures
- repair/recovery procedures
- ticket system

Fault Management - Ticket System

- system provides for:
  - short term memory & communication
  - scheduling and work assignment
  - referrals and dispatching
  - oversight
  - statistical analysis
  - long term accountability

Fault Management - Ticket Usage

- create a ticket on ALL calls
- create a ticket on ALL problems
- create a ticket for ALL scheduled events
- copy of ticket mailed to reporter and mailing list(s)
- all milestones in resolution of problem create a new ticket entry with reference to original

Fault Management - Ticket Example

```
From nearnet-ops-request@nic.near.net Fri Dec 14 13:12:56 1990
Received: from nic.near.net by nic.near.net id aa22499; 14 Dec
Date:     Fri, 14 Dec 90 13:01:42 EST
From: ops@nic.near.net
Subject:  NEARnet Ticket #1582
To: nearnet-ops@nic.near.net, nearnet-outages@nic.near.net,
tmurphy@athena.mit.edu
Status: R
Ticket Number: 1582                    Ticket Status:   open
Ticket Type:   unplanned               Ticket Source:   email
Ticket Scope:  host                    Site/Line:       mit
Ticket Owner:  perfetti                Problem Fixer:
Ticket Opened: 12/14/90 12:56          Problem Started: 12/14/90 11:20
Problem Description:
   User is experiencing difficulty in reaching a host located at Rutgers.  The
   host in question is quartz.rutgers.edu (128.6.4.8).  This problem is being
   investigated.
   For a complete history of this ticket, do "finger ticket-
   1582@nic.near.net".
```

Performance Management

- evaluate the behavior of network elements
- information used in planning
  - interface stats
  - throughput
  - error rates
  - software stats
  - usage
  - queues
  - system load
  - disk space
  - availability per cent
  - response time

Security Management

- security required to operate network and protect managed objects
- security services
  - Kerberos
  - PGP key server
  - secure time
- security tools
  - cops - host configuration checker (www.cert.org)
- distribute security information
  - bug reports
  - bug fixes
  - intruder alerts

Security Management, cont.

- reporting procedure for security events

  e.g. break-ins

- control internal and external gateways

  control firewalls (external and internal)

- security logs
- privacy issues can be a conflict

Configuration and Name Management

- network state information
  - network topology
  - operation status of network elements
    - including resources
  - network element configuration
- control network elements
  - start/stop
  - modification of network attributes
  - addition of new features
- configuration modification
  - allocation and addition of network resources
  - reconfiguration if dictated by link outages

Config. Mgmt. - Network State Info.

- e.g. SNMP driven display

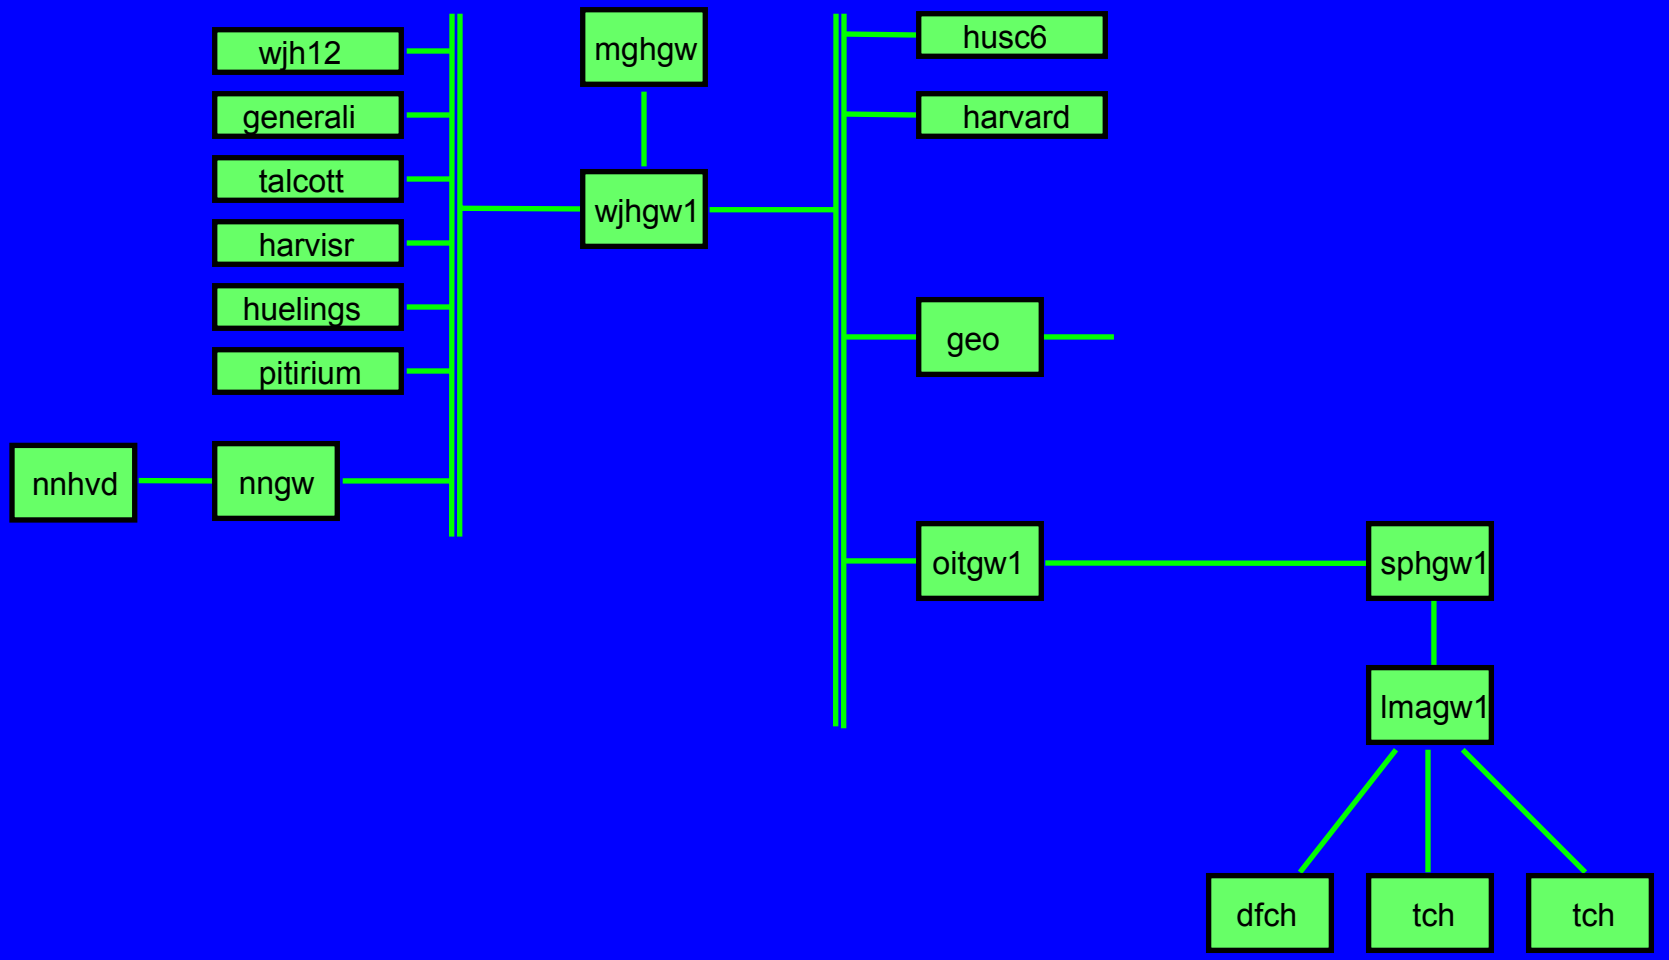

SNMP

- Simple Network Management Protocol
- mostly a query - response system
- little network traffic initiated by agent
- currently only a primitive security system
  - SNMPv2 was to have real security but working group fragmented, SNMPv3 now ready
- uses database defined in MIB
- can have "enterprise" extensions to MIB
- SMI defines structure of MIB
- SMI defines data structure using ASN.1

ASN.1

ISO standard

- specification of Abstract Syntax Notation One (ASN.1)

  defines a language used to describe data types

- specification of Basic Encoding Rules for Abstract Notation One (ASN.1)

  defines a method for unambiguous transmission of data

- machine architecture independent.
- operating system independent.
- network protocol independent.

ANS.1 Data Encoding (TLV)

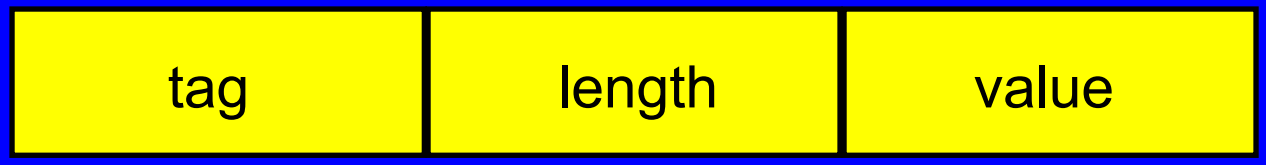

- tag
  - asn.1 data type
- length
  - length in octets
- value
  - value of data element
  - format dependent on type

ANS.1 Data encoding, length

- data element length field

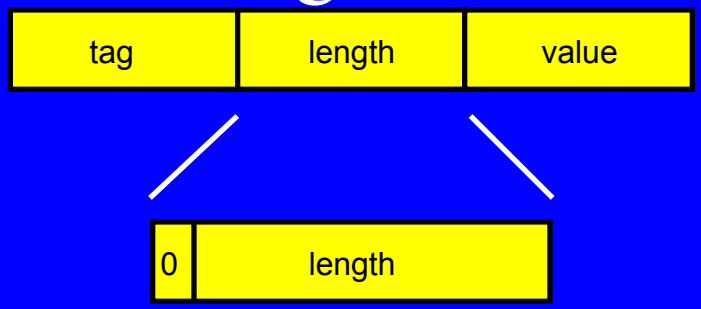

- if element length <= 126 octets

  actual tag value is in length octet with high bit = 0

  (value 127 is reserved)

- if element length > 127 octets

  length made up of chunks of 7 bits per octet

  high bit in all but the last octet = 1

  high bit in the last octet = 0

ASN.1, Object Identifier

- OBJECT IDENTIFIER

  sequence of integers that describe a pathway taken in traversing a tree of options, must be unique

  e.g.

  1.3.6.1.2.1.1.1

  or

  iso org dod internet mgmt mib system sysDescr

  the base of the tree is defined by ISO,

  sections are defined by other authorities

SNMP, MIBII

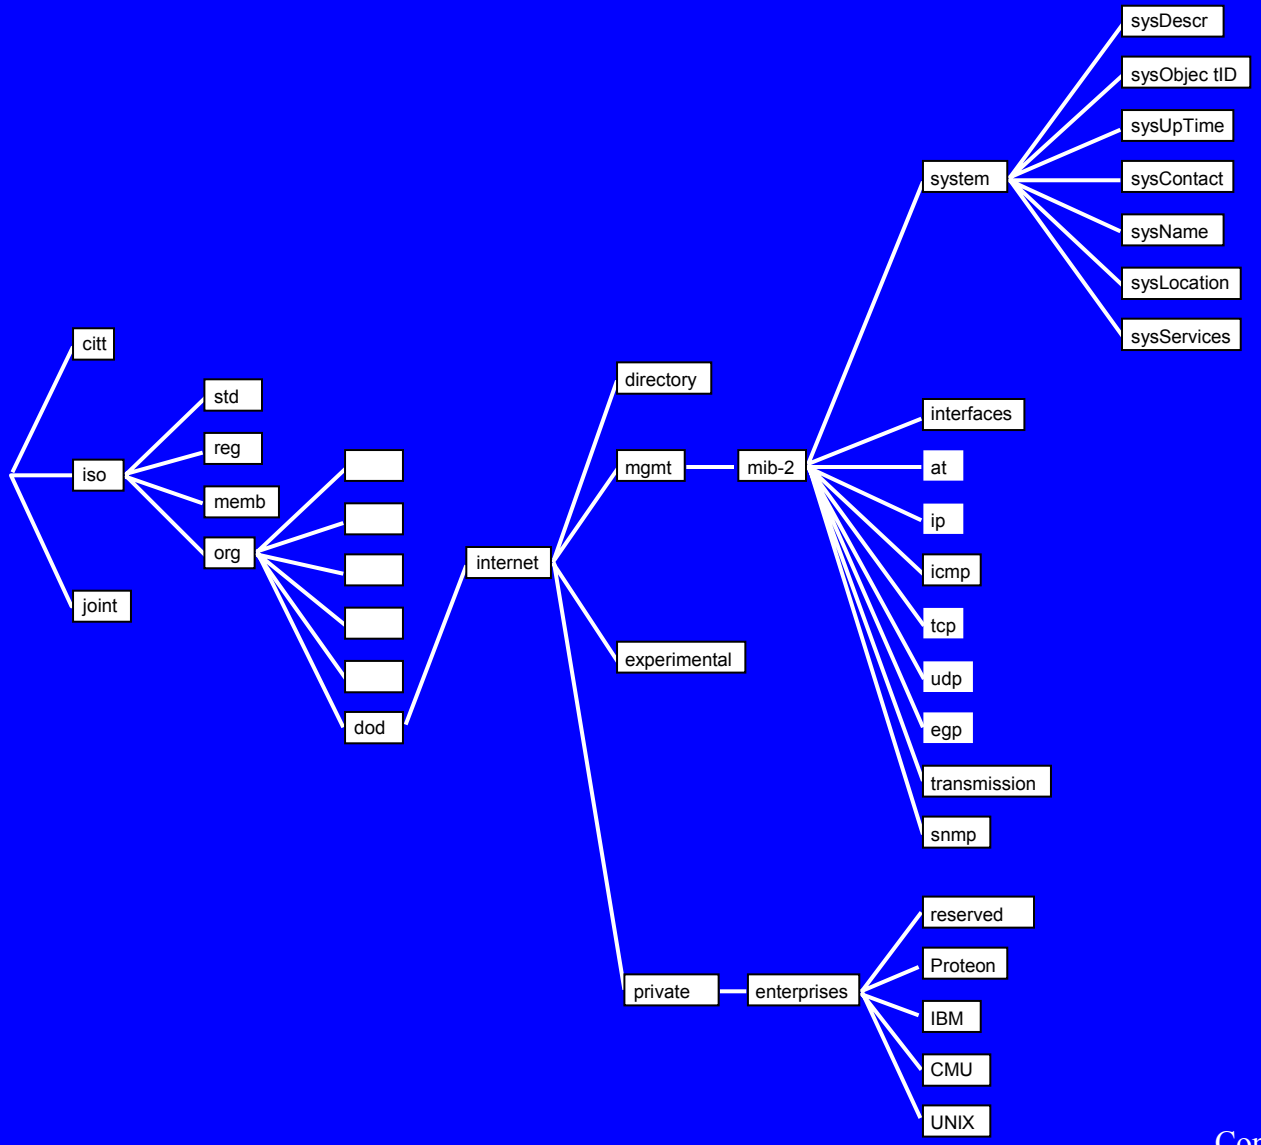

SNMP, cont.

- defines three query messages to get information from an agent being monitored
- defines a set message to be used in managing an agent
- defines a response message for an agent to use in responding to a query or set message
- defines a set of trap messages by which an agent can send notification of a status change to a management station
- defines an inform message for reliable communications

SNMP: Query Messages

- **`GetRequest`**

  request to an agent to return the current value of a specific MIB variable can take more than one variable in one request.

- **`GetNextRequest`**

  request to an agent to return the "next" MIB variable used to walk the tree in an agent

- **`GetBulkRequest`**

  request to an agent to return large blocks of data

SNMP: Set Message

- **`SetRequest`**

  request to an agent to change the values of one or more MIB variables to specific new values.

  if there is an error in the `SetRequest` and one or more variables cannot be set, none will be set

- error conditions

  1/ one or more objects not available for set operation, given access controls

  2/ contents of value field does not correspond to definition

  3/ size of response message would be larger than local limitations

  4/ some other reason a value cannot be altered

SNMP: Response Message

- **Response**

  message from an agent to a NMS in response to a `GetRequest`, a `GetNextRequest` or a `SetRequest`

  used to return requested values or to indicate success or failure of set request

  includes an error status and an error index

SNMP: Trap Message

- trap message:

  message from an agent to a NMS in response to a status change or event in the agent

- trap conditions:

  coldStart

  warmStart

  linkDown

  linkUp

  authenticationFailure

  egpNeighborLoss

  enterpriseSpecific

SNMP: InformRequest

- like a “reliable trap”
- designed to be used between network management stations
- expanding to other uses
- resent until acknowledged

SNMP: Communities

- provides trivial security
- like a password
- community name sent in clear over net with each message
- some agents have more than one community for different access modes

  these are know as "views"
- some agents can link access to community name and IP address of NMS

SNMPv3

- add security to SNMPv2
  - secure SET support
  - protect against
    - modification of information
    - masquerade
    - message stream modification
    - disclosure
  - does not deal with
    - denial of service
    - traffic analysis

SNMPv3, contd.

- three levels of security
  - no authentication, no privacy
  - authentication, no privacy
  - authentication & privacy
- can support more than one security model
  - user-based security model defined
    - security based on “name” of a user
- new message format
  - to add security information
- overview in RFC 2261

Accounting Management

- what do you account for?
- if you count packets sent
  - it can inhibit anonymous ftp & web sites
  - QoS differences in the future
- want to charge "user" of service
  - application dependent determination of "user"

Accounting, Cont.

- could do settlements based in routing information
  - try to minimize size of routing tables
- telco model
  - everyone shares in revenue
  - call an 800 number from a pay phone
    - 800 destinations pays pay phone owner
  - receive a long distance call to your own switch
    - you get fee for local delivery

getoctets

- simple traffic stats collector
- cron-driven shell procedure

```
get-octets router1 router2 router3 ...
```

- figures out interface list for each router
- then gets

  ifInOctets, ifOutOctets, ifInUcastPkts, ifOutUcastPkts
  ifInNUcastPkts, ifOutNUcastPkts, system.sysUpTime

- ftp://conrad.harvard.edu/pub/SNMPoll/octets.tar

  needs cmu snmp package

getoctets, contd.

- makes separate stats file for each interface
  - example filename: 128.103.1.2.WJHgw1
- example data

1997,06,23,160,09,1,00,02,37,EDT,1764089502,1045789221,99138769,92200835,10,628226,758006814
1997,06,23,160,09,1,00,22,37,EDT,1766362487,1047093977,99151676,92213338,10,628281,758126831
1997,06,23,160,09,1,00,42,36,EDT,1768439726,1048266407,99163118,92224546,10,628342,758246748

- processing a bit hard
  - must deal with counter wrap & router reboots
  - sample period must be < 59 min for an Ethernet
- link utilization calculation complex
  - must include link encapsulation etc

getoctets, processing

- UpDate routine

  bug in 32 bit versions of perl (gives bad results)

- example output

week	millions of bits per second				millions of octets		
ending	peak in	peak out	95% in	95% out	in	out	total
1997.06.01	5.0976	0.9330	1.3389	0.4104	18782	13752	32534

Policy Based Management

- want to manage the network not just network elements
- define policy rules to tell network what to do
  - e.g. network operations center gets access to all routers
  - e.g. accounting department gets priority last 3 days of month
  - e.g. max of 10% video on any link
- policy rule gets translated into changes in configurations of devices

Parts of a Policy-Based System

- conceptual parts
  - **policy management tool**
    - used to create policy rules
  - **policy repository**
    - store policy rules
  - **policy consumer**
    - pushes policy rules (or translations) to policy target
  - **policy target**
    - functional element effected by policy rule

Example Policy System

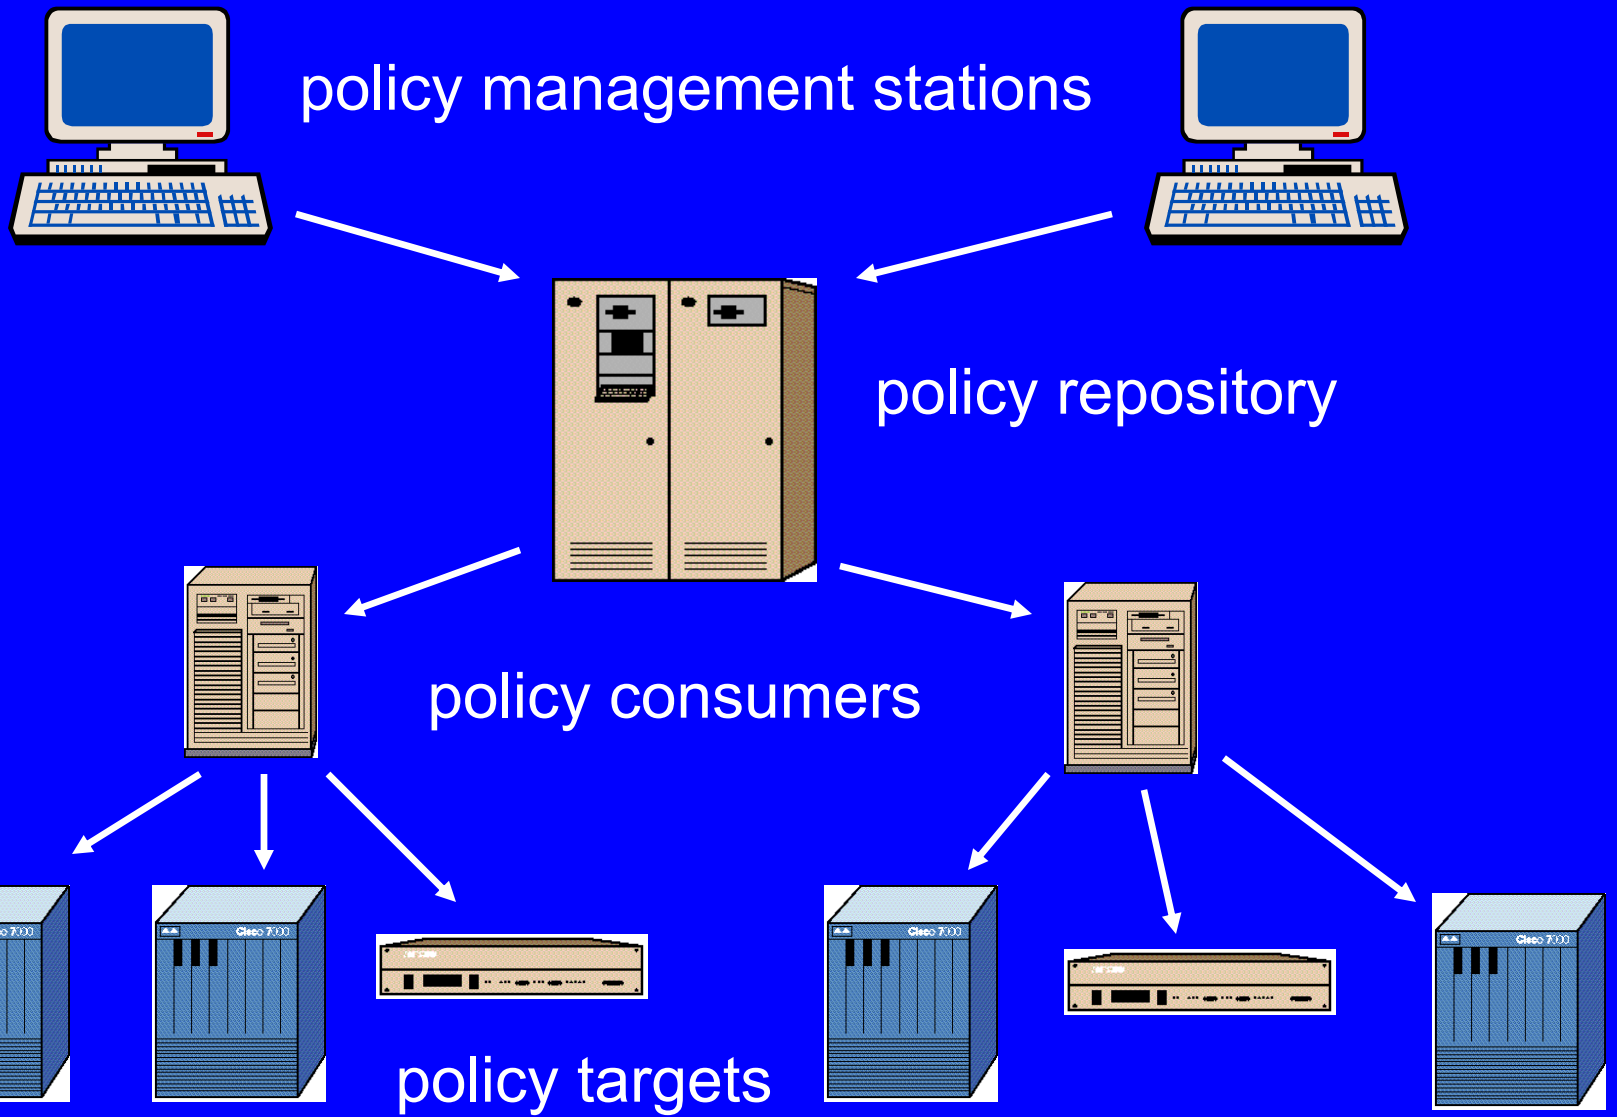

Policy Sequence

- rule defined
  - check for static conflicts
  - put in repository (could use LDAP)
- retrieved by policy consumers (could use LDAP)
  - processed to create configuration for policy target
  - can use conceptual rather than actual interfaces
    - e.g. “backbone interface”, “customer interface”
  - can have time component
  - pushed to policy targets
    - using COPS or SNMP

Policy Sequence, contd.

- policy target

  installs configuration

  may have to translate conceptual to physical interfaces

  e.g. “customer interface” -> interfaces 1, 3 & 8

  e.g. “backbone interface” -> interfaces 2 & 7

Management for Real

- A few basic tools
- echo request

  ping on IP

  function in many protocols - IP, OSI, AppleTalk, XNS

  checks path & basic node function

  can return round trip time

  normally not higher node function

```
newdev> ping noc.barrnet.net
PING noc.barrnet.net (131.119.245.5): 56 data bytes
64 bytes from 131.119.245.5: icmp_seq=0. time=83. ms
64 bytes from 131.119.245.5: icmp_seq=1. time=83. ms
64 bytes from 131.119.245.5: icmp_seq=2. time=83. ms
^C
----noc.barrnet.net PING Statistics----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 83/83/83
```

Management for Real, Cont.

- traceroute - finds path to node with delays
  - function only in some protocols
    - must have error returned on TTL exceeded

```
golem> traceroute gatekeeper.dec.com
traceroute to gatekeeper.dec.com (204.123.2.2), 30 hops max, 38 byte packets
1    38.1.1.1 (38.1.1.1)  0 ms  0 ms  17 ms
2    ipl-node17.camb.ma.cable.psi.net (38.127.88.1)  17 ms  0 ms  17 ms
3    38.146.139.1 (38.146.139.1)  0 ms  17 ms  17 ms
4    leaf.net121.psi.net (38.1.10.15)  50 ms  50 ms  33 ms
5    38.1.2.16 (38.1.2.16)  101 ms  117 ms  133 ms
6    San-Jose5.CA.ALTER.NET (137.39.29.1)  133 ms  283 ms  251 ms
7    Palo-Alto1.CA.ALTER.NET (137.39.29.3)  133 ms  133 ms  100 ms
8    Palo-Alto3.CA.ALTER.NET (137.39.47.7)  117 ms  100 ms  133 ms
9    border-gwl.pa-x.dec.com (204.123.0.241)  117 ms  117 ms  100 ms
10   gatekeeper.dec.com (204.123.2.2)  133 ms 117 ms  133 ms
```

file: pub/net/jacobson/traceroute.tar.Z

Management for Real, Cont.

network monitors/analyzers

- local systems
  - take unit to problem
  - don't depend on working network
  - wide range of cost & function
- remote systems
  - leave unit on problem or key network
  - remote control & viewing of information
  - SNMP standard from IETF RMON working group
- privacy & security issues

Management for Real, Cont.

- management agents
  - SNMP agents in all "gateway" devices
  - SNMP agents in all servers
- need something that knows what it is looking at it
  - not all SNMP variables are the same

Monitoring

- simple monitoring tools do 95% of task

  e.g. `ftp://conrad.harvard.edu/pub/SNMPoll`

- monitor should be both poll & trap based for best reliability

  but just polling will do better than just traps

  and will work fine other than response latency

- simple, terse, messages on problems

Example SNMPoll Error Messages

- interface

```
Date: Mon, 16 Jun 97 09:11 EDT
From: SNMPoll@Conrad.Harvard.EDU
To: HDN-mail@Conrad.Harvard.EDU
Subject: FMD_175_No_Harva 128.103.245.1 down

FMD_175_No_Harva 128.103.245.1 down : Hamgw1 128.103.15.21 Eth 5 down
   at 09:11:03
```

- router

```
Date: Fri, 13 Jun 97 17:17 EDT
From: SNMPoll@Conrad.Harvard.EDU
To: HDN-mail@Conrad.Harvard.EDU
Subject: MEEIgw1 not responding

MEEIgw1 not responding : No Response from 204.166.68.1 at 17:17:37
```

SNMPoll

- command to cause auto configuration
  - needs `config.seed`
- retries if failed poll of router
  - reduce false error messages
- understands hierarchy
  - tries next step “back up” hierarchy if failed poll
- query for many interfaces in same `get_request`
  - minimize network traffic and router load

Things to Look For

- dup protocol addresses
  - very bad in AppleTalk and IPX
- network/link load
- router/bridge
  - CPU load
  - errors
  - drops!!
  - interface resets
  - collisions (if CSMA/CD network)

Things to Do (Defensive)

- filter
- filter
- filter

Route Filtering

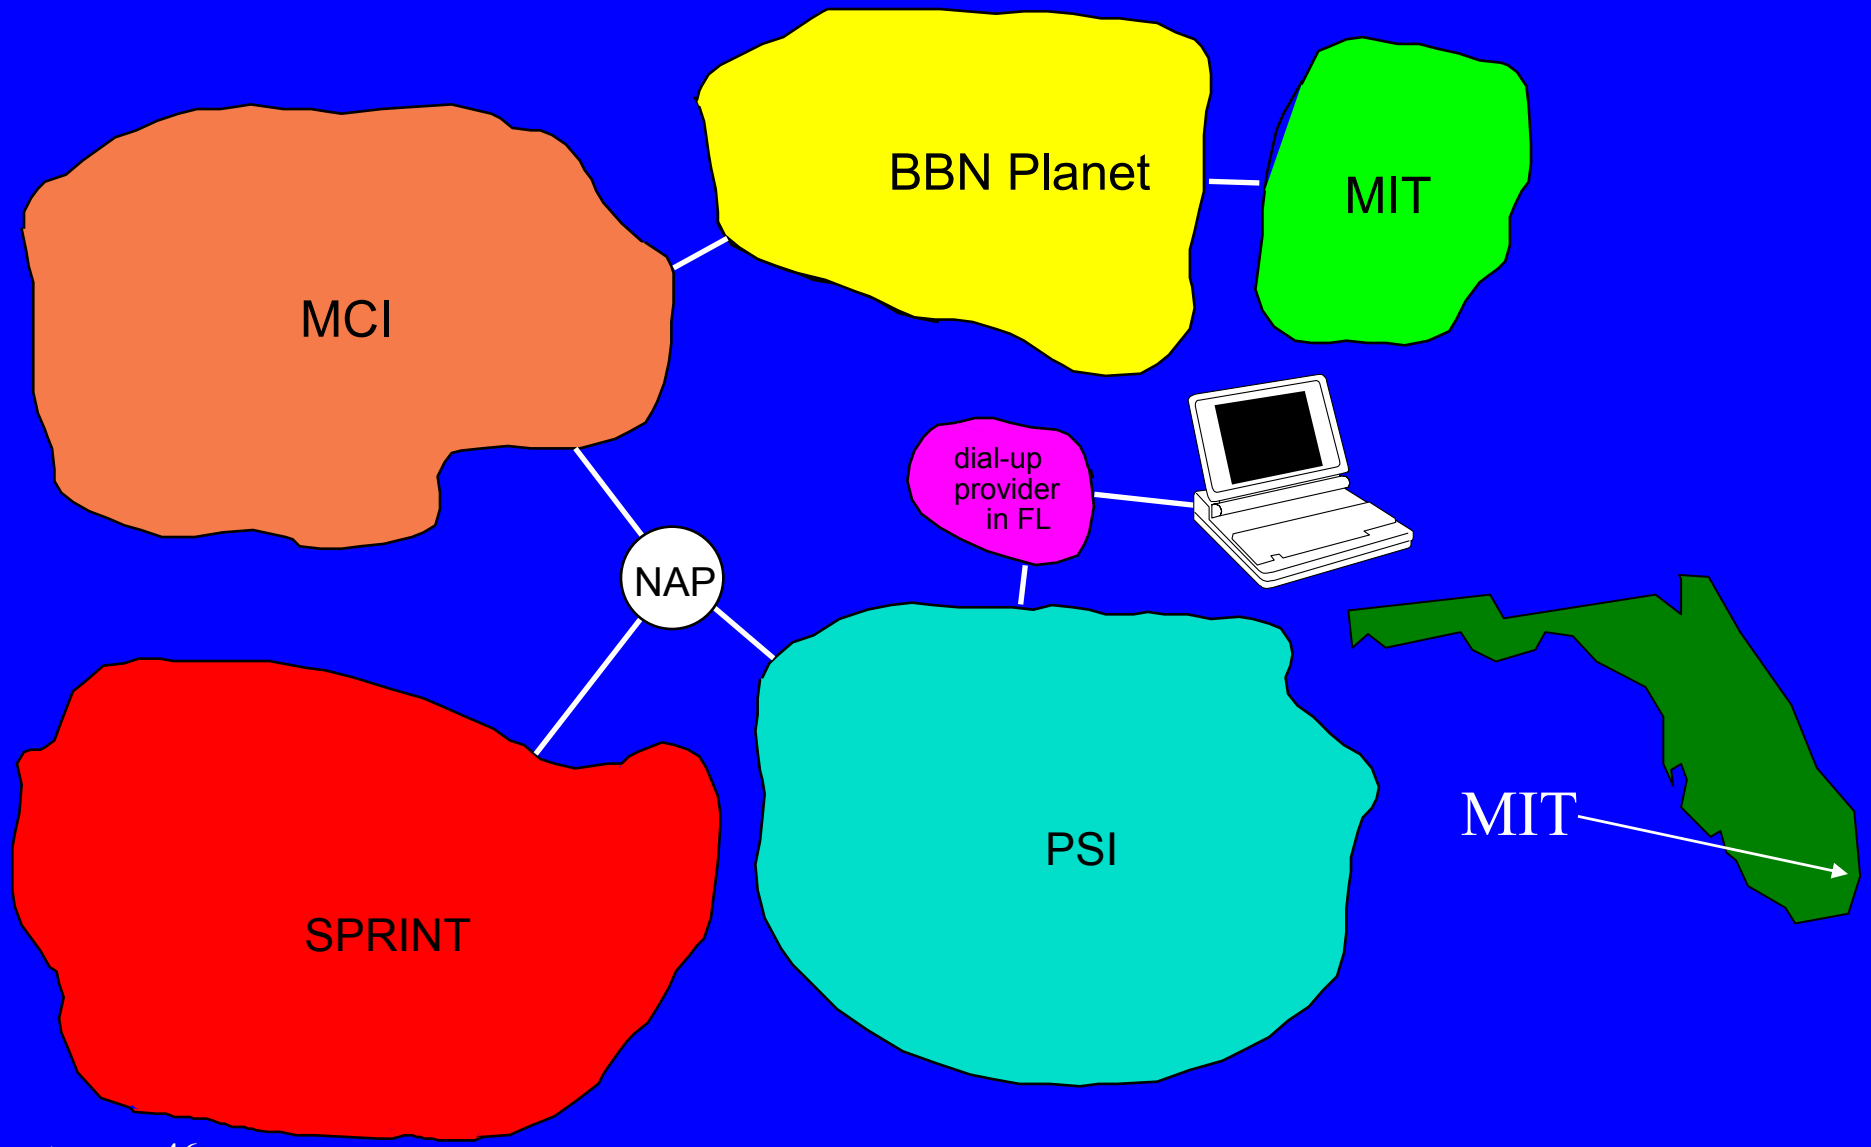

Problems

- we are early in the Internet management game
  - there is still a lot to learn
- little AI
  - not NETVIEW
- prices still high for functionality
- still gurus
- data networks are not "plug and play" with large scale
- nefarious people

More Problems

- not so good at providing simple, easy to understand, warning to non-gurus
- mostly managing elements

  policy-based management should help

  but still monitoring elements - no “big picture”
- needs to be usable by “normal” people
- needs to say when users will complain