

Host naming

Dave Eckhardt
de0u@andrew.cmu.edu

Synchronization

- P3 survey: shell?
- No class Monday
 - Faculty candidate talk
 - Click, From Extensible Router to Versatile Packet Processor
 - 10:00 a.m., 4623 Wean Hall
- No class Friday
 - Spring Carnival (“Mobot” races @ noon)

Synchronization

- Today: Naming (in the Internet)
 - DNS vs. ARP
- Upcoming topics survey
 - Distributed deadlock management?
 - Who is reading about Plan 9? EROS?
 - Kerberos? PGP?
 - Proof-carrying code?

Overview

- Three names for your PC
 - Why?
- Two resolution protocols
 - DNS, ARP
- Security considerations

Three names for ~~your~~ my PC

- PIPER.NECTAR.CS.CMU.EDU

- What's a “nectar”?

- What's a “piper”?

- 128.2.194.80

- 00-20-AF-D9-FD-CA

- *All* are globally unique

- Won't one do?

Questions about names

- Who uses the name?
 - For what?
- Who owns/defines the namespace?
- How long is the name valid?

PIPER.NECTAR.CS.CMU.EDU

- Who?
 - Human beings
- What?
 - Remembering a name for each box
 - Crude service-location mechanism
 - *www*.<organization>
 - Crude *device*-location registry
 - pa-mtlebanon3a-39.pit.adelphia.net
 - p4-2-1-0.r02.mclnva02.us.bb.verio.net

Fun break – hostname schemes

- Animals, birds, dinosaurs
- Cars, wines
- CMU SCS Facilities
 - Desktop machines: astronomical entities
 - Servers: fruits, nuts, vegetables
- Wean cluster: Medication
- MIT AI Lab: Breakfast cereals

PIPER.NECTAR.CS.CMU.EDU

- How long is it valid?
 - Lifetime of “the machine”
- How long is the *binding* valid?
 - See below

128.2.194.80

- Who/what *pairs*
 - IP router
 - End-system hosts

IP routers

- Which link does the packet leave on?
 - (definition of router)
 - Used as table lookup key
 - Short fixed length
 - String would *not* do
 - IP address structure
 - “Network ID”: top bits
 - “Host ID”: bottom bits
 - Network/host division depends on frame of reference

IP routers

- Link parameters (optional)
 - Again, table lookup key
 - Link address (station)
 - Link-level encryption state
 - Link-level scheduling policy

End-system hosts

- Connection management
 - TCP connection *defined by* (IP1, port1, IP2, port2)
 - “only” 65536 TCP connections per host pair
 - Client: my _____ server is x.y.z.w
 - Server: (IP, port) -> security state

128.2.194.80

- How long is it valid?
- Historically: “a long time”
 - 128.2 = CMU.EDU
 - 194 = some chunk of CS
 - 80 = random selection
 - No need to change for “lifetime of machine”

128.2.194.80

- Nothing fails like success
 - IP router table size explosion
- CIDR compresses via hierarchy
 - 12.0.0.0/8 (12.*) belongs to ATT.net
 - 216.218.128.0/17 belongs to he.net (Hurricane Electric)
 - 216.218.132.24/29 belongs to Panasas.com
- Change ISPs, your netblock changes
 - ... “ISP” can be Starbucks 802.11

00-20-AF-D9-FD-CA

- Who assigns?
 - IEEE <http://standards.ieee.org/regauth/oui/>
 - 00-20-AF assigned to 3Com
 - D9-FD-CA assigned by factory

IEEE MAC address

- Globally unique address
- For every “Ethernet” “card”
- “Ethernet”
 - Or 802.11, or ATM, or Frame Relay, or ...
- “card”
 - Semi-permanent expansion card
 - PCMCIA/CompactFlash card
 - Chip on motherboard

IEEE MAC address - Usage

- “Station” identification on “a network”
- Cooperating set of bridges agree on location
 - Which bridge owns which stations
 - Dynamic “spanning tree” algorithm
- Not “routable” outside that network
 - Then why is it *globally* unique?

Link addresses – why?

- Why do we need a MAC address?
 - Can't we just use the IP address?
- IP was designed to be *subnet-independent*
 - ARPAnet, SATnet, ARPA mobile radio network
 - DIX Ethernet, IBM Token Ring, Corvus Omninet, PPP
 - Each link has its own kind of address
 - Differ in size, meaning

Three names for my PC

- `piper.nectar.cs.cmu.edu`
 - For human use
 - Good for “a long time”
 - Maps to IP address
- `128.2.194.80`
 - For use by IP routers and IP protocols
 - Good while attached via a given ISP
 - Typically mapped to link-level address

Three names for my PC

- 00-20-AF-D9-FD-CA
 - Address used by Ethernet link hardware
 - Good for lifetime of interface card
 - Binding to machine is variable
 - Motherboard: pretty permanent
 - PCI card: almost as long
 - PCMCIA/CF card: at least 1 minute

Three names for my PC

- User specifies host name
- Data packet sent to IP address
- Last-hop router must know MAC address
- Two lookup problems
 - Name -> IP address: global, pretty stable
 - IP address -> MAC address: local, variable

Host name lookup

- In the beginning...
 - RFC 606: HOSTS.TXT!
 - One line per host
 - HOST : 128.2.194.80 :
PIPER.NECTAR.CS.CMU.EDU : INTEL-
GATEWAY : NetBSD ::
 - Available by FTP from SRI-NIC.ARPA
 - 10.0.0.51, the first time
 - Good for ~10 years, 1973 – 1983

Problems with HOSTS.TXT

- <http://public.planetmirror.com/pub/textfiles/internet/hosts.txt>
- Size
 - July 23, 1992
 - 22,000 hosts, 1 megabyte
 - Scale *that* up!
- Update frequency

Domain Name System

- RFC 882 (1983)
- Goals
 - Distributed database
 - Frequent updates
 - Cacheing
 - High availability
 - Map name to address even while host is down

DNS concepts

- Resource Record (RR)
 - Name, class (IN = Internet), type, value
 - PIPER.NECTAR.CS.CMU.EDU IN A 128.2.194.80
 - cs.cmu.edu IN NS BLUEBERRY.SRV.cs.cmu.edu
 - cs.cmu.edu IN NS MANGO.SRV.cs.cmu.edu
 - cs.cmu.edu IN NS PEACH.SRV.cs.cmu.edu
 - cs.cmu.edu IN NS BANANA.SRV.cs.cmu.edu

DNS concepts

- TTL = Time-To-Live
 - How many seconds a record will remain valid
 - Promise about stability of mapping
 - CS.CMU.EDU default: 2 days
- Query
 - Question, flags, query id #

DNS Concepts

- Response
 - Question, flags, query id # - from query
 - Result (Ok, “No such domain”, ...)
 - Answer records
 - Answer to your question
 - Answers to questions you *meant* to ask

DNS Concepts

- DNS server
 - Knows “all the answers” for a sub-tree
 - Except for sub-sub-trees it *delegates*
 - Like Unix file system mounts
 - EDU servers delegate CMU.EDU
 - CMU.EDU servers delegate CS.CMU.EDU
- Resolver (library)
 - Consults one or more DNS servers
 - Contains retry logic, “marshalling”

DNS Flow

- `gethostbyname("PIPER.NECTAR.CS.CMU.EDU")`
- Resolver contacts **D.ROOT-SERVERS.NET**
 - EDU IN NS L3.NSTLD.COM (and others)
 - *L3.NSTLD.COM IN A 192.41.162.32*
- Resolver contacts **L3.NSTLD.COM**
 - CMU.EDU IN NS T-NS1.NET.cmu.edu (...)
 - *T-NS1.NET.CMU.EDU IN A 128.2.4.14*

DNS Flow

- Resolver contacts T-NS1.NET.cmu.edu
 - CS.CMU.EDU IN NS PEACH.SRV.cs.cmu.edu
 - *PEACH.SRV.CS.CMU.EDU IN A 128.2.242.81*
- Resolver contacts PEACH.SRV.CS.CMU.EDU
 - PIPER.NECTAR.CS.CMU.EDU IN A 128.2.194.80

Advanced topics

- Flow for LAPIS.PRT.CS.CMU.EDU?
- How do we handle gethostbyaddr()?
 - Map *IP address onto name*

Advanced topics

- Flow for LAPIS.PRT.CS.CMU.EDU?
- How do we handle gethostbyaddr()?
 - Map *IP address onto name*
 - 80.194.2.128.IN-ADDR.ARPA IN PTR
PIPER.NECTAR.CS.CMU.EDU
- IP over DNS

ARP design

- Map IP address onto MAC address
- Within a single “network”
- MAC addresses have internal structure
 - But it's wrong: manufacturer, serial-number
- Two solutions
 - Ask a server
 - Why not?
 - ...?

ARP protocol

- Ask *everybody!*
 - That should include asking the right person
- Ethernet supports broadcast
 - Send packet to all stations on “network”
- WHO-HAS 128.2.194.80 TELL 128.2.254.36
 - Broadcast
- REPLY 128.2.194.80 IS-AT 00-20-AF-D9-FD-CA

Summary

- Three names for three purposes
- Two mapping protocols
 - Totally different according to function