

Networking 3

Misc TCP/IP Issues

IP Port Numbers

Each IP addr is logically divided into logical port numbers in the range 1-65535. Traffic is addressed to a specific port number at the IP -- this allows a computer with one IP to be in multiple conversations at once. As we'll see, specific port numbers are reserved for specific purposes -- e.g. port 80 is used for HTTP traffic.

Each end of a connection has a source and destination port number as well as IP addr.

Broadcast

It's possible to send a packet on the LAN specifically marked as "broadcast", so everyone reads it. You may be able to see a broadcast packet on your hub if all your "receive" lights blink at once.

TCP/IP also has a "broadcast" notion for sending information to an entire subnet.

IP to LAN Translation

How does the router know the right LAN addr to use to contact another machine on its LAN?

i.e. given an IP addr, what is the LAN addr of that machine?

ARP -- address resolution protocol

Broadcast "who has IP addr 171.64.64.250" on local LAN

The owner answers back with it's LAN addr (e.g. its ethernet MAC addr if ethernet is the LAN)

Reverse ARP -- aka RARP

At boot time, a machine broadcasts its ethernet addr. A RARP server notices the broadcast, and replies with an IP addr to use. DHCP has pretty much replaced RARP.

Ping

ping -- sends a little IP query packet to see if a host responds at all -- see if it is up and reachable.

```
nick% ping www.whitehouse.gov
PING a1289.g.akamai.net (171.66.255.138): 56 data bytes
64 bytes from 171.66.255.138: icmp_seq=0 ttl=251 time=3.706 ms
64 bytes from 171.66.255.138: icmp_seq=1 ttl=251 time=2.497 ms
64 bytes from 171.66.255.138: icmp_seq=2 ttl=251 time=1.932 ms
64 bytes from 171.66.255.138: icmp_seq=3 ttl=251 time=2.4 ms
64 bytes from 171.66.255.138: icmp_seq=4 ttl=251 time=2.32 ms
```

^C

```
--- a1289.g.akamai.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.932/2.571/3.706 ms
```

Time To Live -- TTL

Neat bit of design -- every IP packet has a max -hops counter, so it cannot get stuck in a loop -- a robust design, even in the face of router errors. When the packet hits its hop limit, a polite router will send back a "failure" notification, which will identify the IP addr of the router.

Traceroute

Send out packets with TTL values 1, 2, 3, 4, to probe the path to somewhere. With firewalls and whatnot attempting to "hide" internal networks, sometimes traceroute won't work and the path appears to go on endlessly.

The times listed are round-trip times in milliseconds

```
elaine40:/usr/class/cs193i/WWW> traceroute www.sjsu.edu
traceroute to rhea.sjsu.edu (130.65.3.13), 30 hops max, 40 byte packets
 1 leland-gateway (171.64.15.97)  1.277 ms  1.371 ms  0.973 ms
 2 Core2-gateway (171.64.1.233)  0.530 ms  0.624 ms  0.513 ms
 3 Core4-gateway-2 (171.64.3.18)  0.897 ms  0.785 ms  0.926 ms
 4 i2-gateway (171.64.1.225)  1.180 ms  0.770 ms  0.777 ms
 5 STAN.POS.calren2.NET (171.64.1.213)  1.125 ms  0.784 ms  1.084 ms
 6 SUNV--STAN.POS.calren2.net (198.32.249.73)  1.562 ms  1.119 ms  1.438 ms
 7 QSV-QSV-C2-GSR-ATM.CSU.net (137.145.203.209)  1.791 ms  2.051 ms  2.078 ms
 8 SJSU-QSV-ATM.CSU.net (137.145.203.106)  286.774 ms  149.992 ms  150.181 ms
 9 firewall.sjsu.edu (130.65.11.6)  138.667 ms  137.488 ms  147.755 ms
10 cclomnicore.sjsu.edu (130.65.11.2)  134.991 ms  134.259 ms  111.274 ms
11 cclomni-a.sjsu.edu (130.65.5.252)  313.109 ms  142.407 ms  146.947 ms
12 rhea.sjsu.edu (130.65.3.13)  137.231 ms  *  134.619 ms
```

Host Configuration

Things that need to be set for a computer to use TCP/IP...

1. IP addr -- need to know what IP addr to use
2. Subnet mask -- this number encodes what the local set up is for which part of the IP is subnet vs. which part is host addr.
3. Router -- the IP addr of the local router to forward traffic to
4. DNS server addr -- the IP addr of the DNS server(s) to use.

DHCP

Dynamic Host Configuration Protocol

Allows a host, at boot time or whatever, to broadcast a query to a local DHCP server. The DHCP server can reply with all of the above configuration values.

The DHCP may give out an arbitrary IP addr from its supply, or it may use a specific setup based on the LAN addr of the sender.

This is much more convenient way to do configuration for the end user -- just set the use-DHCP checkbox and you're done.

Layered ISO Model

Official layers

1. physical
2. data link -- one packet
3. network -- IP
4. transport -- TCP
5. session --
6. presentation
7. application

Layers

The point of layered design, is that they are as independent as possible -- you can change the implementation at layer 1 or 2, without needing to change the code that's working at layer 4.

Try to make the design at each level independent.

e.g. My cell phone knows how to talk to my computer. Sometimes I connect it with a USB cable, and sometimes I use bluetooth.

Style

The ISO layers were developed by a committee

In contrast, TCP/IP was built up by researchers and tuned and played with by experience. It does not exactly follow the abstract ISO model.

There's a natural style comparison between the "build a prototype and experiment" TCP/IP style vs. the top-down-plan ISO style,

UDP

TCP provides reliable, stream-oriented connections built on IP

UDP provides an unreliable, one-shot transport of a single datagram

Lower overhead, faster -- (no 3-way handshake)

Need to build in reliability yourself -- the datagram may or may not get through

e.g. Could use to send a stream of audio packets with minimum overhead

Special IP Addrs

The IP subnet 192.168.0.0 is set aside as a special "non-routable" network. An organization can use 192.168.0.0 addresses internally, and the routers are not supposed to let those addrs leak outside the local network. That way, another organization can also use 192.168.0.0 addresses, and the two organizations will not conflict. 10.0.0.0 is another non-routable network, but 192.168.0.0 is more commonly used and so is probably more reliable.

NAT

Network Address Translation NAT

RFC 2663

Allow computers on a network to share a single IP addr on the Internet.

Cable/DSL modem has one IP addr and an upstream router as usual

Translating router

Single transfer point between the internal network and the Internet at large -- it has a "LAN" side and an "Internet" side

The router has one IP addr on the Internet side

On the LAN side, the router is 192.168.1.1

The hosts on the LAN have addrs in the range 192.168.1.100, 192.168.1.101, 192.168.1.102,

The router will probably use a DHCP server to distribute configuration to the LAN computers

Translate

Suppose we have a computer on the LAN with IP addr 192.168.1.100

A computer on the LAN operates like any other TCP/IP computer. It forwards its packets to its router in the usual way.

Suppose the router's IP addr on the Internet side is 1.2.3.4

The router takes in traffic from it's LAN side, and then translates the packet on its way out -- the packet used to say "From: 192.168.1.100", gets translated to appear to be from the router itself "From:1.2.3.4". The other end of the connection thinks it is interacting with 1.2.3.4 in the usual way.

However, the router knows to re-send incoming packets in that conversation on its internal network to 192.168.1.100. The router keeps track of which port numbers are currently part of which conversation.

Results

Multiple computers on the LAN side can have conversations at the same time. The router needs to keep these straight.

To the outside world, it appears that all of the traffic is coming from 1.2.3.4 -- just a lot of traffic, since it's actually the sum of all the traffic from the LAN computers.

Internally, the computers think they have IP addrs in the 192.168.1.100 -- which is not the same addr that the computers in the outside world think they are using.

Problem: if a computer on the inside sends what it thinks is its IP addr to a computer outside the LAN, it's not going to work. Mostly, that's not a problem -- regular "call/response" connections work fine.

Incoming calls don't work well -- what LAN computer should it go to? Can set up certain ports to go to certain LAN side computers.

TCP/IP -- Standard, Collective, Cooperative

TCP/IP is just a standard -- an agreed format and protocol for things

It's a free, public, open standard

Vendors voluntarily implement TCP/IP

TCP/IP allows each vendor's equipment to, in one step, interoperate with all the other computers which is irresistible

Compare this to the "balkanized" picture where each vendor only works with their own equipment.

TCP/IP was developed with a small budget in a non-profit way.