TCP - Services

Ping /Traceroute

In theory, the network is now working in Terman, so we can look at the ping and traceroute utilities.

ping -- sends "ping" datagrams to a host, see if it responds

traceoute -- figure out the sequence of routers on the way to a host

security -- both ping and traceroute depend on the hosts or routers involved to reveal some information about themselves. In security conscious networks, the hosts and routers may not give up such information.

Larger NAT Example

Setup: Client in on the LAN at addr 192.168.1.100. NAT router is at 192.168.1.1 on its LAN side, and 1.2.3.4 on its Internet side

Suppose the client wants to make a connection to 2.2.2.2, port 80.

1. The client chooses, at random, port 4000 for its end, and tries to make the connection.
2. The datagram addressing is: From: 192.168.1.100:4000, to: 2.2.2.2:80
3. The datagram is forwarded, as usual, to the local router, in this case the NAT router
4. The router chooses, at random, port 5000 for its outgoing port. The NAT router re-writes the addressing so the traffic appears to be from the router itself -- From: 1.2.3.4:5000, To: 2.2.2.2:80 and sends it on its way. The router remembers that the outgoing port 5000 is being used for this particular connection.
5. Eventually the traffic comes back to the router with addressing From: 2.2.2.2:80, To: 1.2.3.4:5000.
6. The router remembers that port 5000 is for the 192.168.1.100:4000 connection, and so re-writes the datagram back to its "local" form -- From 2.2.2.2:80, To: 192.168.1.100:4000

Summary

The NAT router distinguishes the addressing schemes on its LAN vs. Internet sides

It re-write the addresses as they cross the boundary

It uses the port numbers to keep track of which connection is which.

The address the client thinks it has, 192.168.1.100, is not the address the client appears to have to the rest of the Internet. All of the clients on the LAN appear to be at the single IP address of the NAT router -- 1.2.3.4 in this example.

For the most part the client and server do not need to know that the NAT router is doing something to the datagrams -- they can just send and receive as usual, and the NAT router hides the complexity.

TCP / IP -- Phone Sys Infrastructure

The TCP / IP standard allows computers to identify and exchange bytes with each other.

IP addresses

IP datagrams

TCP virtual circuits / sockets

What are you going to build on top of this basic infrastructure?

Service

A service is a protocol -- two or more computers use the protocol to interact with each other to provide the service. Typically, a service is associated with a port number -- e.g. the HTTP service uses port 80.

RFC

Request for Comments -- the term for the document typically used to define an Internet service standard.

Each RFC gets a number, and you can look at them all at www.ietf.org

Behind every useful two-way dialog between computers, there's a humble old RFC document which allowed it to happen.

Internet Standards

The secret of the Internet

Q: How did the Internet come about?

Q: Standards

1. Standards

Any time two computers want to co-operate to get something done, there needs to be a standard (a "protocol") that they both follow so the interaction works. The Internet is entirely about pairs of computers communicating.

2. Publicly available

The standard needs to be well-defined and publicly available.

There may be test suites that verify that an implementation is true to the standard.

3. Unencumbered -- freely implementable

The technology in the standard, typically, should not be patented, or should have generous licensing terms. Anyone should be free to implement the standard.

The standard may come from a particular vendor, but it only works as a standard when anyone is free to implement it.

4. Well Managed

The standard needs to be managed over time: allowing for improvements without breaking compatibility. Typically, this requires an organization that is not too tied to any one vendor. The interests of a vendor, and the great interests of the standard participants not the same.

Incredible Growth -- further questions

We'll study this more later, but at this point we can notice that well defined standards can bring tremendous growth.

TCP/IP, HTML, HTTP -- these are all standards

Co-operateive Standards vs. Individual Vendors

Microsoft, IBM, Oracle -- these are all dominant vendors in their domains. And yet the Internet did not arise using anything controlled only by them. The Internet arose with the collective participation of everyone through standards. The Internet is about co-operating to create value. Even in the face of a vendor with 90% marketshare, the power of co-operation is greater -- it's an incredible story and an important lesson.

Example of Standardized Services

All the interesting things you can accomplish on the Internet are actually services built on the TCP/IP structure...

DNS

Domain Name Service -- the service computers on the Internet use to look up names like "www.yahoo.com" to find their IP addresses. Implemented as a distributed database.

FTP

File transfer protocol -- move a named file form one computer to another (does not encrypt the password)

SCP

Secure Copy -- like the unix CP command, but over the Internet. Encrypted, so this is safe (unlike FTP).

Ping

Contact a host just to see if it's up and reachable

Finger

Ask a host if it knows about a particular username

Telnet

Establish a command line interface to a login on the server. (SSH is a newer service that encrypts the password and session data.)

SMTP

Simple Internet Mail Transfer Protocol -- Send/forward Internet email

POP

Post Office Protocol -- The owner of an email box retrieves their arrived mail.

IMAP

Internet Mail Access Protocol -- An alternative to POP where the email remains up on the server -- it is not copied down.

HTTP

Hypertext Transfer Protocol -- The "web server" service -- how a client uses a URL like http://www.stanford.edu/class/cs193i/" to contact a server, make a request, get back a web page.