

ICS 351: Today's plan

- Ethernet switching
- Spanning Tree Protocol

Learning Switches

- if a switch gets a packet from A on interface I , it forwards the packet,
- and remembers that A can be reached on interface I
- the next time a packet for A is received on interface I' , it is only forwarded on interface I (unless $I == I'$, and then it is not forwarded)
- if there is no record of communication from A (within the last 60 seconds), the packet is broadcast on all interfaces except I'

Broadcast Storms

- given a network with redundant links
- if the network is connected by hubs, every packet will cause collisions with itself
- if the network is connected by switches, any broadcast packet will live forever
- packets may even be multiplied if there is more than one loop
- this is useless traffic that gets in the way of useful traffic -- a "**broadcast storm**"

Preventing Broadcast Storms

1. have no redundant links in the network, or
2. restrict "broadcast" forwarding by switches:
 - * select a root switch, based on priority, using MAC addresses to break ties in case of equal priority
 - * find a least-cost path to the root, reached via the root port for this switch
 - * for each segment, determine a least-cost switch port to use to reach the root, the designated port for this segment
 - * only forward broadcasts along root ports and designated ports
 - * root ports and designated ports form a *Spanning Tree*

Rapid Spanning Tree Protocol

RSTP

- the regular spanning tree protocol can take tens of seconds to converge after a topology change
- instead, a switch can pre-select alternate ports that also lead to the root bridge
- broadcast data is only sent on alternate ports when it is determined that the root port is disconnected
- similarly for backup paths to individual segments
- switches also actively exchange their information, so one switch can quickly hand off forwarding to another switch

IP address exhaustion

- IPv4 addresses are 32 bits long
- so there should almost be one address per person on the planet (and many people on the planet do not own a computer)
- however, IP addresses must be assigned within networks
- renumbering networks is time-consuming and expensive
- so, there is an effective shortage of IPv4 addresses, especially in some areas
- three techniques are used to deal with this shortage:
 - o use a single address to connect to the Internet multiple computers: Network Address Translation, or NAT
 - o assign an IP address to a computer only when it is needed: Dynamic Host Configuration Protocol, or DHCP
 - o switch to IPv6, with 2^{128} addresses
- the first two techniques are common today, the last may become common in the future

Network Address Translation

- both TCP and UDP use ports to identify different applications on one computer
- most computers use far fewer than the 65,536 ports they could use
- so, multiple computers could share a single IP address, and just agree to use different ports
- since all computers use the same IP address, the combination of <internal port, external port, external IP> is what uniquely identifies the socket and source computer

Network address translation details

- in a NATed network, the computers run as if they were on the Internet
- each interface is assigned an IP address from a reserved space, 10.0.0.0/8, 172.16.0.0/12, or 192.168.0.0/16 -- these are non-routable IP addresses
- the router for this private network also runs DHCP to assign these addresses, is configured as the private computers' default gateway, and must acquire or be configured with the external IP addresses
- this "router" (really a NAT unit which usually does not participate in routing protocols) must examine each packet being forwarded and change some of the headers:
 - for every outgoing packet, the private source address must be replaced with the public source address. The combination of <internal port, external port, internal IP, external IP> is recorded in a table
 - for every incoming packet, the external (destination) IP must be replaced with the correct internal IP address taken from the table
 - if two or more machines on the internal network are using the same port number, the NAT unit generally also changes the port number
- the NAT table can also have static entries, for example pointing to a web server on the inside network
- the only limitation is that there cannot be two servers on the same port (e.g. port 80) on the inside network
- so to the outside, the entire private network behaves as a single computer, with any number of clients and servers

Firewalls

- the extended "router" can also discard any incoming packets for ports not listed in its table, or for statically configured ports
- this prevents outside access to applications (e.g. printers) that would be available inside the networks
- this also prevents outside access to applications that were installed automatically, and that the user is unaware of
- both of these are security risks, so the job of blocking "ports" is called firewalling, and such units are known as firewalls
- although the firewall function can be present in a "router" or a NAT unit, it can also be present in software on each computer
- the Linux IPTables mechanism is general enough to allow general forwarding as well as blocking of packets

Dynamic Host Configuration Protocol

- in our lab, computers have statically-configured IP addresses
- we also manually reconfigure IP addresses when needed
- when there are lots of computers or non-technical users, this becomes very cumbersome
- instead, a network administrator could decide the assignment of IP addresses centrally, and let a central computer distribute IP addresses
- however, the computer is not really on the Internet until it has an IP address
- so DHCP (like ARP) cannot use IP packets
- each address is leased for a given period of time, then must be renewed
- in most DHCP servers, renewal is automatic unless the network administrator decides otherwise
- so the lease expiration simply makes it easier to recycle addresses
- the system administrator decides the lengths of leases, which IP address ranges are available for DHCP, etc

Putting it all together

- many small office / home office (SOHO) networks are connected to the Internet by a single "router"
- this "router" runs no routing protocols, and may use DHCP to obtain its address from the Internet Service Provider
- or, if there are significant servers on the SOHO network, the IP address may be static and manually configured
- this "router" acts as a default router for the SOHO network, forwarding to its own default router (configured, or obtained by DHCP) all packets from the inside to the outside
- this "router" always performs NAT, to allow the sharing of the IP address
- this "router" usually performs some sort of firewalling
- for example, the firewall might by default allow all outgoing connections/streams, initiated by a computer inside, but not connections or streams initiated by outside computers
- but data for valid connections (streams) must still flow in both directions