ICS 451: Today's plan

- Firewall
- NAT
- Home “routers”
- IP routing: RIP, OSPF

Firewall: motivation I

- A good way to attack a computer is to connect to a server running on the computer
  - e.g. a mail server intended to handle local mail
  - or a name server used by local applications
- well-designed servers should:
  - default to being secure
    - and be easy to configure to be secure
  - not be vulnerable to attack
- many servers are not well-designed
- better: block incoming connections altogether!

Firewall: motivation II

- a network administrator controls the router
- but may not control individual computers
  - especially with BYOD (bring your own device)
- users may have little security knowledge
- so: block all incoming connections
- unblock selectively when someone has a need

- at home, unblock only if user knows enough to configure the home “router”

Firewall

- a firewall is a special kind of router
  - forwards packets, but
  - usually does not implement routing protocols
  - only serves as default gateway
- by default, a firewall blocks incoming connections
  - must work at the TCP/UDP/ICMP level
  - default is secure
  - configuration permits access to servers
  - client-only machines are more secure

Firewall challenges

- one more device to configure and maintain
  - configuration may be incorrect
    - especially if the network changes
- does not protect against attacks that attract vulnerable clients to external servers
  - phishing attacks
  - Java/Javascript attacks on browsers
- does not protect against attackers who are behind the firewall
  - e.g. devices compromised by viruses

Intrusion Detection System, IDS

- a firewall that also pattern-matches traffic
- trying to detect attacks
- response may include:
  - shutting down offending traffic
  - shutting down compromised device
  - alerting network administrator

IDS challenges

- same as virus detection: patterns may be obsolete (no protection against zero-day attack)
- needs very high speed pattern matching
  - may need to reassemble TCP streams
- too many false positives
- cannot inspect encrypted payloads

Network Address Translation NAT

- due to scarcity of IPv4 addresses, better to use a single address for many clients
  - hosts “inside” use local IP addresses
  - packets going “out” are re-written to have as source address the IP address of the NAT
    - may need to change source port number too
  - packets coming “in” are re-written to have as destination address the internal IP address of the local device
    - as identified by destination port number
    - must keep a table of such translations

NAT table

- protocol, inside local IP and port, outside local IP and port, remote IP and port
  - ICMP doesn't have ports, so use ID instead
- incoming packet matching protocol, remote IP and port, outside local IP and port, given inside local IP and port and forwarded
  - if no match, packet is dropped (firewall)
- outgoing packet matching all 5 is rewritten and forwarded
  - add to table if no match (maybe only SYN)
- table entries time out after a while

NAT details

- a NAT may manage multiple public IP addresses
- NAT timeouts can vary widely, esp. for UDP
- local port need not be rewritten if unique
- TCP/UDP header checksum must be recomputed when changing IP addresses

- NAT only works with protocols it understands!

Home “router”

- Firewall + NAT
- distinct “inside” and “outside”
- DHCP to get its outside address, default gateway, DNS server(s)
- provides DHCP (and may provide caching DNS) to its inside hosts
- acts as a router in forwarding packets (HW 10)
- does not run any routing protocols, so in that sense not a router

IP routing

- how to build the routing tables
- basically distance vector or link state
  - path vector includes details about the path
- three main protocols:
  - RIP, which uses distance vector
    - used within smaller networks
  - OSPF, which uses link state
    - used within larger networks
  - BGP, which uses path vector
    - used in the Internet backbone

IP routing overview

- administrators decide which routing protocol to use within a domain
  - this is an IGP, Interior Gateway Protocol
  - RIP and OSPF are IGPs
- formally a domain is known as an Autonomous System, AS
- BGP is used between Autonomous Systems
  - and is an EGP, Exterior Gateway Protocol
  - BGP is currently on Version 4

IP routing goals

- for an IGP, the goal is to find the best route to a destination
  - assumes all routers in the AS are trusted and somehow equal
    - links need not be equal
- for an EGP, the goal is to find a route that satisfies internal policy
  - policy may aim for the cheapest route
    - use AS x only if AS y is not available
  - policy may avoid routing through a specific AS
    - for example, a competitor

RIP

- Routing Information Protocol (v1, v2)
  - RIPng for IPv6
- distance vector
  - with split horizon with poisoned reverse
  - infinity is 16
  - messages sent every 30s, time out after 3min
- messages contain the distance to each destination (from the routing table)
  - lower-distance routes added to the routing table

RIP properties

- UDP port 521
- v2 can handle masks
- fast convergence when new links added
- slow convergence when existing links removed
- simple to configure
- supports networks with diameter up to 15

OSPF

- Open Shortest Path First (v2)
  - OSPF v3 for IPv6
- link state (HELLO packets every 10s)
- can split network (AS) into areas:
  - the backbone (area 0) connects to all other areas
  - complete link state distributed within area
  - summary info distributed between areas
  - area border routers span multiple areas
    - and are part of the backbone area

OSPF properties

- runs directly over IP (protocol number 89)
- fast convergence when new links added
- fast convergence when existing links removed
  - link state floods information quickly
- harder to configure if multiple areas
  - different types of routers,
    - e.g. area border routers
- supports large networks (100s of routers)

- IS-IS is similar to OSPF