CSIS 4222

Ch 27: Internet Routing

Ch 30: Packet filtering & firewalls

Internet Firewall

A combination of hardware and software that isolates an organization's internal network from the Internet at large

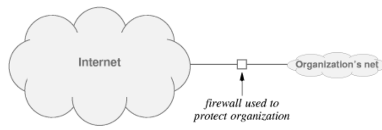

By placing a firewall on each external network connection, an organization can define a secure perimeter.

Internet Firewall

- Used by network administrator to manage traffic flow in and out of the internal network
- Implements a security policy and rejects any traffic that doesn't adhere to it
- Primary means of accomplishing this is through *packet filtering*

Packet Filtering

Filtering decisions typically based on fields in a packet's header:

- IP source or destination address
- TCP or UDP source and destination port
- ICMP message type
- Connection initialization datagrams using the TCP SYN or ACK bits

Examples

- To block all telnet connections
  - Block all TCP segments whose source or destination port number is 23
- To block streaming video
  - Block all UDP segments
- To prevent external clients from connecting to internal servers
  - Block incoming TCP segments with ACK=0 (all other segments have ACK=1)

Packet filtering can be implemented in a router

Specify which packets can pass through and which should be blocked

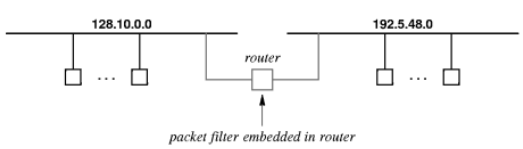

Linux Packet Filtering

The Linux `iptables` program acts as a packet filter

- Used to design a firewall to protect a single computer
- It can filter traffic based on port numbers, addresses, and flags
- It organizes rules into groups called chains
  - *Input*, *output*, and *forward* are built-in chains
- Rules are applied in order, first match is the one used
- A policy specifies how to handle packets that do not match any rules.

Adding Filtering Rules

- Accept incoming TCP packets on interface eth0 from any IP address destined for 92.168.1.1

```
iptables -A INPUT -i eth0 -d 192.168.1.1 -p TCP -j ACCEPT
```

- Reject ping packets from 192.168.1.5

```
iptables -A INPUT -s 192.168.1.5 -p icmp -j REJECT
```

Stateful Firewalls

- A *stateful* firewall allows traffic from inside the network to exit but doesn't allow general traffic from outside to enter
  - Outside packets can enter only if they match a request from within the network
- Keeps track of packet flow
  - Maintains information about recent history of *traffic on a connection*

Stateful Firewalls

Example: Host requests a page from a web server outside the network

- Firewall recognizes SYN packet from host and creates a state w/source and destination IP addresses
- Web server returns a SYN-ACK which the firewall allows to pass through
- State is maintained until connection ends

Stateful Firewalls

Are outside users ever allowed access?

- Firewalls generally must open ports for incoming traffic to web servers, DNS, email
- Create demilitarized zones (DMZ) to isolate these servers from the rest of the network

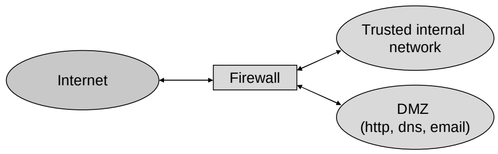

Intrusion Detection Systems (IDS)

- Monitors all arriving packets and notifies the site administrator if a security violation is detected
- Provides an extra layer of security awareness even if a firewall prevents an attack
- Can be configured to watch for specific types of attacks
  - For example, port scanning

Content Scanning and Deep Packet Inspection

- A firewall only examines fields in a packet header
  - Cannot test the payload of a packet for viruses, etc.
  - This requires content analysis:
    - File scanning
    - Deep Packet Inspection (DPI)

Content Scanning

- Take a file as input and looks for byte patterns that indicate a problem
  - Many virus scanners look for strings of bytes known as a fingerprint
  - Virus scanner software searches files for such sequences
- File scanning can make mistakes
  - *false positive*
  - *false negative*

Deep Packet Inspection

- Operates on packets
  - Examines the data in the packet payload
  - Includes the header fields
  - In many cases, the payload cannot be interpreted without examining fields in the packet header
- Disadvantage of DPI is computational overhead

Routing Terminology

Forwarding
- Refers to datagram transfer
- Performed by host or router
- Uses routing table

Routing
- Refers to propagation of routing information
- Performed by routers
- Inserts / changes values in routing table

Routing Issues

A routing algorithm must provide:
- **Correctness and simplicity**: Networks are never taken down; individual parts (e.g., links, routers) may fail, but not the whole network
- **Stability**: Handle topology and traffic changes without aborting jobs, rebooting, etc.
- **Fairness and optimality**: Often in conflict. Fairness is not part of definition of optimality.

Two Forms of Internet Routing

Static routing
- Forwarding table initialized when system boots
- No further changes

Dynamic routing
- Table is initialized when system boots
- Routing software learns routes and updates table
- Continuous changes are possible

Static Routing

Used by most Internet hosts

Typical routing table has two entries for:

- Local network → direct delivery
- *Default* → nearest router

```
olanm@zeus:~$ /sbin/route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
134.210.177.0   0.0.0.0         255.255.255.0   U     0      0        0 eth0
127.0.0.0       0.0.0.0         255.0.0.0       U     0      0        0 lo
0.0.0.0         134.210.177.1   0.0.0.0         UG    1      0        0 eth0
```

Direct delivery

Default

Dynamic Routing

- Used by IP routers
- Requires special software
- Each router communicates with its neighbors by passing routing information
- Uses a route propagation protocol

Dynamic Routing and Routers

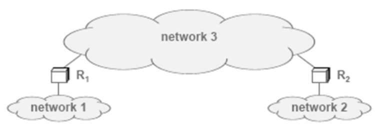

- Router R_1 knows about networks 1 and 3
- Router R_2 knows about networks 2 and 3

Each router exchanges information with other routers

Dynamic Routing and Routers

- Routing software updates the local forwarding table when it learns about changes in routes
- Routers exchange information periodically
- In the example:
  - R_2 will install a route to network 1 and R_1 will install a route to network 2
  - If R_2 crashes, the route propagation software in R_1 will detect that network 2 is no longer reachable and will remove the route from its forwarding table
  - Later, when R_2 comes back on line, the routing software in R_1 will determine that network 2 is reachable again and will reinstall the route

Routing in the Global Internet

- A route propagation protocol allows one router to exchange routing information with another
- But this cannot scale to the entire Internet
  - Routers and networks in the Internet are divided into groups
  - All routers within a group exchange routing information
  - Then, at least one router (possibly more) in each group summarizes the information and passes it to other groups

Routing in the Global Internet

How large is a group?

- To accommodate organizations of various size, no exact group size is dictated

How is routing information represented?

What protocol do routers use within a group?

- Each organization can choose a routing protocol independently

What protocol do routers use between groups?

- Interconnected groups must agree

Autonomous Systems

An autonomous system is a region of the Internet (networks and routers) that is administered by a single authority

Examples:

- UUNet (Verizon) backbone network
- Regional Internet Service Provider
- A big university

Each AS chooses a routing protocol

Internet Routing Protocol Classes

Interior Gateway Protocols (IGPs)

- Used by routers within an autonomous system
- Destinations lie within same AS

- Example protocols
  - RIP (simple, old)
  - OSPF (better)

Exterior Gateway Protocols (EGPs)

- Used between autonomous systems
- Destinations lie throughout Internet

- Example protocols
  - EGP
  - BGP (more recent)

Illustration of IGP / EGP Use

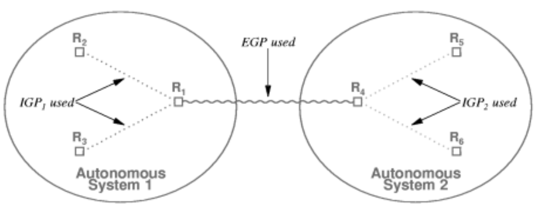

Optimal Routes and Routing Metrics

- Routing software should find all possible paths and then choose one that is optimal
- How to measure the optimal path between any source and destination?

  For a remote desktop application
  - a path with least delay

  For a browser downloading a large graphics file
  - a path with maximum throughput

  For a real-time audio webcast application
  - a path with least jitter

Optimal Routes and Routing Metrics

- Typical Internet routing uses a combination of two metrics:

  Administrative cost and hop count
- Hop count gives the number of intermediate networks on the path to the destination
- Administrative costs are assigned manually
  - Often to control which paths traffic can use
  - Routing software chooses the least cost path

Routes and Data Flow

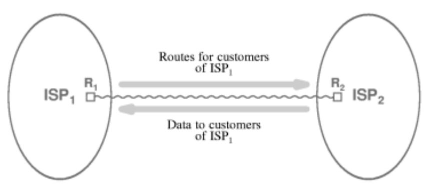

- Each Tier 1 ISP is an autonomous system that advertise its customers' networks to other ISPs.
- After an ISP advertises destination D, datagrams destined for D can begin to arrive

Distance Vector Routing

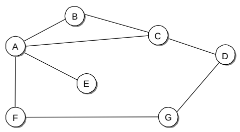

All nodes start by building a local view of what nodes are 1 hop away.

Every node sends its vector to its directly connected neighbors.

F tells A that it can reach G at cost 1. A knows it can reach F at cost 1, so it updated its own vector to indicate that it can reach G at cost 2.

Higher cost routes to G will be ignored, finding a lower cost route will replace the route currently in the vector.

After a few iterations of these exchanges, the routing table ***converges*** to a consistent state.

Periodic updates: Every *t* seconds, send local info to your neighbors. This allows other nodes to know that you are running.

Triggered updates: Every time you learn new info from a neighbor that makes you to update your local vector, send the recomputed vector to all your neighbors.

Internet Routing Protocols (Interdomain)

Border Gateway Protocol (BGP-4)

- Currently the EGP of choice for the Internet
- Provides routing between autonomous systems
- Gives path of autonomous systems for each destination
- Uses reliable transport (TCP)
- Distance vector algorithm

BGP Tracing: http://www.routeviews.org/

Internet Routing Protocols (Intradomain)

Routing Information Protocol (RIP)

- Routing within an autonomous system (IGP)
- Hop count metric
- Distance vector algorithm
- Unreliable transport (uses UDP)
- Implemented by the Unix program `routed`

Link State Routing

- Each node knows the distance to its neighbors
- The distance information (link state) is broadcast to all nodes in the network
- Each node calculates its routing table independently
  - Route calculations based on Dijkstra's shortest-path algorithm

Internet Routing Protocols (Intradomain)

Open Shortest Path First Protocol (OSPF)

- Routing within an autonomous system (IGP)
- More powerful but more complex than RIP
- Can scale to handle a much larger number of routers than other IGPs
- Uses link-state (SPF) algorithm

OSPF Areas and Efficiency

- Allows subdivision of AS into areas
- Link-status information propagated within area
- Routes summarized before being propagated to another area
- Reduces overhead (less broadcast traffic)

OSPF AS Organized in a 2-level Hierarchy

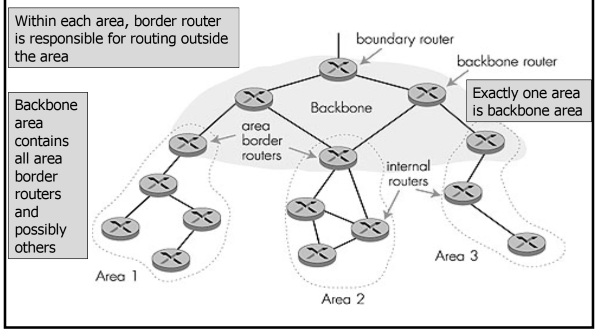

Link-Status in the Internet

- Router corresponds to a node in a graph
- Network corresponds to an edge
- Adjacent pair of routers periodically
  - Test connectivity
  - Broadcast link-status information to area
- Each router uses link-status messages to compute shortest paths

Illustration of OSPF Graph

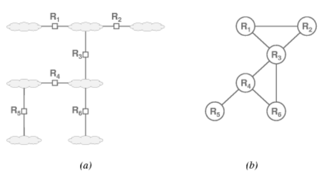

(a) an interconnection of routers and networks, and

(b) an equivalent OSPF graph