

CS558L: Networking Background for Lab Exercises

Instructor: Dr. Wei Ye

<http://www.isi.edu/~weiye/teaching/cs558/>

Outline

- Introduction
- IP routing concepts
- Lab 1: BGP Route Oscillation
- Lab 2: Network security
- Lab 3: TCP flow behavior

Introduction

- What to learn from lab exercises?
 - Hands-on experiences on router and network configuration
 - Traffic measurement and analysis
- Configure Cisco routers
- Configure Linux workstations as routers using the Zebra software toolkit
- This talk gives you networking background knowledge to start with

Next...

- Introduction
- IP routing concepts
- Lab 1: BGP
- Lab 2: Network security
- Lab 3: TCP flow behavior

Autonomous Systems

- Internet is a distributed network managed by individual Autonomous systems (ASes)
- What is an AS?
 - A collection of routers and hosts managed by one organization, e.g.,
 - An Internet service provider (ISP)
 - A large corporate network
 - Each AS has a unique AS number for routing
 - Assigned by the American Registry of Internet Numbers (ARIN)

Intra- and Inter-Domain Routing

- Intra-domain routing
 - Routing within an AS
 - Intra-domain routing protocols are also called Interior Gateway Protocols (IGP)
- Inter-domain routing
 - Routing between different ASes
 - Inter-domain routing protocols are also called Exterior Gateway Protocols (EGP)
 - The only widely used EGP is the Border Gateway Protocol (BGP)

Distance Vector vs. Link State

- Interior protocols has two classes
 - distance vector routing protocols
 - Advertise to neighbors about known routes to all destinations with their distances (metrics)
 - Routing Information Protocol (RIP)
 - Interior Gateway Routing Protocol (IGRP)
 - Enhanced IGRP (EIGRP) } Cisco Proprietary
 - Link state routing protocols
 - Propagate the state of each link (connection to each neighbor) to the whole network
 - Example: Open Shortest Path First (OSPF)

Classful vs. Classless Protocols

- IP addresses are historically divided into Class A, B and C
- Subnet mask
 - Used to further divide a network into sub-nets
 - A Class C address with subnet mask of 255.255.255.224 (0xFFFFE0) specifies a subnet with 32 addresses
- Classful routing protocols
 - Assume fixed subnet mask in each subnet (RIPv1)
- Classless routing protocols
 - Explicitly exchange subnet mask – can be different in different subnets (RIPv2, OSPF)

Special Routes

- Static routes
 - Manually configured routes
 - Won't change unless you re-configure them
- Default route
 - Used when the router has no specific route to an address
 - Can be configured as a static route
- Common Linux commands
 - Use *netstat* to kernel routing table
 - Use *route* to add routes manually

Special Network Interfaces

- Ethernet is a common network interface
 - Use *ifconfig* to check all interfaces on Linux
- Loopback interface
 - A virtual interface that packets don't go out
 - Packets to this interface will appear as received ones
- Null interface
 - All traffic to the null interface is discarded
- Virtual interfaces
 - Associate more than one IP addresses to a real interface (e.g. eth0:1, eth0:2...)
 - Each virtual interface can be on a different net

Route Redistribution

- Router can run multiple routing processes
 - From the same or different protocols
- Different routing processes do not share routing information automatically
- Route redistribution allows the sharing of routing information among all routing processes
 - Can specify how each protocol treats the routes learned from other protocols

Administrative Distance

- A router can learn routes from many sources
 - e.g., static routes, different routing protocols
- Decide which one to use according to the smallest administrative distances

Routing information source	Administrative distance
Directly connected interface	0
Static route	1
External BGP	20
OSPF	110
RIP	120
Internal BGP	200

Next...

- Introduction
- IP routing concepts
- Lab 1: BGP
- Lab 2: Network security
- Lab 3: TCP flow behavior

External and Internal BGP

- External BGP (E-BGP)
 - Routing between different ASes
 - This is the goal of BGP, e.g., routing between different ISPs
- Internal BGP (I-BGP)
 - I-BGP: exchange external routing info with peer BGP routers in the same AS
 - A large network can have multiple BGP routers, and they are usually not close together
 - Necessary in networks that have multiple paths to the Internet

BGP Route Selection

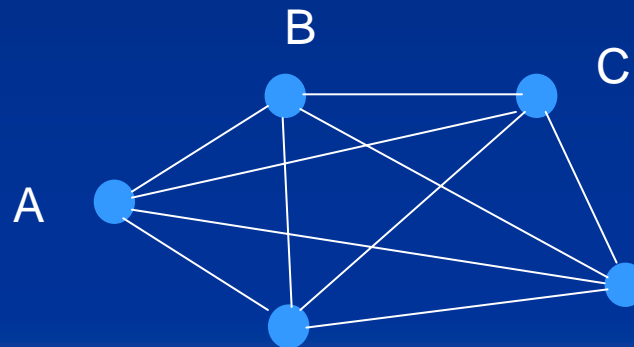
- Based on several parameters
 - Local preference
 - Local measure but shared with I-BGP routers
 - Multi-exit discriminator (MED)
 - Leaves local network and tells neighboring BGP routers in other ASes which link we prefer to receive traffic
 - AS path
 - BGP routing is based on a list of ASes that are traversed in order to reach a destination
 - IGP cost to next-hop router

BGP Route Selection Algorithm

- Brief summary of route selection order
 - Select a route with highest local preference
 - Select a route with minimum AS path length
 - If there are multiple routes to an AS, find out the one with the minimum MED value (No comparison of MED for routes to different ASes)
 - Prefer E-BGP routes over I-BGP routes
 - Select a route with minimum IGP cost to the next-hop router
- Get familiar with the algorithm

I-BGP Router Connection


- An I-BGP router does not advertise a route learned from another I-BGP router
 - If B learns a route from A, it can not tell C about it. C must learn directly from A
 - Requires all BGP routers are fully meshed



- Simple way to prevent loops, but has scaling problem

Route Reflectors

- Route reflectors are allowed to redistribute routes to their clients
 - Only Route reflectors are fully meshed



- Route reflection can have persistent route oscillation – I-BGP routers never converge
 - You need to reproduce it in exercise 1

Next...

- Introduction
- IP routing concepts
- Lab 1: BGP
- Lab 2: Network security
- Lab 3: TCP flow behavior

Denial-of-Service (DoS) Attack

- Flood useless traffic to overwhelm the Victim's CPU, memory and network resources
- Distributed DoS (DDoS)
 - a multitude of compromised systems attack a single target
- IP Spoofing
 - Attackers typically forge, or “spooF”, the IP source addresses of each attack packets to conceal where they are originated (often use random IP addresses)

DoS Attack Types

- Overloading the network
 - Send small packets as quickly as possible
 - Network devices are more limited to packet processing rate than bandwidth
- Overloading the CPU or memory
 - Requires additional processing upon receiving a packet
 - TCP SYN flood: when the victim receives a TCP SYN packet, it must search through all existing connections. If no match is found, allocates a new data structure for the connection

Some Anti-DDoS Techniques

- Blackhole: drop all traffic to the victim
 - Think about the null interface on routers
- Sinkhole: divert all traffic away from target
- Traceback: find out ingress points of an attack
 - Example: ICMP traceback
 - Routers that forward a packet, with a low probability, sends an ICMP packet to destination
 - With enough ICMP packets, the destination is able to reconstruct the real path of the packet

Next... and Last

- Introduction
- IP routing concepts
- Lab 1: BGP
- Lab 2: Network security
- Lab 3: TCP flow behavior

Two major functionalities of TCP

- End-to-end reliability
 - ACK and retransmit
- Congestion control
 - Congestion detection
 - Timeouts on ACKs due to packet drop
 - Back off when congestion is detected
 - Slow start
 - Linear increase/multiplicative decrease
- Refresh your memory about TCP congestion control algorithm

Topics in Exercise 3

- TCP tries to be fair for both short flows (mice) and long flows (elephants)
 - How good is its fairness in practice?
 - Understand effects of TCP congestion control to different flows
- TCP works when everybody follow the same rule in the game
 - What happens if someone does not back off in congestions? e.g., UDP packets
- Service Differentiation
 - Better service: more bandwidth, throughput
 - To do so, router drops more packets from other flows

Random Early Detection (RED)

- TCP detects congestion by ACK losses
 - Packets are dropped by some routers
- RED gateway improves congestion avoidance through queue management
 - Computing average queue size as sign of congestion
 - Early detection if queue size exceeds certain threshold
 - When congestion happens, drops or marks each incoming packet with certain probability
 - Marked packets notifies the connection to slow down
- You will examine in what cases RED helps

How to Get Better Service

- Type of Service (TOS)
 - Four TOS bits in the IP header tell the router how to treat the packet differently
 - Minimum delay, maximum throughput, maximum reliability, Minimum cost – but only one bit can be set at any given time
 - Use *IPtable* in Linux to set the TOS bit
- Routers treat packets without a TOS bit set as normal packets
- Packets with a TOS bit set get better service