Systems and Internet Infrastructure Security

Network and Security Research Center
Department of Computer Science and Engineering
Pennsylvania State University, University Park PA

Interdomain Routing Security

CSE598K/CSE545 - Advanced Network Security
Prof. McDaniel - Spring 2008

Routing redux

PENNSTATE

- The Internet is broken up into Autonomous Systems
- All the hosts in an AS have a single administrative control
- Two types of Routing
  - Intradomain routing
    - Accomplished via OSPF and other protocols
  - Interdomain routing
    - Accomplished only via BGP
      - ASes cooperatively inform each other, for each IP address, in which AS it's located and how to get there.

Routing in a nutshell

PENNSTATE

- The Internet ...

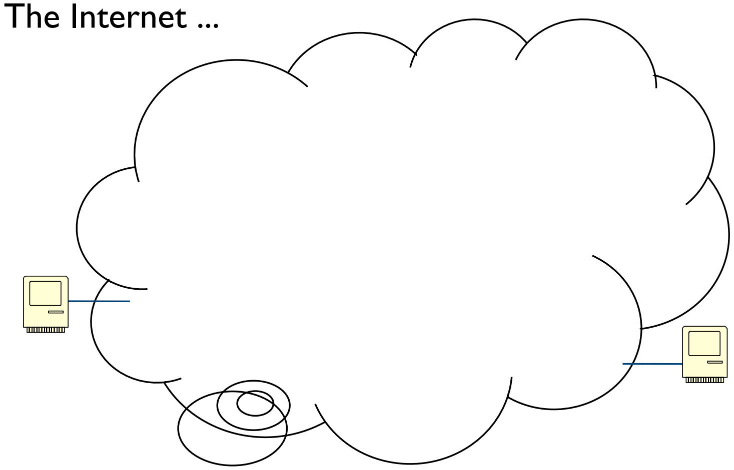

Routing in a nutshell

PENNSTATE

- …is made up of Autonomous Systems (ASes)…

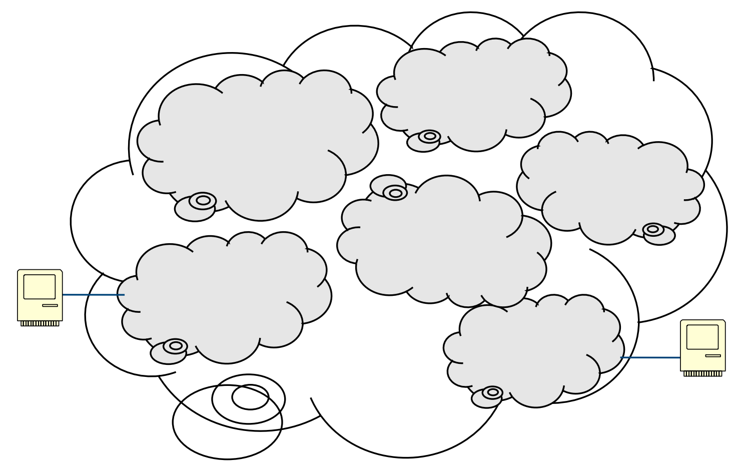

Routing in a nutshell

- …linked at Border Routers.

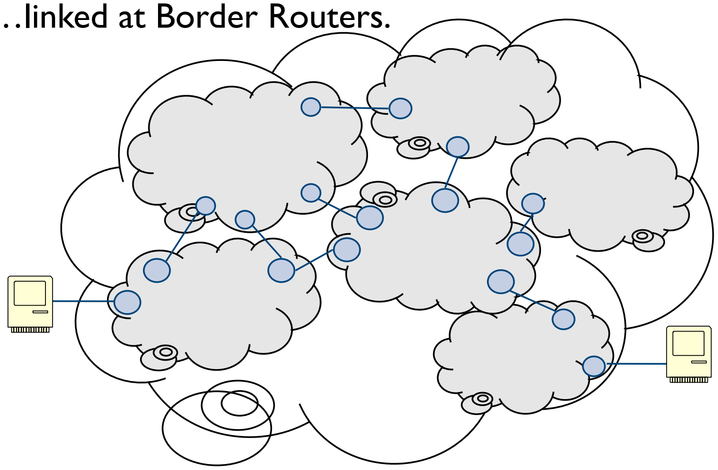

Routing in a nutshell

- The Border Gateway Protocol determines which ASes to follow from source to destination.

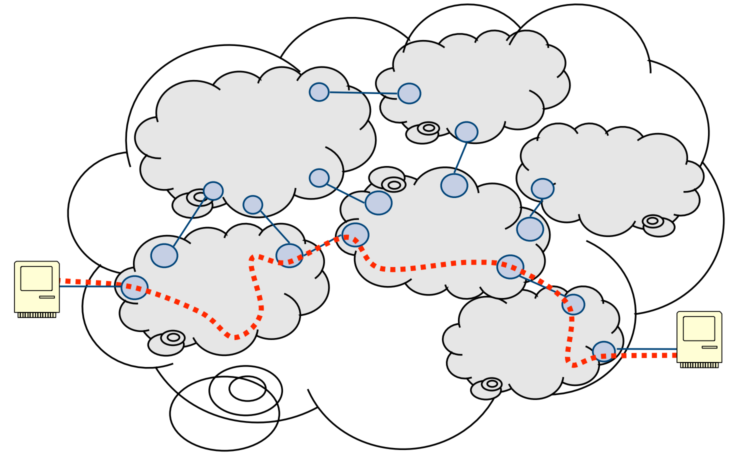

Routing in a nutshell

- Each AS is responsible for moving packets inside it.
- Intra-AS routing is (mostly) independent from Inter-AS routing.

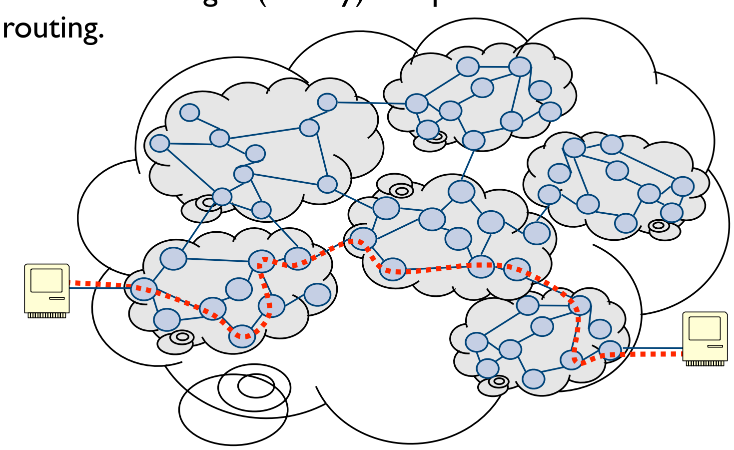

The BGP Protocol

- BGP messages
  - Origin announcements:
    - "I own this block of addresses"
  - Route advertisements:
    - "To get to this address block, send packets destined for it to me. And by the way, here is the path of ASes it will take"
  - Route withdrawals:
    - "Remember the route to this address block I told you about, that path of ASes no longer works"
- Route decisions
  - Border routers receive many origin announcements/ route advertisements, one from each of their peers
  - They choose the "best" path and send their selection downstream
- BGP Attributes
  - BGP messages have additional attributes to help routers choose the "best" path
  - AS_path (above), MED, community strings, …

CIDR Block	Path			Attributes
192.168.28.0/24	768	4014	664	quest:bkup

Routing in a nutshell

- Propagate throughout the network.

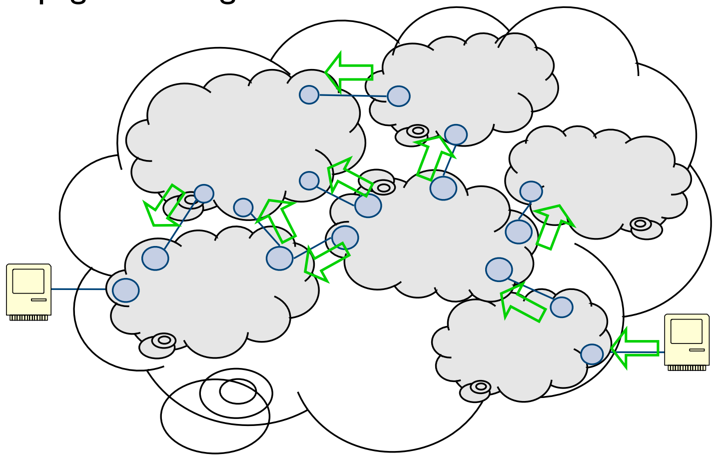

BGP announcements

- Which path gets picked depends on the advertised attributes.

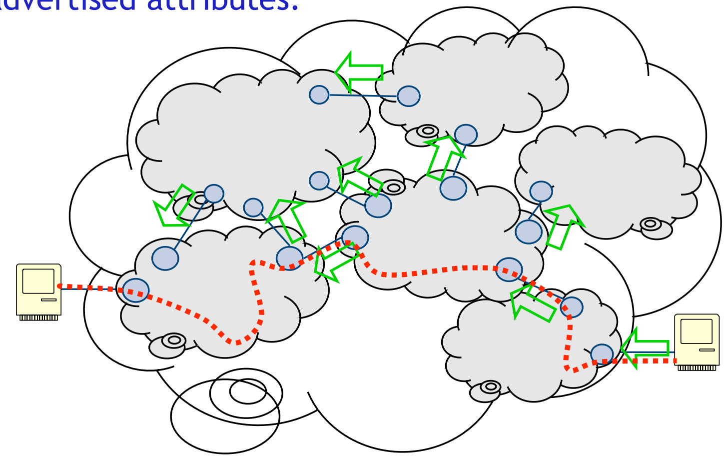

BGP Connection FSM

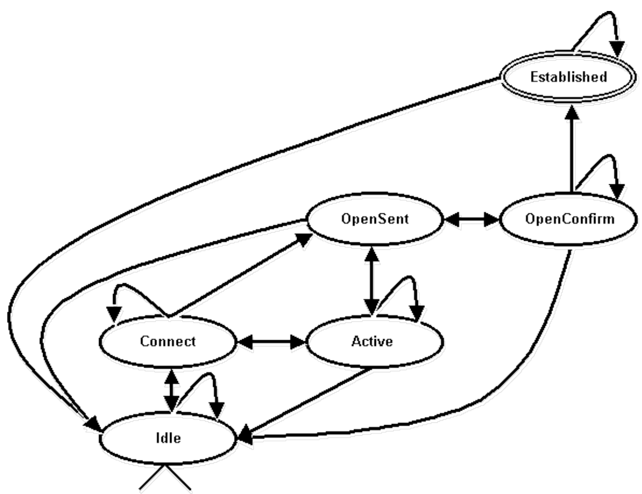

BGP Operation: Connection Setup

- A router is speak BGP with another router, generally physically connected to it, in another AS
  - These two routers are called *BGP peers*
  - Before coming online, the router is in the *Idle* state
- When the router comes on line, it creates a BGP session with its peer
  - BGP runs over TCP, and a TCP connection is made first between the two peers (port 179)
  - The router is in the Connect state during this time
  - When the connection is established, the router moves into the *Established* state

BGP Operation: Information Exchange

- Once the BGP session is active, the peers exchange routing data
  - This information is passed through the UPDATE message
- Contains a list of advertised prefixes, known as network layer reachability information (NLRI), and withdrawn routes
- Prefixes with different policy attributes are sent in separate UPDATE messages
- Route setup can create heavy exchanges of messages and be computationally intensive for the router

BGP Operation: Path Attributes

- ORIGIN: shows whether prefix was learned through interior or exterior routing
- AS_PATH: the ASes that the prefix has passed through during this advertisement
  - BGP is a path vector protocol, and the prefix with the fewest ASes traversed is usually preferred
  - Including AS path vector prevents looping
- NEXT-HOP: the node to send packets back to in order to get them closer to their destination

Other Common Path Attributes

- MULTI-EXIT DISCRIMINATOR: if two ASes connect in multiple locations, the MED can be used by a peer to favour a particular link to improve routing
- LOCAL-PREF: used by the local AS to assign a degree of preference of one link for a given prefix over another
- ATOMIC-AGGREGATE: lets the router know not to deaggregate an advertisement into more specific prefixes
- AGGREGATOR: specifies AS and router that performed aggregation of a prefix

In class exercise

- Fill in the routing tables ...

CIDR Block	AS PATH
192.10.10.0/24	
192.10.0.0/16	
192.10.11.0/16	
192.10.12.0/16	
192.10.14.0/16	

CIDR Block	AS PATH
192.10.10.0/24	
192.10.0.0/16	
192.10.11.0/16	
192.10.12.0/16	
192.10.14.0/16	

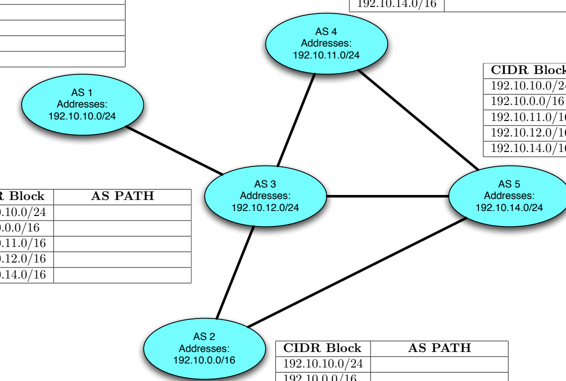

CIDR Block	AS PATH
192.10.10.0/24	
192.10.0.0/16	
192.10.11.0/16	
192.10.12.0/16	
192.10.14.0/16	

CIDR Block	AS PATH
192.10.10.0/24	
192.10.0.0/16	
192.10.11.0/16	
192.10.12.0/16	
192.10.14.0/16	

CIDR Block	AS PATH
192.10.10.0/24	
192.10.0.0/16	
192.10.11.0/16	
192.10.12.0/16	
192.10.14.0/16	

BGP Misconfiguration

PENNSTATE 1855

- One of the largest problems with BGP is misconfiguration
  - Leading cause of instability on the Internet
  - Causes
    - Stupidity
    - Poor configuration tools
    - Under-specified network requirements
  - Often misconfiguration can lurk for months or years before it is detected or its effects felt
    - Changing network topology
    - Unexpected network states

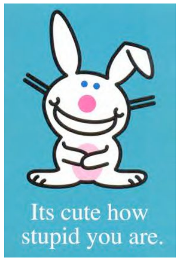

Mahajan et al.

PENNSTATE 1855

- SIGCOMM '02 study of BGP misconfiguration
  - Those instances where configurations caused problems:
    - *unintended suppression* of legitimate advertisement
    - *unintended creation* of illegitimate advertisement
  - Human factors terminology
    - slip - inadvertent errors, e.g., typos
    - mistakes - design errors, e.g.,
- Methodology: use data from RouteViews routing repository collected over 3 years and 23 vantage points located located over the globe.
  - contacted ASes for information on causes

Study Results

PENNSTATE 1855

- Errors detected
  - prefix hijacking - incorrect advertisement of addresses
  - improper route export - exporting routes/paths in violation of stated ISP policies
- Problems are universal, pervasive, and pathological
  - 200-1200 prefixes seeing misconfiguration per day (0.2-1.0% of 2002 table size)
  - 3 in 4 new prefix advertisements result of misconfigurations
  - About 15 hijacks per day (getting much worse)
- Result: constant stream of incorrect information being received by routers.*
- Interesting thought: how to secure in this environment?

*only gets worse after 2002.

Attacks Against BGP

PENNSTATE 1855

- Control Plane
  - Timing
  - Availability
- Data Plane
  - Origin
  - Path

Origin Attacks

- Prefix hijacking
- Prefix destabalization
- Self-deaggregation
- Unauthorized use

- The most serious of the attacks, particularly because they can happen accidentally

Path Attacks

- Path modification
- Path forgery
- Policy modification
- AS forgery

- These attacks can be used to subvert routing and bias the way packets travel through the system

Timing Attacks

- Spoofed OPEN message during negotiation
- TCP SYN attack
- Altering BGP timers
- Forged KEEPALIVE messages while peers are connecting

Availability Attacks

- In-protocol attacks
  - Forged NOTIFICATION messages
  - Syntax errors in BGP messages
  - Forcing route flooding to occur
  - Forged TCP RST packet
- Physical attacks
  - Resetting the router by gaining control of it
  - Link cutting

Prefix Hijacking

PENNSTATE

- An attacker can forge an UPDATE message that claims to originate a known prefix
- For example, my organization could decide to be AT&T for a day, and advertise 12.0.0.0/8
- Outbound route filtering should catch this, but many operators do not perform proper filtering policy within their AS

Prefix Destabilization

PENNSTATE

- By forcing route flapping on a given link, an attacker to a peer can cause BGP dampening to occur
  - Routes that flap are penalized by being suppressed
  - The period of suppression increases depending on how many times the BGP session changes state and length of the prefix (longer prefixes are penalized more than shorter ones)
- Black holes are a major problem of origin attacks

Black Holes are out of sight

PENNSTATE

- If another AS advertises one of our prefixes, bad things happen:

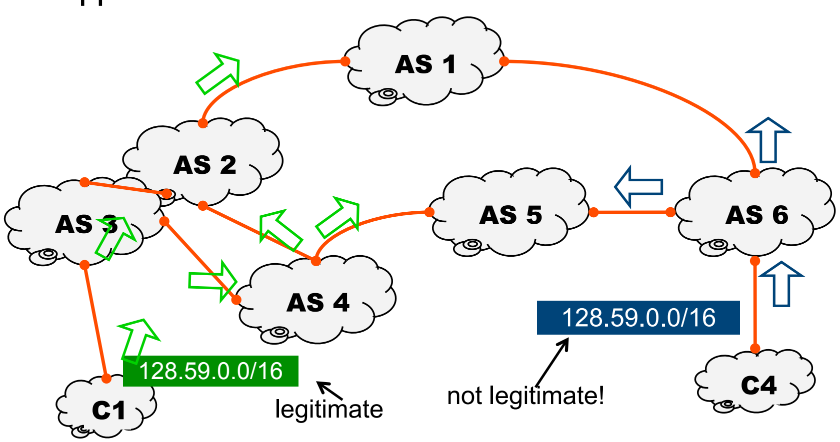

Black Holes are out of sight

PENNSTATE

- Prefix becomes unreachable from the part of the net believing C4's announcement.

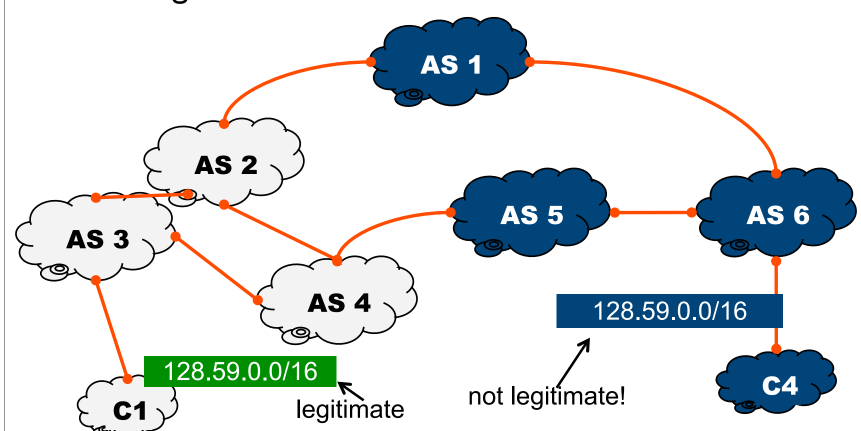

Self-deaggregation

- Within the AS, a prefix can be broken into smaller blocks and advertised as such
- Because of longest-prefix matching, these will be preferred (eg. 12.10.8.0/24 is preferred over 12.0.0.0/8 because it is more specific)
- This is the heart of the AS7007 incident, where much of the Internet lost its routing
- It can also cause a large burden on the routers, because of increase in computation and routing table size

Path Modification

- BGP is a path-vector protocol, so the length of the path is a major factor in accepting a route
- AS path prepending can be used to bias a route (adding the same AS number repeatedly to a route)
- An attacker with the ability to modify the AS path can force traffic to follow patterns it otherwise wouldn't

Path Forgery

- If an AS_PATH attribute is completely forged, the attacker has even more control over traffic
- This can allow for traffic analysis since traffic is engineered in the direction the attacker desires
- This can also lead to black holes, as previously discussed

Policy Modification

- By modifying policy attributes, traffic can be be biased in certain ways and routing can be compromised
- Examples: changing the MED or Local_Pref values can cause suboptimal routing within the peer's or local AS, respectively

TCP SYN Attacks

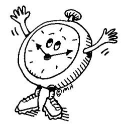

- SYN forgery
  - If the attacker sends a SYN, the peer may think this is a legitimate connection
  - If the attacker guesses the correct SYN ACK, a collision will result, causing the legitimate connection to fail
- SYN-ACK forgery
  - Attacker timing a SYN ACK and sending it during TCP setup can bring down connection
- SYN flood
  - Overwhelm the router resources with SYN packets until it runs out of connections

Spoofing

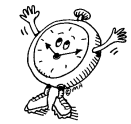

- A forged BGP OPEN message can bring down a connection
  - If a connection is in the process of being opened, an attacker sending an OPEN message can cause a collision
  - Legitimate connection would be terminated
- Similarly, a BGP KEEPALIVE sent while peers are connecting will cause the session to fail
  - If peers are in Connect, Active or OpenSent state

Modifying BGP Timers

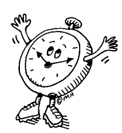

- If the attacker can gain control of timer functionality, messages can be delayed and connections forced closed
  - KeepAlive timer, Hold timer and OpenDelay timer - if altered, messages and the connection itself may be dropped
- KEEPALIVE messages are "heartbeat" messages to ensure the BGP connection exists

Availability Attacks through BGP

- Forged NOTIFICATION message
  - NOTIFICATION is indicative of an error, so whenever this message is passed, the connection is brought down and the peer states change to Idle
- Syntax or parse errors with BGP messages
  - If a packet is malformed, values are invalid or message headers contain errors, the peer will drop the connection

Route Flooding

- Any attack that brings down the causes a connection to bounce will force its peers to dump their routing tables to it
  - These can overwhelm the router depending on the number of routes it receives, and is computationally and bandwidth intensive in any case
- Route flapping also an availability attack
  - Penalized by BGP dampening algorithms that force suppression of the advertisement

Physical Attacks

- Link cutting
  - If the attacker knows the network topology, bringing down certain links (through DoS attacks, or a backhoe) can force traffic into the pattern they desire
- Taking control of the router
  - For example, exploiting a buffer overflow (such as the SNMP attack)
  - Can cause the router to reboot
- Physical destruction of the router
  - As always, network security is dependent on physical security

Conclusions

- There are concerns about BGP's vulnerability, particularly to deliberate attack
- Origin-based attacks are the most pressing concern because of their feasibility
- Lots of research remains to be done in this area
  - We are working on origin and path authentication methods, and discussing security proposals with Cisco
- *Next lecture*: the solutions