Interdomain Routing Security

COS 461: Computer Networks

Jennifer Rexford

Interdomain Routing

- AS-level topology
  - –Nodes are Autonomous Systems (ASes)
  - –Edges are links and business relationships

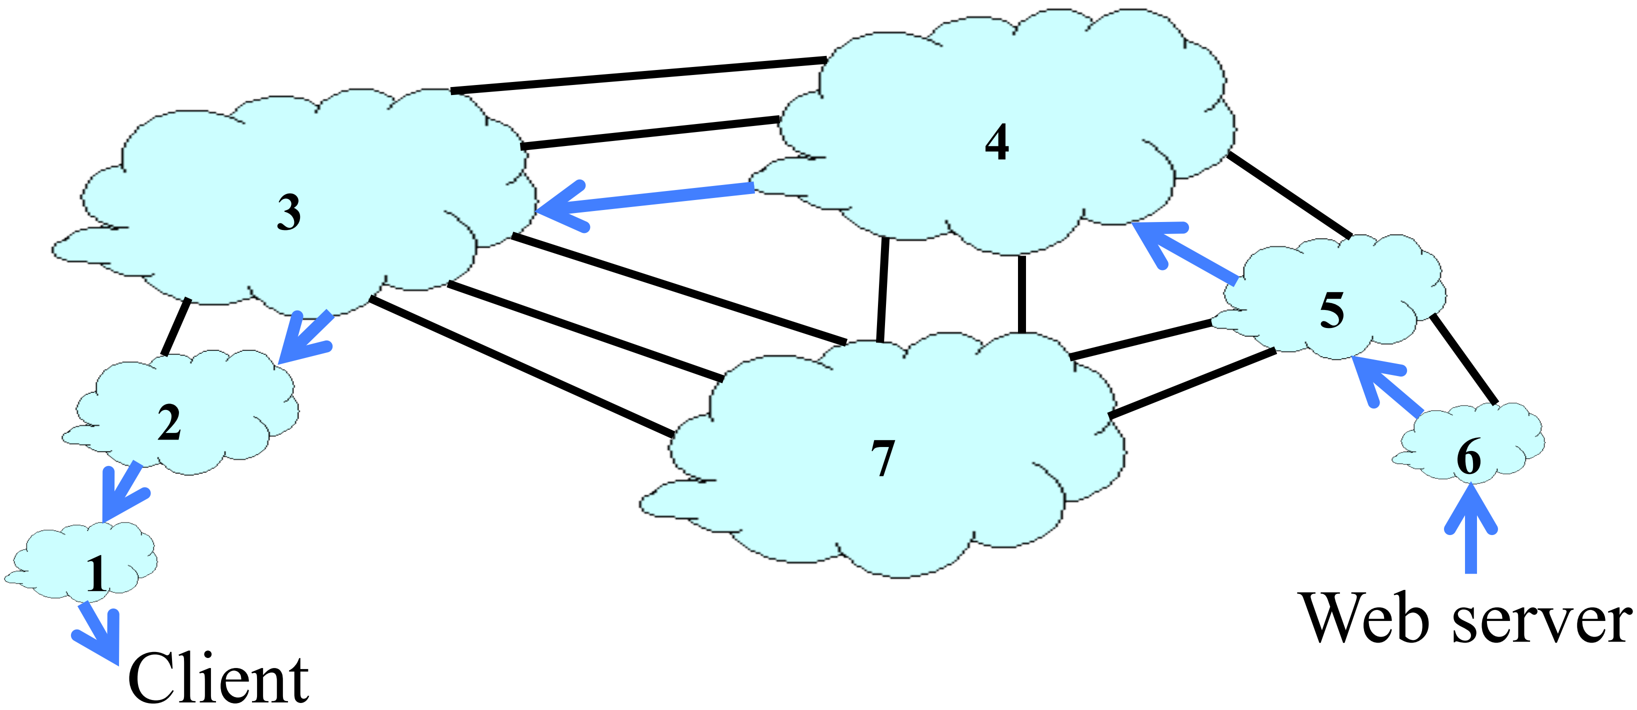

Border Gateway Protocol (BGP)

- ASes exchange reachability information
  - Destination: block of addresses (an "IP prefix")
  - AS path: sequence of ASes along the path
- Policies configured by network operators
  - Path selection: which of the paths to use?
  - Path export: which neighbors to tell?

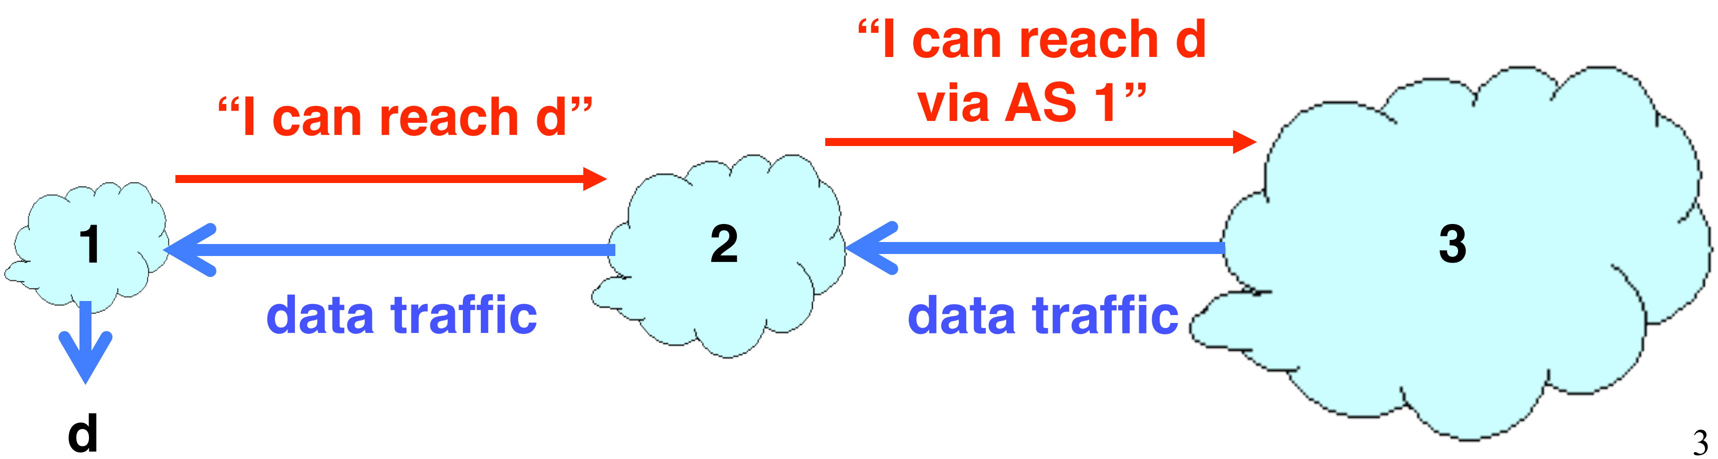

Goals of Today's Lecture

- BGP security vulnerabilities
  - BGP session
  - Prefix ownership
  - AS-path attribute
- Improving BGP security
  - Protective filtering
  - Cryptographic variant of BGP
  - Anomaly-detection schemes
- Data-plane attacks
- Difficulty of upgrading BGP

BGP Session Security

TCP Connection Underlying BGP Session

- BGP session runs over TCP
  - TCP connection between neighboring routers
  - BGP messages sent over TCP connection
  - Makes BGP vulnerable to attacks on TCP

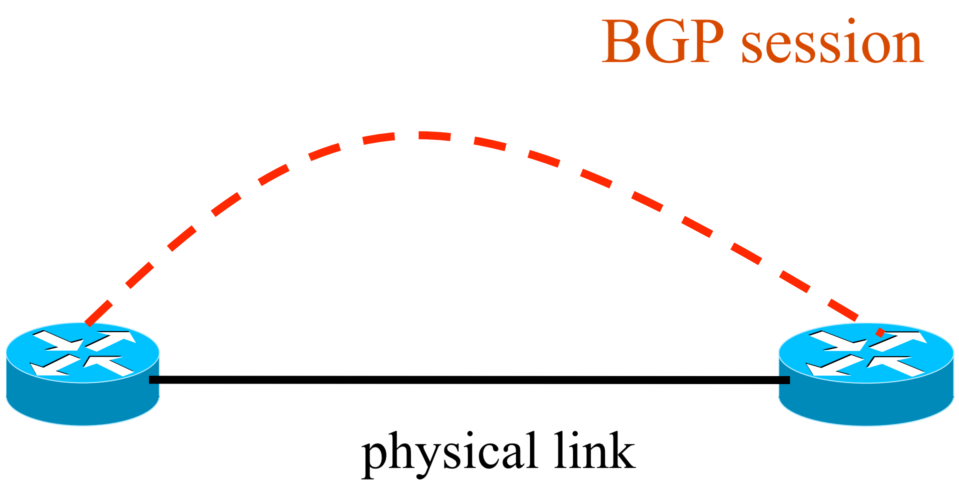

Attacks on Session Security

- Confidentiality
  - Eavesdropping by tapping the link
  - Inferring routing policies and stability
- Integrity
  - Tampering by dropping, modifying, or adding packets
  - Changing, filtering, or replaying BGP routes
- Availability
  - Resetting the session or congesting the link
  - Disrupting communication and overloading the routers

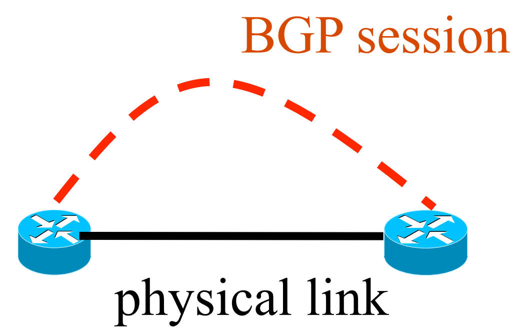

Defending Session Security is Easy

- BGP routing information is propagated widely
  - Confidentiality isn't all that important
- Two end-points have a business relationship
  - Use known IP addresses and ports to communicate
  - Can agree to sign and encrypt messages
- Limited physical access to the path
  - Direct physical link, often in same building
- Low volume of special traffic
  - Filter packets from unexpected senders
  - Filter packets that travel more than one hop
  - Can give BGP packets higher priority

Validity of the routing information: Origin authentication

IP Address Ownership and Hijacking

- IP address block assignment
  - Regional Internet Registries (ARIN, RIPE, APNIC)
  - Internet Service Providers
- Proper origination of a prefix into BGP
  - By the AS who owns the prefix
  - … or, by its upstream provider(s) in its behalf
- However, what's to stop someone else?
  - Prefix hijacking: another AS originates the prefix
  - BGP does not verify that the AS is authorized
  - Registries of prefix ownership are inaccurate

Prefix Hijacking

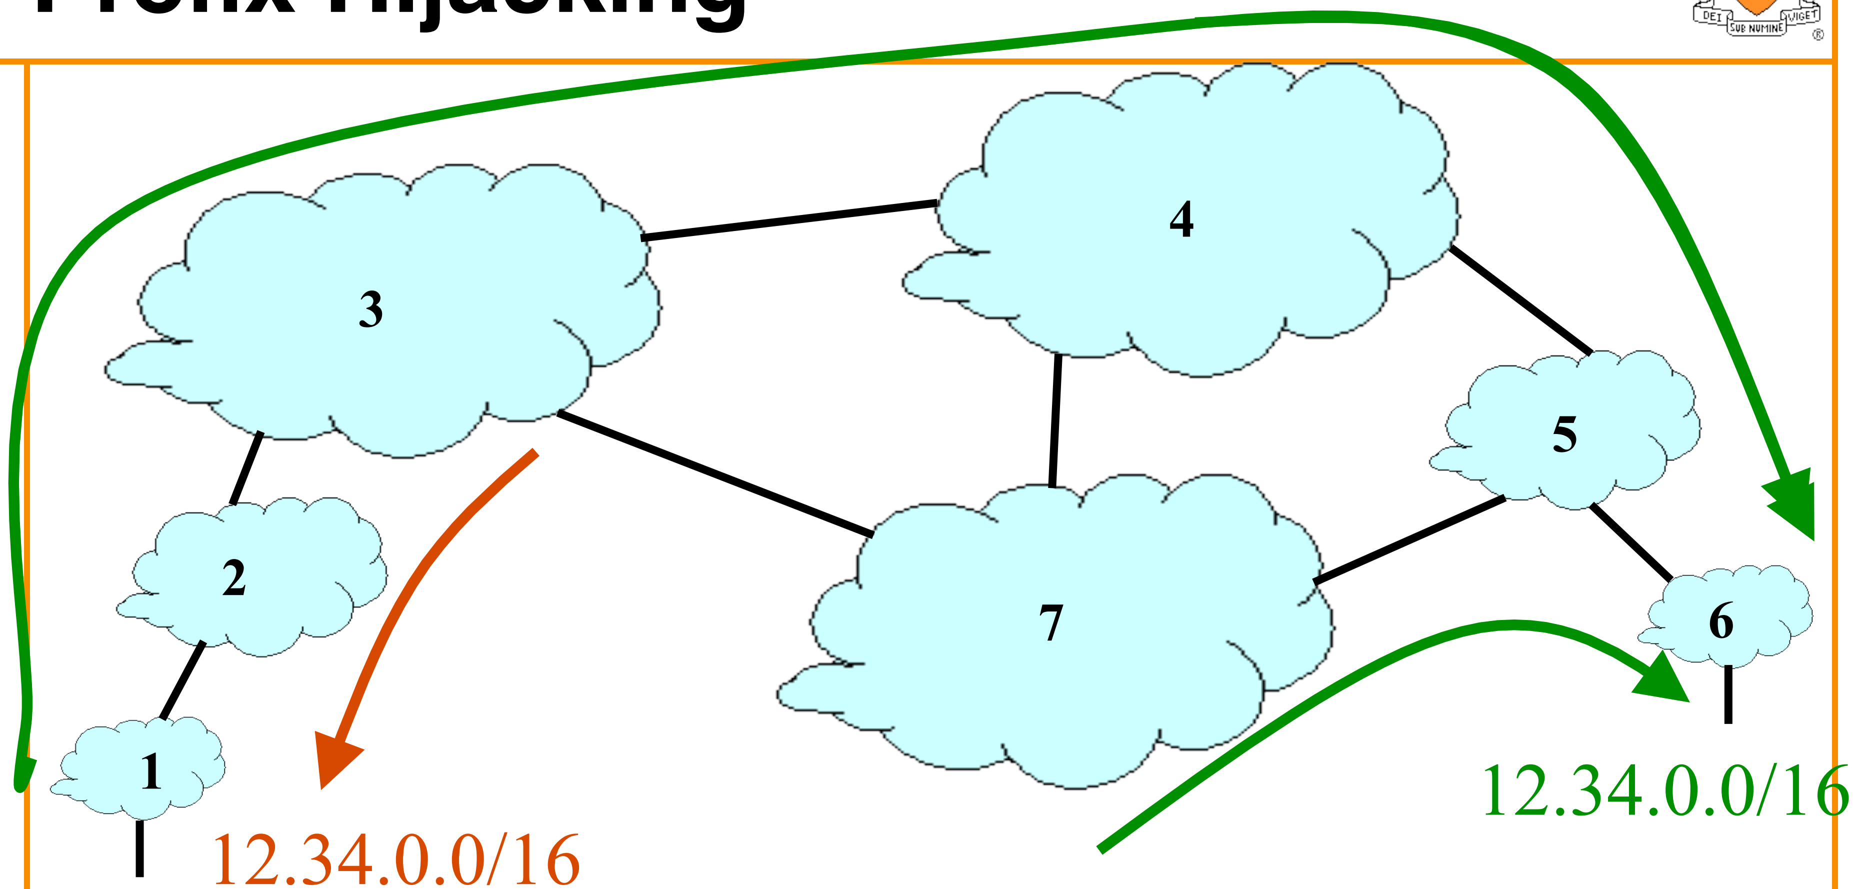

- Consequences for the affected ASes
  - Blackhole: data traffic is discarded
  - Snooping: data traffic is inspected, and then redirected
  - Impersonation: data traffic is sent to bogus destinations

Hijacking is Hard to Debug

- The victim AS doesn't see the problem
  - Picks its own route
  - Might not even learn the bogus route
- May not cause loss of connectivity
  - E.g., if the bogus AS snoops and redirects
  - … may only cause performance degradation
- Or, loss of connectivity is isolated
  - E.g., only for sources in parts of the Internet
- Diagnosing prefix hijacking
  - Analyzing updates from many vantage points
  - Launching traceroute from many vantage points

Sub-Prefix Hijacking

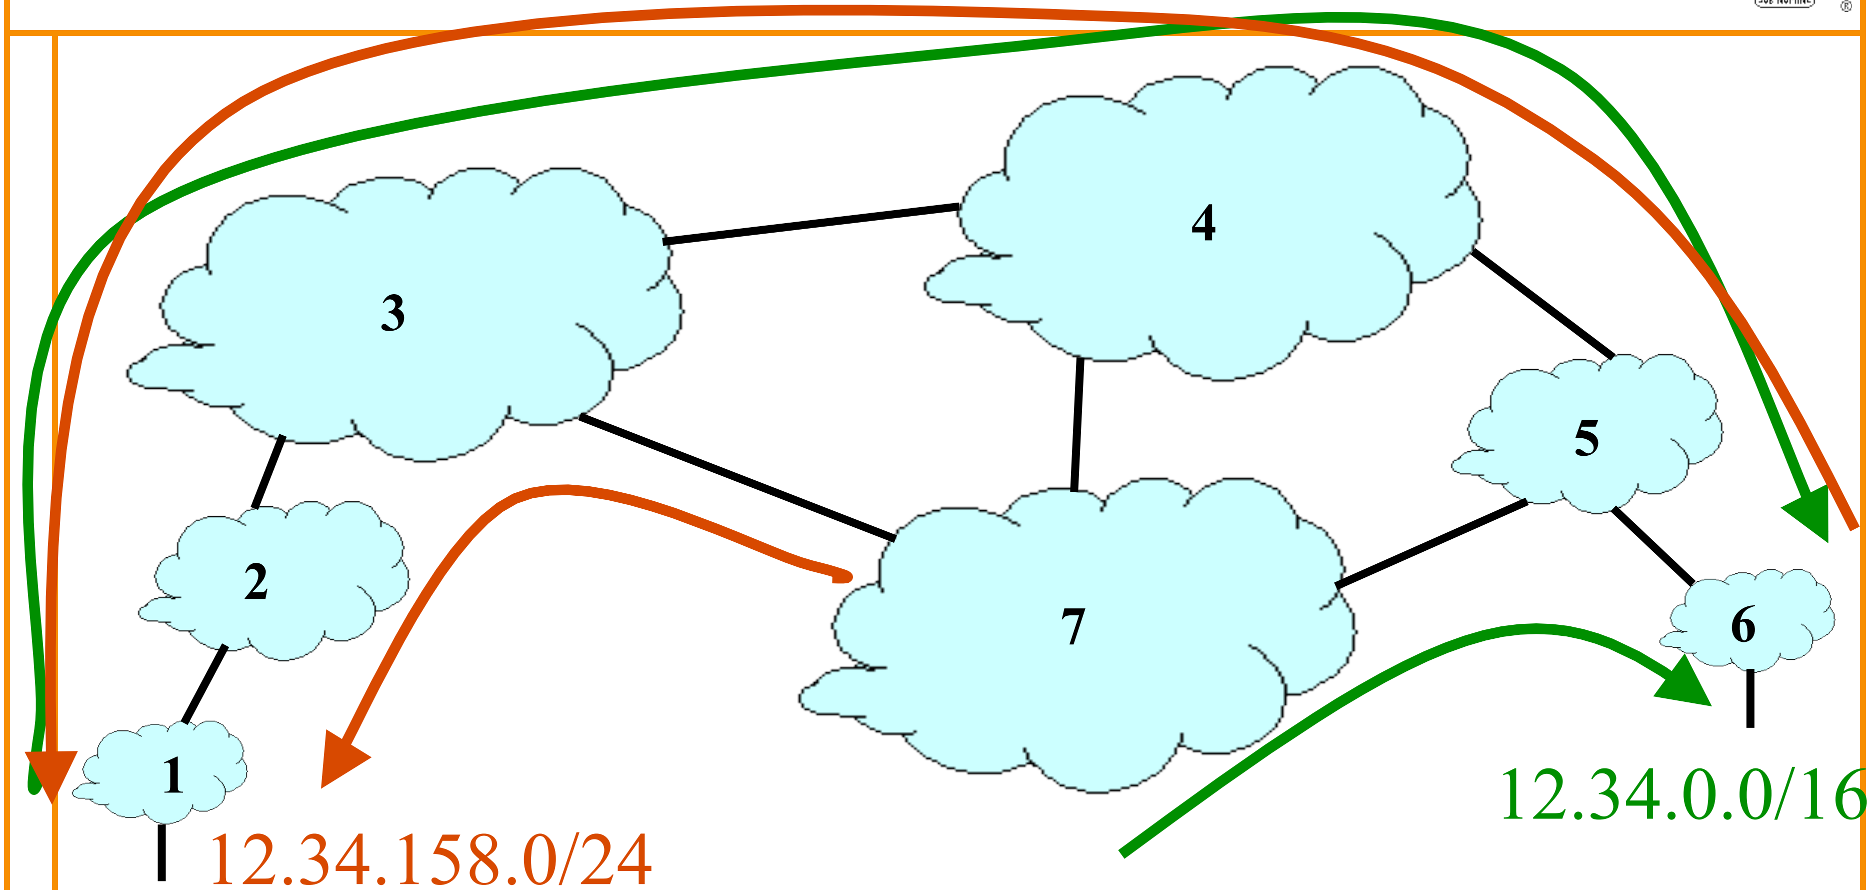

- Originating a more-specific prefix
  - Every AS picks the bogus route for that prefix
  - Traffic follows the longest matching prefix

How to Hijack a Prefix

- The hijacking AS has
  - Router with BGP session(s)
  - Configured to originate the prefix
- Getting access to the router
  - Network operator makes configuration mistake
  - Disgruntled operator launches an attack
  - Outsider breaks in to the router and reconfigures
- Getting other ASes to believe bogus route
  - Neighbor ASes do not discard the bogus route
  - E.g., not doing protective filtering

YouTube Outage on Feb 24, 2008

- YouTube (AS 36561)
  - Web site www.youtube.com
  - Address block 208.65.152.0/22
- Pakistan Telecom (AS 17557)
  - Receives government order to block access to YouTube
  - Starts announcing 208.65.153.0/24 to PCCW (AS 3491)
  - All packets directed to YouTube get dropped on the floor
- Mistakes were made
  - AS 17557: announcing to everyone, not just customers
  - AS 3491: not filtering routes announced by AS 17557
- Lasted 100 minutes for some, 2 hours for others

Timeline (UTC Time)

- 18:47:45
  - First evidence of hijacked /24 route propagating in Asia
- 18:48:00
  - Several big trans-Pacific providers carrying the route
- 18:49:30
  - Bogus route fully propagated
- 20:07:25
  - YouTube starts advertising the /24 to attract traffic back
- 20:08:30
  - Many (but not all) providers are using the valid route

http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

Timeline (UTC Time)

- 20:18:43
  - YouTube starts announcing two more-specific /25 routes
- 20:19:37
  - Some more providers start using the /25 routes
- 20:50:59
  - AS 17557 starts prepending ("3491 17557 17557")
- 20:59:39
  - AS 3491 disconnects AS 17557
- 21:00:00
  - All is well, videos of cats flushing toilets are available

http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

Another Example: Spammers

- Spammers sending spam
  - Form a (bidirectional) TCP connection to a mail server
  - Send a bunch of spam e-mail
  - Disconnect and laugh all the way to the bank
- But, best not to use your real IP address
  - Relatively easy to trace back to you
- Could hijack someone's address space
  - But you might not receive all the (TCP) return traffic
  - And the legitimate owner of the address might notice
- How to evade detection
  - Hijack unused (i.e., unallocated) address block in BGP
  - Temporarily use the IP addresses to send your spam

BGP AS Path

Bogus AS Paths

- Remove ASes from the AS path
  - E.g., turn “701 3715 88” into “701 88”
- Motivations
  - Make the AS path look shorter than it is
  - Attract sources that normally try to avoid AS 3715
  - Help AS 88 look like it is closer to the Internet’s core
- Who can tell that this AS path is a lie?
  - Maybe AS 88 *does* connect to AS 701 directly

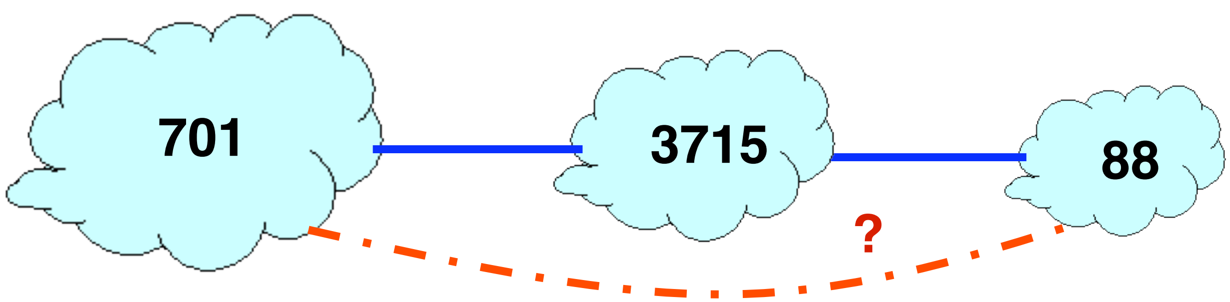

Bogus AS Paths

- Add ASes to the path
  - E.g., turn "701 88" into "701 3715 88"
- Motivations
  - Trigger loop detection in AS 3715
    - Denial-of-service attack on AS 3715
    - Or, blocking unwanted traffic coming from AS 3715!
  - Make your AS look like is has richer connectivity
- Who can tell the AS path is a lie?
  - AS 3715 could, if it could see the route
  - AS 88 could, but would it really care as long as it received data traffic meant for it?

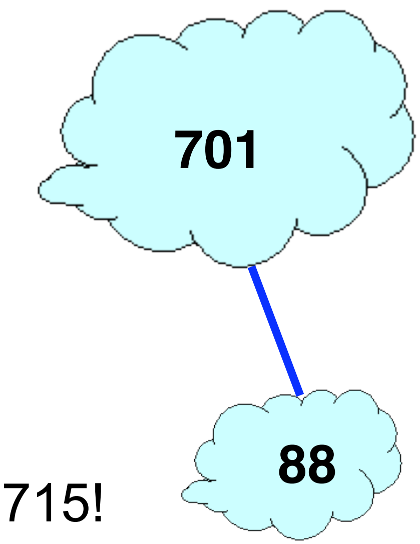

Bogus AS Paths

- Adds AS hop(s) at the end of the path
  - E.g., turns "701 88" into "701 88 3"
- Motivations
  - Evade detection for a bogus route
  - E.g., by adding the legitimate AS to the end
- Hard to tell that the AS path is bogus…
  - Even if other ASes filter based on prefix ownership

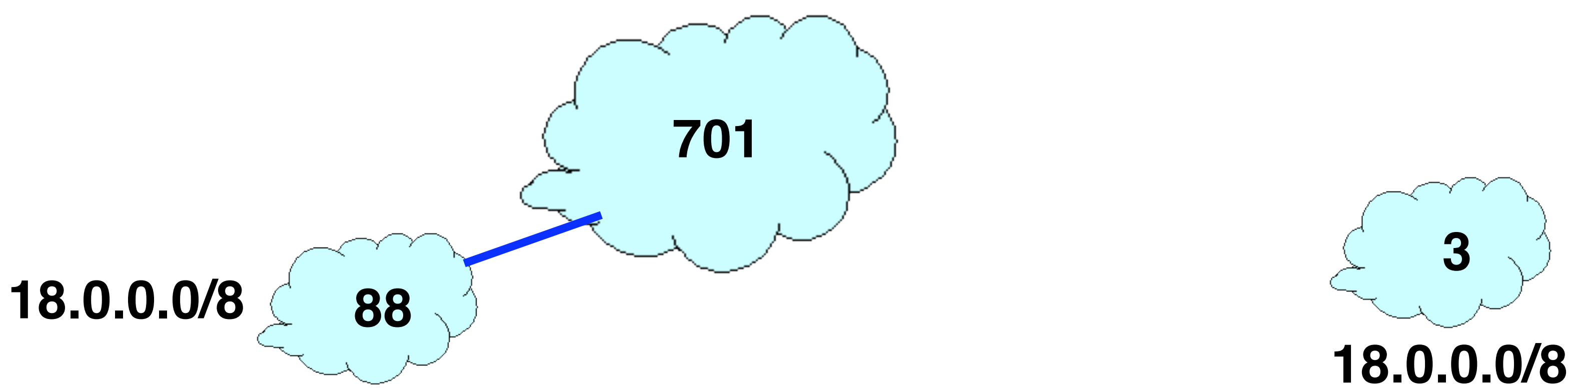

Invalid Paths

- AS exports a route it shouldn't
  - AS path is a valid sequence, but violated policy
- Example: customer misconfiguration
  - Exports routes from one provider to another
- … interacts with provider policy
  - Provider prefers customer routes
  - … so picks these as the best route

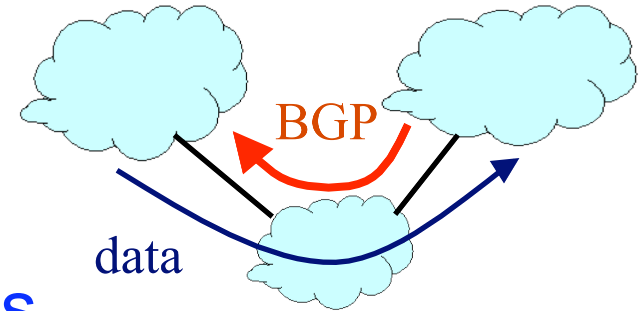

- … leading the dire consequences
  - Directing all Internet traffic through customer
- Main defense
  - Filtering routes based on prefixes and AS path

Missing/Inconsistent Routes

- Peers require consistent export
  - Prefix advertised at all peering points
  - Prefix advertised with same AS path length
- Reasons for violating the policy
  - Trick neighbor into "cold potato"
  - Configuration mistake
- Main defense
  - Analyzing BGP updates
  - … or data traffic
  - … for signs of inconsistency

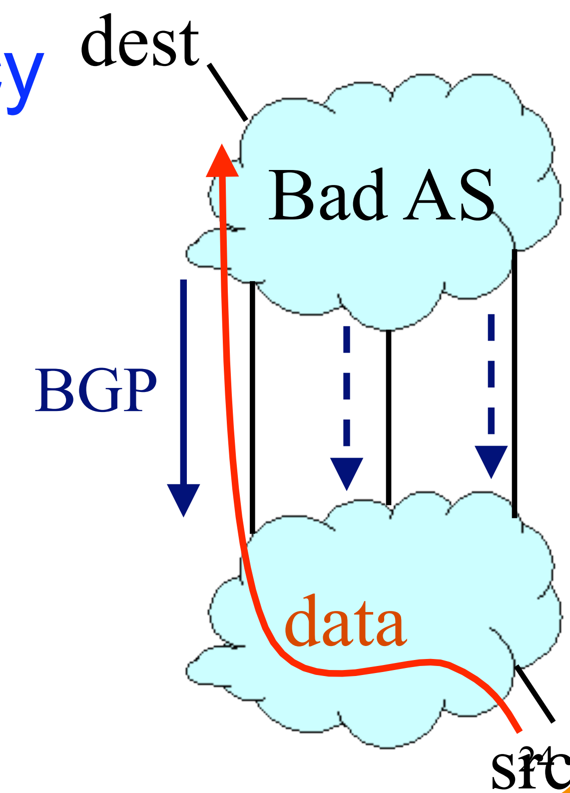

BGP Security Today

- Applying best common practices (BCPs)
  - –Securing the session (authentication, encryption)
  - –Filtering routes by prefix and AS path
  - –Packet filters to block unexpected control traffic
- This is not good enough
  - –Depends on vigilant application of BCPs
    - … and not making configuration mistakes!
  - –Doesn't address fundamental problems
    - Can't tell who owns the IP address block
    - Can't tell if the AS path is bogus or invalid
    - Can't be sure the data packets follow the chosen route

Proposed Enhancements to BGP

S-BGP Secure Version of BGP

- Address attestations
  - Claim the right to originate a prefix
  - Signed and distributed out-of-band
  - Checked through delegation chain from ICANN
- Route attestations
  - Distributed as an attribute in BGP update message
  - Signed by each AS as route traverses the network
  - Signature signs previously attached signatures
- S-BGP can validate
  - AS path indicates the order ASes were traversed
  - No intermediate ASes were added or removed

S-BGP Deployment Challenges

- Complete, accurate registries
  - E.g., of prefix ownership
- Public Key Infrastructure
  - To know the public key for any given AS
- Cryptographic operations
  - E.g., digital signatures on BGP messages
- Need to perform operations quickly
  - To avoid delaying response to routing changes
- Difficulty of incremental deployment
  - Hard to have a “flag day” to deploy S-BGP

Incrementally Deployable Solutions?

- Backwards compatible
  - No changes to router hardware or software
  - No cooperation from other ASes
- Incentives for early adopters
  - Security benefits for ASes that deploy the solution
  - … and further incentives for others to deploy
- What kind of solutions are possible?
  - Detecting suspicious routes
  - … and then filtering or depreferencing them

Detecting Suspicious Routes

- Monitoring BGP update messages
  - Use past history as an implicit registry
- E.g., AS that announces each address block
  - Prefix 18.0.0.0/8 usually originated by AS 3
- E.g., AS-level edges and paths
  - Never seen the subpath “7018 88 1785”
- Out-of-band detection mechanism
  - Generate reports and alerts
  - Internet Alert Registry: http://iar.cs.unm.edu/
  - Prefix Hijack Alert System: http://phas.netsec.colostate.edu/

Avoiding Suspicious Routes

- Soft response to suspicious routes
  - Prefer routes that agree with the past
  - Delay adoption of unfamiliar routes when possible
- Why is this good enough?
  - Some attacks will go away on their own
  - Let someone else be the victim instead of you
  - Give network operators time to investigate
- How well would it work?
  - If top ~40 largest ASes applied the technique
  - … most other ASes are protected, too
  - … since they mostly learn routes from the big ASes

http://www.cs.princeton.edu/~jrex/papers/pgbgp++.pdf

What About Packet Forwarding?

Control Plane vs. Data Plane

- Control plane
  - BGP is a routing protocol
  - BGP security concerns validity of routing messages
  - I.e., did the BGP message follow the sequence of ASes listed in the AS-path attribute
- Data plane
  - Routers forward data packets
  - Supposedly along the path chosen in the control plane
  - But what ensures that this is true?

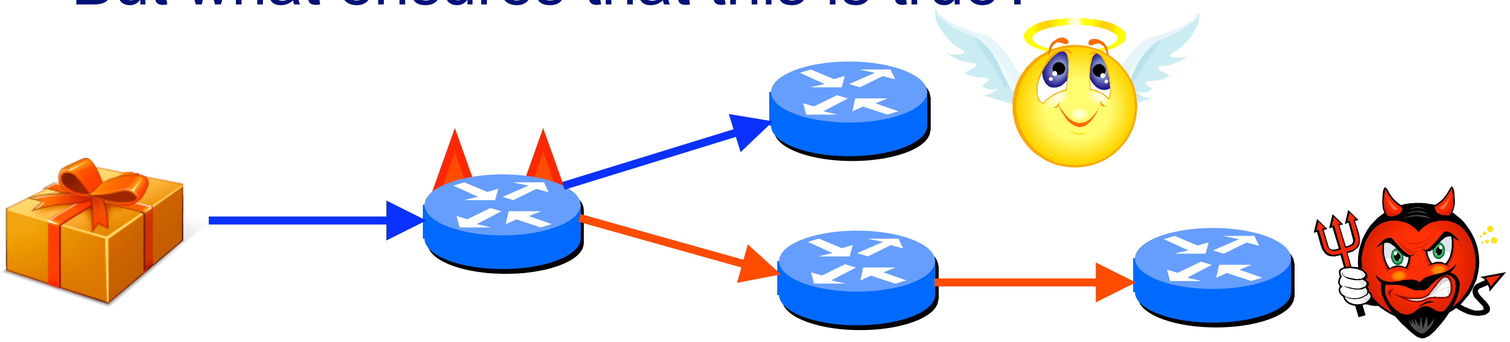

Data-Plane Attacks, Part 1

- Drop packets in the data plane
  - While still sending the routing announcements
- Easier to evade detection
  - Especially if you only drop some packets
  - Like, oh, say, BitTorrent or Skype traffic
- Even easier if you just slow down some traffic
  - How different are normal congestion and an attack?
  - Especially if you let ping/traceroute packets through?

Data-Plane Attacks, Part 2

- Send packets in a different direction
  - Disagreeing with the routing announcements
- Direct packets to a different destination
  - E.g., one the adversary controls
- What to do at that bogus destination?
  - Impersonate the legitimate destination (e.g., to perform identity theft, or promulgate false information)
  - Snoop on the traffic and forward along to real destination
- How to detect?
  - Traceroute? Longer than usual delays?
  - End-to-end checks, like site certificate or encryption?

Fortunately, Data-Plane Attacks are Harder

- Adversary must control a router along the path
  - So that the traffic flows through him
- How to get control a router
  - Buy access to a compromised router online
  - Guess the password
  - Exploit known router vulnerabilities
  - Insider attack (disgruntled network operator)
- Malice vs. greed
  - Malice: gain control of someone else’s router
  - Greed: Verizon DSL blocks Skype to gently encourage me to pick up my landline phone to use Verizon long distance $service ☺

What's the Internet to Do?

BGP is So Vulnerable

- Several high-profile outages
  - http://merit.edu/mail.archives/nanog/1997-04/msg00380.html
  - http://www.renesys.com/blog/2005/12/internetwide_nearcatastrophela.shtml
  - http://www.renesys.com/blog/2006/01/coned_steals_the_net.shtml
  - http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml
  - http://www.theregister.co.uk/2010/04/09/china_bgp_interweb_snafu/
- Many smaller examples
  - Blackholing a single destination prefix
  - Hijacking unallocated addresses to send spam
- Why isn't it an even bigger deal?
  - Really, most big outages are configuration errors
  - Most bad guys want the Internet to stay up
  - … so they can send unwanted traffic

BGP is So Hard to Fix

- Complex system
  - Large, with around 40,000 ASes
  - Decentralized control among competitive ASes
  - Core infrastructure that forms the Internet
- Hard to reach agreement on the right solution
  - S-BGP with public key infrastructure, registries, crypto?
  - Who should be in charge of running PKI and registries?
  - Worry about data-plane attacks or just control plane?
- Hard to deploy the solution once you pick it
  - Hard enough to get ASes to apply route filters
  - Now you want them to upgrade to a new protocol
  - … all at the exact same moment?

Conclusions

- Internet protocols designed based on trust
  - The insiders are good guys
  - All bad guys are outside the network
- Border Gateway Protocol is very vulnerable
  - Glue that holds the Internet together
  - Hard for an AS to locally identify bogus routes
  - Attacks can have very serious global consequences
- Proposed solutions/approaches
  - Secure variants of the Border Gateway Protocol
  - Anomaly detection schemes, with automated response
  - Broader focus on data-plane availability

Backup Slides

Encrypting and Decrypting With Keys

- Encrypt to hide message contents
  - Transforming message contents with a key
  - Message cannot be read without the right key
- Symmetric key cryptography
  - Same secret key for encrypting and decrypting
  - … makes it hard to distribute the secret key
- Asymmetrical (or public key) cryptography
  - Sender uses public key to encrypt message
    - Can be distributed freely!
  - Receiver uses private key to decrypt message

Authenticating the Sender and Contents

- Digital signature for authentication
  - Data attached to the original message
    - … to identify sender and detect tampering
  - Sender encrypts message digest with private key
  - Receiver decrypts message digest with public key
    - … and compares with message digest it computes
- Certificate
  - Collection of information about a person or thing
    - ... with a digital signature attached
  - A trusted third party attaches the signature

Public Key Infrastructure (PKI)

- Problem: getting the right key
  - How do you find out someone's public key?
  - How do you know it isn't someone else's key?
- Certificate Authority (CA)
  - Bob takes public key and identifies himself to CA
  - CA signs Bob's public key with digital signature to create a certificate
  - Alice can get Bob's key and verify the certificate with the CA
- Register once, communicate everywhere
  - Each user only has the CA certify his key
  - Each user only needs to know the CA's public key

Exploiting the IP TTL Field

- BGP speakers are usually one hop apart
  - To thwart an attacker, can check that the packets carrying the BGP message have not traveled far
- IP Time-to-Live (TTL) field
  - Decremented once per hop
  - Avoids packets staying in network forever
- Generalized TTL Security Mechanism (RFC 3682)
  - Send BGP packets with initial TTL of 255
  - Receiving BGP speaker checks that TTL is 254
  - … and flags and/or discards the packet others
- Hard for third-party to inject packets remotely