



Security II: Network Security

Lecture 21

COS 461: Computer Networks

Kyle Jamieson

Today: Network Security

- Last lecture: Foundation Concepts
 - Application layer (Email, Web)
 - Transport layer (TLS/SSL)
 - Network layer (IP Sec)
- **This lecture: Network Infrastructure Security**
 - Naming: Secure DNS (DNS-Sec)
 - Routing: Secure BGP (BGP-Sec)

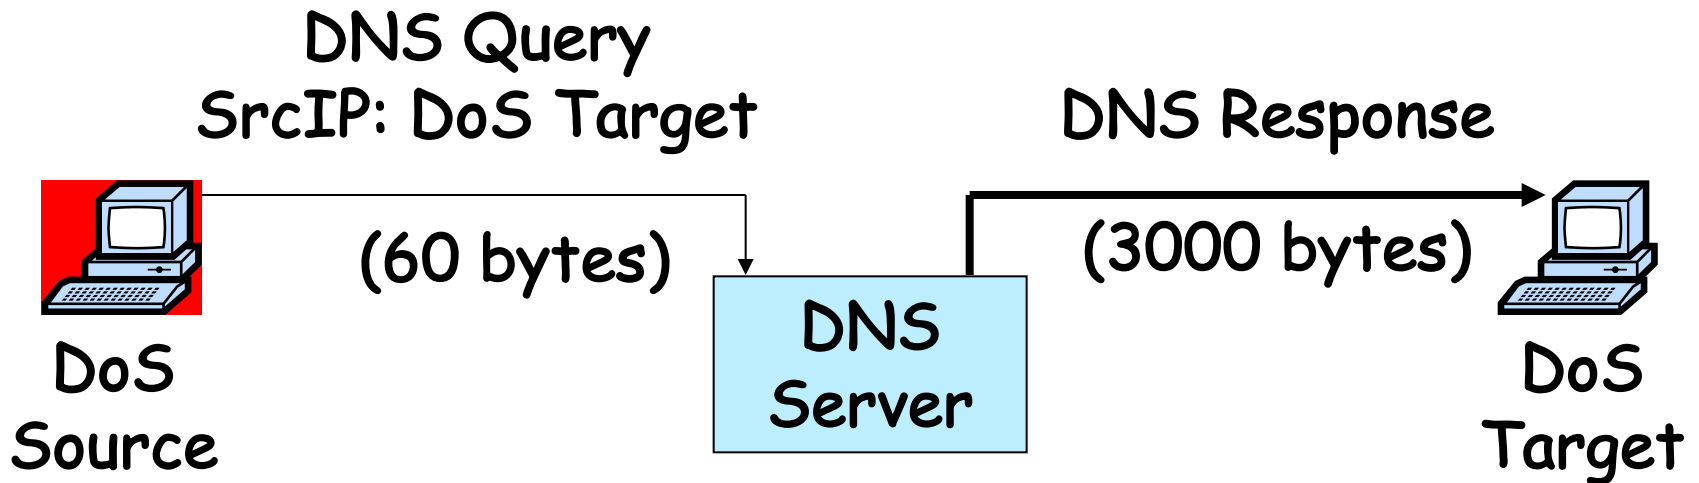
DNS Security

DoS attacks on DNS Availability

- **February 6, 2007**
 - Botnet attack on the 13 Internet DNS root servers
 - Lasted 2.5 hours
 - None crashed, but two performed badly:
 - g-root (DoD), l-root (ICANN)
 - Most other root servers use anycast

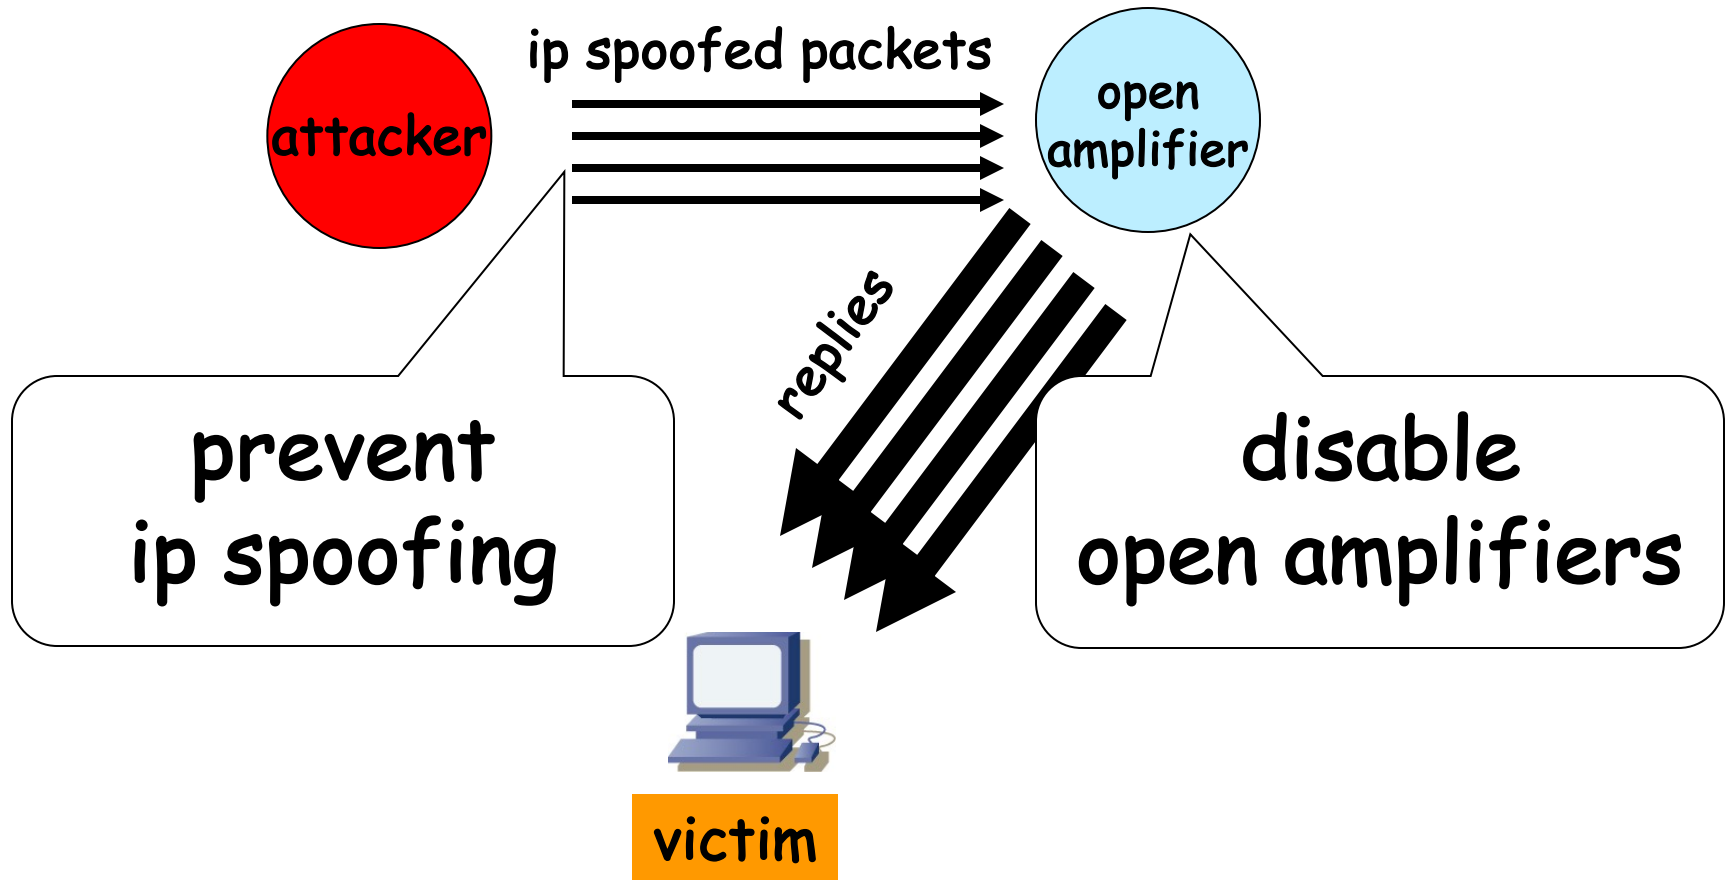
Denial-of-Service Attacks on Hosts

×40 amplification



580,000 open resolvers on Internet (Kaminsky-Shiffman'06)

Preventing Amplification Attacks



DNS Integrity: Cache Poisoning

- Was answer from an authoritative server?
 - Or from somebody else?
- DNS cache poisoning
 - Client (local nameserver) asks for www.evil.com
 - Nameserver authoritative for www.evil.com returns **additional section for (www.cnn.com, 1.2.3.4, A)**
 - Local name server: "Thanks! I won't bother to check what I asked for"

DNS Integrity: DNS Hijacking

- To prevent cache poisoning, client remembers:
 - The domain name in the request
 - A 16-bit request ID (used to demux UDP response)
- DNS hijacking
 - 16 bits: 65K possible IDs
 - What rate to enumerate all in 1 sec? 64B/packet
 - $64 * 65536 * 8 / 1024 / 1024 = 32$ Mbps
- Prevention: also randomize DNS source port
 - Kaminsky attack: this source port... wasn't random

Instead: Let's strongly believe the answer! Enter DNSSEC

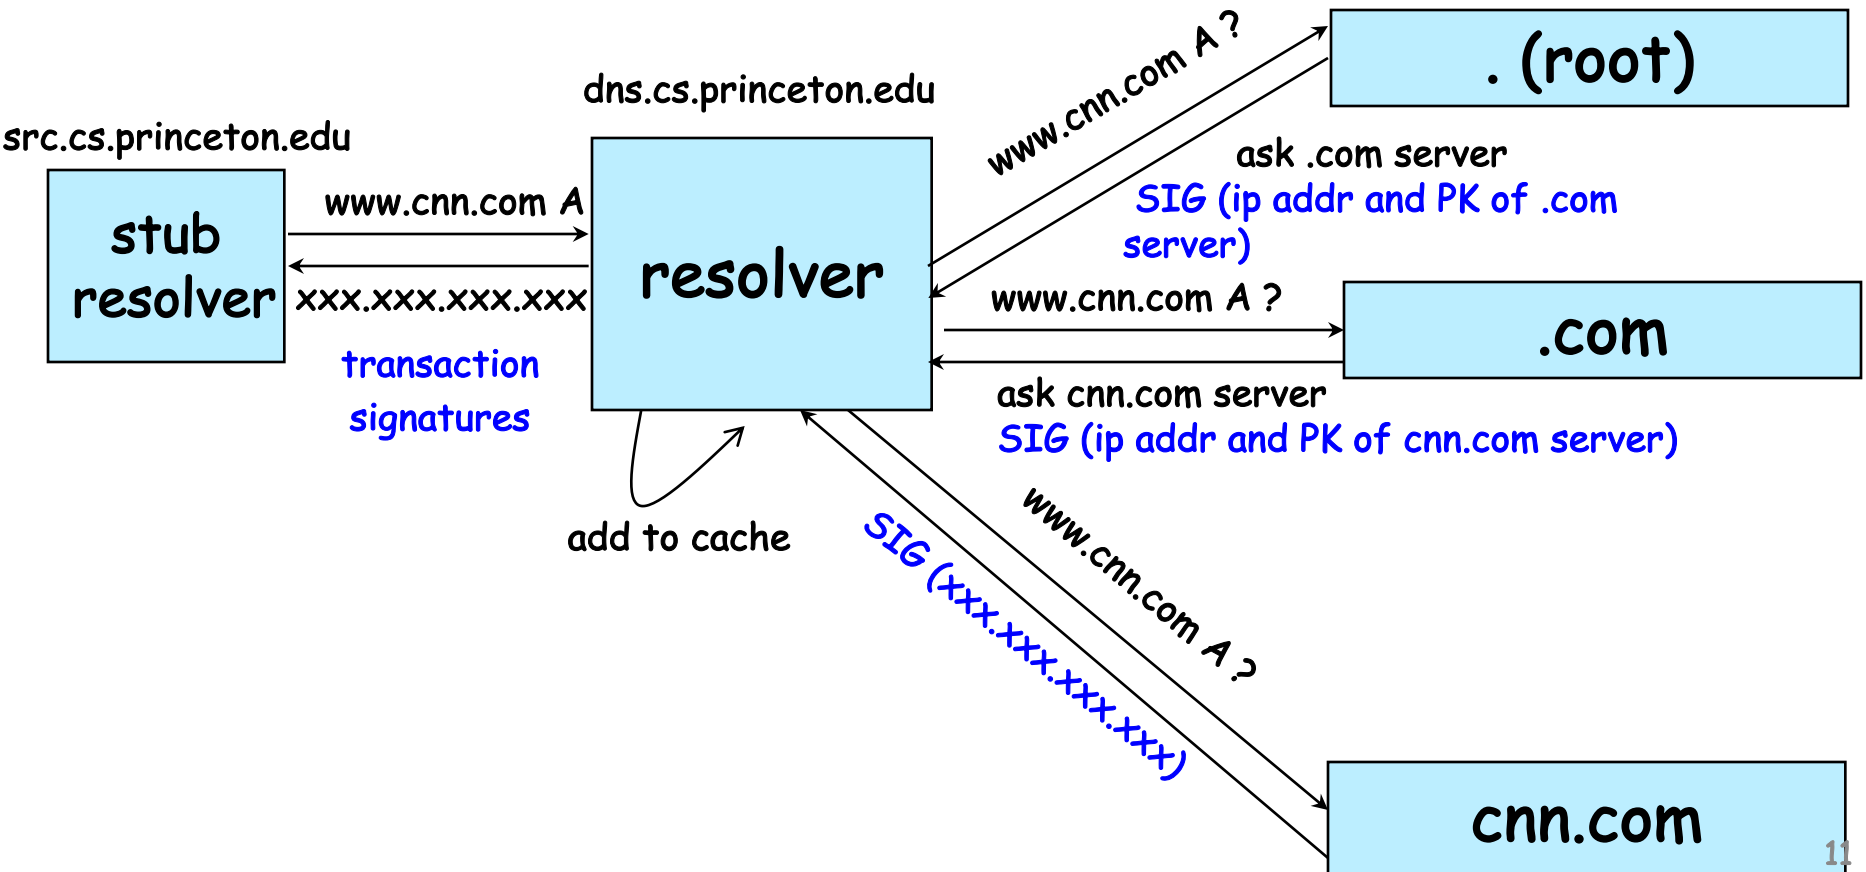
- DNSSEC protects against data spoofing and corruption
- DNSSEC also provides mechanisms to authenticate servers and requests
- DNSSEC provides mechanisms to establish authenticity and integrity

PK-DNSSEC (Public Key)

- The DNS servers sign the hash of resource record set with its private (signature) keys
 - Public keys can be used to verify the SIGs
- Leverages hierarchy:
 - Authenticity of name server's public keys is established by a signature over the keys by the parent's private key
 - In ideal case, only roots' public keys need to be distributed out-of-band

Verifying the Tree

Question: **www.cnn.com** ?




Interdomain Routing Security

Interdomain Routing

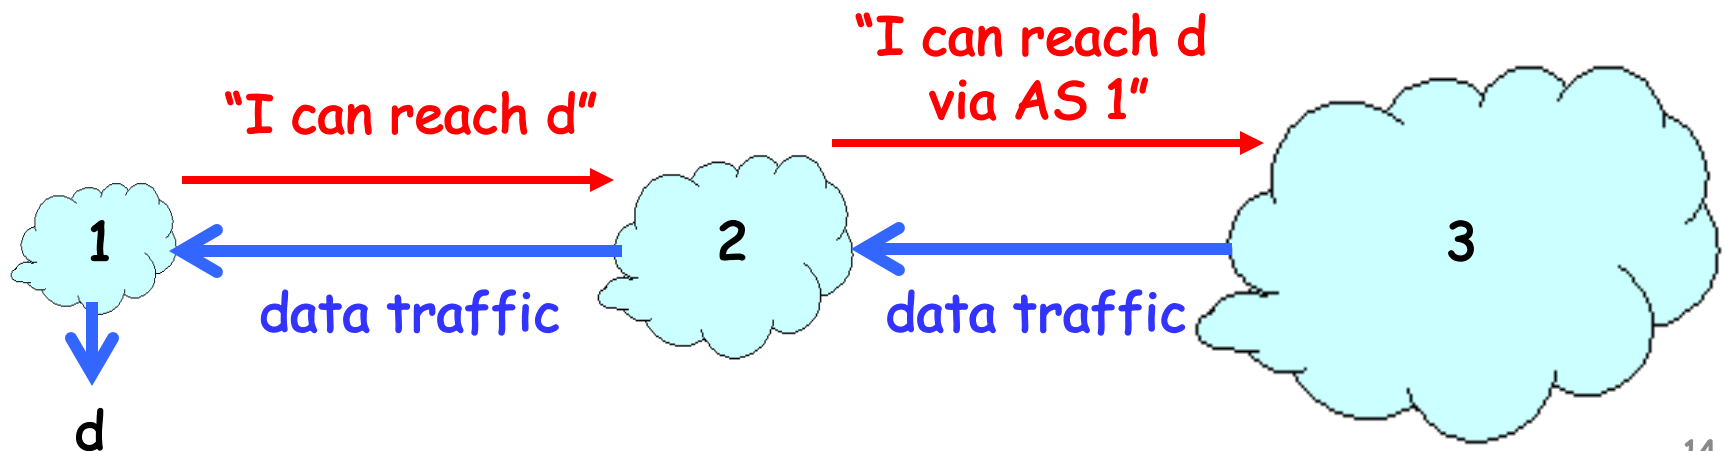
- **AS-level topology**

- Nodes are Autonomous Systems (ASes)
- Edges are links and business relationships



Review: Border Gateway Protocol

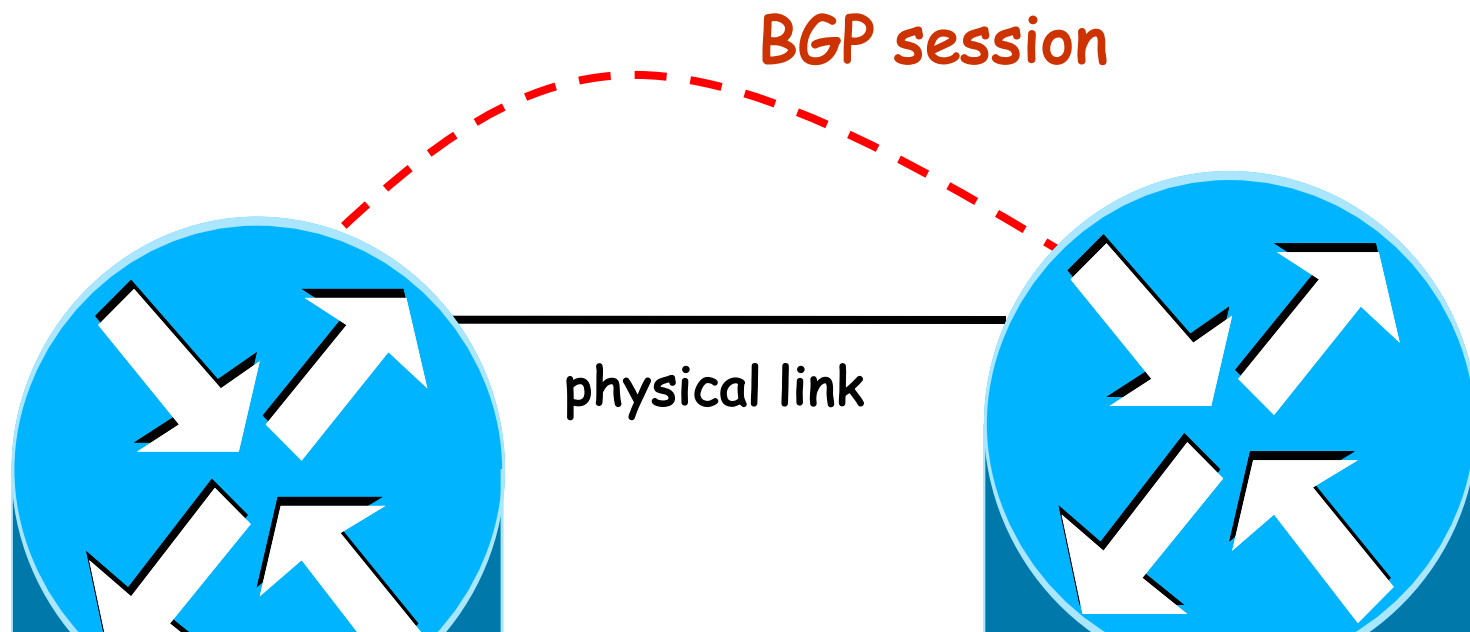
- **ASes exchange reachability information**
 - Destination: Block of addresses (an "IP prefix")
 - AS path: Sequence of ASes along the path
- **Policies configured by network operators**
 - Path selection: Which of the paths to use?
 - Path export: Which neighbors to tell?



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, adding packets
 - Changing, filtering, or replaying BGP routes
- **Availability**
 - Resetting the session or congesting the link
 - Disrupting communication and overloading 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
 - Can give BGP packets higher priority

Validity of routing information:
Origin authentication

IP Address Ownership, Hijacking

- IP address block assignment
 - ICANN -> Regional Internet Registries -> ISPs
- 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




- **Blackhole:** data traffic is discarded
- **Snooping:** data traffic is inspected, then redirected
- **Impersonation:** traffic sent to bogus destinations

Hijacking is Hard to Debug

- The victim AS doesn't see the problem
 - Picks its own route, might not learn the bogus route
- May not cause loss of connectivity
 - Snooping, with minor 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

YouTube Outage on Feb 24, 2008

- YouTube (AS 36561): 208.65.152.0/22
- Pakistan Telecom (AS 17557)
 - Government order to block access to YouTube
 - Announces 208.65.153.0/24 to PCCW (AS 3491)
 - All packets to YouTube get dropped on the floor
- Mistakes were made
 - AS 17557: announce 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 in Asia
- 18:48:00: Several big trans-Pacific providers carrying route
- 18:49:30: Bogus route fully propagated
- 20:07:25: YouTube advertising /24 to attract traffic back
- 20:08:30: Many (but not all) providers are using valid route
- 20:18:43: YouTube announces 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: Internet back up


Another Example: Spammers

- Spammers sending spam
 - Form a (bidirectional) TCP connection to mail server
 - Send a bunch of spam e-mail, then disconnect
- 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
- How to evade detection
 - Hijack unused (i.e., unallocated) address block
 - 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
 - Attract sources that normally try to avoid AS 3715
 - Help AS 88 appear 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 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?




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
 - Directing all 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 traffic for signs of inconsistency



BGP Security Today


- Applying "best common practices"
 - 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 practices
 - 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 data packets follow the chosen route

Proposed Enhancements to BGP

Secure BGP



Origin Authentication + cryptographic signatures



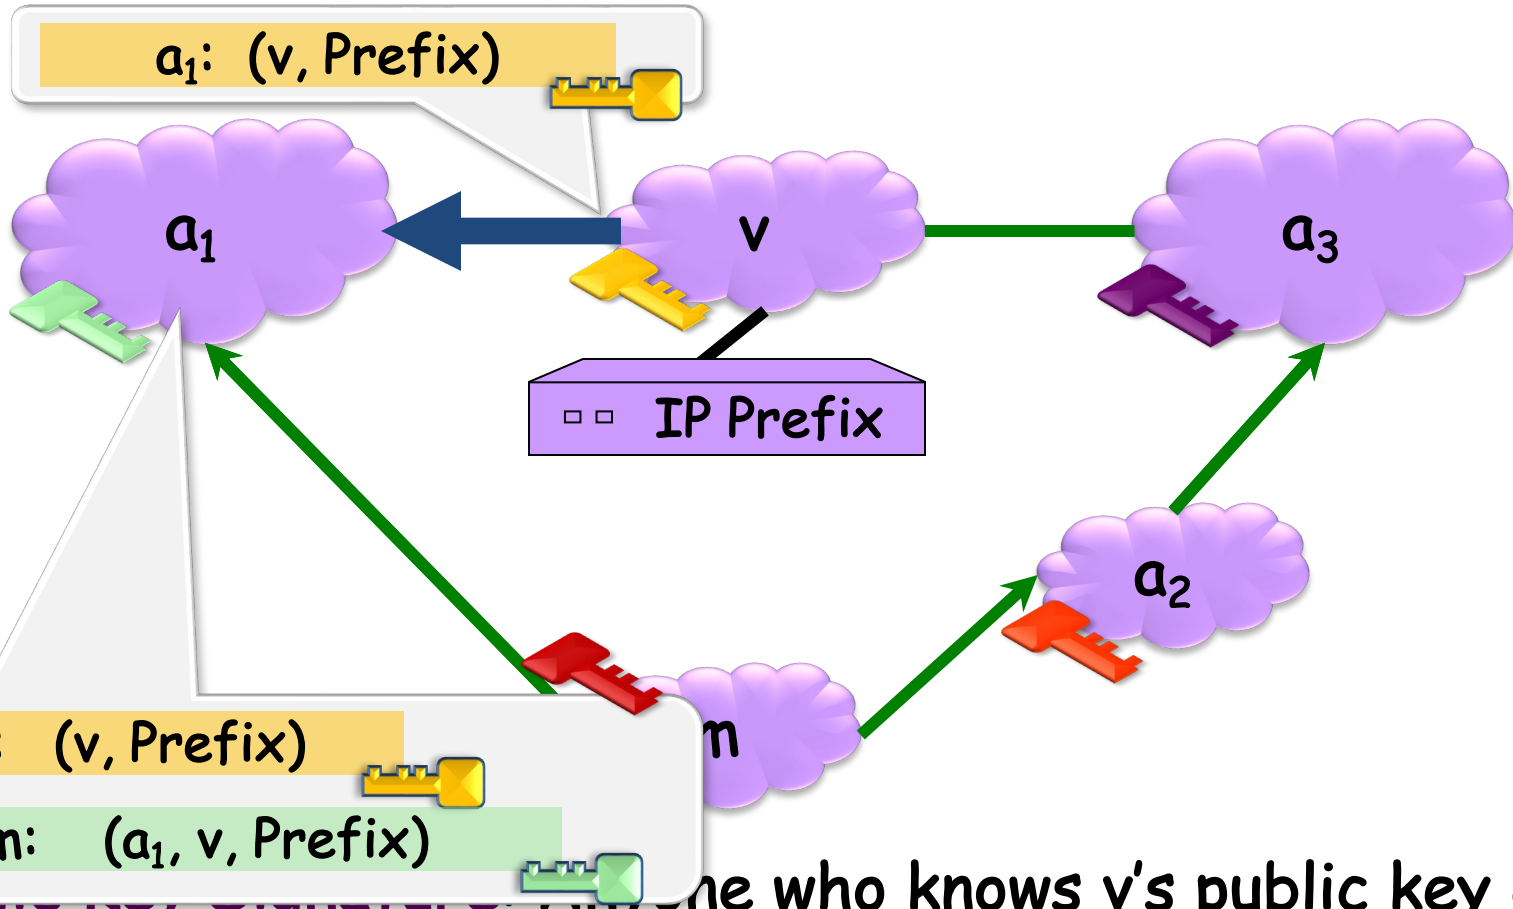
Public Key Signature: Anyone who knows v 's public key can verify that the message was sent by v .



Secure BGP



Origin Authentication + cryptographic signatures



...one who knows v 's public key can verify that the message was sent by v .



"Secure BGP"

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

S-BGP Deployment Challenges

- Complete, accurate registries of prefix "owner"
- 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

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 actors want the Internet to stay up

BGP is So Hard to Fix

- **Complex system**
 - Large, with around 40,000 ASes
 - Decentralized control among competitive ASes
- **Hard to reach agreement on the right solution**
 - S-BGP with PKI, registries, and crypto?
 - Who should be in charge of running PKI & 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

Conclusions

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