eecs489 Computer Networks

Lecture 19:
BGP Security

Security Goals for BGP

Secure message exchange between neighbors
- confidential BGP message exchange
- no denial of service

Validity of the routing information
- origin authentication
  - is the AP owned by the AS announcing it?
- AS path authentication
  - is the AS-Path the actual sequence of ASs traversed?
- AS path policy
  - does the AS path adhere to the routing policies of each AS?

Correspondence of the forwarding path
- does the traffic follow the advertised AS path?

BGP Security Topics

We'll look at an *ad hoc* collection of potential attacks on BGP and an *ad hoc* collection of ways to protect against them

Topics:
- BGP session security
- Origin authentication
- AS path validity
- Secure BGP
- Data forwarding vulnerability
- State of the BGP

TCP and BGP Session

BGP session runs over TCP
- neighboring routers create a TCP stream
- BGP messages sent over TCP
- makes BGP vulnerable to attacks on TCP

Primary types of attacks
- against confidentiality: eavesdropping
- against integrity: tampering
- against infrastructure availability: denial-of-service

Primary means of defense
- message authentication or encryption
- limiting access to physical path between routers
- defensive filtering to block unexpected packets

Attacks Against Confidentiality

Eavesdropping
- monitoring messages on BGP session by tapping the link(s) between the neighbors

Reveals sensitive information
- inference of business relationships
- analysis of network stability

Hard because
- difficult to tap link
  - often, eBGP session traverses just one link
  - and it may be hard to get access to the link
- encryption may obscure message contents
  - BGP neighbors may run BGP over IPSec

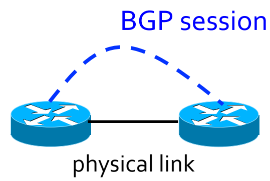

Attacking Message Integrity

Tampering
- man-in-the-middle tampers with the messages
- insert, delete, modify, or replay messages

Leads to incorrect BGP behavior
- delete: neighbor doesn't learn of new route
- insert/modify: neighbor learns bogus route

Hard because
- getting in-between the two routers is hard
- use of authentication (signatures) or encryption
- spoofing TCP packets the right way is not trivial
  - getting past source-address packet filters
  - generating the right TCP sequence number

Denial-of-Service Attacks (I)

Overload the link between the routers
- to cause packet loss and delay
- disrupting the performance of the BGP session

Relatively easy to do
- can send traffic between end hosts
- as long as the packets traverse the link
- (which you can figure out from traceroute)

Easy to defend
- give higher priority to BGP packets
- e.g., by putting packets in separate queue

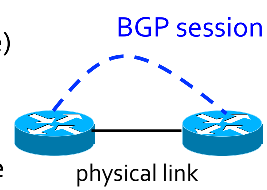

Denial-of-Service Attacks (II)

Third party sends bogus BGP/TCP packets
- FIN/RST to close the BGP session
- SYN flooding to overload the router

Leads to disruptions in BGP
- session reset, causing transient routing changes
- route-flapping, which may trigger flap damping

Hard because
- spoofing TCP packets the right way is not trivial
  - difficult to send FIN/RST with the right TCP header
- packet filters may block SYN flooding
  - filter packets to BGP port from unexpected sources
  - or filter packets destined to router from unexpected sources

Exploiting the IP TTL Field

BGP routers are usually one hop apart
- to thwart an attacker, can check that the packets carrying the BGP message have not traveled far (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

IP Address Ownership and Hijacking

IP address-block assignment by
- Regional Internet Registries (ARIN, RIPE, APNIC)
- or Internet Service Providers

Proper origination of a prefix into BGP
- by the AS who owns the prefix
- or, by its upstream provider(s) on 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

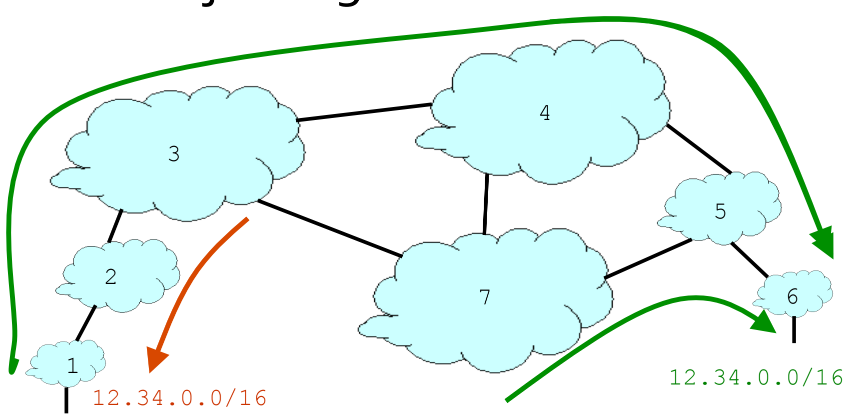

Consequences for the affected ASs
- blackhole: data traffic is discarded
- snooping: data traffic is inspected, and then redirected
- impersonation: data traffic is sent to bogus destinations

Hijacking is Hard to Detect

Legitimate origin AS doesn't see the problem
- picks its own route
- might not even learn of 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

How to diagnose prefix hijacking?
- analyze updates from many vantage points on the Internet
- launch traceroute from many vantage points
- requires access to BGP routers or hosts across the Internet

Sub-Prefix Hijacking

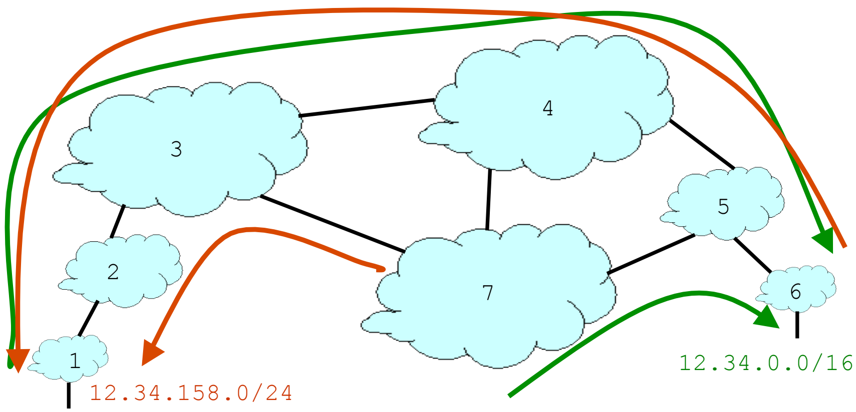

Originating a more-specific prefix
- traffic follows the longest matching prefix
- every AS picks the bogus route for that prefix

How to Hijack a Prefix

The hijacking AS has
- a router with eBGP session(s)
- that is configured to originate the prefix

Ways to get access to a router:
- network operator makes configuration mistake,
- disgruntled operator launches an attack, or
- outsider breaks in to the router and reconfigures

Getting other ASs to believe bogus route
- neighboring ASs do not filter routes, i.e., by allowing only expected prefixes
- specifying filters on peering links is hard

February 24, 2008 YouTube Outage

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 its provider 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

Another Example: Spammers

Spammers sending spam
- form a (bidirectional) TCP connection to a mail server
- send a bunch of spam e-mail
- disconnect

Real IP addresses are relatively easy to trace back

Could hijack someone else'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

Bogus AS Paths

Remove ASs 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

Hard to tell that an AS path is invalid
- maybe AS 88 does connect to AS 701 directly

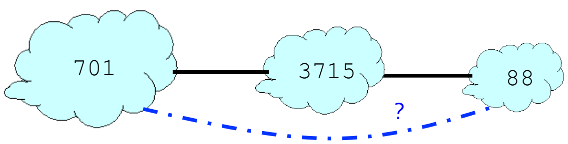

Bogus AS Paths

Add ASs to the path
- e.g., turn "701 88" into "701 3715 88"

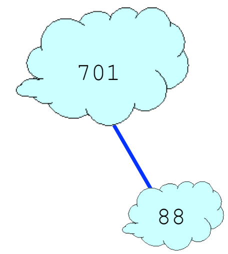

Motivations
- trigger loop detection in AS 3715
  - denial-of-service attack on AS 3715
  - or, blocking unwanted traffic coming from AS 3715!
- make your AS (701) looks like is has richer connectivity

Hard to tell that an AS path is invalid
- 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 of a bogus route by adding the legitimate AS to the end

Hard to tell that the AS path is bogus
- even if other ASs 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

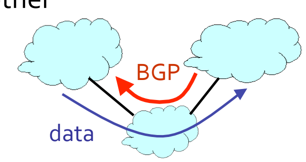

leading to dire consequences
- directing all Internet traffic through customer, who does not have enough resources to handle so much traffic

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
- analyze BGP updates, or data traffic, for signs of inconsistency

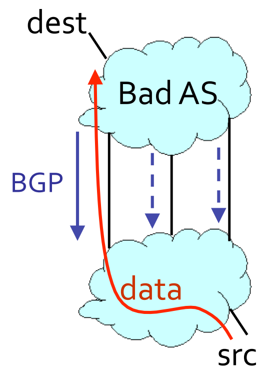

BGP Security Today

Applying best common practices (BCPs)
- securing the session (authentication, encryption)
- filtering routes by prefix and AS path
- filtering packets 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

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

Security provided by S-BGP:
- AS-Path indicates the order ASs were traversed
- no intermediate ASs were added or removed

S-BGP Deployment Challenges

Requires complete, accurate registries
- e.g., of prefix ownership

Requires public-key infrastructure
- to know the public key for any given AS

Requires expensive 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
- impossible to have a "flag day" to deploy S-BGP

Incrementally Deployable Schemes

Monitoring BGP update messages
- use past history as an implicit registry
- e.g., AS that announces each address block
- e.g., AS-level edges and paths

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/

Soft response to suspicious routes
- prefer routes that agree with the past
- delay adoption of unfamiliar routes when possible
  - some (e.g., misconfiguration) will disappear on their own

Routing vs. Forwarding

Routing:
- BGP is a routing protocol
- BGP security concerns validity of routing messages
  - i.e., did the BGP message follow the sequence of ASs listed in the AS-path attribute

Forwarding:
- routers forward data packets
- supposedly along the path chosen by the routing protocol
  - but what ensures that this is true?

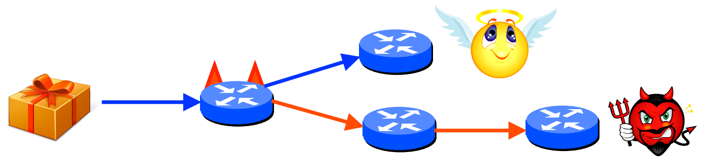

Forwarding Attacks (I)

Drop packets in the data plane
- while still sending the routing announcements

Easier to evade detection
- especially if you only drop some packets
- e.g., BitTorrent or Skype traffic

Even easier if you just slow down some traffic
- how different are normal congestion and an attack (or provider throttling)?
- especially if you let ping/traceroute packets through?

Forwarding Attacks (II)

Direct packets to a different path
- that disagrees with the routing announcements

Direct packets to a different destination
- e.g., one controlled by an adversary

Motivations:
- to impersonate the legitimate destination (e.g., to perform identity theft, or promulgate false information)
- to snoop on traffic before forwarding it to the real destination

How to detect?
- traceroute? longer than usual delays?
- end-to-end checks, like site certificate or encryption?

Forwarding Attacks are Hard

Adversary must control a router along the path
- so that the traffic flows through it

How to get control of 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 steer customers towards its voice products (net neutrality?)

BGP is 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

Many smaller examples
- blackholing a single destination prefix
- hijacking unallocated addresses to send spam

Why isn't it an even bigger problem?
- really, most big outages are configuration errors
- most bad guys want the Internet to stay up
  - so they can send unwanted traffic (e.g., spam, identity theft, denial-of-service attacks, port scans, etc.)

BGP is Hard to Fix

Complex system
- large, with around 50,000 ASs
- decentralized control among competitive ASs
- 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 the PKI and registries?
- worry about forwarding vulnerability or just routing?

Hard to deploy the solution once you pick it
- hard enough to get ASs to apply route filters
- now you want them to upgrade to a new protocol, all at the exact same moment, without incremental deployment plan? Ha!

Conclusions

Internet protocols were designed based on trust

Border Gateway Protocol is very vulnerable
- twigs and twine that hold 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 BGP
- anomaly detection schemes, with automated response
- broader focus on forwarding availability