BGP Problems

BGP has numerous problems

Problems

- Reachability
- Security
- Convergence
- Performance
- Anomalies
- Relevance

Unlike normal routing, policy-based routing does NOT guarantee reachability even if the graph is connected

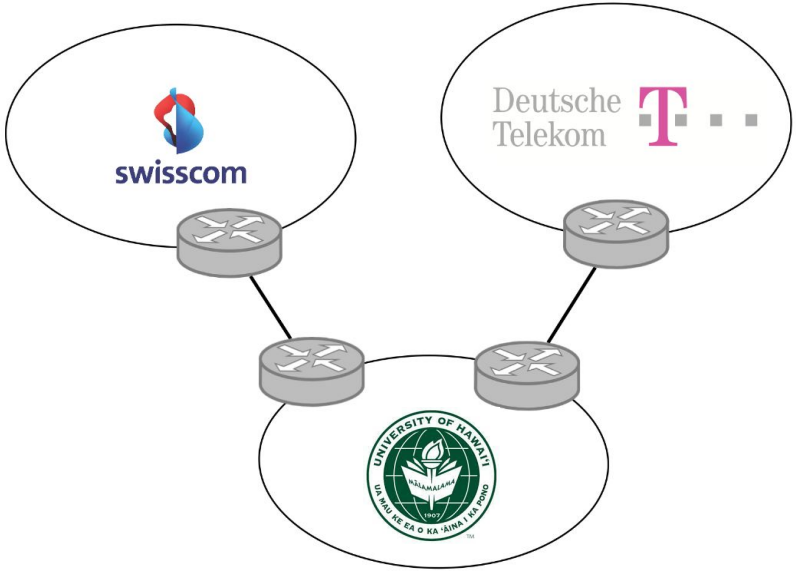

Because of policies,

Swisscom cannot reach DT

even if the graph is connected

BGP attacks on underlying TCP

- BGP session runs over TCP
  - TCP connection between neighboring routers
  - BGP messages sent over TCP connection
  - Makes BGP vulnerable to attacks on TCP
- Main kinds of attacks
  - Against confidentiality: eavesdropping
  - Against integrity: tampering
  - Against performance: denial-of-service
- Main defenses
  - Message authentication or encryption
  - Limiting access to physical path between routers
  - Defensive filtering to block unexpected packets

BGP confidentiality attacks

- Eavesdropping
  - Monitoring the messages on the BGP session
  - … by tapping the link(s) between the neighbors
- Reveals sensitive information
  - Inference of business relationships
  - Analysis of network stability
- Reasons why it may be hard
  - Challenging to tap the link
    - Often, eBGP session traverses just one link
    - … and may be hard to get access to tap it
  - Encryption may obscure message contents
    - BGP neighbors may run BGP over IPSec

BGP message integrity attacks

- 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 the new route
  - Insert/modify: neighbor learns bogus route
- Reasons why it may be hard
  - Getting in-between the two routers is hard
  - Use of authentication (signatures) or encryption
  - Spoofing TCP packets the right way is hard
    - Getting past source-address packet filters
    - Generating the right TCP sequence number

BGP denial of service attacks

- 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

BGP denial of service attacks - part 2

- Third party sends bogus TCP packets
  - FIN/RST to close the 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
- Reasons why it may be hard
  - Spoofing TCP packets the right way is hard
    - Difficult to send FIN/RST with the right TCP header
  - Packet filters may block the SYN flooding
    - Filter packets to BGP port from unexpected source
    - ... or destined to router from unexpected source

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

Many security considerations are absent from the BGP specification

ASes can advertise any prefixes

even if they don't own them!

ASes can arbitrarily modify route content

e.g., change the content of the AS-PATH

ASes can forward traffic along different paths

than the advertised one

BGP's (terrible) security

#1 BGP does not validate the origin of advertisements

#2 BGP does not validate the content of advertisements

BGP's (terrible) security

#1 BGP does not validate the origin of advertisements

#2 BGP does not validate the content of advertisements

IP Address Ownership / 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

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

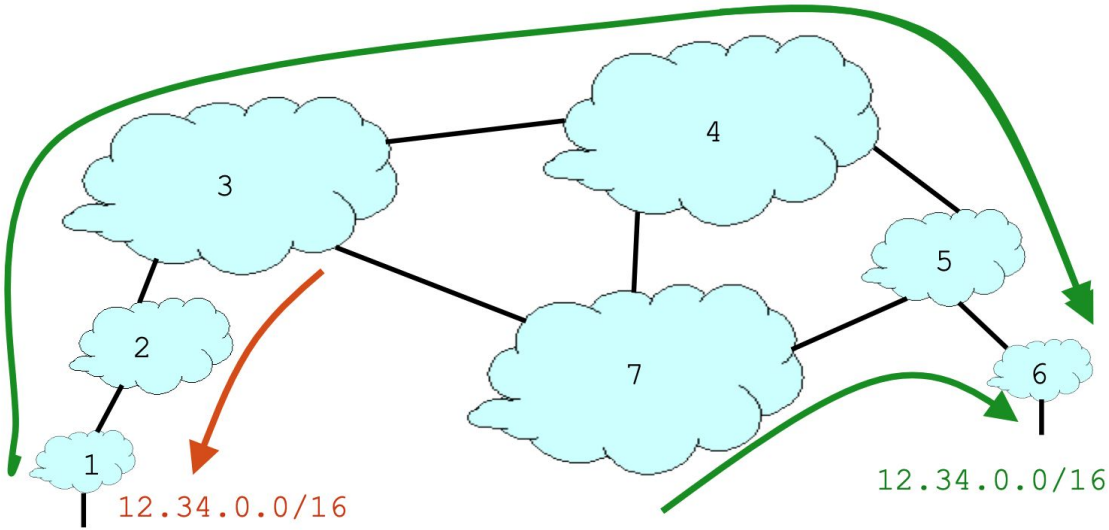

Hijacking is not easy to debug

- 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

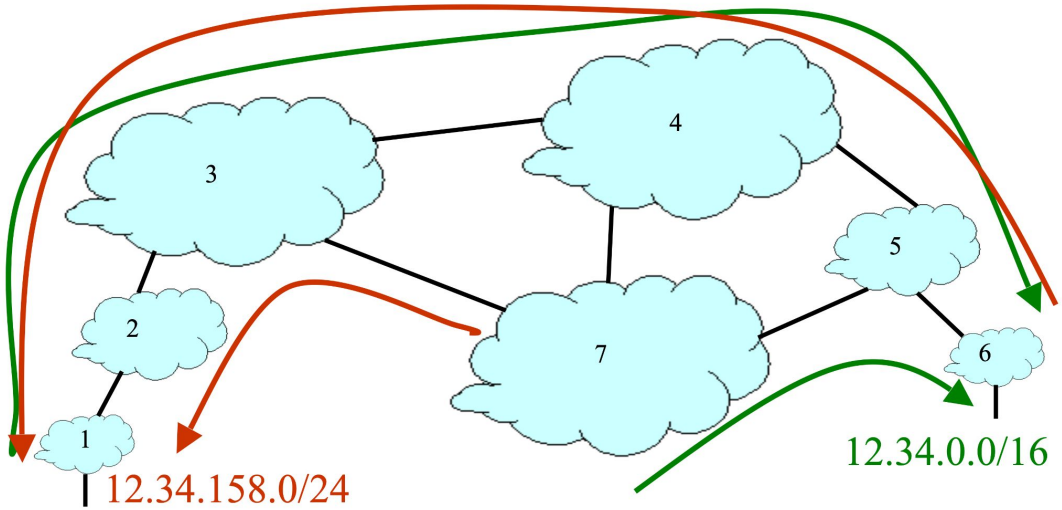

- Originating a more-specific prefix
  - **Every** AS picks the bogus route for that prefix
  - Traffic follows the longest matching prefix

Hijacking How-To

- 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: Feb 24, 2008

- YouTube (AS 36561) owns 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) who shares it with the world
  - All packets to YouTube get dropped on the floor
- 20 minutes later AS36561 (YouTube) starts announcing 208.65.153.0/24.
  - With two identical prefixes in the routing system, BGP policy rules, such as preferring the shortest AS path, determine which route is chosen. This means that AS17557 (Pakistan Telecom) continues to attract some of YouTube's traffic.
- 11 minutes later AS36561 (YouTube) starts announcing 208.65.153.128/25 and 208.65.153.0/25. Because of the longest prefix match rule, every router that receives these announcements will send the traffic to YouTube.
- 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

https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

Another Example: SPAM

- 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
- Profit $$$

BGP's (terrible) security

#1 BGP does not validate the origin of advertisements

#2 BGP does not validate the content of advertisements

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 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

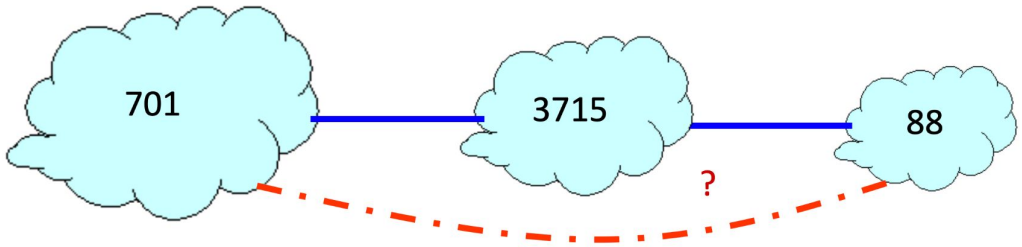

Bogus AS paths

- 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?

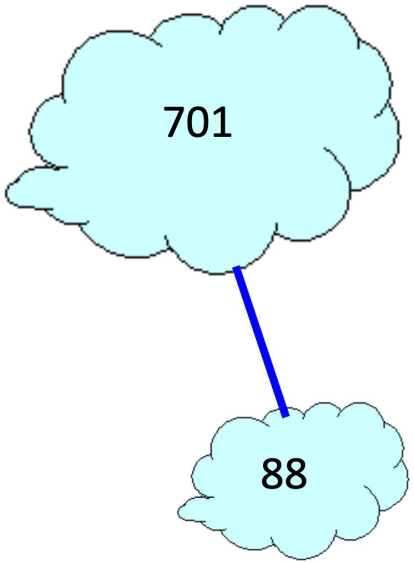

Bogus AS paths

- 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

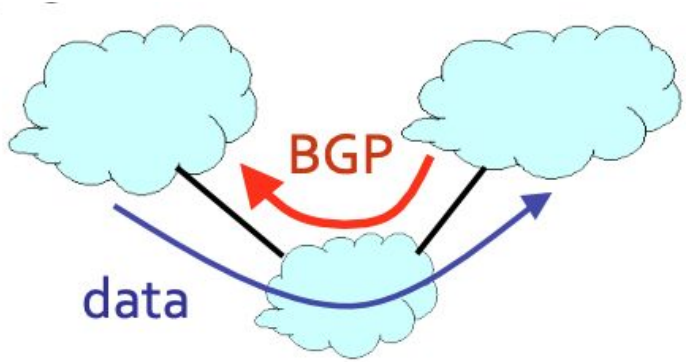

S-BGP: Secure BGP

- 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
- S-BGP can validate
  - AS path indicates the order ASes were traversed – No intermediate ASes were added or removed

S-BGP 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