Interdomain Routing Security

COS 461: Computer Networks

April 22, 2008

Sharon Goldberg

Mostly based on slides by Jennifer Rexford with some changes.

Goals of Today's Lectures

- BGP security vulnerabilities
  - TCP sessions
  - Prefix ownership
  - AS-path attribute
- Improving BGP security
  - Protective filtering
  - Security Enhancements to of BGP
  - Anomaly-detection schemes
- Data-plane attacks
- Difficulty in upgrading BGP

Security Goals for BGP

- Secure message exchange between neighbors
  - Integrity of BGP message exchange
  - No denial of service
- Validity of the routing information
  - Origin authentication
    - Is the prefix owned by the AS announcing it?
  - AS path authentication
    - Is AS path the sequence of ASes the BGP update traversed?
  - AS path policy
    - Does the AS path adhere to the routing policies of each AS?
- Correspondence to the data path
  - Does the traffic follow the advertised AS path?
  - Is it actually arriving at the destination?

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
- Main kinds of attacks
  - 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

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 the new route
  - Insert/modify: neighbor learns bogus route
- Reasons why it may be hard
  - Getting in-between the two routers is hard
  - Spoofing TCP packets the right way is hard
    - Generating the right TCP sequence number
  - Not feasible if (cryptographic) message authentication is used.

Denial-of-Service Attacks, Part 1

- 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

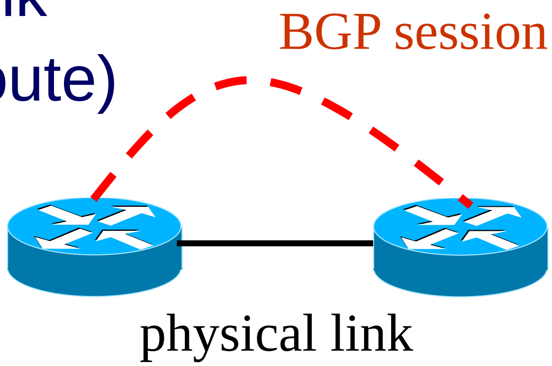

Denial-of-Service Attacks, Part 2

- 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 (port, seq #'s)
  - 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

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

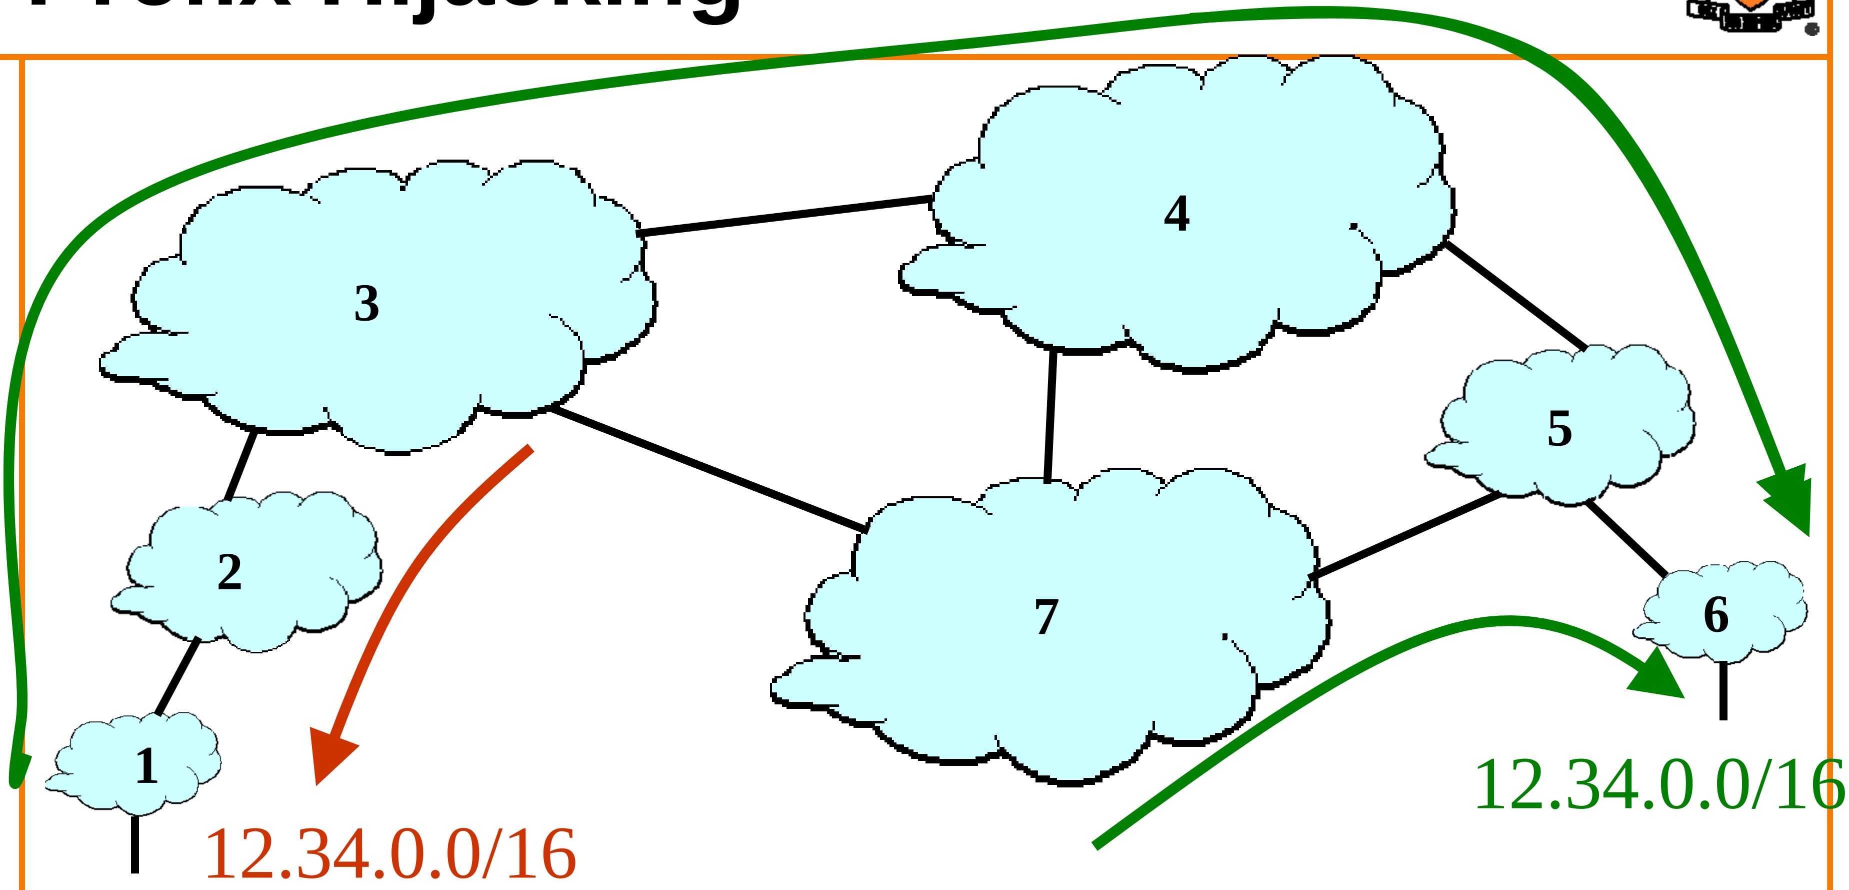

- Consequences for the affected ASes
  - Blackhole: data traffic is discarded
  - Snooping: data traffic is inspected, and then redirected
  - Impersonation: data traffic is sent to bogus destinations

Hijacking is Hard to Debug

- Real origin 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

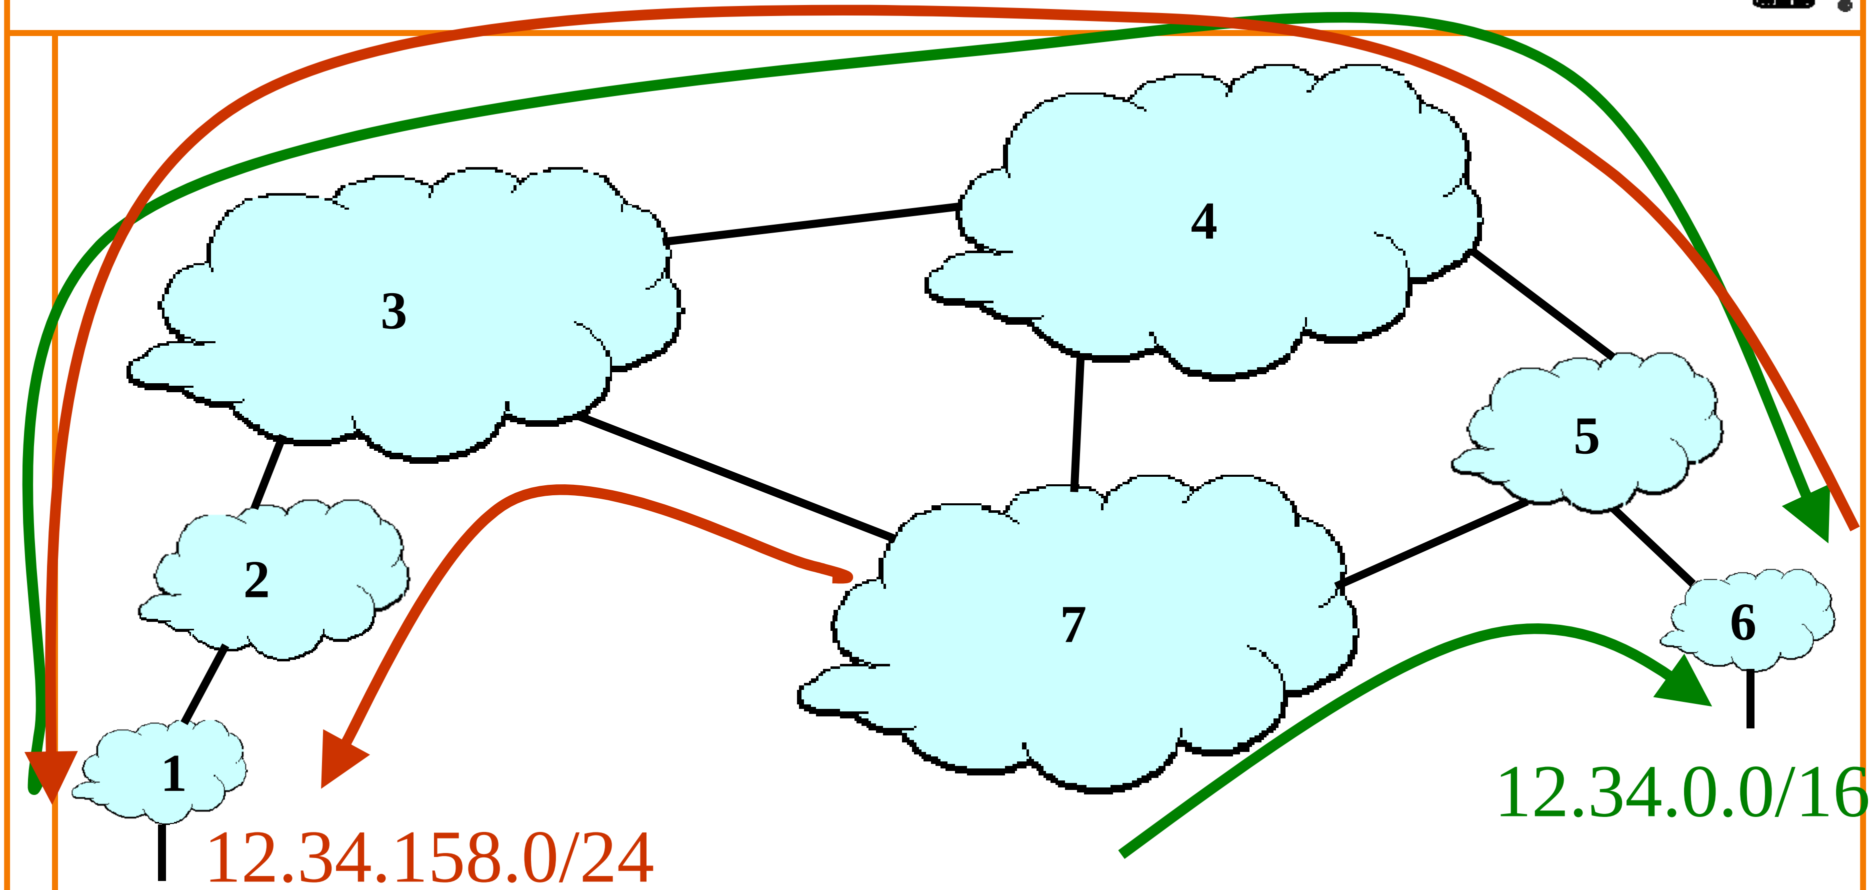

- Originating a more-specific prefix
  - Every AS picks the bogus route for that prefix
  - Traffic follows the longest matching prefix

How to Hijack a Prefix

- The hijacking AS has
  - –Router with eBGP 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 not filtering the routes
  - –… e.g., by allowing only expected prefixes
  - –But, specifying filters on *peering* links is hard

The February 24 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 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

BGP AS Path

Bogus AS Paths

- Path shortening - 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

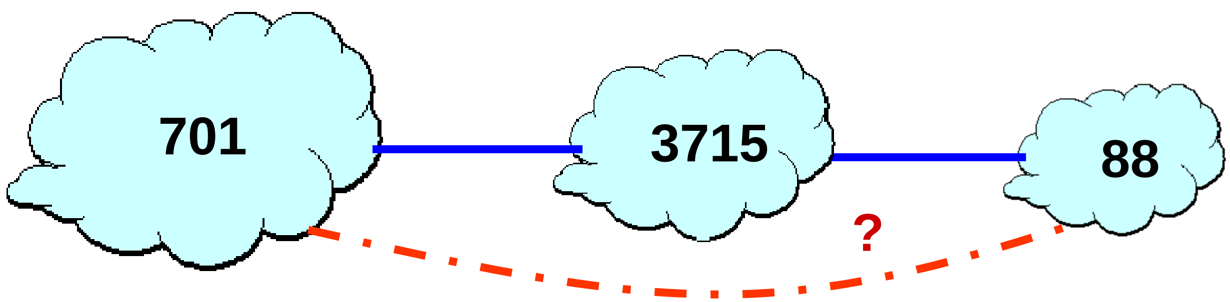

Bogus AS Paths

- 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?

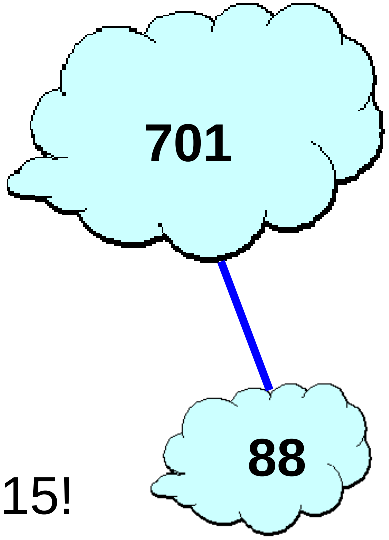

Bogus AS Paths

- 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

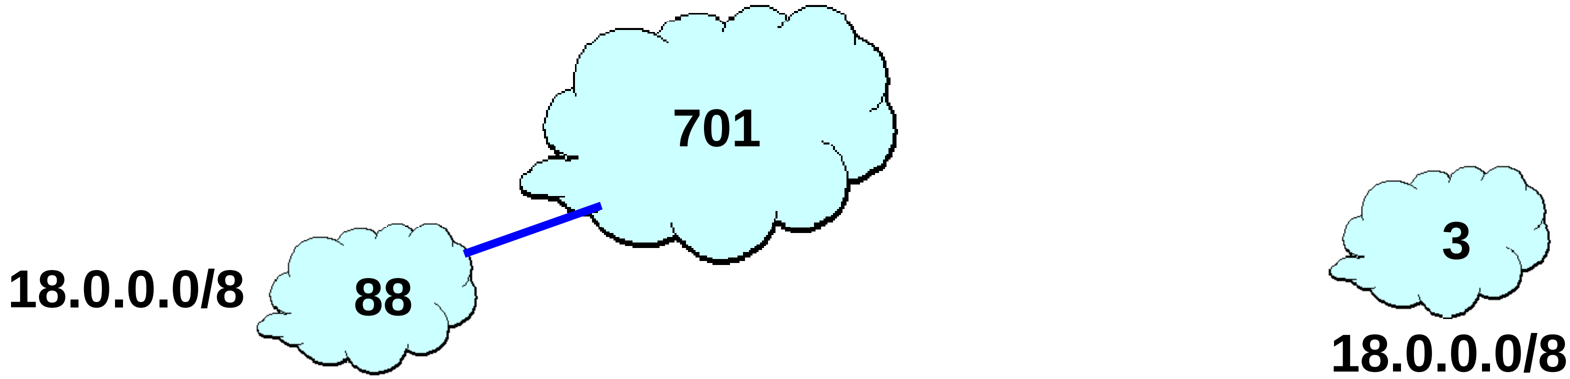

Invalid Paths

- 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

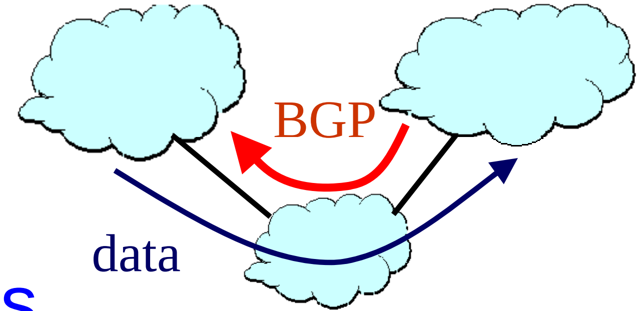

- … leading the dire consequences
  - Directing all Internet traffic through customer
- Main defense
  - Provider filters routes based on business relationships, prefixes and AS path

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

The BGP Man-In-The-Middle Attack

See the pdf from Pilosov and Kapela

ASN 200 announces 10.10.220.0/22 to its providers AS20 and AS30

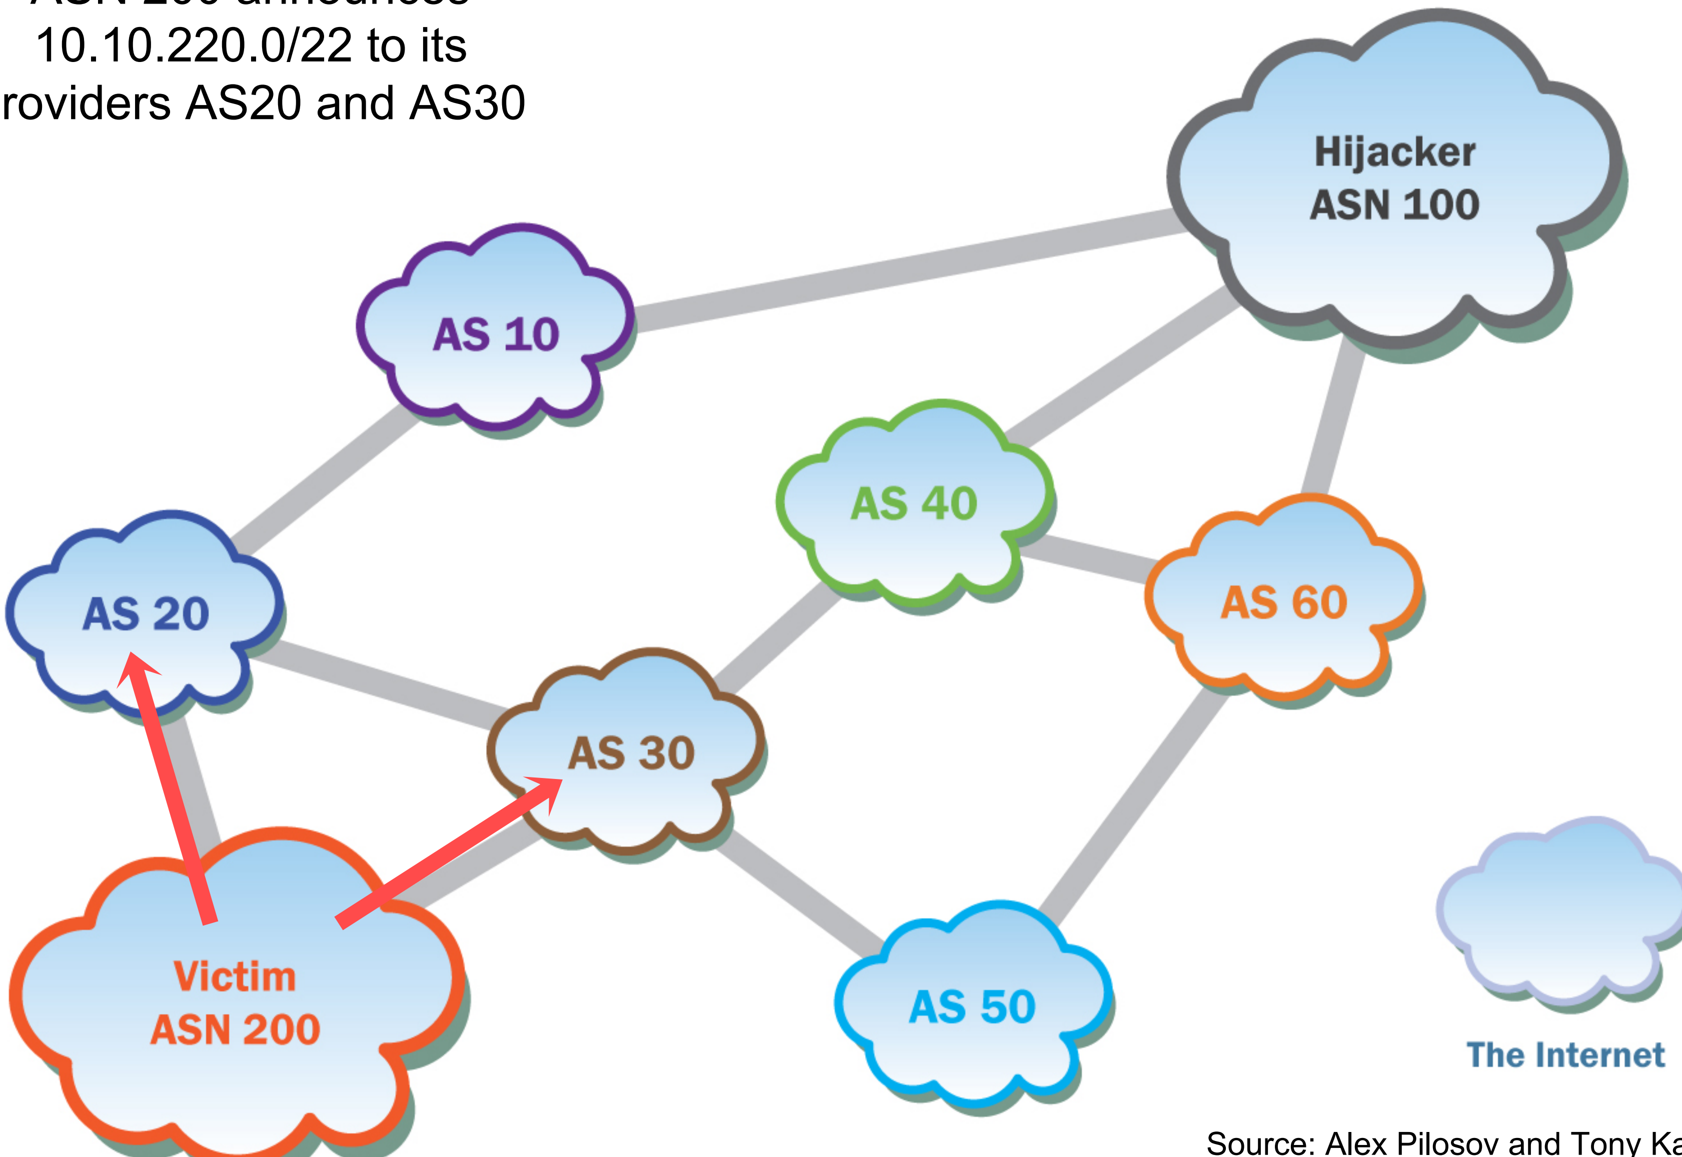

Source: Alex Pilosov and Tony Kapela

Announcement propagates
to the global Internet

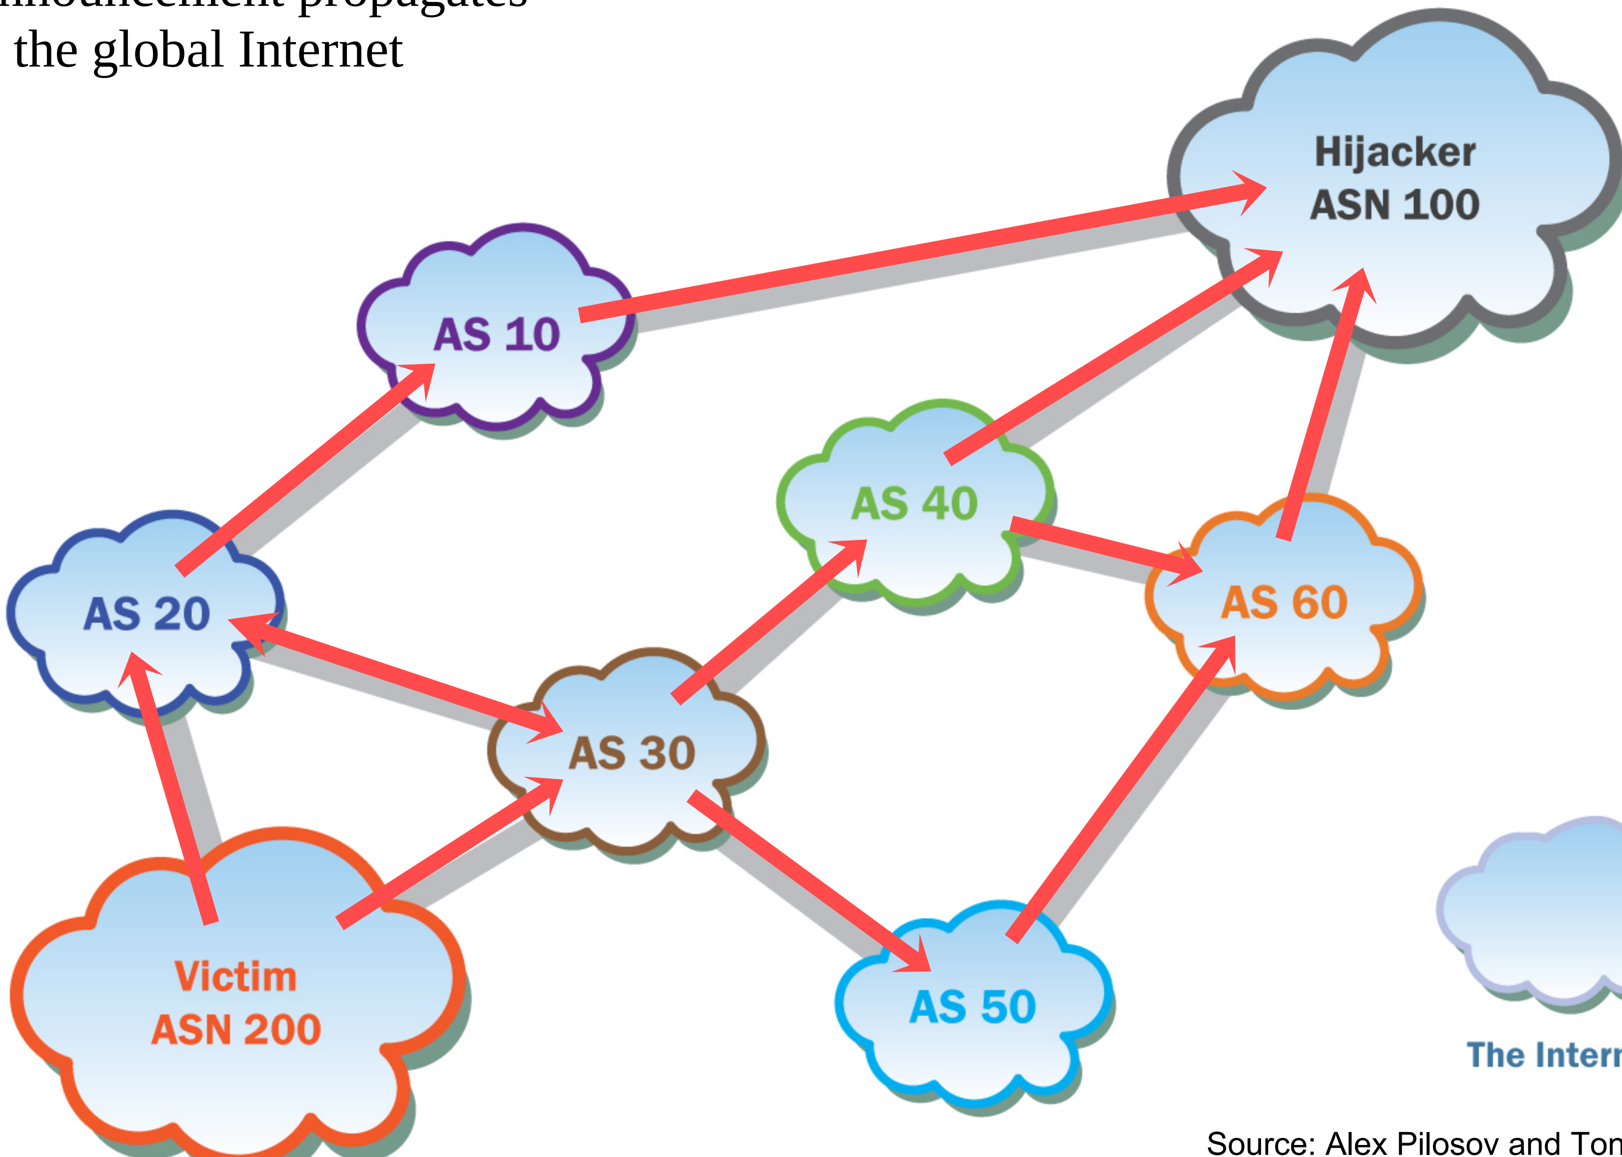

Source: Alex Pilosov and Tony Kapela

Every AS picks its “best” route to 10.10.220.0/22

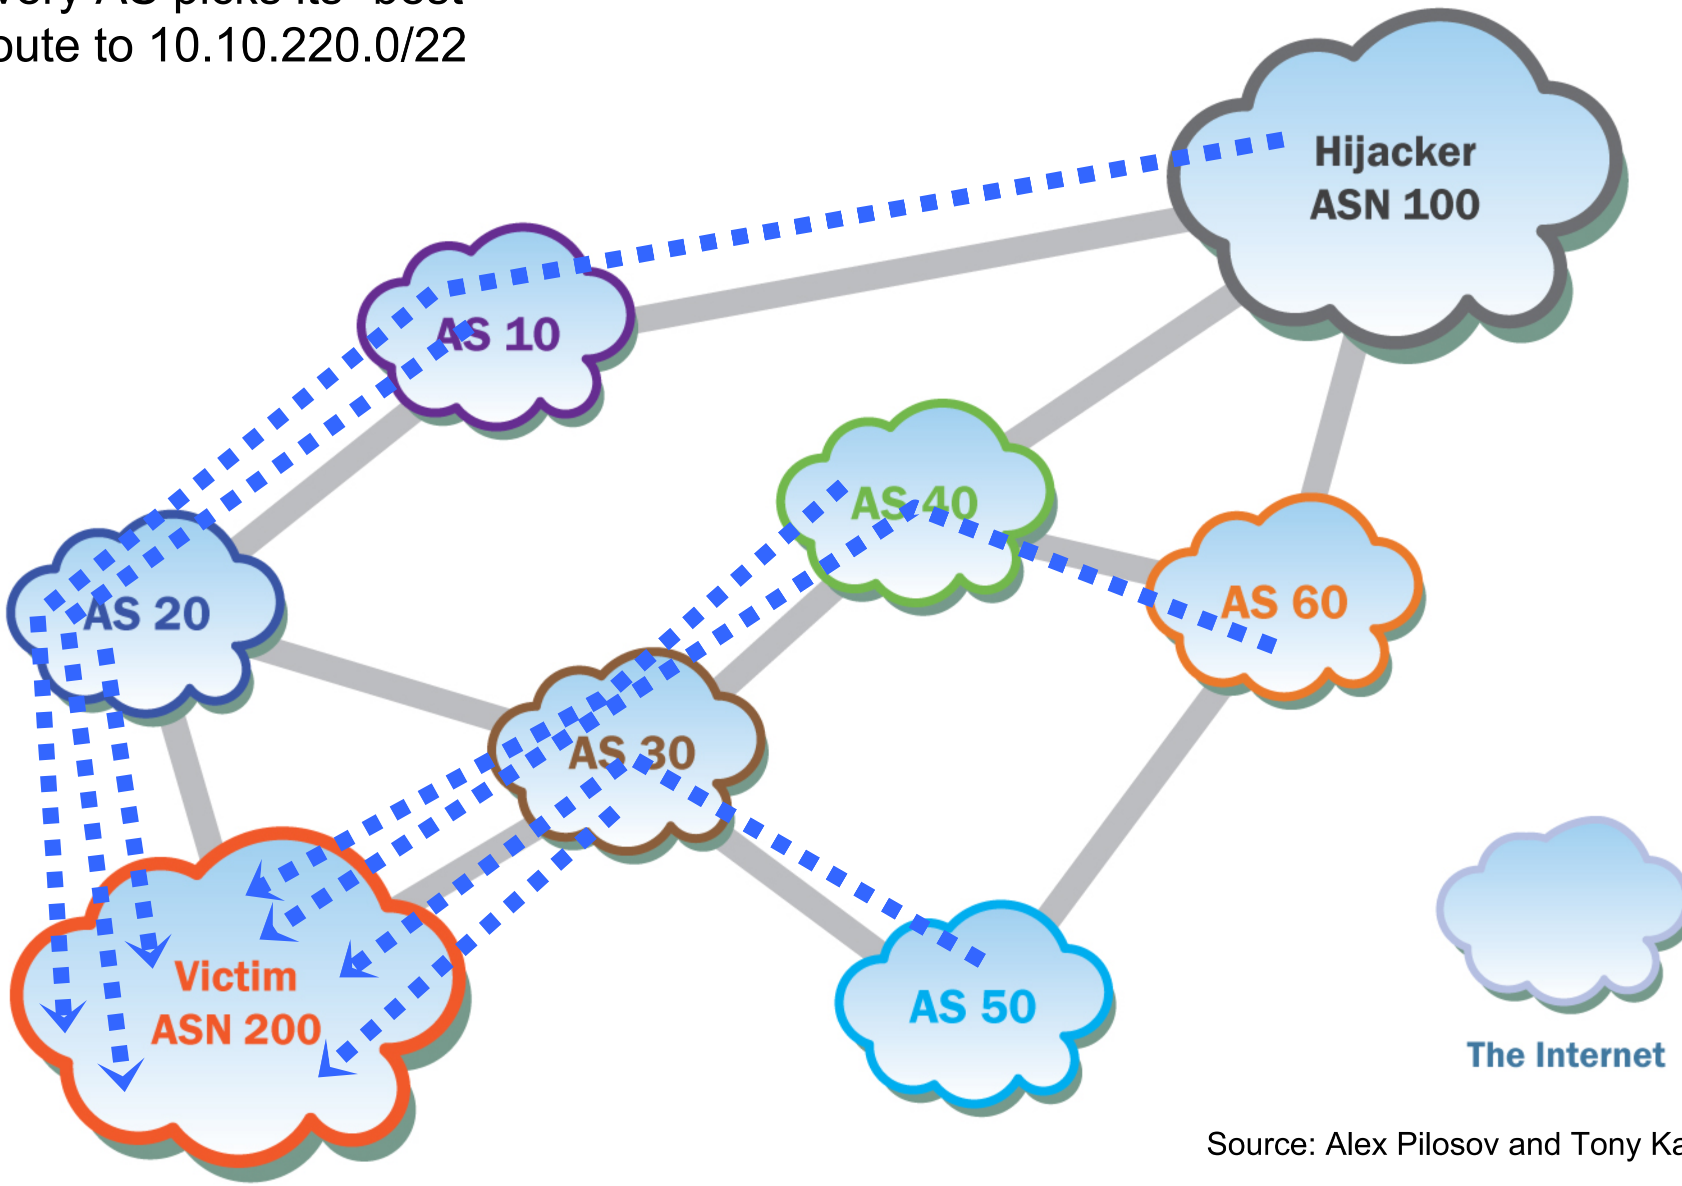

Source: Alex Pilosov and Tony Kapela

Hijacker prepares to attack

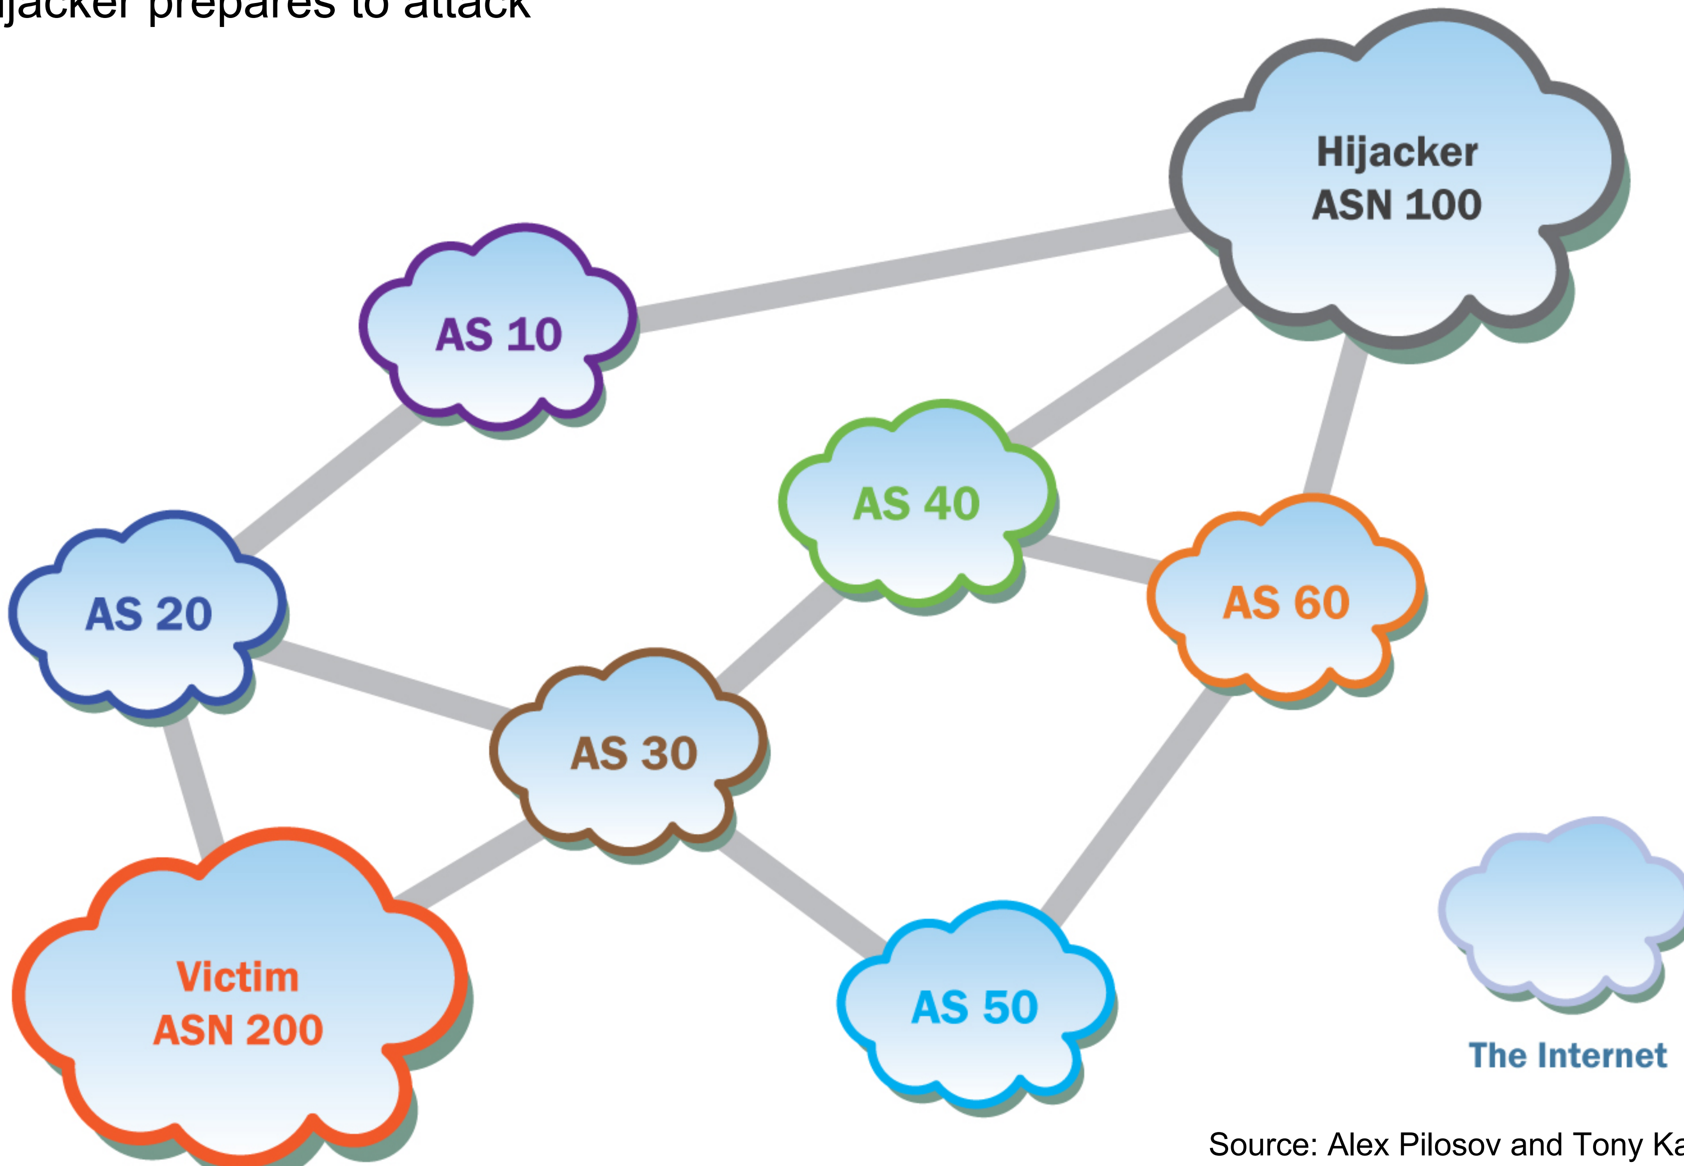

Source: Alex Pilosov and Tony Kapela

ASN 100 wants to preserve its valid route to 10.10.200.0/22 via AS10

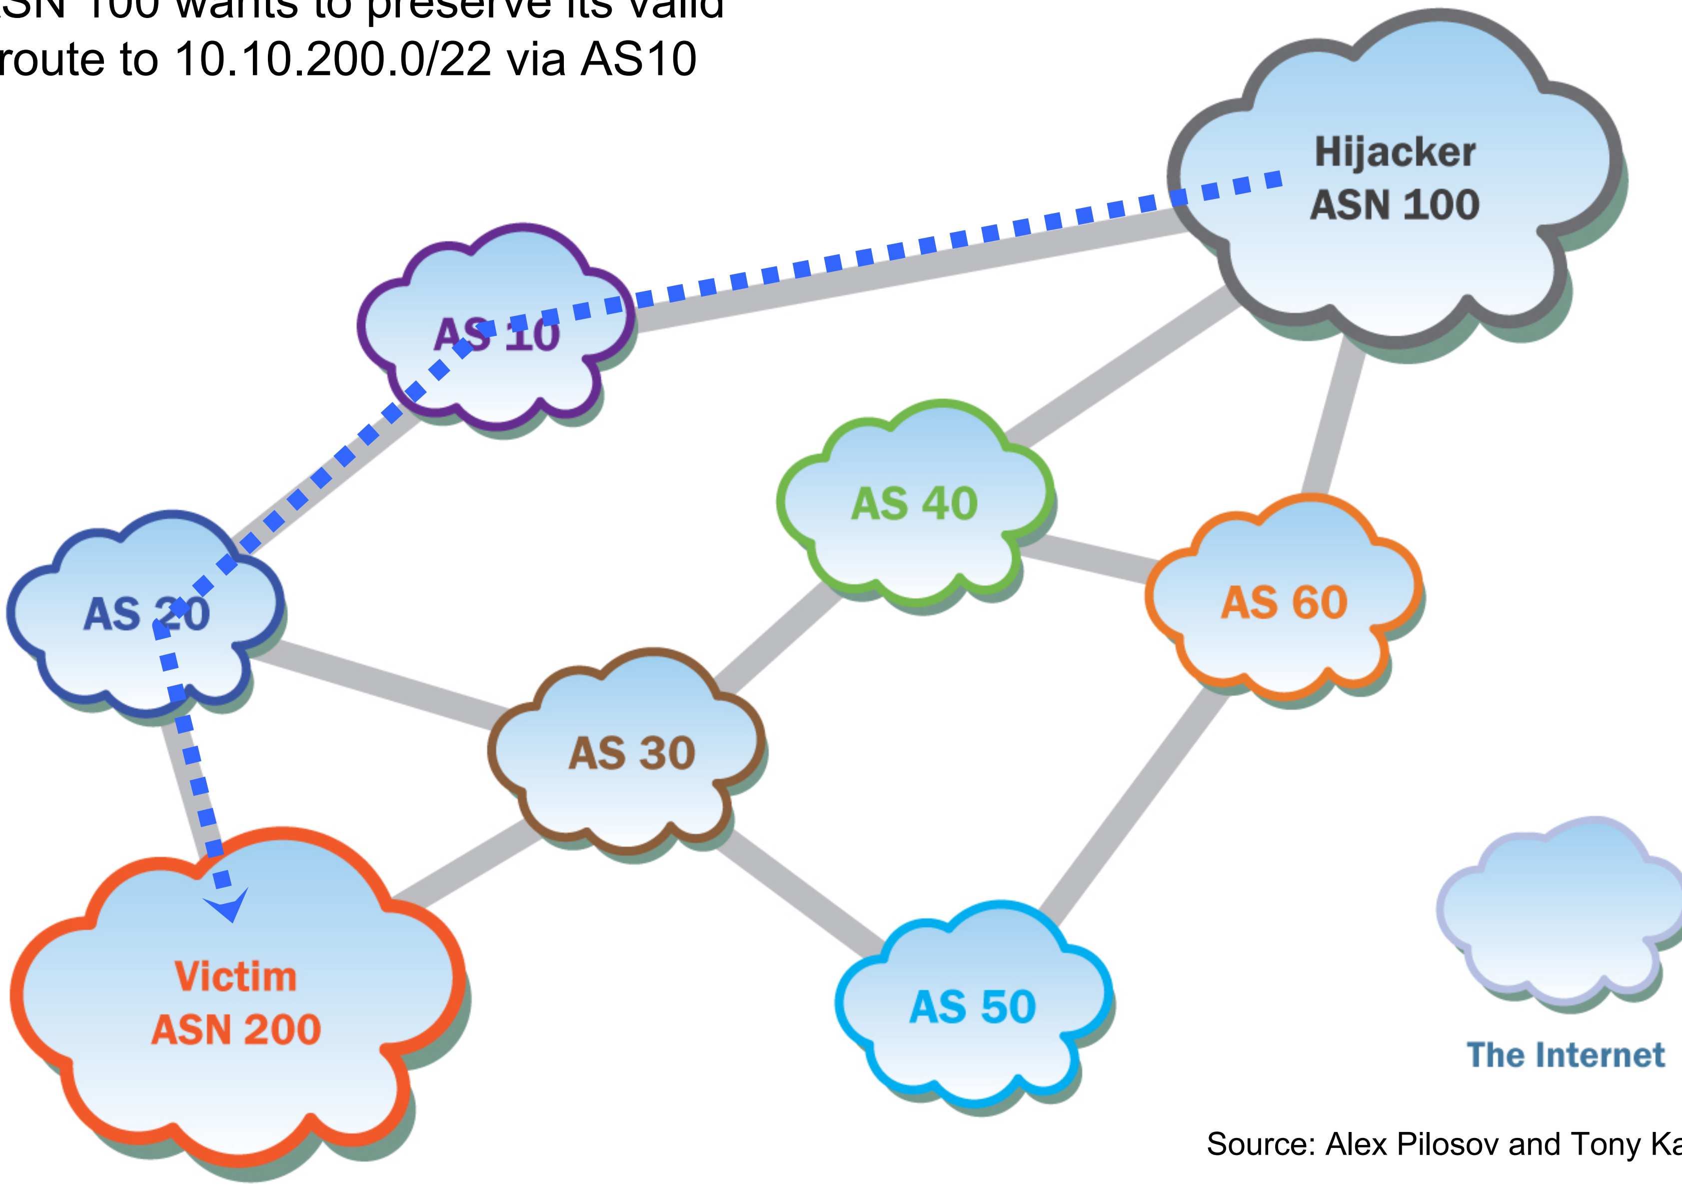

Source: Alex Pilosov and Tony Kapela

To preserve this path, ASN 100 must prepend AS 10, 20, and 200 to the ASPATH of any hijack announcement

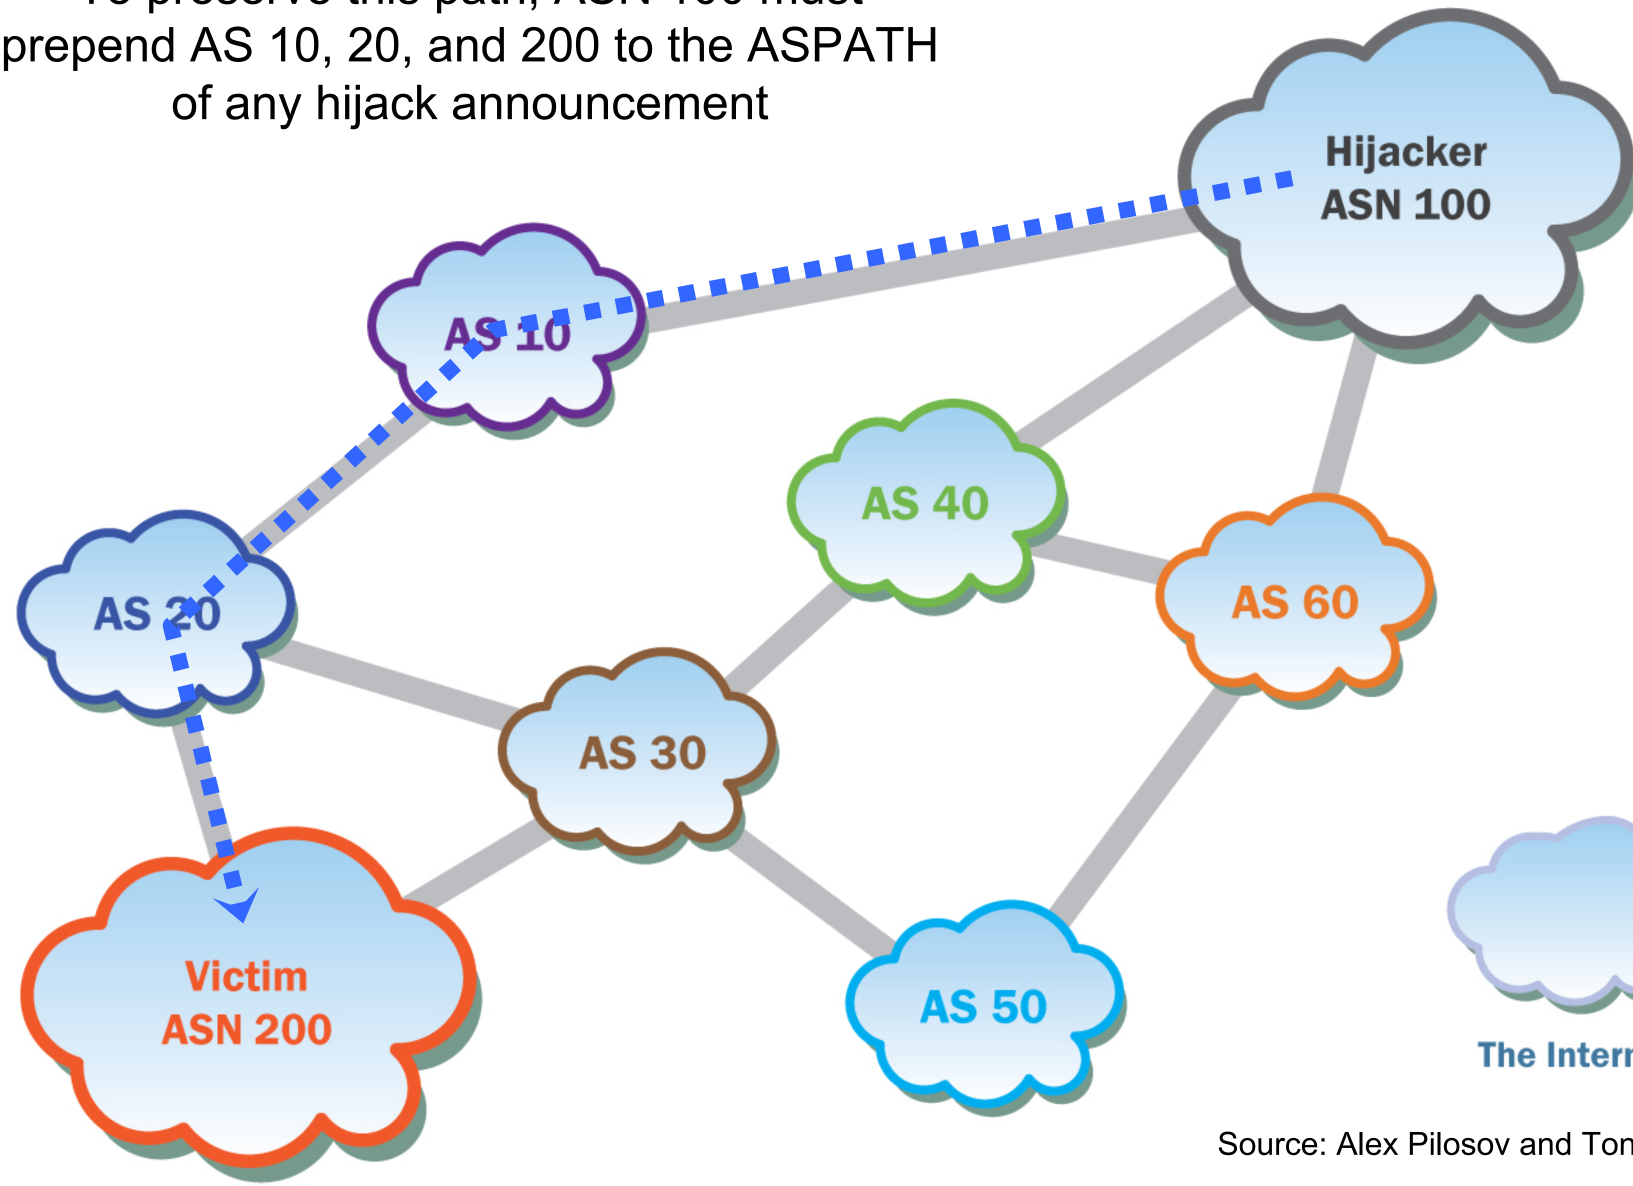

Source: Alex Pilosov and Tony Kapela

AS 100 announces the more-specific prefixes 10.10.220.0/23 and 10.10.222.0/23, prepending AS 10, 20 and 200 to the ASPATH.

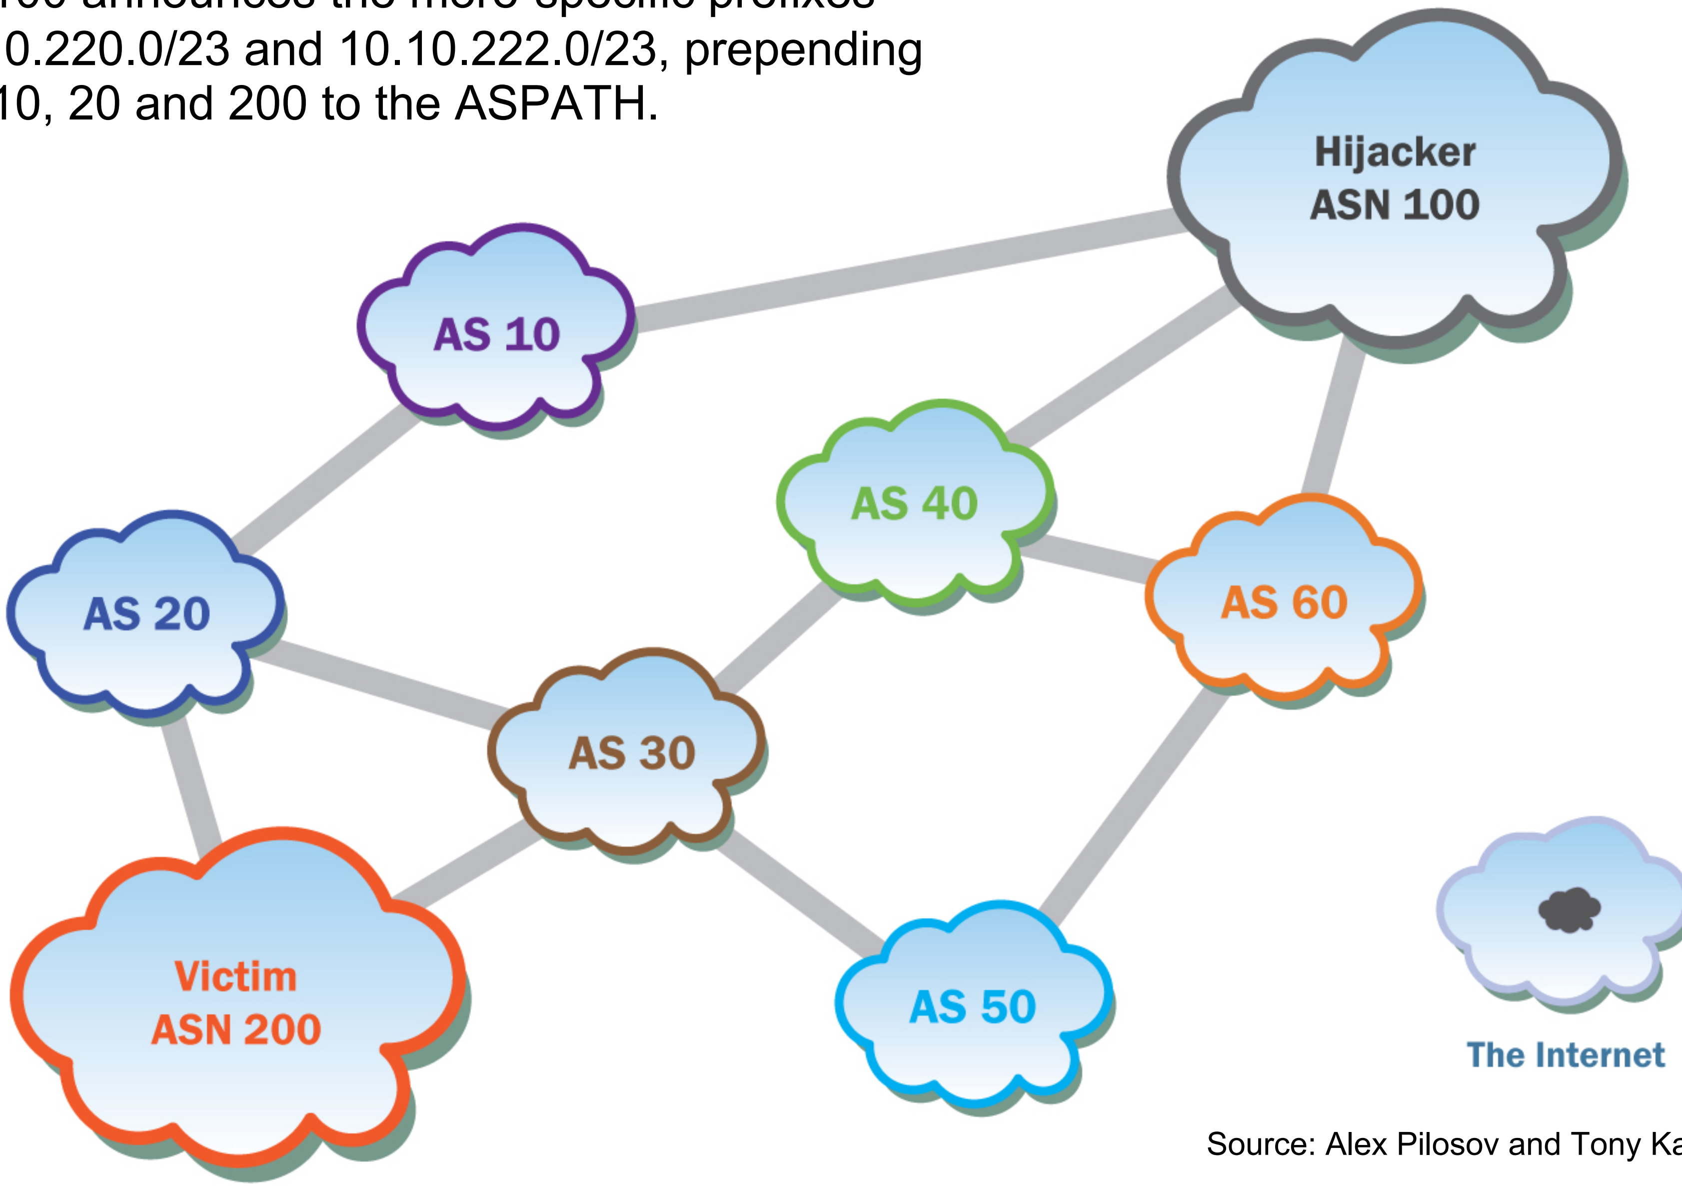

Source: Alex Pilosov and Tony Kapela

Most global traffic for 10.10.220.0/22 now goes to ASN 100, who forwards it onto AS 10 after examining, copying or altering it.

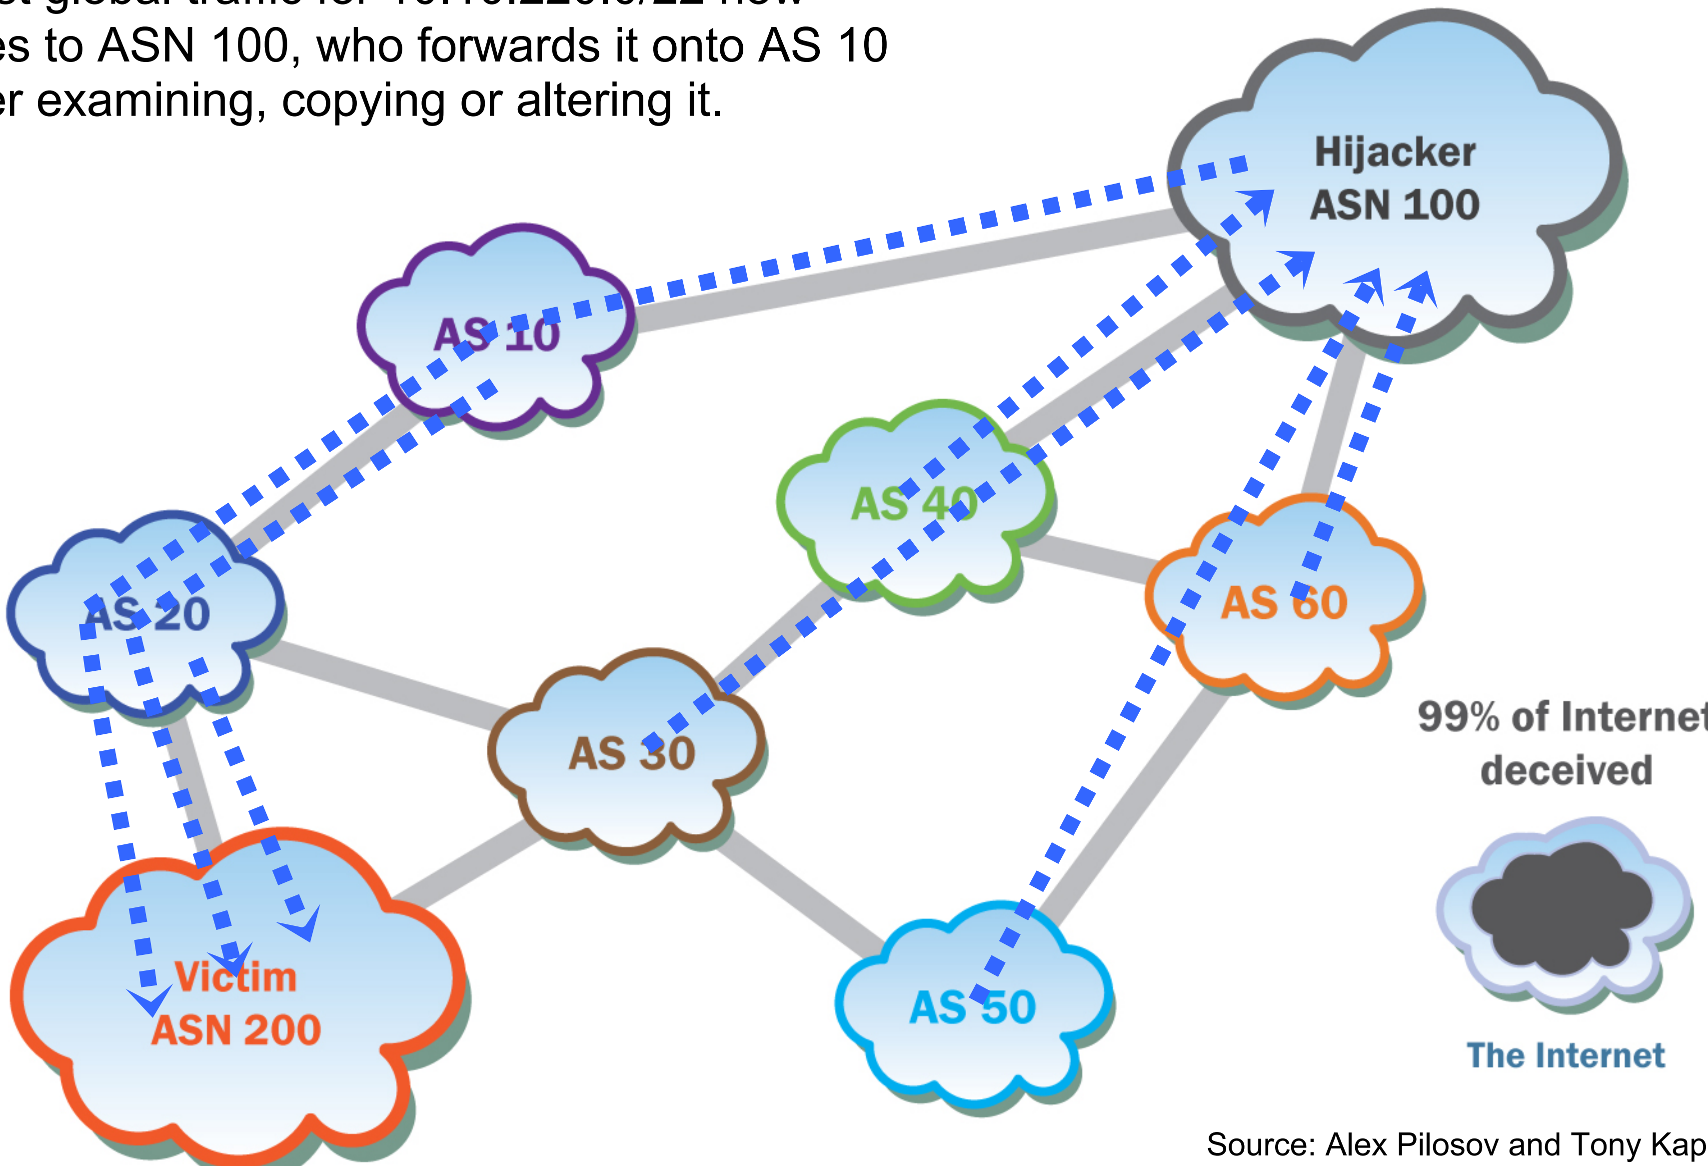

Source: Alex Pilosov and Tony Kapela

Proposed Security Enhancements to BGP

Secure BGP

- Origin Authentication
  - Claim the right to originate a prefix
  - Signed and distributed out-of-band
  - Checked through delegation chain from ICANN
  - Public Key infrastructure approach
- Path Verification
  - Validates that the AS path attribute really indicates
  - … the order ASes traversed by the announcement
  - Uses digital signatures and public key infrastructure

Route Attestations in Secure BGP

If AS a announced path abP then b announced bP to a

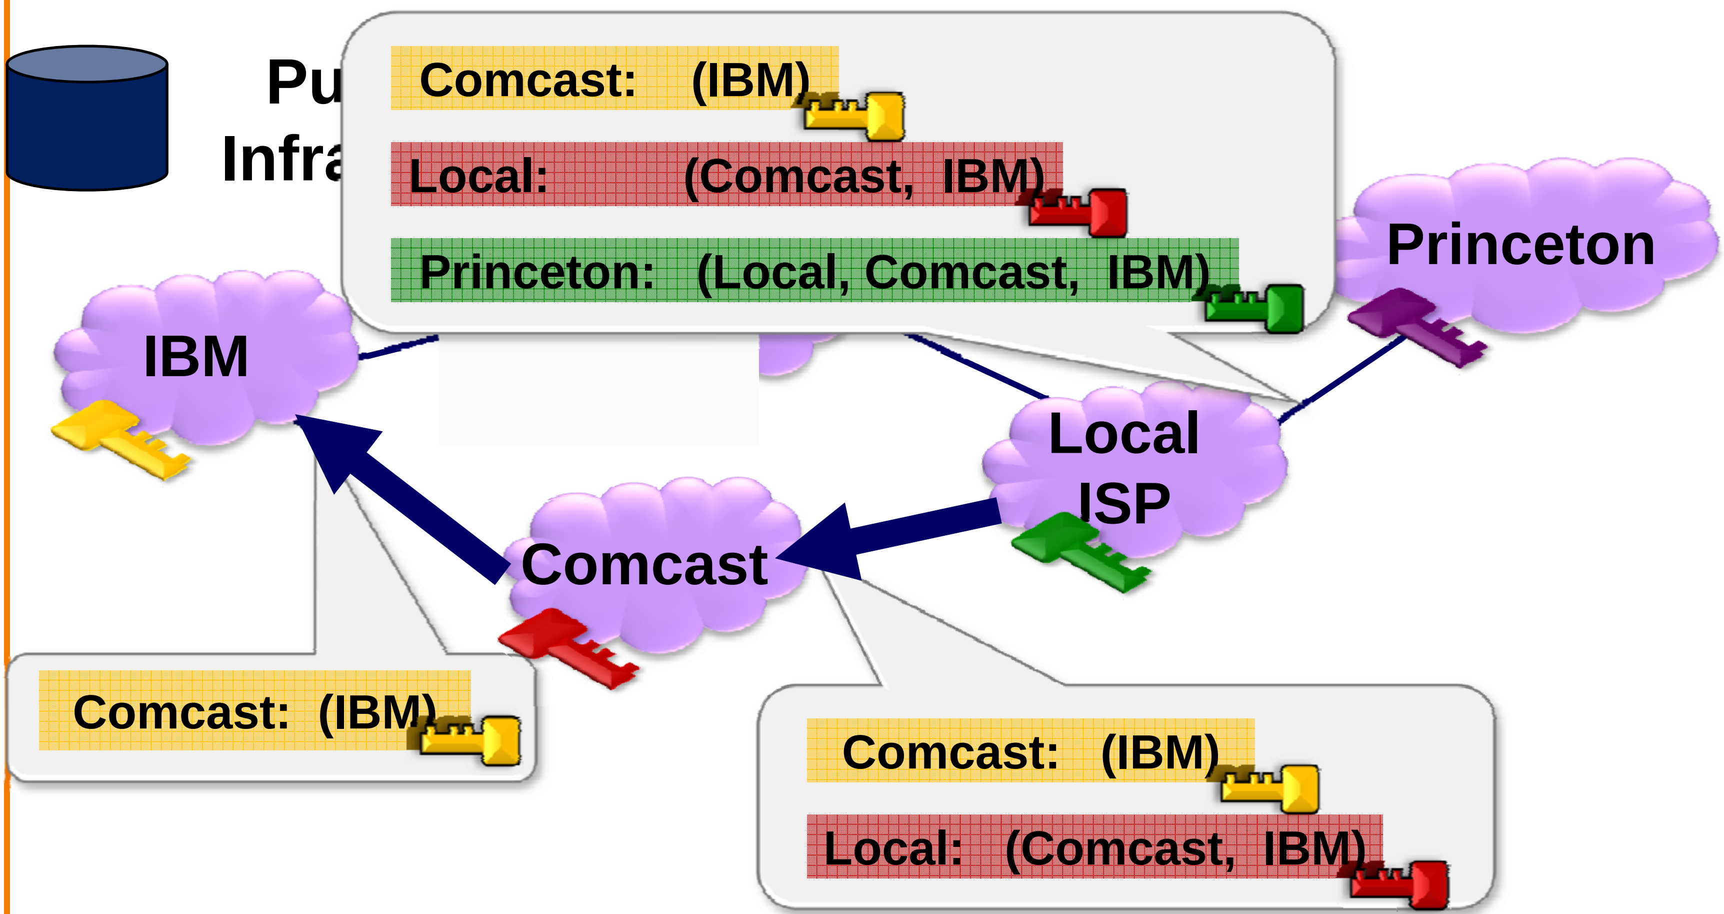

Public Key Signature: Anyone who knows IBM's public key can verify the message was sent by IBM.

Secure BGP Deployment Challenge

- Complete, accurate registries
  - E.g., of prefix ownership
  - What about mobility of prefixes?
- Public Key Infrastructure
  - To know the public key for any given AS
- Efficiency issues
  - E.g., route attestations make BGP messages longer
  - Need to compute public key operations quickly
- Difficulty of incremental deployment
  - Hard to have a “flag day” to deploy S-BGP
  - Expensive (and useless) for a **single** node to upgrade.

Secure Origin BGP

- Origin Authentication
  - As in secure BGP, claim the right to originate a prefix
  - Signed and distributed out-of-band
  - Instead of public key infrastructure, use a web of trust.
- Topology verification
  - Instead of signing messages as they traverse the path
  - .. Maintain a database of AS-level network topology
  - ASes can check that the AS-path attribute is path that
  - …really exists in the network.

Secure Origin BGP

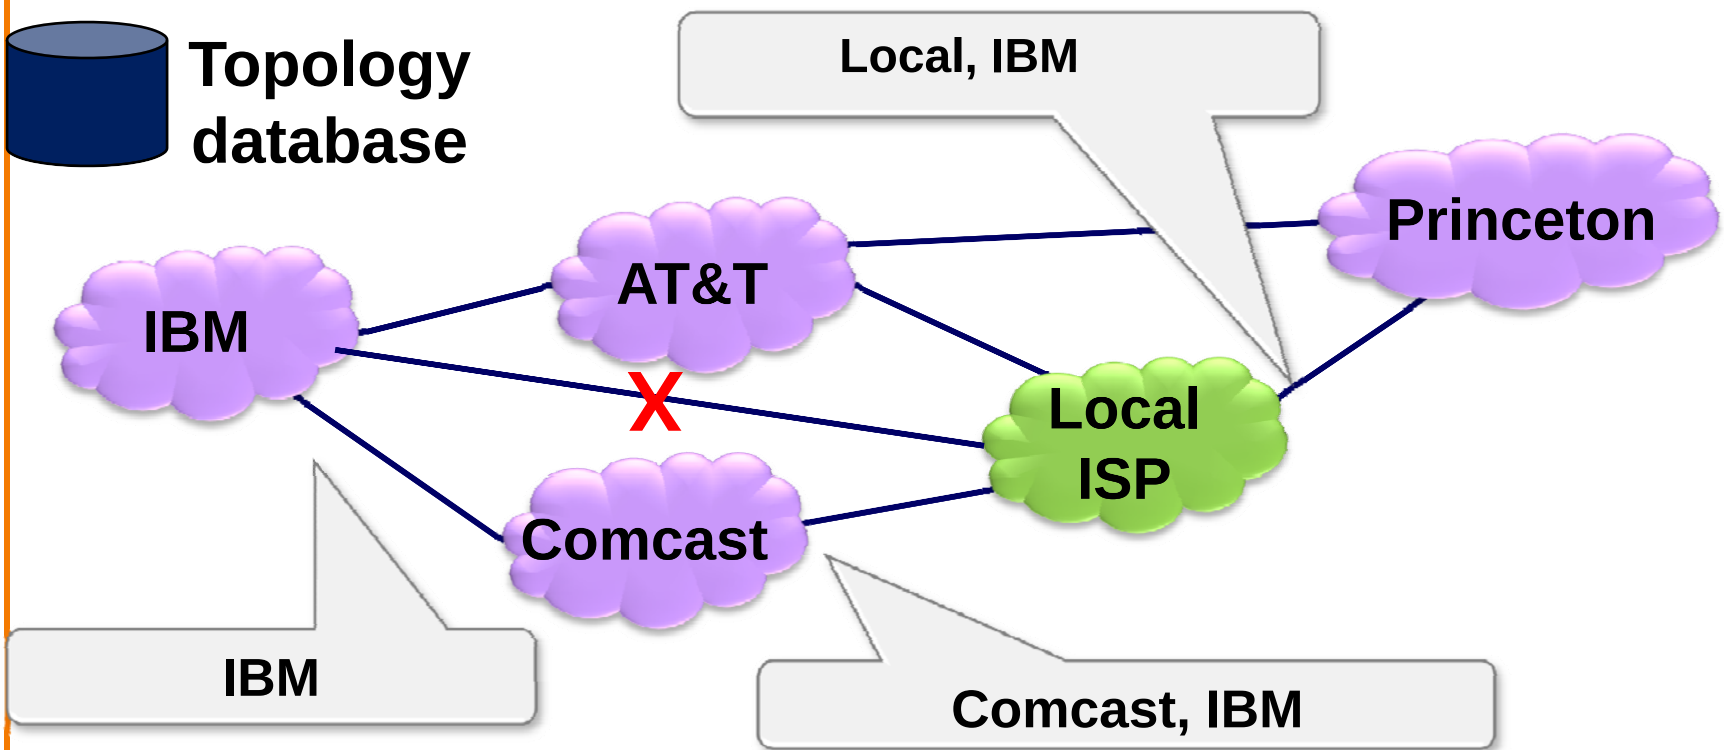

If link between Local ISP and IBM doesn't exist in the topology, then Local ISP will get caught.

But what if it does?

Secure Origin BGP Deployment

- Complete, accurate registries of prefix ownership
  - Mobility of prefixes still and issue
  - Based on Web of Trust, not public key infrastructure
- Efficiency issues
  - Everything is done out of band
  - No crypto on BGP messages
- How hard is incremental deployment?
  - We don't need a "flag day"
  - BUT topology database could reveal private info
- Weaker security guarentee than Secure BGP!
  - Path existing in topology doesn't imply it was announced

Anomaly Detection for BGP

- 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
  - 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
- Security relative to S-BGP, SoBGP?
- What about deployment challenges?

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?

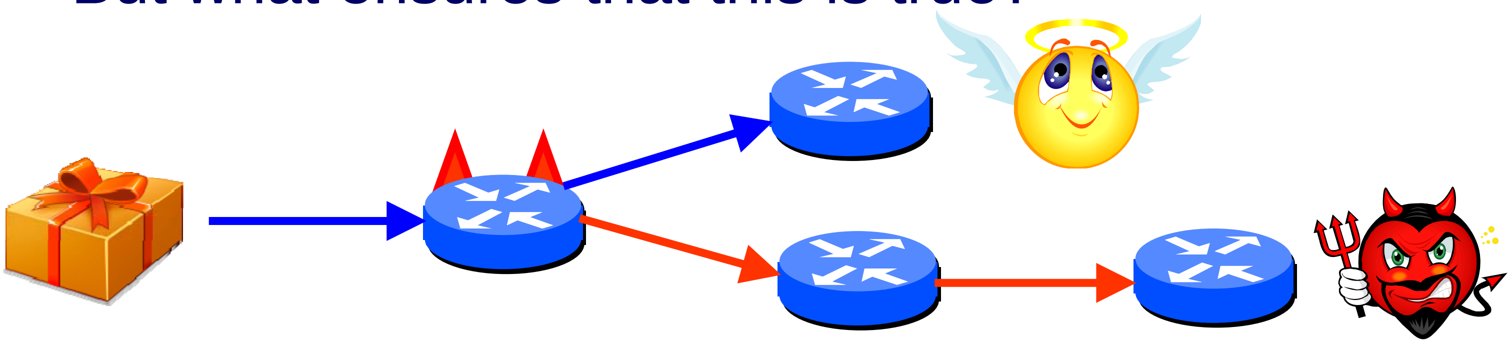

Data-Plane Attacks, Packet Dropping

- 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?

Packet Dropping – Gaming Ping

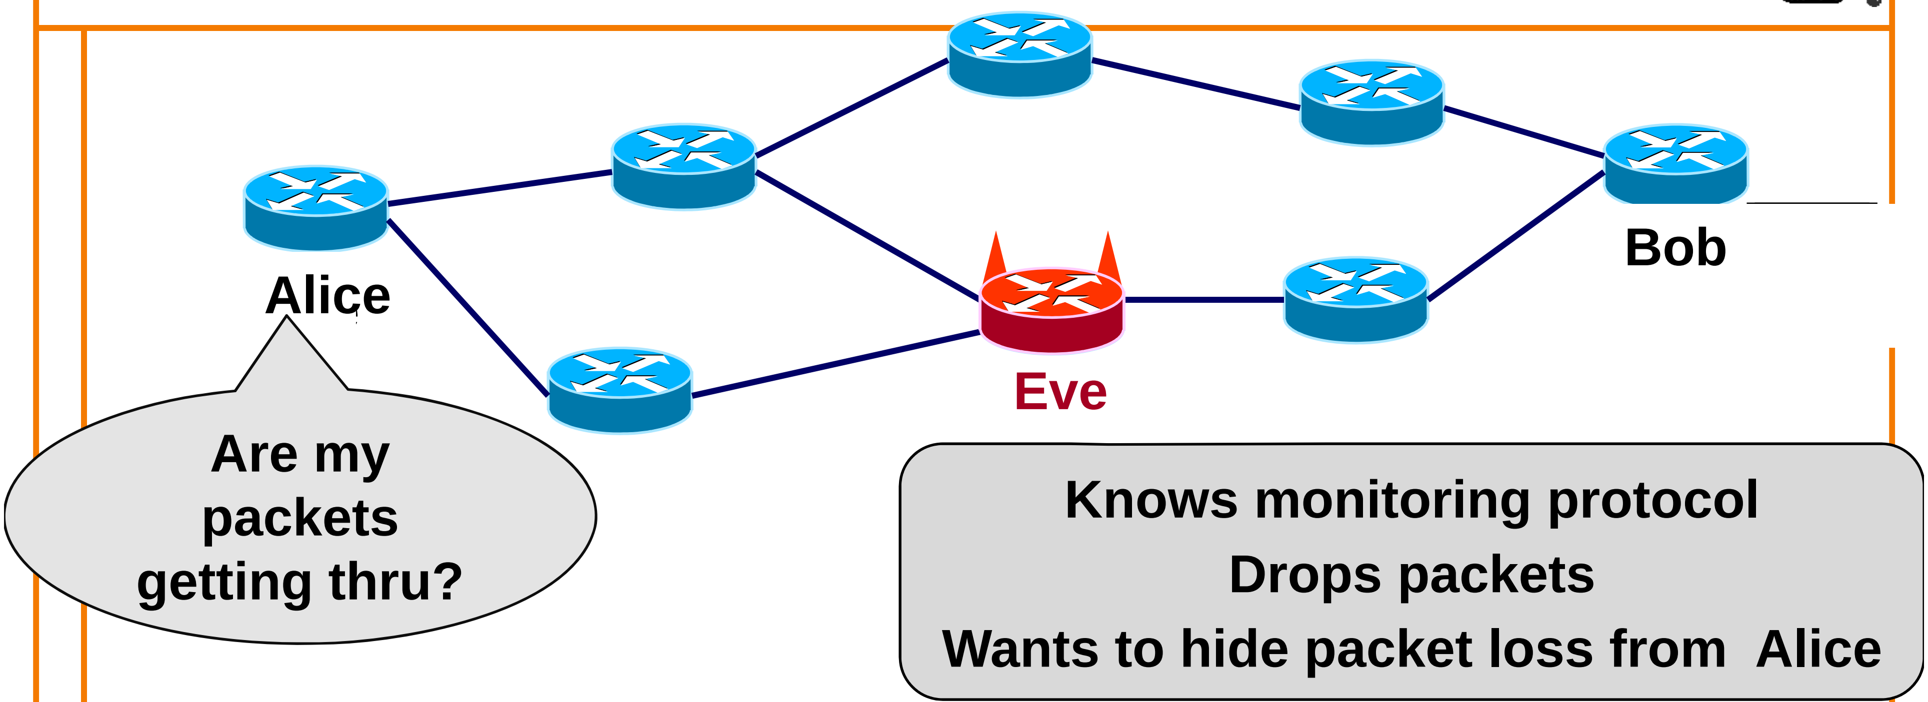

Today's approaches cannot withstand active attack

(ping, traceroute, active probing, marked diagnostic packets)

Data-Plane Attacks, Redirect packets

- 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
- This is really hard to detect?
  - Longer than usual delays? (maybe – if path is long)
  - Traceroute? (can be gamed)
  - Sign **each** packet as goes thru network (impractical)

Redirect Packets - Gaming traceroute

The DEFCON MiTM attack used a trick similar to this one to ‘fool’ traceroute.

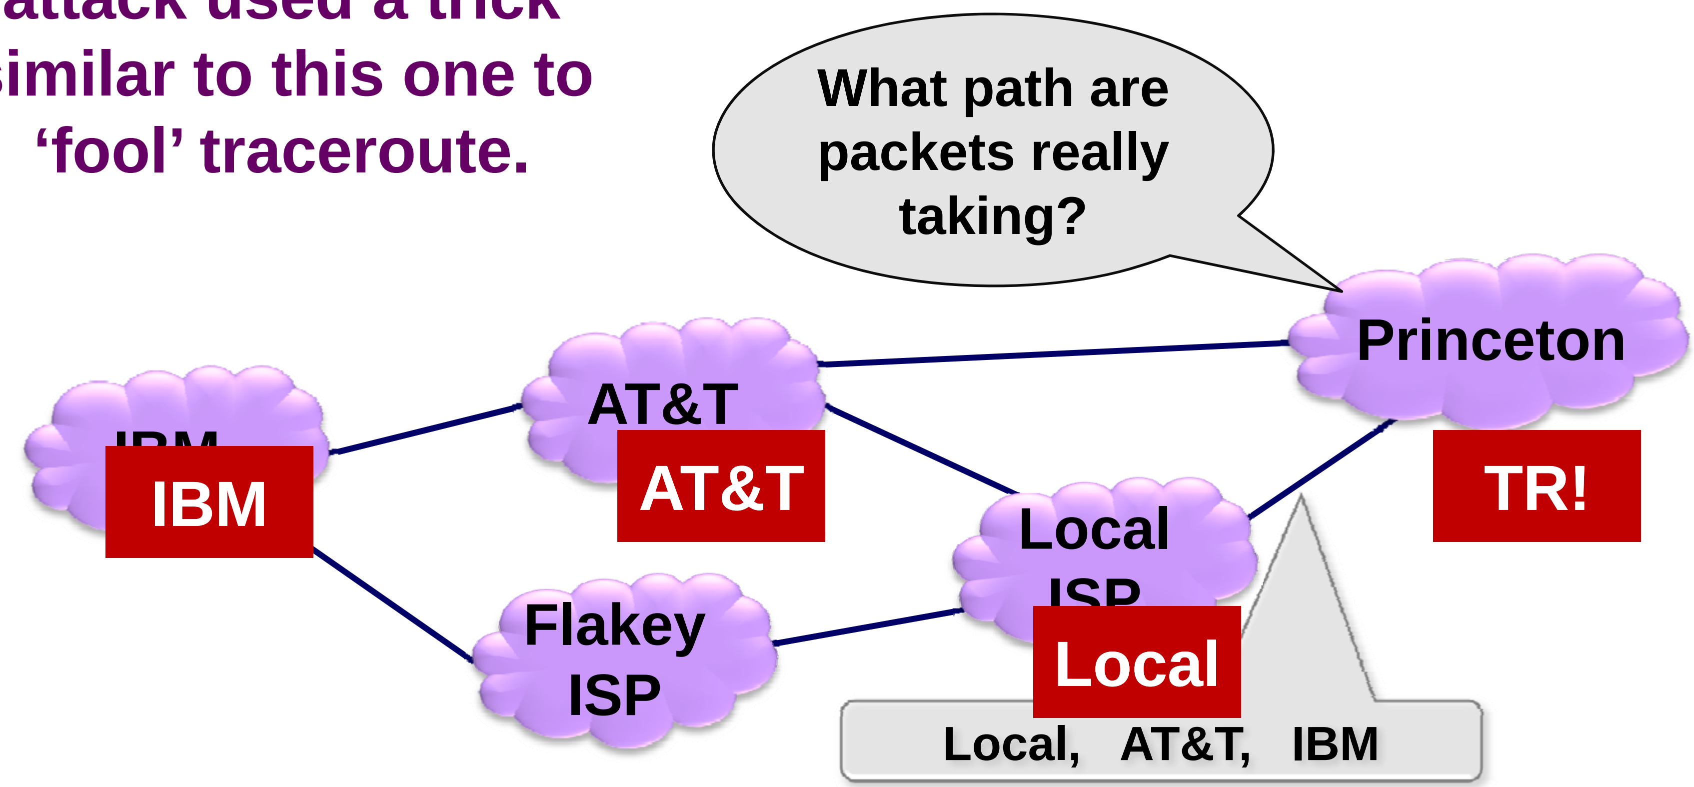

Fortunately, Launching Data-Plane Attacks is Harder

- 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 $ervice ☺

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
- 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 (e.g., spam, identity theft, denial-of-service attacks, port scans, …)

BGP is So Hard to Fix

- Complex system
  - Large, with around 30,000 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 were designed based on trust
  - The insiders are good guys (the military!)
  - 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

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
- Key revocation is also an (ugly) issue