Stealing The Internet

An Internet-Scale Man In The Middle Attack

Presented at NANOG 44
Los Angeles, CA
October, 2008

Tony Kapela
tk@5ninesdata.com

Agenda

- Prior Work
- Hijacking Mechanics
- Route Filtering
- Analysis

Prior MITM Work

- NIST Report July 07: says “it’s possible”
- Paul Francis et al (Cornell): hijack through AS-PATH – >50% interception rate
- UIUC, *A BGP Attack Against Traffic Engineering*, Jintae Kim et al – doesn’t create feasible paths towards target

What's Novel?

- Sub-prefix hijacking is not new
  - I'm well aware of this
- Creating Feasible Return path in-place
  - Possibly novel contribution
- Half-novel
  - TTL increment to hide Layer 3 path
  - Transparent-AS, route-server-client style
  - Hide hijacker from monitoring ASN's

BGP MITM Hijack Concept

- Attacker must have a feasible path back to the victim (sink)
  - Uses AS-PATH loop detection of BGP to create the path (DAG)
- Hijacked route + feasible path to victim permit interception

BGP MITM Setup

1. Plan a viable path to target
2. Note the ASN's seen towards target from the attacker's vantage point
3. Apply as-path prepends naming each of the ASN's intended for viable path
4. Install static routes towards the next-hop of the first AS in viable path
5. Adjourn to Lobby Bar

BGP MITM – First Observe

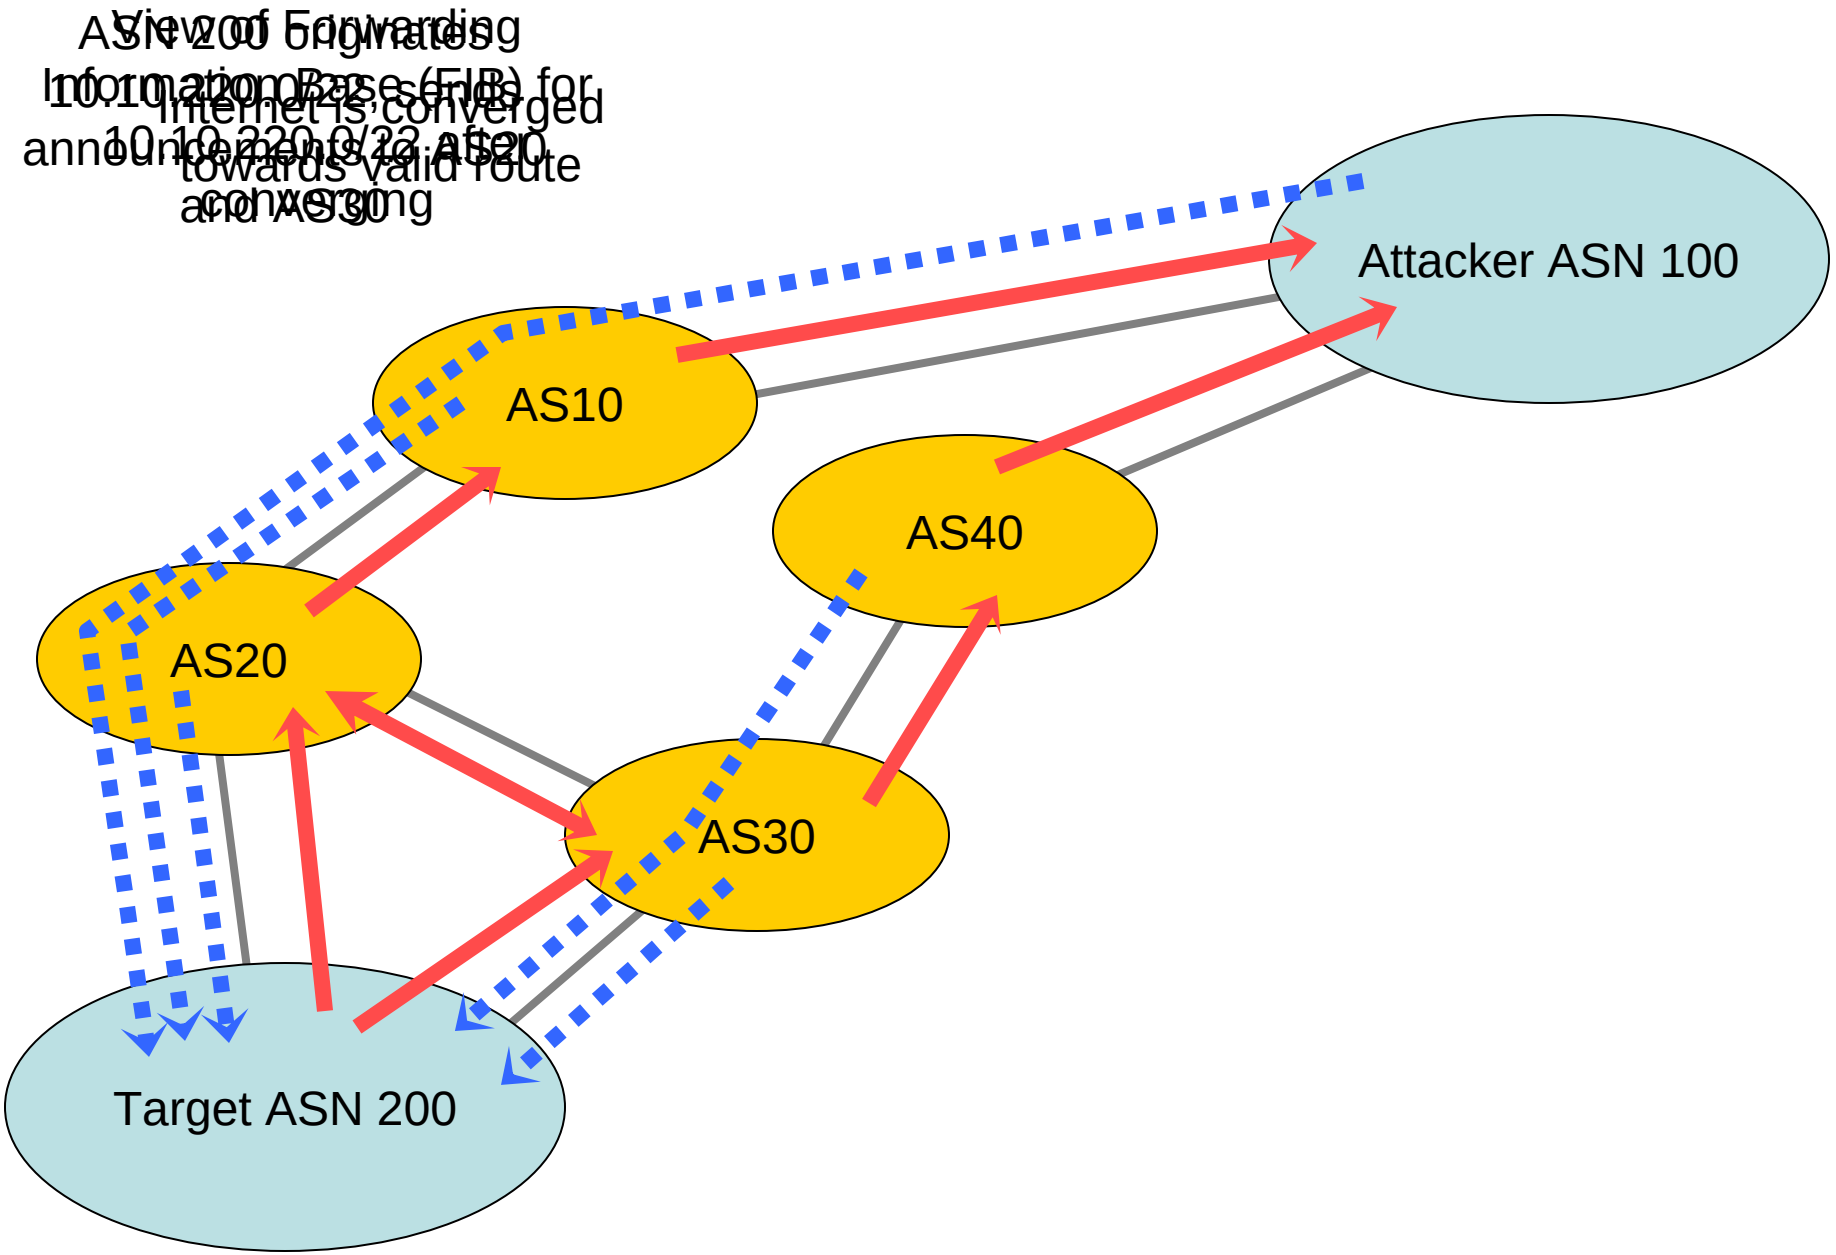

BGP MITM – Plan viable path

We then build the as_path prepend list to include AS10, 20 and 200

ASN 100 sends BGP update for 10.200.0.0/22 via AS10

ASN 100 sends BGP update for 10.200.0.0/22 via AS10

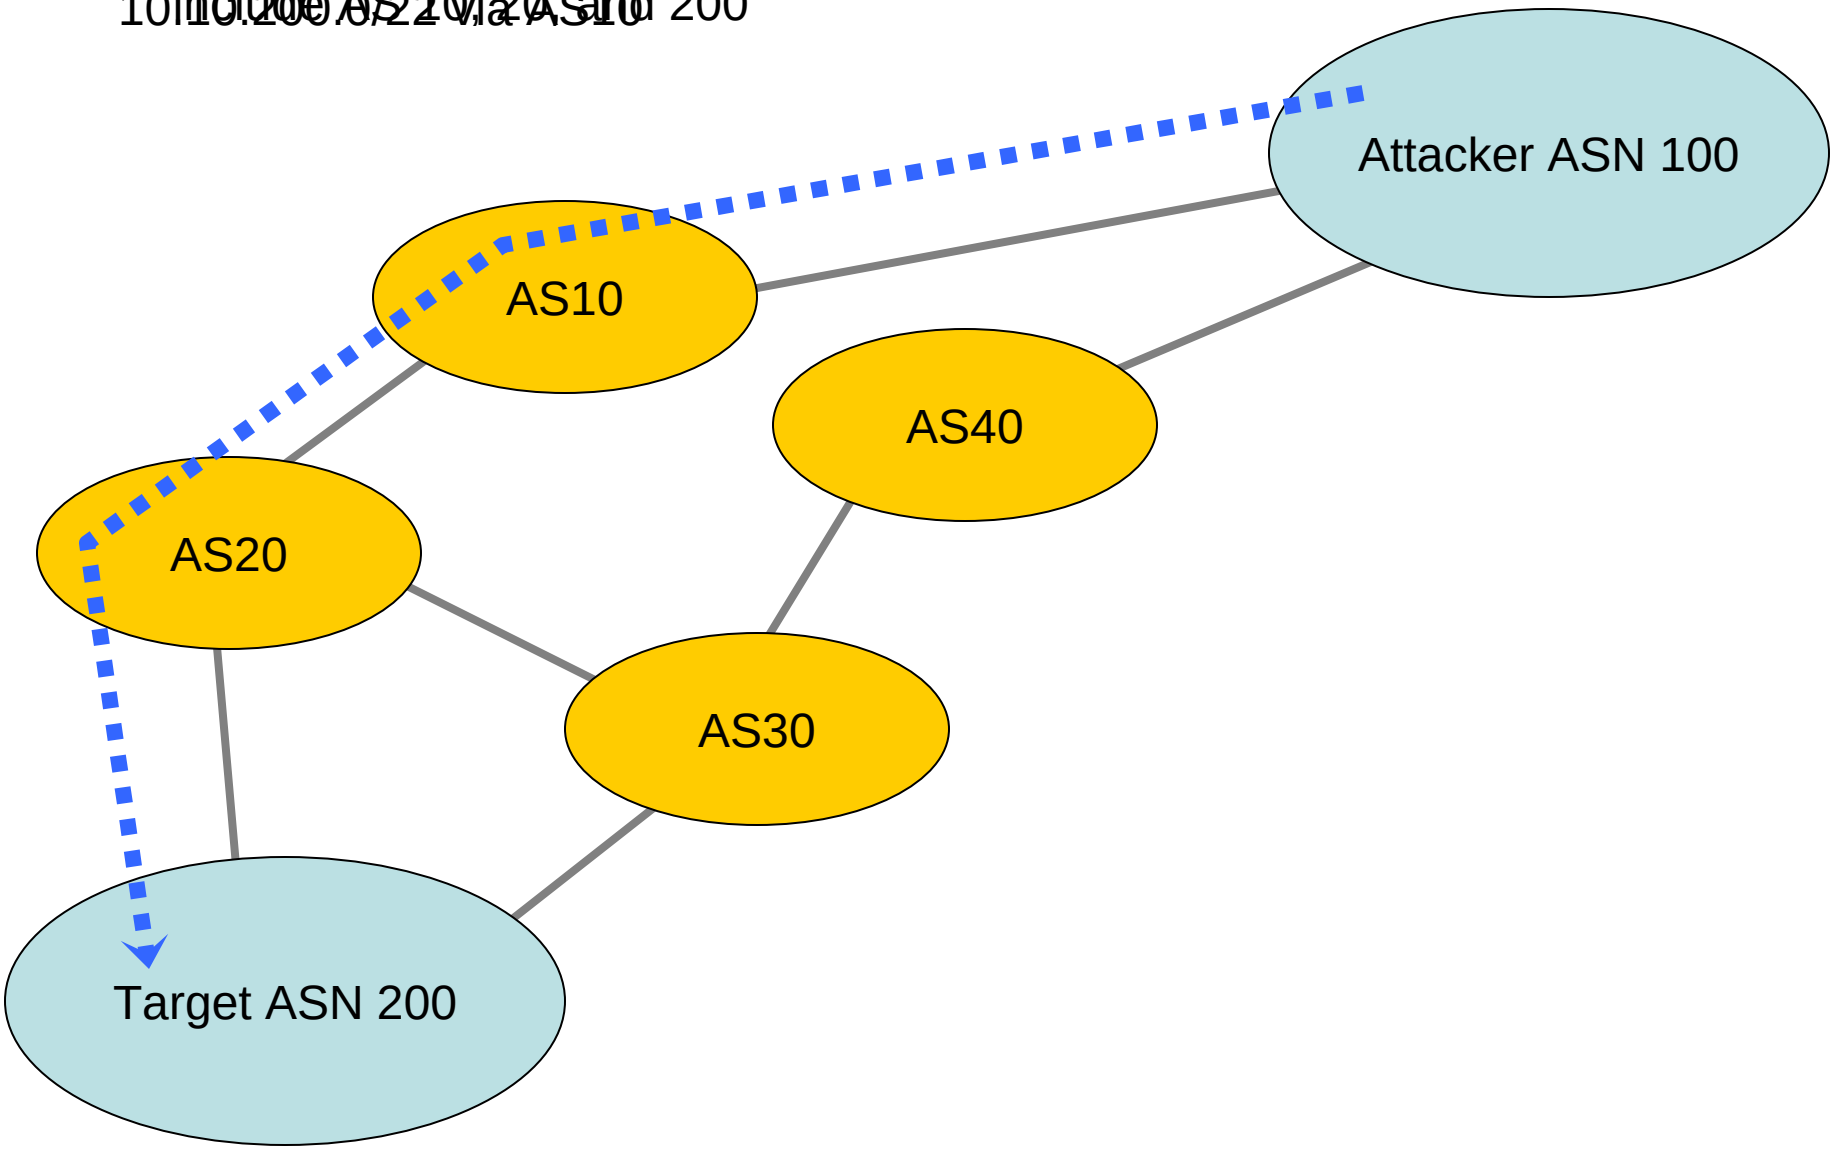

BGP MITM – Setup Routes

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
```

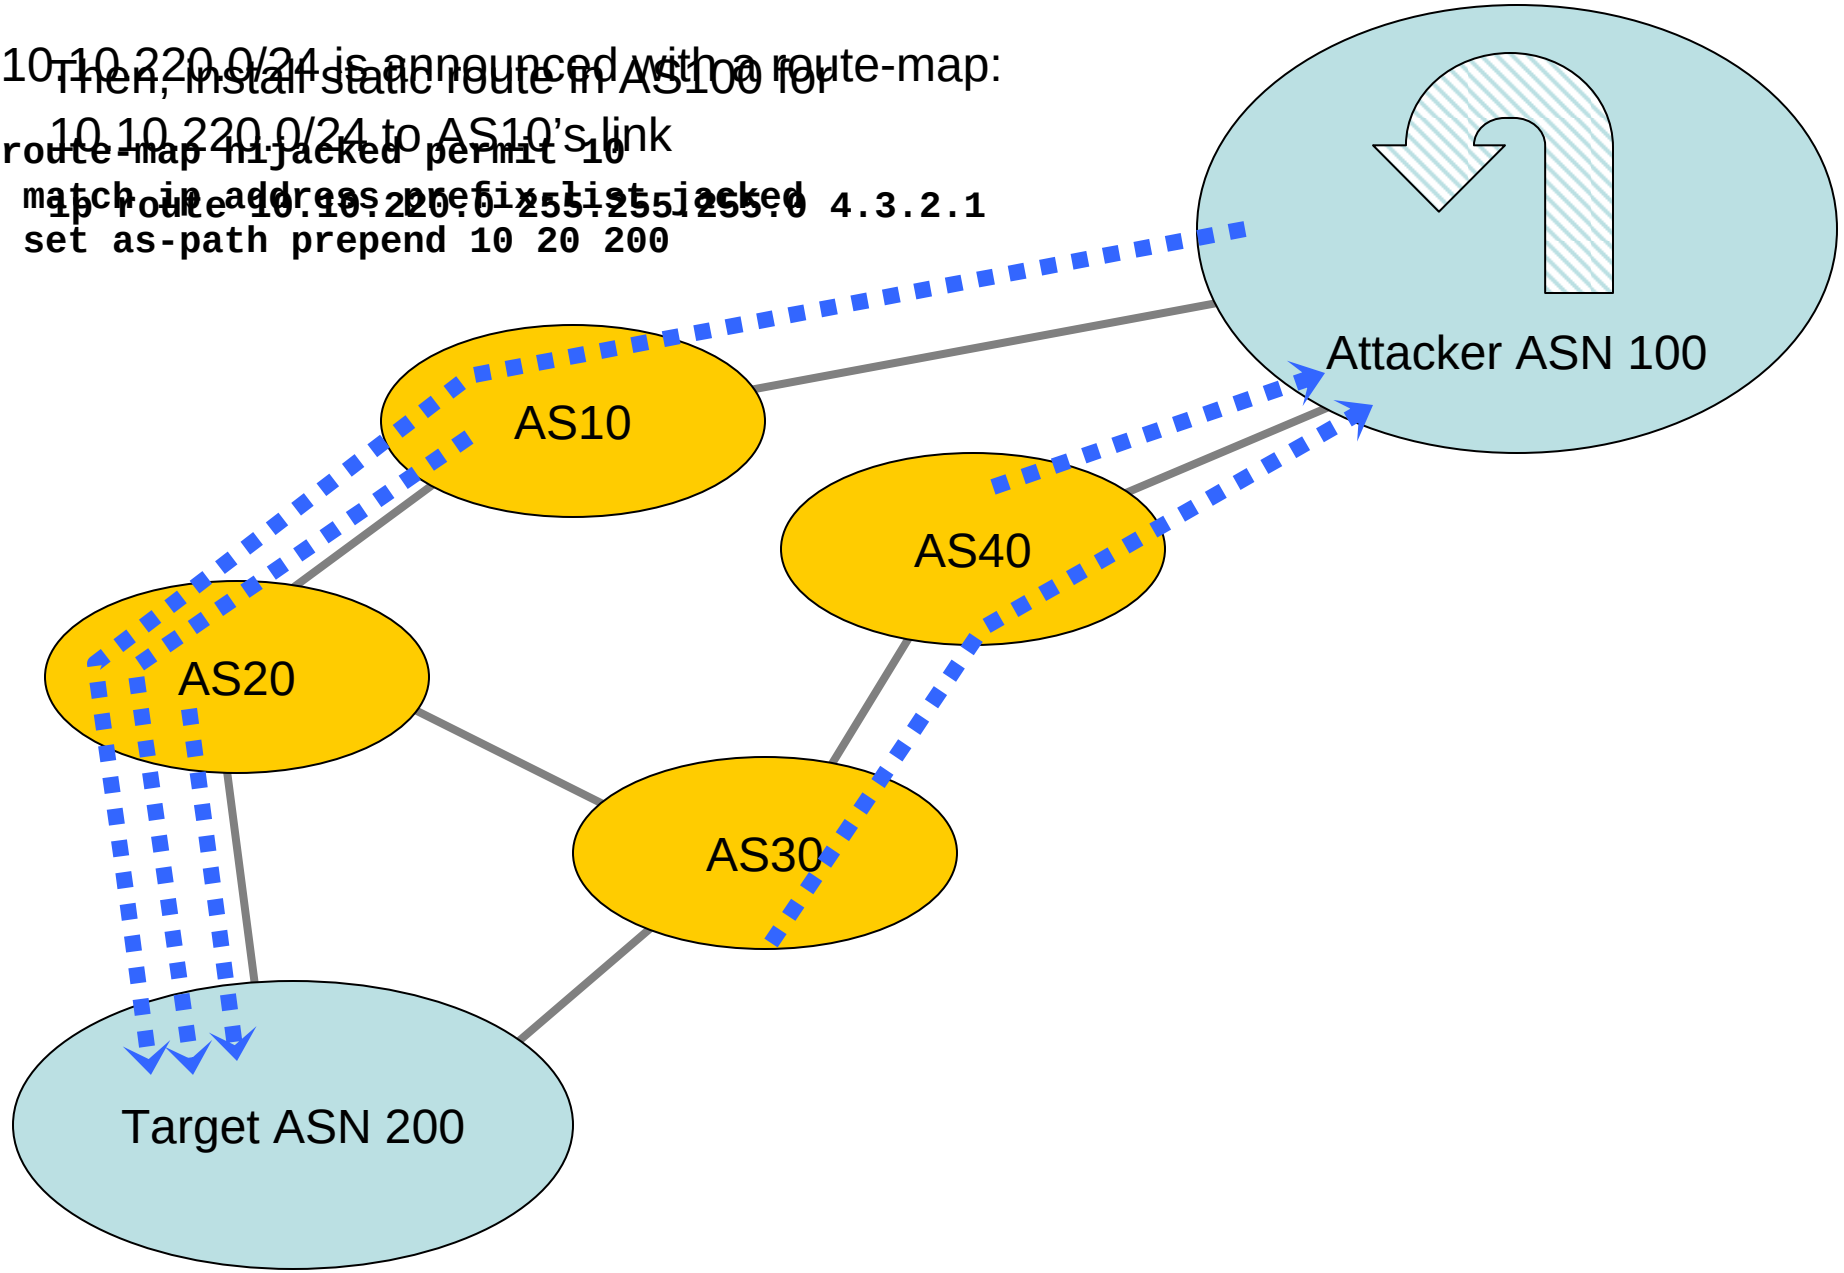

Defcon Hijack Uptake Summary

Timestamp	Plus-t0	Carrying /22	Carrying /24
1218396798	0	252	0
1218396887	80	252	238

Data courtesy Martin Brown of Renesys Corp.

Defcon Prefix Hijacking Statistics

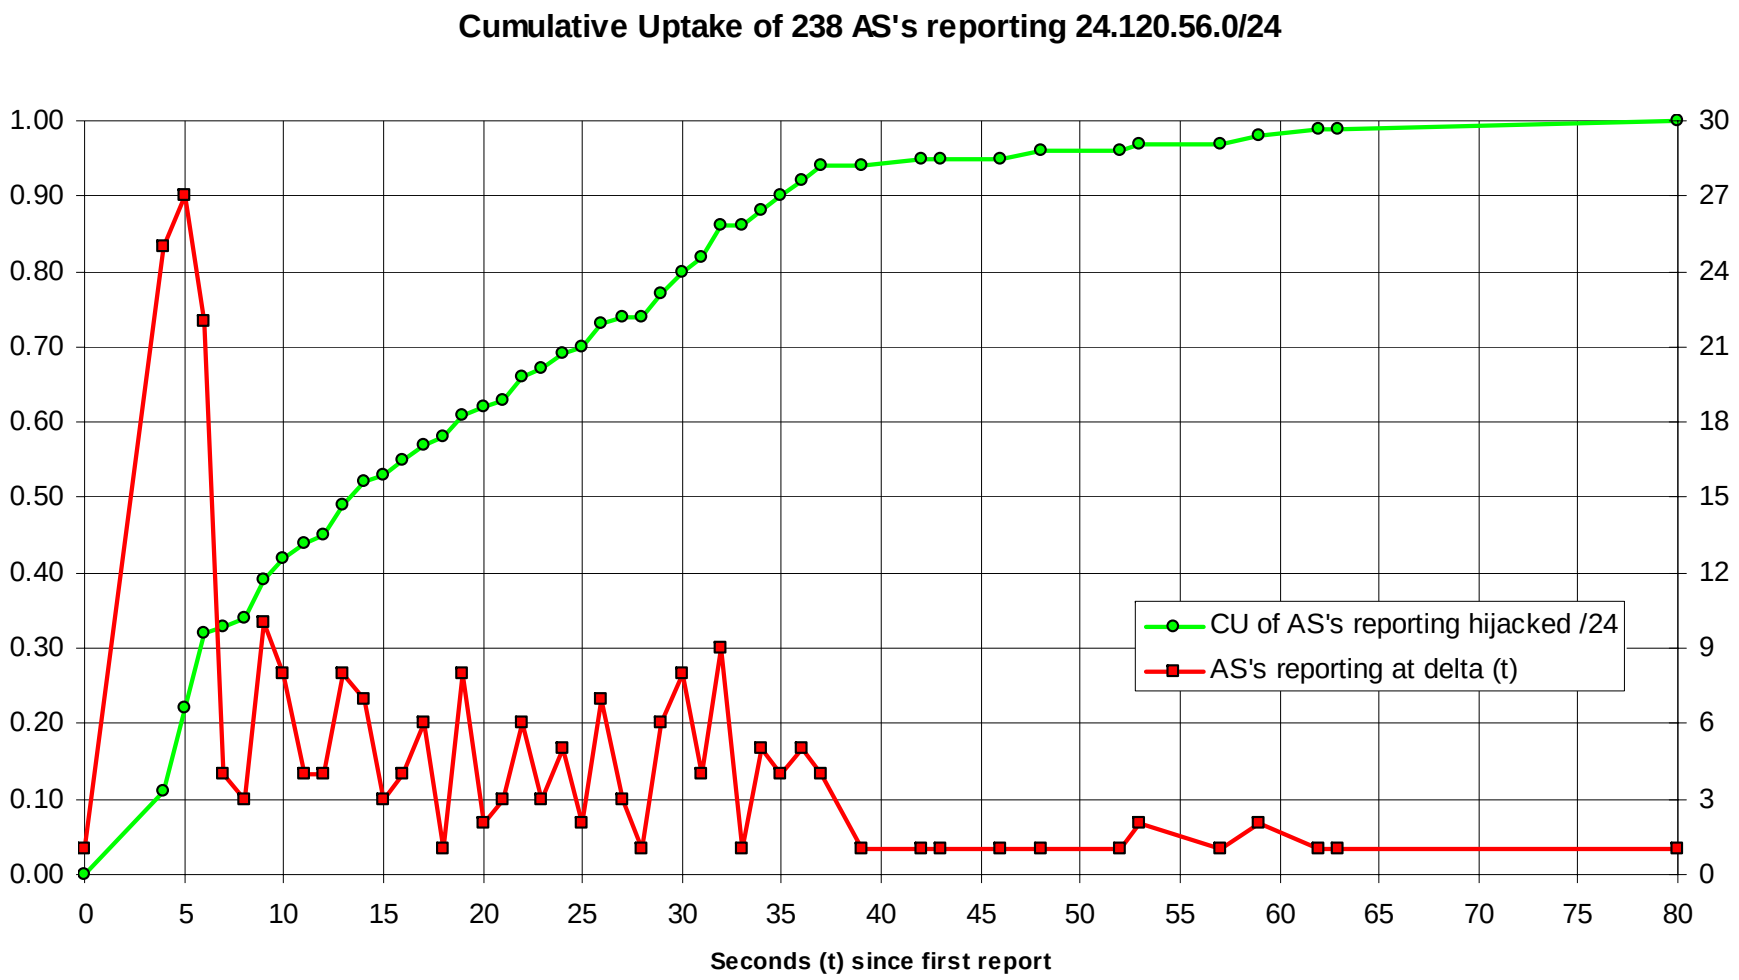

Data courtesy Martin Brown of Renesys Corp.

Observations

- Route propagates (as expected)
  - Nearly everyone accepted
  - Can’t speak to ‘true’ forwarding reality of 30k ASN’s
- Low disruption at “Ramp Up” of hijack
  - “Nearly silent” insertion of eavesdropper
- Definite hit at “Ramp-Down” of hijack
  - FIB micro-loops as expected

Future Of Filtering

- Researchers Welcomed
  - soBGP, sBGP: new features in routing system
  - R-PKI: happens outside routing system
  - Need more creative minds on this problem
- How do we address ‘trust?’
  - Maybe we don’t, build fast alerting systems
  - RIR’s could anchor something

Anonymyzing The Hijacker

- We add value to TTL of packets in transit (iptables)
- Effectively hides hops for the hijacked inbound traffic and ‘viable path’ to target

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
```

Before & After BGP-MITM+TTL

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
```

Anonymizing The Hijacker, More

- Transparent-AS and Route-reflector-client operation
  - Permits attacker to originate prefixes with $whatever for AS-PATH
- AS-PATH now ‘clean’
  - Attacker ASN is simply not present
  - feasible path now looks ‘more correct’

In conclusion

- We saw that BGP MITM can happen nearly invisibly
- We noted the BGP as-path does reveal the attacker unless massaged
- Duh; filter your customers
- Enforce next-as (where you can)

Acknowledgements

- Todd Underwood, Martin Brown (also locally famous) & Renesys Staff
- Latt Mevine (transparent-as)