Global BGP Attacks that Evade Route Monitoring

Henry Birge-Lee
Princeton University

With
Jennifer Rexford (Former CS Department Chair/Current Provost)
Maria Apostolaki (Prof. Electrical and Computer Engineering)

Overview

1. **Motivation and Terminology**
2. Background
3. Launching the Attack in the Real World
4. Attack Viability
5. Countermeasures

What is BGP monitoring and why is it important

1. Using a combination public (e.g., **RIPE NCC RIS**) and/or private data feeds to observe global BGP updates

2. Monitoring can identify BGP attacks for real time responses/mitigations or after-the-fact incident investigation

3. BGP monitoring systems are coming to play a large role in routing security given the relatively slow adoption of high-security interdomain routing protocols (e.g., BGPSEC)

Why are we are interested in stealthy BGP attacks

1. If you can make your attack stealthy you can (potentially) evade detection and mitigation

2. Stealthy attacks are possible and studied by prior work

   Birge-Lee et al. ‘19 SICO https://doi.org/10.1145/3319535.3363197
   Milolidakis et al. ‘23 “On the Effectiveness of BGP Hijackers That Evade Public Route Collectors," https://doi.org/10.1109/ACCESS.2023.3261128
   Morillo et al. ‘21 ROV++ https://dx.doi.org/10.14722/ndss.2021.24438

Exactly what do we mean by “stealthy attack”

1. Many different metrics:
   a. How much of the Internet sees your route?
   b. Do attack detection algorithms think your attack is just background noise?
      N.A. for this talk
   c. Do networks using the route see the route?
      also achieved by our attack
   d. Do BGP monitoring services see the route? (focus of today’s talk)

Terminology

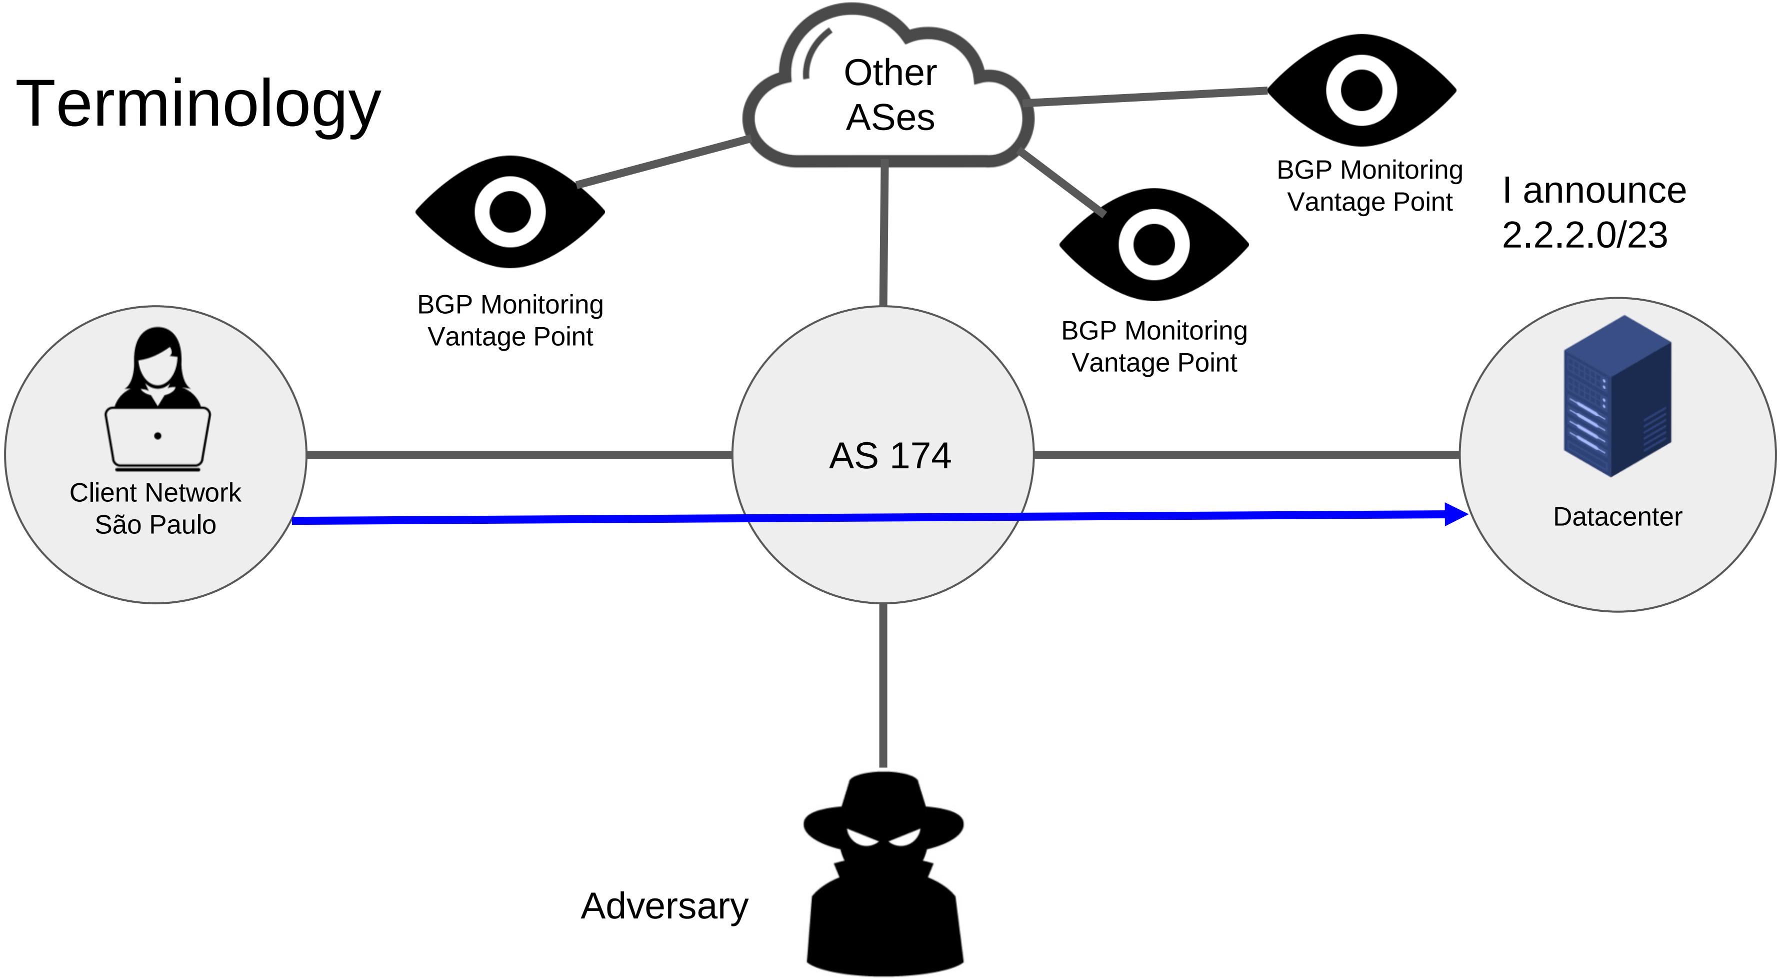
Other ASes
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
I announce 2.2.2.0/23
Client Network São Paulo
AS 174
Datacenter
Adversary

Terminology

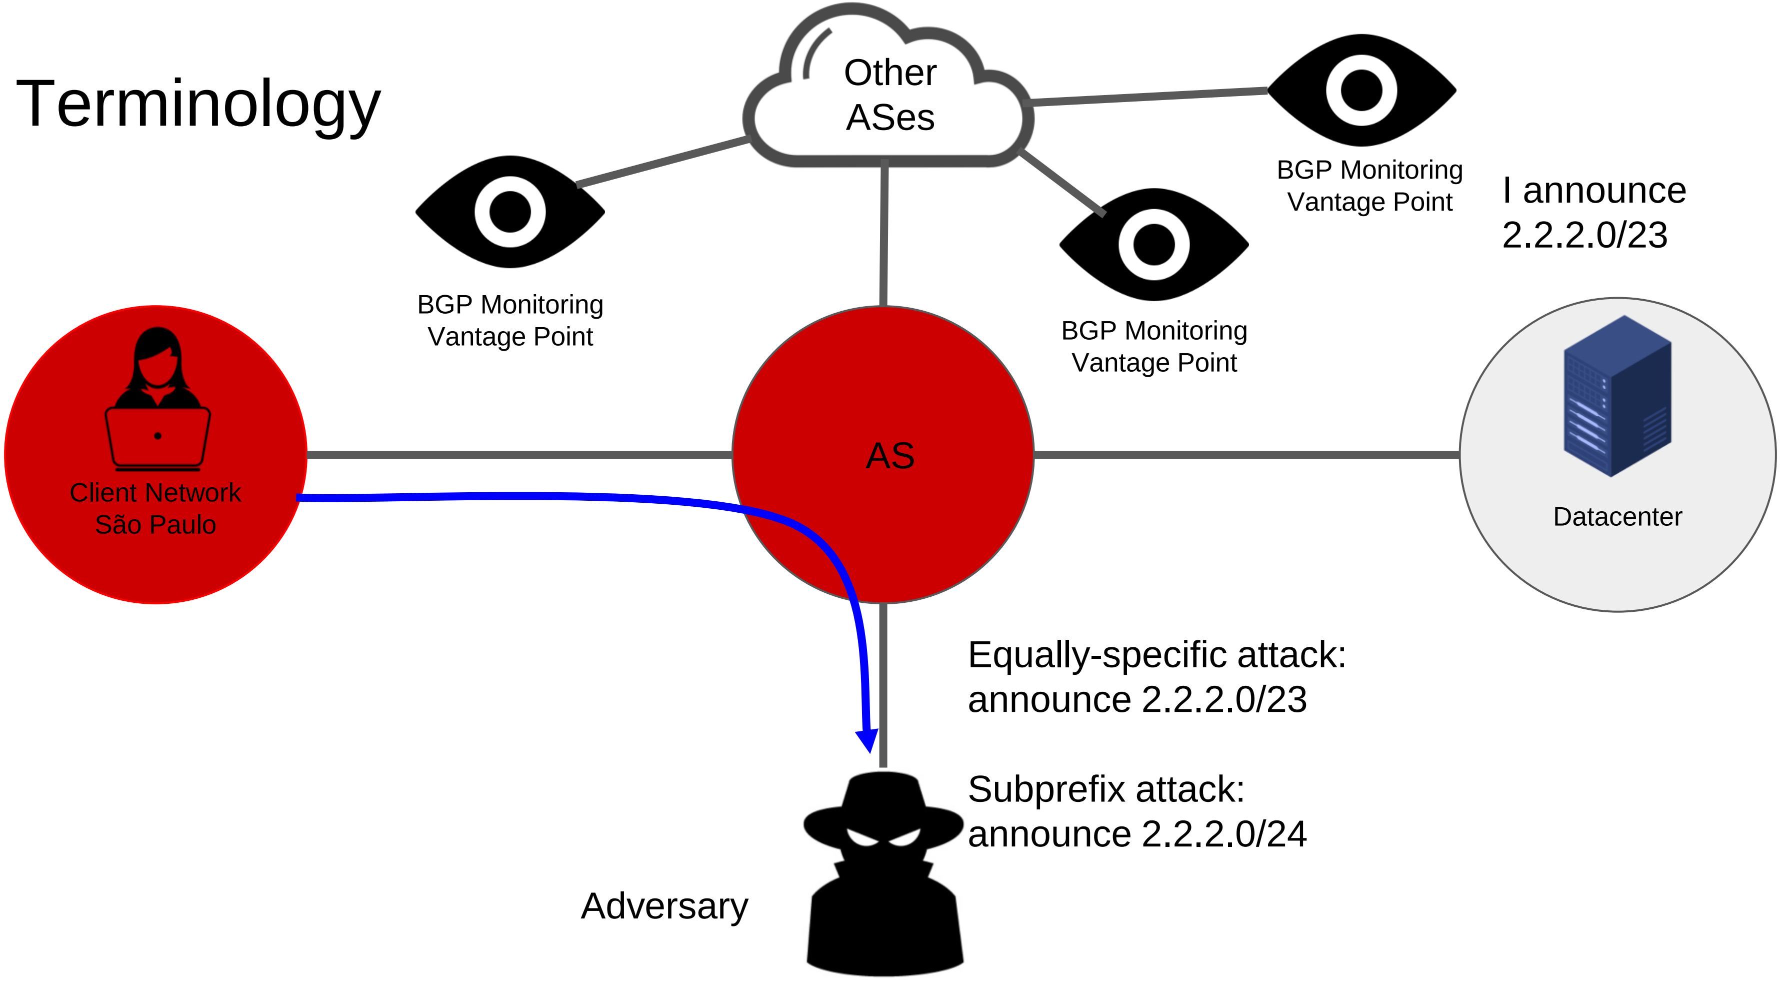
Other ASes
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
Client Network São Paulo
AS
Datacenter
Adversary

I announce 2.2.2.0/23

Equally-specific attack:
announce 2.2.2.0/23

Subprefix attack:
announce 2.2.2.0/24

Overview

1. Motivation and Terminology
2. **Background**
3. Launching the Attack in the Real World
4. Attack Viability
5. Countermeasures

Work to Date on Stealthy Attacks

1. Use an equally-specific BGP attack and shaping your announcement (BGP communities/AS-path poisoning) to avoid public monitors
   a. Upside: Viable (Birge-Lee et al. ‘19 SICO; Milolidakis et al. ‘21 Smart BGP Hijacks)
   b. Downside: Lot of monitors means only a small portion of the uses the malicious route
   c. Downside: Networks using the malicious route (source victims) see it in their tables

2. Use a subprefix attack and stop it from spreading to the source victim
   a. Upside: Source victims have not knowledge of the malicious route they are using (Morillo et al. ‘21 ROV++)
   b. Downside: Subprefix attacks spread all over the Internet, not stealthy at all from public monitors

SICO: Surgical Interception Attacks by Manipulating BGP Communities
https://dl.acm.org/doi/10.1145/3319535.3363197

A. Milolidakis, T. Bühler, K. Wang, M. Chiesa, L. Vanbever and S. Vissicchio, "On the Effectiveness of BGP Hijackers That Evade Public Route Collectors," in IEEE Access, vol. 11, pp. 31092-31124, 2023, https://doi.org/10.1109/ACCESS.2023.3261128

ROV++: Improved Deployable Defense against BGP Hijacking
https://www.ndss-symposium.org/ndss-paper/rov-improved-deployable-defense-against-bgp-hijacking/

Real-world Stealthy Attacks

1. Adversary established direct peering with source victim (Yahoo mail) at DE-CIX
2. Announced malicious subprefix only over direct peering
   a. Avoided all BGP monitors
3. Was eventually caught via coordination from prefix owner and Yahoo
   a. Yahoo confirmed that the networks in their table were not the same as the networks announced by the prefix owner
4. Attack was limited in spread and could not be scaled: only spread to Yahoo

Invisible_Hijacking
https://ripe72.ripe.net/presentations/45-Invisible_Hijacking.pdf

Current beliefs on stealthy attacks

“Although higher-Type attacks [i.e., attacks with more AS prepends] may sometimes be completely stealthy to the infrastructure (e.g., in 21% of the attacks for the Type-4 simulations), such hijackers could not affect while remaining stealthy more than 2% of the Internet.” [1]

[1] A. Milolidakis, T. Bühler, K. Wang, M. Chiesa, L. Vanbever and S. Vissicchio, "On the Effectiveness of BGP Hijackers That Evade Public Route Collectors," in IEEE Access, vol. 11, pp. 31092-31124, 2023, https://doi.org/10.1109/ACCESS.2023.3261128

Current beliefs on stealthy attacks

“Although higher-Type attacks [i.e., attacks with more AS

We achieve an attack that:

affects the vast majority of the Internet
is seen by zero BGP monitors
is not in the routetable of affected networks

remaining stealthy more than 2% of the Internet.” [1]

[1] A. Milolidakis, T. Bühler, K. Wang, M. Chiesa, L. Vanbever and S. Vissicchio, "On the Effectiveness of BGP Hijackers That Evade Public Route Collectors," in IEEE Access, vol. 11, pp. 31092-31124, 2023, https://doi.org/10.1109/ACCESS.2023.3261128

Overview

1. Motivation and Terminology
2. Background
3. **Launching the Attack in the Real World**
4. Attack Viability
5. Countermeasures

Key Attack Insight

- BGP communities are 32-bit tags attached to BGP routes which can impact how routers propagate BGP announcements

- The RFC that introduced BGP communities (RFC 1997) defined several well-known communities including NO_EXPORT which tells routers not to propagate a BGP announcement outside of their own AS

- **The NO_EXPORT community (supported out of the box by routers) prohibits the exporting of a route to BGP monitoring even if a network is a direct peer of a BGP monitoring service**

- Previous work on stealthy BGP attacks assumes if a network is a peered with a monitoring service, it will send all routes it uses to the monitoring service

- This insight overcomes this allowing for much more effective stealthy attacks

Differentiation with previous work

Previous work: if malicious route is used by peer network, monitoring service sees route

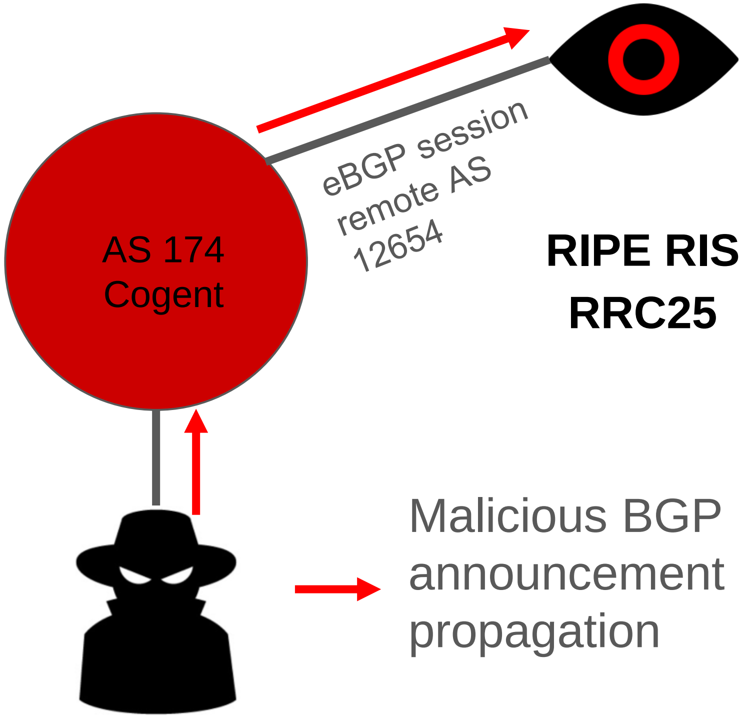

eBGP session remote AS 12654

AS 174 Cogent

RIPE RIS RRC25

Malicious BGP announcement propagation

Our work: peer networks do not share their routes with monitoring services

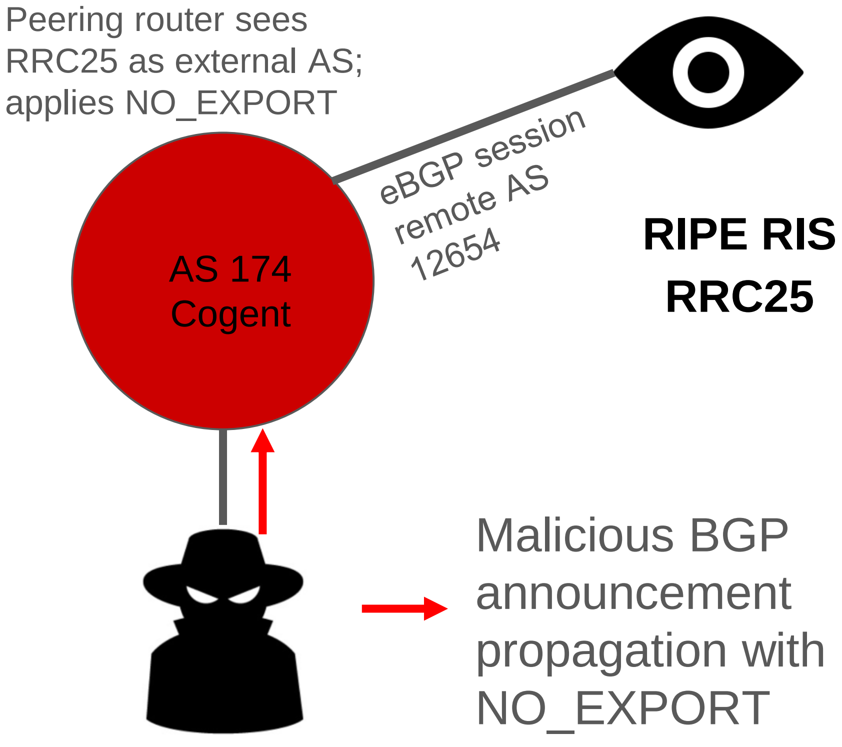

Peering router sees RRC25 as external AS; applies NO_EXPORT

eBGP session remote AS 12654

AS 174 Cogent

RIPE RIS RRC25

Malicious BGP announcement propagation with NO_EXPORT

Attack Steps

1. Announce a malicious sub-prefix of your victim’s prefix to a major transit provider (attacker needs to social engineer or otherwise bypass prefix filters)
2. Attach the RFC NO_EXPORT community to your announcement
    a. Currently BGP monitoring services work by establishing eBGP sessions with peers (e.g., major transit networks that provide route tables to the monitoring service)
    b. To these the routers that run these sessions, the sessions appear to be standard eBGP sessions to what looks like a remote network, thus the RFC NO_EXPORT community applies
3. The malicious route will be installed in a major transit provider (optimally your victim’s upstream)

4. The malicious route will not be sent to BGP monitoring

All real world attacks were conducted ethically

1. Hijack ourselves approach (we controlled IP prefix)

2. All nodes (including “adversary”) were authorized to announce route

3. No real services on prefix used in experiment

Benign Setup

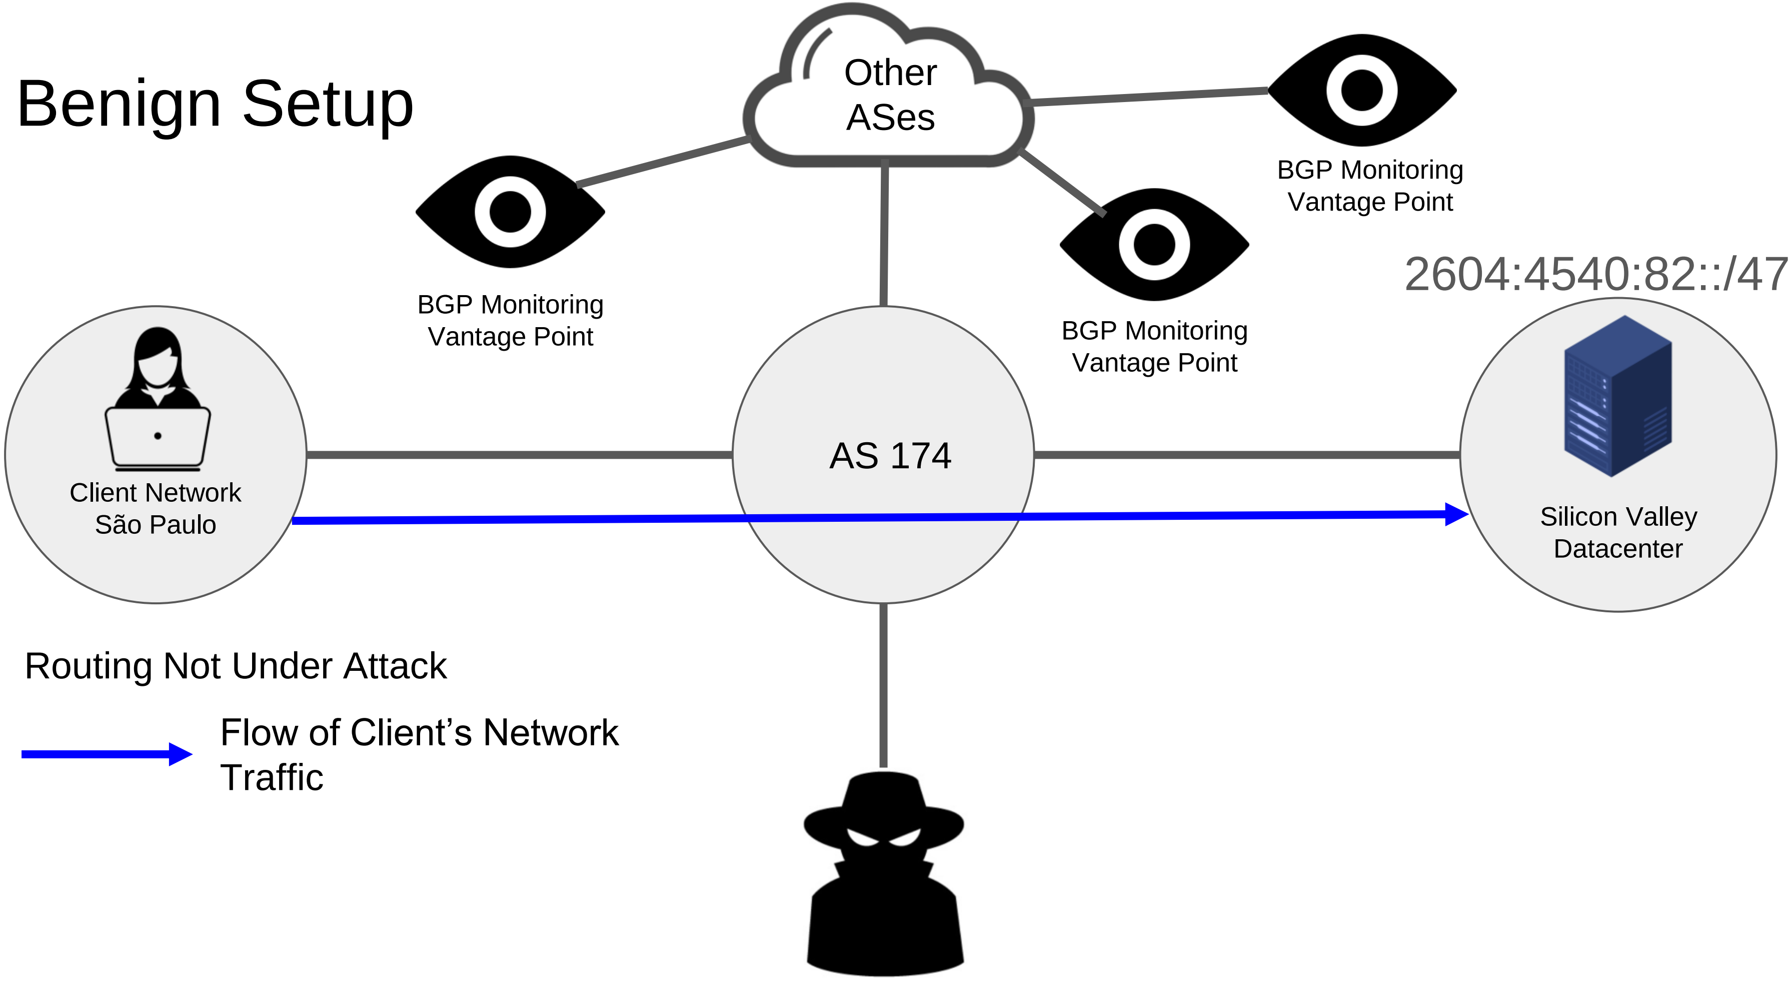

Other ASes

BGP Monitoring Vantage Point

BGP Monitoring Vantage Point

BGP Monitoring Vantage Point

2604:4540:82::/47

Client Network São Paulo

AS 174

Silicon Valley Datacenter

Routing Not Under Attack

→ Flow of Client’s Network Traffic

Standard sub-prefix attack

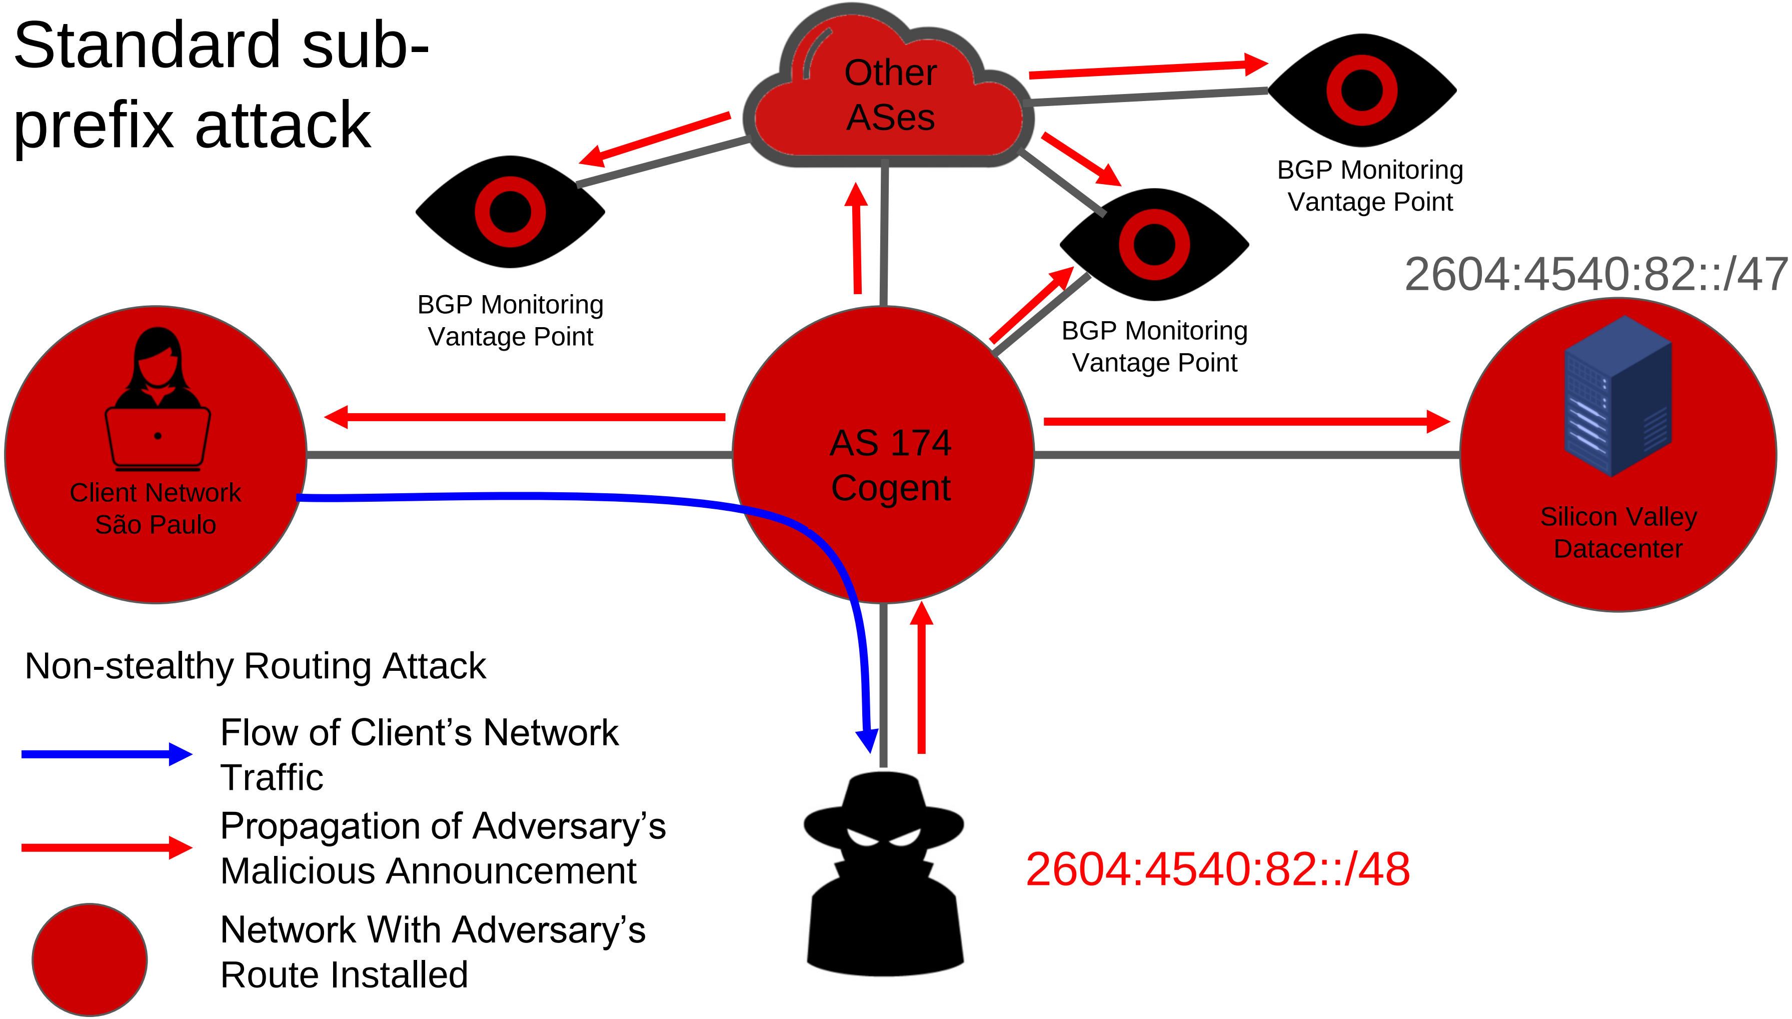
Other ASes
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
2604:4540:82::/47
Client Network São Paulo
AS 174 Cogent
Silicon Valley Datacenter
2604:4540:82::/48

Non-stealthy Routing Attack

→ Flow of Client’s Network Traffic

→ Propagation of Adversary’s Malicious Announcement

● Network With Adversary’s Route Installed

Standard subprefix attack is seen everywhere

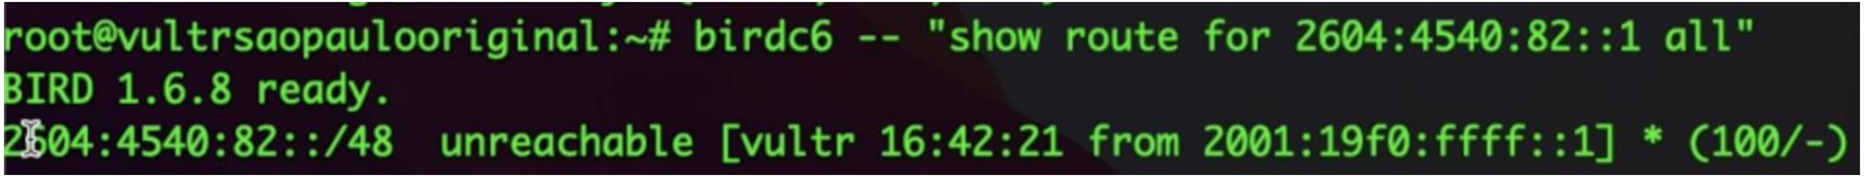

```
root@vultrsaopaulooriginal:~# birdc6 -- "show route for 2604:4540:82::1 all"
BIRD 1.6.8 ready.
2604:4540:82::/48  unreachable [vultr 16:42:21 from 2001:19f0:ffff::1] * (100/-)
```

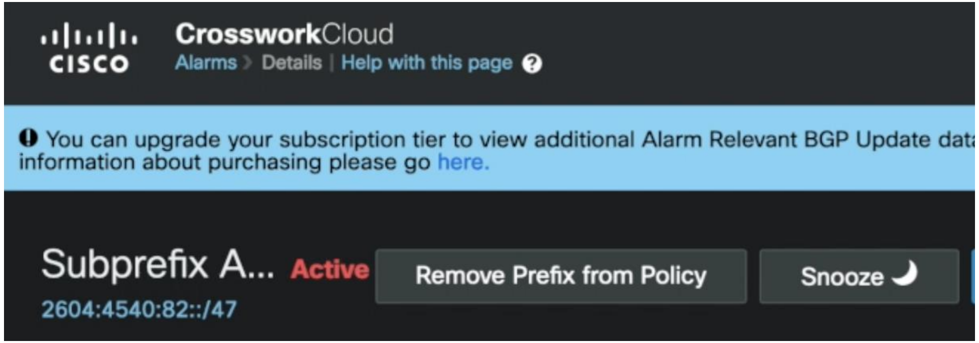
cisco
CrossworkCloud
Alarms › Details | Help with this page ?
Remove Prefix from Policy
Snooze

You can upgrade your subscription tier to view additional Alarm Relevant BGP Update dat
information about purchasing please go here.

Subprefix A... Active
2604:4540:82::/47

```
ris_subscribe_ok {'subscription': {'prefix': '2604:4540:82::/48', 'moreSpecific': True, 'lessSpecific': False},
'socketOptions': {'includeRaw': True, 'acknowledge': True}}
ris_message {'timestamp': 1663691310.63, 'peer': '2001:978:4::b', 'peer_asn': '174', 'id': '25-28178-16349017', 'host':
'rrc25', 'type': 'UPDATE', 'path': [174, 20473], 'origin': 'igp', 'med': 78041, 'announcements': [{'next_hop':
'2001:978:4::b', 'prefixes': ['2a0e:97c6:5226::/48', '2a0e:97c6:506d::/48', '2604:4540:82::/48']}]
```

Stealthy sub-prefix attack

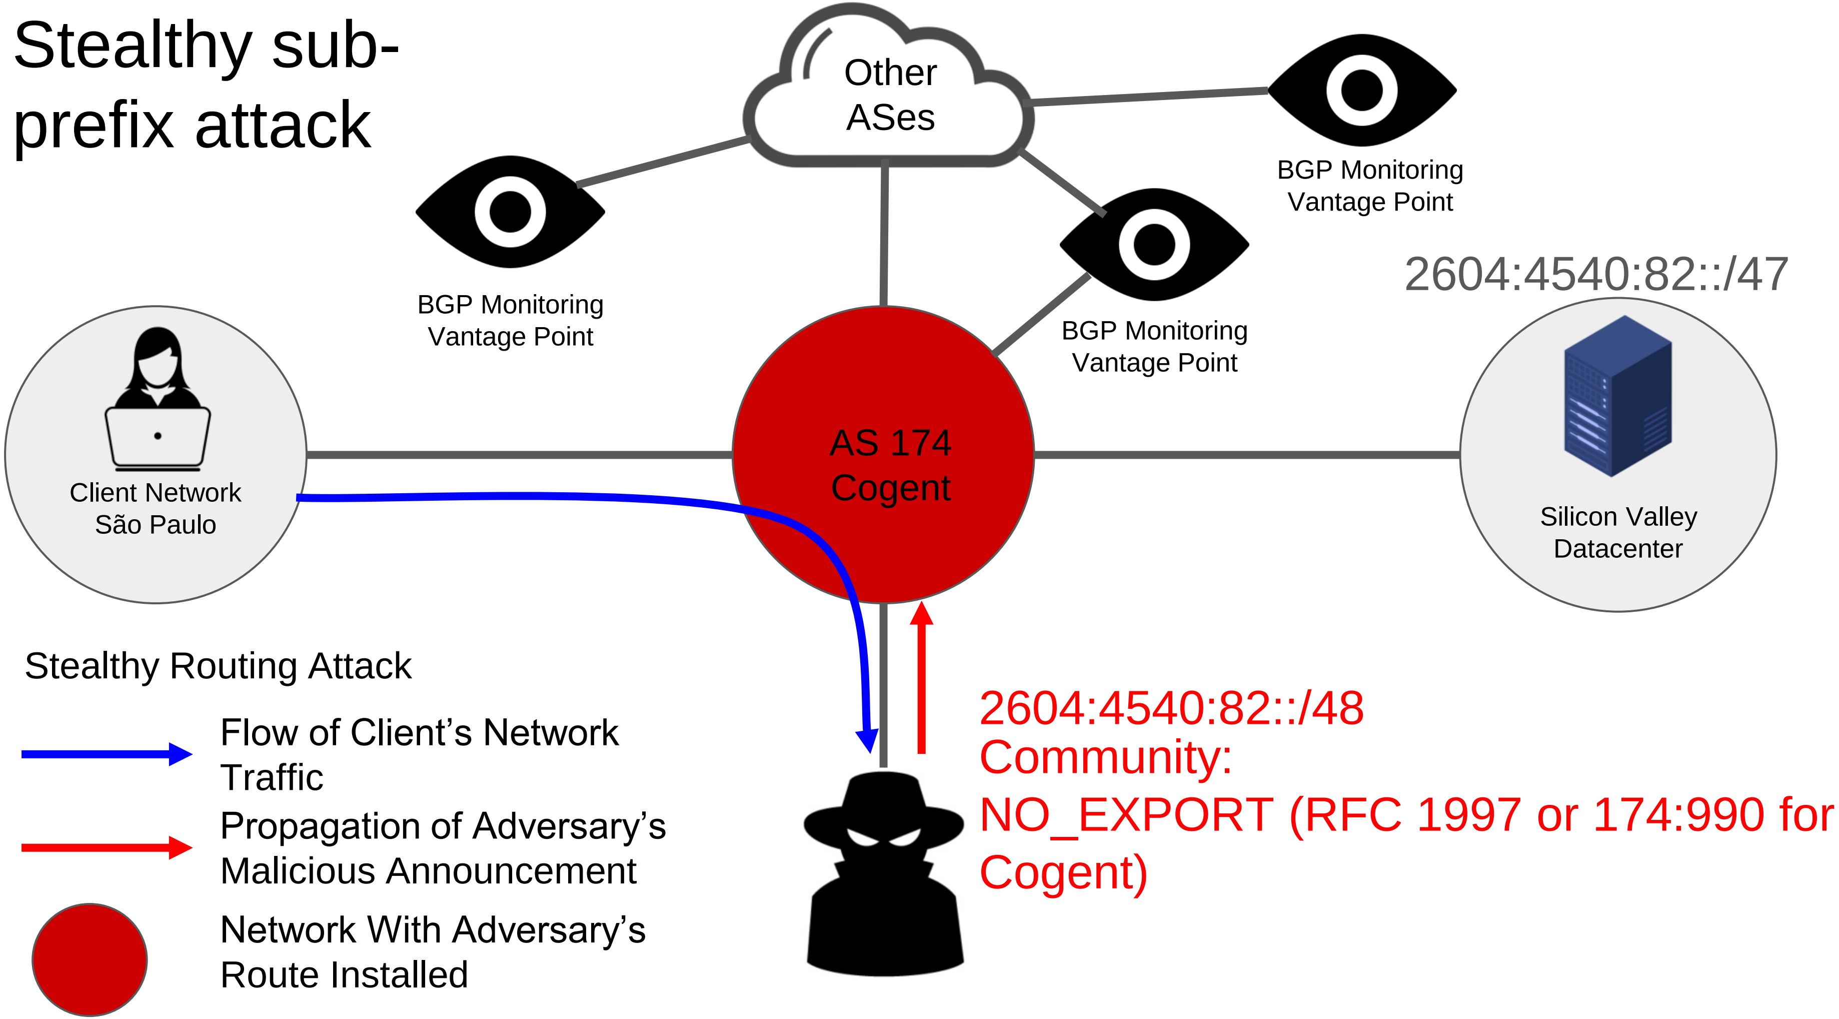
Other ASes
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
BGP Monitoring Vantage Point
2604:4540:82::/47
Client Network São Paulo
AS 174 Cogent
Silicon Valley Datacenter

2604:4540:82::/48
Community:
NO_EXPORT (RFC 1997 or 174:990 for Cogent)

Stealthy Routing Attack

- Flow of Client's Network Traffic
- Propagation of Adversary's Malicious Announcement
- ● Network With Adversary's Route Installed

Stealthy attack is seen nowhere

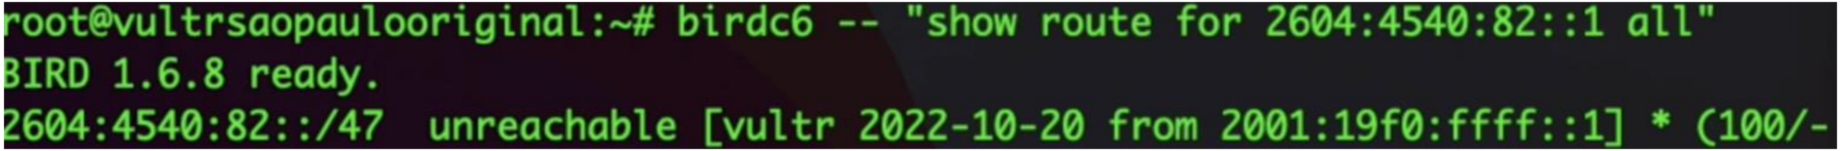

```
root@vultrsaopaulooriginal:~# birdc6 -- "show route for 2604:4540:82::1 all"
BIRD 1.6.8 ready.
2604:4540:82::/47  unreachable [vultr 2022-10-20 from 2001:19f0:ffff::1] * (100/-
```

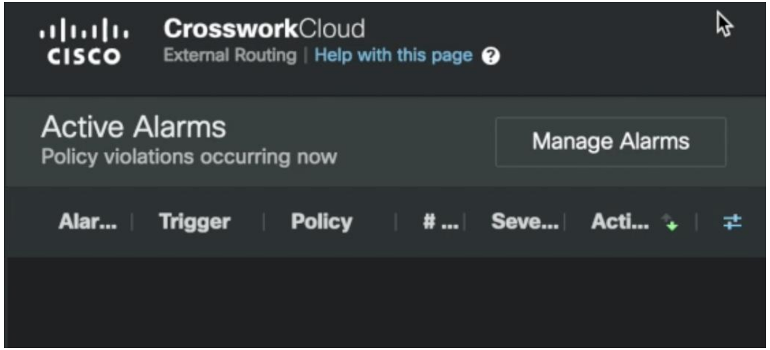

ris_subscribe_ok {'subscription': {'prefix': '2604:4540:82::/48', 'moreSpecific': True, 'lessSpecific': False}, 'socketOptions': {'includeRaw': True, 'acknowledge': True}}

100% of Internet was affected by stealthy route

1. Ran a ping scan to a 1k randomly selected ping hosts
2. Source IP address of the scan was within the hijacked prefix
    a. I.e., Hosts responded to either victim or adversary depending on whether they were affected by the attack
3. In scan, 100% of responses went to adversary
    a. Prefix owner exclusively used Cogent for transit
4. See attack demo video for more: https://drive.google.com/file/d/19Cr0JNRXeYqlWKoPR8K39JlDImlnPmC4/view?usp=sharing

Overview

1. Motivation and Terminology
2. Background
3. Launching the Attack in the Real World
4. **Attack Viability**
5. Countermeasures

How many networks have this behavior

Many thanks to:

RGNet
Edge
Princeton OIT

Network	Cogent	Telia	NTT	Sprint
Supports NO_EXPORT	Yes	Yes	Yes	Yes
BGP Monitoring Session with RIPE RIS	Yes	Yes	Yes	Yes
Suppressed session with RIPE	Yes	Yes	Yes	Yes
Invisible from Thousandeyes	Yes	Yes	Yes	Yes
Invisible from Cisco Crosswork Cloud	Yes	Yes	Yes	Yes

How much of the Internet is hijackable via these networks

1. Ran simulations of AS-level paths between 150 random ASes using the CAIDA topology

2. Just assuming an adversary installed a malicious route in the networks from the previous slide, picked a random victim AS's prefix to hijack, the adversary would affect traffic from 23% source ASes on the Internet (on average)

3. If an adversary instead chooses the top 5 largest ASes by customer cone, (3356, 1299, 174, 2914, 6762), on average it could hijack traffic from 39% of source ASes
    a. previous work on stealthy hijacks 2% of Internet traffic (90th percentile) and median case is < .2%

4. A strategic adversary could do even better (e.g., 100% like in our experiment)

How viable is a subprefix attack

1. RPKI can prevent subprefix attacks

    a. RPKI ROAs specify proper prefix length, announcements for other prefix lengths are blocked

    b. ROAs now cover ~50% of IP prefixes, ~50% still not protected

    c. Of the ~50% of IP prefixes covered by RPKI, about 36% (18% of total) have an improper maxLength attribute leaving them still vulnerable to sub-prefix attacks [1] [2]

2. My 2023 Usenix Security Paper [2] “How Effective is Multiple-Vantage-Point Domain Control Validation?” looked at TLS domains and found only %29.2 had all A record and Nameserver IPs protected against subprefix attacks

    a. About half of these protected domains (15.2%) ran exclusively out of /24 announcements. Only %18.2 of domains were protected because of RPKI.

3. Even with great RPKI progress, subprefix attacks are still highly viable (massive numbers in the wild each year)

[1] Yossi Gilad, Omar Sagga, and Sharon Goldberg. 2017. MaxLength Considered Harmful to the RPKI. In International Conference on Emerging Networking EXperiments and Technologies (Incheon, Republic of Korea) (CoNEXT). Association for Computing Machinery, New York, NY, USA, 101–107. https://doi.org/10.1145/3143361.3143363

Would a major network really propagate a malicious route directly from a customer?

1. Unfortunately: Yes
2. Some examples:
    a. “AS3266: BitCanal hijack factory, courtesy of Cogent, GTT, and Level3” https://mailman.nanog.org/pipermail/nanog/2018-June/096034.html
    b. “I guess AS1299 Arelion doesn't check if the origin AS of an announcement is in the customer's AS-SET but it's pretty normal and understandable.” https://mailman.nanog.org/pipermail/nanog/2022-August/220320.html
    c. “As such, it seems likely a peer or customer of AS6461 [Zayo] was [launching BGP hijacks].” https://mailman.nanog.org/pipermail/nanog/2022-February/217602.html
3. Know your customer is important but only effective to an extent because it does not rely on true cryptographic identifiers.

Overview

1. Motivation and Terminology
2. Background
3. Launching the Attack in the Real World
4. Attack Viability
5. **Countermeasures**

Countermeasures

1. Networks need to not allow customer-applied NO_EXPORT to control exporting behavior to route monitoring
   a. Customer networks should not decide what is sent to BGP monitoring
   b. **Option 1**: Translate RFC NO_EXPORT to a different AS-specific use community that only applies to normal BGP neighbors and not route monitoring
   c. **Option 2**: Run BGP monitoring over iBGP sessions
      i. Currently used by kentik, still appears to not be pervasive enough to catch our sample attacks
      ii. Still vulnerable to NO_ADVERTISE
   d. **Option 3**: Move monitoring to BMP (optimal long term)
   e. **Option 4**: Vender support flag (e.g., route_monitor_session BGP config flag)

Config Changes to Major Networks (option 1 rewriting) (pseudo config for AS 1234)

```
filter route_in {
        if (NO_EXPORT) in bgp_community then {
                bgp_community.delete([NO_EXPORT]);
                bgp_community.add(1234:997);
        }
        if (NO_EXPORT_SUBCONF) in bgp_community then {
                bgp_community.delete([NO_EXPORT_SUBCONF]);
                bgp_community.add(1234:998);
        }
        if (NO_ADVERTISE) in bgp_community then {
                bgp_community.delete([NO_ADVERTISE]);
                bgp_community.add(1234:999);
        }
        accept;}

filter bgp_neighbor_out {
        if (1234:997 or 1234:998 or 1234:999) in bgp_community then {
                reject;
        }
        accept;}

filter ripe_routeviews_out {accept;}
```

Config Changes to Major Networks (option 2 iBGP) (pseudo config for AS 1234)

```
protocol bgp routeviews {
    local as 1234;
    neighbor 10.142.12.6 as 1234; # RIPE/Routeviews could be configured to establish iBGP sessions using peer ASNs
    import none;
    export all;
    filter
    next hop self; # Don’t leak internal nexthop details to RIPE/RotueViews
}
```

RIPE NCC Admins: Please consider iBGP sessions with peers

(this is highly viable and already used by Kentik’s BGP monitoring https://kb.kentik.com/v0/Bd01.htm#Bd01-Router_BGP_Considerations)

Interest at Networks

- **Among US R&E networks: ESNet and Internet 2 are considering deploying countermeasures**
- **Any networks interested in deploying countermeasures for NO_EXPORT please reach out (birgelee@princeton.edu)**
- **See our technical report for more details https://arxiv.org/abs/2408.09622**

Questions

Henry Birge-Lee

Research Software Engineer
Princeton Electrical Engineering and Computer Science
birgelee@princeton.edu
https://henrybirgelee.com