Flexsealing BGP Against Route Leaks: Peerlock Active Measurement and Analysis

NDSS ‘21

Outline

- Introduction
- Background
  - BGP incidents
  - Route leak
- Peerlock/-lite
- Active peerlock deployment measurement
- Evaluation
- Conclusion

Introduction

- **BGP route leaks** frequently cause serious disruptions to inter-domain routing
  - These incidents have plagued the Internet for decades while deployment and usability issues cripple efforts to mitigate the problem

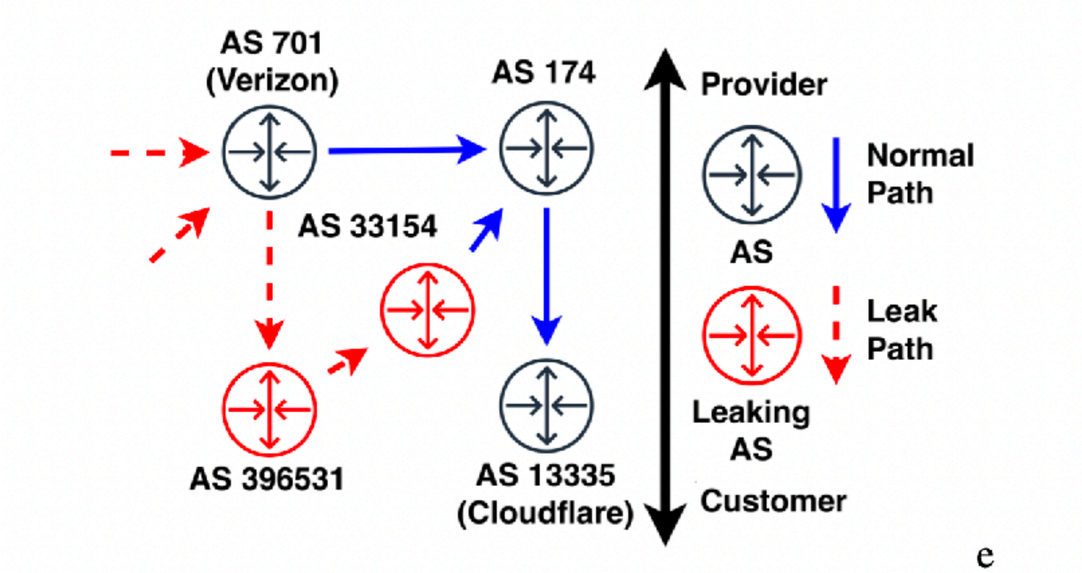

Fig. 1: 2019 Verizon/Cloudflare leak. Other destination services were also affected.

Introduction

- **Peerlock**, presented in 2016, addresses route leaks with a new approach.
  - filtering agreements **between transit providers** to protect their own networks without the need for broad cooperation or a trust infrastructure

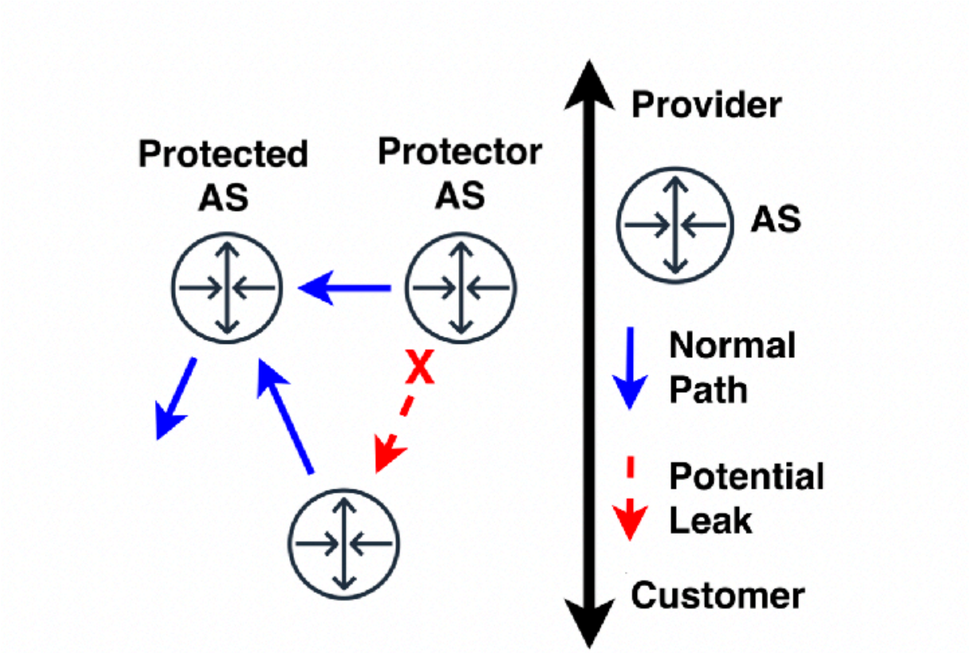

Fig. 3: Simple Peerlock deployment. Protector AS filters updates containing the peer Protected AS from unauthorized propagators.

Background: BGP incidents

- (sub-)prefix hijack
  - the *control plane* message (i.e., BGP announcement) contains invalid origin ASN
  - steer user traffic on the *data plane* onto invalid paths
- route leak
  - the *control plane* message (i.e., BGP announcement) propagates beyond its intended scope
  - steer user traffic on the *data plane* onto unintended paths

Background: Route Leak

- Route leaks are defined in RFC 7908 as the propagation of an advertisement beyond its intended scope
- Six types of leaks
  - **Type 1 - 4**: cover various **valley-free routing violations**
  - **Type 5**: occur when one provider's routes are announced to another with the AS PATH stripped, effectively **re-originating the prefix from the leaker**
  - **Type 6**: an AS announcing **routes used internally to its neighbors**. These routes are often more specific than externally announced routes

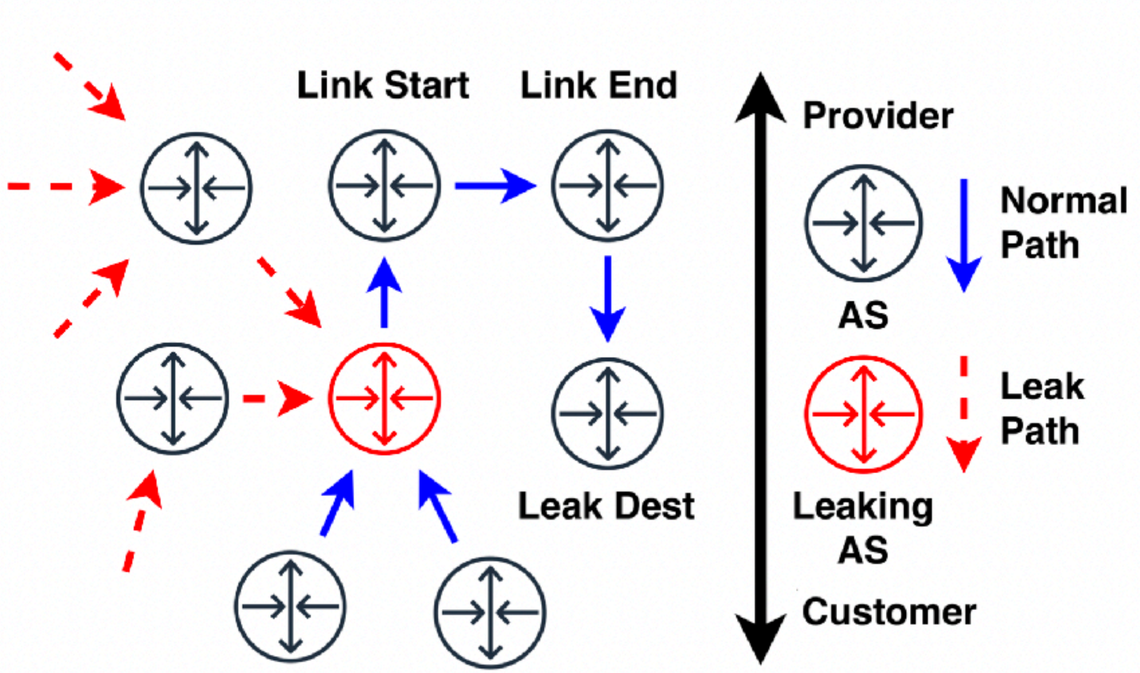

Background: Route Leak

- Route leaks are defined in RFC 7908 as the propagation of an advertisement beyond its intended scope
- Six types of leaks
  - **Type 1 - 4**: cover various **valley-free routing violations**
  - **Type 5**: occur when one provider's routes are announced to another with the AS PATH stripped, effectively **re-originating the prefix from the leaker**
  - **Type 6**: an AS announcing **routes used internally to its neighbors**. These routes are often more specific than externally announced routes

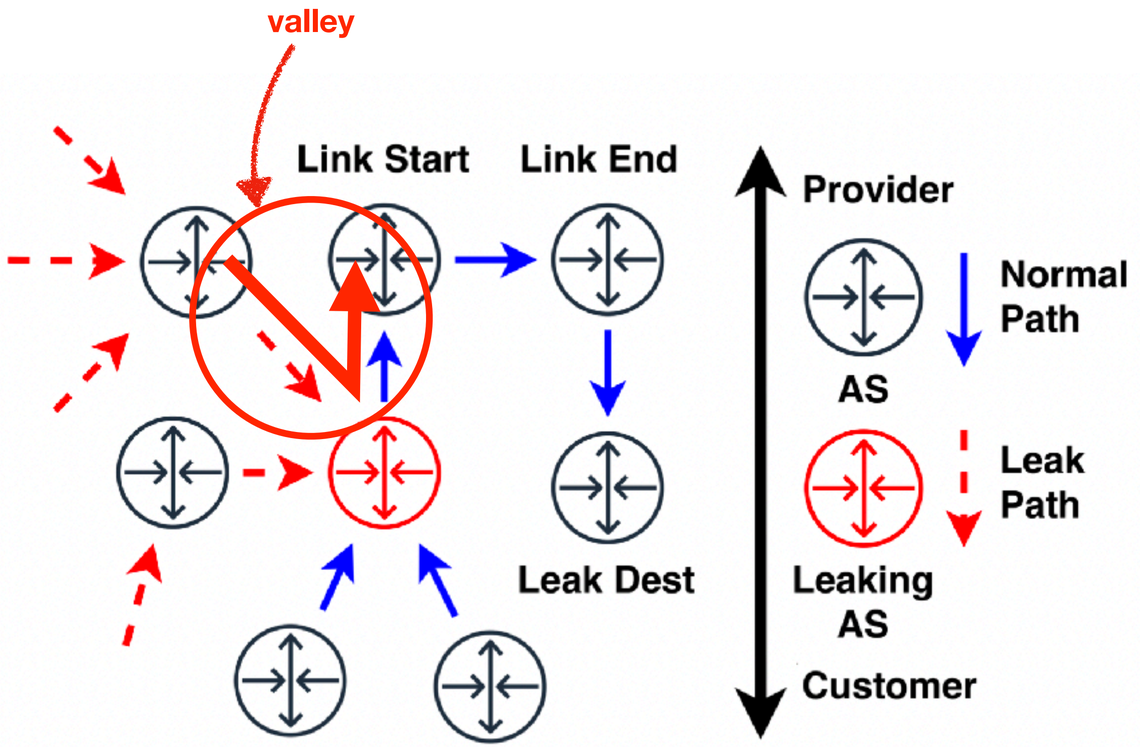

Peerlock

- Leak defense approach presented in 2016

- Peerlock deployment occurs between two neighboring ASes, the ***protector AS*** and ***protected AS.***
  - **protector AS**: rejects any BGP update whose AS PATH contains the protected AS unless received

    1) **directly from the protected AS** or

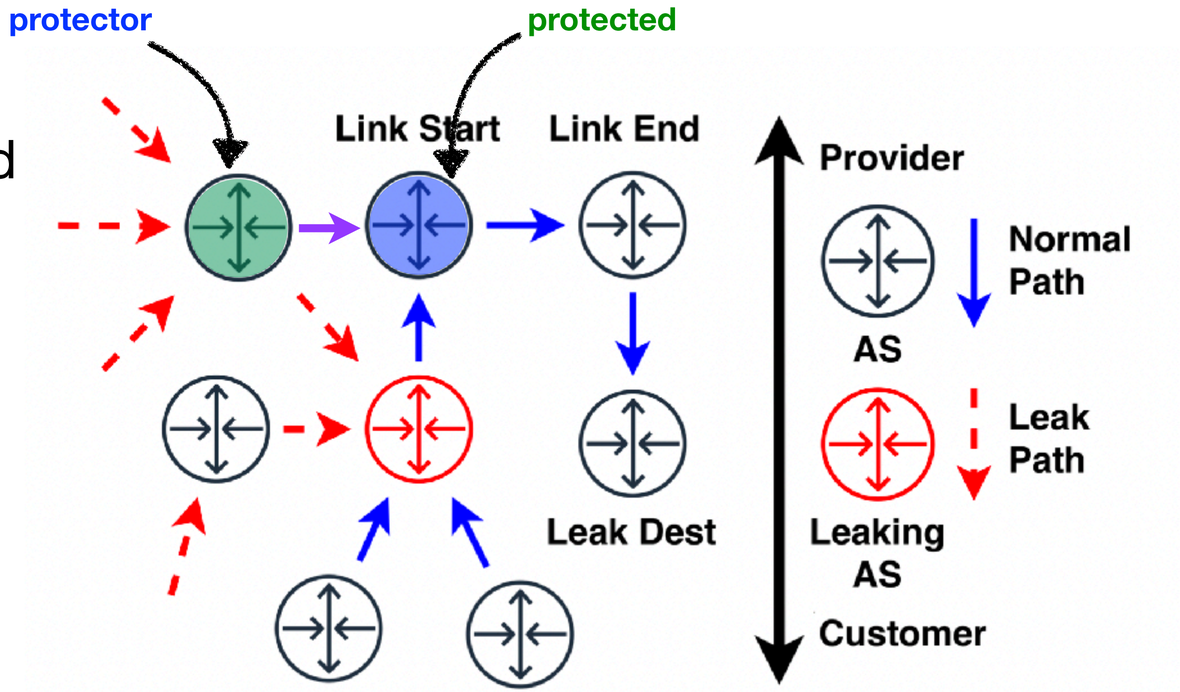

Peerlock

- Leak defense system was presented in 2016

- Peerlock deployment occurs between two neighboring ASes, the ***protector AS*** and ***protected AS.***
  - **protector AS**: rejects any BGP update whose AS PATH contains the protected AS unless received

    1) **directly from the protected AS** or

    2) **from an authorized upstream**, with the protected AS immediately following the authorized upstream in the AS PATH

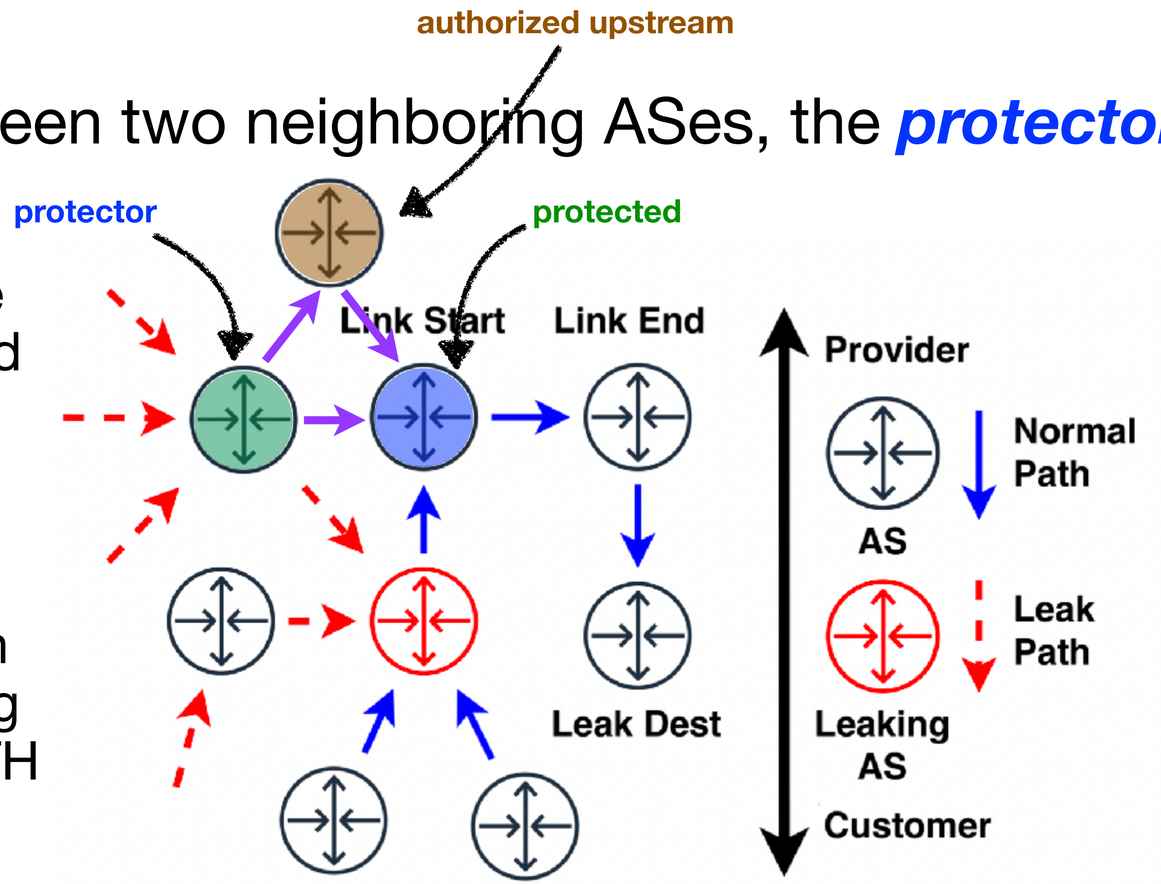

Peerlock

- Leak defense system was presented in 2016

- Peerlock deployment occurs between two neighboring ASes, the ***protector AS*** and ***protected* AS.**
  - The filter prevents the **protector AS** from propagating or steering its traffic onto any leaked route that transits the **protected AS**, **regardless of origin AS, destination prefix**

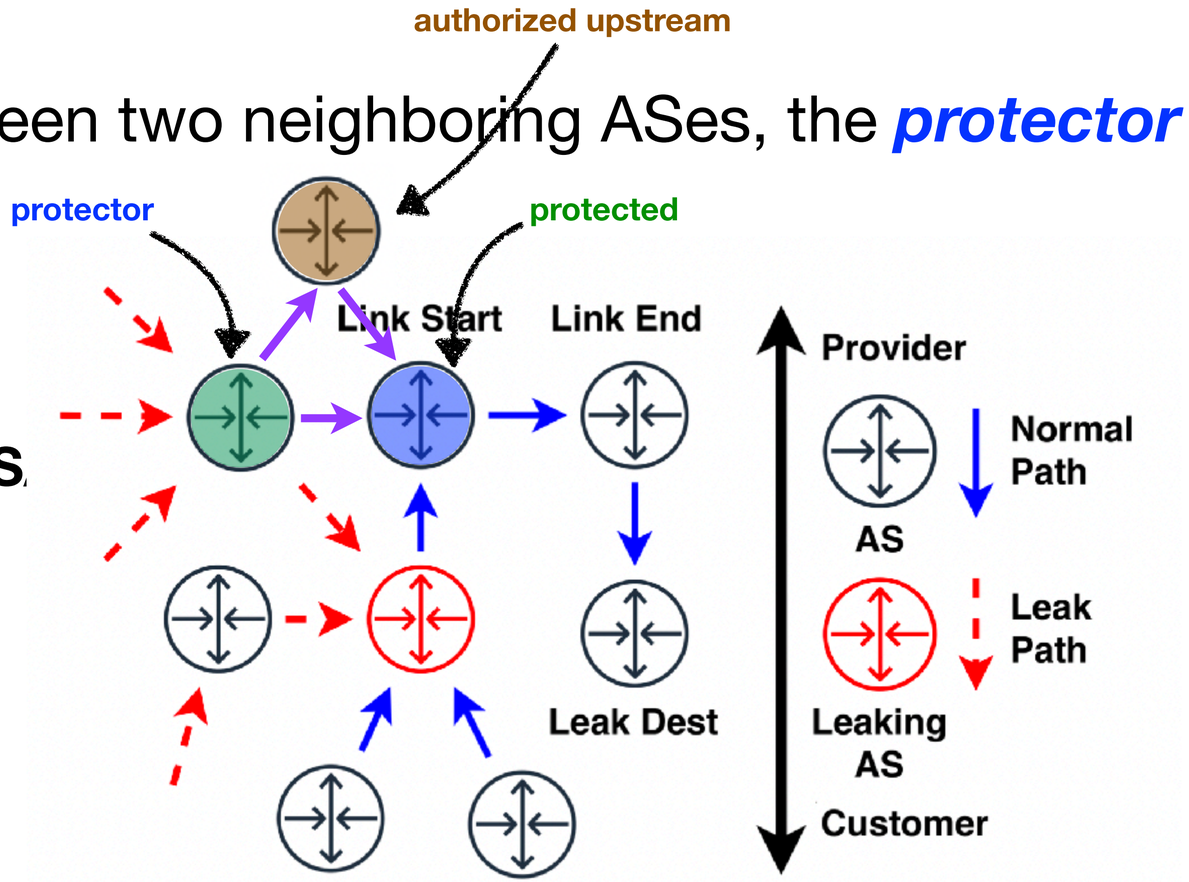

Peerlock-lite

- ***Peerlock-lite*** (or Tier 1 filter, "big networks" filter) is a related technique
  - No agreement between Tier1 AS and provider AS: a provider AS simply checks if a route from a customer includes Tier 1 AS in its path and, if it is true, drops it.
  - based on the assumption that transit providers should never receive a route whose AS PATH includes a Tier 1 AS from a customer (valid assumption under valley-free routing model)

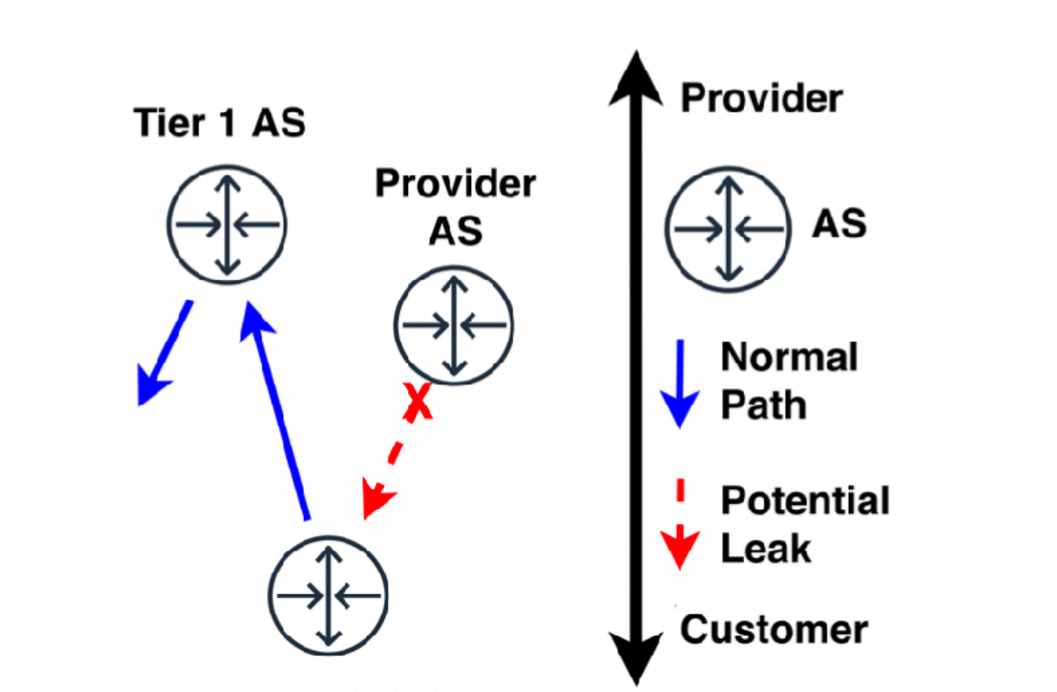

Active Peerlock Deployment Measurement

- Employs a BGP Poisoning technique to mimic route leaks transiting the poisoned AS

- *BGP poisoning* is a technique designed to manipulate the BGP decision process in remote networks
  - Poisons can be used for inbound traffic engineering purposes
  - an **origin AS** of a prefix can *poison* an advertisement of its prefix by including the ASNs of remote networks in the AS PATH
  - the **poisoned ASNs** will be inserted between copies of the origin's ASN
  - This "sandwiching" ensures traffic is routed properly and that the advertisement is valid for ROV filtering purposes

BGP Poisoning

- AS 1, the origin of prefix 1.2.3.0/24, inserts AS 3, the poisoned AS, in AS_PATH of its advertisement

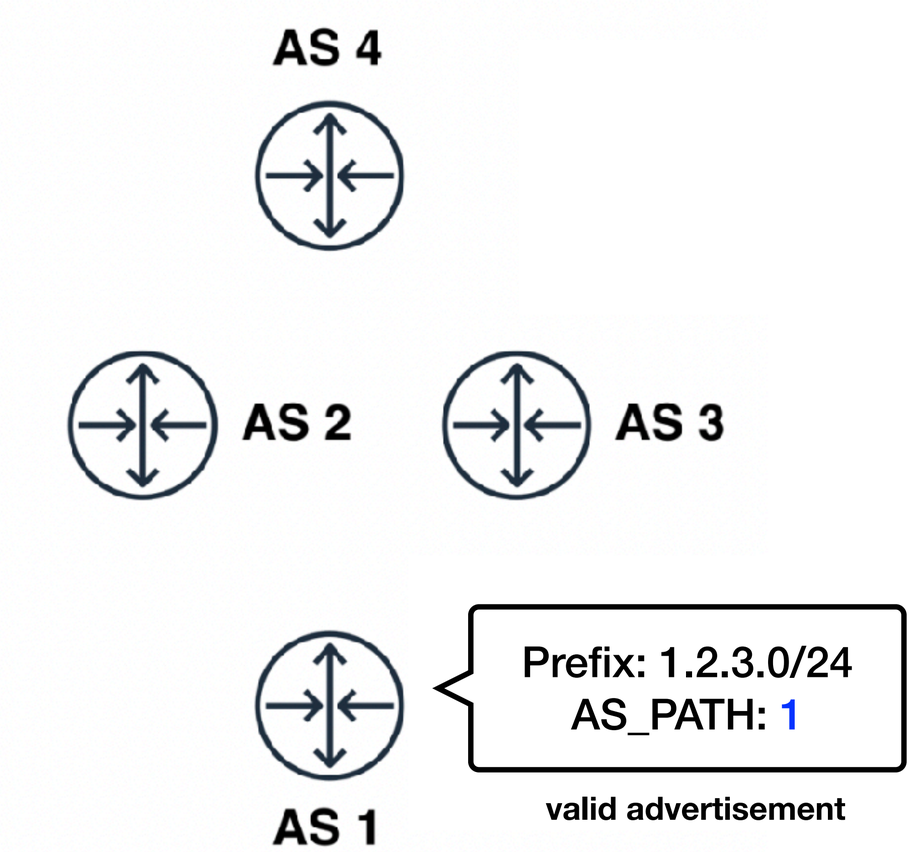

BGP Poisoning

- AS 1, the origin of prefix 1.2.3.0/24, inserts AS 3, the poisoned AS, in AS_PATH of its advertisement

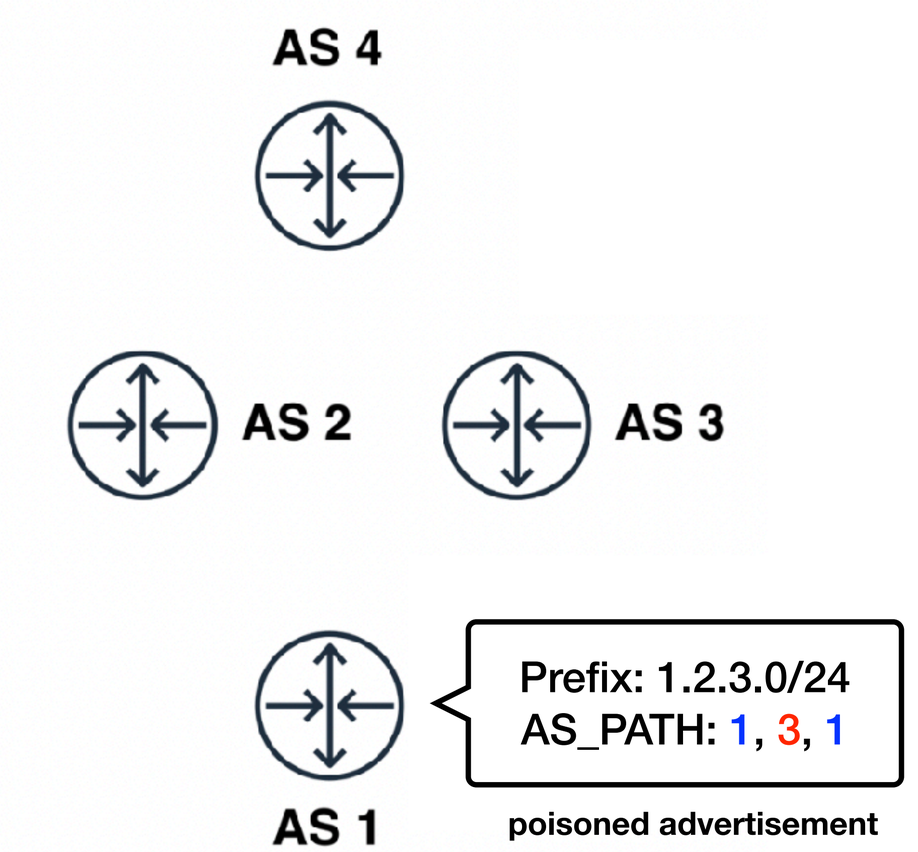

BGP Poisoning

- AS 1, the origin of prefix 1.2.3.0/24, inserts AS 3, the poisoned AS, in AS_PATH of its advertisement
  - poisoned AS will drop the advertisement

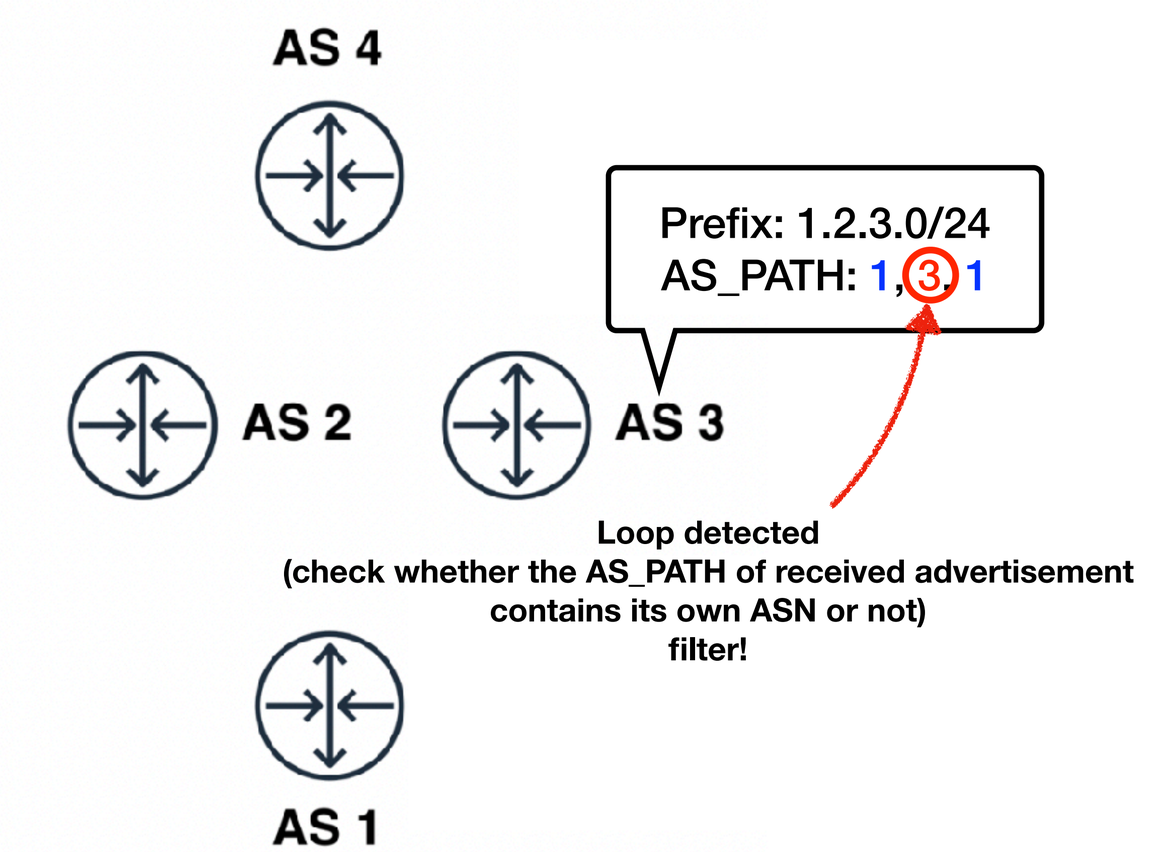

BGP Poisoning

- AS 1, the origin of prefix 1.2.3.0/24, inserts AS 3, the poisoned AS, in AS_PATH of its advertisement
  - poisoned AS will drop the advertisement
  - but, other ASes will forward the advertisement since it contains valid origin AS

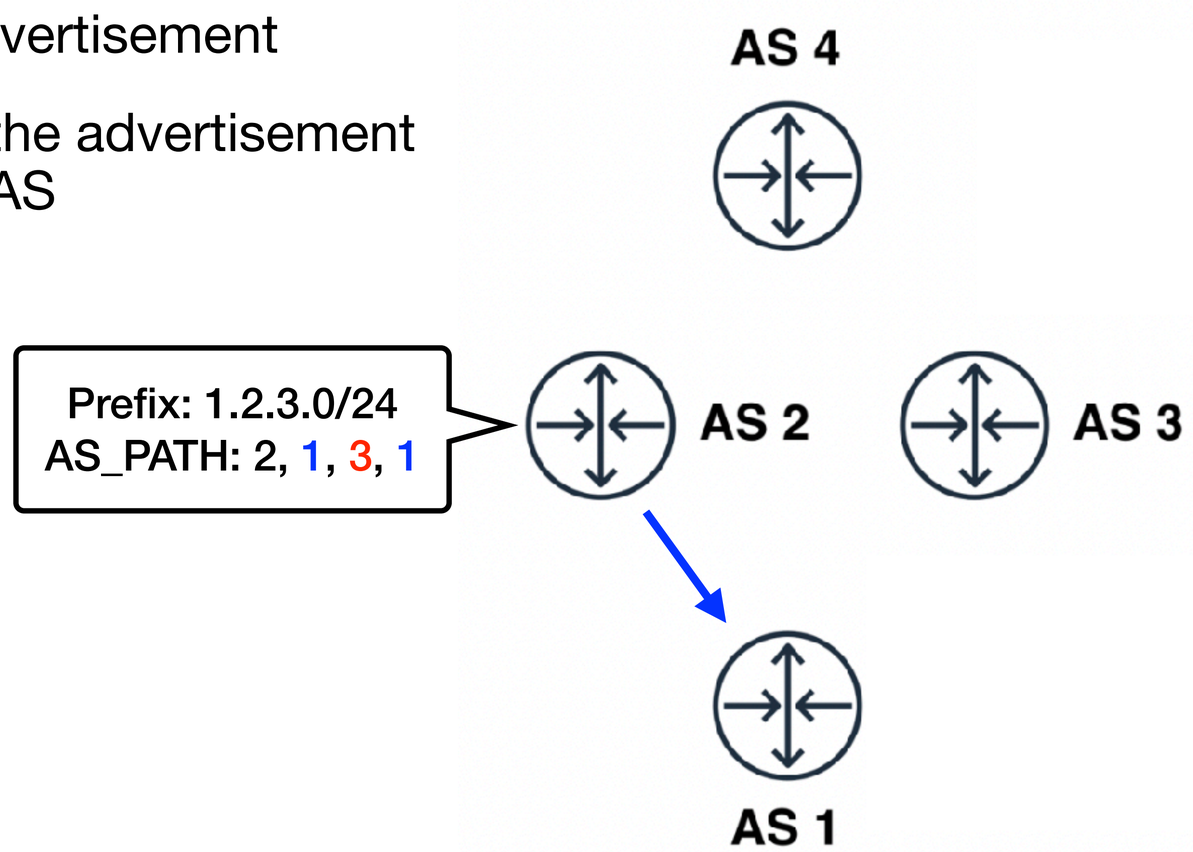

BGP Poisoning

- AS 1, the origin of prefix 1.2.3.0/24, inserts AS 3, the poisoned AS, in AS_PATH of its advertisement
  - poisoned AS will drop the advertisement
  - but, other ASes will forward the advertisement since it contains valid origin AS

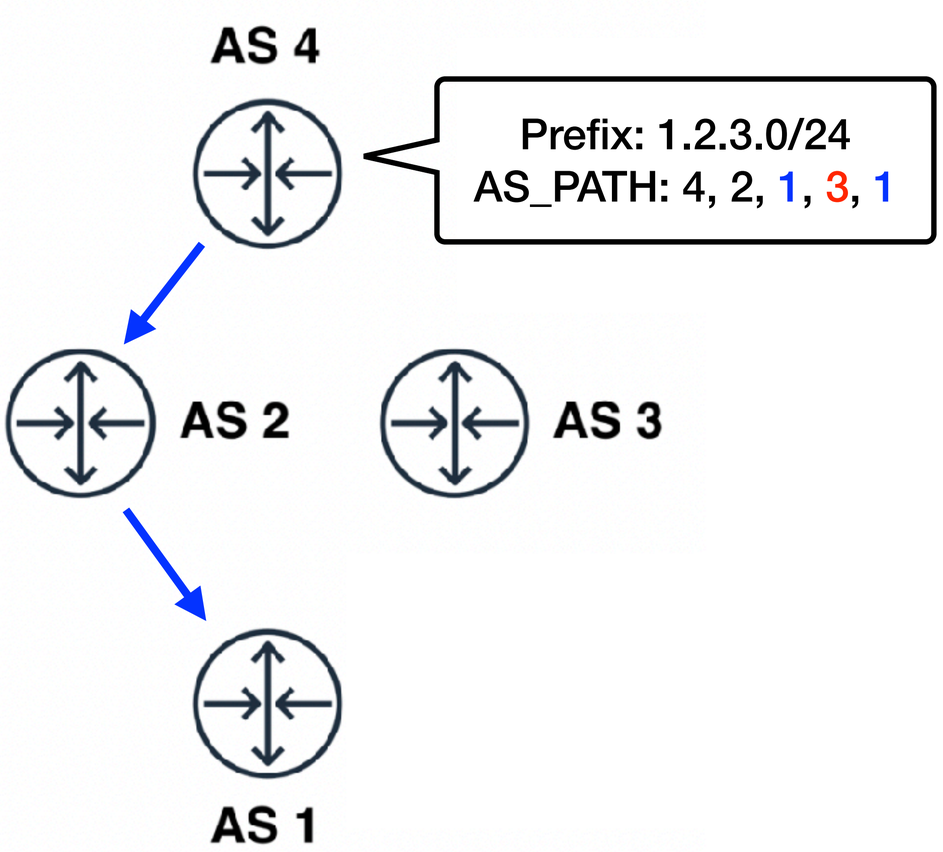

Active Peerlock Deployment Measurement

- send control and target advertisements
  - prefix = 1.2.3.0/24, AS_PATH = [1, **control ASN**, 1] (control advertisement)
  - prefix = 1.2.3.0/24, AS_PATH = [1, **target ASN**, 1] (target advertisement)

- Assumption
  - if target AS is protected by an AS that peerlocks with the target AS,
  - then target advertisement will be dropped by the AS
  - since the update is not directly forward by target ASN nor one of its authorized upstreams

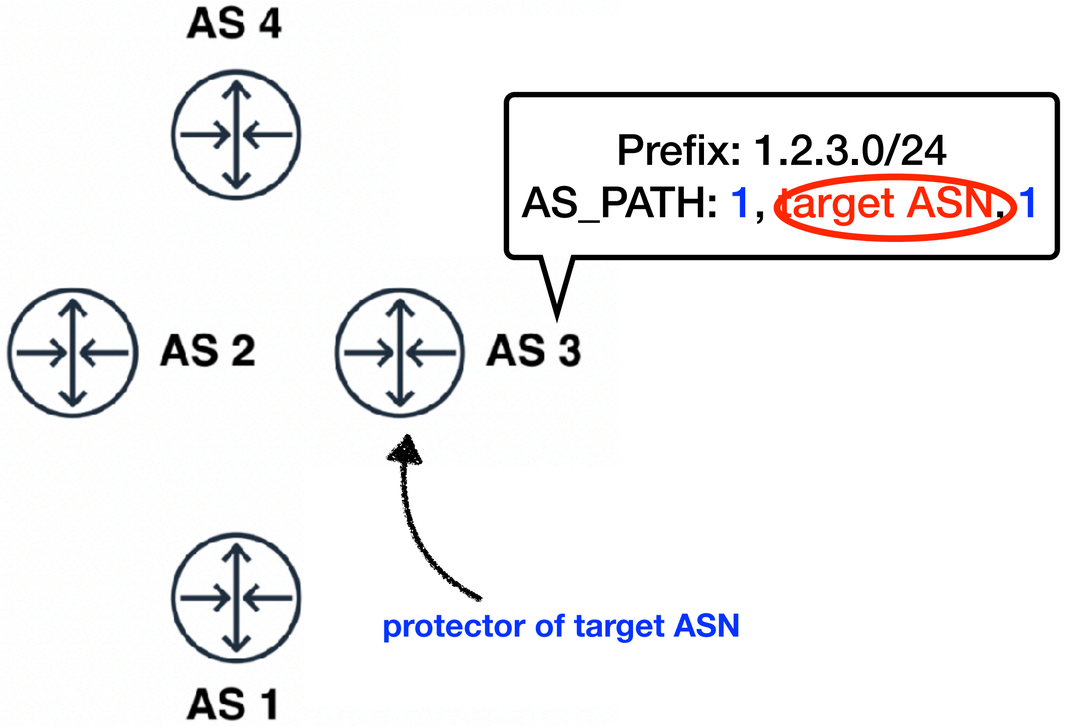

Active Peerlock Deployment Measurement

- send control and target advertisements
  - prefix = 1.2.3.0/24, AS_PATH = [1, **control ASN**, 1] (control advertisement)
  - prefix = 1.2.3.0/24, AS_PATH = [1, **target ASN**, 1] (target advertisement)

- Assumption
  - if target AS is protected by an AS that peerlocks with the target AS,
  - then target advertisement will be dropped by the AS
  - since the update is not directly forward by target ASN nor one of its authorized upstreams

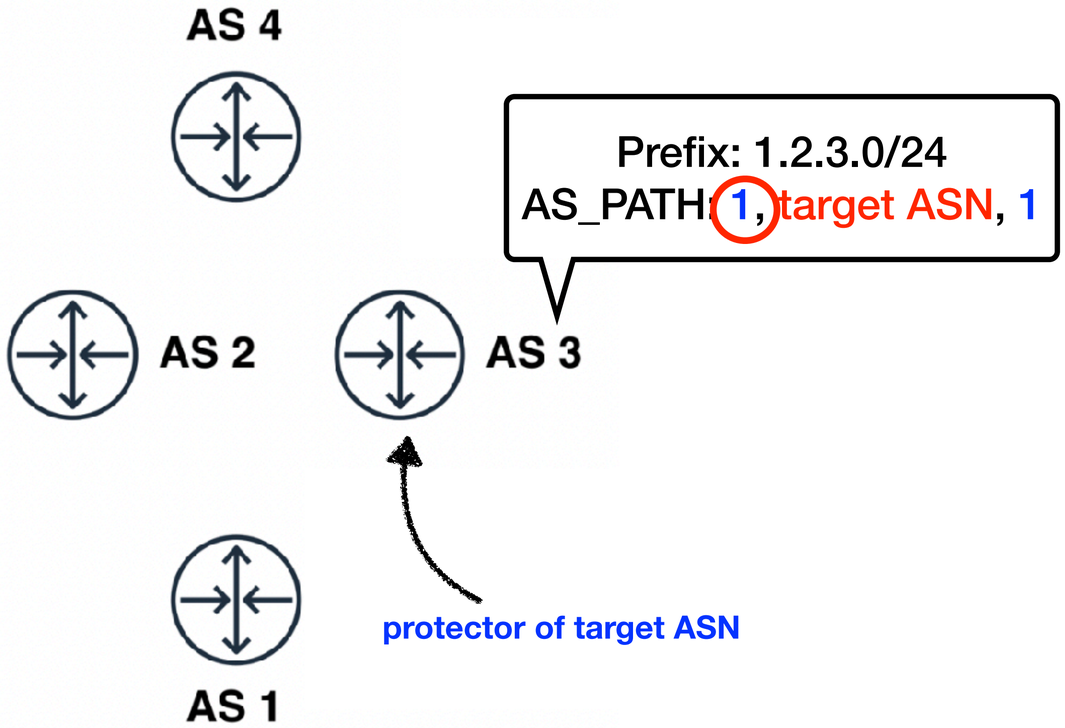

Active Peerlock Deployment Measurement

- send control and target advertisements
  - prefix = 1.2.3.0/24, AS_PATH = [1, **control ASN**, 1] (control advertisement)
  - prefix = 1.2.3.0/24, AS_PATH = [1, **target ASN**, 1] (target advertisement)

- Assumption
  - if target AS is protected by an AS that peerlocks with the target AS,
  - then target advertisement will be dropped by the AS
  - since the update is not directly forward by target ASN nor one of its authorized upstreams

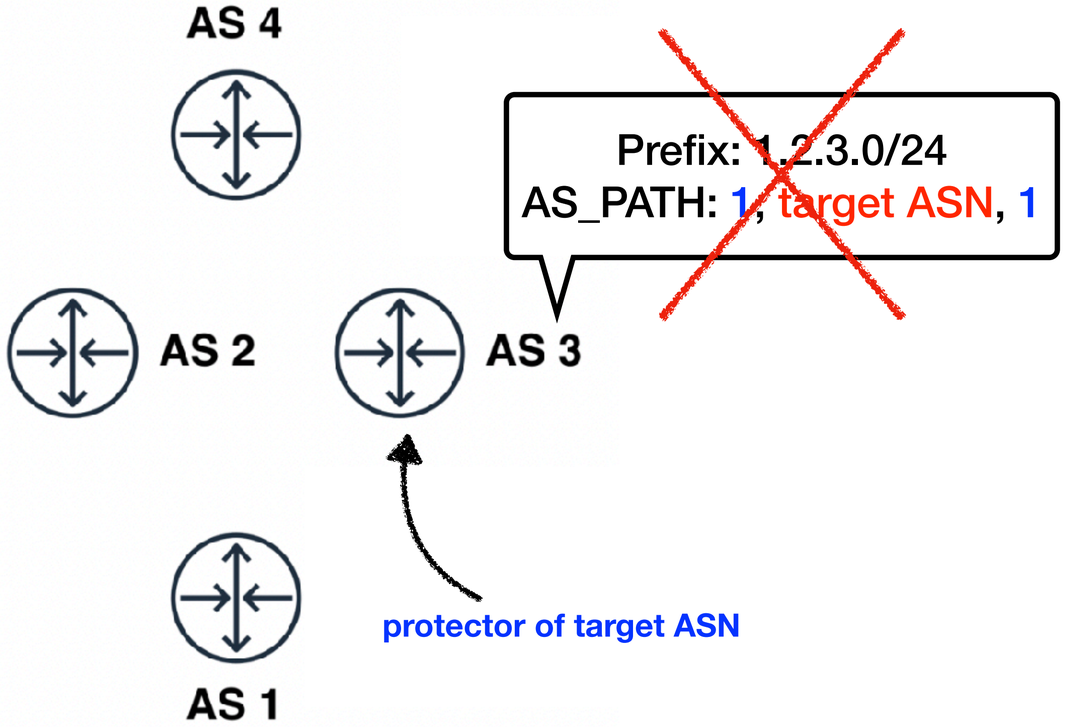

Active Peerlock Deployment Measurement

- send control and target advertisements
  - prefix = 1.2.3.0/24, AS_PATH = [1, **control ASN**, 1] (control advertisement)
  - prefix = 1.2.3.0/24, AS_PATH = [1, **target ASN**, 1] (target advertisement)

- collect those advertisements from global collectors
  - 30 RouteViews and 24 RIPE RIS collectors

- figure out protector ASes of target AS, by comparing the ASes in AS PATH of collected advertisements
  - ASes in AS_PATH are ASes that forward those advertisements
  - if an AS is only included in AS_PATHs of control advertisements, then it could be protector of target ASN

Active Peerlock Deployment Measurement

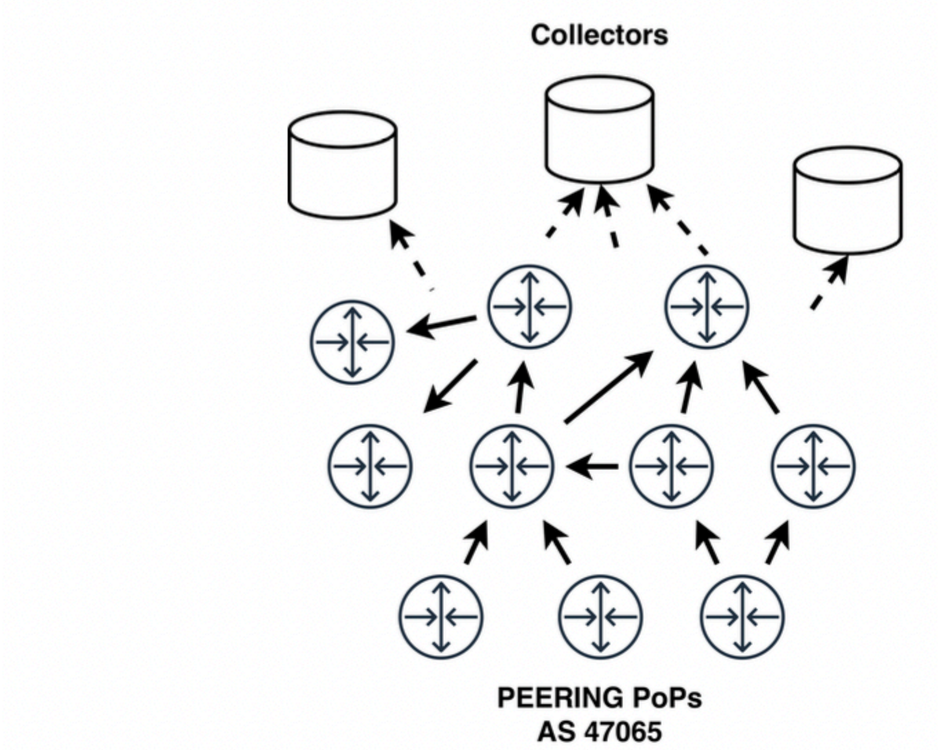

Propagation of control advertisements

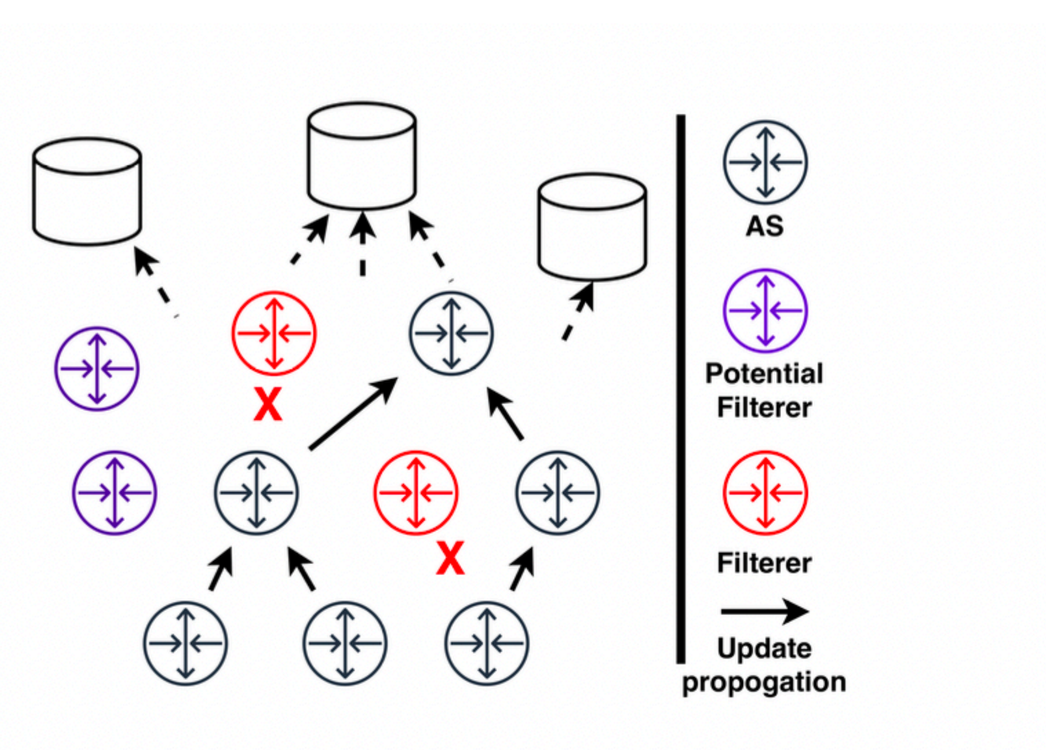

Propagation of target advertisements

Evaluation: Tier 1 leak within clique

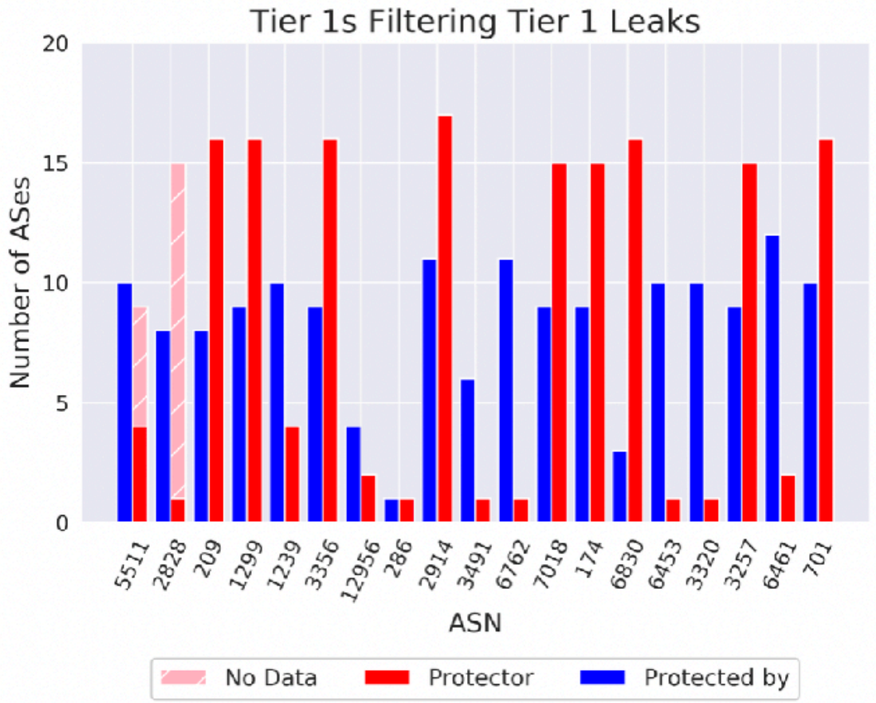

(a) Number of protector/protected rules by ASN. Protector numbers include ASes protecting their own ASN via loop detection.

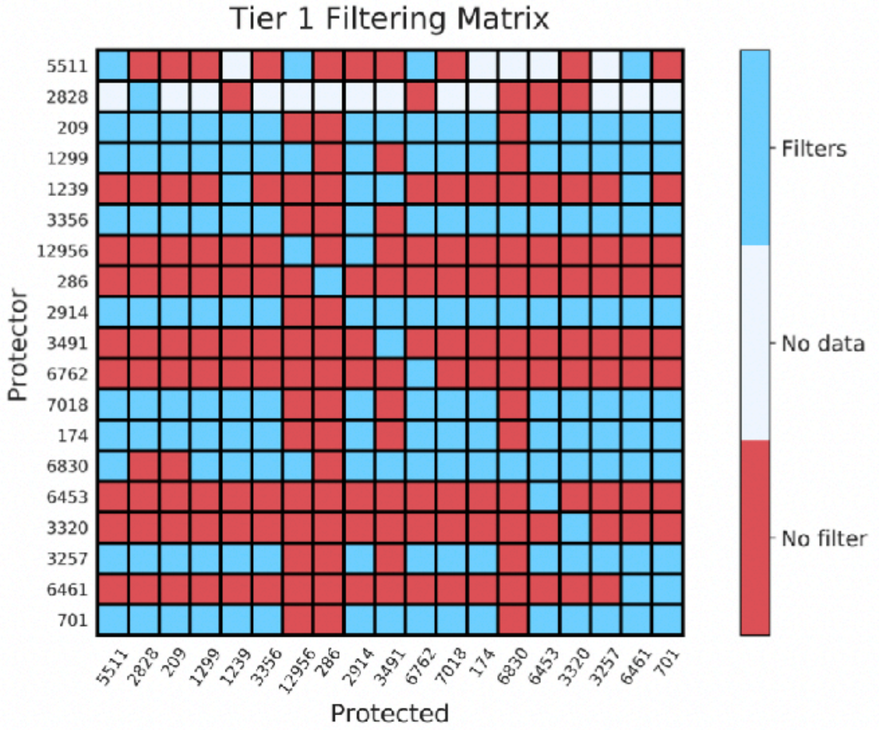

(b) Depiction of Tier 1 protection rules.

Fig. 6: Tier 1s filtering Tier 1 leaks, 2019/2020 measurements.

Evaluation: Peerlock/Peerlock-lite simulation

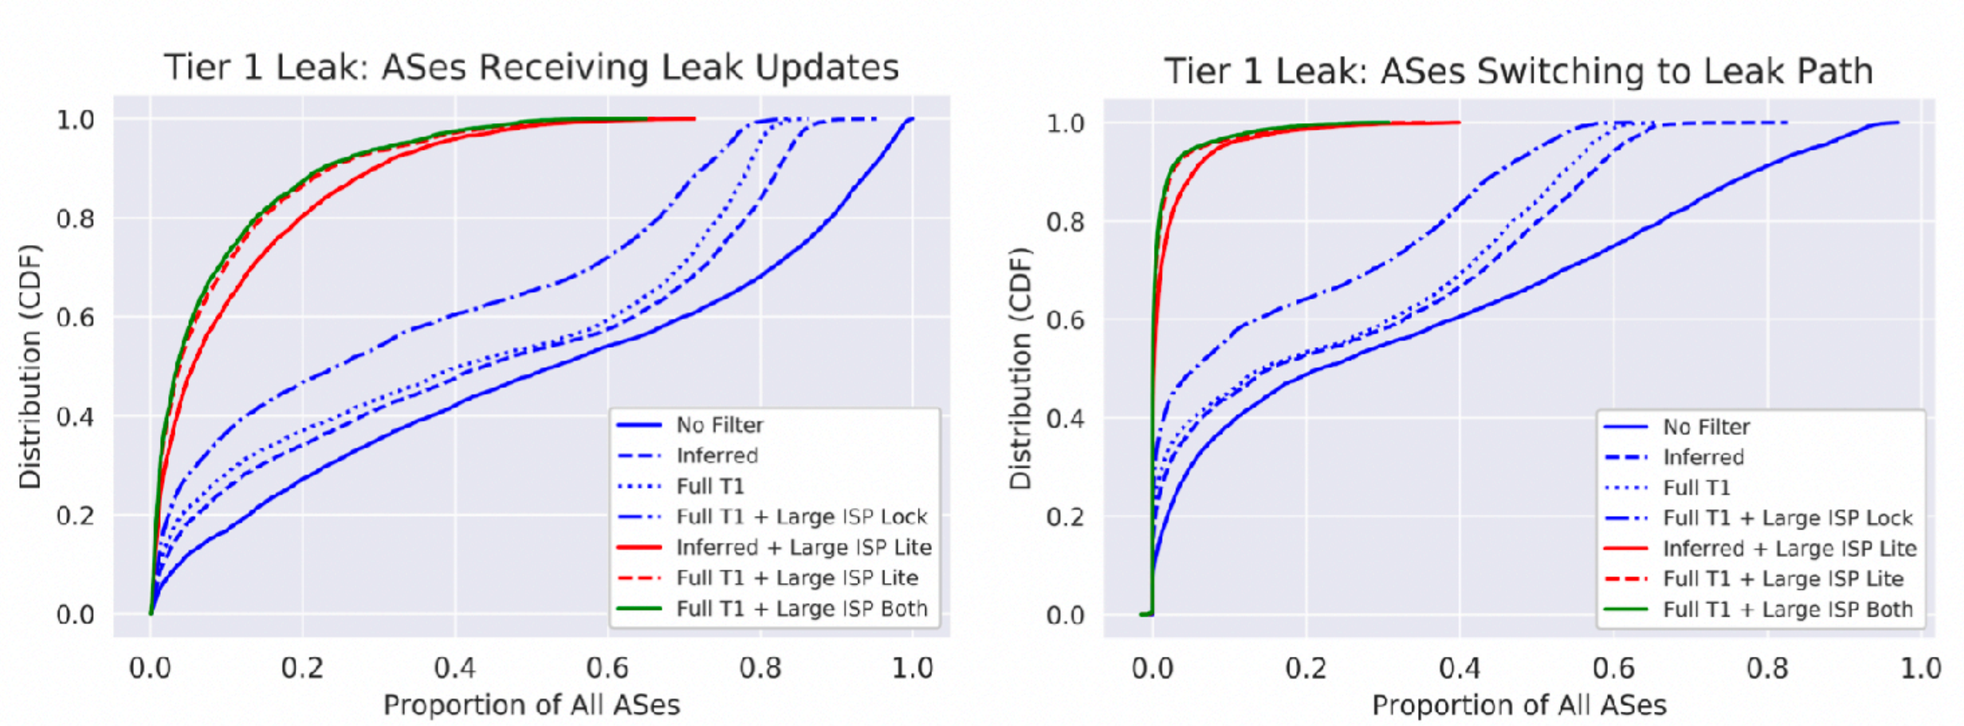

(a) Impact of various deployment scenarios on leak update propagation.

(b) Note increased Peerlock-lite performance for path switching vs. leak update propagation.

Fig. 11: Peerlock/Peerlock-lite simulation results.

Conclusion

- Probes the current deployment of Peerlock/Peerlock-lite on the control plane with active Internet measurements

- find substantial evidence for deployment of these leak defense systems, especially in large transit networks

- measure a rise in Peerlock deployment within the peering clique during our experiments